Thoughts on Teaching

I spent about forty years teaching, first mostly in economics departments and a business school, then in law schools, and came up with a number of ideas about how different parts of the enterprise could be done better. Here are some of them.

An Application of Economics to Giving Exams

After answering the questions you know the answers to you still have some time left for questions you do not know the answers to. Your hope is that something you write will fool the professor grading the exam into thinking you know at least part of the answer, expressed it unclearly, deserve partial credit. Doing that costs you time writing and me time reading and lowers the quality of the information the exam produces, since there is a risk that I will either be fooled into giving you credit you do not deserve or interpret some other student's poorly written answer as entirely bogus when it is not.

My solution to this problem is an economic solution, although I like to claim that it was inspired by the story of Socrates and the Delphic Oracle.

> Well, one day [Chaerephon] went to Delphi, and there he had the impudence to put this question -- do not jeer, gentlemen, at what I am going to say -- he asked, "Is anyone wiser than Socrates?" And the Pythian priestess answered, "No one." Well, I was fully aware that I knew absolutely nothing. So what could the god mean? for gods cannot tell lies. For some time I was frankly puzzled to get at his meaning; but at last I embarked on my quest. I went to a man with a high reputation for wisdom — I would rather not mention his name; he was one of the politicians — and after some talk together it began to dawn on me that, wise as everyone thought him and wise as he thought himself, he was not really wise at all. I tried to point this out to him, but then he turned nasty, and so did others who were listening; so I went away, but with this reflection that anyhow I was wiser than this man; for, though in all probability neither of us knows anything, he thought he did when he did not, whereas I neither knew anything nor imagined I did."

On my exams, knowing what you do not know is worth something — twenty percent. That is the fraction of the points for a question that you get for not doing it. If you suspect that the best bogus answer you can come up with will be worth less than twenty percent, you are better off leaving the question blank and going home early, saving both of us time and me some of the hassle of trying to figure out which answers are or are not entirely bogus.

The Case for Short Exams

Another policy that I adopted many years ago was to try to make my exams short enough so that most students could finish before their time ran out. My original reason was my dissatisfaction with the practice of giving students who could persuade the relevant university officials that they had an invisible handicap, some sort of learning disability, extra time on exams. While some may have suffered from a real problem, I suspected that in many cases all that was special about their situation was having parents willing to pay a professional to produce the needed diagnosis. I did not like being a party to what I regarded as legalized cheating, had no way of preventing it, did have a way of making it ineffective. If everyone can finish the exam before time runs out, having an extra hour is no longer an advantage.

That was my original reason for trying to write sufficiently short exams. After I had been doing it for a while, I concluded that it was a good idea on its own merits. Being able to do things fast is sometimes useful, but in most contexts getting the right answer is more important than getting it quickly. An exam that most students find hard to complete rewards speed by more than I think it should be rewarded.

Another policy I followed was not to write the exam until after the last class. That way I did not have to worry, when answering questions in the final review class, that I might be giving away the answer to an exam question and so rewarding memory, not understanding.

The Use of Old Exams

University libraries often keep a file of old exams and make them available to students. As best I can tell, there are two reasons they do so. One is to help students study for exams they are going to take. The other is to prevent students who have access to old exams from other sources, a friend who took the course the year before or a fraternity that keeps a file of old exams, from having an advantage over students who lack such access. A third possibility is that the practice is part of an elaborate charade: Let the students memorize the answers to last year's exam, give more or less the same exam this year, and everyone can pretend that they have learned something when most of them haven't.

My own practice is to cut out the middleman by webbing some of my old exams and linking to them on the class web page. My usual suggestion is that, after studying for the midterm or final, a student should take one of the webbed exams and use my answers, if I have put them up, to check his. If they are not there, he can at least try to check his answers against the book or his lecture notes.

Should I web answers as well as questions? One of the problems in teaching economics is that, because it deals with features of the world that students are familiar with and uses ordinary language, often with specialized meanings, a student may go through a course thinking he understands everything but the fine points and end up having learned almost nothing. Providing answers makes it easier for a student to tell whether he actually understands the subject by how well his answer fits mine.

