[SQUEAKING]

[RUSTLING]

[CLICKING]

SARA ELLISON: OK, let's go ahead and get started. So last time, we ended-- well, we were talking about functions of random variables, of course. And there we talked in particular, one of the standard examples that I did was the probability integral transformation. And I talked about how this could be useful for generating random draws. Like for instance, if you were interested in writing a program where you wanted to simulate random draws from some distribution, that using the reverse of the probability integral transformation could be useful. And there seemed to be some demand for an example.

So what I'm going to do I'm not going to do exactly that. But I'm going to do an example that is quite related. And then, we'll talk about the relation in a second. Oh, let's start with housekeeping notes. I don't think there's anything sort of controversial here.

Problem set 3 is due on Wednesday. Keep an eye out for problems set 4. Today, more on functions of random variables. We will start talking about moments of distributions today. And I'm going to finish off the lecture today with an extended example of auction theory, basically.

And then Wednesday, we're going to talk about properties of moments, conditional moments and so forth. So here is the example. And just sit tight. There will be a connection to the probability integral transformation example here.

So suppose we start with a random variable x that has a uniform distribution on the unit interval. And we have some function. We just were given this function or came up with this function that we want to use to transform x. So what is the distribution of this transformed random variable?

Well, we know how to do this. Remember, we can figure out first what the CDF of y is. And we do that in steps. We argue what it is in steps. And then once we have the CDF, then we take the derivative, and we get the PDF.

So let's do that. Oh, but first, note that the induced support of y is just the sort of non-negative real values. Yep?

STUDENT: Log 0 is undefined, right?

SARA ELLISON: Pardon me?

STUDENT: Log 0 is undefined?

SARA ELLISON: Yeah. So it doesn't actually matter too much, because the probability at any particular point is equal to 0. So we don't have to worry too much about what happens at any particular point, but yeah.

So the induced support is the following. And we'll use that in a second. So now, let's go through step by step, figure out what the CDF of y is.

So remember, we'll start with the definition of the CDF. And then we plug in this function that we have. So we plug in negative log of x over lambda into the probability statement. We solve for x, standard stuff. And then, we don't actually have to do the integral.

So most of the time at this point, what we do is we write down the PDF of x. And then, we integrate over the relevant area defined by this probability statement. But we're going to take a shortcut today. And that's because x has a uniform distribution.

And so the probability that x-- so x is uniform, 0, 1. The probability that x is greater than any number is just 1 minus that number. That's only because x has a uniform 0, 1 distribution. So when we're dealing with uniform random variables, we can often skip this step where we have to do the integral.

So here's what we have. That's the induced support. That's the CDF of the new random variable. And so we just take the derivative to get the PDF. And here's what we get.

Fine. Everything's standard. Does this distribution look familiar? So we've seen this least a couple times this semester, I think. It might be familiar from other contexts.

STUDENT: Exponential.

SARA ELLISON: Exponential, exactly. So now, recall that I said that if you have a distribution and you want, say you want the function that transforms a uniform 0, 1 to that distribution, what you do is you just take the inverse of the CDF. And you use that function to transform the random variable.

So now let's just go back, take a step back, take the inverse of the CDF that we found above. And let's see if it's the same function that we use to transform the random variable originally. So here's our CDF. And let's just find the inverse function.

So we're just writing x equals 1 minus e to the negative lambda y and solving for y. And we do that in a couple of steps. And then, we plug-in the capital X and Y, because this is now the function we're going to use to transform the random variable x. But that's not the same function that we used in the example. What's going on? Any thoughts about what's going on here?

STUDENT: X and 1 minus x is the same.

SARA ELLISON: Sorry x, and 1 minus x is?

STUDENT: Same.

SARA ELLISON: Same distribution, exactly. So when I said that you can take the inverse of the CDF and use that to transform a uniform 0, 1 random variable to get the distribution or to get of the PDF of the random variable that had that, whose CDF you're using. When I said that, I didn't say that was unique.

So in particular, there might be lots of different functions that are going to transform a uniform 0, 1 random variable into an exponential random variable. And these happen to be two of the functions. And in fact, if you look at the functions, you might be able to convince yourself that these are going to give you the same distribution, if you use these two different functions to transform a uniform 0, 1 random variable.

