Sci & Educ
https://doi.org/10.1007/s11191-018-9970-2

ARTICLE

The Missing Curriculum in Physics Problem-Solving Education

Mobolaji Williams[1]

Abstract Physics is often seen as an excellent introduction to science because it allows students to learn not only the laws governing the world around them, but also, through the problems students solve, a way of thinking which is conducive to solving problems outside of physics and even outside of science. In this article, we contest this latter idea and argue that in physics classes, students do not learn widely applicable problem-solving skills because physics education almost exclusively requires students to solve well-defined problems rather than the less-defined problems which better model problem solving outside of a formal class. Using personal, constructed, and the historical accounts of Schrödinger's development of the wave equation and Feynman's development of path integrals, we argue that what is missing in problem-solving education is practice in identifying gaps in knowledge and in framing these knowledge gaps as questions of the kind answerable using techniques students have learned. We discuss why these elements are typically not taught as part of the problem-solving curriculum and end with suggestions on how to incorporate these missing elements into physics classes.

Keywords Problem-solving · Education · Teaching · Physics · History of physics

1 Introduction: Liberal Arts of the Sciences

Individual practice, it is said, exists at the heart of all physics education. In any physics class, you'll invariably hear the professor introduce her lectures with a caveat: The lectures can give the student some background and context for what he's learning, but it is only by carefully, and often independently, working through as many problems as possible that he will become proficient in the subject.

✉ Mobolaji Williams
mwilliams@physics.harvard.edu

[1] Department of Physics, Harvard University, Cambridge, MA 02138, USA

Published online: 12 May 2018

This point forms the basis of the argument, if made only implicitly, for why an education in physics is valuable outside any career related to physics or even science. Having majored in a subject sometimes touted as "the liberal arts of the sciences" (Notre Dame Department of Physics 2009)—meaning, a science subject with as wide applicability as an education focused on history, literature, and philosophy—graduating physics students often find themselves in the awkward position of ending an undergraduate career where their peers marveled at the incipient intelligence apparently necessary to work in their chosen subject and entering a job market where few employers seem terribly interested in the specifics of that intelligence. The suggested solution is a throwback to the social sphere many physics students were trying to escape through their studies: marketing. You package yourself not only as someone who understands the physical laws of the universe, but also as someone who has spent the past 4 years solving, and thus learning how to solve, problems. A report commissioned by the NSF regarding how physics curricula can be adapted for and justified in an increasing technological world advocates a similar sense of physics' relevance:

> Undergraduate physics education provides students with unique skills and ways of thinking that are of profound value to themselves and to society…Physics students learn to develop conceptual and mathematical approaches to models to help them understand complicated systems and solve complex problems. As a result of learning the inquiry process and ways of thinking used in physics, students with a physics education are prepared for success in complex analytical professional programs such as medicine, business, finance, and law. (Committee on Undergraduate Physics Education Research and Implementation 2013)

Regarding the claim that physics students "learn to develop conceptual and mathematical approaches...to help them…solve complex problems," we can ask whether this sense of physics as conduit to problem-solving mastery is true: that is, do physics students graduate having learned how to solve problems in a way that is generally useful outside of physics?

Given studies on students' experiences solving problems in non-course related contexts (Fortus 2009; Ogilvie 2009; Milbourne 2016), I think it is easier to argue something else, something more in line with what physics students actually practice. Namely, from their dozens of courses with many dozens of problems each, physics students graduate having a solid knowledge of how to solve *homework* and *textbook* problems. However, it is not clear that such a specific type of problem-solving knowledge is at all transferable to less artificial contexts outside or even inside of academia.

In the following sections, we discuss the inadequacies of physics problem-solving education by first using personal examples to argue that the typical qualities of such an education do not prepare students for research or for work outside of typically academic contexts. We then discuss a refined and more complete depiction of the problem-solving process and highlight the steps that that most physics courses miss. We justify these steps by analyzing the historical accounts of how Schrödinger developed the wave equation and of how Feynman developed the path integral formalism of quantum mechanics. We conclude by outlining (and refuting) the possible reasons for not teaching the full problem-solving process, and then we discuss ways to incorporate a full representation of that process into the curriculum.

2 Author's Experience: Problem-Solving in Academia

What often surprises a beginning student about a first research project is the sheer amount of fumbling required before what one may have thought of as the "real work" can begin. The student spends much of his or her time meandering and confused, simultaneously trying to move out of this confusion and to suppress the debilitating intellectual insecurity it engenders. Worse still they might have done quite well in their course work, and thus have entered research with a false sense of confidence. This confidence likely grew from an education that trained them for the clean parts of research, for the solutions to problems which were well-articulated, manifestly present, and known to be important. The fumbling (which is arguably the true "real work" of research) such a student goes through is concerned with finding better ways to represent the governing question, and navigating this process is something that, in my experience, an undergraduate education in physics does not teach.

I remember feeling somewhat betrayed when I realized this. I was a competent student in all the directions that supposedly mattered (i.e., mainly in my course work), and yet I seemed to be woefully unprepared for research. My first research project concerned R-axions in theories of multiple supersymmetry breaking, a somewhat niche area of high energy theoretical physics (The details of the work are not crucial, but see the footnotes[1] for a short description of a supersymmetry, supersymmetry breaking, and R-axions). The task was to investigate the physical properties of such theories and to determine the masses and ultimately the experimental signatures of the associated particles. However, at this time multiple supersymmetry breaking was (and still is) a fairly unexplored topic, and I did not have the requisite background in supersymmetry to even begin the project. Therefore, as a first step, my advisor suggested I work on a toy model of the problem which only required quantum field theory, a pre-requisite of supersymmetry.

Unfortunately, even this toy model proved intractable. During the first months I worked on the problem, I would begin and end my work sessions in confusion. I knew quantum field theory and had previously spent many hours computing decay rates, cross sections, and loop diagrams using the standard methods of quantum field theory, but, for the toy model, it was not clear how or where such techniques could be useful. It was not even clear to me what questions I could generate in model which could serve as reasonable analogs to the questions I solved in the standard quantum field theory texts, or even whether these questions were the ones I needed to answer in order to develop a framework which could be extended to the supersymmetric case.

Those who have successfully made the transition from student to researcher, might recognize something obvious in what it took me many frustrating years to learn: That simply because you solved all the problems in a course, it does not directly follow that you can creatively or flexibly use what you have supposedly learned in a way that informs the way you see the world. The issue is that through your education you learned how to solve problems meant to inculcate you in the techniques of your discipline, but the world's problems bear no fidelity to that representation, that is, the world does not immediately present itself in a way amenable to the methods you spent so long learning.

[1] Supersymmetry is a postulated symmetry of nature which states that for every boson (fermion) in the universe there is a fermion (boson) with the same mass and interaction properties (Wess 2000). Physicists say that the symmetry of a theory is "broken" when the dynamical equations of the theory satisfy the symmetry, but the solutions to those equations (i.e., the physical observations) do not (Weinberg 2005, pp. 163–76). Therefore, if our world is indeed supersymmetric—meaning that it is invariant under appropriate exchanges of bosons and fermions—supersymmetry must be broken at every-day energy scales because it is clear that the world around us is not manifestly supersymmetric. Theories of broken supersymmetry can include a particle called the "R-axion" which is itself associated with another broken symmetry called R-symmetry (Nelson and Seiberg 1994). It should be noted that supersymmetry is a conjectured symmetry and has not been physically observed.