What I do not want the students to do — and am concerned that many may attempt — is to memorize the answers to the questions on past exams in the belief that at least some of them will appear on this year's exam; memorizing an answer does not require you to understand it. That approach goes along with the practice of going through a textbook using a highlighter to mark the five percent of it that you need to learn because you expect to be tested on it, a practice that makes me wonder what the student thinks the rest of the book is there for, other than as an excuse to sell it to him at a high price.

I tell my students how I want them to use the old exams but they may suspect, reasonably enough, that my objectives are not identical to theirs, that advice it is in my interest to give them may not be advice it is in their interest to follow. I therefor also warn the students that I try to avoid putting questions from the webbed exams on the current one, which may be more effective, providing they are paying attention, believe me, and remember.

At some level, my response to all such issues is that it is my job to make it possible for students to learn, theirs to make it happen. If a student chooses to ignore my advice and devote his efforts to memorizing answers in order to get a good grade on the exam, rather than learning ideas in order to understand what the course teaches, that is on him, not me. Along similar lines, I make no attempt to enforce compulsory attendance.

But I would still prefer, so far as I can manage, to teach the course in a way that makes it more likely that students end up understanding the ideas it covers.

Looking around the internet to see what other professors did, I found a wide range of policies. Many universities provide copies of past exams, occasionally but not usually with answers, in some cases only making them available to their students; why they go to the trouble of keeping other people from seeing them I don't know. There are web pages that link to exams from many universities, possibly contributed by students. One university, Texas Christian University, suspended 12 students for "allegedly using the Quizlet app to cheat on their exams."[1] As best I could tell from the news story, the professor said that their offense was failing to tell him that they recognized questions on the exam as ones they had seen and studied on the app.

I can see two different arguments for trying to prevent students from studying past exams. One is paternalistic — that students will try to memorize answers to old exams and it won't work, leaving them with neither a good grade nor understanding of the material of the course. The other is that it will work, that the questions on this year's exam will be sufficiently close to those on past exams so that memorizing the answers to old exams will be an adequate substitute for learning the course. That implies that everything students can be tested on can be tested with a small enough number of questions so a student can memorize the answers to all of them. That suggests that either there is not much content to the course or the professor is too lazy to write a different exam each year.

One commenter responded to my blog post on the subject with:

> The best way to solve your exam problems, as well as many others, is to grant everyone a PhD, MD and JD at birth. Then classrooms would stop being polluted by students who didn't want to learn and those whose chief qualification is rich parents.
>
> Just think how happy professors would be, teaching, like Socrates, only those who were desperate to improve themselves.

Which brings me to …

The Classes I Most Like Teaching

I have twice taught, as part of my university's adult education program, an abbreviated version of my law school seminar on legal systems very different from ours — ten hours, divided into four or five lectures. Interested readers can find recordings of most of the more recent lectures and part of the earlier ones on the class web page. It was great fun.

There were two important differences between that class and most I have taught. The first was that nobody was there who was not interested in what I was teaching, since the class did not meet any

[1] The suspensions were eventually overturned, but it is not clear whether other academic punishments were cancelled.

requirement for getting a degree. The second was that I did not have to grade the students. It thus eliminated the two least attractive features of conventional teaching.

Each year I teach ten or fifteen shorter classes of the same sort, under rather different circumstances. The setting is the Pennsic War, an annual two week long historical recreation event held in a private campground in Pennsylvania. Attendance at the event is upwards of ten thousand people. My classes, most of them about an hour long, deal with medieval historical recreation — how to cook from a period recipe, make hardened leather armor, tell a period story in a way that creates the illusion of a medieval story teller entertaining a medieval audience. They are part of a Pennsic University that offers thousands of classes each year, all taught by volunteer teachers to volunteer students. Nobody is there to get a degree and nobody has to give any grades. Again, more fun than my usual teaching. Which suggests that perhaps there is something wrong with the alternative model employed for most teaching from kindergarten through college.

One comment on that blog post:

> I taught ESL, which was very satisfying, since the students were always motivated; while they got no credit, they desperately needed English for jobs and schooling.
>
> On the other hand, there's nothing worse than trying to teach Baby Physics or Baby Math to pre-law and pre-med students who take baby courses because they have to maintain their GPAs.
>
> Unfortunately, someday you may be forced to hire one of them as a lawyer or be treated by one as a doctor.

