THE REVIEW OF SYMBOLIC LOGIC
Volume 14, Number 4, December 2021

PLANS AND PLANNING IN MATHEMATICAL PROOFS

YACIN HAMAMI
Centre for Logic and Philosophy of Science, Vrije Universiteit Brussel
and
REBECCA LEA MORRIS
Independent Scholar

Abstract. In practice, mathematical proofs are most often the result of careful planning by the agents who produced them. As a consequence, each mathematical proof inherits a plan in virtue of the way it is produced, a plan which underlies its "architecture" or "unity." This paper provides an account of *plans* and *planning* in the context of mathematical proofs. The approach adopted here consists in looking for these notions not in mathematical proofs themselves, but in the *agents* who produced them. The starting point is to recognize that to each mathematical proof corresponds a *proof activity* which consists of a sequence of *deductive inferences*—i.e., a sequence of *epistemic actions*—and that any written mathematical proof is only a report of its corresponding proof activity. The main idea to be developed is that the plan of a mathematical proof is to be conceived and analyzed as the plan of the agent(s) who carried out the corresponding proof activity. The core of the paper is thus devoted to the development of an account of plans and planning in the context of proof activities. The account is based on the theory of planning agency developed by Michael Bratman in the philosophy of action. It is fleshed out by providing an analysis of the notions of *intention*—the elementary components of plans—and *practical reasoning*—the process by which plans are constructed—in the context of proof activities. These two notions are then used to offer a precise characterization of the desired notion of plan for proof activities. A fruitful connection can then be established between the resulting framework and the recent theme of *modularity* in mathematics introduced by Jeremy Avigad. This connection is exploited to yield the concept of *modular presentations* of mathematical proofs which has direct implications for how to write and present mathematical proofs so as to deliver various epistemic benefits. The account is finally compared to the technique of proof planning developed by Alan Bundy and colleagues in the field of automated theorem proving. The paper concludes with some remarks on how the framework can be used to provide an analysis of understanding and explanation in the context of mathematical proofs.

§1. Introduction. Mathematicians often refer to the *plan* of their mathematical proofs. The two following examples are extracted verbatim from articles published in the *Annals of Mathematics*:

Received: April 10, 2019.
2010 *Mathematics Subject Classification*: 00A30, 00A35, 03A05.
Key words and phrases: mathematical proof, proof plan, planning agency, mathematical practice.

doi:10.1017/S1755020319000601

> The plan of the proof of the λ_g conjecture is as follows. We first prove (11) is equivalent to the λ_g conjecture in Section 2.3. The next step is to show the solution (11) satisfies all of our linear relations (13). This result is established in Section 2.4 via known $g = O$ formulas. In Section 2.6, the linear relations are proven to admit at most a one-dimensional solution space in each degree. Together these three steps prove the λ_g conjecture. (Faber & Pandharipande, 2003, p. 108)

> The general plan of the proof is as follows. Let $\mathcal{O}^\infty$ be defined as $\bigcup_{j\in\mathbb{Z}} f^j\mathcal{O}$ in J_+^-. Let $\hat{\mathcal{O}}$ denote the closure in U^+ of the set $\mathcal{O}^\infty$. We begin by proving analogs of (1–5) with J_+^- replaced by $\hat{\mathcal{O}}$. In Theorem 2.6 we prove an analog of (1). In proposition 2.7 we prove an analog of (2). In Corollary 2.8 we prove an analog of (3). In Corollary 2.9 we prove an analog of (4). In Corollary 2.10 we prove (5). Corollary 2.14 proves (6) and shows that $\hat{\mathcal{O}} = J_+^-$. Thus it follows that items (1–5) hold for J_+^- as claimed. (7) is proved in Corollary 2.18. (Bedford & Smillie, 1998, p. 703)

One mathematician and philosopher who pointed out the importance of the notion of plan for mathematical proofs is Henri Poincaré:

> When the logician has resolved each demonstration into a host of elementary operations, all of them correct, he will not yet be in possession of the whole reality; that indefinable something that constitutes the unity of the demonstration will still escape him completely. What good is it to admire the mason's work in the edifices erected by great architects, if we cannot understand the general plan of the master? (Poincaré, 1908, p. 126)

Poincaré considered that analyzing mathematical proofs by focusing only on the validity of their elementary deductive steps was bound to leave out one of their essential dimensions, namely what he refers to as their "unity." For Poincaré, that mathematical proofs have a unity is, presumably, a direct consequence of them inheriting a "general plan" in virtue of being produced by an "architect."

To the best of our knowledge, no attempts have been made so far in the philosophy of mathematics to propose an account of the notion of plan for mathematical proofs. And yet, in addition to its intrinsic interest for the epistemology of proofs in practice, the notion appears to play an important role in philosophical discussions of understanding and explanation in the context of mathematical proofs. For instance, Janet Folina, in a contribution dedicated to mathematical understanding, wrote the following:

> Poincaré and Feferman both draw attention to the importance of the *plan* of the proof; the unity that is found in its "architecture": how the parts are related and connected to one another. Emphasizing unity, and relationships between parts, resonates with Haylock's emphasis on making connections as central to mathematical understanding. In addition, the importance of seeing the whole, or the blueprint, of an argument also calls to mind Michener's view that mathematical

> understanding enables one to avoid getting lost in details.[1] (Folina, 2018, p. 136)

Similarly, the logician and computer scientist John Alan Robinson established a direct connection between the notions of understanding, explanation, and plan with respect to mathematical proofs:

> When we survey a real proof in this high-level, outline way, and fix our attention on its main idea or ideas, we can better intuitively appreciate its overall plan. We understand the proof as an explanation, in a sense, even though our view of it is neither rigorous nor complete. With only the overall plan before it, the mind is not concerned with the details. For the purpose of obtaining a (higher-level) understanding, it even seems essential that the (lower-level) details should be ignored. If too many details enter into the primary sketch, we simply lose sight of the main architecture of the proof—we are unable to see the big picture. (Robinson, 2000, p. 292)

If the notions of understanding and explanation in the context of mathematical proofs are indeed to be accounted for by appealing to the notion of plan, then getting clear on what is meant by the plan of a mathematical proof becomes a prerequisite to the development of an account of understanding and explanation along the lines indicated by these two passages.

The aim of this paper is to provide an account of plans and planning in the context of mathematical proofs. Our approach will consist in looking for these notions not in mathematical proofs themselves, but in the *agents* who produced them. To this end, we will begin by introducing the notion of *proof activity* as the counterpart of the notion of mathematical proof in the realm of action. The development of our account will then be driven by the idea that the plan of a mathematical proof is to be conceived and analyzed as the plan of the *agent*(*s*) who carried out the corresponding proof activity. To provide an account of the notion of plan for proof activities, we will build on the theory of planning agency developed by Michael Bratman (1987) in the philosophy of action. As we will see, the resulting notion of plan for mathematical proofs will turn out to be intimately connected to the recent theme of *modularity* put forward by Jeremy Avigad (2020). Indeed, the notions of plans and modularity can be fruitfully combined to provide a systematic way of generating what we will call *modular presentations* of mathematical proofs. This has direct implications for an important issue raised by Uri Leron (1983) and Leslie Lamport (1995, 2012) concerning the best ways to write and present mathematical proofs so as to maximize various epistemic benefits, such as their comprehensibility, communicability, readability, or rigor. Another intended application of our framework is getting to grips with the notions of understanding and explanation in the context of mathematical proofs. This is one of our main motivations for developing the present account, but for reasons of space our analysis of these notions will be provided in a subsequent paper. We will indicate in the conclusion how we intend to use our framework for this purpose.

---

[1] The references associated to the authors mentioned in this passage and cited in Folina (2018) are Poincaré (1996, 1908), Feferman (2012), Haylock & Cockburn (2008), and Michener (1978).

The paper is organized as follows. In §2, we define the notion of *proof activity* as the counterpart of the notion of mathematical proof in the realm of action. Our main task will be to develop an account of plans and planning in the context of proof activities, and for this we will build on Bratman's theory of planning agency. In §3, we will briefly present Bratman's theory of planning agency, focusing on the key notions of *intentions*, *practical reasoning*, and *plans*. §4, §5, and §6 will then be concerned respectively with the notions of intentions, practical reasoning, and plans in the context of proof activities. Taken together, these three sections provide an account of plans and planning in the context of proof activities, and insofar as we conceive of the plan of a mathematical proof as the plan of the agent(s) who carried out the corresponding proof activity, this yields our desired account of plans and planning in the context of mathematical proofs. In §7, we will discuss the connection between the notions of plan and modularity, and we will show that this connection can be exploited to yield the concept of modular presentations of mathematical proofs, a concept which has direct implications for the issues of how to write and present mathematical proofs as raised by Leron (1983) and Lamport (1995, 2012). In §8, we will compare our account with the framework known as *proof planning* developed by Alan Bundy and colleagues in the field of automated theorem proving (Bundy, 1988; Bundy, Basin, Hutter, & Ireland, 2005). We will end the paper with some concluding remarks on how to use the present framework to analyze the notions of understanding and explanation with respect to mathematical proofs, and on the interest of developing an account of plans and planning in the more general contexts of proof discovery and problem solving.

§2. Mathematical proofs and proof activities. Mathematical proofs have traditionally been conceived as static, agent-free objects. But if one approaches them through their primary epistemic function—namely to bring knowledge of the associated theorem—one is forced to see them in a dynamic way, and to bring the mathematical agent back into the picture. So given a written mathematical proof, what does a mathematical agent need to do to acquire knowledge of the associated theorem? Obviously, staring at the proof, or even reading it, won't do the job. What is required is that the agent *actively* carries out all the deductive inferences in the proof so as to reach an epistemic state in which she knows the theorem.

This perspective can be made precise by introducing some terminology. Taken as a static, written object, a *mathematical proof* is a sequence of *deductive steps*—this is the form in which mathematical proofs are ordinarily found in mathematical textbooks, articles, and monographs. But the two notions of deductive step and mathematical proof have direct counterparts in the realm of *action*. Corresponding to the notion of deductive step is the notion of *deductive inference*,[2] a deductive inference being first and foremost an *action* of an epistemic nature, as has been repeatedly emphasized in the philosophy of logic (see, e.g., Corcoran, 1989; Prawitz, 2012, 2015; Sundholm, 2012) and in epistemology (see, e.g., Boghossian, 2014; Wright, 2014).[3] To the best of our knowledge, no terms have been commonly used to designate the counterpart

[2] On deductive inferences in ordinary mathematical practice, see Hamami (2018, 2019).

[3] We assume here that the agent under consideration is indeed able to carry out the deductive inferences corresponding to the deductive steps in the considered mathematical proof, that is, the agent belongs to the intended audience of the written mathematical proof. For an in-depth discussion of the notion of proof as audience-relative, see Corcoran (1989).

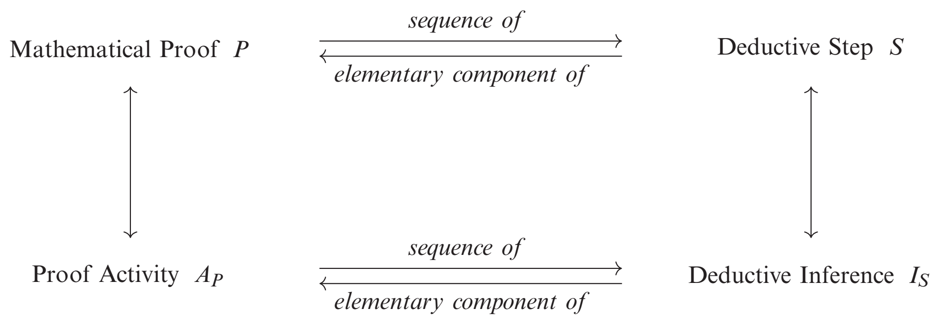

Fig. 1. Mathematical proof, proof activity, deductive step, deductive inference.

of the notion of mathematical proof in the realm of action. To this purpose, we shall introduce the notion of *proof activity*[4]: in the same way as a mathematical proof is a sequence of deductive steps, a *proof activity* is a sequence of *deductive inferences*. In this way, there is thus an *isomorphic* correspondence between a mathematical proof and its corresponding proof activity. From this perspective, what an agent needs to do to acquire knowledge of a theorem based on a mathematical proof of it is to carry out the corresponding proof activity.[5] Figure 1 provides a schematic representation of the notions just introduced, the vertical arrows representing the isomorphic correspondence just mentioned.

To understand the notion of proof activity, it might be useful to consider it in analogy with another familiar activity such as travelling. First, proof activities and travelling activities are both intended to bring the agent from one *state* to another. In the case of proof activities, the objective is to get from one epistemic state to another, while in the case of travelling activities, the objective is to get from one position state to another. Second, these two types of activities proceed through a sequence of *moves*—which are the *actions* constituting the building blocks of these activities—bringing the agent closer and closer to the desired state. For proof activities, these moves are deductive inferences, while for travelling activities these moves would involve walking, driving, flying, etc. Thus, in the same way as a travelling activity brings an agent from her current position state to a desired one through a sequence of spatial moves, a proof activity brings an agent from her current epistemic state to a desired one through a sequence of epistemic moves.

Now, an activity such as travelling most often requires a certain form of *planning*. If you want to go from point A to point B, you can always try to take a random sequence of moves using various means of transport, but this is unlikely to succeed. Rather, you will engage in careful planning to ensure that your various moves will bring you to your final destination. This will involve a sequence of *practical decisions* such as, for instance, deciding which bus(es), train(s), flight(s) to take so that the various connections go

[4] For reflections on proofs and constructions as activities carried out by rational human agents from the perspective of intuitionism, see van Benthem (2018).