Why? Because 1 minus x, this random variable is also going to have a uniform 0, 1 distribution. I didn't prove it. But it's probably not very hard to believe that that's true. And you can prove it if you want.

So what's going on? Both of these functions work. In particular, the reason why both of them work is that, as I said, if x is uniform 0. 1, 1 minus x is as well.

So there is a lot of information in a PDF. A PDF, it's a whole function. It sort of tells you all kinds of things. And sometimes it tells you a lot of things that we don't actually care about.

Perhaps we don't care precisely what the shape of the distribution of our random variable is in every region of the reals or something like that. Maybe what we really want is we want some way to summarize. And I should say sometimes we do care about that, And. Sometimes we don't. But in the times that we don't, maybe what we want to do is just summarize some of the most salient features of a distribution, where it is centered, where it reaches its peak, how spread out it is, whether it's symmetric, how thick its tails are, et cetera.

And so what we do is we define something called a moment or a series of moments of a distribution to help us summarize some of these most salient features. And that's what we're going to talk about next. So here are some distributions. I just drew them. I don't know what they are.

And the one moment, or I guess three different moments that are somewhat related, that you've probably all run across in one context or another, are mean, median, and mode. And they all describe, in different ways, where the distribution is located or centered. So all of these moments are going to be useful in various ways.

So the mode is the point where the PDF reaches its highest value. And I've indicated it on these three distributions. The median is the point above and below which the integral of the PDF is equal to 1/2. So whatever point is the point at which the integral below is equal to 1/2 and above is equal to 1/2 is the median of the distribution. And then the mean, also known as the expectation or expected value of the distribution, is defined in this way. I'll give you a geometric interpretation of it in just a second.

But basically, the way we define the mean of the distribution with the expectation is just the integral over the support of the distribution of x times the PDF of x. So if we need to find the expectation of a continuous random variable, we just integrate the PDF times the value over the support. That's just what I had on the previous slide. And I should also point out, because this comes in handy quite a bit, the discrete analog is just what you would expect it to be.

Instead of integrating, we sum. But we sum up all over all of the possible values of x. So it's just summing up x times the probability of x at that point.

A couple comments about notation and terminology-- I will often use capital E of x to denote the mean of a distribution of the expectation of a distribution. But I'll often also use the Greek letter mu. And so, those are sometimes used interchangeably.

And this is the geometric distribution or geometric interpretation I promised you. So if we think of the PDF literally as a density, the expectation is the balancing point of that density. So this is different from the median of a distribution, because if you think about how I describe the median, the point at which the integral both above and below is equal to 1/2, if you have a median of a distribution and you shift probability on one side of it, the median doesn't change because the integral on that side is still equal to 1/2.

But the expectation might change. You take a chunk of probability and you shift it out, then the expectation of the distribution shifts out as well. And the terminology does get a little more complicating or complicated when we start talking about estimation, because we're going to talk about something called a sample mean, which is quite definitely different from this. It's related, but it's different.

And so then, we've got to keep mean and sample being kind of straight in our heads. But we'll be able to do it. Oh, here.

So here's just a quick example of calculating the expectation of a distribution. So let's suppose we have an exponential PDF. And how do we calculate its expectation? Well, it's just plugging into the formula. It's just going to be the integral over the support of the exponential, which is 0 to infinity of x times the PDF of x. If we didn't put x in here, of course, we know that integral is just equal to 1. But then we put the x in the formula, and that's the formula for calculating the expectation.

So specifically, in the case of the exponential, that's what it looks like. I'm skipping the calculus and just giving you the answer. And it turns out there's a good reason for me to skip the calculus this time. I didn't realize when I picked this example, you have to do integration by parts to get it.

So if you want a little practice on your calculus skills, go ahead and do it. If you don't, you can just take my word for it. This is the expectation of an exponential distribution.

So now, sort of having a few of these things under our belt, we're going to take a little side trip into auction theory. So why what do auctions have to do with probability? Well, there's sort of a number of reasons that I want to take this little side trip. One reason is that last lecture, we talked about order statistics. And it turns out that order statistics are a very useful mathematical tool to analyze auctions. Why is that?