Read then Listen or Listen then Read?

Most of the reading for the course I taught for many years on law and economics consisted of a book I wrote, based on my lecture notes from previous iterations of the course. As in most courses, students were supposed to read each chapter before the first of the classes that discussed it.

I reach the point in the discussion at which I pose a puzzle to see if the students can work out the solution. It occurs to me that any student who has done the required reading already knows the answer, because it is in the book. I make a particularly telling point, summarize an argument with a punch line that I least believe to be witty, and if a student reacts that is evidence that he has not read the assigned chapter, since it contains the same punch line. Large parts of the dramatic effect of the class only work for students who have not done the reading I assigned them.

I could use someone else's book — twenty some years earlier I did — but if there were another book on the subject I was happy with I would not have had to write mine. Alternatively I could use my book and create an entirely new explanation of the ideas for class. I doubt I could do it, and if I could I wouldn't; I would rather spend my time understanding and explaining some new set of ideas.

An easier alternative would be to reverse the order, assign each chapter to be read after the relevant class instead of before. That way the class introduces the ideas, the reading fills in details, reinforces what was discussed in class, gives the student a second chance to make sense of things he did not understand the first time through.

And I can deliver at least some of my punch lines to students who haven’t already seen them.

Another, perhaps better, possibility was suggested by a student; he said that he usually read a chapter after the first class in which the material was discussed but before the second. That way the material was fresh when he first heard it in class and he could use the second class to raise any questions that the reading had left him with.

One commenter on my blog raised the obvious question: Why, if the material I was covering was in the book, did I need to repeat it in my lecture? That question had occurred to me long before. After I wrote *Price Theory*, my first textbook, I tried the experiment of telling the students to read a chapter, then come into class and tell me what they didn't understand or wanted to discuss; if nobody had anything to discuss we could go home. The students did not like it, so I reverted to the more usual approach.

The puzzle of why they did not like it is related to my old puzzle of why the mass lecture, where there are too many students for significant interaction with the lecturer, did not vanish after the invention of the printing press. One possible answer is that some people learn better by listening than by reading. If so, the sensible arrangement would seem to be for some students to read the book and skip the lectures, some to attend the lectures and skip the book — a conclusion that I gather at least some students have worked out for themselves.

At Oxford and Cambridge, teaching is mostly done by a tutorial system, a weekly meeting with a tutor and one to three students. There are lectures but, according to some Oxford students I discussed the question with, most students don't attend them.[2] That sounds a good deal closer to my ideal than the usual U.S. system.

Curving Grades: Gordon Tullock's Solution

Ideally, a grade in a course one professor teaches represents the same level of accomplishment as the same grade in a course someone else teaches. Often it doesn't, since different professors have different standards. The usual solution to that problem is to impose a grade curve on the professors, a fixed distribution of grades, sometimes with some flexibility in the details. That makes it hard to get an A in a class that only the best students take, easy in one populated mostly by the worst students. Not only is that unfair, it gives students an incentive to take only the easy classes.

My friend and ex-colleague Gordon Tullock, one of the most ingenious people I have had the pleasure of knowing, had an elegant solution to this problem. You start with some measure of student ability — in a law school, you could use grades in the first-year classes that everyone has to take. At the beginning of the second year, base the curve for each class on the distribution of first year grades of the students taking it. In each semester thereafter, recalculate student ability, adding in the information provided by the previous semester's classes, and repeat.

Most undergraduate colleges do not have a first year where everyone takes mostly the same classes, so need some other way of starting the process, perhaps SAT score or high school GPA. Alternatively they could start everyone even, giving all the first quarter classes the same curve, and use the grades students get thereafter to gradually improve the measure of student ability. Mechanically speaking it should be easy enough to do, in a world of computers and spreadsheet programs, but I do not know of any school that has done it.

[2] The discussion occurred after a talk I gave on education to the Karl Popper Society at Oxford. The talk and discussion are webbed.