[5] In this case, the written mathematical proof appears as a *recipe*—i.e., a set of *instructions*—for carrying out a proof activity. On the view of mathematical proofs as *recipes* or *blue-prints* see Sundholm (2012, p. 948) and Tanswell (2019). For an analysis of the relation(s) between text and action in mathematical and scientific scholarly texts, with a particular focus on *instructional* texts, see Chemla & Virbel (2015).

smoothly. Your travelling activity—the actual sequence of moves bringing you from point *A* to point *B*—would then be the direct result of planning, and as a consequence would possess a specific plan. It is in this sense that we would say, for instance, that *her* plan to go to Hawaii in the summer is to first take a shuttle to the local airport, then to take a flight to Honolulu International Airport, and finally to take a taxi to reach her final destination.

The exact same observation holds for proof activities. Carrying out a random sequence of deductive inferences is unlikely to bring the mathematician to an epistemic state in which she will know the desired theorem. Rather, to have any chance of reaching such an epistemic state, she will have to engage in some form of planning, as the mathematician Timothy Gowers (2002) pointed out: "What determines whether [mathematicians] can use their expertise to solve notorious problems is, in large measure, a matter of careful planning" (Gowers, 2002, p. 128). Of course, when searching for a sequence of deductive inferences leading to the desired theorem, the mathematician will encounter numerous dead ends, will engage in trial and error, and will often have to retreat from a given path, for instance after having explored and then given up a certain proof strategy. This happens as well in travelling activities, for instance when driving and realizing that the road you intended to take is closed for some reason, forcing you to go backward and take another road. When the mathematician has finally found a sequence of deductive inferences leading to the desired theorem, she will only report the *successful* sequence of moves leading to the final end in the written proof, and will omit all the twists and turns of the search process. But this search process would have involved careful planning all along, and so the final successful proof activity—the sequence of moves which indeed brought the mathematician to the desired epistemic state—will as well be the direct result of planning, and as a consequence will possess a specific plan. To come back to the quotes provided at the beginning, it is in this sense that we would say, for instance, that the plan of Faber & Pandharipande (2003) to prove the λ_g conjecture is to first prove it to be equivalent to a certain equation, then to show that the solution of this equation satisfies certain linear relations, and finally to show that these linear relations admit at most a one-dimensional solution space.

The main idea to be developed in this paper is that when we attribute a plan to a mathematical proof, we indeed attribute a plan to the *agent*(*s*)[6] who carried out the corresponding proof activity.[7] From this perspective, the written mathematical proof is

[6] In the following, we will usually drop the plural option, and only write 'agent' in the singular. One should keep in mind, however, that a mathematical proof might be produced by a *group* of agents, in which case 'agents' in the plural should be used.

[7] One might find it surprising that the notion of plan is attached here to the agent who produced the proof and not to the proof itself, but a moment of reflection reveals that this is a natural way of approaching the issue. To see this, imagine a mathematical proof that has been produced by a mathematician, and imagine that the exact same proof has been found as well by a computer through a random search. In this situation, we will be willing to say that there is a plan behind the proof produced by the mathematician, but we will be reluctant to attribute a plan to the proof produced by the computer. The reason, presumably, is that we would expect the mathematician to have reached the proof through some form of planning involving various practical decisions, while we know that there is no such thing behind the proof produced by the computer—that this proof was found through a random search was just "pure luck." Now, the key observation is that the proof is *exactly the same* in the two cases, but our inclination to attribute a plan to it differs depending on how and by whom we

nothing more than a report of this proof activity, recording the successful sequence of deductive inferences that the agent went through. An analysis of plans and planning in the context of mathematical proofs calls then for an analysis of plans and planning in the context of proof activities. It turns out that these notions of plans and planning have been extensively studied in the philosophy of action, most notably by the philosopher Michael Bratman who has developed over the years a full-fledged theory of planning agency. Our account builds directly on Bratman's theory of planning agency that we now turn to.

§3. Bratman on plans and planning agency. Bratman's theory of planning agency (Bratman, 1987) was originally developed to provide an account of the notion of *intention* as used in characterizing both mind and action—i.e., an account of what it means to *intend* to act as well as to act *intentionally*[8] —in which the notion of intention is conceived as being intimately connected with the phenomena of plans and planning. Central to the theory is the idea that intentions are elements of larger plans, in that they are the building blocks of plans, and for this reason are central to the functioning of our human agency:

> According to the planning theory, intentions of individuals are plan states: they are embedded in forms of planning central to our internally organized temporally extended agency and to our associated abilities to achieve complex goals across time, especially given our cognitive limitations. One's plan states guide, coordinate, and organize one's thought and action both at a time and over time. (Bratman, 2014, p. 15)

It is important to emphasize that a plan is not here to be understood as a mere recipe or procedure, and that having a plan to A does not simply mean possessing a recipe or procedure to A. Rather, a plan is to be conceived as a kind of *mental state* involving some sort of *commitment*: when an agent has a plan to A, this means that she indeed *plans* to A. In this sense, plans are "intentions writ large" (Bratman, 1987, p. 29), and possess some of the distinctive properties of intentions. According to Bratman, plans thus understood play a central role in our capacity for intrapersonal and extrapersonal coordinations, and permit our prior deliberations to influence our later conduct.

Because we are agents with *limited* cognitive resources, our plans are usually *partial* and *hierarchical*. Plans are partial in the sense that they are only partially specified at any given point in time. If you decide in the spring to go to Hawaii during the summer, you might not decide right away on the exact dates of your trip, on which means of transportation you will adopt, or on which hotel you are going to stay in—all of this would have to be specified eventually, but this can be done at a later time. Furthermore, plans are hierarchical in the sense that the intentions that comprise them can embed more specific intentions, which themselves might also embed more specific intentions,

expect the proof to have been produced. This speaks for attributing the notion of plan not to the proof itself but to the agent who produced it. For an account in which the notion of plan is attributed to the proof itself, see Bundy (1988) and Bundy et al. (2005)—this account is compared to our own in §8.

8 Bratman (1987) proceeds by first developing an account of what it means to intend to act, which he then uses to develop an account of what it means to act intentionally. In this work, we will be primarily concerned with the former notion.

and so on, and so forth. Thus, at a later stage, your plan to go to Hawaii during the summer might contain the intention to go there by plane, which itself might embed more specific intentions regarding the company you might use, the airport you might depart from, or the means of transport you might adopt to get to the airport.

The partiality of our plans generates, in turn, a demand for *practical reasoning*: since our plans are partial, they will need to be filled in appropriately as time goes by, practical reasoning is then required to specify our plans further whenever we need or wish to do so. Such practical reasoning takes as input one or several prior intentions in a given plan and yields as output one or several more specific intentions arranged in a *subplan*. For instance, while planning your trip to Hawaii, you might be led at some point to reason from your intention to fly to Hawaii to more specific intentions regarding, among others, the exact flight you will take as well as the class you will fly in. Our capacity for practical reasoning is thus essential for our plans to function properly in their coordinating role as well as in their capacity to shape later conduct, and for this reason lies at the heart of our planning agency.

In the next three sections, we will work toward an account of the notion of plan for proof activities by following the lead of Bratman's theory of planning agency. This means that we will proceed by first specifying the notions of intention and practical reasoning in the context of proof activities in §4 and §5, and we will then use those elements to specify the desired notion of plan for proof activities in §6.

§4. Intentions in proof activities. What are intentions in the particular context of proof activities? We will refer to them as *proving intentions*, and to figure out what they are we will begin by looking at an example of an actual proof activity. We propose to consider, as a running example, the proof activity corresponding to the following mathematical proof of the well-known fact that an equivalence relation on a nonempty set yields a partition of this set, a proof taken from the seventh edition of John B. Fraleigh's introductory book on abstract algebra (Fraleigh, 2003, pp. 7–8):[9]

THEOREM (Equivalence Relations and Partitions). *Let S be a nonempty set and let $\sim$ be an equivalence relation on S. Then $\sim$ yields a partition of S, where*

$$\bar{a} = \{x \in S \mid x \sim a\}.$$

Proof. We must show that the different cells $\bar{a} = \{x \in S \mid x \sim a\}$ for $a \in S$ do give a partition of S, so that every element of S is in some cell and so that if $\bar{a} \in \bar{b}$, then $\bar{a} = \bar{b}$. Let $a \in S$. Then $a \in \bar{a}$ by the reflexive condition (1), so a is in *at least one* cell.

Suppose now that a were in a cell $\bar{b}$ also. We need to show that $\bar{a} = \bar{b}$ as sets; this will show that a cannot be in more than one cell. There is a standard way to show that two sets are the same:

Show that each set is a subset of the other.

We show that $\bar{a} \subseteq \bar{b}$. Let $x \in \bar{a}$. Then $x \sim a$. But $a \in \bar{b}$, so $a \sim b$. Then, by the transitive condition (3), $x \sim b$, so $x \in \bar{b}$. Thus $\bar{a} \subseteq \bar{b}$. Now we show that $\bar{b} \subseteq \bar{a}$. Let $y \in \bar{b}$. Then $y \sim b$. But $a \in \bar{b}$, so $a \sim b$ and, by symmetry (2), $b \sim a$. Then by transitivity (3), y $\sim a$, so $y \in \bar{a}$. Hence $\bar{b} \subseteq \bar{a}$ also, so $\bar{b} = \bar{a}$ and our proof is complete. □

[9] This example is simple enough such that we can consider it here in its entirety, and yet rich enough with respect to its structure so that it can be used to illustrate the phenomena we are interested in. Hereafter, we will refer to this example as 'Fraleigh's proof'.

This particular example happens to be particularly informative for our present purpose, due to the fact that the author wrote down explicitly some of his proving intentions at different stages of his proof.[10] To highlight them, we have underlined in the above text all the phrases expressing proving intentions.[11] It is then noteworthy that the expressions of proving intentions occurring in this example are all written according to the same syntactical construction: they are all built from the phrase 'we show that'—sometimes with some modal variations—followed by a mathematical proposition.

This is to be expected. Any proof activity must necessarily begin with the intention to show—prove, establish—the theorem at hand,[12] and as the proof activity proceeds, this intention gives rise to more specific proving intentions, in the same way as your intention to go to Hawaii in the summer gives rise to more specific intentions regarding dates, means of transport, places to stay, etc. Thus, in the above example, the very first expression of a proving intention corresponds to the intention to show the theorem at hand, namely to show that the different cells $\bar{a} = \{x \in S \mid x \sim a\}$ for $a \in S$ do give a partition of S. This initial intention gives rise, as the proof activity proceeds, to more specific proving intentions, such as the intentions to show that $\bar{a} = \bar{b}$ as sets, to show that $\bar{a} \subseteq \bar{b}$, and to show that $\bar{b} \subseteq \bar{a}$. The (partial) plans described in the two quotes at the very beginning of this paper proceed in the exact same way.

Proving intentions of this type can be analyzed as a specification of the following schema:

$$\textbf{To show } C \ from \ [P_1, \dots, P_n],$$

where $P_1, \dots, P_n$ and C are placeholders for ordinary mathematical propositions[13] — $P_1, \dots, P_n$ and C will be referred to as the *hypotheses*[14] and the *conclusion* of the

[10] This feature is, of course, what motivated us to consider this example in the first place. The reasons why certain proving intentions have been made explicit in this particular example is, presumably, to serve some pedagogical purposes. In research contexts, explicit expressions of proving intentions are often used to structure the presentation of mathematical proofs. But in most cases, the reader is expected to figure out by herself the proving intentions of the author(s) at the various stages of the proof.

[11] This example shows that the text of a mathematical proof does *not* always consist in a mere list of deductive steps, but might contain some extra information as well. As we are seeing here, the analysis of such extra textual components might appear especially valuable in revealing particular epistemological dimensions of mathematical proofs.

[12] In practice, the proving intention the author is starting with can sometimes be more involved when the author intends to show the theorem in *such and such a way*, for instance by using—or avoiding the use of—certain concepts, techniques, methods, etc. This is particularly true when it comes to proving a theorem that has already been proved. These aspects can be integrated by refining the present framework. But insofar as they are not of primary relevance for the aim of this paper, a proper treatment of such refinements is left for another occasion.

[13] By an 'ordinary mathematical proposition' we mean a mathematical proposition from the vernacular language of ordinary mathematical practice.

[14] By the hypotheses of a proving intention, we mean here the hypotheses that pertain to the current *context of discourse*. For instance, the two hypotheses of the very first proving intention in Fraleigh's proof are "S is a nonempty set" and "$\sim$ is an equivalence relation on S." It is to be noted that the hypotheses of proving intentions are only rarely written explicitly in the texts of mathematical proofs, in the same way as the premisses of deductive steps are most often left implicit as well. Hypotheses of proving intentions, as well as premisses of deductive steps, must usually be inferred from the current context of discourse.

considered proving intention. We will refer to those as proving intentions of type *'to show'*. For convenience, we will often use the following more compact notation:

$$P_1, \dots, P_n \;\Rightarrow\; C.$$

Using this notation, the initial proving intention in Fraleigh's proof can be represented as:

$$S \neq \emptyset, \sim \text{ is an eq. rel. on } S \;\Rightarrow \sim \text{ yields a partition of } S \text{ where } \bar{a} = \{x \in S \mid s \sim a\},$$

while the proving intention expressed in the second paragraph can be represented as:[15]

$$S \neq \emptyset, \sim \text{ is an eq. rel. on } S,\; a \in S,\; a \in \bar{b} \;\Rightarrow\; \bar{a} = \bar{b}.$$

In general, expressions of proving intentions of type 'to show' can be identified in the text of a mathematical proof through the occurrence of what we might call a *proving verb*—e.g., 'we show that', 'we prove that', 'we establish that'—sometimes with some modal variations—e.g., 'we must show that', 'we need to show that', 'we will show that', 'we now show that'—and followed by a mathematical proposition.

