

Reflections of a Textbook Author

N. Gregory Mankiw

Harvard University

March 6, 2019

Abstract

In this essay I reflect on textbook writing after three decades participating in the activity. I address the following questions: What perspective should textbooks take? What is the best approach to teaching microeconomics? What is the best approach to teaching macroeconomics? How does the content of the introductory course evolve? How much material should textbooks include? Are textbooks too expensive? How is digital technology changing the market for textbooks? Who should become a textbook author?

I am grateful to Denis Fedin, Rohit Goyal, Catherine Mankiw, Deborah Mankiw, Rohan Shah, Jane Tufts, and Nina Vendhan for their suggestions on an early draft.

For better or worse, I have spent much of my career writing and revising undergraduate textbooks. My intermediate macro text is now in its tenth edition, and my introductory text is going into its ninth. In this essay, I reflect on what I have learned about this activity over the past several decades. These reflections will, I hope, be of some interest to those who are thinking about becoming textbook writers themselves and to the larger number of instructors who use textbooks in their courses.

At the outset, I should confess to a perhaps peculiar fondness for economics textbooks. When I entered college, I did not intend to major in economics. But I remember picking up the textbook a friend was using for a course (it was Lipsey and Steiner's principles text) and being fascinated by the material. The use of straightforward mathematics, of the sort everyone learns in high school, to shed light on how society functions was, to me, novel, elegant, and compelling. I started taking economics courses the next semester, and my career began.

As a student, I often enjoyed reading textbooks. When I wanted to learn about some field but could not fit the relevant course into my schedule, I would find a leading text for that subject and read it on my own. Some of my fondest memories of college involve finding a comfortable chair in the library and reading a good textbook. Yes, I was (am) a nerd.

During my years in the economics profession, many of my most rewarding moments have come from my role as a textbook author. I don't know the exact number, but I would guess that, including translations, about 4 million copies of my books are in print. It is common for me to meet a random person—such as a host in a restaurant—who, after hearing my unusual last name, asks me if I am the author of her textbook. And many young economists have told me that reading my principles text got them interested in the field (much as the Lipsey and Steiner book did for me many years ago).

I have organized this essay around several questions about textbook writing and the textbook industry that often come up in conversations with colleagues. If a particular question does not strike you as interesting, feel free to skip that section and go to the next. My goal here—as when I write my textbooks—is to convey the information that readers want to learn in the most efficient way possible. I hope that some readers will even enjoy the experience, especially those who are nerds like me.

What perspective should textbooks take?

I have always thought that instructors, especially in introductory courses, are like ambassadors for the economics profession. The role of ambassadors is not to represent their own views but to act as agents for their principal. Just as ambassadors are supposed to faithfully represent the perspective of their nations, the instructor in an introductory course (and intermediate courses as well) should faithfully represent the views shared by the majority of professional economists.

On some topics, economists are in near unanimous agreement. Most economists agree that closing off a nation's trade with tariffs and other restrictions reduces the welfare of most citizens, that rent control is a highly inefficient way to help those at the bottom of the economic ladder, and that most tax cuts do not generate sufficient growth to increase tax revenue. I feel comfortable, as an instructor and textbook author, representing these perspectives as the consensus judgments of the profession.

Yet many issues are more controversial. Economists disagree about whether a higher minimum wage is desirable, how distortionary the tax system is, and whether monetary policy should be rule-based or discretionary. In such cases, I feel obliged to explain both sides of the

debate. I try my best to present both sides fairly, and I avoid revealing my own views (though I am surely imperfect in this endeavor). For example, the last chapter of my principles text considers six debates over macroeconomic policy. For each issue, I present brief essays both pro and con, trying to argue each side with the verve that its advocates would muster.

Deciding which views command a consensus and which are still hotly debated is, to some extent, a judgment call. Having spent about 40 years in the economics profession, attending dozens of conferences and speaking with hundreds of economists, I feel confident making those judgments. Nonetheless, while writing the manuscripts of my texts, I benefited from the input of a great editor and from pre-publication reviews by many instructors. They were invaluable at pointing out when the views being expressed were biased or idiosyncratic. That is, they told me when I wasn't being a faithful ambassador, allowing me to correct course before the book went to print.

In the most recent edition of my principles text, I include a new feature that helps ensure I fulfill the role of faithful ambassador. Over the past several years, the IGM Economics Experts Panel has surveyed several dozen leading economists on various applied topics. When my book discusses material related to one of these topics, an "Ask the Experts" box appears in the margin presenting the IGM survey results. Students can then see when economists agree, when they disagree, and when they are generally uncertain about a particular conclusion.

