FLIPPED LEARNING

A Guide for Higher Education Faculty

Robert Talbert

Foreword by Jon Bergmann

Stylus

STERLING, VIRGINIA

Published by Stylus Publishing, LLC.
22883 Quicksilver Drive
Sterling, Virginia 20166-2102

Library of Congress Cataloging-in-Publication Data

[To come]

13-digit ISBN: 978-1-62036-431-4 (cloth)
13-digit ISBN: 978-1-62036-432-1 (paperback)
13-digit ISBN: 978-1-62036-433-8 (library networkable e-edition)
13-digit ISBN: 978-1-62036-434-5 (consumer e-edition)

Printed in the United States of America

All first editions printed on acid-free paper
that meets the American National Standards Institute
Z39-48 Standard.

First Edition, 2017

10 9 8 7 6 5 4 3 2 1

PREFACE

What This Book Is About

This book is about *flipped learning,* an approach to the design and instruction of classes that involves the following simple changes to the way classes are usually conducted:

- Instead of having students gain their first exposure to new concepts and material *in class* (often through a lecture), we set up ways for students to have that first exposure *prior* to class, along with ways to guide them through that first exposure.
- Because all this is taking place *before* class, it frees up large chunks of time *during* class that can now be spent on the activities where students typically need the most help, such as applications of the basic material and engaging in deeper discussions and creative work with it.

This "flipping" of the contexts of student work—first exposure to new concepts before class, then deeper and higher work during class instead of the other way around—is, as we said, a simple idea. But it is also a profound one, and it has the potential to transform education, particularly higher education, in ways we are only beginning to understand.

I began using flipped learning to organize and teach my classes several years ago out of necessity (more on that in a minute) and gradually expanded the scope of my flipped learning designs until, today, every class I teach uses flipped learning: high-level topics courses for majors in my discipline as well as general education courses, large sections and small sections, and face-to-face courses and online courses. The results have been undeniable. Students in these classes not only learn the content of their course as well as or better than in a traditional course setup, they learn to take more responsibility for their work, make greater strides toward becoming self-regulating learners (a concept we will explore in detail in the chapters ahead), work harder and smarter during class time, and are generally *happier* with the course. In fact,

in my department (where many of my colleagues are also on board), students now openly complain when a course *doesn't* use flipped learning. As some of our students say, "How are we supposed to learn anything if all the teacher does is *lecture* during class?"

Students *get it*: Higher education today does not have to look like the higher education of yesteryear, or even of a decade ago, when the only reasonable way to guide students through first exposure to new concepts was through having everyone in the same room on a fixed schedule where the information could be disseminated by an expert. In fact, today's students are acutely aware that in a highly competitive job marketplace, their future as professionals depends on their ability to learn on their own, to be able to pick up new skills and ideas on an as-needed basis and then put them to work on difficult, creative applications that matter to them. Someone who *needs* a traditional classroom setup to learn anything will be left behind. They need a higher education that "gets it" as much as they do.

And lest we think of higher education only as job preparation, the notion of *lifelong learning* also is predicated on the idea of having not only the ability but the *taste* for learning new things on one's own, letting one's curiosity and interests guide one into new ideas and skills, and translating passions and interests into actual knowledge and abilities. Higher education, by its very structure, ought to communicate the vital importance of lifelong learning—but often it fails, again by its very structure, which communicates the opposite of what we want: that learning take place during set times of the day, a set number of times of the week, and only after an expert tells you what you need to know. But those of us who are, or at least like to think that we are, lifelong learners know a different story. Learning is unbounded, deeply personal, deeply rooted in our native human ability to reason for ourselves. The structures higher education places around this process *used* to serve the learning process but are today increasingly *served by* the learning process, so that lifelong learning is not the point any more—attending class and taking tests (that are taken straight from a lecture) and getting high scores is the point. Higher education needs a reboot to come back to the core principle of lifelong learning.