In the end, I was able to determine the masses of the particles in the toy model—a success which was just small part of a much larger investigation which was never completed—but what allowed me to make what little progress I did was an abandonment of the habits which were so instilled from traditional education. I had to learn to engage in a problem-solving process much deeper than the one to which I was accustomed by learning how to ask myself questions when confused and explore a subject area beyond the confines of my central question.

That course work does not necessarily prepare students for open-ended challenges as they occur in research was shown by Fortus (2009). When analyzing how professors of physics and post-docs and first-year graduate students in science education (all of whom had at least a BA in physics) solved well-defined and "real world" Newtonian mechanics problems—the latter of which were defined as problems requiring the solver to make extensive assumptions and choose for themselves the relevant concepts and algorithms necessary for solution—Fortus found that participants were not successful in transforming the unfamiliar context of the real-world problem into a recognizable problem archetype, and ultimately solving it, unless they had extensive prior experience in dealing with such open-ended problems. For the participants, this experience was not garnered through their prior physics course work—all of which was extensive—but through their prior work in scientific research.

Given the results of this study, my argument may strike some readers as a weak one to make. Of course physics classes do not prepare students for open-ended problems as they exist in research. Only research can do that. However, here, research is just an example of a context where the general problem-solving techniques a physics student apparently learns throughout his or her education are shown to not at all be applicable to more realistic problem contexts. And the fact that this problem-solving education is shown to be inapplicable even in a less artificial space still within the domain in which it is obtained (namely academia) is disconcerting. For if indirect problem-solving education proves to be inadequate for problems in research, a prototypical academic activity, how might it fare in preparing students for work outside of academic contexts?

3 Author's Experience: Problem-Solving in Industry

In my first year in graduate school, I completed freelance work for Reasoning Mind, an educational software nonprofit which seeks to bring the relatively higher math education standards of Russian elementary schools to the US school system. It was a data analysis project, and at the start of the work I was given a list of questions to answer and a few gigabytes of excel spreadsheets which apparently held the answers to those questions. It was a strange experience to be so completely lost. I had spent the prior semester in a data analysis class and the year before learning Python in ways very similar to how I learned during my previous physics education. In homework assignments or in the course textbook, there would be an assigned problem, as explicit and as clear as possible to prevent the student from being confused, and we students would solve the problem using the techniques we were learning. But for my work in Reasoning Mind, the "problems" I was tasked with solving were not at all as clearly framed as the problems in my homework assignments. In homework assignments I was asked to "Write a for loop which prints the location of each vowel in each word of a sentence," or "Write a function which implements Euler's algorithm for pendulum motion," whereas for the data analysis project I was asked "What is the average performance trajectory for a student throughout the curriculum?" or "Do students who run out of time on questions early in the curriculum catch up later on?"

Therefore, before I could begin any explicit work on the project, I had to teach myself how to answer such questions, going through the slow and uncertain work of attempting to represent these more general queries in ever more concrete forms so that I could interpret and answer them through the lens of computational techniques which seemed to make no reference to them in the first place. Moreover, my physics education, an education supposedly doubling as training in problem solving, had mostly supplied me with heuristics which seemed concerned with solving an altogether different species of problem. For how can you "work backwards" when you don't even have a notion of what the answer should be? How could you "guess and check" when the question isn't even framed quantitatively?

In short, I was running into the same wall I had encountered when I had first started research, namely that the problems I needed to solve were distinguished in scope and in the precision of their definitions from the problems I had learned how to solve. The fact that I encountered this discrepancy in a context wholly removed from my foundational physics education suggests that not only did that education not entirely prepare me to solve problems outside of physics, but also, given my previous data science and programming courses, this pedagogical lack was perhaps a symptom of technical education in general.

I should mention that my experience does not appear to be representative of how all graduates of physics degree programs view their education. In interviews with six physics graduates, who had moved on to careers in medical physics, computer programming, and data analysis, the researchers found that a majority saw their education in problem-solving and analytical thinking as the most important of the skills they developed (Sharma et al. 2008). In a later more comprehensive survey of 108 physics graduates (about 70% of whom entered the non-academic workforce after graduation), the researchers again found that graduates on average saw problem-solving as their most developed skill (O'Byrne et al. 2008).

However, the employers of these physics graduates saw their skills with somewhat more nuance. In the earlier study, two employers stated that "whilst physics graduates are generally competent at carrying out tasks or experimental work when the procedure is given, they found that they do less well when forced to start from scratch," although these employers still stated that physics graduates were better equipped than other graduates in tackling problems outside their field of expertise.

4 Problem-Solving: Inside and Outside the Classroom

What was the source of my dissonant problem-solving experiences both in research and in industry? Phrased differently, why was the real-world application of the problem-solving process not accurately reflected in my tacit education in that that process?

The thesis of this article is that the principal missing element concerns questions, specifically a failure to account for the role of questions in the so called "problem-solving process." The very phrase "problem-solving process" is predicated on a false idea of why problems are solved or even how they are created. It seems to suggest that the goal of implementing the process is to solve the problem and that one begins with the problem explicitly stated and then needs to implement heuristics and techniques in order to find the solution.

Instead, it is more accurate to think of the problem as a deliberately chosen framework imposed on some sense of incompleteness. Something confusing or unclear in a knowledge network or some need materially unrealized is brought closer to clarity or realization, respectively, by first expressing this incompleteness in a way which better allows for the possibility of a solution. This framework is what we most often think of as the problem itself. This conception of a problem as centered around incompleteness is similar to Andreas Faludi's conception of a problem as a "subjective state of tension" (Faludi 1973, pp. 82–4), except here we consider this state of tension to exist apart from and prior to any definition of a problem. That is, we consider confusion and tension to not, in of themselves, represent problems unless they are articulated as such.

What then is the true problem-solving process? The status of *How to Solve it* (Polya 2014) as a useful and comprehensive reference for problem-solving heuristics suggests we could consider its commentary on the problem-solving process as archetypal of general views on that process. In the text, Polya describes the process as consisting of four phases: First, one has to understand the problem. Second, one has to plan a solution. Third, one has to carry out the plan. And, finally, one must look back and review the completed solution.

Polya's text aims to provide general problem-solving heuristics in all contexts, but since we are primarily discussing problem-solving in physics we can also turn to canonical introductions in physics to get a sense of physics problem-solving in particular. A 2014 AIP report (Tesfaye et al. 2014), cited *Fundamentals of Physics* (Halliday et al. 2010) as comprising a plurality of textbooks used by American high school teachers for Advanced Placement Physics C, the equivalent of a calculus-based 1st year college course. And yet, while *Fundamentals of Physics* includes on the order of a thousand physics problems, there is never a discussion of the general process by which students should solve problems. More advanced introductions are better on this account. *Introduction to classical mechanics* (Morin 2008) is often used in introductory physics courses where the incoming students have had more in-depth physics and mathematics exposure. For example, Morin's is the main text in the first semester advanced introductory physics course for physics majors at Harvard University, and its older inspiration, *Introduction to mechanics,* (Kleppner and Kolenkow 2013)) is used as the main texts in similar courses at MIT (Burgasser 2008), Rice University (PHYS 111 001 – Rice's Course Schedule 2015), and UC Santa Barbara (Fratus 2016).