This perspective of instructor as ambassador raises the question of what instructors should do if they hold views far from the mainstream of the economics profession. If you are an Austrian or Marxist economist, for example, what should you do if asked to teach an introductory course? In my view, there are only two responsible courses of action. One is to sublimate your own views and spend most of the course teaching what the mainstream believes,

even if you disagree with it. Because many introductory students will take only one or two courses in economics throughout their educations, it would be pedagogical malpractice, in my judgment, to focus on an idiosyncratic minority viewpoint. The other responsible course of action is to avoid teaching introductory (and even intermediate) courses entirely. In a more advanced elective, there is nothing wrong with teaching an idiosyncratic minority viewpoint, as long as students know what they are getting.

For the same reason, I am wary of so-called “niche” textbooks that offer a peculiar angle on the field. For example, asymmetric information, political economy, and behavioral economics are important topics in mainstream economics. But if you want to reframe the entire field of economics through the lens of asymmetric information, political economy, or behavioral economics, your views become more idiosyncratic. Reasonable people can disagree about the dividing line between idiosyncratic and cutting edge. Indeed, for knowledge and pedagogy to evolve, instructors and textbook authors must be willing to try new things. But if the emphasis and range of topics in your introductory course look very different from comparable courses taught at most schools, you should probably question whether you are faithfully executing your job as ambassador.

Even within these boundaries, however, instructors and textbook authors have some discretion. Mainstream economics is so extensive that introductory covers can cover only the most essential topics, and reasonable people can disagree about what is essential. Moreover, the order of topics matters as well. Psychologists tell us that we are all subject to anchoring bias—the tendency to weigh information we receive early more heavily than information we receive late. Let’s therefore move on to discuss more concretely how micro and macro should be taught.

What is the best approach to teaching microeconomics?

Supply and demand are at the heart of how market economies work. When teaching the introductory course, therefore, it is important to develop and apply the tools of supply and demand as fully and consistently as possible. This tenet was my guiding beacon as I drafted my principles text.

But doesn't everyone agree with this? Haven't supply and demand always been at the center of the introductory course? Surprisingly, no. The first edition of Paul Samuelson's great text, published in 1948 and 608 pages long, did not introduce supply and demand curves until page 447. That is in part because Samuelson, writing in the shadow of the Great Depression, began his book by emphasizing Keynesian macroeconomics. As the book was revised over many editions, standard microeconomic tools became more prominent. But even today, many introductory courses do not develop the framework of supply and demand as fully as they should.

In particular, welfare economics is sometimes not given sufficient coverage. The basic tools of welfare economics are consumer surplus and producer surplus, which are natural extensions of supply and demand. Consumer surplus is the area between the price and the demand curve, and producer surplus is the area between the price and the supply curve, so these concepts are best taught soon after students have learned about supply and demand. Teaching welfare economics early and prominently in an introductory course has three advantages.

First, it gives students a deeper understanding of where supply and demand curves come from and how they are similar to each other. When explaining consumer surplus, one develops the idea of consumers' willingness to pay and shows how this underpins the demand curve. Similarly, when explaining producer surplus, one draws the link between producers' costs and the supply curve. When students learn these ideas together, they recognize that producer surplus and consumer

surplus, like supply and demand curves, are parallel constructs.

Second, the tools of welfare economics enable students to understand market efficiency. If there is one thing that separates economists from mere mortals, it is an appreciation of the power of markets as a mechanism for allocating scarce resources. Economists have known this lesson at least since Adam Smith introduced the metaphor of the invisible hand, and it should be a key topic in any introductory course. After all, it explains the biggest economic development of the last century—the success of decentralized capitalist economies and the failure of centrally planned communist ones. The best way to teach market efficiency is with the tools of welfare economics. Using not much more than supply and demand curves, students can learn that the market equilibrium maximizes the size of the economic pie as measured by the sum of producer and consumer surplus.

Third, after introducing the basic concepts of welfare economics, supply and demand curves can be used to address a greater range of policy questions. How do taxes affect market efficiency, and which kinds of taxes impose the smallest deadweight losses? Who wins and who loses when a country opens itself to international trade, how do the gains of the winners compare with the losses of the losers, and how do tariffs affect economic well-being? How do externalities, such as pollution, affect the efficiency of market outcomes, and how can government policy remedy the market failure? These questions are important, easy for students to understand, and well addressed using the tools of supply, demand, producer surplus, and consumer surplus.