I believe flipped learning is an idea whose time has come. It's a way of doing higher education that shows unique promise for simultaneously getting it back to its intellectual roots while ushering it into the future.

At the same time, flipped learning is, at its heart, a very *simple* idea. How can one possibly write an entire book about it? Well, it turns out that implementing this simple idea is not, itself, all that simple. In fact, this book began with that first attempt at flipped learning that was a pretty massive failure.

Productive Failure

The year was 2009, and I was on the mathematics faculty at a small liberal arts college. As is the case in small colleges, I also did a lot of administrative and service work. One of those administrative jobs was to manage our college's "3+2" engineering program. Like other similar programs at small colleges throughout the United States, students in the 3+2 program attended my college for their first three years to complete the basic science, math, and general education courses needed for an engineering degree; then they transferred to a partner institution, a larger university with an engineering program to take their specialized engineering courses. After five years they received two degrees, one in applied mathematics from my college and one in engineering from our partner.

This was a great program but tricky to manage because of the omnipresent scheduling constraints in small colleges. For example, we could only offer one section of Calculus 2 per year. If there were any schedule conflicts between my college's courses and our partners', it was a serious problem. So to help ameliorate this problem, I proposed a course on computer programming for mathematics majors, using the software MATLAB, that would satisfy a basic course requirement at our partner university.

To ensure sufficient enrollment each time it was offered, the course was made a requirement for *all* math majors, not only the engineering students but also students who don't normally take scientific computing courses, like preservice primary and secondary education majors. The audience therefore included many students who weren't typically very interested in—or comfortable with—the idea of computer programming. To give the full range of students the time and space to learn the subject, I proposed a three-credit course meeting three times a week for 50 minutes each time.

By the time the course was debated and modified by our dean, it had been whittled down to a *one*-credit course—meeting once a week for 75 minutes. This was a serious, if not crippling, problem because I had just one concept in my mind about how a college course should be designed and managed. That setup is the one that I encountered exclusively in college as well as for most of my courses in graduate school, and it goes like this:

- During class meetings, the students in the class get their first contact with the material through a (hopefully well-constructed and logically sequenced) lecture that includes lots of instructor examples, and maybe if there is time, there will be some practice time made available.
- After class is over, the students are tasked with going back to their dorms, homes, and apartments to work on bigger and better

things—things like homework, projects, programming, assignments, and so on.

But immediately I realized that this model simply won't work for a first course on computer programming, aimed at a general audience, done once a week for 75 minutes. First, computer programming is a skill that cannot be learned by listening to a lecture—one has to *do* programming, make mistakes, fix the mistakes, and reflect on those mistakes in order to really learn the subject. I knew this from my own computer programming courses in the past, which had been taught exclusively via lecture. But because I was a skilled learner back in those days and loved technology and had a technical background, it was no problem for me to make the time outside of class for the *real* learning experiences. Second, the students in *my* class were not like I was. They weren't dumb by any means, but they were not expert learners and did not in many cases possess even a basic level of comfort with technology, much less skill. Third, even if I were to lecture flat out with no breaks for 75 minutes once a week, and even if the students in the course could absorb all that lecture and put it to good use in their own time, we still wouldn't be able to learn everything we needed to learn in the course.

I needed a different model.

So I did some research, and I eventually came across an article called "The Inverted Classroom in Software Engineering" (Gannod, Burge, and Helmick 2008), by three professors in the Computer Science Department at Miami University of Ohio. Reading their paper and how their software engineering course was set up, I knew this was an immediate fit, and I modeled my course after theirs.

These professors at Miami had devised something called the "inverted classroom" (which I later came to realize had been invented by some of their colleagues in the Economics Department, eight years earlier). In the inverted classroom, the direct transmission of information that one normally associates with a classroom lecture still happened, but it happened in a different context: *outside*, actually *before*, the class meeting in the form of prerecorded video and readings. Then, because this activity was relocated outside the class meeting, the meeting was wide open to be spent on active learning activities. Not only this, but having the lectures prerecorded allowed students to pause and replay as often as they needed without having to lose themselves in note taking.