In *Introduction to classical mechanics*, there is again no specific description of the problem-solving process, but we can infer the text's sense of that process from the suggested problem-solving strategies in the first chapter. Among these strategies are to "Draw a diagram, if appropriate," to "Write what you known, and what you're trying to find," and to "Solve things symbolically." All of these are heuristics found in, or similar to those found in, Polya's *How to Solve it.* This similarity in heuristics suggests that Morin too conceptualizes the problem-solving process as something similar to the four-stage model put forward Polya's text.

However, Polya's account of the problem-solving is incomplete by default because the heuristics he advocates are too precisely tuned to their apparent objective. That is, both Polya's and Morin's heuristics and the associated problem-solving processes they suggest, are mostly concerned with ways to solve problems which are already well-formulated and where the most crucial step of the process, recognizing that a potential problem even exists, has already been performed for the student. In this way, the problem-solving process is mischaracterized by many educational treatments of problem-solving because these treatments rarely discuss how one could develop the problem or motivating question without the already provided context typically doing so.

So, instead of turning to educational accounts of the problem-solving process, we will attempt to find a more complete description of the process by outlining the stages one goes through to solve a problem in a real-world, albeit still scientific, context.

Constructed examples are sometimes unreliable arenas in which to explore ideas about procedure because one could possibly develop an example which, although illustrating one's point, does not have a legitimate mirror in the world. We will try to correct for this by, in the next section, connecting our depiction of the problem-solving process to historical accounts of Schrödinger and Feynman developing the wave equation and path integrals, respectively.

Confining our examples to scientists may appear to be an unnecessary limitation, but if the problem-solving education of science students has any universal value, it is primarily a value in the direction of simulating the knowledge and techniques actual scientists use to solve problems. More generally, students are taught statistics from statisticians, mathematics from mathematicians, and history from historians not merely because the academic practitioners of a discipline often have the best grasp of the discipline, but also because students' engagement with the discipline, either outside or inside academia, requires employing tacit skills and knowledge which only practitioners can model and eventually impart to students.

Thus, the student of statistics learns not only how to compute linear regressions and correlations, but also of the manifold ways that statistical data, without an account of the underlying analysis, can mislead and thus be misused. The student of real analysis learns not merely the standard proof that differentiability implies continuity and the fact that the converse implication has no proof, but also of the unreliability of intuition and the larger necessity of justifying claims, both mathematical and not, with rigorous argument. The student of history learns not only about America's mid-twentieth century excursions into foreign statecraft, but also that there is a larger context to existing cultural norms and political relationships, and that a country's conception of itself exists as just one point in one narrative that is still being lived and written.

Much in the same way that statistics, mathematics, and history courses simulate for students the best practices for engaging with data, quantitative arguments, and current events, physics courses simulate how best to analyze the underlying physical laws of a phenomena, and also, considering physics' role as a cultivator of problem solving ability, the best general practices for solving problems. So although not all physics majors will become physicists, it is still useful to consider their education as preparation for work in this direction.

Say we have a proverbial investigator who is solving a problem concretely in her work using methods she learned during her education. How and when does she get to the point that she is solving problems of the kind that she might have encountered in a modern college classroom? Ever before the investigator can solve a problem of the type found in a typical textbook or an assignment, she must create the problem. This involves first recognizing that something in her environment, intellectually or physically rendered, deserves inspection. This recognition in turn leads her to admit to some measure of personal confusion, incompleteness in her view or engagement with the world, or some need which must be fulfilled in order to achieve a larger goal. In short, there is a "state of tension" (Faludi 1973, pp. 82–4) which suggests that there is something in her space of knowledge that is missing and could potentially be filled. This space is often felt intellectually as a sort of nagging lack, ever before any concrete problem formulation is devised. But once a need is recognized, she can then work to formulate a framework (using the techniques and language that she has previously learned) to address it. This is the second stage of the true problem-solving process, and it is rarely a oneshot activity. There are always more possible frameworks than one could ever possibly investigate, and some are better than others.

Thus, the stage of developing a framework (i.e., creating the problem) exists somewhat cyclically with the stage of implementing the framework (i.e., solving the problem). Out of many possible problem frameworks, the investigator chooses an optimal one—perhaps judged by its solubility, esthetics, or how well it can address the original gap in knowledge—and then implements the framework to check whether it is as useful as initially perceived. If the framework proves to be less useful than desired, one returns to the problem development stage and the process begins again. After a potential problem is recognized, a problem framework is developed and the investigator can then work to solve the problem in the manner typically presented in course textbooks. This might involve her using heuristics, borrowed techniques, or even something simply algorithmic. It is this part of the problem-solving process that education most successfully models and teaches. But in a more realistic scientific or real-world context, the other parts of the process are just as necessary, for without them one could never reach this more concrete stage. We summarize this characterization of the problem-solving process in Fig. 1.

As a comparison, a similar problem-solving process was presented by Bransford and Stein (1993). Called the IDEAL model, it takes problem solving to consist of (a) *I*dentifying the problem, (b) *D*efining and representing the problem, (c) *E*xploring strategies for solving the problem, (d) *A*cting on the strategies, and (e) *L*ooking back and evaluating your results. Thus, the IDEAL scheme is similar to that in (Polya 2014). However, the scheme depicted in Fig. 1 differs from the IDEAL scheme (and most other characterizations of the problem-solving process) in that it incorporates the intellectual state which precedes the identification or definition of a problem. Ever before a problem is recognized to even exist, there is a state of confusion or incompleteness which is only acknowledged and responded to upon the decision to define a problem framework. It is these dual steps of recognizing and then articulating one's

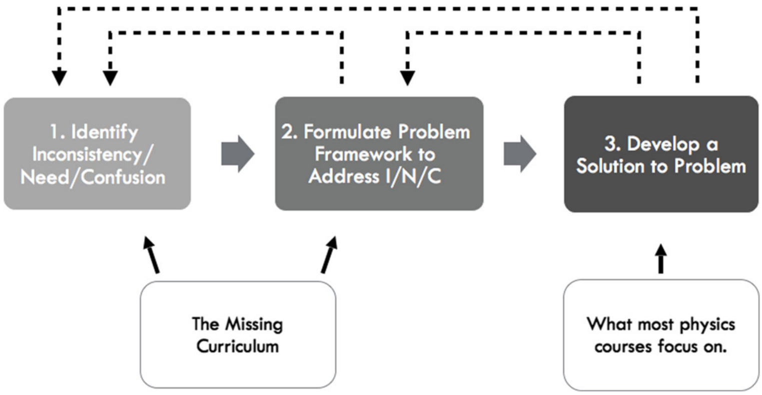

Fig. 1 The steps of the problem-solving process: identify a gap in knowledge or a need; formulate a problem framework to address gap; develop a solution to the problem framework. Physics courses provide extensive practice in solving well-formulated problems but often do not teach students to formulate problems themselves, or to identify and articulate the confusion or gap in knowledge which is the source of the problem. Although the steps are described as a linear sequence, as the investigator's understanding of the inconsistency or confusion develops, she returns to the earlier stages again and again to find better ways to frame, or even identify, the situation she is trying to understand

initial state of confusion as an appropriate problem framework which are rarely considered in most characterizations of the problem-solving process.

5 Historical Examples of the Process

We can see this model implemented in the way scientists throughout history solved problems. We will take Schrödinger's construction of the wave equation and Feynman's development of path integrals as examples and track each scientist's work from the start of their investigations to the point where they were solving problems akin to what one finds in a modern physics textbook.

