

Quantifying critical thinking: Development and validation of the physics lab inventory of critical thinking

Cole Walsh,^{1,*} Katherine N. Quinn,¹ C. Wieman,² and N. G. Holmes¹

¹Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, New York 14853, USA

²Department of Physics and Graduate School of Education, Stanford University, Stanford, California 94305, USA



(Received 21 January 2019; published 31 May 2019)

Introductory physics lab instruction is undergoing a transformation, with increasing emphasis on developing experimentation and critical thinking skills. These changes present a need for standardized assessment instruments to determine the degree to which students develop these skills through instructional labs. In this article, we present the development and validation of the physics lab inventory of critical thinking (PLIC). We define critical thinking as the ability to use data and evidence to decide what to trust and what to do. The PLIC is a 10-question, closed-response assessment that probes student critical thinking skills in the context of physics experimentation. Using interviews and data from 5584 students at 29 institutions, we demonstrate, through qualitative and quantitative means, the validity and reliability of the instrument at measuring student critical thinking skills. This establishes a valuable new assessment instrument for instructional labs.

DOI: [10.1103/PhysRevPhysEducRes.15.010135](https://doi.org/10.1103/PhysRevPhysEducRes.15.010135)

I. INTRODUCTION

More than 400 000 undergraduate students enroll in introductory physics courses at post-secondary institutions in the United States each year [1]. In most introductory physics courses, students spend time learning in lectures, recitations, and instructional laboratories (labs). Labs are notably the most resource intensive of these course components, as they require specialized equipment, space, facilities, and occupy a significant amount of student and instructional staff time [2,3].

In many institutions, labs are traditionally intended to reinforce and supplement learning of scientific topics and concepts introduced in lecture [2,4–7]. Previous work, however, has called into question the benefit of hands-on lab work to verify concepts seen in lecture [2,6–10]. Traditional labs have also been found to deteriorate students' attitudes and beliefs about experimental physics, while labs that aimed to teach skills were found to improve students' attitudes [10].

Recently, there have been calls to shift the goals of laboratory instruction towards experimentation skills and practices [11–13]. This shift in instructional targets

provides renewed impetus to develop and evaluate instructional strategies for labs. While much discipline-based education research has evaluated student learning from lectures and tutorials, there is far less research on learning from instructional labs [2,4–6,14,15]. There exist many open research questions regarding what students are currently learning from lab courses, what students could be learning, and how to measure that learning.

The need for large-scale reform in lab courses goes hand in hand with the need for validated and efficient evaluation techniques. Just as previous efforts to reform physics instruction have been unified around shared assessments [16–19], larger scale change in lab instruction will require common assessment instruments to evaluate that change [4–6,15]. Currently, few research-based and validated assessment instruments exist for labs. The website PhysPort.org [20], an online resource to support physics instructors with research-based teaching resources, has amassed 92 research-based assessment instruments for physics courses as of this publication. Only four are classified as evaluating lab skills, all of which focus on data analysis and uncertainty [21–24].

To address this need, we have developed the physics lab inventory of critical thinking (PLIC), an instrument to assess students' critical thinking skills when conducting experiments in physics. The intent was to provide an efficient, standardized way for instructors at a range of institutions to assess their instructional labs in terms of the degree to which they develop students' critical thinking skills. Critical thinking is defined here as the ways in which

*cjw295@cornell.edu

one uses data and evidence to make decisions about what to trust and what to do. In a scientific context, this decision making involves interpreting data, drawing accurate conclusions from data, comparing and evaluating models and data, evaluating methods, and deciding how to proceed in an investigation. These skills and behaviors are commonly used by experimental physicists [25], but are also relevant to students regardless of their future career paths, academic or otherwise [12]. Being able to make sense of data, evaluate whether they are reliable, compare them to models, and decide what to do with them is important whether thinking about introductory physics experiments, interpreting scientific reports in the media, or making public policy decisions. This fact is particularly relevant given that only about 2% of students taking introductory physics courses at U.S. institutions will graduate with a physics degree [1,26].

In this article, we outline the theoretical arguments for the structure of the PLIC, as well as the evidence of validity and reliability through qualitative and quantitative means.

II. ASSESSMENT FRAMEWORK

In what follows, we use arguments from previous literature to motivate the structure and format of the PLIC. We also compare the goals of the PLIC to existing instruments.

A. Design concepts

As discussed above, the goals of the PLIC are to measure critical thinking to fill existing needs and gaps in the existing literature on lab courses [25,27–29] and reports on undergraduate science, technology, engineering, and math (STEM) instruction [13–15]. Probing the *what to trust* component of critical thinking involves testing students' interpretation and evaluation of experimental methods, data, and conclusions. Probing the *what to do* component involves testing what students think should be done next with the information provided.

Critical thinking is considered content dependent: "Thought processes are intertwined with what is being thought about" [[30], p. 22]. This suggests that the assessment questions should be embedded in a domain-specific context, rather than be generic or outside of physics. The content associated with that context, however, should be accessible to students with as minimal physics content knowledge as possible to ensure that the instrument is assessing critical thinking rather than students' content knowledge. We chose to use a single familiar context to optimize the use of students' assessment time.

We chose to have students evaluate someone else's data rather than conduct their own experiment (as in a practical lab exam) for several reasons related to the purposes of standardized assessment. First, collecting data requires either common equipment (which creates logistical

constraints and would inevitably reduce its accessibility across institutions) or an interactive simulation, which would create additional design challenges. Namely, it is unclear whether a simulation could sufficiently represent authentic measurement variability and limitations of physical models, which are important for evaluating the *what to trust* aspect of critical thinking. Second, by evaluating another person's work, students are less likely to engage in performance goals [31–33], defined as situations "in which individuals seek to gain favorable judgments of their competence or avoid negative judgments of their competence" [[32], p. 1040], potentially limiting the degree to which they would be critical in their thinking. Third, it constrains what they are thinking critically about, allowing us to precisely evaluate a limited and well-defined set of skills and behaviors. By interweaving test questions with descriptions of methods and data, the assessment scaffolds the critical thinking process for the students, again narrowing the focus to target particular skills and behaviors with each question. Fourth, by designing hypothetical results that target specific experimental issues, the assessment can look beyond students' declarative knowledge of laboratory skills to study their enacted critical thinking.

The final considerations were that the instrument be closed response and freely available, to meet the goal of providing a practically efficient assessment instrument. The need for a closed-response assessment facilitates automated scoring of student work, making it more likely to be used by instructors. A freely available instrument facilitates its use at a wide range of institutions and for a variety of purposes.

B. Existing instruments

A number of existing instruments assess skills or concepts related to learning in labs, but none sufficiently meet the design concepts outlined in the previous section (see Table I).

The specific skills to be assessed by the PLIC are not covered in any single existing instrument (Table I, columns 1–4). Several of these instruments are also not physics specific (Table I, column 6). Given that critical thinking is considered content dependent [30] and assessment questions should be embedded in a domain-specific context, these instruments do not appropriately assess learning in physics labs. The physics measurement questionnaire (PMQ), Lawson test of scientific reasoning, and critical thinking assessment (CAT) use an open-response format (Table I, column 5). This places significant constraints on the use of these assessments in large introductory classes, the target audience of the PLIC. A closed-response version of the PMQ was created [44], but limited tests of validity or reliability were conducted on the instrument. The Critical Thinking Assessment Test (CAT) and the California Critical Thinking Skills Test are also not freely available to instructors (Table I, column 7). The cost associated with

TABLE I. Summary of the target skills and assessment structure of existing instruments for evaluating lab instruction.

Assessment	Skills or concepts assessed				Structure or features		
	Evaluating data	Evaluating methods	Evaluating conclusions	Proposing next steps	Closed response	Physics specific	Freely available
Lawson test of scientific reasoning [34]		X	X				X
Critical thinking assessment test (CAT) [35,36]	X		X				
California critical thinking skills test [37]			X		X		
Test of scientific literacy skills (TOSLS) [38]	X	X	X		X		X
Experimental design ability test (EDAT) [39]		X			X		X
Biological experimental design concept inventory (BEDCI) [40]		X			X		X
Concise data processing assessment (CDPA) [21]	X				X	X	X
Data handling diagnostic (DHD) [41]	X				X	X	X
Physics measurement questionnaire (PMQ) [42]	X	X				X	X
Measurement uncertainty quiz (MUQ) [43]	X	X		X	X	X	X
Physics lab inventory of critical thinking (PLIC)	X	X	X	X	X	X	X

using these instruments places significant logistical burden on dissemination and broad use, failing our goal to support instructors at a wide range of institutions.

The measurement uncertainty quiz (MUQ) is the closest of these instruments to meeting the PLIC in assessment goals and structure, but the scope of the concepts covered narrowly focuses on issues of uncertainty. For example, respondents taking the MUQ are asked to propose next steps to specifically reduce the uncertainty. The PLIC allows respondents to propose next steps more generally, such as to consider extending the range of the investigation or testing additional variables.