Adam Smith on Laptops in the Classroom

Professors have a mixed view of student laptops in the classroom. They are useful tools for taking notes and, connected to the internet, can be used to quickly research things relevant to classroom discussion. But they can also be used to exchange email or instant messages, play games, do any of a wide variety of things unrelated to and distracting from what is supposed to be going on in the classroom.

Some years ago I had the pleasure of sitting in on a class taught by a colleague who did a brilliant job of keeping his students' interest and attention — and mine — while covering material usually considered less than entrancing. I was at the back of the classroom and so could see quite a lot of laptop screens. With one or two brief exceptions there was no color on them, which I took as evidence that they were being used to take notes, not to browse the web or play games.

Whether to permit students to use laptops connected to the web during class is a new variant on the old question of whether class attendance should be compulsory, since the net lets one be physically in one place, virtually in another, attend class while corresponding with your friends or reading the newspaper. There are, of course, other ways of doing that — some of us remember reading concealed books during boring high school classes, or retreating into thoughts unrelated to what we were supposed to be learning — but the new technology provides a more convenient tool for the purpose.

On the subject of compulsory attendance, I cannot improve on the words of Adam Smith:

> No discipline is ever requisite to force attendance upon lectures which are really worth the attending, as is well known wherever any such lectures are given. Force and restraint may, no doubt, be in some degree requisite in order to oblige children, or very young boys, to attend to those parts of education which it is thought necessary for them to acquire during that early period of life; but after twelve or thirteen years of age, provided the master does his duty, force or restraint can scarce ever be necessary to carry on any part of education.
>
> (Wealth of Nations Book V Chapter 1 Part 3 Article II)

As demonstrated by my colleague.

How not to Teach Math — or Economics

A conversation I once had with my younger son, frustrated over his undergraduate math course, reminded me of my long standing objection to how math, and for that matter economics, are often taught. Theorems are proved with a rigor that is more than the students really need — especially in economics, where rigorous proofs can be applied to the real world only by combining them with non-rigorous models. The rigor is not only more than the student needs, it is more than any save the ablest students can understand. It is one thing to follow a proof step by step. It is a different and much more difficult thing to hold the proof in your head and understand why it is right.

My usual example of the problem is the failure to teach students of calculus why the fundamental theorem, that integrating and taking a derivative are inverse operations, is true. It is possible to give a non-rigorous but intuitively persuasive proof of the theorem in about five minutes, one that any student who understands what the two operations are can follow and has a good chance of

remembering. One of the commenters on my blog did it in 38 words,[3] but it is easier with a drawing.

As best I can tell, very few of the students who take calculus, even at a good school, are ever shown the proof; I would be surprised if more than one in fifty, a year after taking the course, could reproduce the more rigorous proof that they were, presumably, taught. To check the former impression, I asked my wife for her experience. Her response was that she was taught calculus twice, the first time at a good suburban high school but by an incompetent teacher, the second time at a top liberal arts college. To the best of her memory, she was never shown the proof. I got more recent evidence interviewing high school seniors who had applied to Harvard, something I have done several times as an alumni volunteer. They are all bright students, many have taken AP Calculus and gotten a top grade. I ask them if they can show me why integrating and taking a derivative are inverse operations, why the derivative of the integral of a function is the function. I do not think any yet has been able to.

It is common at good schools to complain against cookbook mathematics, memorizing the sequence of steps to solve a problem without ever understanding why it works. It is, I think, an almost equally serious mistake to present a branch of mathematics in the form in which professional mathematicians structure it after all of the original work in that particular field is done. Not only is it a form in which almost no student not qualified to become a professional mathematician can understand it, it is a form that gives a highly misleading picture of how mathematics, or other forms of theory, are actually done.

I am not a mathematician but I am an economist and know by direct observation how the original parts of my work were done. The process did not start with a step by step proof but with an intuition of how some set of ideas fit together, what characteristics the solution to a problem ought to have. Only after I had groped my way to what was (hopefully) the right answer did I, or someone else, go back and make the argument rigorous. About forty years after my first book was published, I wrote a third edition. Part of what it consisted of was filling in the blanks, working out in more depth and more detail ideas whose essence I understood then and still believe, in most cases, were correct.