An obvious thing to do when you're talking about auctions, is to model bids in an auction as an iid random sample. And the winning bid, which is typically the highest bid, as the nth order statistic from that random sample. So we will do that.

Another reason that I want to take this little side trip into auction theory is that most of this class, we're going to be sort of giving you tools that help you empirically analyze questions in social science. And here is an example of probability tools that help you theoretically analyze some questions in social science. So it's not really the main focus of the class. I mean, the main focus of the class is sort of data analysis and empirically analyzing questions in social science. But I did want to show you one example, where you could take tools from probability and help you do some theoretical analysis of questions in social science.

And furthermore, we also will compute some expectations along the way. So we'll get a little practice with that. So here are some things that are sold at auction.

Over here, there's a Picasso. So if you have a Picasso sitting around and you'd like to liquidate that and maybe use it to pay for your MIT tuition or something, what would you do? You probably wouldn't just put it up on Craigslist at a posted price. You'd probably go to Christie's or Sotheby's and auction it off.

We can talk about the decision between posted price and auction in a second. But start thinking about that question. Why would you decide to sell that at auction, as opposed to a posted price?

Down here are some livestock. Livestock in the US is typically sold at auction before going to the butcher shop. This is a picture of radio spectrum. OK so it used to be in the United States and pretty much all countries, radio spectrum was sort of controlled by the government and given away for free to radio stations and television stations and things like that.

And it was given away typically, with conditions like, you have to put on programming that's useful for the public good or something like that. But here, it was just given away. And then at some point, I think it was in the 80s or 90s maybe, I'm not sure, the US government decided, we're sitting on a gold mine. We have hugely valuable asset here.

And furthermore, we don't actually know. We know it's hugely valuable. We don't know exactly how valuable it is. Why don't we start auctioning it off?

And so, the United States started running spectrum auctions. And they still do for various parts of the spectrum. And now, this is sort of common practice among lots of countries.

Down here is a picture of an F-16. So who produces these, Lockheed Martin? I'm not sure.

So Lockheed Martin does not sell F-16s at auction, obviously. But in fact, they are acquired through an auction mechanism, kind of a reverse auction. So when the US government wants to buy the sixth generation fighter jets, what do they do? They write down the specifications that they want in the sixth generation fighter jets. And then they let all of the people, all of the firms who would like to produce those fighter jets bid.

So the bids can be very complicated. But they boil down to here's how we can fulfill your specifications, and here's what we'll charge. And then the US government buys the fighter jets from the lowest bidder. So that's called a procurement auction. And it's kind of a reverse auction. So instead of the highest bidder winning and purchasing, the lowest bidder wins and sells the product.

I went on eBay a couple of days ago when I was making up this slide. And I found some collectible beer tray tin signs. I guess they're just beer trays. I don't know.

Has anyone heard of Tech beer? Did you know Tech beer existed? No, I've never heard of Tech beer. I didn't know if it was like maybe the official MIT beer or something like that.

But if you're inclined, you can buy memorabilia like Tech beer trays on eBay. And I think the starting bid was only $8 there. So you can get them pretty cheaply.

Up on the very top, where it says Amica, do you know what that is? You've seen a zillion of them. So you they should look familiar.

It's an add-on, you don't necessarily know, but Google. Yeah, turns out it's an add on Google. So one thing which you may or may not know, is every time you type in a search term on Google, an auction is run. Every time anyone anywhere in the world types in a search term on Google, an auction is run.

And that auction determines which ads you are served. So I went on Google and I typed, auto insurance. And that was one of the ads I was served. Someone paid Google for the ability to serve me an ad when I typed in auto insurance.

Does anyone know what this is in the upper right corner?

STUDENT: Tulips.

SARA ELLISON: Tulip bulbs, right. So I have one slide just for fun on the tulip bulbs in just a second. So we'll get to those. But basically, in tulip bulbs not only were auctioned off in the 1600s in Amsterdam, but sort of a new auction mechanism was devised to handle the auctioning off of tulip bulbs. And I'll mention more about this in a second.

Does anyone know what this is?

STUDENT: Old map

SARA ELLISON: Old map, but actually it's what's on the map I'm trying to represent, not the map itself.

STUDENT: Roman Empire.

