Characterizing decision-making opportunities in undergraduate physics coursework

Barron J. Montgomery
Physics Department, Stanford University, 450 Serra Mall, Stanford, California 94305, USA

Argenta M. Price
Doerr School of Sustainability, Stanford University, 397 Panama Mall, Stanford, California 94305, USA

Carl E. Wieman
Physics Department and Graduate School of Education, Stanford University, 382 Via Pueblo, Stanford, California 94305, USA

(Received 2 February 2024; accepted 9 July 2024; published 29 July 2024)

A major goal of physics education is to develop strong problem-solving skills for students. To become expert problem solvers, students must have opportunities to deliberately practice those skills. In this work, we adopt a previously described definition of problem solving that consists of a set of 29 decisions made by expert scientists. We quantified the amount of practice undergraduate physics students get at making each decision by coding the decisions required in assignments from introductory, intermediate, and advanced physics courses at a prestigious university. A research-focused capstone course was the only example that offered substantial practice at a large range of decisions. Problems assigned in the traditional coursework required only a few decisions and routinely reduced potential opportunities for students to make other decisions. In addition, we modified traditional physics coursework to offer more decision-making practice. We observed that this increased the number of decisions students actually made in solving the problems. This work suggests that to better prepare undergraduates for solving problems in the real world, we must offer more opportunities for students to make and act on problem-solving decisions.

DOI: 10.1103/PhysRevPhysEducRes.20.020103

I. INTRODUCTION

Scientists and engineers face intricate problems without clear-cut answers, so one of the main goals of science, technology, engineering, and mathematics (STEM) educators has been to teach problem-solving skills. The field of physics problem solving has been studied extensively. A review of the physics problem-solving literature since 1980 by Ince lists 23 studies of the strategies used by students in solving standard course problems and an additional 4 studies that compare the problem-solving strategies of experts and introductory students. [1] They also list 16 studies that have students practice specific problem-solving strategies, usually showing improvement in problem-solving performance. A review by Maries and Singh categorizes the findings of problem-solving research (in physics and other fields) into three factors that underpin effective problem solving: information processing and cognitive load, knowledge organization, and metacognition (including reflection). [2] One weakness of nearly all of the research on physics problem solving, the exception being the work by Heller *et al.* using context-rich problems, is that it deals with solving standard "textbook" type problems [3]. Such problems have an artificial simplified context and provide all the information needed to solve the problem. This makes them quite different from authentic problems that a scientist will encounter in the real-life performance of their work. Thus, while the strategies taught in much of this research can be useful, they are likely to be insufficient when applied to more complex authentic problems.

Problem solving of complex science problems requires more than routine following of procedures; instead, it requires adaptive expertise or applying knowledge and practices to solve novel problems in one's area of expertise [4]. Teaching adaptive expertise has long been a challenge in education, possibly because the specific skills required to become an adaptive expert in problem solving have not been well defined. Price *et al.* [5] addressed this deficiency by characterizing the problem-solving process of expert scientists and engineers through cognitive task analysis interviews [6]. They identified a list of 29 decisions that comprised the problem-solving process. Some examples of these decisions include defining the problem's essential features or concepts that apply,

TABLE I. Problem-solving decisions characterized by Price *et al.* [5] Note that the problem-solving process involves iteration; numbers are for reference, not meant to imply a sequential order of decisions.

Decision title	Description
D1: Importance	What is important in the field?
D2: Fit	What problems and opportunities fit the solver's expertise?
D3: Goals	What are the goals of the problem? What are criteria for and constraints on the solution?
D4: Features	What are important features of the problem and underlying concepts that apply?
D5: Mental model	Which mental model to apply? How to organize existing knowledge when solving this problem?
D6: Narrow	How to narrow down the problem?
D7: Related problems	What are related problems and aspects of their solution process that will help?
D8: Potential solutions	What are potential solutions?
D9: Solvable?	Is the problem solvable?
D10: Simplifications	What approximations or simplifications are appropriate?
D11: Decompose	How to decompose into subproblems with individually solvable pieces?
D12: What is difficult?	Which aspects of the problem are most difficult or uncertain?
D13: Info needed	What information is needed to solve the problem?
D14: Priorities	How to prioritize among competing considerations?
D15: Plan	What is the plan to find or collect needed information?
D16: Calculations	What calculations or data analysis are needed?
D17: Represent	What is the best way to represent and organize information to provide clarity and insights?
D18: Info believable?	How believable is the available (provided or collected) information?
D19: Results vs predictions	How do results compare to predictions based on the mental model?
D20: Any anomalies?	Are there any significant anomalies: If yes, how to follow-up on them?
D21: Conclusions appropriate?	What conclusions are appropriate based on information and analysis?
D22: Choose best solution	What is the best solution?
D23: Reflect on assumptions	Were assumptions and simplifications appropriate?
D24: Reflect on knowledge	Is any additional knowledge or information needed?
D25: Reflect on approach	How well is problem-solving approach working?
D26: Reflect on solution	How adequate is the chosen solution?
D27: Implications	What are broader implications of the results or solution?
D28: Who is audience?	Who is the audience for communication?
D29: How to present?	What is the best way to present the work?