We now introduce another type of proving intentions that we will refer to as proving intentions of type *'to infer'* and which correspond to intentions of carrying out deductive inferences. The particularity of proving intentions of this type is that they can be directly fulfilled by carrying out the corresponding deductive inferences, that is, by performing an *action*. This means that whenever an agent encounters a proving intention of the type 'to infer' in the execution of her plan, she will be led to carry out the corresponding deductive inference. We will use the following specific schema for analyzing proving intentions of type 'to infer':

$$\textbf{To infer } C \; from \; [P_1, \dots, P_n],$$

where $P_1, \dots, P_n$ and C are placeholders for ordinary mathematical propositions—$P_1, \dots, P_n$ and C will be referred to as the *premisses* and the *conclusion* of the considered proving intention. We will also use the compact notation introduced above to write proving intentions of type 'to infer'.

Proving intentions of types 'to show' and 'to infer' constitute the building blocks or atoms of plans for proof activities. We now turn to the issue of practical reasoning which concerns the step by step process by which plans for proof activities are constructed.

§5. Practical reasoning in proof activities. What does practical reasoning in the context of proof activities consist in? In Bratman's theory of planning agency, practical reasoning is required on the part of the agent to fill in her plans as time goes by, and proceeds by turning a given intention in the agent's plan into one or several more specific intentions arranged in a subplan. In the context of proof activities, *practical reasoning* is reasoning taking as input a given proving intention in the agent's plan

[15] The definitions associated to various terms in the hypotheses and the conclusion of a given proving intention are often present in the current context of discourse, and so should normally be listed as hypotheses. For making proving intentions more readable, we will usually not list the relevant definitions in the hypotheses of the proving intentions we are considering. Thus, in the following proving intention, we will not, for instance, list the definitions of $\bar{a}$ and $\bar{b}$, although those are part of the current context of discourse. One should, nonetheless, keep in mind that this is just a way to ease presentation, and that the relevant definitions *are* part of the hypotheses.

and yielding as output one or several more specific proving intentions arranged in a subplan. Practical reasoning is thus the central process by which plans are constructed: each construction step of a plan is *always* the result of an instance of practical reasoning.

Interestingly, in the example of Fraleigh's proof, not only has the author mentioned explicitly some of his proving intentions at different stages of his proof, he has also written explicitly in the text some of the key steps of practical reasoning underlying his proof activity. The first example occurs at the very beginning of the proof where the author tells us that he intends to show that "the different cells $\bar{a} = \{x \in S \mid x \sim a\}$ for $a \in S$ do give a partition of S," an intention that he immediately transforms into the further intention to show that "every element of S is in some cell and so that if $a \in \bar{b}$, then $\bar{a} = \bar{b}$." This instance of practical reasoning takes as input and yields as output the following proving intentions:

Input: $S \neq \emptyset, \sim$ is an eq. rel. on $S \Rightarrow \sim$ yields a partition of S where $\bar{a} = \{x \in S \mid s \sim a\}$,

Output: $S \neq \emptyset, \sim$ is an eq. rel. on $S \Rightarrow$ for all $a \in S$, a is in some cell and if $a \in \bar{b}$, then $\bar{a} = \bar{b}$.

Another example occurs later on in the proof, when the author expresses the intention to show that " $\bar{a} = \bar{b}$ as sets." We can see that this proving intention is then transformed into two more specific intentions, namely to show that "$\bar{a} \subseteq \bar{b}$" and to show that "$\bar{b} \subseteq \bar{a}$." This instance of practical reasoning takes as input and yields as output the following proving intentions:

Input: $S \neq \emptyset, \sim$ is an eq. rel. on S, $a \in S$, $a \in \bar{b} \Rightarrow \bar{a} = \bar{b}$,

Output: $S \neq \emptyset, \sim$ is an eq. rel. on S, $a \in S$, $a \in \bar{b} \Rightarrow \bar{a} \subseteq \bar{b}$,
$S \neq \emptyset, \sim$ is an eq. rel. on S, $a \in S$, $a \in \bar{b} \Rightarrow \bar{a} \supseteq \bar{b}$.

These two examples correspond to very basic forms of practical reasoning, but of course practical reasoning in proof activities can be much more involved. We will not attempt in this paper to provide a full specification of what practical reasoning consists in in the context of proof activities, as this would simply amount to offering a full-blown theory of proof discovery. We will nonetheless isolate one basic form of practical reasoning which is widely used in practice and which is likely to be present in the plan of almost any mathematical proof. We will then illustrate the fact that practical reasoning can be much more advanced by presenting an example of a more sophisticated instance of practical reasoning taken from Pólya (1949).

5.1. A simple form of practical reasoning in proof activities. One simple form of practical reasoning proceeds through the application of *mathematical methods*. In the present context, a mathematical method is defined as a *rule* that transforms a proving intention into one or several other proving intentions arranged in a subplan, and which, at least in some cases, can lead to a plan that can be successfully executed. Because mathematical methods are rules, most mathematical methods can be entirely characterized by a *schema* specifying how the rule transforms its input into its output.

More specifically, the schema of a mathematical method can be put into the following form:[16]

To show C, *from* $[P_1, \dots, P_n]$:

$$\begin{array}{lll} 1. & \textbf{Show/Infer } C_1, & \textit{from } [P_{1,1}, \dots, P_{1,n_1}], \\ 2. & \textbf{Show/Infer } C_2, & \textit{from } [P_{2,1}, \dots, P_{2,n_2}], \\ \vdots & & \\ k. & \textbf{Show/Infer } C_k, & \textit{from } [P_{k,1}, \dots, P_{k,n_k}], \end{array}$$

where all the occurrences of Cs and Ps—with or without subscripts—are placeholders for ordinary mathematical propositions.

From this perspective, the first example of practical reasoning in Fraleigh's proof mentioned above can be conceived as an application of the mathematical method consisting in replacing a term by its definition within the mathematical proposition one intends to show, a method that can be characterized by the following schema[17]:

To show $X[\Phi]$, *from* $[P_1, \dots, P_n]$:

$$\begin{array}{ll} 1. \textbf{ Show } X[\Phi_{\mathsf{Def}}/\Phi], & \textit{from } [P_1, \dots, P_n], \\ 2. \textbf{ Infer } X[\Phi], & \textit{from } X[\Phi_{\mathsf{Def}}/\Phi] \textit{ and } \Phi_{\mathsf{Def}}, \end{array}$$

where Φ is a placeholder for a mathematical term, Φ_{Def} is a placeholder for the mathematical definition of this term, and $X[\Phi_{\mathsf{Def}}/\Phi]$ is the mathematical proposition resulting from X by replacing every occurrence of Φ in X by its definition Φ_{Def}. The second example of practical reasoning in Fraleigh's proof can also be conceived as an application of a mathematical method, this time of the method consisting in showing that two sets are equal by showing that each is a subset of the other, a method that can be characterized by the following schema[18]:

To show $A = B$, *from* $[P_1, \dots, P_n]$:

$$\begin{array}{ll} 1. \textbf{ Show } A \subseteq B, & \textit{from } [P_1, \dots, P_n], \\ 2. \textbf{ Show } B \subseteq A, & \textit{from } [P_1, \dots, P_n], \\ 3. \textbf{ Infer } A = B, & \textit{from } A \subseteq B \textit{ and } B \subseteq A, \end{array}$$

where A and B are placeholders for expressions denoting sets.[19] We shall now make two important remarks on the notion of mathematical method.

---

16 The output of a mathematical method is a subplan, and as such possesses a *tree structure*. For expository reasons, we will only consider in the following mathematical methods yielding subplans of depth one, which can then be represented as *lists*. All that we say in the following about mathematical methods can be straightforwardly adapted to handle cases of mathematical methods yielding subplans of depth greater than one.

17 When we reported earlier the input and output of this instance of practical reasoning, we did not mention the intention of type 'to infer' that we are now listing as the second item in the schema. The reason for this is that this proving intention was not mentioned explicitly in the text of Fraleigh's proof, nor was the deductive step corresponding to it. Yet, such a deductive step *must* be present for completing the proof activity, although it was left implicit by Fraleigh. This is why it must be listed in the schema of this mathematical method.

18 The remark in footnote 17 also holds for the schema of this mathematical method.

19 Notice that, in this particular case, it does not matter in which order steps 1 and 2 are carried out. In our definition, a mathematical method will, however, always yield a subplan whose order of execution of its steps is fully determined. It is, nonetheless, possible to represent mathematical methods for which the order of (some of) its steps does not matter through an equivalence class of equivalent mathematical methods in the sense defined here.

First, mathematical methods come at different levels of *generality* in terms of their *range of application*. In this respect, the two examples of mathematical method just discussed can be applied widely to many fields of mathematics. This is most often the case for mathematical methods of a logical or set-theoretical nature. For instance, the following basic mathematical method of conjunction introduction can be applied to any proving intention fitting its input schema:

To show $X \wedge Y$, *from* $[P_1, \dots, P_n]$:

1. **Show** X,	*from* $[P_1, \dots, P_n]$,
2. **Show** Y,	*from* $[P_1, \dots, P_n]$,
3. **Infer** $X \wedge Y$,	*from* X *and* Y,

where X and Y are placeholders for ordinary mathematical propositions. Other mathematical methods have a much more restricted range of application. To see this, it suffices to consider the well-known method of mathematical induction for natural numbers, which can be formulated as follows:

To show $\forall n H(n)$, *from* $[P_1, \dots, P_n]$:

1. **Show** $H(0)$,	*from* $[P_1, \dots, P_n]$,
2. **Show** $\forall p(H(p) \to H(p+1))$,	*from* $[P_1, \dots, P_n]$,
3. **Infer** $\forall n H(n)$,	*from* $H(0)$ *and* $\forall p(H(p) \to H(p+1))$,

where $H(\_)$ is a placeholder for an expression taking as argument another expression denoting a concrete or arbitrary natural number. Such a mathematical method can only be applied within the context of number theory. Similarly, the following mathematical method from Fraleigh (2003, p. 132) to show that a map between two groups is an isomorphism[20] can only be applied within the context of group theory:

To show $\phi : G \to G'$ is an isomorphism, *from* $[P_1, \dots, P_n]$:

1. **Show** ϕ is a homomorphism, [21]	*from* $[P_1, \dots, P_n]$,
2. **Show** $\text{Ker}(\phi) = \{e\}$, [22]	*from* $[P_1, \dots, P_n]$,
3. **Show** ϕ maps G onto G',	*from* $[P_1, \dots, P_n]$,
4. **Infer** $\phi : G \to G'$ is an isomorphism,	*from* ϕ is a homomorphism *and* $\text{Ker}(\phi) = \{e\}$ *and* ϕ maps G onto G' *and* theorem T[23] *and* the definition of an isomorphism,

where G and G' are both placeholders for the name of a group, ϕ is a placeholder for the name of a map of the first group into the second, and e is a placeholder for the name of the identity element of the first group. The question of the generality of

20 A map ϕ of a group $\langle G, \circ \rangle$ into a group $\langle G', * \rangle$ is an *isomorphism* if ϕ is a one-to-one function mapping G onto G' such that $\phi(a \circ b) = \phi(a) * \phi(b)$ for all $a, b \in G$.

21 A map ϕ of a group $\langle G, \circ \rangle$ into a group $\langle G', * \rangle$ is a *homomorphism* if $\phi(a \circ b) = \phi(a) * \phi(b)$ for all $a, b \in G$.

22 When a map ϕ of a group $\langle G, \circ \rangle$ into a group $\langle G', * \rangle$ is a homomorphism, the *kernel* of ϕ is defined by $\text{Ker}(\phi) = \{x \in G \mid \phi(x) = e'\}$ where e' is the identity element of $\langle G', * \rangle$.

23 Theorem T states that a homomorphism ϕ of a group $\langle G, \circ \rangle$ into a group $\langle G', * \rangle$ is a *one-to-one* map if and only if $\text{Ker}(\phi) = \{e\}$. Theorem T is called Corollary 13.18 in Fraleigh (2003, p. 131).

mathematical method is an important one which has already received some attention in the philosophy of mathematics, most notably from Avigad (2006) and Rav (1999).

Second, mathematical methods can proceed *forward*, *backward*, or in a *mixed fashion*. A mathematical method proceeds *forward* if all the proving intentions of type 'to show' it yields as output only differ from the proving intention it takes as input with respect to their *hypotheses*. A mathematical method proceeds *backward* if all the proving intentions of type 'to show' it yields as output only differ from the proving intention it takes as input with respect to their *conclusion*. A mathematical method proceeds in a *mixed fashion* if it does not proceed either forward or backward. The basic idea behind this terminology, which is common in interactive and automated theorem proving, is the following: a mathematical method proceeds forward when it makes a forward step *toward* the conclusion one intends to show by extending the current set of hypotheses; a mathematical method proceeds backward when it makes a backward step *from* the conclusion one intends to show by transforming it into something that would hopefully be simpler to show. All the examples of mathematical methods that we have seen so far happen to proceed backward, but this is not always the case. Consider for instance the stage in Fraleigh's proof that the agent finds herself at right after processing the phrase "Let $x \in \bar{a}$" in the third paragraph of the proof. At this stage, the proving intention that is considered as input into further practical reasoning is the following:

$$S \neq \emptyset, \sim \text{ is an eq. rel. on } S, a \in S, a \in \bar{b}, x \in \bar{a} \;\Rightarrow\; x \in \bar{b}.$$

Although the output of the practical reasoning is not written explicitly in the text of Fraleigh's proof, we can see from the way the proof proceeds that the output is given by the following subplan:

1. **Infer** $x \sim a$, *from* $x \in \bar{a}$ *and* $\bar{a} = \{x \in S \mid s \sim a\}$,
2. **Show** $x \in \bar{b}$, *from* $[S \neq \emptyset, \sim$ eq. rel. on $S, a \in S, a \in \bar{b}, x \in \bar{a}, x \sim a]$.

