Introduction to Proofs Over Easy: A Low-Cost Alternative to the Flipped Classroom

Lewis D. Ludwig

Abstract

A modification of the flipped-classroom technique is presented, which is referred to as the "over-easy" method. In it, students use reading notes to engage in proof-writing by first learning to read mathematics, then transferring that skill to their proof-writing. A description of the technique including several examples is provided and several references from cognitive psychology supporting its advantages. The article concludes with the benefits and challenges of the over-easy method and other applications of it.

Difficulty Level: High; Course Level: Transitional

1 Background and Context

Eleven years ago when I arrived at Denison University, a small national liberal arts college of 2,100 students, I was asked to teach our introduction-to-proofs course. It is a bridge course, required for the math major and minor, that helps students transition from calculus to the higher-level proof-based courses. Traditionally we offer one section in the fall with 15-25 students who are math majors or minors in their sophomore year. The verbal/math SAT scores for these students average in the upper 1200s to lower 1300s.

I have taught the course five times in the last eleven years. When first assigned to it, like any young faculty member I was eager to challenge my students and help prepare them for this critical juncture in their mathematical careers. I carefully chose a text that was challenging, yet readable. I made detailed lecture notes by carefully following the text and adding interesting examples and exercises from other sources. When classes started, I gave interactive lectures, which included a presentation of my notes at the board and a series of well-crafted questions to make sure students were understanding and participating in the discussion. As we covered various proof techniques I would provide propositions that needed to be proved and would stand at the board coaching the class through the process arriving at a completed proof. All the while students took detailed notes.

At the end of the semester, my course evaluations were positive. Students liked how the class was structured, liked the interaction in the classroom, and felt they had learned a great deal. My classroom observations from senior members in the department were glowing, emphasizing how well I interacted with students and kept them engaged. The next time I taught the class, I took the same approach and got similar reviews from students and colleagues.

However, something didn't seem quite right. It felt like I was doing almost all the work. I did all the reading, found interesting examples, assimilated the new material, and then transferred this knowledge to students. They were doing little information-gathering or higher-level thinking aside from taking notes in class and chiming in when I asked questions. Moreover, when I had the same students in subsequent advanced courses such as real analysis, I found they had difficulty following the longer proofs from the text. They could create their own proofs of small propositions for homework, but struggled to read and comprehend a detailed proof from the text such as the Bolzano-Weierstrass or Heine-Borel theorem.

After a good deal of reflection and talking with others at Denison and beyond, it was clear that the course should not only teach students how to *write* mathematical proofs, but how to *read* mathematical proofs as well. At the time, my daughter was in early elementary school. A large portion of her time and effort was given to *reading* instruction. While she knew her letters and had spelling tests, her writing instruction did not get serious until several years later, after she could read.

It seemed the same should be true for my proofs course. Before students can become effective writers of mathematics, they first need to become proficient readers of mathematics. This notion is supported by the work presented by Fitzgerald and Shanahan [4]. They provide research showing the strong link between developing reading and writing skills. For example, they look at the domain knowledge about substance and content. This includes the "meanings or ideas that are constructed through the context of the connected text." They also discuss the importance of procedural knowledge which "refers to knowing how to access, use, and generate knowledge in any of the areas ..." They argue that reading and writing are connected through domain knowledge and procedural knowledge, and that one informs the other.

Armed with this knowledge, I began looking for ways to adjust my students' learning experience in my proofs class so that they developed their reading skills as well as their writing skills. In addition, I wanted them to have a more active role in their learning by taking ownership of the gathering and assimilation of information. Over the past nine years this has developed into what I call the "over-easy" classroom – a modification of the flipped classroom.

2 Description and Implementation

2.1 Motivation

Before I describe the method, let me first provide motivation. Considering my prior technique of creating class notes and delivering an interactive lecture, my colleagues from psychology would describe my students as passive learners. True, I did engage students in class by asking leading questions, but they did nothing before class to prepare for this. Moreover, their learning was dependent on me, the instructor. I was treating them as empty vessels, waiting to be filled when they came to class. The filling came in the form of their passively taking notes and answering questions.