When my principles text was first published, some instructors told me that my extensive and early coverage of welfare economics was innovative. Yet when writing the book, I never thought of it that way, because we had long taken this approach at Harvard. If there is anything novel about it, the credit goes not to me but to Marty Feldstein, who spent many years in charge of Harvard's introductory course, including when I taught a section under his leadership in 1985. (I took over the

role of course head from Marty in 2005.) My views on the pedagogical importance of welfare economics trace back to his, which in turn stem from his interest in practical issues of public policy.

What is the best approach to teaching macroeconomics?

My career as a textbook writer began in 1988. My department chair had asked me to teach intermediate macroeconomics, a required course for Harvard economics majors. I happily accepted the assignment and continued teaching intermediate macro for the next 15 years (stepping away only when asked to take over the introductory course). As I prepared for the course by surveying existing texts, I realized that none of them fully satisfied me. While many were excellent books, I felt that they did not provide the right balance between long-run and short-run perspectives, between classical and Keynesian insights.

In particular, I decided that the leading texts of the time were too focussed on Keynesian economics. (The exception was Robert Barro's book, which was relentlessly classical.) Although my training as a student of Alan Blinder, Stan Fischer, and Larry Summers was heavily Keynesian and my early research was often called "new Keynesian," I thought that the then-standard texts gave insufficient attention to the ideas of classical macroeconomics.

When my intermediate text was published in 1991, it found a ready audience. Many instructors apparently shared my views about the right mix of topics. (At about the same time, Andy Abel and Ben Bernanke published their own intermediate macro text, which also gave a larger role to classical models.) Shortly after my book came out, publishers started approaching me to ask whether the organizational ideas that were transforming the intermediate macro course might be applied to the introductory level. I said they could and, after some initial hesitation about taking on another big project, decided to try to do it myself.

What does it mean to give a greater role to classical macroeconomics? Put simply, it means giving more attention to the forces that shape the economy in the long run. On the real side of the economy, this means spending more time on the theory of economic growth, the role of financial markets in equilibrating saving and investment, and the determinants of the natural rate of unemployment. On the nominal side of the economy, it means spending more time on classical monetary theory, such as the quantity theory of money, the Fisher effect, the causes of hyperinflation, and purchasing-power parity.

Both my intermediate macro text and the macro portion of my principles text follow the organizational strategy of teaching these long-run classical ideas before introducing short-run Keynesian ideas. This approach has several advantages.

First, notwithstanding Keynes's famous quip about asymptotic mortality, long-run issues are extraordinarily important for human welfare. Consider: In 1900, Japan had less than half the income per person of Argentina. Now Japan has more than twice the income of Argentina. Why is that? What does it mean for the future? What can poor countries do to replicate the Japanese experience of rapid growth? What can rich countries do to avoid Argentina's fate of relative decline? These questions get students excited about studying economics. And they are questions best addressed with long-run, classical models.

Second, classical macroeconomics is more closely linked to the lessons of microeconomics. After students have learned about how the forces of supply and demand govern market economies, they are ready to apply these tools to the questions of macroeconomics. This is precisely what the classical model does. The classical model is built on the foundations of supply and demand—for labor, for loanable funds, and for money.

Third, the theory of short-run fluctuations is more easily understood after a solid grounding

in the economy's long-run equilibrium. According to standard theories, the business cycle represents a transitory deviation of the economy from its trend growth path. From this perspective, it is natural to study the determinants of trend growth before studying what pushes the economy temporarily away from that trend.

Fourth, short-run fluctuations are more complex than long-run growth. This follows simply from the classical dichotomy—the theoretical tenet that nominal variables (such as the money supply and the price level) do not influence real variables (such as real GDP and unemployment). The classical dichotomy allows long-run macroeconomics to be broken up into smaller, more easily digestible pieces. After absorbing each of these pieces, students are better prepared to study the short-run business cycle, to which the classical dichotomy does not to apply.

Fifth, the macroeconomic theory of the short run is more open to debate than the macroeconomic theory of the long run. Although I believe that the traditional model of aggregate supply and aggregate demand remains the best framework for understanding the business cycle, not all economists concur. By contrast, few economists today dispute the ideas of classical economics as a description of the economy's long-run equilibrium. In my view, it is more pedagogically sound to begin the study of macroeconomics on the firm ground of consensus.

This change in teaching strategy toward early and more extensive coverage of classical economics is now well entrenched and is most clearly seen in the textbooks for higher level courses. At the intermediate level, my text and the text by Abel and Bernanke adopted this approach in their first editions, and they both remain among the best-selling books for this course. The same is true for the more recent intermediate macro text by Chad Jones. At the graduate level, David Romer's superb text *Advanced Macroeconomics* teaches the theory of growth before the theory of the business cycle. It is no surprise that an increasing number of principles texts are adopting this

approach as well.