5.1 Schrödinger and the Wave Equation

In 1900, Max Planck, in an attempt to explain the spectrum of the frequencies of light produced by a perfectly emitting heated body (i.e., a "black body"), gave an ad hoc derivation of what is now known as the "black-body radiation formula". Seeking to better justify this result through combinatorial methods that had been employed by Boltzmann, Planck re-derived the formula this time from the assumption that the total energy of the light consisted of integer multiples of $h\ \nu$, where ν is the frequency of light and h was a numerical quantity that we now call Planck's constant. In this assumption, historians recognize the beginnings of a theoretical quantum description of nature, but, at the time, Planck had no such recognition; the discrete energy assumption was introduced more out of an esthetic desire for an improved derivation rather than out of any more fundamental physical insight, and thus Planck was reluctant to interpret physically the implications of his formalism (Kuhn 1987, pp. 92–110).

Interpretation and understanding came in 1905 through Einstein's work on the photoelectric effect (Pais 1982, pp. 379–82). Attempting to explain why electrons are ejected from a metal when light shines on it, Einstein fully committed to Planck's discrete energy postulate and further postulated that not only did *emitted* light come in discrete quantities, but in fact all light was constitutively discrete and at sufficiently small scales exhibited particle like-properties.

In 1924, Louis de Broglie published his dissertation where he argued the converse of Einstein's photoelectric effect postulate. Namely, taking Einstein's work on photons and the newly developed theory of special relativity, de Broglie argued that similar to the way electromagnetic waves can have particle-like properties, particles can have wave-like properties (Mehra and Rechenberg 1982, pp. 582–603). He thus established that a quantum particle of a certain momentum p carried with it a wave with wavelength λ given by $\lambda = h/p$, where h is again Planck's constant.

This was the state of affairs when, in November of 1925, Peter Debye urged Schrödinger to review de Broglie's recently published thesis in a biweekly research seminar. To prepare for the review, Schrödinger worked through de Broglie's thesis and sent letters to Einstein and other physicists expressing his awe of the work and his desire to obtain a better sense of what it really meant to consider a quantum particle as a wave. In a letter to Alfred Landé, Schrödinger wrote:

> …I have been intensely concerned these days with Louis de Broglie's ingenious theory. It is extraordinarily exciting, but still has some very grave difficulties. I have tried in vain to make for myself a picture of the phase wave of the electron in the Kepler orbit. Closely neighboring Kepler ellipses are considered as rays. This, however, gives horrible 'caustics' or the like for the wave fronts. (Moore 1994, p. 192)

In this correspondence, we see Schrödinger recognizing the conceptual potency of de Broglie's results while also struggling to understand and fully utilize that potency. It seems Schrödinger's main road block was that he did not know the correct question to ask.

This question came during the seminar where Schrödinger presented de Broglie's work. In the seminar, Debye mentioned that de Broglie's wave-particle duality ideas were interesting but woefully imprecise. In his opinion, any system that truly had wave-like properties should also have a wave equation. Thus, Schrödinger's task was set: He needed to find a wave equation for quantum particles. Weeks later he obtained what is currently known as the Klein-Gordon equation. Weeks later he applied the equation to the study of the hydrogen atom and used it to compute the associated spectrum of energy levels. But the computed energy levels did not match the confirmed predictions of the Bohr's old quantum theory. Consequently, Schrödinger developed a new equation (what is currently known as the three-dimensional time-independent Schrödinger equation) and found that it reproduced the correct spectrum when applied to the hydrogen atom (Moore 1994, pp. 192–207).

In Schrödinger's experiences, we see the hallmarks of what we previously described as the problem-solving process. There was first the recognition of a gap in knowledge (Schrödinger's tentative search for a more concrete picture of de Broglie's wave-particle ideas and Debye's suggestion that a wave equation is necessary to precisely describe the duality); an imposition of a problem framework to address that gap (Schrödinger's development of the Klein-Gordon equation for the hydrogen atom); a solution of that framework (Schrödinger's calculation of the Klein-Gordon hydrogen atom spectrum which incorrectly predicted the known results); and when that solution was found to be inadequate, a repetition of the process from the problem framework stage (Schrödinger's development of a non-relativistic version of the wave equation and his later application to the hydrogen atom to obtain the correct spectrum).

5.2 Feynman and the Path Integral

In the early 1940s, Richard Feynman progressed through similar stages as he developed his path integral formulation of quantum mechanics and applied it to the problem of removing oscillator degrees of freedom from quantum particle interactions. The subsequent historical discussion is drawn from (Mehra 1994, Chapters 5 and 6) unless otherwise cited.

For context, by the end of the 1920s, quantum mechanics had just concluded a rapid few years of progress during which it had largely matured into its modern manifestation. By the end of that time, the quantum theories of the day were Matrix Mechanics, formulated around matrix operators and the Heisenberg uncertainty principle, and Wave Mechanics, based on Schrödinger's wave equation. The two theories accurately predicted the properties of various quantum systems and were shown by Schrödinger to be equivalent (Schrödinger 1926). With physicists feeling as though they well understood the quantum physics of particles, they next set their sights on understanding the quantum physics of the electromagnetic field and began developing what is now known as quantum electrodynamics (Schweber 1994, pp. 76–92).

Quantum electrodynamics was specifically treated by Dirac in the last chapter of his *Principles of Quantum Mechanics* (Dirac 1930), and also by Walter Heitler in *The Quantum Theory of Radiation* (Heitler 1936). However, as Feynman realized by the end of his undergraduate degree in 1939 (Mehra, The beat of a different drum, Mehra 1994), there were problems with both treatments.

> All I could remember in Heitler and Dirac... was that they could not solve the problems. They were getting infinities, and the last sentence in Dirac's book was that some new ideas were needed. That's the main thing I got from Dirac; that new ideas were needed. This, to me, meant that I did not have to study the old [theories].

Here, Feynman referred to the fact that the first calculations of the quantum properties of the electromagnetic field led to some physical quantities (like the energy of an electron) being equated to divergent integrals, that is, nominally infinite quantities. This quote reveals that from early on in his education, Feynman recognized fundamental problems inherent in the basic theories of quantum electrodynamics. Moreover, his dissatisfaction with the existing representations of the problem proved to be pivotal to his later work as it was only by attempting to develop an original solution that he was led to incorporating the principle of least action into quantum mechanics.

Now, Feynman also knew that there were “infinities” in classical electromagnetism as well; computing the classical electromagnetic energy of an electron due its interaction with its own field also yields a divergent quantity. Recognizing the divergence problems in both quantum and classical theories of electromagnetism, Feynman then hypothesized that the electromagnetic field was fundamentally an incorrect physical quantity, and he believed that the solution to both divergent problems lay in a formulation of electrodynamics in which the electromagnetic field did not consist of independent degrees of freedom, but was completely determined by the positions and velocities of charged particles. The idea was that without an independent electromagnetic field, particles could only act on other particles (that is, never on themselves) and thus there would be no divergent self-interactions.

In his work with John Wheeler, Feynman formulated such a classical theory of electromagnetism. This “action-at-a-distance” theory presented the electric and magnetic fields as arising solely from the motion of particles and gave finite values for the self-energy of an electron (Wheeler and Feynman 1945). The theory was later shown to be flawed because it could not reproduce standard results in classical electrodynamics, but it was still an important starting point for Feynman’s later work.