SARA ELLISON: The Roman Empire, exactly. So the Roman Empire was actually auctioned off, believe it or not. So let me go to the next slide. And I'll give you a couple details. So this is again, these next two slides are just for fun.

So the Praetorian Guard, which was sort of the elite military officers in the Roman Empire, found themselves without an emperor. And they decided to auction off the Roman Empire to the highest bidder. This is in 193 CE.

So Marcus Didius Severus Julianus was the highest bidder. And he ended up paying 25,000 sesterces per soldier for the Roman Empire. So he served as emperor for 66 days before being executed. I don't know precisely what happened, whether there was a follow-on auction after he was executed. If there was, I can bet that the bids were lower. That's my guess.

STUDENT: Did they not have democracy yet?

SARA ELLISON: Well, I don't know. I'm sort of not an expert on this time in history in the Roman Empire. That was a time when the Praetorian Guard had a lot of power. So they just decided to exploit that power and auction off them. Yes?

STUDENT: They killed them. They already had democracy. And then they got overthrown by some by Caesar with the army. But apparently, this sparked the Civil War.

SARA ELLISON: Oh, the execution sparked a Civil War. Oh, interesting, good. So the word auction actually comes from the Latin, "auctio," meaning increase. And that is a description or a partial description of one of the very common auction mechanisms, where the price increases until all but one bidder drops out.

And the highest bidder is in an auction in ancient Rome, was called the "emptor." That's the Latin term for it. And obviously, we are familiar with that term from the phrase, "Buyer beware." a caveat emptor.

Auctions were used very extensively in ancient Rome. They were used to sell plunder, household effects, slaves, wives, commodities, et cetera. And now to the tulip bulbs-- so as I said in 1600s, traders from the Ottoman Empire brought tulip bulbs back to Holland. They were unknown in Holland. So no one knew exactly what the demand was going to be.

And furthermore, supply was fixed in the short run. There was no way that they could generate more tulip bulbs in the short run because it takes 7 to 12 years to go from a tulip seed-- you know what existed, anyhow a tulip seed to a tradable bulb. So whatever the traders brought back from the Ottoman Empire, that was what they had.

So it turns out demand was really high. And what the traders did is they invented basically, an auction mechanism called the Dutch Auction, where you start at a high price. And you decrease the price until someone buys the tulip bulbs. And I don't remember off the top of my head what tulip bulbs ended up selling for. But it was sort of thousands of dollars in current dollars or something, I mean, huge amounts of money. You have a question?

STUDENT: So what if multiple people pick a bid? Do they start going up again?

SARA ELLISON: Oh, so you're asking, how are ties resolved? I don't know. Maybe if two people put up their hand at exactly the same time, then the auctioneer picks one of them at random.

I actually have a board game called, Merchants of Amsterdam. Has anyone ever played this? No. So it has this little thing, where you a timer kind of thing. It's to do the Dutch mechanism or the Dutch auction mechanism.

So you turn the timer. And then basically, everyone's watching the price of the commodity tick down. And then you're all sort of like, sitting around the timer. And whoever presses the timer first wins the auction.

At least in that case, someone has actually pressing it. And that's how the ties get resolved. And for people who are interested in financial innovation, the Tulip Panic, the Dutch Tulip Panic. So I should say prices went up very high, but there was also sort of a spectacular crash. And that was why it was called the tulip panic.

But for people interested in financial innovation, options and future contracts were also pioneered during this period to help deal with the interesting financial setting that was caused by these tulip bulbs. OK, yes?

STUDENT: How long does it take for that bulb to grow?

SARA ELLISON: That I don't know. Does anyone know? How long did the Dutch tulip bulb bubble last?

STUDENT: I think it lasted like a year, year and a half.

SARA ELLISON: Year and a half. OK, cool. OK, so those couple slides were mostly for fun. But let's go back to the question of selling a product at posted price versus auctioning.

So here now is a slide of things that are sold at posted price. You buy fruit at posted prices. You would never go into the Stop and Shop and sort of engage in some kind of an auction in the produce aisle or something like that. You just pay whatever price is posted. When I bought my iPhone, I just bought it at posted price. When I buy iced tea every day, I just go in. I buy whatever the price is posted.

Clothing is the same. The new Macklemore and Ryan Lewis CD, if you buy that, you will just buy it at a posted price. I heard it wasn't very good. That's what the reviews say. But I'm not an expert. Yeah.