The execution of this subplan leads the agent to first carry out the deductive inference with conclusion "$x \sim a$" and premisses "$x \in \bar{a}$" and "$\bar{a} = \{x \in S \mid s \sim a\}$," and then to consider the subsequent proving intention. This instance of practical reasoning can be conceived as an application of the mathematical method sometimes known as the 'principle of abstraction', a method that can be characterized by the following schema:

To show C, *from* $[P_1, \dots, P_n, x \in A]$:

1. **Infer** $Q(x)$, *from* $x \in A$ *and* $A = \{x \mid Q(x)\}$,
2. **Show** C, *from* $[P_1, \dots, P_n, x \in A, Q(x)]$.

Many other mathematical methods proceed forward. One of the most illustrative examples is the mathematical method that consists in transforming one of the hypotheses by replacing one of the terms occurring in it by its definition, a method that can be characterized by the following schema:

To show C, *from* $[P_1, \dots, P_n]$:

1. **Infer** $P_i[\Phi_{\mathsf{Def}}/\Phi]$, *from* $P_i[\Phi]$ *and* Φ_{Def},
2. **Show** C, *from* $[P_1, \dots, P_n, P_i[\Phi_{\mathsf{Def}}/\Phi]]$,

where Φ is a placeholder for a mathematical term, Φ_{Def} is a placeholder for the mathematical definition of this term, and $P_i[\Phi_{\mathsf{Def}}/\Phi]$ is the mathematical proposition resulting from P_i by replacing every occurrence of Φ in X by its definition Φ_{Def}. Finally, all the mathematical methods that correspond to a (forward) logical rule of inference proceed forward, as for instance the mathematical method corresponding to an application of *modus ponens*:

To show C, *from* $[P_1, \ldots, P_n, X, X \to Y]$:

1. **Infer** Y, *from* X *and* $X \to Y$,
2. **Show** C, *from* $[P_1, \ldots, P_n, X, X \to Y, Y]$.

It is interesting to notice that, in Fraleigh's proof, the mathematical methods involved sometimes proceed forward and sometimes backward. This is extremely common, and most ordinary mathematical proofs indeed involve a mixture of mathematical methods proceeding forward, backward, and in a mixed fashion.[24]

It is often the case that several different mathematical methods can be applied to the same proving intention.[25] In these situations, before a proving intention can be transformed into further proving intentions by the application of a mathematical method, it must be decided which mathematical method is to be applied to it. To account for this, we introduce a dedicated epistemic capacity that we refer to as the *resolution expertise*, and that we define as follows: the *resolution expertise* is a process that takes as input a proving intention in the agent's plan and yields as output the mathematical method in the agent's *toolbox* that the agent judges to have the best chance of success at fulfilling the proving intention taken as input. The agent's *toolbox* is here defined as the collection of all the mathematical methods that the agent knows.

Two important remarks are in order with respect to the notion of resolution expertise. First, it is important to notice that the resolution expertise does not have to consider all the possible mathematical methods in the agent's toolbox while deciding which method to apply to a given proving intention: (i) it is, of course, not necessary to consider the mathematical methods for which the considered proving intention is not an instance of their input schema, and (ii) it is not necessary to consider the mathematical methods for which the agent knows that they cannot successfully lead to a fulfilling of the considered proving intention. Second, we will not attempt in this work to explain how the resolution expertise proceeds, i.e., by which mechanisms an agent can decide which mathematical method has the better chances of successfully fulfilling a given proving intention in the context of a proof activity. This is a highly complicated matter which touches directly on the issue of problem-solving and proof discovery in mathematics.[26]

---

24 We did not provide a specific example of a mathematical method proceeding in a mixed fashion. One can easily build such an example by simply considering the mathematical method resulting from the combination of two mathematical methods proceeding respectively forward and backward.

25 For instance, faced with a proving intention whose conclusion is a universal proposition of the form $\forall n P(n)$, an agent can transform it using the method of mathematical induction we have already seen, but can also transform it into two proving intentions using various forms of case splitting—e.g., the intention to establish it for the even numbers on the one hand, and for the odd numbers on the other.

26 In the development of their prover, Ganesalingam & Gowers (2017) have been led to implement mechanisms that do exactly what the resolution expertise is supposed to achieve. In particular, Ganesalingam & Gowers (2017) have faced the very same issue of deciding which tactic to apply to a given goal, when multiple tactics are available. Their solution is discussed in §5.10 of their paper. From our perspective, any account of how to choose which tactic or method to apply to a given goal or proving intention which purports to be faithful to how ordinary mathematical agents proceed—which is the case for the one

To sum up, one simple form of practical reasoning consists in applying a mathematical method to a given proving intention. In case several mathematical methods in the agent's toolbox can be applied to the same proving intention, the agent needs to choose which method to apply, a decision which requires a dedicated epistemic capacity that we have called the resolution expertise. This form of practical reasoning is widely used in practice and is to be found in the plan of almost any mathematical proof.

5.2. A more sophisticated form of practical reasoning in proof activities. The simple form of practical reasoning just discussed is often qualified as 'routine' in mathematical practice. But many cases of practical reasoning do not proceed this way, and will often turn out to be much more involved. To illustrate this, we will now present a more sophisticated instance of practical reasoning that comes from an example previously discussed by Pólya (1949) and which concerns a mathematical proof of Carleman's inequality.

Carleman's inequality and its proof are given below as they appear in Pólya (1949, pp. 684–685):

THEOREM *If the terms of the sequence* a_1, a_2, a_3, ... *are non-negative real numbers, not all equal to* 0, *then*

$$\sum_{n=1}^{\infty}(a_1a_2a_3\dots a_n)^{1/n} < e\sum_{n=1}^{\infty}a_n.$$

Proof. Define the numbers c_1, c_2, c_3, ... by

$$c_1c_2c_3\dots c_n = (n+1)^n$$

for $n = 1, 2, 3, \dots$ We use this definition, then the inequality between the arithmetic and the geometric means, and finally the fact that the sequence defining e, the general term of which is $[(k+1)/k]^k$, is increasing. We obtain

$$\begin{aligned}
\sum_{n=1}^{\infty}(a_1a_2\dots a_n)^{1/n} &= \sum_{n=1}^{\infty}\frac{(a_1c_1a_2c_2\dots a_nc_n)^{1/n}}{n+1}\\
&\le \sum_{n=1}^{\infty}\frac{a_1c_1+a_2c_2+\cdots+a_nc_n}{n(n+1)}\\
&= \sum_{k=1}^{\infty}a_kc_k\sum_{n\ge k}\frac{1}{n(n+1)}\\
&= \sum_{k=1}^{\infty}a_kc_k\sum_{n=k}^{\infty}\left(\frac{1}{n}-\frac{1}{n+1}\right)\\
&= \sum_{k=1}^{\infty}a_k\frac{(k+1)^k}{k^{k-1}}\frac{1}{k}\\
&< e\sum_{k=1}^{\infty}a_k.
\end{aligned}$$

□

proposed by Ganesalingam & Gowers (2017)—would constitute a potential candidate for specifying further the notion of resolution expertise.

https://doi.org/10.1017/S1755020319000601 Published online by Cambridge University Press

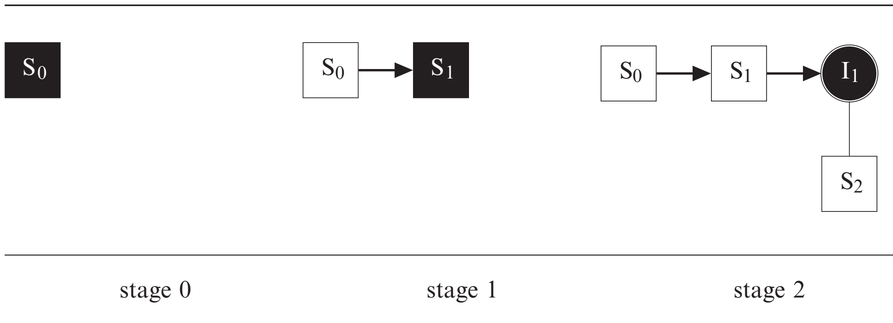

Fig. 2. Stage by stage progression for the initial stages of planning underlying Pólya's proof.

We will focus here on the initial stages of planning that underlie the proof activity associated with this proof, a graphical representation of which is provided in Figure 2.

As usual, the agent's plan at the beginning of the proof activity is composed of a single proving intention, which is the intention to show the theorem at hand:

$$\mathrm{S}_0: \begin{array}{l} a_1,\ a_2,\ a_3, \ldots \text{ are non-negative} \\ \text{real numbers,} \\ a_1,\ a_2,\ a_3, \ldots \text{ are not} \\ \text{all equal to } 0 \end{array} \quad \Rightarrow \quad \sum_{n=1}^{\infty}(a_1 a_2 a_3 \ldots a_n)^{1/n} < e \sum_{n=1}^{\infty} a_n.$$

This intention of type 'to show' is then transformed into a further intention by including an additional hypothesis—the definition of the c_i sequence—thus leading to the following subplan:

1. (S_1) **Show** $\sum_{n=1}^{\infty}(a_1 a_2 a_3 \ldots a_n)^{1/n} < e \sum_{n=1}^{\infty} a_n$, *from*
   - $a_1, a_2, a_3, \ldots$ are non-negative real numbers,
   - $a_1, a_2, a_3, \ldots$ are not all equal to 0,
   - $c_1 c_2 c_3 \ldots c_n = (n+1)^n$ for $n = 1, 2, 3, \ldots$

This proving intention S_1 is itself transformed into further intentions through the generation of the following subplan:

1. (I_1) **Infer** $\sum_{n=1}^{\infty}(a_1 a_2 \ldots a_n)^{1/n} = \sum_{n=1}^{\infty} \frac{(a_1 c_1 a_2 c_2 \ldots a_n c_n)^{1/n}}{n+1}$, *from*
   - $c_1 c_2 c_3 \ldots c_n = (n+1)^n$ for $n = 1, 2, 3, \ldots$
2. (S_2) **Show** $\sum_{n=1}^{\infty}(a_1 a_2 a_3 \ldots a_n)^{1/n} < e \sum_{n=1}^{\infty} a_n$, *from*
   - $a_1, a_2, a_3, \ldots$ are non-negative real numbers,
   - $a_1, a_2, a_3, \ldots$ are not all equal to 0,
   - $c_1 c_2 c_3 \ldots c_n = (n+1)^n$ for $n = 1, 2, 3, \ldots,$
   - $\sum_{n=1}^{\infty}(a_1 a_2 \ldots a_n)^{1/n} = \sum_{n=1}^{\infty} \frac{(a_1 c_1 a_2 c_2 \ldots a_n c_n)^{1/n}}{n+1}$.

https://doi.org/10.1017/S1755020319000601 Published online by Cambridge University Press

The two transformations just described correspond to the two instances of practical reasoning underlying the moves from stage 0 to stage 1 and from stage 1 to stage 2 as represented in Figure 2. These are the two initial construction steps of the plan underlying Pólya's proof.

The interesting instance of practical reasoning here is the one behind the transition from stage 0 to stage 1 in Figure 2, that is the one with input S_0 and with output S_1. A distinctive aspect of this transformation is the introduction of the c_i sequence, which does not appear to be a 'routine' step. Indeed, one might really wonder in this case where this c_i sequence comes from, that is, how the author of the proof came up with this transformation of his initial proving intention. Pólya (1949) points out that this step is very likely to appear to the reader as a 'deus ex machina'.[27] He then provides a detailed reconstruction of the thinking process that led to the introduction of the c_i sequence. From our perspective, Pólya is explaining the instance of practical reasoning which consists in transforming S_0 into S_1, and which is the very first step in the construction of the plan underlying his proof. We will now summarize this instance of practical reasoning, drawing on Hardy et al. (1934), Hardy, Littlewood, & Pólya (1949), and Steele (2004, pp. 27–30).

Pólya begins with the proving intention S_0. To make progress, he considers what he knows that could be helpful. One theorem he knows is the Arithmetic-Geometric Mean Inequality, which states that for every sequence of non-negative real numbers $b_1, b_2, \dots, b_n$, $(b_1 b_2 \dots b_n)^{1/n} \leq (b_1 + b_2 + \cdots + b_n)/n$, with equality if and only if $b_1 = b_2 = \cdots = b_n$. As one of the hypotheses in S_0 is that $a_1, a_2, \dots, a_n$ is a sequence of non-negative reals, the Arithmetic-Geometric Mean Inequality can be applied to the a_i sequence. Moreover, the summand in the expression $\sum_{n=1}^{\infty} (a_1 a_2 \dots a_n)^{1/n}$ has the same form as the expression on the left hand side of the Arithmetic-Geometric Mean Inequality. So applying it seems to be a good idea, and when he does so, he obtains the following:

$$\sum_{n=1}^{\infty}(a_1 \dots a_n)^{1/n} \leq \sum_{n=1}^{\infty} \frac{a_1 + \cdots + a_n}{n} = \sum_{k=1}^{\infty} a_k \sum_{n=k}^{\infty} \frac{1}{n}.$$

As $\sum \frac{1}{n}$ diverges, however, this is a dead-end. Nevertheless, he can make further progress by analyzing why the proof attempt failed. There is an important heuristic to follow when attempting to prove one inequality by applying another: try applying the inequality when it is (nearly) sharp (Steele, 2004, p. 26). In Pólya's aborted proof attempt, however, the Arithmetic-Geometric Mean Inequality is not applied where it is (nearly) sharp. This is because the inequality is sharp when and only when the terms it is applied to are equal, yet the terms of the a_i sequence can be "very unequal" (Pólya, 1949, p. 687). This is what caused the proof attempt to fail.