As a sign of how times have changed, imagine asking a group of introductory students the following question: If Americans decided to save a larger fraction of their income, how would this change affect the economy? The answer I learned as a freshman in 1977, studying macroeconomics from Paul Samuelson's celebrated text, was based on the Keynesian cross and the paradox of thrift: Higher saving rates depress aggregate demand, reduce national income, and in the end fail to result in higher quantities of saving. By contrast, the first answer I teach as an instructor today is based on classical growth theory: Higher saving means more investment, a larger future capital stock, and a higher level of national income. Most economists now agree that both answers have some degree of truth, depending on the circumstances, and that students need to learn both perspectives to understand debates about public policy.

How does the content of the introductory course evolve?

The previous two sections describe the perspective on micro and macro that I had when initially writing my principles book many years ago. While this overall pedagogical approach has remained constant over nine editions, many details have changed over time.

One reason textbooks evolve is that knowledge advances. For example, one of the most important developments in economics over my lifetime has been the rise of behavioral economics, the subfield that combines economics and psychology to challenge economists' conventional assumption of rational decision making. Behavioral economics did not exist when I was a student, and it was not mentioned in the first edition of my principles text. But as the subfield has become established, I would have failed in my role as the profession's ambassador if I did not add a discussion of it. And I did add a section on behavioral economics in the third edition.

Keeping up with the changing field of economics is important not only because students deserve access to state-of-the-art thinking but also because the evolution of thinking influences policy debates. Consider policies aimed at getting people to save more. Conventional models that assume rational decision making emphasize incentives. According to these models, simply expanding access to IRAs and 401k plans should encourage people to save more by reducing the taxation of capital income. Yet, according to some recent research, policymakers may be better able to pursue their goals by focusing on behavioral mechanisms. In particular, because consumers exhibit inertia in their behavior, the default matters. By making enrollment in 401k plans automatic, rather than an option that requires deliberate action, policymakers can promote greater saving.

Another reason textbooks evolve is that the world changes. The financial crisis of 2008 offers a vivid illustration. It has long been known that banks operate with limited capital and that, because of leverage, this capital can disappear quickly if their risky bets turn bad. Yet early editions of my principles book did not discuss bank capital or leverage; I thought these topics were best left for more advanced courses in financial institutions or money and banking. The events of 2008 upended that judgment, and now these topics are discussed prominently.

Although changing world events will always require courses and textbooks to evolve, my goal when writing my books is to provide students with a set of tools robust enough to be applied to new and unexpected challenges. The recent Trump administration is a case in point. Among the most significant economic policies enacted by the administration have been sizable tariffs on many imported goods. Tariffs are a topic I have discussed early and thoroughly in my principles book since the first edition. Perhaps I should send Mr. Trump a thank-you note for making Chapter 9 so relevant to the current generation of students.

As knowledge increases and events raise new challenges, it seems natural that the

curriculum should expand. But we should remember that students' capacity to absorb information does not grow at the same time. That brings me to my next topic.

How much material should textbooks include?

For many years, Otto Eckstein ran the introductory course at Harvard. Unfortunately, I never met him as he passed away just before I joined the faculty. But I have heard one of his aphorisms. Apparently, Otto often told section leaders, "The less you teach them, the more they learn." What I believe he meant by this is that instructors should avoid overwhelming introductory students with too much information all at once.

When I tell my non-economist friends about my principles text, I sometimes say that it is a short book—it's only 800 pages. They think I am joking, but I am not. Many principles texts are closer to 1000 pages, and some use a double-column format to cram more words onto each page. The amount of material I have left out would be enough in itself to make a normal-length book.

It is easy to understand why textbooks are so long. A publisher incurs large fixed costs when producing a textbook, as it hires people to prepare everything from test banks to websites. To protect this investment, the publisher sends the manuscript to numerous reviewers, who offer crucial feedback to the publisher and author. At the same time, however, the reviewers make plugs for their own pet topics. Reviewer A wants a discussion of the Herfindahl index, reviewer B wants a presentation of the kinked demand curve model, reviewer C wants a definition of the Gini coefficient, and so on. The editor, who has little training in economics, then recommends adding material on the Herfindahl index, kinked demand curves, the Gini coefficient, and the rest of the long list. The easiest path for the author is to follow the advice. When the book is published, professors A, B, C, etc.—and especially their students—all agree that it is too long.

As economists, we teach our students about scarcity. As instructors and textbook authors, we should remember that student time is a scarce resource. We must avoid making our courses encyclopedic. That means taking out all of the easily ignored details and stressing the big ideas. The main goal of the introductory course is not to produce future economists but to produce well-informed citizens. Any topic that a person does not need to understand to intelligently follow the news is a plausible candidate for omission.