deciding what information is necessary, choosing suitable calculations, and evaluating possible solutions. The full list of decisions is provided in Table I and Table S1 (Supplemental Material [7]). Deliberate practice theory [8] proposes that to become an expert in a skill such as problem solving, students must repeatedly practice and receive feedback on the individual subskills of problem solving. We argue that the 29 problem-solving decisions are the subskills students need to deliberately practice to become expert problem solvers. Work by Burkholder and Wieman has shown that chemical engineering students only enhance their decision-making skills for solving complex problems after intensive capstone courses or internships [9]. Meanwhile, work by Holmes *et al.* found that introductory physics students can improve their critical thinking skills when provided with "structured agency," the act of instructing students to make relevant decisions during their labs [10]. These studies support the notion that to achieve problem-solving expertise, students must repeatedly practice making decisions.

Many studies have characterized and trained the solving of "textbook-style" problems. Various features of the problem-solving process described in this literature are captured by particular decisions. Larkin and Reif [11] characterized the problem-solving processes of one novice and one expert solving standard mechanics "textbook" problems. They concluded that "the problem-solving processes of experts, unlike those of novices, exhibit the following two characteristics: (1) The knowledge of experts is organized into coherent methods, rather than merely into individual principles or equations. (2) Experts approach problem solving by a process of successive refinements. Thus they tend first to consider gross problem features, describing them rather vaguely by words or pictures. Only later do experts consider a problem in greater detail by using more mathematical language" [11]. Thus, experts make decisions D4 (features) and D5 (mental model) when applying their organized knowledge and make D10 (simplifications) and D17 (represent) in their iterative problem-solving process. Similarly, Chi *et al.* [12] found that novices and experts categorize problems differently: according to surface features vs underlying concepts and solution approaches. This categorization relies on the solver's schema or mental model of the problem, which is reflected in making D4 (features) and D5 (mental model). A feature of a problem noted by Simon is whether it is well defined (also called well structured) or ill defined [13]. Problems categorized as ill defined require solvers to make additional decisions compared to problems categorized as well defined. This would include D3 (goals) and likely D13 (info needed). Ill-defined problems are also described as having multiple possible solutions and solution paths [14], so they would require solvers to make D8 (potential solutions) and D25-26 (reflect on approach and reflect on solution). Mayer and Wittrock [15] focus on the transfer of problem-solving skills. This is connected to D7 (related problems). The importance of metacognition highlighted in the review by Maries and Singh [2] is captured through the reflection decisions (D23-26). Heller *et al.* [3] have presented "context-rich" problems for physics instruction. These call on students to make more decisions than in typical physics course problems, and they have outlined a five-step solution strategy that involves making many of the 29 decisions we present here.

Intentionally making any of the 29 decisions requires the appropriate application of relevant knowledge. Thus, the acquisition of relevant knowledge is an essential part of the learning process. In our studies of expert problem solvers, we saw their knowledge was organized so as to be useful in making decisions. Choosing the appropriate mental model to apply to a given problem was therefore an important decision (D5). This implies that knowledge is best learned in the context of how it is actually to be used in making problem-solving decisions. This is supported by the concept of situated cognition of Brown *et al.*, in which they argue that useful knowledge can only be learned by situating it in the context in which it will be used [16]. The work of Schwartz and Bransford also supports this idea, as they have shown that students learn a concept more effectively if, before being told the relevant idea, they first struggle with a problem that is solved by using that idea [17].

It is unreasonable to expect all students at every level should practice making all 29 decisions. Some decisions are clearly only appropriate for learners at the graduate level. It is also clear that it makes more sense to focus on a limited set of decisions for introductory students. The question of when it is best to introduce which decisions is a matter for future research, but it seems evident that one could hope that a student is prepared to make most of the decisions by the time they complete an undergraduate degree in physics.

The primary question we address in this work is whether undergraduate physics majors are getting sufficient practice at making problem-solving decisions in their courses. Although some STEM programs offer final-year hands-on experiences or research opportunities that would involve authentic decision making, primary learning opportunities lie in traditional courses that make up the bulk of time spent in the major. Therefore, we wanted to investigate how much decision-making practice students get from solving problems assigned in their traditional courses. A preliminary version of this work was presented at the Physics Education Research Conference [18].

In this work, we analyzed homework assignments from multiple physics courses to see where decisions were encountered (students likely to make the decision in the process of solving), prompted (a subset of encountered: the problem statement directs students to make the decision), or reduced (a potential opportunity to practice a decision that is circumvented or removed by the problem statement). We also assessed the accuracy of our decision coding by conducting think-aloud interviews. Our research questions are as follows: (i) What problem-solving decisions are students practicing in homework problems used in various levels of physics courses, and do those change over the duration of the course? (ii) Can problems be designed to offer more decision-making practice?