In a sense, I was just transferring knowledge to my students. They had little or no responsibility in the information-gathering phase of their learning. This limited the amount of time and energy they spent assimilating the information. To truly engage in their learning, students need to fully understand the material and place it in the context of their existing knowledge [1] or domain knowledge, as described by Fitzgerald and Shanahan [4]. So my motivation for developing the over-easy method was two-fold: improve my students' ability to read and assimilate information, and have them take greater ownership of their learning starting with the information-gathering stage.

2.2 Method: The Why and How

First, why the over-easy classroom? Traditionally, a flipped classroom involves a video lecture that students watch before coming to class. The information-gathering stage of the learning process is done before class instead of in a lecture, giving ownership of the process to the student. In class, the instructor focuses on clarifying the information gathered by having students work problems in class alone or in groups. This allows for immediate feedback and often closer instructor-student interaction [3]. We will see that the over-easy method is neither as time nor resource intensive as creating videos, but is just a modification of the development and use of traditional class notes. Compared to a traditional flipped class, it is over-easy.

So how does it work? As described in the introduction, I used to take detailed notes from the textbook and supplement them with additional exercises and questions. With the over-easy method I do the same thing, but without providing the answers or details. That is, I create a handout that is a skeletal outline of what would previously be considered my class lecture notes and have the students fill in the details as homework before class. I call these handouts "reading notes." There are two kinds of reading notes: information-gathering and assimilation of knowledge, or proof-reading practice.

Let's consider a typical set of reading notes that focuses on information-gathering and assimilation of knowledge in my introduction-to-proofs class. Suppose I plan to cover the section on indexed sets on Wednesday. By the prior Monday, students will receive a handout that is partially depicted in Figure 1[1]. Outside of class, the students use the text to complete at least 70% of the handout by the beginning of the following class, which in this case would be Wednesday.

Section 1.4 Indexed Collections of Sets

1. Give a working definition and example of:
   (a) union of an indexed collection of sets
   (b) intersection of an indexed collection of sets.
2. Find the union and intersection of each of the following families of sets.
   (a) For $n \in \mathbb{N}$, let $A_n = \{1, 2, 3, \dots n\}$.
   (b) For $n \in \mathbb{Z}$, let $B_n = [n, n+1)$.
   (c) For $n \in \mathbb{N}$, let $C_n = [0, \frac{1}{n})$.
3. Give an example of an indexed collection of sets $\{A_\alpha : \alpha \in \Delta\}$ such that $A_\alpha \subseteq (0, 1)$, and for all $\alpha \in \Delta$ and $\beta \in \Delta$ we have $A_\alpha \cap A_\beta \neq \emptyset$ but $\bigcap_{\alpha \in \Delta} A_\alpha = \emptyset$.

Figure 1: A portion of a sample reading assignment.

I begin class on Wednesday by randomly putting students in pairs with their reading notes in hand. Students spend the first 10-15 minute of class comparing their notes, checking for agreement in their responses or discussing disagreements. I circulate through the room, checking for completion and listening to the discussions. If a student asks a question, I reply by first asking what his/her partner thinks about the question. (Initially, students are more comfortable asking me than their peers for help. The reminders help students better engage with one another.) Often times students are able to resolve their own questions, but I design the reading questions so that the students will be significantly challenged by roughly 30% of the questions, hence the 70%-completion threshold. As I circulate, I note which questions are giving difficulty and will center the day's activities around them.