One risk when simplifying matters for students is oversimplification, losing too much of the nuance that economists bring to an issue. But given the difficulty some students have learning basic economics, it is a bigger risk to overcomplicate the analysis early in the course. To be sure, students should leave the course with a nuanced understanding of the field. But first we must achieve clarity. To keep students engaged and on the right track, we need to start simple and add complexity as their understanding grows.

Consider how a child learns mathematics. A toddler starts by counting: 0, 1, 2, 3, and so on. These whole numbers, however, are highly simplistic and fail to solve all sorts of problems. To solve $x + 2 = 0$, we need negative numbers; to solve $2x - 1 = 0$, we need fractional numbers; to solve $x^2 - 2 = 0$, we need irrational numbers; and to solve $x^2 + 1 = 0$, we need imaginary numbers. Counting with whole numbers sweeps all that nuance under the rug! But that is okay. An instructor who tries to teach too much too early runs the risk of intimidating or overwhelming the student.

Learning economics is similar. The first chapter of my introductory textbook introduces ten big principles of economics. They are aimed at the student who has never studied economics before. The goal is to give a brief and understandable introduction to the field, a general sense of

where we are heading, and a foundation for the material to come. Adding a lot of detail at this early stage in the course would sacrifice too much clarity.

Subsequent chapters of my book build on these ten principles, and the discussion becomes fuller and more refined. In the second chapter, students start to learn about economic methodology, the distinction between positive and normative statements, and the value and risks of simplification in model building. More elements are added as various substantive topics are discussed. The micro section of the book ends with a chapter on asymmetric information, political economy, and behavioral economics. The macro section of the book ends with a chapter on several open debates about macroeconomic policy. Even though students do not leave chapter one with a full appreciation of the subtleties of economic thinking, they should finish the book with it. And for those students who major in economics, future courses will build on the foundation set in the introductory course.

When writing my text, I made many hard decisions about what to exclude and, sometimes, had to battle my editors over these choices. The biggest debate concerned my decision to omit the Keynesian cross (sometimes called the income-expenditure model), which has been at the center of teaching macroeconomics since Samuelson introduced it into the introductory course with his textbook in 1948. I am not opposed to the substance of this model, and I include it in my intermediate text, but I felt that the added level of complexity was unnecessary for introductory students.

As a college freshman in my first economics course, I found the Keynesian cross confusing. (I recall puzzling about whether $Y = C + I + G$ was an identity, an equilibrium condition, or both.) And this model did not connect well with other things I was learning in economics; only later did I realize that the absence of cogent “microfoundations” was a problem

for much of macroeconomics. As a result, when I started writing my own principles text, I looked for a better approach. My goal was to explain the concepts illustrated by the Keynesian cross model without introducing a complex geometric apparatus that tended to confuse introductory students more than it enlightened. Most students who take the introductory course are not going to become economics majors, and those who do will likely explore the model in sufficient depth in intermediate macroeconomics.

As I see it, the essence of the Keynesian cross is four related ideas: (1) An economy can end up at an equilibrium where it produces less than its potential output. (2) Changes in the aggregate demand for goods and services can affect that equilibrium. (3) Fiscal policy is one thing that can change aggregate demand and therefore the equilibrium. (4) Fiscal policy has a multiplier effect so that a \$1 increase in government spending can potentially change aggregate demand by more than \$1. In my principles text, each of these ideas is discussed using the model of aggregate demand and aggregate supply. Students get the same intellectual payoff at a smaller pedagogical cost.

When authors get feedback from editors and prepublication reviewers on their textbook drafts, they have to decide how much to stick to their guns and how much to defer to the advice of others. There is no easy answer here. On the one hand, always deferring to advice is a surefire recipe for producing a book that looks just like every other one out there. It kills creativity and pedagogical improvement. On the other hand, sticking to your guns, if taken to an extreme, is sheer pig-headedness. It precludes you from learning. Finding the right balance is perhaps the most important judgment call that an author makes.

In the case of the Keynesian cross, I stuck to my guns and am glad I did, but other times I made what, with the benefit of hindsight, I now consider mistakes. In the first edition of my

principles text, against advice, I used the phrase “charlatans and cranks” to describe some advisers to Ronald Reagan, who told him that broad-based income tax cuts would have such large supply-side effects that the tax cuts would increase tax revenue. I did not find the claim credible and believed that most economists shared my view. The book made clear that the critique applied to a specific reason to favor the tax cuts and not necessarily to the policy of cutting taxes. (There is a big difference between rejecting a policy and rejecting one argument made by some proponents of the policy. A person can favor tax cuts for many reasons besides the belief that tax cuts are self-financing.) In the second edition of the text, I took out the phrase “charlatans and cranks” because an editor and some readers of the first edition told me—correctly, I then appreciated—that it was too inflammatory for a textbook description of a policy debate (though the substantive analysis of supply-side economics remained about the same). I wish I had listened to the editor the first time she offered the advice.