Alfred Marshall, arguably the figure most responsible for the creation of neo-classical economics, commented in a letter on the relation between mathematical and verbal arguments in his field:

> But I know I had a growing feeling in the later years of my work at the subject that a good mathematical theorem dealing with economic hypotheses was very unlikely to be good economics: and I went more and more on the rules---(1) Use mathematics as a short-hand language, rather than as an engine of inquiry. (2) Keep to them till you have done. (3) Translate into English. (4) Then illustrate by examples that are important in real life. (5) Burn the mathematics. (6) If you can't succeed in 4, burn 3. This last I did often.[4]

There is much to be said for that policy. Mathematics is a more precise language than English, but also a language farther from the intuition of almost all of us. If you have the math and cannot translate it or think of anything in the real world it corresponds to, it is quite likely that the reason is you do not understand it. I sometimes referee journal articles. Occasionally I get one where, if

[3] "One can start with Rieman's definition of the derivative, draw little rectangles under the curve, and show that derivating the cumulative integral is equivalent in the limit to looking at the height of the curve on that point."

[4] Which left me wondering how much of the economics of the next century went into Marshall's fireplace.

you translate the math into words, it makes no sense, is arguably insane. The author or authors presumably had doctorates in the field. But they were manipulating symbols, not ideas.

When my daughter transferred to Chicago after two years at Oberlin, she was seriously considering majoring in economics. After taking an economics course, she decided on her alternative major. The reason was not that she does not like economics; she audited several of my courses while a home schooled student of high school age, one of my articles contains an idea that I credit in a footnote to her, and since graduating college she has edited several of my books. The reason was that the course was mostly about mathematics not economics.

I discussed her experience with a professor at that university of whom I have a high opinion, someone on the short list of people who, when they disagree with me, cause me to seriously consider that I may be making a mistake. He also had a daughter taking economics at the same school. He agreed with my daughter's judgement, that the courses were teaching mathematical rigor instead of economic intuition.

What matters is not remembering but understanding. If you have memorized a proof but cannot explain why the result is true, you have been wasting your time.

WoW Economics

You are teaching economics at a large university and looking for a new way of getting the interest of your students. It occurs to you that a substantial fraction probably play or have played World of Warcraft. It also occurs to you, since you too play WoW, that the game contains a complex economy that poses all sorts of interesting questions for an economist. You announce a new course — WoW economics — and get a gratifying large enrollment. Now what?

WoW has markets and prices, including an auction house with many buyers, many sellers, and a wide range of products for sale. Prices are readily observed: starting prices, buyout prices, relative prices at one time, changes over time. Actual sales prices are a bit harder, but students who are buying and selling things in the auction house could probably be persuaded to keep track of prices paid and received and make the information available to the rest of the class.

Consider one simple puzzle — relative prices of ore. Low level players mine copper and tin, higher level iron, mithril, thorium, … . What determines relative prices? One's first guess might be that prices are higher the higher the level of the ore, since a high level player can mine low level ores but not the other way around, making the potential supply higher at the lower level. It sometimes turns out that way, but not always.

One reason is that ore is a joint product. A lot of ore, I suspect a majority, is produced by players who are wandering around killing monsters and doing quests, happen to see a vein, and mine it. As the population of a server accumulates more and more high level characters, more and more time is being spent wandering in high level areas with high level ore, producing an increase in the supply of thorium, a decrease in the supply of copper.

High level characters need high level gear made from high level metal and higher level characters have more gold, so one might expect the demand to be higher for higher level metal. What complicates that is the process of skilling up — becoming a better and better smith (or jeweler or leatherworker), able to make better and better stuff. Before you can make swords out of mithril you must first make them out of iron, before that copper (I simplify, as any WoW enthusiast will realize). If your high level character decides to take up smithing he must start at the bottom. Even

if he himself is a miner, he is not spending his time in the low level areas where copper is mined, so goes to the auction house instead, bidding up the price of copper.

Skilling up complicates the situation in another way as well. To go from making swords of copper to making them of bronze, you have to make a lot of copper swords, or daggers, or armor, or … . You cannot use all of it yourself, so have to sell most of what you make. That sometimes means turning fifty silvers worth of material into something you can only sell for forty, implicitly paying ten for the additional skill from making it. That works for low and medium level gear, made mostly by people who are still skilling up. But to make the highest level of gear a craftsman has to already be all the way up in skill, so gets no additional skill by making it, hence will only make things if he expects to sell them at a profit.