Let's consider the questions from Figure 1. Question 1 asks for students to give a working definition. For a working definition, students should read the definition from the text and rewrite it as if they were explaining it to a friend. I discourage verbatim copying from the text as very little assimilation occurs if they do this. It is important for them to provide examples, different from the text, for each definition. This helps in

[1] Ample white space is usually provided so students can write on the handout. It has been removed to save space.

assimilation and helps students put the new information into their existing framework of knowledge. That is, how does the new information relate to things they already know? For example, students already know about finite unions. How can they extend this understanding to infinite unions? What are the similarities? What are the differences? Anderson et al. [2] would refer to this as the second level of the cognitive domain of Bloom's taxonomy of learning objectives: understanding. I expect everyone to be able to sufficiently answer question 1 as it is straightforward information-gathering. A quick check with their peers will confirm this and I have saved five minutes of lecturing.

Question 2 is more involved. Here I am expecting students to put the new definitions into practice. This is the third level of cognitive domain in Bloom's Taxonomy: application. Most students will independently get 2(a). The majority will understand 2(b) and about 30-40% will get 2(c) since it is dealing with new definitions and a limiting process. By pairing the students at the start of class, 2(a) is confirmed, 2(b) is clarified, and 2(c) will be something we discuss as a whole class. The process also demonstrates instructional scaffolding where support is given during the learning process that is tailored to the needs of the learner in order to help the student achieve his/her learning goal [5]. In addition, I have saved at least another five minutes of lecturing.

Question 3 requires an even higher level of understanding. Unlike question 2 where I asked them to apply the new definitions to three concrete examples, in question 3 they must create their own examples. This encompasses the highest level of Bloom's taxonomy, the creation of new ideas. The observant student will notice that a slight modification of 2(c) will lead to a useful example, but generally only 20% or fewer students see this. This question will play a large part in my discussions with the whole class. By having students consider this question on their own and with a partner, I need less time to explain the notation, etc., and can focus on problem-solving techniques.

This is roughly how the over-easy method works for information-gathering. With the over-easy method, I spend the same amount of class time on the material as when I lectured. However, the time I used to spend on definitions and basic examples is now allocated to more challenging ideas and concepts.

The reading notes that involve proof-reading practice are slightly different; an example is in Figure 2. This is the typical proof, attributed to Euclid, found in most introduction-to-proof texts. While it is concise and beautiful to those fluent in mathematics, it is a challenge for students new to mathematical proofs. The proof itself is only four lines long and takes about 30 seconds to read. However, students soon realize that the ability to read words does not guarantee understanding of the logical argument.

The questions for this type of exercise mimic how a mathematician would read a mathematical proof. First, do I know all the definitions? This in the intent of question 1. Students need to remind themselves that 1 is not considered prime. Next, what type of proof technique is being invoked? The notes explain the proof-by-contradiction technique to the students and then ask what the two outcomes are that establish the contradiction. What assumption is being made? If the students understand the two outcomes, infinitely many primes or finitely many primes, then question 2 should answer question 3. What role does P play? Herein lies the beauty of Euclid's argument. To understand the proof, the student must realize that P may or may not be prime, but in either case, the conclusion still follows. The last two questions secure the contradiction and conclude the proof.

I would start class with students' comparing their answers to the handout in Figure 2. The 70%-completion threshold still applies. Generally by the end of their discussion with their partners, most students will get 1-3 and 5 fairly easily. Questions 4 and 6 are the stumbling blocks where I will focus the class discussion.

With my prior interactive-lecture method, I would spend nearly the whole class on the proof. A good portion of time was spent just copying down the proof. With the reading practice, we can cover the proof in class in about 20 minutes with a fairly high level of confidence that everyone understands what is going on. To ensure this, we spend the rest of the 30 minutes in class tackling other proof-by-contradiction problems. This is usually done with a handout containing several additional statements that the students try to prove.