To try to fix the failed proof attempt, Pólya thus needs to make the terms of the sequence he applies the Arithmetic-Geometric Mean Inequality to more equal. One way he can do this is by introducing an auxiliary c_i sequence to multiply the a_i sequence by. For this to work, c_n should be asymptotically proportional to n when $\sum_{n=1}^{\infty} a_n$ is near the boundary of convergence (Hardy et al., 1934, p. 249).

[27] Proofs containing deus ex machinas are often said to be "unmotivated" (in fact, the title of Pólya's paper was "With, or without, motivation?"). Recently Morris (2020) has built on Pólya's work to offer an analysis of motivated proofs.

To see what other constraints there may be on the choice of the auxiliary sequence and to obtain a more precise definition of it, Pólya starts a new proof attempt involving an indeterminate c_i sequence. First, he rewrites $\sum_{n=1}^{\infty}(a_1 a_2 \ldots a_n)^{1/n}$ so that it involves the $c_i a_i$ sequence rather than just the a_i sequence. In particular, he obtains:

$$\sum_{n=1}^{\infty}(a_1 a_2 \ldots a_n)^{1/n} = \sum_{n=1}^{\infty}\left(\frac{c_1 a_1 c_2 a_2 \ldots c_n a_n}{c_1 c_2 \ldots c_n}\right)^{1/n}.$$

Next he tries applying the Arithmetic-Geometric Mean Inequality again and finds (Hardy et al., 1934, p. 249):

$$\begin{aligned}\sum_{n=1}^{\infty}\left(\frac{c_1 a_1 c_2 a_2 \ldots c_n a_n}{c_1 c_2 \ldots c_n}\right)^{1/n} &\leq \sum_{n=1}^{\infty}(c_1 c_2 \ldots c_n)^{-1/n}\frac{1}{n}\sum_{k\leq n} c_k a_k \\ &= \sum_{k=1}^{\infty} a_k c_k \sum_{n\geq k}\frac{1}{n}(c_1 c_2 \ldots c_n)^{-1/n}.\end{aligned}$$

To make further progress, he must find a way to evaluate the expression

$$c_k \sum_{n\geq k}\frac{1}{n}(c_1 c_2 \ldots c_n)^{-1/n}.$$

Again he considers what he knows that might be helpful. He recalls that the sum of a telescoping series can be evaluated, in particular $\sum_{i=k}^{\infty}\left(\frac{1}{b_i}-\frac{1}{b_{i+1}}\right)=\frac{1}{b_k}$. So if he sets $c_1 c_2 c_3 \ldots c_n = (n+1)^n$ for $n = 1, 2, \ldots, n$, he can evaluate $c_k \sum_{n\geq k}\frac{1}{n}(c_1 c_2 \ldots c_n)^{-1/n}$ and find it to be equal to $\frac{c_k}{k}$. Pólya has thus found a way of defining the c_i sequence that helps him make further progress with the proof attempt. So he now adds this definition of the c_i sequence as a new hypothesis. This is how the initial proving intention S_0 is turned into the proving intention S_1.

This example illustrates that practical reasoning in the context of proof activities can be highly sophisticated, involving trial and error, backtracking, changes of strategies, etc. Providing a detailed account of such higher forms of practical reasoning in proof activities goes far beyond the scope of this paper. The main message we want to convey here is this: however complex this process of practical reasoning might be, the plan of a proof activity is *always* going to be the result of a step by step construction where each step consists in an instance of practical reasoning, that is, in the transformation of a given proving intention in the agent's plan into one or several more specific intentions arranged in a subplan.

§6. Plans in proof activities. We now have all the elements to define the notion of plan for proof activities: an agent's *plan* for a proof activity is an *ordered tree*[28] such that (1) each *node* is a proving intention, (2) the *root* is the proving intention corresponding to the theorem at hand, and (3) each *set of ordered children* of a given parent node is a subplan that has been obtained from the parent node through an instance of practical reasoning. The *execution* of a plan for a proof activity by a given agent leads to the

28 In the mathematical sense of the term: an *ordered tree* is a rooted tree where each node comes equipped with an ordering of its children.

actual *realization* of the proof activity, that is, to the actual *carrying out* of the sequence of deductive inferences—the actions—constituting the proof activity. As we mentioned earlier, the written mathematical proof constitutes then a *report* of the corresponding proof activity, where the order of the deductive steps as they figure in the mathematical proof corresponds exactly to the order in which the deductive inferences are carried out in the corresponding proof activity.

The notion of plan for proof activities is best illustrated with an example. The most convenient way to describe the plan of a proof activity is through its *dynamics*, by presenting step by step the intertwined process of its construction and execution. We will now do so in some detail for the first steps of this process in the case of the proof activity corresponding to Fraleigh's proof—Figure 3 will accompany our description with a graphical representation of this step by step progression.

As we already saw, any proof activity begins with the intention to show the theorem at hand. At the initial stage of the proof activity corresponding to Fraleigh's proof, the agent's plan is thus constituted by the following single proving intention:

$$\mathrm{S}_0: \quad S \neq \emptyset,\ \sim \text{ is an eq. rel. on } S \quad \Rightarrow \quad \begin{array}{c} \sim \text{ yields a partition of } S \\ \text{where } \bar{a} = \{x \in S \mid s \sim a\}. \end{array}$$

This proving intention of type 'to show' is then transformed into further proving intentions through an application of the mathematical method consisting in replacing a term by its definition within the mathematical proposition one intends to show. This yields the following subplan:

1. (S_1)	**Show** for all $a \in S$, a is in some cell and if $a \in \bar{b}$, then $\bar{a} = \bar{b}$,	*from* [$S \neq \emptyset$, $\sim$ is an eq. rel. on S],
2. (I_{17})	**Infer** $\sim$ yields a partition of S where $\bar{a} = \{x \in S \mid s \sim a\}$,	*from* for all $a \in S$, a is in some cell and if $a \in \bar{b}$, then $\bar{a} = \bar{b}$ *and* the definition of a partition of a set.

The execution of the plan now leads the agent to address the proving intention S_1, which corresponds to stage 1 in Figure 3. The proving intention S_1 is, in turn, transformed into further intentions through a combination of the mathematical methods of conjunction introduction and universal quantifier introduction, yielding the following subplan:

1. (S_2)	**Show** a is in some cell,	*from* [$S \neq \emptyset$, $\sim$ is an eq. rel. on S, $a \in S$],
2. (S_3)	**Show** if $a \in \bar{b}$, then $\bar{a} = \bar{b}$,	*from* [$S \neq \emptyset$, $\sim$ is an eq. rel. on S, $a \in S$],
3. (I_{16})	**Infer** for all $a \in S$, a is in some cell and if $a \in \bar{b}$, then $\bar{a} = \bar{b}$,	*from* $a \in S$ *and* a is in some cell *and* if $a \in \bar{b}$, then $\bar{a} = \bar{b}$.

The execution of the plan now leads the agent to address the proving intention S_2, which corresponds to stage 2 in Figure 3. This time, a forward mathematical method is applied. Its first step is a proving intention 'to infer' corresponding to the chaining of first deducing the mathematical proposition "$a \sim a$" from the reflexive condition in the definition of $\sim$, and then deducing the mathematical proposition "$a \in \bar{a}$" from the definition of $\bar{a}$. Its second step is a proving intention 'to show' which is obtained from S_2 by adding the deduced mathematical proposition "$a \in \bar{a}$" to the set of hypotheses of S_2:

1. (I_1) **Infer** $a \in \bar{a}$, *from* the reflexive condition in the definition of $\sim$ *and* the definition of $\bar{a}$,
2. ($S_{2,1}$) **Show** a is in some cell, *from* [$S \neq \emptyset, \sim$ is an eq. rel. on $S, a \in S, a \in \bar{a}$].

The execution of the plan leads the agent to address the proving intention I_1—we are now in stage 3 in Figure 3—which is fulfilled by carrying out the corresponding deductive inference. This brings the agent to stage 4, where she is led to address the proving intention $S_{2,1}$. This proving intention 'to show' is then transformed into the following subplan constituted of a single proving intention 'to infer', which corresponds to the application of a forward mathematical method which consists in chaining an exploitation of the definition of a cell, a conjunction introduction, and an existential introduction:

1. (I_2) **Infer** a is in at least one cell, *from* $a \in \bar{a}$ *and* the definition of a cell.

The execution of the updated plan leads the agent to address the proving intention I_2—we are here in stage 5 in Figure 3—which is fulfilled by carrying out the corresponding deductive inference. The agent is then brought to stage 6 in which she is led to address the proving intention S_3. This is a complete description of the first seven stages of planning underlying the proof activity corresponding to Fraleigh's proof. In the appendix, we provide a full analysis of this example, with a list of the deductive steps constituting Fraleigh's proof, a list of all the proving intentions of type 'to show' involved in this example, and a graphical representation of all the stages involved in the planning underlying the corresponding proof activity. The complete execution of this plan leads then to the realization of the entire proof activity corresponding to Fraleigh's proof.[29]

This completes our account of plans and planning in the context of proof activities. Insofar as we conceived of the plan of a mathematical proof as the plan of the agent who carried out the corresponding proof activity, this yields our desired account of plans and planning in the context of mathematical proofs. In the next two sections, we will see how our notion of plan is connected to the theme of modularity put forward by Avigad (2020), and we will then compare our account to the framework of proof planning developed by Alan Bundy and colleagues (Bundy, 1988; Bundy et al., 2005) in the field of automated theorem proving.

§7. Plans and modularity in the presentation of mathematical proofs. Our notion of plan is intimately connected to the *modularity* of mathematical proofs in the sense of Avigad (2020), i.e., the idea that mathematical proofs can be seen as being "divided into

[29] We chose here Fraleigh's proof as a simple example to illustrate the functioning of our framework. But the more complex a proof is, the more interesting the plan becomes. We invite the reader who wants to get a sense of what plans for mathematical proofs are to work out for him or herself the first stages of planning of one of his or her favorite proofs, and to thereby reconstruct at least part of the plan underlying it.

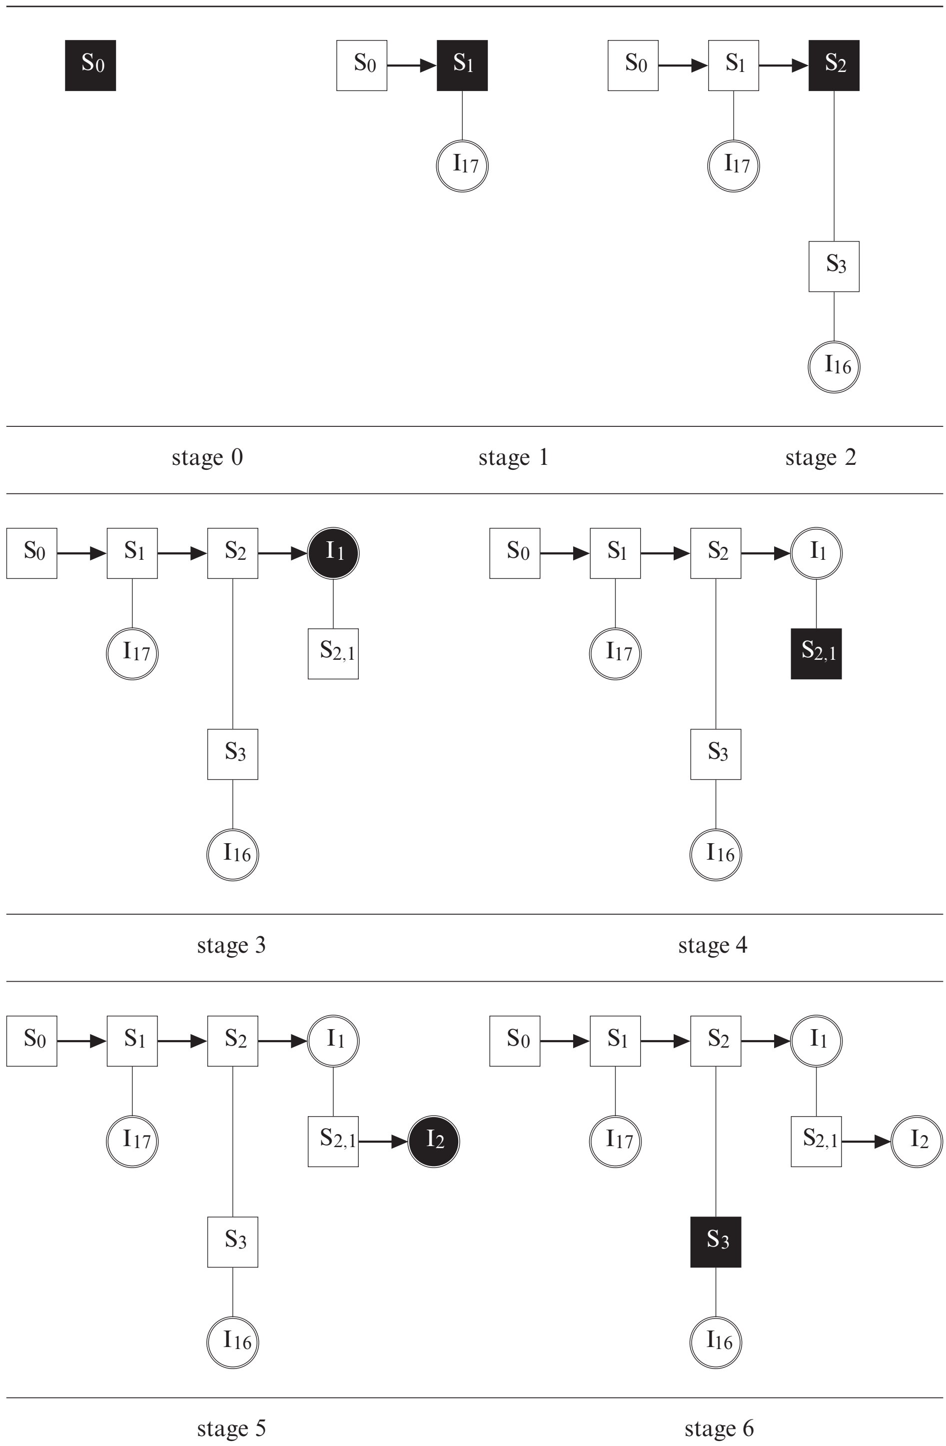

Fig. 3. Step by step progression for the first seven stages of planning underlying the proof activity corresponding to Fraleigh's proof.

https://doi.org/10.1017/S1755020319000601 Published online by Cambridge University Press

components, or *modules*, with *dependencies* between them" (Avigad, 2020, p. 5).[30,31] In this section, we will see that combining the notions of plan and modularity opens the way to what we will call *modular presentations* of mathematical proofs, that is, presentations where parts of the proof are encapsulated into *modules*. Furthermore, we will see that the possibility of modular presentations of mathematical proofs is precisely what Leron (1983) and Lamport (1995, 2012) have exploited in their respective alternative proposals on how to write and present mathematical proofs, and that the benefits of such alternative forms of proof presentation—an increase in comprehensibility and communicability for Leron, and an increase in rigor, readability, and understandability for Lamport—can be traced back to benefits associated to modularity.