As these examples suggest, World of Warcraft provides multiple examples of the economics of joint production. Gems, in the base version of the game, are found when mining ore; each time you mine a lump of mithril ore you have some probability of finding an aquamarine An increase in the price of mithril due to increased demand would result in more being mined, increasing the supply of aquamarines and driving down their price, so a high price of mithril due to an increase in demand will be associated with a low price of aquamarine — but a high price of mithril due to a decrease in supply would have the opposite effect. In a later version of the game gems could also be produced by prospecting ore you had mined, at the cost of not being able to refine that ore into metal, so a high price for mithril bars produced by a change in either supply or demand would raise the cost of prospecting mithril ore hence the cost, and price, of aquamarine. The game is full of such relationships, providing lots of opportunities both to explore theoretical issues and to test the conclusions empirically.

Patterns of prices in the game change over time, sometimes in comprehensible ways. At one point a new option opened up in the game, a type of character that got to start not at level 1 but at level 55. Your brand new level 55 Death Knight fights like a level 55. But if he takes up smithing it will be at skill level 1. The effect on the relative prices of low and high end ore and metal is left as an exercise for the reader.

For another example of WoW economics, consider the opportunities for arbitrage, both across goods and across time, and the implications thereof.[5] Any miner can, at the cost of a little time, convert a lump of iron ore into a bar of iron, a bar of iron and a lump of coal into a bar of steel. The result should be a predictable relation between the market prices of lumps of ore and bars of metal or of iron, coal and steel. Students can look on the auction house and see if the pattern holds. Similarly, any predictable pattern of price changes over time — some things being more expensive on the weekend, say, when more players are online — should open opportunities for enterprising players to buy low, sell high, and make a profit. The result of players doing so should be to raise low prices, lower high prices, and eliminate the pattern. Does it happen? To what degree is arbitrage restricted by transaction costs, the five percent cut that the auction house takes of any transaction and the time cost of buying, selling, producing, watching prices?

[5] One commenter on my blog pointed me at an entertaining account by Christopher Marks, his son, of arbitrage and other trading strategies, mostly on Runescape, used to produce both in-game and real-world money. The title of his talk is "Everything I know about Economics I learned by playing video games." His final comment is a reference to an economist hired by Valve to analyze in-game economics. "So far he has most of it right but he's still a bit of a noob."

As several of these suggest, WoW not only provides interesting economics, it also provides an environment for learning to do empirical work, test theories, extract information from a noisy signal.

There is no antitrust law in WoW, which makes it a good place to observe collusive behavior by sellers. My wife, who spends much more time in the auction house buying and selling than I ever did, has observed both an attempt to corner a market and an attempt, at least partly successful, to form a cartel, one she was invited to join. Her refusal was met by a threat to drive her out of the market by underselling her. The organizer of the cartel had apparently not read McGee's classic article on the myth of predatory pricing.[6] It had not occurred to her that if she was selling, at an artificially low price, ten times as many gems as the interloper, she was also losing money ten times as fast. It took only a few days for her to discover the flaw in the strategy and abandon it.

In World of Warcraft, as in the outside world, a critical issue for cartelization is the number of producers. If making some high end item requires a very rare recipe and only three players have it, it may be practical for them to get together and agree on a common (high) pricing policy. With forty crafters that does not work as well.

The other critical issue is elasticity of supply. If many people can make the item but it requires some rare item as an input, a seller who can corner the market on that input, buy all units as soon as they show up, can have an effective monopoly of supply. That works if the input is a rare drop, something produced occasionally as a side effect of doing something else. It does not work if the input is something that is deliberately produced, since the high price for the input will give other players an incentive to produce more of it.

It might occur to a sufficiently clever WoW monopolist that if he can monopolize a key input he doesn't need to monopolize the product it is used to make. He can collect his monopoly profit in the price of the input, letting other players make the product and sell it at a price reflecting what they have had to pay him. He will then be well on his way to understanding a point that generations of anti-trust scholars got wrong — why using vertical integration to extend a monopoly is, while possible, generally pointless.[7]

These are a few examples I have come across of economics in the World of Warcraft. It should not be hard to come up with enough more to fill a quarter.