After one reading of this article, I knew that this was the model I was looking for. Students would watch videos about certain feature sets of the software prior to the weekly class meeting. [MATLAB maintains a collection of professional-quality instructional videos for this purpose (see MathWorks

2016)], so I didn't need to *make* any video content, just link to what had already been made and add some viewing notes. Then I created weekly programming labs in which students would partner and apply the basic feature set of the week—which they learned through watching the videos—to solve an applied problem using MATLAB. Students were assessed on the quality of their solutions as well as on a handful of timed assessments.

This seemed very promising. Students would get to know the software on a basic level before class, thereby liberating tons of time in class to do creative application work. Thus would the flipped learning idea be implemented solidly in my course. This was going to be *awesome*!

Except, as it turned out, it was not so awesome. In fact, it was a resounding failure on a number of fronts:

- I gave students well-crafted video resources, but that's all I did—*gave them* with instructions to "watch this video before class." It turns out this is not enough. Students tried to watch the videos but quickly got lost, did not learn what they needed to learn to be ready for class, and did not retain what they had learned from previous weeks.
- Because students weren't learning from the video preclass assignments, they were showing up to class unready to work on applications of the basic material. Those labs that I so painstakingly constructed for class time? They were nonstarters for the students.
- Students asked me to reteach the material from the videos because they didn't understand them. I refused—maybe I should say "declined"—because I assumed that the reason students were asking was because they had simply refused to watch the videos out of laziness. Perhaps some students *were* lazy, but it's more likely that many of them simply didn't know how to watch a video in the sense of *learning* something from it. Things got tense in the class as a result.
- Students became frustrated that I "wasn't teaching the material." I was mad at the students because, in my mind, they weren't working on their preparatory material hard or well enough. The community atmosphere of the class collapsed in the second week of the semester and never recovered. There was no trust, no sense of mutual help, and honestly very little got learned that semester by anybody.

In short, it was one of the worst experiences in my 20-year teaching career, and it was mostly my fault, and by all rights I should have dropped this model then and there and never looked back.

However, I remained convinced that although going back to the traditional model of instruction might make students (and myself) *happier*

it would not result in improved learning, because, again, one cannot learn computer programming by listening to someone talk about computer programming. One has to *try* to program, with the time and space to make and repair mistakes while working on an interesting problem that requires basic skills. And those basic skills can, and therefore should, be acquired independently prior to class time, while reserving precious class time for me to *help* students as they do the hard and frustrating work of getting computer code to behave correctly.

In other words, I faced a choice after that initial disaster of a semester: Either try to figure out what was missing from my teaching in the class and then try again or abandon the idea altogether. I chose the former, and it was one of the best decisions of my career.

The Audience and Structure of This Book

The model that I have described, called the "inverted classroom" by the professors at Miami (and still often called by that name), is the pedagogical model we will call *flipped learning* throughout this book. As mentioned earlier, it is called "flipped" because of the reversal or flipping of the contexts of what students do in their own time versus their time with the entire class. This book, to a great degree, is an extended message to you about what I have figured out in the decade or so since that first attempt, in hopes that you can learn from my mistakes without actually *making* those same mistakes.

I strongly believe that flipped learning is the future of higher education, and therefore all of us in higher education need to be devoting time, attention, and effort in bringing this model into our classes. I also believe that adopting flipped learning cannot happen effectively without a solid set of basic resources and a community to help practice it. In my experiences as a workshop facilitator, speaker, and writer on flipped learning, I've found that although there is a strong interest in flipped learning in higher education, there are few "one-stop shop" resources for how to actually implement it, and only isolated pockets of communities of practice. This situation isn't sustainable. So one of my main hopes for this book is that it would be such a resource for instructors and a starting point for a robust community of practice about flipped learning in higher education.