Reading practice: There are Infinitely Many Primes

Theorem 1 *There are infinitely many primes.*

Proof. Suppose that $p_1 = 2 < p_2 = 3 < \ldots < p_r$ are all the primes. Let $P = p_1 p_2 ... p_r + 1$ and let p be a prime dividing P; then p cannot be any of p_1, p_2, ..., p_r, otherwise p would divide the difference $P - p_1 p_2 \ldots p_r = 1$, which is impossible. So this prime p is still another prime, and p_1, p_2, ..., p_r would not be all of the primes. QED

1. Define what it means for a number to be prime.

   This proof technique is called *proof by contradiction*. The author assumes there are only a finite number of primes and based on the premise comes to a contradiction. Thus, the assumption must be incorrect and the alternative true.

2. For this proof by contradiction there are **only** two possible outcomes. What are they?
3. How many primes does the author assume exist?
4. Could P be prime? Sometimes? Always? Never? EXPLAIN.
5. Why can't p divide $P - p_1 p_2 \ldots p_r = 1$?
6. What is the contradiction?

Figure 2: An example of proof-reading practice.

They work in pairs on the questions for about 20 minutes while I circulate. The last 10 minutes is used to discuss their findings with the class.

A typical set of fact-gathering reading notes (Figure 1) is three to four pages. So in 10 minutes, my students can cover roughly 70% of the material I would have normally covered in a 50-minute interactive lecture. As described, a similar time savings occurs with the reading practice notes (Figure 2). By condensing the time spent on the more mundane information-gathering portion of the class, we have more time to tackle the issues and concepts that the students find most challenging. The over-easy method makes better use of class time and students are significantly more engaged in their learning. Some of the better class days are when the reading notes contain some challenging questions that no one got. My students arrive in class complaining, and I love it! I now have them emotionally hooked on their learning. They are no longer passively absorbing information, but now have a vested interest in seeing how a problem is solved. Making an emotional connection with learning is very powerful [1].

3 Outcomes

As anyone who has stood in front of a classroom knows, there is no one way to teach. However, I have incorporated the over-easy method into all my classes for majors (multivariable calculus on up). In them, at least 80% of the class meetings are conducted with this method. While I am a firm believer in the over-easy method, it does have its limitations. I will now discuss some its benefits and several of the challenges. I will also provide several approaches that address them.

3.1 Some Benefits

Students become better readers of mathematics. For example, in the subsequent real analysis course, the proof of the Heine-Borel theorem is still challenging, but the students are not as overwhelmed as they were in the past. When faced with a new statement to prove, instead of just reading the words, students now transfer the skills they learned from the reading practice notes to better understand the new statement. The skills include considering concrete examples of the statement, the definitions needed, the proof technique used, the assumptions made, and the key element(s) of the proof. All the skills are developed and reinforced with the over-easy method.

Students became better writers of proofs. By developing their domain knowledge and procedural knowledge as described by Fitzgerald and Shanahan [4], my students no longer came to me saying "I have no idea how to get started on this problem." Instead, they would say things like "I tried this" or "I have this idea." That is, they were reading the questions and understanding the domain knowledge, and then trying to generate their own procedural knowledge. With one or two suggestions from me, they were usually on their way to a correct proof. Prior to the over-easy method, I often felt like I had to reteach the lecture during my office hour. This is no longer the case.

As I refined the over-easy process, I found I was able to ask my students more challenging and involved questions. Moreover, instructors, myself included, found that they were able to hit the ground running in courses like abstract algebra or real analysis. As another piece of interesting anecdotal proof, members of our Career Exploration and Development Department who help our students craft resumes and cover letters recently noted that they enjoy working with the students from my class as they are "such fine writers."

Misconceptions are more easily identified and rectified before the midterm or final. Prior to this method, students would often have misconceptions that did not manifest themselves until a graded situation. With the over-easy method, many misconceptions are caught and corrected as the students compare their results with one another and the class. This also helps the students better organize their knowledge so they can more readily retrieve it for use and application [1].

Students are more willing to ask questions. In the past, one of my least answered questions to the class was "Are there any questions?" With the reading notes, students have a clear idea of what they need to understand and seek help when they do not.