STUDENT: This is the US, right? In other countries--

SARA ELLISON: So that's right. There are cultural differences between what is sold at posted prices and what is auctioned. That's absolutely right.

STUDENT: [INAUDIBLE]

SARA ELLISON: And sometimes you can negotiate a posted price. Which a negotiation is a little bit different than selling something at auction. But it has some of the elements of an auction, for sure.

STUDENT: I thought it was interesting. Buying vegetables from a vegetable seller in India is a reverse auction. You keep reducing the price until he refuses to sell at a lower price.

SARA ELLISON: So he has a posted price, right? Is that what you're saying? And then you say, no, I'll give you a half of that. And then he says, OK. And then you say, no. I'll give you a quarter of that.

STUDENT: Well, no-- well, you start with the lowest.

SARA ELLISON: Oh, OK. Fine. Yeah, yeah. So it's a standard bargaining. Yeah, exactly. Yeah, yeah. That's right. So that's a very important point that there are cultural differences. We'll get to that in a second, across cultures about what's auctioned and/or sort of haggled over and what's sold at a posted price. And then yeah, I guess the other point is that haggling isn't quite the same as an auction. But it has many of the same characteristics.

So in the US, houses are also sold at posted prices, which has always been a little puzzling for me. And it doesn't-- a house doesn't quite fit in to, in terms of its characteristics, it doesn't quite fit into what other things are sold at posted prices in the US. Now, of course, there often is a bargaining that goes on after in the process of buying a house.

STUDENT: Bidding wars, as well, if you have a certain price [INAUDIBLE]

SARA ELLISON: That's right. I tried to talk my parents into auctioning their house off. And it didn't-- They said you and all your crazy economics ideas, no way. OK, and actually, I don't know if you guys have been on eBay recently, if you buy things on eBay. It's true that most things on eBay today are sold at posted prices, not through an auction mechanism. And that's a big change. We'll talk about that, as well.

OK, so what determines whether a seller will decide to sell an item at a posted price or auction it off? So based on your own experience, based on the products that we saw here, what would you say are some of the factors that a seller might consider? Yes.

STUDENT: Supply.

SARA ELLISON: Pardon me.

STUDENT: Supply.

SARA ELLISON: Supply meaning-- can you elaborate a little bit?

STUDENT: Like if the thing is common, then it does make sense to auction it because you can buy it anywhere. But if the thing is really unique, and you want to auction it because you don't. And if it's hard to determine the value of the good, then you use the auction.

SARA ELLISON: Yeah, absolutely right. So if you have a product that's unique and/or hard to determine its value, then you're more likely to auction it off than sell it at a posted price. And I guess maybe this is sort of an obvious point, but the reason that that's the case is because when you auction off a product, then you're getting a lot of price information for free-- or sorry, a lot of value, a lot of information in the valuation of that product for free.

You know, you're getting not only-- I mean, depending on the mechanism-- you get not only the winning bid but you may get lots of information on the entire distribution of valuations for this product. You also hinted at another thing which was that there could be, relatedly, there could be an asymmetry of information between the buyer and the seller. There could be situations where there is a lot of information about the value of a product out there, but the seller doesn't have that information.

So go back to the case of the spectrum auction. When the US government is auctioning off a slice of the radio spectrum, they know that the people bidding on that slice of the spectrum have done huge amounts of research. And they sell products that are going to be using the spectrum. They have lots of engineers and experts predicting what they'll be able to sell their product for and what products they'll be able to invent that take advantage of that spectrum and so forth. And the US government doesn't have, isn't privy to any of that research. So there's really a big informational asymmetry there between the seller and the buyer.

Other, let me see what I have here. Well, actually that's-- OK. So we've been talking mostly about this sort of second point, information. The seller receives free information and the value of the good and, in fact, possibly information on the whole distribution. But how about transactions costs?

So it's sort of a little ridiculous to think about me going into the Stop and Shop and bargaining over every single thing that I buy in the Stop and Shop, or having to engage in some kind of an auction on every single thing that I'm buying in my trip to the Stop and Shop.