With the self-energy divergences of the *classical* theory of electromagnetism eliminated (although incorrectly so) Feynman’s next task was to use his formalism as a new basis for quantum physics. However, he faced one major problem: The classical theory he developed with Wheeler was grounded in Lagrangian mechanics and a principle of least action, while the quantum systems of the day were based on analogs of Hamiltonian mechanics. Now, there *is* a way to transform the Lagrangian of a system into the system’s corresponding Hamiltonian and also a way to write the principle of least action in terms of a Hamiltonian, but Feynman decided what was needed was a formulation of quantum mechanics which employed a principle of least action in terms of Lagrangians.

In the fall of 1941, 2 years into his graduate studies at Princeton, Feynman met Professor Herbert Jehle at a bar. Feynman talked with him about the problems he was working on, and in particular asked him if he knew of any way of doing quantum mechanics with a principle of least action. Jehle expressed personal ignorance in this direction, but also told Feynman about one of Dirac’s papers where Lagrangians were shown to be relevant to quantum mechanics. The next day, the two of them visited the Princeton library where they looked at Dirac’s paper. Dirac had indeed discussed the relevance of the Lagrangian to quantum physics, but there were ambiguities in his presentation, and there was no analytical derivation concerning the role the principle of least action might play. In the library with Jehle, Feynman clarified one ambiguity on a blackboard by showing that the operator which took a wave function from position x_1 at a time t_1 to a new position x_2 at a time $t_1 + \epsilon$ (for infinitesimal time ϵ) was proportional to the exponential of $i\epsilon/\hbar$ times the Lagrangian. He also showed that from this evolution operator one could derive the time-dependent Schrödinger equation. Feynman had thus found an explicit way to employ Lagrangians, but he still lacked a way to incorporate the action (i.e., the time integral of the Lagrangian) into quantum physics.

A few days after his meeting with Jehle, Feynman considered how the evolution operator might change if one were interested in evolution over long, rather than infinitesimal, times. By multiplying many of Dirac's evolution operators together for a sequence of times, Feynman found that the overall argument of the resulting exponential was a sum of the Lagrangians over the times considered. That is, in the continuum limit, he found the argument of the exponential was the classical action. Through this calculation Feynman had derived what we now call the path integral of quantum mechanics: he had demonstrated that to find the evolution operator for long times, one would need to integrate the exponential of the classical action over all forward-paths connecting the starting and ending points of the possible trajectories.

By this point, Feynman had succeeded in his goal of finding a way to incorporate actions and Lagrangians into quantum mechanics. But his larger goal of showing that such an incorporation allowed one to solve the divergence issues in quantum electrodynamics still lay remote. According to his initial program, he needed to show that it was possible to use this new formulation of quantum mechanics to eliminate the electromagnetic field from particle interactions. The specific problem of charged particles interacting through an electromagnetic field was analytically too difficult to solve, so Feynman considered a toy problem which contained many of the properties of the real problem. He considered a system of two particles which are not coupled to one another but interact through each of their individual couplings to a harmonic oscillator. The two particles represented two electrically charged particles and the harmonic oscillator represented for the electromagnetic field. With his newly formulated path integral representation of quantum mechanics, Feynman showed that it was indeed possible to eliminate (specifically, "integrate out") the harmonic oscillator degree of freedom from the system and thus reduce it to one which only included direct interactions between particles. He had thus provided a proof-of-concept demonstration of his larger goal of eliminating the electromagnetic field from quantum particle interactions.

We should mention two ironies in this historical development. First, Feynman as an undergraduate disliked employing Lagrangians and the Euler-Lagrange equations to solve mechanics problems because he felt they obscured a physical understanding of the system considered. But Lagrangians ended up being fundamental to his new approach to quantum mechanics (Gleick 1992, pp. 60–1). Second, his new approach had as its overriding goal, the elimination of the electromagnetic field in quantum electrodynamics and more generally, one could argue, the elimination of all fields from quantum field theory. But Feynman was not successful in either of these directions, and after his path integrals were generalized, they were instead found to be an exceedingly natural (if not the most natural) context to study quantum field theory with fields still considered fundamental (Weinberg 1995, Chapter 9).

In Feynman's path to path integrals, we see the basic steps by which a problem is solved in a real scientific context. We can take Feynman's solution of the toy model—where a harmonic oscillator was eliminated from a system of two particles—as the main problem he solved. But before he solved this problem, and before he even concretely formulated it as a problem one could solve, his main preoccupation was with the apparent contradictions of quantum electrodynamics. Namely, he first recognized that the existing formulations were producing unphysical "infinities," and he sought to fix these problems by developing a framework for classical and quantum physics where the electromagnetic field was not present. In his attempt to create such a framework, he was led to another problem of presenting quantum mechanics in the language of a principle of least action. The scheme he used to solve this latter problem was extrapolated from work by Dirac, and from it Feynman developed the path integral for quantum mechanics. With the path integral, he then used a toy model to investigate a conceptually and analytically simplified solution to the problem of eliminating oscillator degrees of freedom from particle interactions.

5.3 Inferences from Historical Examples

The steps of the problem-solving process that both Schrödinger and Feynman implemented were to first identify a point of confusion; second, to formulate a problem framework to address the confusion; third, with the framework established, to develop a solution by implementing what is typically known as the problem-solving process. When the obtained solution failed to solve the original problem, they returned to the problem formulation stage and tried to develop a different framework.

Key to the progress both Schrödinger and Feynman made towards the solutions of their respective problems were the questions they asked. Some of these questions were asked within the explicit context of the concrete problems they eventually solved, but the most important questions lay at the start of their investigations when they were formulating the problems themselves. Questions such as "Is there a way to frame the wave properties of the electron through a wave equation?" and "How can one study quantum mechanics through an action principle?" were the starting points to their works and without these initial questions, they would never have reached the more epistemically solid areas of work where more routine mathematical manipulations were applicable.

6 Reasons for the Missing Curriculum

The missing curriculum in problem-solving education exists in the first two steps of the problem-solving process depicted in Fig. 1, two steps which we have seen were crucial in Schrödinger's development of the wave equation and Feynman's development of path integrals. The problem-solving process as it is conveyed in typical physics courses today focuses almost exclusively on teaching the student to solve explicitly framed problems and rarely considers how those problems were developed or how to, for oneself, move from ill-defined queries to explicitly defined ones.

This is not a new claim. Two decades ago, Mazur (Mazur 1996), for example, mentioned the artificiality of physics problems, and Root-Bernstein (Root-Bernstein 1989) recognized that scientific work more legitimately concerns the recognition of problems rather than merely the solution of existing problems. More recently, Fortus (2009) has shown that even students with a BA in physics are incapable of transforming ambiguous problem contexts into well-defined problem contexts. So, given the prior recognition of the limitations of current educational models, why do courses still focus mostly on evaluating student answers rather than teaching them how to ask questions?

A cynical possibility is that the reason for the curricular gap is philosophical and deliberate rather than technical and unintentional. That is, maybe there is something about systems of education and how they exist in the larger structure of employment and capitalism that makes a focus on answers evaluated by figures of authority, rather than inquiries initiated by those lower in the institutional hierarchy, an eminently desirable model of education (Schmidt 2001, pp. 161–79).