C. PLIC content and format

The PLIC provides respondents with two case studies of groups completing an experiment to test the relationship between the period of oscillation T of a mass hanging from a spring given by

$$T = 2\pi\sqrt{\frac{m}{k}}, \tag{1}$$

and the variables on which it depends: the spring constant k and the mass m . The model makes a number of assumptions, such as that the mass of the spring is negligible, though these are not made explicit to the respondents.

The mass on a spring content is commonly seen in high school and introductory physics courses, making it relatively accessible to a wide range of respondents. All required information to describe the physics of the problem, including relevant equations, are given at the beginning of the assessment. Additional content knowledge related to simple harmonic motion is not necessary for answering the questions. This claim was confirmed through interviews with students (described in Secs. III and IV B).

This mass on a spring scenario is also rich in critical thinking opportunities, which we had previously observed during real lab situations [45].

Respondents are presented with the experimental methods and data from two hypothetical groups:

- (1) The first hypothetical group uses a simple experimental design that involves taking multiple repeated measurements of the period of oscillation of two different masses. The group then calculates the means and standard uncertainties in the mean (standard errors) for the two data sets, from which they calculate the values of the spring constant k for the two different masses.
- (2) The second hypothetical group takes two measurements of the period of oscillation at many different masses (rather than repeated measurements at the same mass). The analysis is done graphically to evaluate the trend of the data and compare it to the predicted model. That is, plotting T^2 vs m should produce a straight line through the origin. The data, however, show a mismatch between the theoretical prediction and the experimental results. The second group subsequently adds an intercept to the model, which improves the quality of the fit, raising questions about the idealized assumptions in the given model.

PLIC questions are presented on four pages: one page for group 1, two pages for group 2 (one with the one-parameter fit, and the other adding the variable intercept to their fit), and one page comparing the two groups. There are a combination of question formats, including five-point Likert-scale questions, traditional multiple choice questions, and multiple response questions. Examples of the different question formats are presented in Table II. The multiple choice questions each contain three response choices from which respondents can select one. The multiple response

TABLE II. Question formats used on the PLIC along with the types and examples of questions for each format.

Format	Types	Examples
Likert	Evaluate data	How similar or different do you think group 1's spring constant (k) values are?
	Evaluate methods	How well do you think group 1's method tested the model?
Multiple choice	Compare fits	Which fit do you think group 2 should use?
	Compare groups	Which group do you think did a better job of testing the model?
Multiple response	Reasoning	What features were most important in comparing the fit to the data?
	What to do next	What do you think group 2 should do next?

questions each contain 7–18 response choices and allow respondents to select up to three. The decision to use this format is discussed further in the next section. An example multiple choice and multiple response question from the fourth page of the PLIC is shown in Appendix B, Fig. 5 and the entire assessment can be found at Ref. [20].

III. EARLY EVALUATIONS OF CONSTRUCT VALIDITY

An open-response version of the PLIC was iteratively developed using interviews with students and written responses from several hundred students, all enrolled in physics courses at universities and community colleges. These responses were used to evaluate the construct validity of the assessment, defined as how well the assessment measures critical thinking, as defined above.

Six preliminary think-aloud interviews were conducted with students in introductory and upper-division physics courses, with the aim to evaluate the appropriateness of the mass-on-a-spring context. In all interviews, students demonstrated familiarity with the experimental context and successfully progressed through all the questions. The nature of the questions, such that each is related to but independent of the others, allowed students to engage with subsequent questions even if they struggled with earlier ones; a student is able to answer and receive full credit for later questions whether or not they provided the correct response to earlier questions. All students expressed familiarity with the equipment, physical models, and methods but did not have explicit experience with evaluating the limitations of the model. The interviews also indicated that students were able to progress through the questions and think critically about the data and model without taking their own data.

The interviews revealed a number of ways to improve the instrument by refining wording and the hypothetical scenarios. For example, in an early draft, the survey included only the second hypothetical group who fit their data to a line to evaluate the model. Interviewees were asked what they thought it meant to “evaluate a model.” Some responded that to evaluate a model one needed to identify where the model breaks down, while others

described obtaining values for particular parameters in the model (in this case, the spring constant). The latter is a common goal of experiments in introductory physics labs [25]. In response to these variable definitions, the first hypothetical group was added to the survey to better capture students' perspectives of what it meant to evaluate a model (i.e., to find a parameter). This revision also offered the opportunity to explore respondents' thinking about comparing pairs of measurements with uncertainty.

The draft open-response instrument was administered to students at multiple institutions described in Table III in four separate rounds. After each round of open-response analysis, the survey context, and questions, underwent extensive revision. The surveys analyzed in rounds 1 and 2 provided evidence of the instrument's

TABLE III. Summary of the number of open-response PLIC surveys that were analyzed for early construct validity tests and to develop the closed-response format. The level of the class (first-year or beyond-first-year) and the point in the semester when the survey was administered are indicated.

Institution	Class level	Survey	Round			
			1	2	3	4
Stanford University	FY	Pre	31
		Post	...	30
		Pre	170
University of Maine	FY	Mid	120	93
		Post	189
		Pre	25	1
Foothill College	BFY	Post	3
		Pre	...	19	40	...
		Post	...	19	78	...
University of British Columbia	FY	Pre	10	107
		Post	38
University of Connecticut	FY	Pre	22
		Post	3
Cornell University	FY	Pre	89
		Post	99
		Pre	35
St. Mary's College of Maryland	BFY	Post	29
		Pre	2
St. Mary's College of Maryland	FY	Pre	2
		Post	3

ability to discriminate between students' critical thinking levels. For example, in cases where the survey was administered as a post-test, only 36% of the students attributed the disagreement between the data and the predicted model to a limitation of the model. Most students were unable to identify limitations of the model even after instruction. This provided early evidence of the instrument's dynamic range in discriminating between students' critical thinking levels. The different analyses conducted in rounds 1 and 2 will be clarified in the next section.

Students' written responses also revealed significant conceptual difficulties about measurement uncertainty, consistent with other work [46–48]. The first question on the draft instrument asked students to list the possible sources of uncertainty in the measurements of the period of oscillation. Students provided a wide range of sources of uncertainty, systematic errors (or systematic effects), and measurement mistakes (or literal errors) in response to that prompt [49]. It was clear that these questions were probing different reasoning than the instrument intended to measure, and so were ultimately removed.

The written responses also exposed issues with the data used by the two groups. For example, the second group originally measured the time for 50 periods for each mass and attached a timing uncertainty of 0.1 sec for each measured time (therefore 0.02 sec per period). Several respondents thought this uncertainty was unrealistically small. Others said that no student would ever measure the time for 50 periods. In the revised version, both groups measure the time for five periods.

The survey was also distributed to 78 experts (faculty, research scientists, instructors, and post-docs) for responses and feedback, leading to additional changes and the development of a scoring scheme (Sec. IV C). Experts completing the first version of the survey frequently responded in ways that were consistent with how they would expect their students to complete the experiment, which some described in open-response text boxes throughout the survey. We intended for experts to answer the survey by evaluating the hypothetical groups using the same standards that they would apply to themselves or their peers, rather than to students in introductory labs. To make this intention clearer, we changed the survey's language to present respondents with two case studies of hypothetical "groups of physicists" rather than "groups of students."

IV. DEVELOPMENT OF CLOSED-RESPONSE FORMAT

To develop and evaluate the closed-response version of the PLIC, we coded the written responses from the open-response instrument, conducted structured interviews, and used responses from experts to generate a scoring scheme. We also developed an automated system to communicate with instructors and collect responses from students, facilitating ease of use and extensive data collection.

A. Generation of the closed response version

In the first round of analysis, common student responses from the open-response versions were identified through emergent coding of responses collected from the courses listed in Table III, round 1 ($n = 66$ test responses). The surveys analyzed in round 1 only include those given before instruction. In the second round, a new set of responses were coded based on this original list of response choices, and any ideas that fell outside of the list were noted. If new ideas came up several times, they were included in the list of response choices for subsequent survey coding. Combined, open response surveys from 134 students across four institutions were used to generate the initial closed-response response choices (Table III, rounds 1 and 2). The majority of the open responses coded in round 2 were captured by the response choices created in round 1, with only one response per student, on average, being coded as *other*. Half of the *other* responses were uninterpretable or unclear (for example, "collect more data" with no indication of whether more data meant additional repeated trials, additional mass values, or additional oscillations per measurement). A third round of open-response surveys were coded (Table III, round 3) to evaluate whether any of the newly generated response choices should be dropped (because they were not generated by enough students overall), whether response choices could be reworded to better reflect student thinking, or if there were any common response choices missing. Researchers regularly checked answers coded as *other* to ensure no common responses were missing, which led to the inclusion of several new response choices.