And again, you know, so there are cultural differences. There are places where if I go to the market, I might, in fact, haggle over everything. But there's big transactions costs associated with that. And so you would think that maybe small ticket items, the seller is not likely to lose that much. I'm not likely to lose that much. If I pay slightly higher price, it might not be worth it to either of us to engage in an auction. It might just be-- we might be able to minimize our transactions costs by just doing the transaction at a posted price.

STUDENT: I don't know if this information is currently available, but historically were posted prices coming quite much later than say, like, an auction type mechanism? So the barter system I'm sure [INAUDIBLE].

SARA ELLISON: Yeah. So you bring up a very interesting point. So the question was, in various markets do auctions kind of precede posted prices? And I think that's an interesting question. I can't say with great certainty-- I'm sort of not a historian of prices. So I can't say with great certainty if that's true. But there are reasons to believe why it could.

And this sort of second bullet point suggests one reason, which is that when a market is new, there's a lot of information asymmetry. And then once a market becomes more mature, the sellers know what they're going to be able to sell this product for. And so yeah, there are definitely reasons to believe that kind of evolution can happen. And actually, that could be a part of what's going on in eBay, which we'll talk about in a few minutes.

So what do these two, sort of keeping in mind this sort of information story and the transactions cost story, what are the goods that are likely to be auctioned instead of sold at posted prices? Unique goods are likely to be auctioned because people don't know, the sellers don't know what they're worth. Expensive goods might be more likely to be sold at auction than cheaper goods, just because the transactions costs might be a fixed cost that doesn't scale with the size of the transaction. But any loss or foregone surplus from uncertainty might scale with the size of the transaction.

Goods where characteristics are costly to assess, goods where buyers know more than the sellers, goods with heterogeneity and buyer valuation. So these are the kinds of things that we would expect to be auctioned instead of sold at posted prices. OK, so now let's consider a simple model to illustrate this point about information in particular.

So we're going to start with N potential buyers of some good. And let's suppose that their valuations are independent. So I'm auctioning off some good. And all of the potential buyers have a valuation. All the valuations are independent, and they're distributed uniformly on the unit interval. So they have a uniform 0, 1 distribution of valuations.

Now, the seller has a choice. Seller can offer the good at no cost at a posted price, or the seller can auction it off. And the seller knows the distribution of the valuations. So knows this uniform 0, 1 piece but doesn't know what the individual potential buyers, what their valuations are. So let's consider the posted price first.

So the seller is going to set a price, p, and sell the good if there's any potential buyer who has valuation greater than or equal to p. Is that agreed? Makes sense. OK, so we can write down an expression for the expected profit in this case. And that's just going to be equal to the price times the probability that there is at least one person with a valuation above p.

Where did I get this expected profit? Well, let me-- OK, so let me just first emphasize what the notation means here. So this is what I get if I sell the product. And this is one-- Oh, yes. The second term there is 1 minus the CDF of the n-th order statistic from a uniform 0, 1, evaluated at p.

Does everyone understand where this came from? Nope. OK, so remember the definition of the n-th order statistic. So n-th order statistic from a uniform 0, 1 distribution is just the highest value, highest realization, from IID random sample, from this uniform 0,1. And so basically in order to sell this good, I need to have the n-th order statistic-- there can be other order statistics that are greater than P, as well. But I need to have at least one-- in other words, the n-th order statistic greater than p.

And so this is, in fact, this is the probability that I sell the good. And this is how I calculate that probability. Yes.

STUDENT: This probability is basically just a number of people who are ready to purchase it for a price p or more.

SARA ELLISON: The number of people who are willing to purchase it, no. It's just the probability. So I'm going to-- I can get the distribution of the n-th order statistic from this uniform 0, 1 distribution. And I'm just saying what's the probability that the n-th order statistic is greater than or equal to some value p?

STUDENT: Think of it this way, that p, number of people-- sorry, number of people who are ready to buy it at the price of p or more, the probability of that, out of n people over here. If, suppose, 10 of us are ready to purchase it at p or more, so then that [INAUDIBLE] the price is the expected.

SARA ELLISON: I don't think that's quite the right intuition. So all I care about is the probability that one person is going to buy it if I post the price at p. That's all I care about. I don't care about how many other people. And furthermore, we know how to find the distribution of that highest price because it's just the n-th order statistic. So I think that the intuition that you're coming up with is kind of heading in the right direction, but I don't think it's precisely correct.