II. METHODS

A. Problems analyzed

We analyzed homework assignments and exams from a series of physics courses required for physics majors at a selective university. These included an introductory mechanics course for very well-prepared freshmen ("intro"), an intermediate statistical mechanics course for juniors and seniors ("intermediate"), and an advanced Lagrangian mechanics course for seniors ("advanced"). The intro and intermediate courses were taught through active learning methods in which students attended mini-lectures and solved problems in groups during class time, while the advanced course was a traditional lecture. The homework problems analyzed for this project were all expected to be solved outside of class. For comparison, we also evaluated the project from a research-project-based senior capstone course ("capstone"), homework problems from another introductory mechanics course that were designed to have students follow a template of problem-solving decisions ("template"), and introductory physics problems that the research team redesigned to incorporate a greater variety of decisions ("decision-rich"). Our research team comprised an undergraduate physics student, an education researcher who specializes in problem solving, and a physics professor who is also an expert in problem-solving research.

B. Decision coding of homework and exam problems

We analyzed the text of each problem to identify which of the problem-solving decisions were present. We followed an iterative coding process in which the student researcher characterized decisions in a subset of problems, then the team discussed and refined the definitions for each decision in the context of actual problems (Table S1 in the Supplemental Material [7]), then the student characterized additional problems and repeated the process. As a first pass to identify decisions, the student researcher solved each problem, documenting every decision made. The team then discussed the problem statement and the student's solution to determine which decisions were necessary to solve the problem and which were reduced by the problem phrasing. Additionally, for select problems, solutions provided by the course instructor were reviewed to ensure a holistic characterization that took into account how the instructors intended for the problems to be solved. The final coding of decisions was based on the problem statement, not on the individual's solution process. Approximately 75% of problems were discussed and agreed upon by at least two members of the team.

Each decision that we identified in a problem statement was classified as either "encountered," "prompted," or "reduced." Encountered decisions were those that students would have to make in order to solve the problem, prompted decisions were a subset of encountered decisions that students were explicitly directed to make by the problem instructions, and reduced decisions were those where a potential decision opportunity was circumvented because the problem statement or accompanying instructions made a decision for the student. On occasion, some decisions overlapped or were interdependent. For example, D15 (plan) was often the same as D16 (calculations) in the context of these problems, because the only information that needed to be collected were quantities to be calculated, so planning how to collect that information was the same as deciding which calculations to carry out. In such cases, we coded both of these decisions as encountered but noted their potential overlap. We chose to still include both overlapping decisions in our data because they are separate decisions in other contexts, such as in the capstone project. On any given problem, most decisions were not given any code (see Table II for average numbers per decision). This was because we did not have sufficient information in the problem statement to conclude a student would be likely to make such a decision when solving the problem. Some decisions are easier to infer from a problem statement and student work than others, reflection decisions being particularly challenging unless they are explicitly prompted or removed.

Figure 1 shows examples of our coding for intro mechanics problems. In Fig. 1(a), D4 (features) is encountered when the student must decide which (if any) coefficient of friction to use, and D16 (calculations) will be encountered in the process of solving for $a_{\min}$. Meanwhile, D17 (represent) is reduced by the problem statement telling the students to draw a force diagram, and D15 (plan) and D11 (decompose) are reduced by the problem specifying exactly how and in what order of steps the student should go about solving. Notice that although D17 (represent) is reduced, it is also prompted because the question prompts the student to define a coordinate system. In Fig. 1(b), we modified the problem to be" decision rich" by reincorporating the decisions that were previously reduced. This was done by removing certain elements of the problem statement that made decisions for the students and prompting students to make D15 (plan) and D10 (simplifications).

We calculated the percentage of problems from each course in which each decision was encountered, prompted, or reduced (see Fig. 2). If a decision was encountered in one part of a problem but reduced in a different part of a problem, we counted the decision as being both encountered and reduced. If a decision was encountered or reduced more than once in a problem, we only counted it once.