While it was relatively straightforward to categorize students' ideas into closed-response choices, students typically listed multiple ideas in response to each question. A multiple response format was employed to address this issue. Response choices are listed in neutral contexts to address the issue of different respondents preferring positive or negative terms (e.g., the number of masses, rather than many masses and few masses).

The closed-response questions, format, and wording were revised iteratively in response to more student interviews and collection of responses from the preliminary closed-response version, which was administered online to the institutions listed in Table III, round 4. To ensure reliability across student populations, subsets of these students were randomly assigned an open-response survey. Additional coding of these surveys were checked against the existing closed-response choices to again identify additional response choices that could be included and to compare student responses between open- and closed-response surveys. The number of open-response surveys analyzed from these institutions is summarized in Table III, round 4.

B. Interview analysis

Two sets of interviews were conducted and audio recorded to evaluate the closed-response version of the

TABLE IV. Demographic breakdown of students interviewed to probe the validity of the closed-response PLIC.

Category	Breakdown	Set 1	Set 2
Major	Physics or related field	8	3
	Other STEM	4	6
Academic level	Freshman	6	3
	Sophomore	3	2
	Junior	1	3
	Senior	0	1
	Graduate student	2	0
Gender	Women	6	5
	Men	6	4
Race or ethnicity	African American	2	1
	Asian or Asian American	3	3
	White or Caucasian	5	4
	Other	2	1

PLIC. A demographic breakdown of students interviewed is presented in Table IV.

In the first set of interviews, participants completed the closed-response version of the PLIC in a think-aloud format. One goal of the interviews was to ensure that the instrument was measuring critical thinking without testing physics content knowledge, using a broader pool of participants. All participants were able to complete the assessment, including nonphysics majors, and there was no indication that physics content knowledge limited or enhanced their performance. The interviews identified some instances where wording was unclear, such as statements about residual plots.

A secondary goal was to identify the types of reasoning participants employed while completing the assessment (whether critical thinking or other). We observed three distinct reasoning patterns that participants adopted while completing the PLIC [50]: (i) *selecting all* (or almost all) possible choices presented to them, (ii) *cueing* to keywords, and (iii) carefully considering and *discerning* among the response choices. The third pattern presented the strongest evidence of critical thinking. To reduce the select all behavior, respondents are now limited to selecting no more than three response choices per question.

In the second set of interviews, participants completed the open-response version of the PLIC in a think-aloud format, and then were given the closed-response version. The primary goal was to assess the degree to which the closed-response options cued students' thinking. For each of the four pages of the PLIC, participants were given the descriptive information and Likert and multiple choice questions on the computer, and were asked to explain their reasoning out loud, as in the open-response version. Participants were then given the closed-response version of that page before moving on to the next page in the same fashion.

When completing the closed-response version, all participants selected the answer(s) that they had previously generated. Eight of the nine participants also carefully read through the other choices presented to them and selected additional responses they had not initially generated. Participants typically generated one response on their own and then selected one or two additional choices. This behavior was representative of the *discerning* behavior observed in the first round of interviews and provided evidence that limiting respondents to selecting no more than three response choices prompts respondents to be discerning in their choices. No participant generated a response that was not in the closed-response version, or expressed a desire to select a response that was not presented to them.

While taking the open-response version, two of the nine participants were hesitant to generate answers aloud to the question "What do you think group 2 should do next?", however, no participant hesitated to answer the closed-response version of the questions. Another participant expressed how hard they thought it was to select no more than three response choices in the closed-response version because "these are all good choices," and so needed to carefully choose response choices to prioritize. It is possible that the combination of open-response questions followed by their closed-response partners could prompt students to engage in more critical thinking than either set of questions alone. Future work will evaluate this possibility. The time associated with the paired set of questions, however, would likely make the assessment too long to meet our goal for an efficient assessment instrument.

C. Development of a scoring scheme

The scoring scheme was developed through evaluations of 78 responses to the PLIC from expert physicists (faculty, research scientists, instructors, and post-docs). As discussed in Sec. II C, the PLIC uses a combination of Likert questions, traditional multiple choice questions, and multiple response questions. This complex format meant that developing a scoring scheme for the PLIC was nontrivial.

1. Likert and multiple choice questions

The Likert and multiple choice questions are not scored. Instead, the assessment focuses on students' reasoning behind their answers to the Likert and multiple choice questions. For example, whether a student thinks group 1's method effectively evaluated the model is less indicative of critical thinking than why they think it was effective.

The Likert and multiple choice questions are included in the PLIC to guide respondents in their thinking on the reasoning and what to do next questions. Respondents must answer (either implicitly or explicitly) the question of how well a group tested the model before providing

their reasoning. The PLIC makes these decision processes explicit through the Likert and multiple choice questions.

2. Multiple response

There were several criteria for determining the scoring scheme for the PLIC. First and foremost, the scoring scheme should align with how experts answer. Expert responses indicated that there was no single correct answer for any of the questions. As indicated in the previous sections, all response choices were generated by students in the open-response versions, and so many of the responses are reasonable choices. For example, when seeing data that are in tension with a given model, it is reasonable to check the assumptions of the model, test other possible variables, or collect more data. An all-or-nothing scoring scheme, where respondents receive full credit for selecting *all* of the correct responses and *none* of the incorrect responses [51], was therefore inappropriate. There needed to be multiple ways to obtain full points for each question, in line with how experts responded.

We assign values to each response choice equal to the fraction of experts who selected the response (rounded to the nearest tenth). As an example, the first *reasoning* question on the PLIC asks respondents to identify what features were important for comparing the spring constants k from the two masses tested by group 1. About 97% of experts identified “the difference between the k values compared to the uncertainty” (R1) as being important. Therefore, we assign a value of 1 for this response choice. About 32% identified “the size of the uncertainty” (R2) as being important, and so we assign a value of 0.3 for this response choice. All other response choices received support from less than 12% of experts and so are assigned values of 0 or 0.1, as appropriate. These values are valid in so far as the fraction of experts selecting a particular response choice can be interpreted as the relative correctness of the response choice. These values will continue to be evaluated as additional expert responses are collected.

Another criteria was to account for the fact that respondents may choose as many as zero to three response choices per question. Accordingly, we sum the total value of responses selected and divide by the maximum value of the number of responses selected:

$$\text{Score} = \frac{\sum_{n=1}^i V_n}{V_{\max_i}}, \quad (2)$$

where V_n is the value of the n th response choice selected and V_{\max_i} is the maximum attainable score when i response choices are selected. Explicitly, the values of V_{\max_i} are

$$\begin{aligned} V_{\max_1} &= \text{Highest value}, \\ V_{\max_2} &= (\text{Highest value}) + (\text{Second highest value}), \\ V_{\max_3} &= (\text{Highest value}) + (\text{Second highest value}) \\ &\quad + (\text{Third highest value}). \end{aligned} \quad (3)$$

If a respondent selects N responses, then they will obtain full credit if they select the N highest valued responses. In the example above, a respondent selecting one response must select R1 in order to receive full credit. A respondent selecting two responses must select both R1 and R2 to receive full credit. A respondent selecting three responses must select R1 and R2, as well as the third highest valued response to receive full credit.

The scoring scheme rewards picking highly valued responses more than it penalizes for picking low-valued responses. For example, respondents selecting R1, R2, and a third response choice with value zero will receive a score of 0.93 on this question. The scoring scheme does not necessarily penalize for selecting more or fewer responses. For example, a respondent who picks the three highest-valued responses will receive the same score as a respondent who picks only the two highest-valued responses. This is true even if the third highest-valued response may be considered a poor response (with few experts selecting it). Fortunately, most questions on the PLIC have a viable third response choice.

The respondent’s overall score on the PLIC is obtained by summing scores on each of the multiple response questions. Thus, the maximum attainable score on the PLIC is 10 points. Using this scheme, the 78 experts obtained an average overall score of 7.6 ± 0.2 . The distribution of these scores is shown in Fig. 3 in comparison to student scores (discussed in Sec. V G). Here and throughout, uncertainties in our results are given by the 95% confidence interval (i.e., 1.96 multiplied by the standard error of the quantity).

D. Automated administration

The PLIC is administered online via Qualtrics as part of an automated system adapted from Ref. [52]. Instructors complete a course information survey (CIS) through a web link (available at Ref. [53]) and are subsequently sent a unique link to the PLIC for their course. The system handles reminder emails, updates, and the activation and deactivation of pre- and postsurveys using information provided by the instructor in the CIS. Instructors are also able to update the activation and deactivation dates of one or both of the pre- and post-surveys via another link hosted on the same web page as the CIS. Following the deactivation of the postsurvey link, instructors are sent a summary report detailing the performance of their class compared to classes of a similar level.