So far I have been discussing teaching at the college or graduate level, that being what I have spent most of my career doing, but the ideas are relevant to for K-12 schooling as well. Following out one of the comments on my blog, I found a web site describing the use of World of Warcraft to teach sixth graders, including lessons on math, writing, and much else. The project seems to have started a few years before I first posted on the subject and is still going.

WoW Statistics

If the economics course works, you might talk to a colleague in the statistics department about the possibilities for a similar course in his field. WoW provides a lot of opportunities to apply

---

[6] John S. McGee, "Predatory Price Cutting: The Standard Oil (NJ) Case," Journal of Law and Economics, vol. 2 (October 1958), p. 137. https://www.journals.uchicago.edu/doi/10.1086/4665471. McGee was testing Aaron Director's argument that the predatory pricing Rockefeller was accused of would not have worked and concluded that it was correct.

[7] For details of the argument, see my *Law's Order*, Chapter 16, p. 253.

statistical methods: testing the implications of WoW economics, estimating from data the probability of various uncertain outcomes, solving with data the vitally important question of what selection of gear, rotation of attacks, use of special abilities, will maximize the damage a character does or minimize the damage he takes.

For one example of using World of Warcraft to teach statistical concepts, consider the question of how to tell whether a process is actually random. Human beings have very sensitive pattern recognition software, so sensitive that it often sees patterns that are not there. There is a tradeoff, as any statistician knows, between type 1 and type 2 errors, between seeing something that isn't there and failing to see something that is. In the environment humans evolved in, there were good reasons to prefer the first sort of error to the second. Mistaking a tree branch for a lurking predator is a less costly mistake than misidentifying a lurking predator as a tree branch. One result is that gamblers routinely see patterns in random events: hot dice, a loose slot machine, a run of luck.

Players in *World of Warcraft* see such patterns too. The situation is made more interesting by the fact that purportedly random events might not be. It is usually safe to assume that the dice which you have used in the past will continue to produce the same results, about a 1/6 chance of each of the numbers 1-6, in the future. In a computer game it is always possible that the odds have changed, that the latest update increased the drop rate for the items you are questing from one in four to one in two. It is even possible that some mischievous programmer has introduced serial correlation into otherwise random events, that the dice really are sometimes hot and sometimes cold.

At one point I was on a quest which required me to acquire five copies of an item dropped by a particular sort of creature. Past experience suggested a drop rate of about one in four. I killed four creatures, got four drops, and began to wonder if something had changed.

The question was one to which statistics, specifically Bayesian statistics, was applicable, and it provides an opportunity to contrast that approach with the more familiar approach of classical statistics. Consider first the latter. The null hypothesis is that the drop rate has not changed, that each creature I kill has one chance in four of dropping what I want. The alternative hypothesis is that the latest update has raised the rate to one in one. A confidence result tells us how likely it is that, if the null hypothesis is true, the evidence for the alternative hypothesis will be at least as good as it is. Elementary probability theory tells us that, if the null hypothesis is correct, the chance of getting four drops out of four is only one in 256. Hence my experiment confirms the alternative hypothesis at better than the .01 level.

Does that mean that the odds that the drop rate has been raised to one in one are better than a hundred to one? That is how, in my experience, people commonly interpret such results, as when an IPCC report explained that "very high confidence represents at least a 9 out of 10 chance of being correct; high confidence represents about an 8 out of 10 chance of being correct."

It does not mean that. 1/256 is not the probability that the drop rate has changed, it is the probability that I would get four drops out of four if it had not changed. To get from there to the probability that it had, which is what would be relevant if I wanted know at what odds to bet someone that the fifth kill would give me my final drop, I need additional information. I need to know how likely it is that, just prior to my doing the experiment, the drop rate has been changed. That prior probability, plus the result of my experiment, plus Bayes Theorem, gives me the posterior probability that I want.