Instructors who teach introductory economics face the difficult but inevitable task of deciding what material to include and what to omit. The right answer will vary from instructor to instructor, depending on the length of term, the backgrounds of the students, and the instructor’s personal judgments about the importance of various topics. My book reflects my judgments, but I also try to write the book in a way that can accommodate a variety of instructor needs and perspectives.

In recent years, to increase the book’s flexibility, I have written several “modules” that instructors using the digital version of the book can assign at their discretion. These modules are additional short chapters on material not in the main book. So far they include one on the economics of health care, one on the European Union, and even one on the Keynesian cross. Most likely, I will write more of these modules in the coming years, so instructors will have a

library of material to choose from. But I encourage users of my book not to take too much from this library. Remember: The less you teach them, the more they learn.

Are textbooks too expensive?

A frequent complaint about textbooks is that they are too expensive. As I write this essay in 2018, the list price of the complete, hardcover version of my principles book is \$249.95. This is not unique: The main competitors of my book are similarly priced. And the phenomenon extends beyond the field of economics. The consumer price index for educational books and supplies, adjusted for overall inflation, has roughly tripled since 1980.

A common argument used to explain the high price of textbooks involves the principal-agent problem between student and instructor. The instructor chooses the book, often oblivious to its price. The student has little choice but to buy the book. As a result, the publisher has substantial market power and sets the price much above cost, resulting in exorbitant profits. Or so the argument goes.

I am skeptical of this story. To be sure, there is a principal-agent problem between instructor and student, and price is much above marginal cost. But there are large fixed costs associated with publishing, so I am not convinced that price is much above average total costs. Indeed, textbook publishing has not been all that profitable in recent years. The publisher of my principles book, Cengage, declared bankruptcy several years ago (and then reorganized under Chapter 11). From 2000 to 2018, while the overall stock market roughly doubled in price, the stock price of the large educational publisher Pearson fell by about two-thirds.

Nonetheless, the perception that publishers are unjustly enriching themselves at students' expense is common. In an editorial on April 25, 2008, *The New York Times* said that college

students are “rightly outraged” about textbook prices and called for reform, including legislation to regulate various industry practices. The paper’s reaction seemed odd to me. After all, the *Times* is a for-profit company in the business of providing information. If it really thought that some type of information (e.g., textbooks) was vastly overpriced, wouldn't the *Times* view this as a great business opportunity? Instead of merely editorializing, why not enter the market and offer a better product at a lower price? The *Times* knows how to hire writers, editors, printers, and so on. There are no barriers to entry in the textbook market, and the *Times* has the advantage of a good brand name. My guess is that the *Times* business managers would not view starting a new textbook publisher as an exceptionally profitable venture, undermining the premise of its editorial writers. In any event, since calling attention to the high prices, the *Times* has not entered the textbook market.

Moreover, while the principal-agent theory can potentially explain *high* textbook prices, it cannot explain *rising* prices. The same principal-agent problem that exists today existed in 1980, but textbook prices were, according to the consumer price index, lower then. Something else must be going on. I am not entirely sure of the answer, but let me suggest two hypotheses, both of which may have some validity.

One hypothesis is that textbook prices are not really rising as quickly as people think. When people complain about textbook prices, they often look at the price of the complete, hardcover version, as I did earlier. That is in part because they remember buying those when they were students. When I took introductory economics in the 1970s, I bought the complete, hardcover Samuelson text for macro and the complete, hardcover Lipsey and Steiner text for micro (and in each course used only half the book).

But that is not what students are buying today. Because I have the data on my principles book, let me use it as an example; the numbers are likely similar for other titles. For the most recent edition, the complete, hardcover book accounts for only about 5 percent of units sold and only about 8 percent of revenue generated. The vast majority of students are accessing the book in some other form (not even counting those who use pirated copies). Some buy the paperback “splits” that include only the subset of chapters they need for their courses. Some buy loose-leaf versions of the book, which they can put into three-ring binders. Some buy the digital book, which they can read on their computer, tablet, or smartphone. All of these options are less expensive than the complete, hardcover book.

In other words, as publishers have introduced newer, less expensive products over time, consumers have substituted away from older, more expensive products. This phenomenon reflects precisely what we teach our students when we explain that the consumer price index overstates inflation. My sense is that the upward bias from product substitution is particularly large when people measure the rising cost of educational books and supplies.