V. STATISTICAL RELIABILITY TESTING

Below, we use classical test theory to investigate the reliability of the assessment, including test and question difficulty, time to completion reliability, the discrimination of the instrument and individual questions, the internal consistency of the instrument, test-retest reliability, and concurrent validity.

A. Data sources

We conducted statistical reliability tests on data collected using the most recent version of the instrument. These data include students from 27 different institutions and 58 distinct courses who have taken the PLIC *at least once* between August 2017 and December 2018. We had 41 courses from Ph.D. granting institutions, 4 from master's granting institutions, and 13 from two- or four-year colleges. The majority of the courses (33) were at the first-year level with 25 at the beyond-first-year level.

Only valid student responses were included in the analysis. To be considered valid, a respondent must

- (1) click submit at the end of the survey,
- (2) consent to participate in the study,
- (3) indicate that they are at least 18 years of age,
- (4) and spend at least 30 sec on at least one of the four pages.

The time cutoff (criteria 4) was imposed because it took a typical reader approximately 20 sec to randomly click through each page, without reading the material. Students who spent too little time could not have made a legitimate effort to answer the questions. Of these valid responses, pre- and postresponses were matched for individual students using the student ID and/or full name provided at the end of the survey. The time cutoff removed 181 (7.6%) students from the matched dataset.

We collected at least one valid survey from 4329 students and matched pre- and postsurveys from 2189 students. The demographic distribution of these data are shown in Table V. We have collapsed all physics, astronomy, and engineering physics majors into the “physics” major category. Instructors reported the estimated number of students enrolled in their classes, which allow us to estimate response rates. The mean response rate to the presurvey was $64\% \pm 4\%$, while the mean response rate to the postsurvey was $49\% \pm 4\%$.

As part of the validation process, 20% of respondents were randomly assigned open-response versions of the PLIC during the 2017–2018 academic year. As such, some students in our matched dataset saw both a closed-response and open-response version of the PLIC. Table VI presents the number of students in the matched dataset that saw each version of the PLIC at pre- and postinstruction. Additional analysis of the open-response data will be included in a future publication, but in the analyses presented here, we focus solely on the closed-response version of the PLIC

TABLE V. Percentage of respondents in valid pre-, post-, and matched datasets broken down by gender, ethnicity, and major. Students had the option to not disclose this information, so percentages may not sum to 100%.

	Pre	Post	Matched
Total	3635	2883	2189
Gender			
Women	39%	39%	40%
Men	59%	60%	59%
Other	0.5%	0.3%	0.4%
Major			
Physics	18%	19%	20%
Engineering	43%	44%	45%
Other science	31%	29%	29%
Other	5.7%	5.8%	5.7%
Race or ethnicity			
American Indian	0.7%	0.6%	0.5%
Asian	25%	26%	28%
African American	2.8%	2.6%	2.4%
Hispanic	4.3%	5.3%	4.8%
Native Hawaiian	0.4%	0.4%	0.4%
White	62%	61%	61%
Other	1.3%	1.6%	1.2%

(1911 students completed both a closed-response presurvey and a closed-response postsurvey).

B. Test and question scores

In Fig. 1 we show the matched pre- and postsurvey distributions ($N = 1911$) of respondents' total scores on the PLIC. The average total score on the presurvey is 5.25 ± 0.05 and the average score on the postsurvey is 5.52 ± 0.05 . The data follow an approximately normal distribution with roughly equal variances in pre- and postscores. For this reason, we use parametric statistical tests to compare paired and unpaired sample means. The pre- and postsurvey means are statistically different (paired t test, $p < 0.001$) with a small effect size (Cohen's $d = 0.23$).

In addition to overall PLIC scores, we also examined scores on individual questions (Fig. 2, i.e., question

TABLE VI. Number of students in the matched dataset who took each version of the PLIC. The statistical analyses focus exclusively on students who completed both a closed-response presurvey and closed-response postsurvey.

Presurvey version	Postsurvey version	N
Closed response	Closed response	1911
	Open response	105
Open response	Closed response	144
	Open response	29

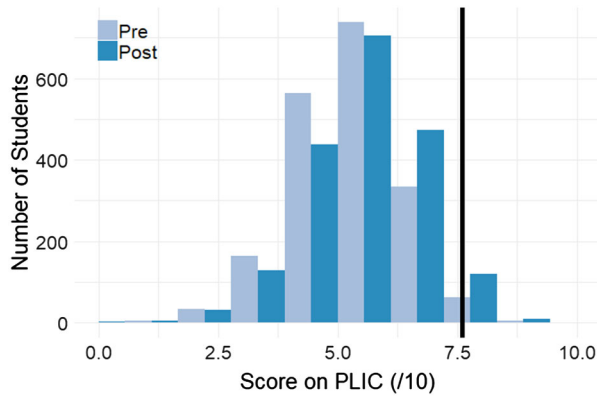


FIG. 1. Distributions of respondents’ matched pre- and post-scores on the PLIC ($N = 1911$). The average score obtained by experts is marked with a black vertical line.

difficulty). All questions differ in the number of closed-response choices available and the score values of each response. We simulated 10 000 students randomly selecting three responses for each question to determine a baseline difficulty for each question. These random guess scores are indicated as black squares in Fig. 2.

The average score per question ranges from 0.30 to 0.80, within the acceptable range for educational assessments [54,55]. Correcting p values for multiple comparisons using the Holm-Bonferroni method, each of the first seven questions showed statistically significant increases between pre- and post-test at the $\alpha = 0.05$ significance level.

The two questions with the lowest average score (Q2E and Q3E) correspond to what to do next questions for group 2. These questions correspond to two of the four lowest scores through random guessing. Furthermore, the highest valued responses on these questions involve changing the fit line, investigating the nonzero intercept, testing

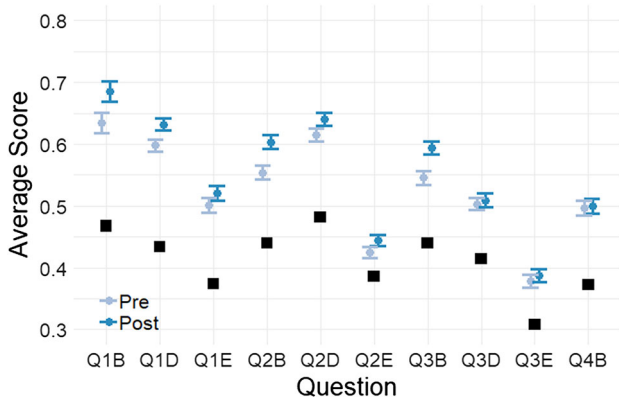


FIG. 2. Average score (representing question difficulty) for each of the 10 PLIC questions. The average expected score that would be obtained through random guessing is marked with a black square. Scores on the first seven questions are statistically different between the pre- and postsurveys (Mann-Whitney U, Holm-Bonferroni corrected $\alpha < 0.05$).

other variables, and checking the assumptions of the model. It is not surprising that students have low scores on these questions, as they are seldom exposed to this kind of investigation, particularly in traditional labs [27]. Although these average scores are low (question difficulty is high), they are still within the acceptable range of [0.3, 0.8] for educational assessments [54,55].

C. Time to completion

We examined the relationship between a respondent’s score and the amount of time they took to complete the PLIC. The total assessment duration was defined as the time elapsed from the moment a respondent opened the survey to the moment they submitted it. This time then encompasses any time that the respondent spent away or disengaged from the survey.

The median time for completing the PLIC is 16.0 ± 0.4 min for the presurvey and 11.5 ± 0.3 min for the postsurvey. The correlation between total PLIC score and time to completion is $r < 0.03$ for both the presurvey and the postsurvey, suggesting that there is no relationship between a student’s score on the PLIC and the time they take to complete the assessment.

D. Discrimination

Ferguson’s δ coefficient gives an indication of how well an instrument discriminates between individual students by examining how many unequal pairs of scores exist. When scores on an assessment are uniformly distributed, this index is exactly 1 [54]. A Ferguson’s δ greater than 0.90 indicates good discrimination among students. Because the PLIC is scored on a nearly continuous scale, all but three students received a unique score on the presurvey, and all but two students received a unique score on the postsurvey. Ferguson’s δ is equal to 1 for both the pre- and postsurveys.

We also examined how well each question discriminates between high- and low-performing students using question-test correlations (that is, correlations between respondents’ scores on individual questions and their score on the full test). Question-test correlations are greater than 0.38 for all questions on the presurvey and greater than 0.42 for all questions on the postsurvey. None of these question-test correlations are statistically different between the pre- and postsurveys (Fisher transformation) after correcting for multiple testing effects (Holm-Bonferroni $\alpha < 0.05$). All of these correlations are well above the generally accepted threshold for question-test correlations, $r \geq 0.20$ [54].