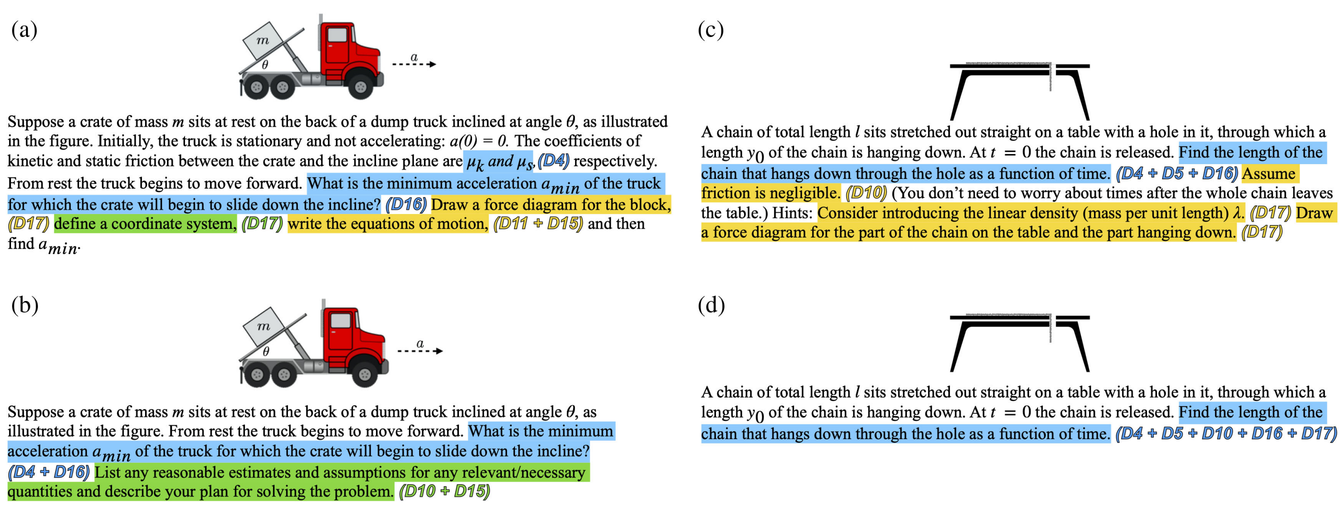

FIG. 1. Original and decision-rich problems used in student interviews. Reduced decisions are notated with yellow highlight, encountered decisions notated with blue highlight, and prompted decision notated with green highlight. (a) Original "truck problem" from intro course. (b) Decision-rich version of the truck problem, modified by the research team to require more decisions. (c) Original "chain problem" from intro course. (d) Decision-rich version of chain problem.

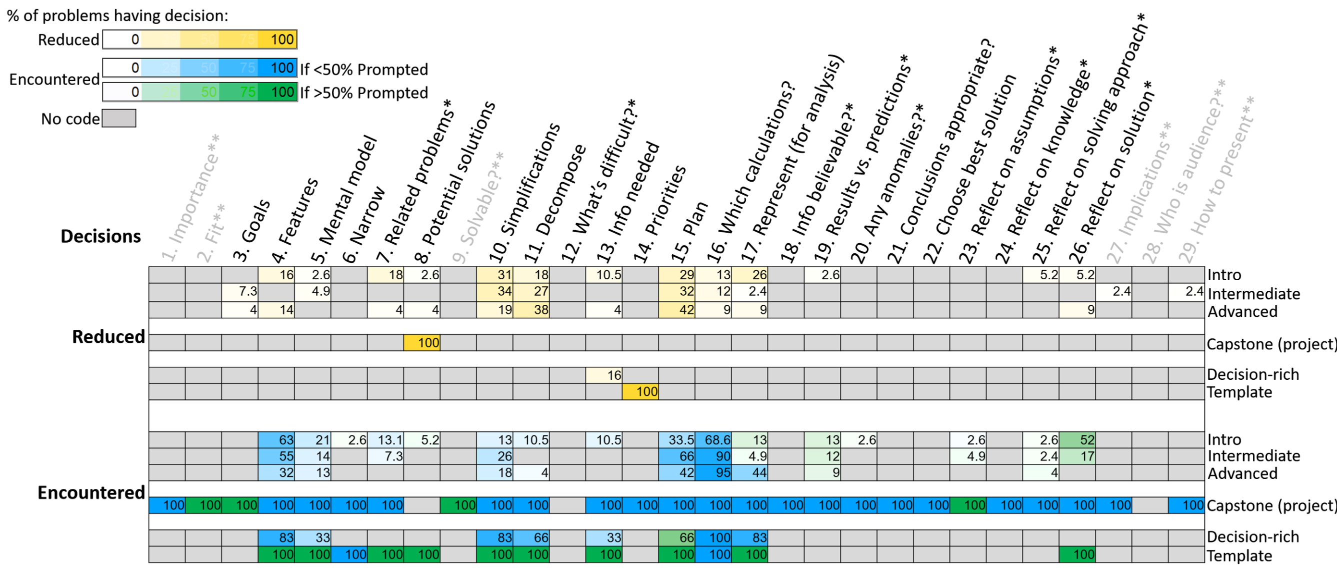

FIG. 2. Percentage of problems in which each decision was reduced (yellow) or encountered (blue or green). Green shading means that decisions were coded as encountered because they were explicitly prompted >50% of the time; blue shading was encountered without (<50%) prompting. Gray shading means the decision was not coded as reduced or encountered in any problems analyzed for that course. Homework problems were analyzed from traditional physics courses: intro mechanics ($n = 38$), intermediate statistical mechanics ($n = 41$), and advanced Lagrangian mechanics ($n = 21$). In contrast, three problem-solving opportunities expected to involve more decisions were also analyzed: the final project of a capstone course (analyzed through retrospective interview, $n = 1$), intro course problems modified by the research team to explicitly include more decisions (decision rich, $n = 6$), and problems from a different intro course that were designed to have students practice problem-solving decisions by following a structured template (template, $n = 2$). Decisions labeled in gray font marked with ** were not expected to be relevant in the context of traditional coursework. Decisions marked with * involve reflection and were more difficult to identify based on problem statement so may be variable depending on student and undercounted here.

C. Decision coding of capstone project

The capstone project was analyzed by conducting a retrospective cognitive task analysis interview [6] in which the student researcher shared the work produced for a term project and described his process for completing the project. The interviewer coded the interview transcript for the decisions mentioned and the research team discussed the coding.

D. Think-aloud interviews for coding validation

When solving problems during the initial coding iteration, the student researcher occasionally noted decisions they made but were not necessary to solve the problem (usually reflection on strategy or solution, for their own interest). We chose not to include those as "encountered" because we could not predict how likely other students would be to make such decisions. Noting the limitation that other students may also decide to make decisions that the problem statement does not necessarily require, and seeking to validate our analysis of the problem text, we conducted five think-aloud interviews with senior physics undergraduate students. Student volunteers were recruited from advanced or intermediate physics courses because we wanted a high likelihood of observing successful solution processes on difficult decision-rich problems. This research was approved by the Stanford IRB (protocol 48785).