Kuhn had a different perspective. He argued that an education in physics, framed primarily around learning the solutions to a large set of canonical problems—and not at all focused on the historical questions which led to those problems, how one might learn to ask such questions, or the competing interpretations of phenomena which might have led to a different set of questions—is the main reason science progresses so much more rapidly than other academic disciplines, and also why it experiences scientific revolutions as conclusively as it does (Kuhn 2012, Chapter 4).

Scientists who are trained and, in a sense, indoctrinated in a certain physical world view, then take that view for granted as they study physical phenomena. Inevitably they encounter phenomena which cannot be encompassed by this representation of the world, and in response they label this "incommensurable phenomena" as anomalies. These anomalies remain on the fringe of the subject until they are resolved by slightly extending the framework of the subject, or until they are revealed to have a scope much larger than anticipated. In this latter case, what Kuhn calls a "crisis" occurs and the scientific community (or at least parts of that community) becomes cognizant of a limit to their knowledge—a limit which was, indeed, always there—and they are then forced to develop a new physical framework to push back the boundaries of those limits.

This is part of the reason, Kuhn argues, that although scientific revolutions are perhaps the most dramatic form of progress observed in science, educational systems do not train future scientists to recognize when they are happening or to even initiate them, but rather to perform the more incremental and daily work termed "normal science." In a sense, students are trained to be soldiers, who work out the consequences of principles they have been taught are correct, as opposed to revolutionaries who question those principles and try to build new ones. The idea is that these principles do not need to be routinely questioned, and from the perspective of the scientific community at large it is simply more efficient for most practitioners to work out the consequences of existing principles than to be distracted by a needless search for some different foundation. In fact, however, the two modes of working are not entirely distinct as it is by working out the consequences of existing principles that a scientist learns of the areas where new principles are needed.

Thus, I am not arguing against Kuhn's point. Rather, I think there are better ways to train scientists to perform and fulfill the normal science function that Kuhn claimed is central to scientific progress. Namely, besides the subject and content knowledge of physics this training claims to provide, it also claims to provide students with the habits and heuristics needed to solve problems. But this second claim is rarely realized due to a misguided focus on only a single aspect of the problem-solving process, and, because the subject-based knowledge is most effectively developed through solving problems, the first goal is often not realized as well.

7 Learning to Formulate Questions

7.1 Becoming Effective Questioners

To effectively work through the often implicit and ambiguous initial parts of the problem-solving process, students need to become effective questioners. There are a number of ways to do this, and each way concerns modifying the types of problems students are assigned. In describing these modifications, I will use examples from my own teaching experiences and contextualize them with studies from the research literature.

In past courses, I have assigned tasks which require students to practice developing and answering their own questions about physical systems. Instead of giving students a physical system and then a problem concerning that system, I would give students the physical system and then ask them to generate several questions of their own concerning that system. When and where such tasks are included in the traditional assignments depends on the pre-requisites for the course and how much material the course has already covered. If an instructor is drawing from recently learned material in a new course, students may not have built the technical skills or the vocabulary to answer their own questions or even to formulate such questions in a way that they are answerable. But they could still practice these skills using material from the list of prerequisites for the class. For example, in a course on mechanical oscillations and waves, the first problem set I assigned included a task where students were asked to develop and answer three questions about a ball being launched from the top of a hill (Fig. 2), a task which was well within their skill-set given that classical kinematics was a prerequisite for the course. For this question, students were graded on the depth of their question (for example, asking "What is the angle at which the ball hits hill?" resulted in more points than the question "What is the acceleration due to gravity?") and the technical

5. **Projectile Motion on a Hill**

Consider the physical situation depicted in the figure below. A rock is thrown at a speed v_0 and at an angle θ (from the horizontal) from the peak of a hill which slopes downward at an angle ϕ.

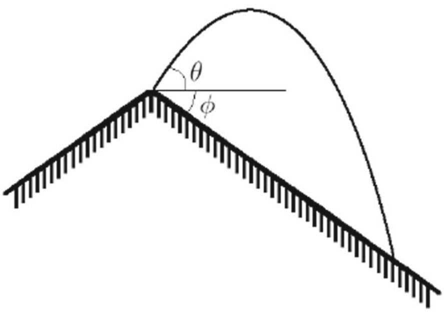

Figure 1: Projectile Motion

Provide **two** physics questions we could ask about this system. Answer both of your questions using your understanding of two-dimensional projectile motion.

Fig. 2 Example "What is the question?" problem. The objective of the problem is for students to practice asking and answering their own questions about physical systems correctness of their solution. As the course progressed, I drew on recently learned class material to develop these “What is the question?”-type problems.

These types of problems are similar to Jeopardy problems (Heuvelen and Maloney 1999; Rebello et al. 2007) where students are given an equation or figure and are asked to state the physical context in which the equation or figure might be applied. In both cases, the purpose of the inverted structure of the problem is to promote the flexible manipulation and representation of physics concepts in a way akin to how experienced physicist organize and apply their knowledge.

Besides assigning new types of problems, an instructor could be flexible in accepting the solutions a student is allowed to turn in for an assignment. As I was encouraged to do in my first research project, a physicist who finds that a main problem is too difficult to solve directly often turns to a simpler, but related, problem in order to obtain a qualitative sense of the properties of the first problem. Such problems are often termed “toy models” and there are numerous historical examples of physicists defaulting to toy models in the face of an insoluble desired problem and yet still finding something interesting: Feynman’s thesis was geared towards solving a toy model of quantum electrodynamics (Feynman and Brown 2005); As his first project in physics, Freeman Dyson was tasked with studying a toy model of what is now termed the Lamb Shift (Dyson 1948; Schweber 1994, pp. 497–500); and arguably the most famous soluble model in statistical mechanics is a toy model of magnetism first solved in one-dimension by Ernst Ising (Brush 1967; Taroni 2015).

To encourage students to pose and solve such toy problems as a general problem-solving strategy, an instructor could be flexible in accepting solutions by allowing students who find assigned problems intractable to solve a related problem of their own and to then explain how their solution is related to the potential solution of the original one. In students’ proposal of the new problem, they would be practicing creating problem frameworks to address gaps in knowledge (i.e., the second step of the problem-solving process depicted in Fig. 1); and in explicating how the proposed problem relates to the original, they would be practicing how to articulate their confusion through the technical language they previously learned. Of course, problems in course assignments are often chosen so that students practice specific skills, so this flexibility in accepted solutions would have to be allowed for certain instructor-chosen problems.

Problems in physics classes are often structured as “instruction manual” questions where the complexity of the problem is mitigated through a step-by-step outline of its intermediary parts. For each step, the student is asked to complete a specific task and each subsequent task builds on the one which came before. Such problems are good for showing students how the techniques they learn have larger applications in longer physics derivations or studying systems outside the purview of the course, but such problems are also detrimental in that they rarely convey the motivations for choosing the mediating steps and thus they lead students to underestimate the difficulty in developing such steps for themselves. By reducing a difficult problem to a simple series of steps, these instruction-manual questions belie the complexity of the original problem and suggest that the detailed calculations the student is required to perform, rather than the questions which set up such calculations, are the true work of the problem.

Such precisely defined problems differ greatly from the partially defined (also called “ill-structured) problems students often encounter outside their education. In “the structure of ill-structured problems” (1973), Simon provided a good description of the wider relevance of partially defined problems:

> In general, the problems presented to the problem solvers by the world are best regarded as [ill-structured problems]. They become [well-structured problems] only in the process of being prepared for the problem solvers. It is not exaggerating much to say that there are no [well-structured problems;] only [ill-structured problems] that have been formalized for problem solvers.