E. Internal consistency

While the PLIC was designed to measure critical thinking, our definition of critical thinking acknowledges that it is not a unidimensional construct. To confirm that the PLIC does indeed exhibit this underlying structure, we performed a principal component analysis (PCA) on

question scores. PCA is a dimensionality reduction technique that finds groups of question scores that vary together. In a unidimensional assessment, all questions should group together, and so there would only be one principal component to explain most of the variance.

We performed a PCA on matched presurveys and postsurveys and found that six principal components were needed to explain at least 70% of the variance in 10 PLIC question scores in both cases. The first three principal components can explain 45% of the variance in respondents' scores. It is clear, therefore, that the PLIC is not a single construct assessment. The first three principal components are essentially identical between the pre- and postsurveys (their inner products between the pre- and postsurvey are 0.98, 0.97, and 0.94, respectively).

Cronbach's α is typically reported for diagnostic assessments as a measure of the internal consistency of the assessment. However, as argued elsewhere [21,56], Cronbach's α is primarily designed for single construct assessments and depends on the number of questions as well as the correlations between individual questions.

We find that $\alpha = 0.54$ for the pre-survey and $\alpha = 0.60$ for the post-survey, well below the suggested minimum value of $\alpha = 0.80$ [57] for a single construct assessment. We conclude, then, in accordance with the PCA results above, that the PLIC is not measuring a single construct and Cronbach's α cannot be interpreted as a measure of internal reliability of the assessment.

F. Test-retest reliability

The repeatability or test-retest reliability of an assessment concerns the agreement of results when the assessment is carried out under the same conditions. The test-retest reliability of an assessment is usually measured by administering the assessment under the same conditions multiple times to the same respondents. Because students taking the PLIC a second time had always received some physics lab instruction, it was not possible to establish test-retest reliability by examining the same students' individual scores. Instead, we estimated the test-retest reliability of the PLIC by comparing the presurvey scores of students in the same course in different semesters. Assuming that students attending a particular course in one semester are from the same population as those who attend that same course in a later semester, we expect preinstruction scores for that course to be the same in both semesters. This serves as a measure of the test-retest reliability of the PLIC at the course level rather than the student level.

In all, there were six courses from three institutions where the PLIC was administered prior to instruction in at least two separate semesters, including two courses where it was administered in three separate semesters. We performed ANOVA evaluating the effect of semester on average prescore and report effect sizes as η^2 for each of the six courses. When the between-groups degrees of freedom

TABLE VII. Summary of test-retest results comparing presurveys from multiple semesters of the same course.

Class	Semester	N	Pre avg.	Comparisons
Class A	Fall 2017	59	5.2 ± 0.3	$F(2, 363) = 0.192$
	Spring 2018	92	5.3 ± 0.2	$p = 0.826$
	Fall 2018	215	5.24 ± 0.14	$\eta^2 = 0.001$
Class B	Fall 2017	79	5.5 ± 0.2	$F(2, 203) = 0.417$
	Spring 2018	36	5.7 ± 0.4	$p = 0.660$
	Fall 2018	91	5.5 ± 0.2	$\eta^2 = 0.004$
Class C	Fall 2017	90	5.8 ± 0.3	$F(1, 207) = 4.54$
	Fall 2018	119	5.50 ± 0.17	$p = 0.034$ $\eta^2 = 0.021$
Class D	Spring 2018	40	6.3 ± 0.3	$F(1, 54) = 0.054$
	Fall 2018	16	6.2 ± 0.7	$p = 0.818$ $\eta^2 = 0.001$
Class E	Fall 2017	142	5.0 ± 0.2	$F(1, 429) = 1.37$
	Fall 2018	289	4.79 ± 0.12	$p = 0.153$ $\eta^2 = 0.005$
Class F	Spring 2018	95	5.0 ± 0.2	$F(1, 182) = 1.89$
	Fall 2018	89	5.1 ± 0.2	$p = 0.171$ $\eta^2 = 0.010$

are equal to 1 (i.e., there are only two semesters being compared), the results of the ANOVA are equal to those obtained from an unpaired t test with $F = t^2$. The results are shown in Table VII.

The p values indicate that presurvey means were not statistically different between semesters for any class other than class C, but these p values have not been corrected for multiple comparisons. After correcting for multiple testing effects using either the Bonferroni or Holm-Bonferroni method at $\alpha < 0.05$ significance, presurvey means were not statistically different for class C.

Given the possibility of small variations in class populations from one semester to the next, it is reasonable to expect small effects to arise on occasion, so the moderate difference in presurvey means for class C is not surprising. As we collect data from more classes, we will continue to check the test-retest reliability of the instrument in this manner.

G. Concurrent validity

We define concurrent validity as a measure of the consistency of performance with expected results. For example, we expect that from either instruction or selection effects, performance on the PLIC should increase with greater physics maturity of the respondent. We define physics maturity by the level of the lab course that respondents were enrolled in when they took the PLIC.

To assess this form of concurrent validity, we split our matched dataset by physics maturity. This split dataset included 1558 respondents from first-year (FY) labs,

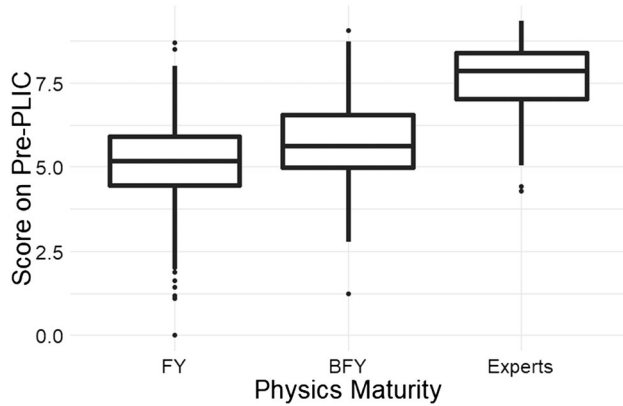


FIG. 3. Boxplots comparing presurvey scores for students in FY and BFY labs, and experts. Horizontal lines indicate the lower and upper quartiles and the median score. Scores lying outside $1.5 \times \text{IQR}$ (inter-quartile range) are indicated as outliers.

353 respondents from beyond-first-year (BFY) labs, and 78 physics experts. Figure 3 compares the performance on the presurvey by physics maturity of the respondent. In Table VIII, we report the average scores for respondents for both the pre- and postsurveys across maturity level. The significance level and effect sizes between pre- and postmean scores within each group are also indicated.

There is a statistically significant difference between pre- and postsurvey means for students trained in FY labs with a small effect size, but not in BFY labs.

An ANOVA comparing pre-survey scores across physics maturity (FY, BFY, expert) indicates that physics maturity is a statistically significant predictor of presurvey scores [$F(2, 1986) = 203, p < 0.001$] with a large effect size ($\eta^2 = 0.170$). The large differences in means between groups of differing physics maturity, coupled with the small increase in mean scores following instruction at both the FY and BFY level, may imply that these differences arise from selection effects rather than cumulative instruction. This selection effect has been seen in other evaluations of students' lab sophistication as well [58].

Another measure of concurrent validity is through the impact of lab courses that aim to teach critical thinking on

TABLE VIII. Average scores across levels of physics maturity, where N is the number of matched responses, except for experts who only filled out the survey once. Significance levels and effect sizes are reported for differences between the pre- and postsurvey means.

	N	Pre avg.	Post avg.	p	d
FY	1558	5.15 ± 0.05	5.45 ± 0.06	<0.001	0.215
BFY	353	5.69 ± 0.12	5.81 ± 0.13	0.095	0.089
Experts	78	7.6 ± 0.2			

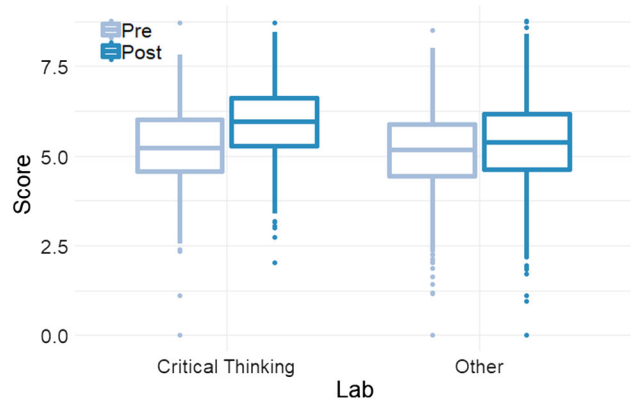


FIG. 4. Boxplots of students' total scores on the PLIC grouped by the type of lab students participated in. Horizontal lines indicate the lower and upper quartiles and the median score. Scores lying outside $1.5 \times \text{IQR}$ are labeled as outliers.

student performance. We grouped FY students according to the type of lab their instructor indicated they were running as part of the course information survey. The data include 273 respondents who participated in FY labs designed to teach critical thinking as defined for the PLIC (CTLabs) and 1285 respondents who participated in other FY physics labs. Boxplots of these scores split by lab type are shown in Fig. 4.