Each student solved 2–3 problems, at least one taken from the "intro" course homework and one modified to be decision rich. The students were instructed to think out loud and explain every step in solving. We coded decisions students made during these interviews by carefully observing the students' problem-solving process, and two researchers discussed each interview. Any time we noticed the student explicitly making a decision, we coded which of the 29 decisions it was. For example, when student B was solving the original truck problem, they stated that they were assuming that static friction was large enough such that the box would not initially slip, which represents D10 (simplifications). This analysis allowed us to compare our original coding of the problem to the decisions actually made by this small sample of volunteers.

III. RESULTS AND DISCUSSION

We found that only a few problem-solving decisions were consistently encountered in the three courses analyzed, while almost as many decisions were reduced (see Fig. 2 and Table II). The most commonly encountered decisions in undergraduate physics homework were D4 (features), D15 (plan), and D16 (calculations). Decisions

TABLE II. Maximum and average number of decisions per problem that were encountered (Enc.), reduced (Red.), or prompted (Prompt.). Note that prompted decisions are a subset of encountered.

Course	Number of problems	Max Enc.	Max Red.	Avg Enc.	Avg Red.	Avg Prompt.
Intro	38	8	6	3.3	1.8	0.92
Intermediate	41	6	4	2.8	1.2	0.53
Advanced	21	6	4	2.4	1.7	0.33
Decision rich	6	7	1	5.5	0.16	0.83
Template	2	12	1	12	1	10
Capstone (project)	1	26	1	26	1	4

D19 (results vs predictions) and D26 (reflect on solution) were the most likely to be explicitly prompted. Meanwhile, the most commonly reduced decisions were D10 (simplifications), D11 (decompose), and D15 (plan). On average, students only encountered 2.8 decisions per problem in the traditional courses and had 1.6 decisions reduced per problem. When comparing across course levels, all three courses were similar. By simple two-tailed t tests comparing the numbers of encountered, reduced, or prompted decisions in each course, there were no statistically significant differences after correcting for multiple comparisons ($p > 0.05$ for most comparisons, $p = 0.04$ comparing encountered decisions in intro and advanced and $p = 0.01$ comparing prompted decisions in intro and advanced; given nine total comparisons, none of these are statistically significant). This suggests that there is not a general shift toward more decision making in more advanced courses.

In contrast, the other examples show that physics problems and courses can be designed so that students encounter many more decisions. The project from the research-focused capstone course offered a far wider diversity of decisions—students got to practice making 26 of the 29 decisions. Homework problems could also include a larger number of decisions. The problem-solving template problems consistently prompted ten different decisions per problem and also required two more that were unprompted. The decision-rich problems rarely reduced or prompted decisions and instead required students to make an average of 5.5 decisions per question. This provides a good estimate of the number of problem-solving opportunities that a traditional physics homework problem could be readily modified to provide.

One might hope that courses would shift toward more decision making over time within the course. We plotted the average number of decisions encountered or reduced over time during the quarter. There is variability but clearly no trend toward more decisions (see Fig. 3). While it is possible for courses to deliberately increase decision making over time, these courses did not. We also coded the exam questions from these courses to see whether exams provided different problem-solving decision opportunities than homework. In general, the decisions required for exams and homework problems were similar, except for the intermediate course where notably more decisions were encountered on the midterm and final than on the homework. This may explain why the instructor for that course told us they were surprised at how poorly students performed on the exams—students had not been given sufficient practice at making the necessary decisions before encountering them on the exam.

As stated in the methods, many decisions were not coded in any category for any particular problem. This is because they were neither explicitly prompted nor reduced but we did not believe making them was required in order for a student to solve the problem. Although most of these absent decisions are not required for solving, students could decide to put in extra effort and make them regardless. Some decisions that involve reflection, like D7 (related problems) and D23-26 (reflect), were particularly difficult to identify based on the problem statement unless explicitly prompted or reduced. We also rarely coded D5 (mental model) because, in the context of these courses, the relevant mental model was usually known because of the course's topic that week. If students were given these problems in an exam that covered multiple units or outside of a course structure, we predict they would be more likely to make D5 (mental model), as we saw in the think-aloud interviews.