The connection between plans and modularity is to be found in the *hierarchical structure* of plans. According to Bratman, plans are hierarchical in the following sense:

> [O]ur plans typically have a *hierarchical structure*. Plans concerning ends embed plans concerning means and preliminary steps; and more general intentions (for example, my intention to go to a concert tonight) embed more specific ones (for example, my intention to hear the Alma Trio). As a result, I may deliberate about parts of my plan while holding other parts fixed. (Bratman, 1987, p. 29)

This hierarchical structure is manifest in our definition of plans for proof activities, since the agent's plan always has an ordered tree structure. Avigad (2020) has pointed out the following connection between hierarchical structure and modularity:

> Modularity is often associated with an additional property:
>
> - Organization into modules can be *hierarchical*: within a module, components can be divided into smaller *submodules*, and so on.
>
> This is not a necessary feature of a modular system, in that one can have modular designs that are essentially flat. But a hierarchical design only makes sense in terms of a modular presentation, and, conversely, the most modular description of a system is often obtained via a hierarchical conception of its components. As mathematical theories and proofs also have a hierarchical structure, this is an issue that is worth keeping in mind. (Avigad, 2020, p. 5)

This relation between hierarchical structure and modularity provides a way to generate modular presentations of mathematical proofs. First, a modular description of a plan $\mathcal{P}$ underlying a given mathematical proof P can be obtained from $\mathcal{P}$ by simply 'omitting' one or several subtrees of $\mathcal{P}$. More specifically, any subtree of $\mathcal{P}$ can be conceived as a module which can be described by stating the proving intention of type 'to show' at its root, and which can be viewed as encapsulating the information as to how the proving intention is fulfilled, that is, as to how one shows that its conclusion follows

30 Avigad does not restrict his attention to the modularity of mathematical proofs, but addresses more generally the modularity of various 'pieces of mathematics' (Avigad's terms), such as theories, definitions, and methods.

31 Avigad & Morris (2016) analyze in detail how proofs of Dirichlet's theorem in number theory became more modular over a period of approximately 100 years.

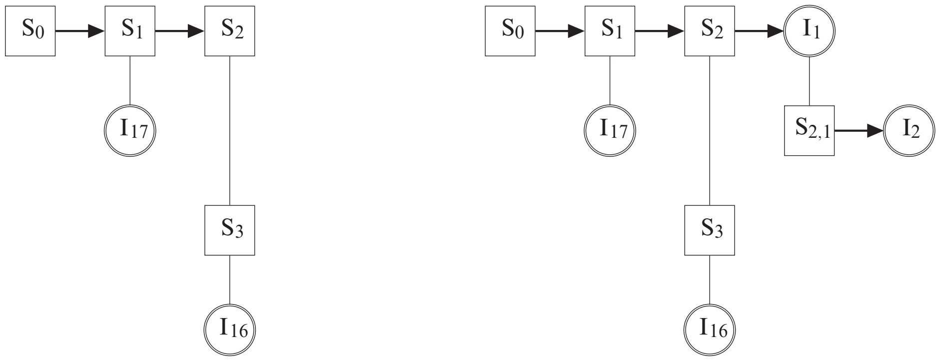

Fig. 4. Two modular descriptions of the plan underlying Fraleigh's proof.

from its hypotheses.[32] As an illustration, the (partial) plans reached at stages 2 and 5 of Fraleigh's proof (cf. Figure 3) and reported in Figure 4 can be seen as modular descriptions of the full plan obtained once the proof activity is completed.[33]

A modular presentation of a proof P can then be obtained from a modular description of its underlying plan $\mathcal{P}$ by replacing all the deductive steps in each module by the description of the module itself—i.e., by the proving intention figuring at its root.[34] This amounts, for each module, to encapsulating the information as to how the considered proving intention is fulfilled, i.e., to only mention *what* is being shown without saying *how* it is shown. As an illustration, the first modular description of the plan underlying Fraleigh's proof displayed on the left in Figure 4 would lead to the following modular presentation of the proof:

Proof. We must show that the different cells $\bar{a} = \{x \in S \mid x \sim a\}$ for $a \in S$ do give a partition of S, so that every element of S is in some cell and so that if $a \in \bar{b}$, then $\bar{a} = \bar{b}$. First, <u>we show that every element of S is in some cell</u>. Then, <u>we show that if $\bar{a} \in \bar{b}$, then $\bar{a} = \bar{b}$</u>. We conclude that the cells $\bar{a} = \{x \in S \mid x \sim a\}$ for $a \in S$ do give a partition of S. □

In this proof, the two underlined proving intentions S_2 and S_3 correspond to the descriptions of two modules providing information on what is being shown, namely that "every element of S is in some cell" and that "if $a \in \bar{b}$, then $\bar{a} = \bar{b}$." S_2 and S_3 encapsulate then information as to how these propositions are shown, i.e., information regarding the deductive steps required to fulfill these proving intentions. The second modular description displayed on the right in Figure 4 would lead to the following modular presentation of Fraleigh's proof:

32 Notice that there are always multiple modular descriptions of $\mathcal{P}$, since each considered set of modules will yield a particular modular description of $\mathcal{P}$.

33 For the full plan associated to Fraleigh's proof, see Figure A1 in the appendix.

34 There are thus as many modular presentations of P as there are modular descriptions of $\mathcal{P}$.

https://doi.org/10.1017/S1755020319000601 Published online by Cambridge University Press

Proof. We must show that the different cells $\bar{a} = \{x \in S \mid x \sim a\}$ for $a \in S$ do give a partition of S, so that every element of S is in some cell and so that if $\bar{a} \in \bar{b}$, then $\bar{a} = \bar{b}$. First, we show that every element of S is in some cell. Let $a \in S$. Then $a \in \bar{a}$ by the reflexive condition, so a is in *at least one* cell. Then, we show that if $a \in \bar{b}$, then $\bar{a} = \bar{b}$. We conclude that the cells $\bar{a} = \{x \in S \mid x \sim a\}$ for $a \in S$ do give a partition of S. □

Here the only difference with the previous modular presentation is that information that was encapsulated before in module S_2 is now provided, that is, information as to how the proving intention S_2 is fulfilled.[35]

The possibility of modular presentations of mathematical proofs appears to be intimately connected with two alternative proposals by Leron (1983) and Lamport (1995, 2012) on how to write and present mathematical proofs. The objectives of these proposals are to increase their comprehensibility and facilitate their communication for Leron, and to make them more rigorous (less error-prone) as well as easier to read and understand for Lamport.

Leron conceives of his approach as a 'structural method' of proof presentation, a method that he describes as follows:

> The basic idea underlying the structural method is to arrange the proof in *levels*, proceeding from the top down; the levels themselves consist of short autonomous "modules," each embodying one major idea of the proof.
>
> The top level gives in very general (but precise) terms the main line of the proof. The second level elaborates on the generalities of the top level, supplying proofs for unsubstantiated statements, details for general descriptions, specific constructions for objects whose existence has been merely asserted, and so on. If some such subprocedure is itself complicated, we may choose to give it in the second level only a "top-level description," pushing the details further down to lower levels. And so we continue down the hierarchy of subprocedures, each supplying more details to plug in holes in higher levels, until we reach the bottom where [...] all the leaks are plugged and the proof is watertight. (Leron, 1983, pp. 174–175)

According to Leron, the main benefit of the structural method is that the "ideas behind the formal proofs are better communicated" insofar as "the main idea is given in the top level (or two), auxiliary ideas are packaged in autonomous modules, and the interconnections between the separate ideas are made explicit through the structure diagram" (Leron, 1983, p. 183). Thus, it is by isolating and identifying the main ideas of a mathematical proof that the structural method increases its comprehensibility and communicability, as compared to standard forms of proof presentation.

Lamport describes his method of proof presentation as being "based on hierarchical structuring" which he sees as "a successful tool for managing complexity" (Lamport, 1995, p. 600). The main idea, as with Leron, is to write and present a mathematical

35 In the two modular presentations provided above, we took the liberty to slightly rewrite the original text of Fraleigh's proof. In particular, we have now mentioned explicitly the deductive step corresponding to inference I_{17}. We have followed Fraleigh's original choice not to mention the deductive step corresponding to inference I_{16} as well as the proving intention $S_{2,1}$ (in the case of the second modular presentation).

proof in *levels*, where levels can themselves be described in terms of lower levels. The result of such a presentation is what Lamport refers to as a 'structured proof'. Lamport provides several examples of structured proofs where various levels are represented in a hierarchical manner on the page, using indentation and a specific numbering scheme to identify each level and its position in the overall hierarchical structure. According to Lamport, the capacity of a structured proof to make the structure of a proof manifest, and to identify precisely each deductive step within this structure, is precisely what makes them more readable and understandable. Moreover, structured proofs are better than traditional ones for avoiding errors since they force the writer to provide more details than usual (Lamport, 1995, pp. 606–607), and to make explicit the logical status of each sentence in the proof (Lamport, 2012, pp. 43–44).

From the perspective developed in this section, both Leron's and Lamport's proposals on how to write and present mathematical proofs amount to providing what we have called modular presentations of mathematical proofs. Their notions of levels correspond, in our framework, to the levels in the plan underlying any given mathematical proof, and their descriptions of each level correspond to the descriptions of the modules lying at this level, that is, descriptions encapsulating information relative to the deductive steps occurring at lower levels in the plan's hierarchical structure. Furthermore, the various benefits of such alternative presentations as claimed by Leron and Lamport correspond exactly to some of the benefits of modularity as identified by Avigad (2020): the fact that a modular presentation increases the comprehensibility, communicability, readability, understandability, of a mathematical proof corresponds to what Avigad refers to as *comprehensibility*, namely the idea that: "[w]hen a system is modular, it is easier to understand, explain, and predict its behavior" (Avigad, 2020, p. 6); the fact that a modular presentation increases the rigor of a mathematical proof by making it easier to spot potential errors corresponds to what Avigad refers to as *reliability* and *robusteness*, namely the idea that: "[a]n appropriately modular description makes it possible to assess and test components of a system individually, to localize a problem to the behavior of one component, and to detect problems that would otherwise be lost in an overabundance of detail" (Avigad, 2020, p. 6).

To sum up, our notion of plan has intimate connections with the modularity of mathematical proofs in the sense of Avigad (2020). These connections can be exploited to generate what we have called modular presentations of mathematical proofs. As we saw through the works of Leron and Lamport, modular presentations of mathematical proofs can deliver various epistemic benefits such as increasing their comprehensibility, reliability, and robustness. Insofar as a given mathematical proof always admits many different modular presentations, it would be particularly interesting to investigate, in further works, which modular presentations maximize these epistemic benefits.

§8. Comparison with the framework of proof planning. Although the notions of plans and planning in the context of mathematical proofs have not been investigated in the philosophy of mathematics, they have received some attention in the field of automated theorem proving. Of particular interest is the framework of *proof planning* developed by Alan Bundy and colleagues (Bundy, 1988; Bundy et al., 2005). Since it is situated in computer science and artificial intelligence, this framework deals primarily with proofs and proving for *artificial agents*, but Bundy (1991) has argued that it also possesses explanatory power for understanding various aspects of proofs and proving for *human agents*. We will now compare our account of plans and planning in the context of mathematical proofs with the framework of proof planning with respect to the specific objective of providing an account of these notions that aims to be faithful to the reality of ordinary mathematical practice.