We fit a linear model predicting postscores as a function of lab treatment and prescore:

$$\text{PostScore}_i = \beta_0 + \beta_1 \times \text{CTLabs}_i + \beta_2 \times \text{PreScore}_i. \quad (4)$$

The results are shown in Table IX. We see that, controlling for prescores, lab treatment has a statistically significant impact on postscores; students trained in CTLabs perform 0.52 points higher, on average, at post-test compared to their counterparts trained in other labs with the same prescore.

H. Partial sample reliability

The partial sample reliability of an assessment is a measure of possible systematics associated with selection effects (given that response rates are less than 100% for most courses) [56]. We compared the performance of matched respondents (those who completed a closed-response version of both the pre- and postsurvey) to

TABLE IX. Linear model for PLIC postscore as a function of lab treatment and prescore.

Variable	Coefficient	p
Constant, β_0	4.6 ± 0.3	<0.001
CTLabs, β_1	0.52 ± 0.15	<0.001
Prescore, β_2	0.25 ± 0.05	<0.001

TABLE X. Scores from respondents who took both a closed-response pre- and post-survey (matched dataset) and students who only took one closed-response survey (unmatched dataset). p values and effect sizes d are reported for the differences between the matched and unmatched datasets in each case.

Survey	Dataset	N	Avg.	p	d
Pre	Matched	1911	5.25 ± 0.05	0.011	0.090
	Unmatched	1376	5.15 ± 0.06		
Post	Matched	1911	5.52 ± 0.05	<0.001	0.242
	Unmatched	797	5.22 ± 0.09		

respondents who completed only one closed-response survey. The results are summarized in Table X. Means are statistically different between the matched and unmatched datasets for both the pre- and postsurvey.

These biases, though small, may still have a meaningful impact on our analyses, given that we have neglected certain lower performing students (particularly for the tests of concurrent validity). We address these potential biases by imputing our missing data in Appendix A. This analysis indicates that the missing students did not change the results.

VI. SUMMARY AND CONCLUSIONS

We have presented the development of the PLIC for measuring student critical thinking, including tests of its validity and reliability. We have used qualitative and quantitative tools to demonstrate the PLIC's ability to measure critical thinking (defined as the ways in which one makes decisions about what to trust and what to do) across instruction and general physics training. There are several implications for researchers and instructors that follow from this work.

For researchers, the PLIC provides a valuable new tool for the study of educational innovations. While much research has evaluated student conceptual gains, the PLIC is an efficient and standardized way to measure students' critical thinking in the context of introductory physics lab experiments. As the first such research-based instrument in physics, this facilitates the exploration of a number of research questions. For example, do the gender gaps common on concept inventories [17] exist on an instrument such as the PLIC? How do different forms of instruction impact learning of critical thinking? How do students' critical thinking skills correlate with other measures, such as grades, conceptual understanding, persistence in the field, or attitudes towards science?

For instructors, the PLIC now provides a unique means of assessing student success in introductory and upper-division labs. As the goals of lab instruction

shift from reinforcing concepts to developing experimentation and critical thinking skills, the PLIC can serve as a much needed research-based instrument for instructors to assess the impact of their instruction. Based on the discrimination of the instrument, its use is not limited to introductory courses. The automated delivery system is straightforward for instructors to use. The vast amounts of data already collected also allow instructors to compare their classes to those across the country.

In future studies related to PLIC development and use, we plan to evaluate patterns in students' responses through cluster and network analysis. Given the use of multiple response questions and similar questions for the two hypothetical groups, these techniques will be useful in identifying related response choices within and across questions. We also intend to investigate how students' conceptual understanding of mechanics, as measured by other mechanics diagnostic assessments, relates to performance on the PLIC. Also, the analyses thus far have largely ignored data collected from the open-response version of the PLIC. Unlike the closed-response version, the open-response version has undergone very little revision over the course of the PLIC's development and we have collected over 1000 responses. We plan to compare responses to the open- and closed-response versions to evaluate the different forms of measurement. The closed-response version measures students' abilities to critique and choose among multiple ideas (accessible knowledge), while the open-response version measures students' abilities to generate these ideas for themselves (available knowledge). The relationship between the accessibility and availability of knowledge has been studied in other contexts [59] and it will be interesting to explore this relationship further with the PLIC.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1611482. We are grateful to undergraduate researchers Saaj Chattopadhyay, Tim Rehm, Isabella Rios, Adam Stanford-Moore, and Ruqayya Toorawa who assisted in coding the open-response surveys. Peter Lepage, Michelle Smith, Doug Bonn, Bethany Wilcox, Jayson Nissen, and members of CPERL provided ideas and useful feedback on the manuscript and versions of the instrument. We are very grateful to Heather Lewandowski and Bethany Wilcox for their support developing the administration system. We would also like to thank all of the instructors who have used the PLIC with their courses, especially Mac Stetzer, David Marasco, and Mark Kasevich for running the earliest drafts of the PLIC.

TABLE XI. Number of PLIC closed-response pre- and post-surveys collected.

	Postsurvey missing	Postsurvey completed
Presurvey missing	245	797
Presurvey completed	1376	1911

APPENDIX A: MULTIPLE IMPUTATION ANALYSIS

In the analyses in the main text we used only PLIC data from respondents who had taken the closed-response PLIC at both pre- and postsurvey. As discussed in Sec. V H, the missing data biases the average scores toward higher scores. It is likely that this bias is due to a skew in representation toward students who receive higher grades in the matched dataset compared to the complete dataset [60]. This bias is problematic for two reasons, one of interest to researchers and the other of interest to instructors. For researchers, the analyses above may not be accurate when taking into account a larger population of students (including lower-performing students). Instructors using the PLIC in their classes who achieve high participation rates may be led to incorrectly conclude that their students performed below average because their class included a larger proportion of low-performing students. Using imputation, we quantified the bias in the matched dataset to more accurately represent PLIC scores across a wider group of students and to be transparent to instructors wishing to use the PLIC in their classes.

Imputation is a technique for filling in missing values in a dataset with plausible values for more complete datasets. In the case of the PLIC, we imputed data for respondents who completed the closed-response version of either the pre- or postsurvey, but not both (2173 respondents). The number of PLIC closed-response surveys collected is summarized in Table XI for both pre- and postsurveys. The 245 students in Table XI who are missing both pre- and postsurvey data represent the students who completed at least one open-response survey, but no closed-response surveys. Without any information about how these students perform on a closed-response PLIC survey or other useful information such as their grades or scores on standardized assessments, these data cannot be reliably imputed.

We used MICE [61] with predictive mean matching (pmm) to impute the missing closed-response data. These methods have been discussed previously in Refs. [62,63] and so we do not elaborate on them here. For each respondent, we used the levels of the lab they were enrolled

TABLE XII. Average scores for datasets containing only matched students, all valid responses, and the complete set of students including imputed scores.

Dataset	N	Pre avg.	Post avg.
Matched	1911	5.25 ± 0.05	5.52 ± 0.05
Valid presurveys	3287	5.21 ± 0.04	
Valid postsurveys	2708		5.43 ± 0.05
Imputed	4046	5.20 ± 0.04	5.42 ± 0.04

in (FY or BFY), the type of lab they were enrolled in (CTLabs or other), and the score on the closed-response survey they completed to estimate their missing score. MICE operates by imputing our dataset M times, creating M complete datasets, each containing data from 4084 students. Each of these M datasets will have somewhat different values for the imputed data [62,63]. If the calculation is not prohibitive, it has been recommended that M be set to the average percentage of missing data [63], which in our case is 27. After constructing our M imputed datasets, we conducted analyses (means, t tests, effect sizes, regressions) on these datasets separately, then combined our M results using Rubin's rules to calculate standard errors [64].

Using our imputed datasets, we now demonstrate that the results shown above concerning the overall scores and measures of concurrent validity are largely the same as with the imputed data set. We did not examine measures involving individual questions or their correlations (question scores, discrimination, internal consistency) as the variability between questions makes the predictions through imputation unreliable. The test-retest reliability measure included all valid closed-response pre-survey responses, and so there is much less missing data, making the imputation unnecessary.

1. Test scores

Mean scores for the matched, total, and imputed datasets are shown in Table XII. The presurvey mean of the imputed dataset is not statistically different from the presurvey mean of the matched dataset or the valid presurveys dataset. Similarly, the postsurvey mean of the imputed dataset is not statistically different from the postsurvey mean of the valid-post surveys dataset. There is, however, a statistically significant difference in post-survey means between the matched and imputed datasets (t test, $p < 0.05$) with a small effect size (Cohen's $d = 0.08$). Additionally, as with the matched dataset, there is a statistically significant difference between pre- and postmeans in the imputed dataset (t test, $p < 0.001$). However, the effect size is smaller (Cohen's $d = 0.16$) than in the matched dataset.