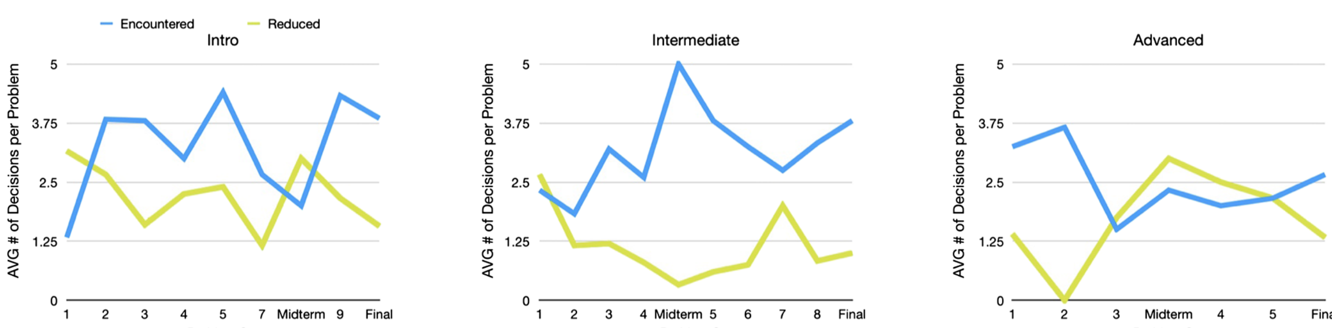

FIG. 3. Average number of decisions encountered (blue) or reduced (yellow) per problem over the course of the quarter for intro, intermediate, and advanced courses. Midterm and final exams were also analyzed and are included here.

Other decisions, such as D1 (importance), D2 (fit), and D27 (implications) are unlikely to be relevant to undergraduate-level problem-solving (these decisions, were also not universally represented in the analysis of expert problem solving by Price *et al.* [5]).

We conducted think-aloud interviews to check the accuracy of our coding, identify hidden decisions that students make but the problem statement does not necessarily require, and get more insight into students' problem-solving processes. We had students solve both original and decision-rich versions of problems taken from the introductory course. Our interview data showed that students did make more decisions on the decision-rich problems: on average, students made 7 decisions when solving the decision-rich truck problem as opposed to 4.5 when solving the original version (see Table III). The interviews confirmed that any decision we coded as encountered was indeed encountered by all of the students. For example, we anticipated that D4 (features) and D16 (calculations) would be encountered when solving any version of the truck problem, and every student ended up making those decisions. Decisions we coded as reduced were also consistently reduced, except when students chose to ignore instructions. For example, student B ended up making an unprompted D17 (represent) to represent the problem in a different way because he ignored the decision-reducing instruction to draw a free-body diagram.

However, some specific decisions made in the interviews, particularly reflection decisions, were highly student dependent. For example, in solving the decision-rich truck problem, students C and D (Table III) made three reflection decisions: D7 (related problems), D25 (reflect on approach), and D26 (reflect on solution), while student E only made D26 (reflect on solution). It is also possible for students to not make any reflection decisions as seen with student A solving the original truck problem. This variability in making reflection decisions is consistent with variability in reflective practices seen in other literature, such as variability in reflection on solution seen by Gjerde *et al.* [19] We also saw student-dependent variation in other decisions. For example, when student A solved the original truck problem, they requested the equation for calculating the force of friction given the coefficient of friction, thus they ended up making D13 (info needed) while no other student made that decision.

The interviews also revealed that students often ignore instructions in the problem. In an extreme example, student B ignored all of the given instructions for the original truck problem and solved the problem by treating gravity as an effective potential. It was not until after they solved the problem that they read the instructions which specified that students solve by drawing a force diagram and determining the equations of motion. This means that decisions coded as "prompted" will not necessarily be made by all students, as can be seen with decisions 10 (simplification) and 15 (plan) for the decision-rich problem. One caveat is that the student volunteers for this study were all senior physics majors, so their process and tendency to ignore instructions may not be representative of students in intro courses.

TABLE III. Encountered, reduced, and prompted decisions from student interviews. Data from the truck problem are presented because all students solved a version of this problem.

Problem and student	Encountered decisions	Reduced decisions	Prompted decisions
Original truck			
Preinterview coding	4, 16	11, 15, 17	17
Student A	4, 5, 13, 16	11, 15, 17	17
Student B	4, 10, 16, 17, 26		
Decision-rich truck			
Preinterview coding	4, 11, 16, 17		10, 15
Student C	4, 5, 7, 11, 16, 17, 25, 26		
Student D	4, 5, 7, 11, 16, 17, 25, 26		
Student E	4, 5, 10, 11, 16, 17, 26		

IV. CONCLUSION

Our results show that the traditional physics courses (intro, intermediate, and advanced) we analyzed only provide students with practice at making a few key problem-solving decisions: D4 (features), D15 (plan), and D16 (calculations). Other decisions that are important for problem solving in the real world are either not encountered or are reduced by being explicitly made for the students. This is mostly due to the nature of the assigned coursework. Often problem scenarios were constrained and provided exactly the information needed, had a single well-defined (sometimes explicitly instructed) pathway to solve the problem, made statements of assumptions to make, or included other "hints." Instructors may unintentionally eliminate decisions for a number of reasons, such as helping guide students, reducing ambiguity to increase student comfort and ensure consistent problem interpretation, reducing difficulty or cognitive load, or easing grading. Unfortunately, if decisions are reduced consistently, this limits the preparation students receive for becoming scientific problem solvers. Since the current curriculum for physics majors typically involves many traditional physics courses and very few research-focused courses (usually as capstone experiences), physics undergraduate students may not be receiving adequate preparation for solving real problems that they will encounter after they graduate.