I rarely see the comment "We need to do more exercises in class" on my course evaluations. This was a very common response when I used my interactive-lecture method. Since I had to present definitions and examples, we had less time as a class to go over exercises. Now the bulk of class time is spent working on examples and exercises.

3.2 Some Challenges

How do you hold students accountable? If not everyone gives a solid effort on the reading notes, the system crumbles. In a real analysis class of juniors and seniors who were familiar with the system, I asked them how I should hold them accountable for the reading notes. The class agreed that it was in their self-interest to do the work and I need not worry about holding them accountable. This turned into a New Year's resolution. By the fifth week, five of my twelve students were not consistently doing the reading notes which totally derailed the whole system from the pairing-and-sharing down to the individual student's ownership of his/her learning. I used this as a learning experience. Now my students get two "get-out-of-jail-free cards" for the semester. That is, they can miss two reading assignments over the semester. For n assignments missed beyond that, I deduct $\sum_{k=1}^{n} 2^k$ points from the total grade of about 400 points. Students quickly catch on that this is a series they don't want to challenge. Moreover, as an incentive, if the students do not use their get-out-of-jail-free cards, they can cash them in on the final exam to skip one question.

Students must buy into this method. I often tell my junior colleagues that teaching is 57% salesmanship. You constantly have to remind students why you are using the method. You can't just tell them it is good for them like broccoli. I share my pedagogical methods and reasonings with my students so they better understand why we are doing what we are doing.

The method requires more work from students. My classes meet four days a week for fifty minutes per class. I average two to three sets of reading notes per week. In addition, I have weekly homework and a test every four weeks. Since adopting the over-easy method, I have shortened the weekly homework sets since many of the questions are now covered in the reading notes.

The method requires timely planning. Due to the larger work load and responsibility now placed on the students, I always give them at least two days to complete a reading exercise. This takes more planning and foresight than preparing a lecture the night before. If a typical set of lecture notes took me one hour to prepare, the over-easy method adds roughly an extra 30 minutes of preparation from me. But the time gained in instruction and the fact that I can easily edit the notes for successive classes makes the extra work well worth it.

4 Extending the Method

The over-easy method can easily be modified for other classes. The reader can see how this method can be used in other proof-based courses like real analysis or abstract algebra, but I have used it in other courses such as multi-variable calculus and linear algebra. The questions in multivariable calculus and linear algebra are often more computation-based, but the general idea is the same. In linear algebra, I often use true or false questions to test students' understanding of the definitions and theory, for example: "The columns of any 4×5 matrix are linearly dependent (EXPLAIN)." To receive full credit, students need the correct response and a supporting argument. This question tests the students' understanding of linear dependence and dimension.

While these are ideas that I have developed and refined over the past decade, they are not unique. I know David Pengelley of New Mexico State University uses a similar technique in a variety of classes and levels of students. While no system is perfect or will work for everyone, I have found the over-easy method to be a way to engage my students in their learning and help them become better readers and writers of mathematics.

References

[1] Ambrose, S., Birdges, M., Dipetro, M., Lovett, M., Norman, M., and Mayer, R., *How Learning Works: Seven Research-Based Principles for Smart Teaching*, Jossey-Bass, San Francisco, 2010.

[2] Anderson, L., Krathwohl, D., Airasian, P., Cruikshank, K., Mayer, R., Pintrich, P., Raths, J., and Wittrock, M. *Taxonomy for Learning, Teaching and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives*, Longman, New York, 2000.

[3] Berrett, D., "How 'Flipping' the Classroom Can Improve the Traditional Lecture," *Chronicle of Higher Education*, February 19, 2012.

[4] Fitzgerald, J., and Shanahan, T., "Reading and Writing Relations and Their Development," *Educational Psychologist*, 35 (2000) 39–50.

[5] Simons, K., and Klein, J., "The impact of scaffolding and student achievement levels in a problem-based learning environment," *Instructional Science*, 35 (2007) 41-72.