TABLE XIII. Scores across levels of physics maturity, where N is the number of responses in the imputed dataset (except for experts). Significance levels and effect sizes are reported for differences in pre- and post-test means.

	N	Pre avg.	Post avg.	p	d
FY	3428	5.12 ± 0.04	5.36 ± 0.05	<0.001	0.173
BFY	656	5.62 ± 0.10	5.74 ± 0.10	0.052	0.088
Experts (nonimputed)	78	7.6 ± 0.2			

TABLE XIV. Linear model for PLIC postscore as a function of lab treatment and prescore using the imputed dataset.

Variable	Coefficient	p
Constant, β_0	4.3 ± 0.3	<0.001
CTLabs, β_1	0.40 ± 0.13	<0.001
Pre-score, β_2	0.27 ± 0.06	<0.001

2. Concurrent validity

Our imputed dataset contains 3428 respondents from FY labs and 656 respondents from BFY labs. Because the experts only filled out the survey once, there is no missing data and imputation is unnecessary. In Table XIII, we report again the average scores for students enrolled in FY and BFY level physics lab courses and experts.

As in the matched dataset, there is a statistically significant difference in pre- and postmeans for students trained in FY labs, but the effect size is smaller. Again, there is no statistically significant difference in pre- and postmeans for students trained in BFY labs. Again, an ANOVA comparing presurvey scores across physics maturity indicates that physics maturity is a statistically significant predictor of presurvey scores [$F(2, 6210.5) = 218.3$, $p < 0.001$] with a large effect size ($\eta^2 = 0.101$).

Our imputed dataset contained a total of 505 students who participated in FY CTLabs and 2923 students who participated in other FY physics labs. The linear fit for postscores using prescores and lab type as predictors (see Eq. (4) again found that students trained in CTLabs

outperform their counterparts in other labs. Coefficients and significance levels are reported in Table XIV.

3. Summary and limitations

In this Appendix, we have demonstrated via multiple imputation that PLIC scores may be, on average, slightly lower than those reported in the main article. This is likely due to a skew in representation toward higher-performing students in the matched dataset and is more prevalent on the postsurvey than the presurvey. This bias does not, however, affect the conclusions of the concurrent validity section of the main article. There is a statistically significant difference in prescores between students enrolled in FY and BFY labs, as well as experts. Students trained in CTLabs score higher on the post-survey, on average, than students trained in other physics labs after taking pre-survey scores into account.

The main limitation to imputing our data in this way stems from the reliability of the imputed values. As briefly mentioned above, we lack information about students' grades or their scores on other standardized assessments, which have been shown to be useful predictors of student scores on diagnostic assessments like the PLIC [62]. Without this information, the reliability of the imputed dataset is limited. Estimating missing PLIC scores using the predictor variables above (level and type of lab a student was enrolled in and their score on one closed-response survey) likely provides a better estimate of population distributions than simply ignoring the missing data, but much of the variance in scores is not explained by these variables alone.

APPENDIX B: SAMPLE PLIC QUESTIONS

A full version of the PLIC can be found [20]. Here, we show an excerpt from the PLIC, providing examples of the multiple choice and multiple response questions that can be found on the assessment.

Summary

Group 1:

Method: 10 trials of 5 oscillations for 2 different masses and compared k values:

With 30 g mass: $k = 4.0857 \pm 0.08067$ N/m

With 50 g mass: $k = 4.1872 \pm 0.06745$ N/m

Group 2:

Method: 2 trials of 5 oscillations for 10 different masses and compared the data to a model equation:

With intercept set to 0: $k = 4.0535 \pm 0.07005$ N/m

With intercept not fixed at 0: $k = 4.6274 \pm 0.1587$ N/m

Click [here](#) to see the intro again, [here](#) to see Group 1's lab book, [here](#) to see Group 2's initial lab book notes, and [here](#) to see Group 2's second set of lab book notes.

Which group do you think did a better job of testing the model?

- Group 1
- Group 2
- Both the same

What features were most important for comparing the two groups?

Please select no more than 3 options.

- | | |
|--|--|
| <input type="checkbox"/> The number of bounces of the spring per trial | <input type="checkbox"/> How they accounted for human error |
| <input type="checkbox"/> The number of repeated trials for each mass | <input type="checkbox"/> How close the results are to what the model predicts |
| <input type="checkbox"/> The number of masses tested | <input type="checkbox"/> The size of the uncertainty (or variability between data) |
| <input type="checkbox"/> How they tested other possible variables | <input type="checkbox"/> Their analysis and calculations |
| <input type="checkbox"/> How they controlled variables | <input type="checkbox"/> How clear, organized, or detailed their lab notes are |
| <input type="checkbox"/> The equipment they used (i.e. stopwatches) | <input type="checkbox"/> Other (Please describe) |

FIG. 5. Excerpt from page four of the PLIC describing the methods of the two hypothetical groups of students. The final multiple choice questions ask respondents which group they think did a better job of testing the model, while the multiple response question asks students to elaborate on their reasoning.

[1] P.J. Mulvey and S. Nicholson, Physics enrollments: Results from the 2008 survey of enrollments and degrees, (2011), statistical Research Center of the American Institute of Physics, www.aip.org/statistics.

[2] M.-G. Séré, Towards renewed research questions from the outcomes of the european project labwork in science education, *Sci. Educ.* **86**, 624 (2002).

[3] R. T. White, The link between the laboratory and learning, *Int. J. Sci. Educ.* **18**, 761 (1996).