From (Fortus 2009), it is apparent that students with prior experience only solving standard physics course problems are not able to solve problems requiring extensive assumptions and with multiple possible solutions. In a study concerning problem solving strategies for astronomy students (Shin et al. 2003), researchers noted that partially defined problems required students to use case-based reasoning and apply their previous experiences. Therefore, students without those experiences found these partially defined problems too ambiguously defined to be soluble. Moreover, for unfamiliar contexts, such partially defined problems required students to plan and monitor their progress, metacognitive strategies not typically needed for problems with less open solutions. Milbourne similarly found that integrating content-knowledge with self-regulatory practices were necessary for advanced high school students to successfully tackling problems in their first research projects (Milbourne 2016). Also work from the last decade (Ogilvie 2009) suggests that such partially defined problems can lead to students engaging more deeply with the content of a course by forcing them to employ qualitative analyses and outlining of sub-problems to solve problems rather than superficial strategies like memorization and equation matching.

Therefore, to more completely represent the process of solving problems as they exist in research (or generally beyond formal instruction), instructors should work partially defined problems into a course. These partially defined problems are characterized by allowing for multiple approaches and answer and thus are closer to short essays than they are to the more mechanical derivations in a typical physics class. Modern sources for such problems are numerous. Thompson's *Thinking Like a Physicist* (Thompson 1987) has a good collection of such problems for physics, and XKCD creator Randall Munroe has assembled an extensive collection of answers to seemingly absurd physics questions largely framed as partially defined problems submitted by his readers (Munroe 2014).

Again, one must consider timing when assigning such problems. Students who have not learned how to use their acquired knowledge to solve well-defined problems will not be able to solve partially defined problems. As Fortus explained (Fortus 2009), the skills used for solving well-defined problems are necessary (although not sufficient) for solving partially defined ones. Thus, such partially defined problems could only reasonably be assigned later in the course. For example, in the oscillations and waves course, the first few weeks of the course were devoted to solving problems which were very explicit in purpose and structure, so that students could tell what questions they were answering and what techniques were relevant. But later in the course, I assigned a problem which did not reference techniques or even provide an explicit question, but rather asked students to analyze a physical model and use it to explain a classroom demonstration.

Also, there is a necessary comment about class size. These types of tasks work in situations where the students-to-teaching staff ratio is small enough so that the careful grading such assignments require is feasible. Since these problems allow students to submit a larger spectrum of solutions than those associated with typical single-solution problems, graders will have to devote more time to reviewing student solutions than is the case for more single-solution problem solving tasks.

7.2 Objections to Course Modifications

There is a straight forward objection to these modifications: The current model of physics education has a very specific purpose which can be explained through Bloom's taxonomy of learning (Bloom and Engelhart 1956). Higher order knowledge—analyzing the various parts of a subject, being able to synthesize new problems within the subject, and developing new ways to understand the subject—is, of course, the laudable goal of any physics education, but in order to get there the student must first proceed through the beginning stages of simple recall, rudimentary understanding, and applications in well-defined contexts.

This claim is sensible, but it should lead us to ask why physics education stops short. That education may begin with simple recall of basic concepts and equations and proceed towards understanding the relationship between those elements, but it usually concludes with teaching students to apply their gained knowledge to contexts which are externally prescribed and set. But if the benefits of such an education really do extend beyond physics itself, then the practices through which those benefits can best be reaped should be completely reflected in the undergraduate education in problem solving.

The authors of *Make It Stick* (Brown et al. 2014) advocate a teaching philosophy which is so sensible as to seem self-evident. The idea is that if teachers want students to perform in a certain way in contexts outside their education (e.g., in students' non-science careers or in their research), teachers must eventually evolve the qualities of the students' practice to coincide with the qualities of their intended performance. If classrooms are any indication of how we want students to ultimately apply their knowledge, it would appear that we want them to be quite skilled at solving problems which are clearly worded and presented in a way that implies a single solution. But, of course, what physics programs really want (or should want) is for students to be able to intelligently contend with the unavoidable messiness of the problems they will encounter after their education. Realizing this desire involves not only teaching students the way we want them to practice, but also testing students in ways that incentivize these methods of practice. Often if an assessment can be successfully navigated without employing the higher-levels skills touted as important for learning, students will default to the lower-level skills which at minimum ensure their favorable evaluation (Ramsden 1997). Consequently, akin to the problem modifications suggested above, and as advocated in Chapter 7 of (Bowden and Marton 2003), educators should strive to create assessments which include problems that are open, not explicitly technical, and which are novel to the students.

8 Conclusion: a Problem Half-Solved

Polya's *How to Solve It* (Polya 2014) is a text geared primarily towards learning how to solve mathematics problems—or even more narrowly towards solving mathematics competition problems—but it holds esteem as a general supplement to lessons on problem solving in many technical domains and therefore can be used as model for how many such courses view problem solving. In one of the early sections of the text, Polya begins an outline of the problem-solving process by exhorting the necessity of first understanding the problem. He writes, "It is foolish to answer a question that you do not understand. It is sad to work for an end that you do not desire."

I think this is a good starting point for Polya's purpose and a good lens through which to see our own. When one has a problem, the fundamental issue is of course understanding, but what one is trying to understand or clarify is not specifically the problem itself but something epistemically deeper which lies beyond the problem, a situation for which the stated problem is merely a proxy or a convenient means for articulation.

For Schrödinger, the well-defined problem was "What is the energy spectrum of hydrogen as computed from the wave equation?" but what he was trying to understand was a correct and conceptually motivated way to mathematically formulate the principles of quantum mechanics. For Feynman, the well-defined problem was "Can we represent a system of two quantum particles coupled by an oscillator as an interaction between the particles themselves?" but what he was trying to understand was whether it was possible to correct the divergence problems in quantum electrodynamics. In each case, the problem each physicist proposed served as a concrete handle for an idea which was otherwise difficult to engage with. The issue with most curricular answers to the problem of "problem-solving education" is their failure to understand, or, more generously incorporate, this fact of a situation beyond the problem, and their resulting mischaracterization of the real process through which problems are developed and solved.

The early twentieth century education philosopher John Dewey once said (Dewey 1938, p. 108) "A problem well put is half solved." So is it with the problems which comprise most of a physics education, and, as a result, students often learn only half the skills they need to be truly effective problem solvers. Physics problems are primarily pedagogical in nature and may consequently require simplifications and idealization of realistic analogs, but if these problems are to effectively model the problem-solving process, they must model all aspects of that process—recognition of a possible problem, problem formulation, and then problem solution—and not merely the aspects which are currently definitive of problem-solving ability.

Acknowledgements The author is grateful to Suzanne Smith and Harvard's Graduate Writing Center for help in editing an early draft of this manuscript. The author is also grateful to Amy Gilson and Ajoke Williams for comments on later drafts.

Compliance with Ethical Standards

Conflict of Interest The author declares no conflict of interest.

References

Bloom, B. S., & Engelhart, M. D. (1956). *Taxonomy of educational objectives: The classification of educational goals*. New York: David Mckay Co..

Bowden, J., & Marton, F. (2003). *The university of learning: beyond quality and competence*. London: Routledge.