There are three kinds of instructors who might find this book useful. I think you will find yourself in at least one of these categories:

1. *The curious.* These are college and university instructors who have heard about flipped learning (or perhaps the more popular term, the *flipped*

classroom; we'll discuss terminology more in chapter 1) and are interested in learning more. For you, the hope is you'll use the parts of this book of most interest to you and come away with a sound idea of what flipped learning is and how it might be useful to your students. I also hope that you'll consider graduating to the next category.

2. *The newbies.* These are college and university instructors who are just beginning to practice flipped learning. Maybe you've tried it and it didn't go so well (see my earlier story) and are curious to know what might be missing. Or maybe you are having a great experience with it so far and want to know how to expand your use. Either way, the hope is that you'll find practical tools and research results to help take your use of flipped learning to the next level, and that you will ultimately make it a permanent part of your teaching practice.
3. *The veterans.* These are college and university instructors who have used flipped learning for at least one entire course and are looking to hone their craft. For you, I hope this book provides a bigger picture, some good tips you can use, and some research background that can deepen your understanding of the practice. I also encourage you to parlay your experiences into scholarship that you can share publicly and so build the community of practice. We'll discuss this last item in the final chapter.

I've tried to write this book with each audience in mind, and with a balance of *research* on flipped learning and its theoretical bases, *course design* concepts that will help you set up courses to use flipped learning effectively, *tips and examples* from real-life usage of flipped learning, and *practical considerations* such as obtaining buy-in from students and getting students to do the preclass activities.

The book is *not* addressed toward K–12 (primary or secondary) schoolteachers, unless you are a teacher working with Advanced Placement or International Baccalaureate students or a similar "college-like" audience. This is because the K–12 world already has a great flipped learning ecosystem in place, with several excellent books on the subject (e.g., Bergmann and Sams 2012) and strong communities of practice. It doesn't need this book. The K–12 resources are well worth examining if you are a college or university instructor. However, despite their excellent quality, their applicability to higher education is limited due to the sometimes vast differences between the assumptions and contexts of higher education and K–12 education. For example, in higher education we often assume that students have 24/7 access to a high-speed Internet connection through their institutions; no such assumption can be made for students in a public high school.

Included in our audience is any person involved in education at a higher level—not only professors at four-year institutions, but also community and two-year college instructors, Advanced Placement or International Baccalaureate instructors in high schools, adjunct instructors who teach classes across multiple institutions or online, even parents with kids who are hungry to learn more. I will often use the terms *colleges* and *universities*; these are always intended as synonyms for any institution of higher education at any level.

This book is divided into three main parts:

- **Part One: What Is Flipped Learning?** focuses on understanding exactly what flipped learning is, what it isn't, and why instructors have chosen to use it (chapter 1); the research basis and theoretical frameworks that support the concept of flipped learning (chapter 2); and case studies that show different ways that flipped learning is used across academic disciplines (chapter 3).
- **Part Two: Flipped Learning Design** is all about the practical process of designing (or redesigning) a course around flipped learning. We'll discuss course design paradigms in general and specifically how the model offered by Dee Fink can be adapted to a flipped learning environment (chapter 4), look at a *seven-step model* for flipped learning design that stresses the development of appropriate learning objectives (chapter 5), and describe a sustainable model for preclass activities that students actually complete (chapter 6).
- **Part Three: Teaching and Learning in a Flipped Learning Environment** transitions from course design to issues of how to make flipped learning a permanent part of your teaching practice. We'll look at variations on the basic flipped design of classes to apply flipped learning to both online and low-tech courses, as well as the idea of "partial flipping" (chapter 7). Then, finally, we look at how to make flipped learning a sustainable and productive everyday part of your career, addressing issues that commonly arise with students as well as ways to parlay flipped learning experiences into your scholarship and service portfolios (chapter 8).

If you don't read the entire book, it's recommended that you read all the way through Part One and then dip into the other chapters as needed.

Flipped learning is an exciting and effective way to teach, and I can't wait for you to learn more about it.