- [4] A. Hofstein and V. N. Lunetta, The role of the laboratory in science teaching: Neglected aspects of research, *Rev. Educ. Res.* **52**, 201 (1982).
- [5] A. Hofstein and V. N. Lunetta, The laboratory in science education: Foundations for the twenty-first century, *Sci. Educ.* **88**, 28 (2004).
- [6] National Research Council, *Committee on High School Science Laboratories: Role and Vision, America's Lab Report: Investigations in High School Science*, edited by S. R. Singer, M. L. Hilton, and H. A. Schweingruber (National Academies Press, Washington, DC, 2006).
- [7] R. Millar, *The Role of Practical Work in the Teaching and Learning of Science, High School Science Laboratories: Role and Vision* (National Academy of Sciences, Washington, DC, 2004).
- [8] C. Wieman and N. G. Holmes, Measuring the impact of an instructional laboratory on the learning of introductory physics, *Am. J. Phys.* **83**, 972 (2015).
- [9] N. G. Holmes, J. Olsen, J. L. Thomas, and C. E. Wieman, Value added or misattributed? A multi-institution study on the educational benefit of labs for reinforcing physics content, *Phys. Rev. Phys. Educ. Res.* **13**, 010129 (2017).
- [10] B. R. Wilcox and H. J. Lewandowski, Developing skills versus reinforcing concepts in physics labs: Insight from a survey of students beliefs about experimental physics, *Phys. Rev. Phys. Educ. Res.* **13**, 010108 (2017).
- [11] S. Olson and D. G. Riordan, Engage to Excel: Producing one million additional college graduates with degrees in science, technology, engineering, and mathematics, Technical Report (President's Council of Advisors on Science and Technology, Washington, DC, 2012).
- [12] L. McNeil and P. Heron, Preparing physics students for 21st-century careers, *Phys. Today* **70**, No. 11, 38 (2017).
- [13] J. Kozminski, H. J. Lewandowski, N. Beverly, S. Lindaas, D. Deardorff, A. Reagan, R. Dietz, R. Tagg, M. Eblen-Zayas, J. Williams, R. Hobbs, and B. M. Zwickl, AAPT committee on laboratories, *AAPT Recommendations for the Undergraduate Physics Laboratory Curriculum* (AAPT, College Park, MD, 2014).
- [14] J. L. Docktor and J. P. Mestre, Synthesis of discipline-based education research in physics, *Phys. Rev. ST Phys. Educ. Res.* **10**, 020119 (2014).
- [15] S. R. Singer, N. R. Nielsen, and H. A. Schweingruber, *Discipline-based Education Research: Understanding and Improving Learning in Undergraduate Science and Engineering* (National Academy Press, Washington DC, 2012).
- [16] R. R. Hake, Interactive-engagement versus traditional methods: A six-thousand-student survey of mechanics test data for introductory physics courses, *Am. J. Phys.* **66**, 64 (1998).
- [17] A. Madsen, S. B. McKagan, and E. C. Sayre, Gender gap on concept inventories in physics: What is consistent, what is inconsistent, and what factors influence the gap?, *Phys. Rev. ST Phys. Educ. Res.* **9**, 020121 (2013).
- [18] A. Madsen, S. B. McKagan, and E. C. Sayre, How physics instruction impacts students beliefs about learning physics: A meta-analysis of 24 studies, *Phys. Rev. ST Phys. Educ. Res.* **11**, 010115 (2015).
- [19] S. Freeman, S. L. Eddy, M. McDonough, M. K. Smith, N. Okoroafor, H. Jordt, and M. P. Wenderoth, Active learning increases student performance in science, engineering, and mathematics, *Proc. Natl. Acad. Sci. USA* **111**, 8410 (2014).
- [20] <https://www.physport.org/>.
- [21] J. Day and D. Bonn, Development of the concise data processing assessment, *Phys. Rev. ST Phys. Educ. Res.* **7**, 010114 (2011).
- [22] D. L. Deardorff, Introductory Physics Students' Treatment of Measurement Uncertainty, Ph.D. thesis, North Carolina State University (2001).
- [23] S. Allie, A. Buffler, B. Campbell, and F. Lubben, First-year physics students perceptions of the quality of experimental measurements, *Int. J. Sci. Educ.* **20**, 447 (1998).
- [24] K. A. Slaughter, Mapping the transition—content and pedagogy from school through university, Ph.D. thesis, University of Edinburgh (2012).
- [25] C. Wieman, Comparative cognitive task analyses of experimental science and instructional laboratory courses, *Phys. Teach.* **53**, 349 (2015).
- [26] S. Nicholson and P. J. Mulvey, Roster of physics departments with enrollment and degree data, 2013: Results from the 2013 survey of enrollments and degrees, (2014), statistical Research Center of the American Institute of Physics, www.aip.org/statistics.
- [27] N. G. Holmes, C. E. Wieman, and D. A. Bonn, Teaching critical thinking, *Proc. Natl. Acad. Sci. U.S.A.* **112**, 11199 (2015).
- [28] B. M. Zwickl, N. Finkelstein, and H. J. Lewandowski, Incorporating learning goals about modeling into an upper-division physics laboratory experiment, *Am. J. Phys.* **82**, 876 (2014).
- [29] B. M. Zwickl, D. Hu, N. Finkelstein, and H. J. Lewandowski, Model-based reasoning in the physics laboratory: Framework and initial results, *Phys. Rev. ST Phys. Educ. Res.* **11**, 020113 (2015).
- [30] D. T. Willingham, Critical thinking: Why is it so hard to teach?, *Arts Educ. Pol. Rev.* **109**, 21 (2008).
- [31] W. Adams, Development of a problem solving evaluation instrument; Untangling of specific problem solving skills, Ph.D. thesis, University of Colorado—Boulder (2007).
- [32] C. S. Dweck, Motivational processes affecting learning, *Am. Psychol.* **41**, 1040 (1986).
- [33] C. S. Dweck, *Self-theories: Their role in motivation, personality, and achievement* (Psychology Press, Philadelphia, PA, 1999).
- [34] A. E. Lawson, The development and validation of a classroom test of formal reasoning, *J. Res. Sci. Teach.* **15**, 11 (1978).
- [35] B. Stein, A. Haynes, and M. Redding, Project cat: Assessing critical thinking skills, in *National STEM Assessment Conference*, edited by D. Deeds and B. Callen (NSF, Springfield, MO, 2006), pp. 290–299.
- [36] B. Stein, A. Haynes, M. Redding, T. Ennis, and M. Cecil, Assessing critical thinking in stem and beyond, in *Innovations in e-learning, instruction technology, assessment, and engineering education* (Springer, Dordrecht, Netherlands, 2007), pp. 79–82.

- [37] P. A. Facione, Using the California Critical Thinking Skills Test in Research, Evaluation, and Assessment, Technical Report (Millbrae, CA, 1990).
- [38] C. Gormally, P. Brickman, and M. Lutz, Developing a test of scientific literacy skills (TOSLS): Measuring undergraduates evaluation of scientific information and arguments, *CBE Life Sci. Educ.* **11**, 364 (2012).
- [39] K. Sirum and J. Humburg, The experimental design ability test (EDAT), *Bio. J. Col. Biol. Teach.* **37**, 8 (2011).
- [40] T. Deane, K. Nomme, E. Jeffery, C. Pollock, and G. Birol, Development of the biological experimental design concept inventory (BEDCI), *CBE Life Sci. Educ.* **13**, 361 (2014).
- [41] K. A. Slaughter, Mapping the transition-content and pedagogy from school through university, Ph.D. thesis, University of Edinburgh, 2012.
- [42] T. S. Volkwyn, S. Allie, A. Buffler, and F. Lubben, Impact of a conventional introductory laboratory course on the understanding of measurement, *Phys. Rev. ST Phys. Educ. Res.* **4**, 010108 (2008).
- [43] D. L. Deardorff, Introductory physics students' treatment of measurement uncertainty, Ph.D. thesis, North Carolina State University (2001).
- [44] R. F. Lippmann, Students' understanding of measurement and uncertainty in the physics laboratory: Social construction, underlying concepts, and quantitative analysis, Ph.D. thesis, University of Maryland, United States—Maryland (2003).
- [45] N. G. Holmes, Structured quantitative inquiry labs: Developing critical thinking in the introductory physics laboratory, Ph.D. thesis, University of British Columbia (2015).
- [46] N. G. Holmes and D. A. Bonn, Quantitative comparisons to promote inquiry in the introductory physics lab, *Phys. Teach.* **53**, 352 (2015).
- [47] A. Buffler, S. Allie, and F. Lubben, The development of first year physics students' ideas about measurement in terms of point and set paradigms, *Int. J. Sci. Educ.* **23**, 1137 (2001).
- [48] A. Buffler, F. Lubben, and B. Ibrahim, The relationship between students views of the nature of science and their views of the nature of scientific measurement, *Int. J. Sci. Educ.* **31**, 1137 (2009).
- [49] N. G. Holmes and C. E. Wieman, Assessing modeling in the lab: Uncertainty and measurement, in *2015 Conference on Laboratory Instruction Beyond the First Year of College* (American Association of Physics Teachers, College Park, MD, 2015), pp. 44–47.
- [50] K. N. Quinn, C. E. Wieman, and N. G. Holmes, Interview validation of the physics lab inventory of critical thinking (PLIC), in *2017 Physics Education Research Conference Proceedings* (American Association of Physics Teachers, Cincinnati, OH, 2018), pp. 324–327.
- [51] A. C. Butler, Multiple-choice testing in education: Are the best practices for assessment also good for learning?, *J. Appl. Res. Mem. Cog.* **7**, 323 (2018).
- [52] B. R. Wilcox, B. M. Zwickl, R. D. Hobbs, J. M. Aiken, N. M. Welch, and H. J. Lewandowski, Alternative model for assessment administration and analysis: An example from the E-CLASS, *Phys. Rev. Phys. Educ. Res.* **12**, 010139 (2016).
- [53] <http://cperl.lassp.cornell.edu/plic>.
- [54] L. Ding and R. Beichner, Approaches to data analysis of multiple-choice questions, *Phys. Rev. ST Phys. Educ. Res.* **5**, 020103 (2009).
- [55] R. L. Doran, *Basic Measurement and Evaluation of Science Instruction* (ERIC, Washington, DC, 1980).
- [56] B. R. Wilcox and H. J. Lewandowski, Students epistemologies about experimental physics: Validating the colorado learning attitudes about science survey for experimental physics, *Phys. Rev. Phys. Educ. Res.* **12**, 010123 (2016).
- [57] P. Engelhardt, An introduction to classical test theory as applied to conceptual multiple-choice tests, in *Getting Started in PER*, edited by C. Henderson and K. A. Harper (American Association of Physics Teachers, College Park, MD, 2009).
- [58] B. R. Wilcox and H. J. Lewandowski, Improvement or selection? A longitudinal analysis of students views about experimental physics in their lab courses, *Phys. Rev. Phys. Educ. Res.* **13**, 023101 (2017).
- [59] A. F. Heckler and A. M. Bogdan, Reasoning with alternative explanations in physics: The cognitive accessibility rule, *Phys. Rev. Phys. Educ. Res.* **14**, 010120 (2018).
- [60] J. M. Nissen, M. Jariwala, E. W. Close, and B. Van Dusen, Participation and performance on paper-and computer-based low-stakes assessments, *Int. J. STEM Educ.* **5**, 21 (2018).
- [61] S. van Buuren and K. Groothuis-Oudshoorn, mice: Multivariate imputation by chained equations in R, *J. Stat. Softw.* **45**, 1 (2011).
- [62] J. Nissen, R. Donatello, and B. Van Dusen, Missing data and bias in physics education research: A case for using multiple imputation, [arXiv:1809.00035](https://arxiv.org/abs/1809.00035).
- [63] S. Van Buuren, *Flexible imputation of missing data* (Chapman and Hall/CRC, Boca Raton, FL, 2018).
- [64] D. B. Rubin, Inference and missing data, *Biometrika* **63**, 581 (1976).