The technique of proof planning rests on the idea that the automatic search for a formal proof can be *guided* using a *proof plan*.[36] By contrast with more local approaches—which proceed by deciding, at any given stage, how the proof search should move forward—the technique of proof planning aims to implement a *global* approach where proof search is conducted from a higher-level description of the proof toward its more specific details. According to Bundy, such a global approach "seems to accord more with the intuitions of human mathematicians that they first make a global plan of a proof and then fill in the details" (Bundy, 1996, p. 261). Furthermore, the idea that proof plans can be used to guide proof search is also directly inspired by the way human mathematicians prove theorems:

> We believe that human mathematicians can draw on an armoury of such proof plans when trying to prove theorems. It is our intuition that we do this when proving theorems, and the same intuition is reported by other experienced mathematicians. One can identify such proof plans by collecting similar proofs into families having a similar structure, e.g., those proved by 'diagonalization' arguments. Many inductive proofs seem to have such a similar structure. (Bundy, 1988, p. 112)

The main idea behind proof planning is thus to capture the *common structure* of a family of proofs in a computational representation which, in turn, could be exploited to find other proofs belonging to the same family—e.g., the development of proof planning for inductive proofs was built on the observation that many inductive proofs had the very same structure, and that identifying this structure can be used to guide the search for other inductive proofs.

Technically, the framework of proof planning involves the notions of *tactic*, *method*, and *critic*. A *tactic* is an algorithmic procedure which constructs part of a proof through the application of a sequence of rules of inference in an LCF-style theorem prover[37]. A *method* is a partial specification of a tactic, which is composed of a set of *preconditions* stipulating conditions for the application of the tactic, and a set of *effects* which will hold if the application of the tactic is successful. A *critic* captures a possible pattern of failure to a given method in a proof plan, and provides corresponding instructions on how to patch the failed proof.[38] A *proof plan* is then defined as:

> [T]he method for one of the top-level tactics, i.e., [...] the specification of a strategy for controlling a whole proof, or a large part of one. This super-method is so constructed that the preconditions of each of its sub-methods are either implied by its preconditions or by the effects

[36] The technique was originally applied to the search of inductive proofs, but then found applications in various other domains. See Bundy et al. (2005) for an overview of the technique and its applications.

[37] The acronym LCF stands for 'Logic of Computable Functions', a logical system due to Dana Scott (Scott, 1993) and implemented by the pioneers of the LCF approach to theorem proving in the so-called Stanford LCF theorem prover (Milner, 1972).

[38] For a description of the functioning of critics see Ireland (1992) and Ireland & Bundy (1996).

> of earlier sub-methods. Similarly, its effects are implied by the effects of its sub-methods. If the preconditions of a method are satisfied then its tactic is applicable. If the tactic application succeeds then its effects are satisfied. The original conjecture should satisfy the preconditions of the plan; the effects of the plan should imply that the conjecture has been proved. (Bundy, 1988, p. 116)

In practice, the approach of proof planning requires a *proof planner* which represents the common structure of a family of proofs as a *general-purpose* tactic, and which, given a conjecture to be proved, will produce a specific *proof plan* for this conjecture, that is, a *customized* tactic which can be used to guide a proof search for this conjecture within a particular tactic-based theorem prover.[39]

As witnessed by the first quote from Bundy (1988) reported above, the framework of proof planning originated from observations similar to the ones motivating our own account, and as a consequence Bundy's notion of proof plan and our own notion of plan bear certain similarities. Upon closer scrutiny, however, Bundy's notion of proof plan appears closer to the notion of method we introduced in §5.1. This parallel is particularly salient in the following passage:

> Mathematicians recognise families of proofs which contain common structure. These families are sometimes named, e.g., diagonalization arguments, but, more often, are not. We propose representing such common structures with proof plans. (Bundy, 1991, p. 180)

Such common structures are often referred to as 'mathematical methods' in mathematical practice—mathematicians commonly speak of the 'diagonalization method', or the 'method of mathematical induction'—and it is this specific use of the term method that we have tried to capture in §5.1. Furthermore, Bundy insists that a proof plan should always "carry some expectation of success" (Bundy, 1988, p. 112), although success is, of course, not guaranteed. In our account, we implemented this aspect directly into the notion of method, by requiring that a method always hold some promise of success in fulfilling the proving intentions it is applied to. In a sense, the framework of proof planning, in its current formulation, represents proof plans as consisting of the application of a *single* method to the conjecture or proposition to be proved which provides, at once, a *complete* plan for the whole proof activity. From the perspective of our account, the notions of method and plan would then coincide in this case.

But the main difference between our account and the framework of proof planning lies in the role attributed to practical reasoning. The framework of proof planning always represents proof plans as being complete and fully determined, which means that there is no need for practical reasoning in the building of proof plans. By contrast, practical reasoning is an essential element of our account since it is central to the step by step process by which plans are constructed. Furthermore, and as we have seen in the examples discussed in §5, plans for mathematical proofs are often built

[39] For concrete implementations of proof planners see Bundy, van Harmelen, Horn, & Smaill (1990) and Dixon & Fleuriot (2003).

from the combination of various and sometimes quite different instances of practical reasoning, while Bundy's proof plans are relatively homogeneous and monolithic. From our perspective, providing a proper analysis of practical reasoning appears then as an essential task in the prospect of articulating an account of plans and planning for mathematical proofs that aims to be faithful to the reality of mathematical practice.

Finally, it is interesting to notice that there is room for a potentially fruitful cross-fertilization between the epistemology of mathematics and the various fields concerned with mathematical reasoning in computer science and artificial intelligence on the issues discussed in this paper. On the one hand, a fine-grained analysis of plans and planning in the contexts of mathematical proofs and proof activities, and the development of epistemological models thereof, might yield important insights to help build more efficient technologies in interactive and automated theorem proving. On the other hand, not only do the already existing technologies in these fields embody important insights concerning the nature of proof activities for human agents that ought to find their way into the epistemology of mathematics, these developments also suggest that formal methods could be used to provide a computational model of plans and planning for mathematical proofs and proof activities that would be tuned to the specificities of human agency in mathematical practice.

§9. Conclusion. The aim of this paper was to provide a philosophical account of plans and planning in the context of mathematical proofs. Our approach has consisted in analyzing the plan of a mathematical proof as the plan of the agent(s) who carried out the corresponding proof activity. This led us to develop an account of plans and planning in the context of proof activities, for which we built on Bratman's theory of planning agency. The plan of an agent for a proof activity has been defined as an ordered tree whose nodes are proving intentions, whose root is the proving intention corresponding to the theorem at hand, and where each set of ordered children consists of a subplan obtained from the parent node through an instance of practical reasoning. The execution of this plan leads then to the realization of the proof activity, i.e., to the carrying out of a sequence of actions—the deductive inferences—in the same way as the execution of one's travel plan leads to a sequence of moves constituting the actual travelling. From this perspective, the written mathematical proof is nothing more than a report of its corresponding proof activity—it is thereby analogous to a travel diary reporting the moves of a travelling activity. This is why it is natural to talk about the plan "underlying" or "lying behind" a mathematical proof. In a sense, a plan always precedes its execution, that is, the activity it gives rise to, and a fortiori any report of this activity. In this work, we have proposed to shift the focus from the written mathematical proofs to their corresponding proof activities, and to the form of planning agency underlying them. The resulting framework thus constitutes an *action*-based account of plans and planning in the context of mathematical proofs.

One intended application of the present framework is to the philosophical study of understanding and explanation in the context of mathematical proofs. As we saw in the introduction, several authors have pointed out a direct connection between these two notions and the one of plan for mathematical proofs. The approach developed in this paper suggests the following characterization: *understanding* a mathematical proof amounts to *grasping* its underlying plan, that is, grasping the plan of the agent who carried out the corresponding proof activity. Given our definition of the notion of plan, grasping the plan underlying a proof activity would then require (1) grasping its *ordered tree structure*, that is how its various *proving intentions* fit into the tree, and (2) grasping the *practical reasoning* behind each parent–children connection. In case an agent fails to understand a mathematical proof by failing to grasp one or more of these components, one might then investigate the specific type(s) of explanation that could be provided so as to enable the agent to grasp those components, and thereby to understand the proof. This proposal should be able to account for two presumed features of understanding raised by Folina (2018) and Robinson (2000) in the passages quoted in the introduction of this paper. First, Folina's quote suggests that understanding a proof may have to do with "how the parts [of the proof] are related and connected to one another." In our framework, the connections between the different parts of a proof are precisely to be found in the plan underlying it, and more specifically in its hierarchical organization. If understanding a proof amounts to grasping its underlying plan, we then automatically get that understanding a proof amounts to grasping how the parts of the proof are connected to one another, which means that Folina's suggestion would appear as a direct consequence of our characterization. Second, both Folina's and Robinson's quotes suggest that understanding is related to the capacity to grasp the "architecture," "blueprint," "overall plan" of a mathematical proof so as to "not [be] concerned with" or "avoid getting lost in" the details. This resonates strongly with our discussion of plans and modularity: because a plan always has an ordered tree structure, details as to how a given proving intention is being fulfilled can always be encapsulated into modules, leading to different possible modular descriptions of the plan. From this perspective, grasping the "architecture," "blueprint," "overall plan" of a mathematical proof would then amount to grasping a certain modular description of the underlying plan. These considerations constitute only some indications of the potential interest of our account for the study of understanding and explanation in the context of mathematical proofs. A detailed analysis of these notions based on the present framework will be the object of a subsequent paper.

Finally, it seems plain that the phenomena of plans and planning play a central role in the contexts of proof discovery and problem solving. This, of course, is not a novel observation, and Pólya (1945, 1954a,1954b, 1962), in his seminal work on those matters, has already emphasized the central importance of plans and planning in these contexts. In this respect, one of the central themes for Pólya is that of the *flexibility* of one's plans:

> When the problem-solver debates his plan of the solution with himself, this plan is usually more "fluid" than "rigid," it is more felt than formulated. In fact it would be foolish of the problem-solver to fix his plan prematurely. A wise problem-solver does not commit himself to a rigid plan. Even at a later stage, when the plan is riper, he keeps it ready for modification, he leaves it a certain flexibility, he reckons with unforeseen difficulties to which he might be obliged to adapt his plan. (Pólya, 1954b, p. 154)

Pólya's remarks raises many interesting issues, such as how to maintain the right level of flexibility of your plans, and how to revise your plans in face of difficulties in completing or executing them. These issues deserve further attention, and call, more generally, for the development of an account of plans and planning in the contexts of proof discovery and problem solving.

§10. Appendix: Fraleigh's proof: Plan, proving intentions, and deductive steps. In §6, we described the first seven stages of planning underlying Fraleigh's proof activity. In this appendix, we are completing this description with a presentation of the entire planning underlying this proof activity. For reasons of space, we shall not provide here a stage by stage description of this proof activity, nor shall we spell out the mathematical methods involved in the various instances of practical reasoning occurring in it. Rather, we are providing in Figure A.1 the overall plan resulting from the entire planning underlying Fraleigh's proof activity. The proving intentions of type 'to show' appearing in Figure A.1 are listed in Table A.1. The proving intentions of type 'to infer' appearing in Figure A.1 are listed in Figure A.2, which is a Fitch style representation of the various deductive steps in Fraleigh's proof, where the notation Ii is used to denote the ith deductive step in Fraleigh's proof—the notation I_i in Figure A.1 being used for the proving intention of type 'to infer' corresponding to the deductive step Ii. In Figure A.2, the deductive steps for which there is a * sign as superscript are ones that are left implicit in the original text of Fraleigh's proof. The interested reader

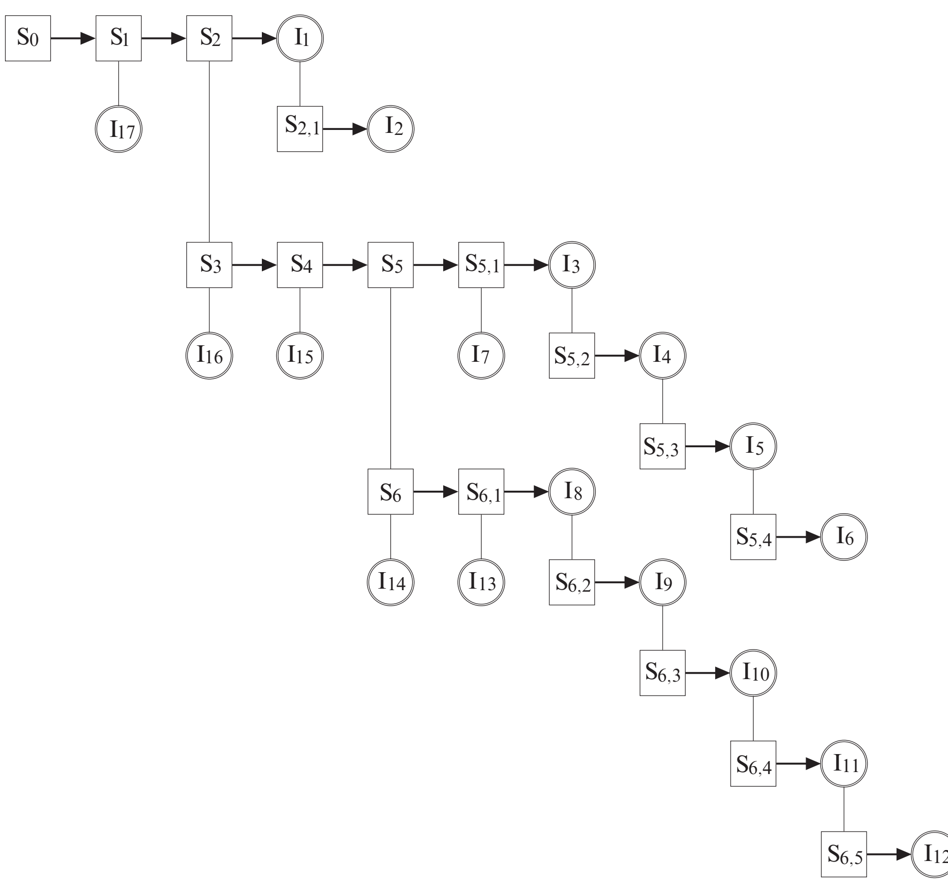

Fig. A.1. The plan underlying the proof activity corresponding to Fraleigh's proof.