Bransford, J. D., & Stein, B. S. (1993). *The ideal problem solver. A guide for improving thinking, learning, and creativity* (2nd ed.). New York: Freeman.

Brown, P. C., Roediger, H. L., & McDaniel, M. A. (2014). *Make it stick*. Cambridge: Harvard Univesity Press.

Brush, S. G. (1967). History of the Lenz-Ising model. *Reviews of Modern Physics, 39*(4), 883.

Burgasser, A. (Fall 2008). *8.012 Physics 1: Classical Mechanics.* Massachusetts Institute of Technology: MIT OpenCourseWare, https://ocw.mit.edu.

Committee on Undergraduate Physics Education Research and Implementation. (2013). *Adapting to a changing world—challenges and opportunities in undergraduate physics education.* National Research Council. Washington D.C.: National Academic Press.

Dewey, J. (1938). *Logic-the theory of inquiry* (p. 108). New York: Henry Holt and Company.

Dirac, P. A. (1930). *The principles of quantum mechanics*. Oxford: Oxford University Press.

Dyson, F. J. (1948). The electromagnetic shift of energy levels. *Physical Review, 73*(6), 617.

Faludi, A. (1973). *Planning theory*. Oxford: Pergamon Press.

Feynman, R. P., & Brown, L. M. (2005). *Feynman's thesis: a new approach to quantum theory*. Singapore: World Scientific Publishing Co..

Fortus, D. (2009). The importance of learning to make assumptions. *Science Education, 93*(1), 86–108.

Fratus, K. (2016). *Physics 20*. Retrieved from Physics UC Santa Barbara. http://web.physics.ucsb.edu/~fratus/SIMS16/

Gleick, J. (1992). *Genius: the life and science of Richard Feynman*. New York: Vintage.

Halliday, D., Resnick, R., & Walker, J. (2010). *Fundamentals of physics extended*. Hoboken: John Wiley and Sons.

Heitler, W. (1936). *The quantum theory of radiation*. Oxford: Oxford University Press.

Heuvelen, A. V., & Maloney, D. P. (1999). Playing physics jeopardy. *American Journal of Physics, 67*(252).

Kleppner, D., & Kolenkow, R. (2013). *An introduction to mechanics*. Cambridge: Cambridge University Press.

Kuhn, T. S. (1987). *Black-body theory and the quantum discontinuity, 1894–1912*. Chicago: University of Chicago Press.

Kuhn, T. S. (2012). *The structure of scientific revolutions*. Chicago: University of Chicago Press.

Mazur, E. (1996). The problem with problems. *Optics and Photonics News, 7*(6), 59–60.

Mehra, J. (1994). *The beat of a different drum: The life and science of Richard Feynman*. Oxford: Clarendon Press.

Mehra, J., & Rechenberg, H. (1982). *The historical development of quantum theory: The quantum theory of Planck, Einstein, Bohr and Sommerfeld. Its foundation and the rise of its difficulties (1900–1925)* (Vol. 1). New York: Springer-Verlag.

Milbourne, J. D. (2016). Self-regulation in the midst of complexity: a case study of high school physics students engaged in Ill-structured problem solving. PhD dissertation, North Carolina State University.

Moore, W. J. (1994). *A life of Erwin Schrödinger*. Cambridge: Cambridge University Press.

Morin, D. (2008). *Introduction to classical mechanics: with problems and solutions*. Cambridge: Cambridge University Press.

Munroe, R. (2014). *What if? Serious scientific answers to absurd hypothetical questions*. New York: Hachette UK.

Nelson, A. E., & Seiberg, N. (1994). R-symmetry breaking versus supersymmetry breaking. *Nuclear Physics B, 416*(1), 46–62.

Notre Dame Department of Physics. (2009). *Career tracks*. Retrieved from Physics at Notre Dame: http://physics.nd.edu/undergraduate-program/career-tracks/

O'Byrne, J., Mendez, A., Sharma, M., Kirkup, L., & Scott, D. (2008). Physics graduates in the workforce: does physics education help? *Congress of the Austrailian Institute of Physics*.

Ogilvie, C. (2009). Changes in students' problem-solving strategies in a course that includes context-rich, multifaceted problems. *Physical Review Special Topics - Physics Education Research, 5*(2).

Pais, A. (1982). *Subtle is the lord: the science and the life of Albert Einstein*. Oxford: Oxford University Press.

PHYS 111 001 - Rice's Course Schedule (2015). Retrieved from Rice University: https://courses.rice.edu/admweb/!SWKSCAT.cat?p_action=COURSE&p_term=201510&p_crn=10258

Polya, G. (2014). *How to solve it*. Princeton: Princeton University Press.

Ramsden, P. (1997). The context of learning in academic departments. In F. Marton, D. Hounsell, & N. J. Entwistle (Eds.), *The experience of learning* (Vol. 2, pp. 198–216). Edinburgh: Scottish Academic Press.

Rebello, N. S., Cui, L., Bennett, A. G., Zollman, D. A., Ozimek, D. J., Rebello, N. S., et al. (2007). Transfer of learning in problem solving in the context of mathematics and physics. In D. H. Jonassen (Ed.), *Learning to solve complex scientific problems* (pp. 223–246). London: Routledge.

Root-Bernstein, R. S. (1989). How scientists really think. *Perspectives in Biology and Medicine, 32*(4), 472–488.

Schmidt, J. (2001). *Disciplined minds*. Lanham: Rowman and Littlefield.

Schrödinger, E. (1926). On the relation between the quantum mechanics of Heisenberg, Born, and Jordan and that of Scrodinger. In *Collected papers on wave mechanics* (pp. 45–61). New York: Chelsea Publishing Company.

Schweber, S. (1994). *QED and the men who made it: Dyson, Feynman, Schwinger, and Tomonaga*. Princeton: Princeton University Press.

Sharma, M., Pollard, J., Mendez, A., Mills, D., O'Byrne, J., Scott, D., … Low, D. (2008). What does a physics undergraduate education give you? A perspective from Australian physics. *European Journal of Physics, 29*, 59–72.

Shin, N., Jonassen, D. H., & McGee, S. (2003). Predictors of well-structured and ill-structured problem solving in an astronomy simulation. *Journal of Research in Science Teaching, 40*(1), 6–33.

Simon, H. A. (1973). The structure of ill-structured problems. *Artificial Intelligence, 4*(3–4), 181–201.

Taroni, A. (2015). Statistical physics: 90 years of the Ising model. *Nature Physics, 11*(12), 997.

Tesfaye, C. L., White, & Susan. (2014). *High school physics textbooks, resources and teacher resourcefulness.* AIP Statistical Research Center. College Park: American Institute of Physics.
Thompson, N. (Ed.). (1987). *Thinking like a physicist: physics problems for undergraduates*. Bristol: Hilger.
Weinberg, S. (1995). *The quantum theory of fields: foundations* (Vol. 1). Cambridge: Cambridge University Press.
Weinberg, S. (2005). *The quantum theory of fields: applications* (Vol. 2). Cambridge: Cambridge University Press.
Wess, J. (2000). From symmetry to supersymmetry. In G. Kane & M. Shifman (Eds.), *The supersymmetric world: the beginnings of the theory* (pp. 67–86). Singapore: World Scientific Pub Co Inc..
Wheeler, J. A., & Feynman, R. P. (1945). Interaction with the absorber as the mechanism of radiation. *Reviews of Modern Physics, 17*(2-3), 157.