How much do students in fact pay? Again, I have good data only on my own principles text, but it is likely representative. In 2018, the average wholesale price received by the publisher for all versions of my book was \$102. With a typical bookstore markup, students paid about \$130. That price is about half of what one would infer by looking at the list price of the complete, hardcover edition. (And it is, incidentally, close to what a student at Harvard pays in tuition for each hour of lecture. By comparison, the book seems like a bargain.)

But this explanation is probably not the whole story. My second hypothesis to explain the high cost of textbooks involves the nature of educational publishing in a digital age. That is my next topic.

How is digital technology changing the market for textbooks?

Over the past several decades, information technology has transformed many industries, especially those in the business of providing information. *The New York Times*, for example, already has more digital subscribers than print subscribers, and its CEO has predicted that, ten years from now, the print version will be gone. The same forces are affecting educational publishing. Increasingly, producing physical books is considered less central to the business model than producing online content. And it may well be true that, at some point in the future, textbooks will be available only in digital format.

For students, the digital book offers a better price. My complete principles book, whose hardback version lists for \$249.95, can be accessed in digital format for about \$130. For that, a student gets an online subscription for the duration of the course. Moreover, in some ways, the digital book is more convenient, because the student does not have to lug a heavy book around campus. As long as the student has some electronic device and an internet connection, which is now almost always true, the book is readily available. On the other hand, the subscription disappears at the end of the course, so the student does not have the book to resell or use as a reference or door stop.

More important, digital books offer a range of pedagogical possibilities that are impossible in a traditional book. For example, graphs can be animated. The student can see them built up, step by step, much as an instructor would do in a classroom. In addition, videos can be embedded to introduce topics, explain difficult concepts, or show problems worked out. For students who are more auditory learners (such as the 20 percent of the population with a language-based learning disability like dyslexia), these videos provide a quicker route to

understanding. And with text-to-speech capabilities, the computer can read the book out loud, which may also be useful for these students.

In addition, digital textbooks often come with online homework systems. Instructors can assign problem sets, which students answer online. The problems sets are automatically and quickly graded, with the scores recorded. The system provides feedback to the student faster than a traditional human teaching assistant can. And an instructor in a large course can assign regular homework without the college having to incur the cost of hiring many teaching assistants to grade the work.

Yet these features of digital textbooks are not free to produce. Their cost is included in the price of the textbook, so it is ultimately paid by students. These extra features complicate the price comparison between a modern text and its counterpart a generation ago. Again, this is a familiar problem with the consumer price index. Unmeasured quality improvement causes the index to overstate inflation.

Is the pedagogical value of these features worth the extra cost? I believe so. But these innovations will surely be subject to study by educational researchers in the years to come. The unsuccessful ones will be dropped, and the better ones will survive. The textbook of the future could look very different from the textbook of today.

As textbooks become increasingly digital, a natural question is whether they should be open source. A digital textbook is an example of a club good: It is excludable, but it is not rival in consumption. The marginal cost of one more person accessing a digital book is close to zero. Why not put it in the public domain so everyone can enjoy the benefits?

The problem is paying for the fixed costs. The same issue arises for many products in the digital economy: software, movies, video games, newspapers, and magazines. Once produced,

these goods are not rival in consumption, but their production nonetheless requires substantial upfront resources. In all these cases, the solution is either to charge for access or to generate revenue through advertising. For digital textbooks, the market has settled on the first option. (Could a good digital textbook succeed if it were free to students but had ads? I am not sure. I am also not sure, given that ads are distracting, whether such a product would be a step in the right direction.)

One possibility is to have the fixed costs of production paid by a foundation grant (I am looking at you, Bill Gates) and then make the digital book freely available. This is similar to the common suggestion that newspapers like *The New York Times* should move from for-profit to non-profit status and then be supported by charitable donors, much like National Public Radio. Yet I am skeptical that this reform would improve on the status quo of the textbook market. After all, the current for-profit educational publishers are not that profitable, and there is no reason to think that a non-profit entity would find cost savings that have eluded existing publishers. I am afraid that the only way to substantially cut costs would be to reduce quality, which would not be in the students' interests. One can imagine shifting the cost of textbooks from students to charitable donors (as we already do with much educational spending). But that can be done by simply helping students pay for their books, rather than getting involved in trying to improve the production and distribution of those books.

Who should become a textbook author?