Within a course, instructors should find it useful to pay attention to the decisions encountered or reduced in their assignments. Research is needed to determine an optimal number or sequence of decisions for students to practice in an assignment, but too many encountered decisions could make a problem more difficult than the instructor expects, while not enough practice at decisions could cause difficulty later. For example, D5 (mental model) was rarely encountered in the homework problems we analyzed because each week was typically dedicated to the same concept, but the same problem if given on an exam might have required students to decide which mental model from concepts throughout the course was most appropriate to apply. This lack of practice could explain why students sometimes do unexpectedly poorly on exams. One might expect that instructors would limit decisions to make problems easier early in a course, or in a more introductory course, then increase opportunities for problem-solving decisions later. However, our analysis did not reveal any trend toward more decision making later in a course or in more advanced courses. The capstone course was a notable exception and provided students with potentially their only experience at making nearly every decision from the problem-solving decisions framework. That course was designed to simulate an authentic research environment. Requiring lots of decisions like this should be a design feature of a good capstone course.

The courses we analyzed are standard requirements for undergraduate physics majors and span a range of difficulty levels, therefore, we believe this work will likely generalize to other universities that have similar courses and assign textbook-style homework problems as the main problem-solving activity. However, our analysis was limited to mostly problems from a few courses at a selective university. Other courses or textbooks may provide students with more decision practice, particularly if they involve a research element or include real-world problems that do not provide all necessary information upfront. Our analysis of the template problems from a problem-solving-focused intro course and the decision-rich problems we redesigned shows that it is possible to create solvable introductory problems that still provide students with opportunities to deliberately practice a reasonable portion of the problem-solving decisions. Interviews with students solving decision-rich problems confirmed that students will generally make more decisions if the problem statement does not reduce as many. In addition, we did not analyze problems solved during class, so it is possible that in-class problem solving in the active learning classes involved more decisions. Indeed, group problem-solving activities described in the literature, such as by Heller [3] or Mason and Singh [20] have been structured to encourage students to make reflection decisions.

There are a variety of ways to design problems that require more decisions to solve. First, if starting with existing problems, eliminate the reduction of decisions. This can be done by removing information or instructions that provide the answers to decisions and, as suitable, replacing these with prompts to make a particular decision. Second, include insufficient or extraneous information in problems, so students have to recognize or seek out some relevant information. A special and valuable case of this is having students make reasonable assumptions and simplifications. A desirable feature that is more challenging to achieve is to create problems that have multiple possible solution paths. Starting with a real-world scenario in creating the problem can be valuable, in that the real-world context is naturally more complex and realistic.

When designing a problem, it is important to consider both how many decisions are encountered and how many decisions are reduced. Other research on problem solving in physics has demonstrated that there can be unintended consequences of reducing decisions about planning D15 (plan) and D17 (represent) [21,22]. In cases where instructors feel students need more support in their problem-solving process, they could opt to prompt a decision rather than eliminate it for the students [23]. Decision prompting has been demonstrated to be a viable approach to incorporate decision making into introductory physics labs as well [10]. However, our interviews also provided useful insight into prompting. We saw many examples of students ignoring explicit instructions, so instructors should not rely only on prompts to provide decision practice. Interestingly, we noticed that students who ignored instructions could end up making either more or fewer decisions than intended by not making prompted decisions but still making reduced decisions. For more structure, instructors could use a consistent template of problem-solving decisions that students are expected to make and train students to respond to those prompts by grading accordingly. In fact, the course that used the template problems needed to train students to follow the template. Online dynamic problems could also provide promising approaches for structuring problems that require and support students in making more decisions.

Our analysis has some limitations. First, it is impossible to eliminate the subjectivity of a problem solver during the process of characterizing problem-solving decisions. We attempted to limit inconsistency by discussing a large portion of the problems and comparing student and instructor solutions. A second limitation is that our analysis only coded for decisions that were clearly encountered or reduced. As seen from our interviews, students could put in a conscious effort to make additional decisions that were absent in our coding. Third, some decisions are easier to infer from a problem statement than others. Our student interviews, while limited by being from a small sample of upper-level students at a selective university, provided some additional insight. Decision 25 (reflect on problem-solving approach) and D26 (reflect on solution) are encountered far more frequently than what we inferred from the problem statements alone. These decisions can differ from student to student, meaning that some students reflect more than others. We also hypothesize that D5 (mental model), D7 (related problems), D10 (simplifications), and D13 (info needed) are highly dependent on a student's level of experience. While more experienced students have a higher number of past encounters with related problems and can decide that one mental model for solving is more appropriate than another, less experienced students may end up having to look up and decide what info is needed. We observed some variability in making these decisions in our interviews with upper-level students, which was likely related to specific content expertise, but we did not interview introductory students so cannot say whether different populations of students would be more or less likely to make particular decisions when solving a given problem.