Suppose we determine by reading the patch notes of past patches or by getting a Blizzard programmer drunk and interrogating him, that any particular drop rate has a one in ten thousand chance of being changed in any particular patch. Further assume that my four drops out of four occurred the first time I played after a patch was installed. The probability of getting my result via a change in the drop rate is then .0001 (the probability of the change) times 1 (the probability of the result if the changed occurred — for simplicity I am assuming that if there was a change it raised the drop rate to 1). The probability of getting it by random chance without a change is .9999 (the probability that there was no change) x 1/256 (the probability of the result if there was no change). The second number is about forty times as large as the first, so the odds that the drop rate is still the same are about forty to one.

A student who spent his spare time playing *World of Warcraft* is more likely to find that explanation interesting than if I were making the same argument in the context of an imaginary series of coin tosses, as I usually do.[8]

Faking it

The son of friends of ours is required by his teacher to spend twenty minutes a day reading and report on doing so. The theory, presumably, is that since regular reading correlates with desirable outcomes, the way to get those outcomes is to compel children to read. Since someone who obeys the rules will be cutting the book into twenty minute chunks instead of reading right through it, the likely effect is to teach him that reading is a chore to be done only under compulsion.

Years earlier, we observed an analogous mistake in a different context. Our home schooled daughter, considering a career as a librarian, volunteered to work without pay at a large local library. After a week they thanked her and told her that her term of volunteering was over. Pretty clearly, their assumption was that she was volunteering because her high school required her to or to get something to claim on her college application, and it was now someone else's turn. Wanting to volunteer to do useful things is evidence of desirable personality traits. Volunteering because someone will reward you for it is not. She found a smaller library that actually had a use for her services and worked there for a couple of years.

For a third example, consider one of the problems I have observed in my experience of historical recreation in the Society for Creative Anachronism. Too often, people in the SCA are convinced that being historically accurate is something you do because other people are pressuring you to do it or in the hope of getting rewards and status. The result is documentation that consists not of trying to figure out how something was actually done in period but of trying to find some excuse for claiming that whatever you want to do qualifies.

In each case the mistake is the same, the attempt to create the effect without its proper cause. To fake it.

One commenter on my blog responded:

> Trying to get the cause by pursuing the effect is one of the characteristic behaviors of the villains in *Atlas Shrugged*, in domains ranging from economics to sex.

[8] A blog post where I did it that way.

Doing it Right: Scav at the University of Chicago

Every year, a sizable fraction of the student body at Chicago spends most of several days on a scavenger hunt claimed to be the largest such event in the world. It starts with a list of some two or three hundred things to be found, done, solved. Competition is by teams, often although not always associated with a dorm, typically a hundred to two hundred people — students, alumni, friends and allies — dividing up the problems of the hunt.

The lists of requirements for past hunts can be found online. For example, from 1999:

> **Item 240.** A breeder reactor built in a shed, and the boy scout badge to prove credit was given where boy scout credit was due. [500 points]

And a pair of physics students did it.

Most of the items are less difficult than that one. Examples:

22. Our plants have been listening to the Plantasia album for a few months now, but we're worried they're getting bored. Record 5 minutes of My Bulblet, My Bulblet and Me, an advice podcast by plants, for plants. Submit to mbmbam@uchicago.edu before Judgment. [6 points]

36. Get a basketball player to perform a baptism. Get dunked on, baby. [8 points for a collegiate basketball player, 21 points for an NBA/WBNA basketball player]

37. We were told that all RAs are given the option to live in Vue53 over the summer for the low price of $500. We were also told that this means that there's just a bunch of RAs living in a dorm alone together for the whole summer. We were disappointed to hear that no one has capitalized on this by making a Big Brother-esque reality show about these RAs living together. [8 points]

Some things must be done outside of Chicago, so a team will need a road trip group. For example:

43. At the Lew Wallace Study and Museum [In Crawforsville, IN], the Traveling Circus picks up two souvenirs for home: Flat Lew and a seed packet from the library outside. [8 points]

Some are puzzles that require research to solve:

30. According to a 1987 rent advertisement at the Givins Castle, the castle could be rented by a school or what other institution that required large grounds? [4 points]

No grades, no curriculum, no classes. Just people doing things for fun — many of which will turn out to be educational.