If you are thinking about writing a textbook, the most important question to ask yourself is: Do you enjoy the process of writing and revising (and revising and revising and...)? Not just tolerate it, but really enjoy it? I have known economists who are great teachers and great

researchers, but they do not necessarily enjoy the process of turning their pedagogy and discoveries into words on a page. Engaging students face to face or devising a clever model or a compelling empirical strategy is very different from stringing words together in a way that engages the reader and explains clearly the subject matter at hand.

Here is a good test: Do you like reading books about writing? Books like *The Element of Style* by William Strunk and E. B. White, or *On Writing Well* by William Zinsser? I love reading these kinds of books, because they give you the tools and strategies for becoming a better writer. But if you don't, then you are probably not really interested in the craft of writing, and you may well not have the inclination or skills to write a successful textbook.

Here is another test: Do you like tinkering with something you have written to make it clearer or briefer? I actually enjoy the triennial revisions of my textbooks, not only because they allow me to update my texts for the ever-changing world but also because they give me the chance to go through the manuscript and tinker some more. I can change "the curve is upward sloping" to "the curve slopes upward," saving one word and two syllables! If that edit does not strike you as a life-affirming victory, you are not a writer at heart.

You might think that a lack of interest in writing can be remedied by having a good editor. (And I have been lucky enough to work with a great one, Jane Tufts, throughout most of my career as a textbook writer.) But there is only so much an editor can do. A good editor can turn a B manuscript into a B+ or an A- book, but she can't turn a C manuscript into an A book. Moreover, in many cases, editors are best at identifying what does not work in a manuscript and then leave it up to the author to figure out how to fix it. The relationship between author and editor works well only if the author has a true interest in writing.

The second prerequisite for being a successful textbook author is the willingness to devote the time it takes to write and revise the book. I have seen great economists who, having been successful in the profession, think they can quickly write down what they know and expect the result to be a good textbook. But in fact, your first draft will likely be highly flawed as a pedagogical tool and will become better only with many revisions. Over the years, I have had not only a great editor work with me on the revisions but also many generations of Harvard students, who read through my manuscripts to identify passages that are confusing, ambiguous, redundant, or verbose.

Devoting sufficient time to writing is important in large part because the textbook market, especially for the introductory course, is like a tournament. You are competing with many other authors, and the rewards, both personal and pecuniary, are distributed highly unevenly. In other words, the benefits of writing the 10th best textbook for a course are small. If you are not ready to put a lot of time into writing and revising, there is little to be gained from entering the tournament. Poets may write for the sheer joy of creation and be fulfilled even if no one ever reads their poems, but textbook authors are satisfied only if their work is read by generations of students.

The tournament-like nature of the textbook market has made me wary of writing my texts with coauthors. For my research papers, writing with coauthors is the norm. Not only are there substantial synergies from comparative advantage and specialization, but I also learn a lot working with others who bring different skills and perspectives. By contrast, writing a textbook is a simpler activity: The product is judged mainly by the quality of exposition. The greater the number of authors, the more likely the product will, by the law of large numbers, be of average quality. But an average competitor is unlikely to win the tournament. Perhaps this logic explains

why the three best-selling economics textbooks during my lifetime—Paul Samuelson’s, Campbell McConnell’s, and my principles book—were initially written by single authors rather than teams. Committees are useful in many endeavors, but producing great writing is not one of them.

The third prerequisite for becoming a successful textbook author is the willingness to focus on your audience. I recall a conversation I had with an author who had written a text that seemed too mathematical for most undergraduates but not sufficiently rigorous for graduate students. When I asked him for whom he had written the book, he sheepishly admitted, “I really wrote it for myself. This is the mathematical level at which I like thinking about things.” I had some sympathy with his response, because that is also the level at which I like thinking about things. But the answer was not good enough. Not surprisingly, his book, while admirable in many ways, was not a commercial success.

Keeping the audience in mind is a challenge for many economists. Much of our writing, especially the writing we do during our graduate training, is aimed at other economists. With this background, it is easy for us to slip into jargon or to rely on implicit assumptions that are natural for economists but not for laypeople. This problem is sometimes called the “curse of knowledge.” Once you know a subject matter well, it is hard to imagine what it is like to not know it. Overcoming this curse is crucial to being a good teacher or textbook writer.

The person I always kept in mind while writing my principles text was my mother. She is not a college graduate but has always been interested in following the news and financial markets. As I wrote, I regularly asked myself, “How would Mom react to this passage? Would she understand it? Would she find it engaging?”

So those are the three conditions to judge whether you may be a candidate for writing a successful textbook: a deep interest in the craft of writing, a willingness to spend the time needed to write a great book, and an understanding of your audience. If, after sufficient soul searching, you think you fit the bill, then get to work. I wish you luck.

But not too much luck, because I will be out there competing with you.