This work raises important empirical questions for future research: what is the optimal number of decisions for students of different experience levels to practice in a single assignment, especially as they are learning both content knowledge and problem-solving skills? Is there an optimal sequence of which decisions to practice at earlier levels? Ericsson's work on deliberate practice [8] provides some guidance by showing in multiple contexts how students need repeated practice, with feedback, performing desired subskills in order to develop expertise in a field. If the goal of physics programs is to help students develop expertise in physics problem solving, students must therefore be provided with repeated practice at all of the problem-solving subskills, which Price *et al.* [5] argue are defined by the problem-solving decisions. We expect it is insufficient to give students their only opportunity to practice these decisions in a capstone course as they complete the major. Although it is not feasible to have an introductory-level student practice all of the decisions in a single problem, they could be supported by gradually increasing the number of decisions required throughout a course. Instructors could also provide support in terms of prompted scaffolding that fades over time. Our work shows that increasing course difficulty does not necessarily increase the number of opportunities students will have to make decisions, so problem-solving decisions need to be explicitly planned at all levels. When students are tasked with solving problems in their future careers (whether in physics or not), they will need to make nearly all the problem-solving decisions during their solution process. Therefore, undergraduate programs need to consciously give students more opportunities to practice making these decisions during their coursework.

ACKNOWLEDGMENTS

This work was funded by a Howard Hughes Medical Institute professor's grant to CEW. We also thank our student interviewees and the instructors of courses analyzed.

[1] E. Ince, An overview of problem solving studies in physics education, J. Educ. Learn. Arch. **7**, 191 (2018).
[2] A. Maries and C. Singh, Helping students become proficient problem solvers part I: A brief review, Educ. Sci. **13**, 156 (2023).
[3] P. Heller, R. Keith, and S. Anderson, Teaching problem solving through cooperative grouping. Part 1: Group versus individual problem solving, Am. J. Phys. **60**, 627 (1992).
[4] G. Hatano and K. Inagaki, Two courses of expertise, in *Child Development and Education in Japan* (W. H. Freeman/Times Books/Henry Holt & Co, New York, NY, 1986), pp. 262–272.
[5] A. M. Price, C. J. Kim, E. W. Burkholder, A. V. Fritz, and C. E. Wieman, A detailed characterization of the expert problem-solving process in science and engineering: Guidance for teaching and assessment, CBE Life Sci. Educ. **20**, ar43 (2021).
[6] R. Hoffman and G. Lintern, Eliciting and representing the knowledge of experts, in *The Cambridge Handbook of Expertise and Expert Performance* (Cambridge University Press, New York, NY, 2006), pp. 203–222.
[7] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevPhysEducRes.20.020103 for supplementary table S1: Definitions of problem-solving decisions in the context of physics problems, used for coding.
[8] K. A. Ericsson, The influence of experience and deliberate practice on the development of superior expert performance, in *The Cambridge Handbook of Expertise and Expert Performance* (Cambridge University Press, New York, NY, 2006), pp. 683–703.
[9] E. W. Burkholder and C. E. Wieman, Comparing problem-solving across capstone design courses in chemical engineering, in *Proceedings of the IEEE Frontiers in Education Conference, Uppsala, Sweden* (IEEE, New York, 2020), 10.1109/FIE44824.2020.9273820.
[10] N. G. Holmes, B. Keep, and C. E. Wieman, Developing scientific decision making by structuring and supporting student agency, Phys. Rev. Phys. Educ. Res. **16**, 010109 (2020).
[11] J. H. Larkin and F. Reif, Understanding and teaching problem solving in physics, Eur. J. Sci. Educ. **1**, 191 (1979).
[12] M. T. H. Chi, P. J. Feltovich, and R. Glaser, Categorization and representation of physics problems by experts and novices, Cogn. Sci. **5**, 121 (1981).
[13] H. A. Simon, The structure of ill structured problems, Artif. Intell. **4**, 181 (1973).

[14] D. H. Jonassen, Instructional design models for well-structured and iII-structured problem-solving learning outcomes, Educ. Technol. Res. Dev. **45**, 65 (1997).

[15] R. E. Mayer and M. C. Wittrock, Problem-solving transfer, in *Handbook of Educational Psychology* (Prentice Hall International, London, England, 1996), pp. 47–62.

[16] J. S. Brown, A. Collins, and P. Duguid, Situated cognition and the culture of learning, Educ. Res. **18**, 32 (1989).

[17] D. L. Schwartz and J. D. Bransford, A time for telling, Cognit. Instr. **16**, 475 (1998).

[18] B. J. Montgomery, A. M. Price, and C. E. Wieman, How traditional physics coursework limits problem-solving opportunities, presented at PER Conf. 2023, Sacramento, CA, 10.1119/perc.2023.pr.Montgomery.

[19] V. Gjerde, B. Holst, and S. D. Kolstø, Integrating effective learning strategies in basic physics lectures: A thematic analysis, Phys. Rev. Phys. Educ. Res. **17**, 010124 (2021).

[20] A. Mason and C. Singh, Using categorization of problems as an instructional tool to help introductory students learn physics, Phys. Educ. **51**, 025009 (2016).

[21] A. F. Heckler, Some consequences of prompting novice physics students to construct force diagrams, Int. J. Sci. Educ. **32**, 1829 (2010).

[22] E. Kuo, N. R. Hallinen, and L. D. Conlin, When procedures discourage insight: Epistemological consequences of prompting novice physics students to construct force diagrams, Int. J. Sci. Educ. **39**, 814 (2017).

[23] E. W. Burkholder, J. K. Miles, T. J. Layden, K. D. Wang, A. V. Fritz, and C. E. Wieman, Template for teaching and assessment of problem solving in introductory physics, Phys. Rev. Phys. Educ. Res. **16**, 010123 (2020).