https://doi.org/10.1017/S1755020319000601 Published online by Cambridge University Press

Table A.1. *Proving intentions of type 'to show' in the plan underlying the proof activity corresponding to Fraleigh's proof*

S_0: $S \neq \emptyset$, $\sim$ eq. rel. on S	$\Rightarrow \sim$ yields a partition of S where $\bar{a} = \{x \in S \mid s \sim a\}$
S_1: $S \neq \emptyset$, $\sim$ eq. rel. on S	$\Rightarrow$ for all $a \in S$, a is in some cell and if $a \in \bar{b}$, then $\bar{a} = \bar{b}$
S_2: $S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$	$\Rightarrow a$ is in some cell
S_3: $S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$	$\Rightarrow$ if $a \in \bar{b}$, then $\bar{a} = \bar{b}$
S_4: $S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b} \Rightarrow \bar{a} = \bar{b}$	
S_5: $S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b} \Rightarrow \bar{a} \subseteq \bar{b}$	
S_6: $S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b} \Rightarrow \bar{a} \supseteq \bar{b}$	

S_2:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$	$\Rightarrow$	a is in some cell
$S_{2,1}$:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{a}$	$\Rightarrow$	a is in some cell

S_5:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b}$	$\Rightarrow$	$\bar{a} \subseteq \bar{b}$
$S_{5,1}$:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b}$, $x \in \bar{a}$	$\Rightarrow$	$x \in \bar{b}$
$S_{5,2}$:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b}$, $x \in \bar{a}$, $x \sim a$	$\Rightarrow$	$x \in \bar{b}$
$S_{5,3}$:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b}$, $x \in \bar{a}$, $x \sim a$, $a \sim b$	$\Rightarrow$	$x \in \bar{b}$
$S_{5,4}$:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b}$, $x \in \bar{a}$, $x \sim a$, $a \sim b$, $x \sim b$	$\Rightarrow$	$x \in \bar{b}$

S_6:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b}$	$\Rightarrow$	$\bar{a} \supseteq \bar{b}$
$S_{6,1}$:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b}$, $y \in \bar{b}$	$\Rightarrow$	$y \in \bar{a}$
$S_{6,2}$:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b}$, $y \in \bar{b}$, $y \sim b$	$\Rightarrow$	$y \in \bar{a}$
$S_{6,3}$:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b}$, $y \in \bar{b}$, $y \sim b$, $a \sim b$	$\Rightarrow$	$y \in \bar{a}$
$S_{6,4}$:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b}$, $y \in \bar{b}$, $y \sim b$, $a \sim b$, $b \sim a$	$\Rightarrow$	$y \in \bar{a}$
$S_{6,5}$:	$S \neq \emptyset$, $\sim$ eq. rel. on S, $a \in S$, $a \in \bar{b}$, $y \in \bar{b}$, $y \sim b$, $a \sim b$, $b \sim a$, $y \sim a$	$\Rightarrow$	$y \in \bar{a}$

is invited to proceed through the various stages of planning underlying the proof activity corresponding to Fraleigh's proof as depicted in Figure A.1, and to realize that the execution of this plan leads then to a sequence of deductive inferences which corresponds to the sequence of deductive steps reported in Fraleigh's proof.

§11. Acknowledgments. We would like to thank Jeremy Avigad, Michael Bratman, Sébastien Gandon, Fenner Tanswell, and Keith Weber for helpful comments on earlier versions of this paper. We would also like to thank audiences at the Savoir-Faire Mathématiques Workshop at the Université Clermont Auvergne (September 2017), the Ninth French Philosophy of Mathematics Workshop at the Archives Henri Poincaré (October 2017), the Topics in Logic Seminar at Stanford University (April 2018), the Seventh Workshop on Logic, Rationality and Intelligent Interaction at Stanford University (June 2018) and the Proofs and Representations Workshop at LMU Munich (July 2018) for insightful questions and comments. Finally we would like to thank two anonymous reviewers for their helpful feedback and suggestions. Rebecca Morris completed most of this work while holding a postdoctoral scholarship at the Suppes Center for History and Philosophy of Science at Stanford University. Yacin Hamami was supported by a postdoctoral fellowship of the Research Foundation—Flanders (FWO).

S is a nonempty set	
$\sim$ is an equivalence relation on S	
$a \in S$	
$a \in \bar{a}$	(I1)
a is in at least one cell	(I2)
a is in a cell $\bar{b}$	
$x \in \bar{a}$	
$x \sim a$	(I3)
$a \sim b$	(I4)
$x \sim b$	(I5)
$x \in \bar{b}$	(I6)
$\bar{a} \subseteq \bar{b}$	(I7)
$y \in \bar{b}$	
$y \sim b$	(I8)
$a \sim b$	(I9)
$b \sim a$	(I10)
$y \sim a$	(I11)
$y \in \bar{a}$	(I12)
$\bar{b} \subseteq \bar{a}$	(I13)
$\bar{b} = \bar{a}$	(I14)
if $a \in \bar{b}$, then $\bar{a} = \bar{b}$	(I15*)
for all $a \in S$, a is in some cell and if $a \in \bar{b}$, then $\bar{a} = \bar{b}$	(I16*)
$\sim$ yields a partition of S where $\bar{a} = \{x \in S \mid x \sim a\}$	(I17*)

Fig. A.2. Deductive steps in Fraleigh's proof.

BIBLIOGRAPHY

Avigad, J. (2006). Mathematical method and proof. *Synthese*, **153**(1), 105–159.

———. (2020). Modularity in mathematics. *The Review of Symbolic Logic*, **13**(1), 47–79.

Avigad, J. & Morris, R. L. (2016). Character and object. *The Review of Symbolic Logic*, **9**(3), 480–510.

Bedford, E. & Smillie, J. (1998). Polynomial diffeomorphisms of $\mathbf{C}^2$: VI. Connectivity of J. *Annals of Mathematics, Second Series*, **148**(2), 695–735.

Boghossian, P. (2014). What is inference? *Philosophical Studies*, **169**(1), 1–18.

Bratman, M. E. (1987). *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press. Reissued by CSLI Publications, Stanford, CA, 1999 (citations are to the latter edition).

———. (2014). *Shared Agency: A Planning Theory of Acting Together*. New York: Oxford University Press.

Bundy, A. (1988). The use of explicit plans to guide inductive proofs. In Lusk, E. and Overbeek, R., editors. *9th International Conference on Automated Deduction*. Lecture Notes in Computer Science, Vol. 310. Berlin: Springer-Verlag, pp. 111–120.

———. (1991). A science of reasoning. In Lassez, J.-L. and Plotkin, G., editors. *Computational Logic: Essays in Honor of Alan Robinson*. Cambridge, MA: MIT Press, pp. 178–198.

———(1996). Proof planning. In Drabble, B., editor. *Proceedings of the Third International Conference on Artificial Intelligence Planning Systems*. Menlo Park, CA, pp. 261–267.

Bundy, A., Basin, D., Hutter, D., & Ireland, A. (2005). *Rippling: Meta-Level Guidance for Mathematical Reasoning*. Cambridge Tracts in Theoretical Computer Science, Vol. 56. Cambridge: Cambridge University Press.

Bundy, A., van Harmelen, F., Horn, C., & Smaill, A. (1990). The OYSTER-CLAM system. In Stickel, M. E., editor. *10th International Conference on Automated Deduction*. Lecture Notes in Artificial Intelligence, Vol. 449. Berlin: Springer-Verlag, pp. 647–648.

Chemla, K. & Virbel, J. (2015). *Texts, Textual Acts and the History of Science*. Cham: Springer.

Corcoran, J. (1989). Argumentations and logic. *Argumentation*, **3**(1), 17–43.

Dixon, L. & Fleuriot, J. (2003). Isa Planner: A prototype proof planner in Isabelle. In Baader, F., editor. *Automated Deduction–CADE*-19. Lecture Notes in Computer Science, Vol. 2741. Berlin: Springer-Verlag, CADE, pp. 279–283.

Faber, C. & Pandharipande, R. (2003). Hodge integrals, partition matrices, and the λ_g conjecture. *Annals of Mathematics, Second Series*, **157**(1), 97–124.

Feferman, S. (2012). And so on...: Reasoning with infinite diagrams. *Synthese*, **186**(1), 371–386.

Folina, J. (2018). Towards a better understanding of mathematical understanding. In Piazza, M. and Pulcini, G., editors. *Truth, Existence and Explanation*. Cham: Springer, pp. 121–146.

Fraleigh, J. B. (2003). *A First Course in Abstract Algebra* (seventh edition). London: Pearson Education.

Ganesalingam, M. & Gowers, T. (2017). A fully automatic theorem prover with human-style output. *Journal of Automated Reasoning*, **58**(2), 253–291.

Gowers, T. (2002). *Mathematics: A Very Short Introduction*. Oxford: Oxford University Press.

Hamami, Y. (2018). Mathematical inference and logical inference. *The Review of Symbolic Logic*, **11**(4), 665–704.

———. (2019). Mathematical rigor and proof. *The Review of Symbolic Logic*. To appear.

Hardy, G. H., Littlewood, J. E., & Pólya, G. (1934). *Inequalities*. Cambridge: Cambridge University Press.

Haylock, D. & Cockburn, A. (2008). *Understanding Mathematics for Young Children: A Guide for Foundation Stage and Lower Primary Teachers*. London: SAGE Publications Ltd.

Ireland, A. (1992). The use of planning critics in mechanizing inductive proofs. In Voronkov, A., editor. *International Conference on Logic for Programming Artificial Intelligence and Reasoning*. Lecture Notes in Computer Science, Vol. 624. Berlin: Springer-Verlag, pp. 178–189.

Ireland, A. & Bundy, A. (1996). Productive use of failure in inductive proof. *Journal of Automated Reasoning*, **16**(1–2), 79–111.

Lamport, L. (1995). How to write a proof. *The American Mathematical Monthly*, **102**(7), 600–608.

———. (2012). How to write a 21st century proof. *Journal of Fixed Point Theory and Applications*, **11**(1), 43–63.

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

Michener, E. R. (1978). Understanding understanding mathematics. *Cognitive Science*, **2**(4), 361–383.

Milner, R. (1972). *Logic for computable functions*: *Description of a machine implementation*. Technical Report STAN-CS-72-288, A.I. Memo 169, Stanford University.

Morris, R. L. (2020). Motivated proofs: What they are, why they matter and how to write them. *The Review of Symbolic Logic*, **13**(1), 23–46.

Poincaré, H. (1900/1996). Intuition and logic in mathematics. In Ewald, W., editor. *From Kant to Hilbert*: *A Source Book in the Foundations of Mathematics* (*Volumes I and II*), Vol. 19. Oxford: Clarendon Press, pp. 1012–1020.

———. (1908). Science et Méthode. Paris: Ernest Flammarion. Translated to English by Francis Maitland as *Science and Method*. London: Thomas Nelson & Sons, 1914. (Citations are to translation.)

Pólya, G. (1945). *How to Solve It*: *A New Aspect of Mathematical Method*. Princeton, NJ: Princeton University Press.

———. (1949). With, or without, motivation. *The American Mathematical Monthly*, **56**(10), 684–691.

———(1954a). *Mathematics and Plausible Reasoning*: *Induction and Analogy in Mathematics*, Vol. 1. Princeton, NJ: Princeton University Press.

———(1954b). *Mathematics and Plausible Reasoning*: *Patterns of Plausible Inference*, Vol. 2. Princeton, NJ: Princeton University Press.

———(1962). *Mathematical Discovery*: *On Understanding*, *Learning*, *and Teaching Problem Solving* (*Two Volumes*). New York: John Wiley & Sons, Inc.

Prawitz, D. (2012). The epistemic significance of valid inference. *Synthese*, **187**(3), 887–898.

———. (2015). Explaining deductive inference. In Wansing, H., editor. *Dag Prawitz on Proofs and Meaning*. Cham: Springer, pp. 65–100.

Rav, Y. (1999). Why do we prove theorems? *Philosophia Mathematica*, **7**(3), 5–41.

Robinson, J. A. (2000). Proof = guarantee + explanation. In Hölldobler, S., editor. *Intellectics and Computational Logic*. Applied Logic Series, Vol. 19. Dordrecht: Kluwer Academic Publishers, pp. 277–294.

Scott, D. S. (1993). A type-theoretical alternative to ISWIM, CUCH, OWHY. *Theoretical Computer Science*, **121**(1–2), 411–440.

Steele, J. M. (2004). *The Cauchy-Schwarz Master Class*: *An Introduction to the Art of Mathematical Inequalities*. Cambridge: Cambridge University Press.

Sundholm, G. (2012). “Inference versus consequence” revisited: Inference, consequence, conditional, implication. *Synthese*, **187**(3), 943–956.

Tanswell, F. (2019). Go forth and multiply! On actions, instructions and imperatives in mathematical proofs. Unpublished manuscript.

van Benthem, J. (2018). Constructive agents. *Indagationes Mathematicae*, **29**(1), 23–35.

Wright, C. (2014). Comment on Paul Boghossian, "What is inference". *Philosophical Studies*, **169**(1), 27–37.

CENTRE FOR LOGIC AND PHILOSOPHY OF SCIENCE
VRIJE UNIVERSITEIT BRUSSEL
BRUSSELS B-1050, BELGIUM
E-mail: yacin.hamami@gmail.com

INDEPENDENT SCHOLAR
E-mail: email@rebeccaleamorris.com