And so yeah, I think it's more useful to just think you need at least one person who has valuation above the price you post. And what's the distribution of the highest value? Well, it's just the n-th order statistic.

And I should step back just for a moment and mention, how does this formula for expectation here square with the formula we saw for expectation? So obviously, the formula that you saw most often for expectation at the beginning of the lecture was this sort of integral. It was an expectation of a continuous random variable. This is a discrete random variable here. So we have to go back to when I flashed up the equation for computing the expectation from a discrete random variable. So think about that.

OK, so we're just summing up the sort of probability function for each value of x times x to get the expectation. And here, this distribution has only two points. It's a discrete distribution that has only two points. One is the probability that we get 0 from this transaction. And that's the probability-- that's basically the CDF of the n-th order statistic evaluated at p. So it's the probability. It's 1 minus that probability.

But then that gets multiplied by p-- or sorry, it gets multiplied by 0. And so that term goes away. And so basically we have two terms in the expectation. One term is the probability that you don't sell it times 0, which is what you get if you don't sell it. So that term goes away. Plus the probability that you sell it times what you get if you do sell it. And that's this equation.

OK, so and now-- Oh, I said just apply the discrete formula. Note that the first term is 0. OK, so we've got an expression for the expected profit. And if I go back to my notes on order statistics, I can remind myself what the CDF for the n-th order statistic from a uniform 0, 1 distribution is. And so I can plug that in and I get this expression for expected profits.

OK, so now for a little economics, if this didn't confuse you enough already. Yep.

STUDENT: Explain the p to the power n.

SARA ELLISON: This is just-- if you go back and you look on your notes about where we talked about order statistics and we talked about the CDF of order statistics, basically what I'm doing is I'm just getting the formula for the CDF of the n-th order statistic and plugging it in. OK, so now for a little bit of economics. So we've got an expression for expected profit.

But as we often do, as economists often do, we assume that the seller is going to do what's in his best interest, even though we realize that doesn't always happen. In order to solve our models, this is typically the assumption we make. We figure that the seller is going to do what's in its best interest, in other words, set p to maximize his profits given that he knows what distribution the valuations are coming from.

So how do we find what that optimal p is? Well, it's just a little calculus. So we take the derivative of the expected profit function with respect to price. We set that equal to 0. And then we solve for price. So by the way, this isn't-- we know this isn't an economics class. And we don't have an economics prerequisite. We're not going to hold you guys responsible for profit maximizing behavior of sellers. But I think it's good to be exposed to this kind of thinking regardless. It's very sort of powerful way of thinking and solving problems.

OK, so what we do is we just take the derivative. We set it equal to 0. And we solve for the optimal price and that's what it is. So we get that it's the n-th root of 1 over N plus 1. Remember N is the number of bidders-- or sorry, the number of potential buyers, not bidders, in this case.

STUDENT: Number of potential what?

SARA ELLISON: Number of potential buyers. And then we can also take this optimal price, plug it back into the expected profit formula. And we get the expected profit under that optimal price is this value here. OK, so now we can just plug-in different ends and we can see what price the seller is going to set for different ends and we can see what profit he makes in the different ends.

OK, so here's a little table. N goes from 1 to 9. So if you only have one potential buyer and you know that his valuation is drawn from uniform 0, 1 distribution, the optimal price that you're going to set is 1/2. You're going to sell the product half of the time. That makes sense, right? You're going to be out of luck half the time.

But then because you sell it half the time and you get profit 1/2 when you do sell it, your expected profit is 1/4. And then the reasoning applies as you go down the table. To get this table, of course, I'm just plugging into the formulas for optimal pricing that we derive before. Yes.

STUDENT: But why is there a pi next to the [INAUDIBLE]?

SARA ELLISON: Sorry, that's the notation we use for profit. So in a probability class, it's not always-- pi is not always going to stand for profit. But yeah, in this particular example, pi just means profit. So that's an economics thing.

OK, so the posted price is rising as the number of potential buyers goes up. Does this make sense? Sort of intuitively, you have 100 buyers. We know something about the n-th order statistic when you have 100 buyers. It's going to be really close to 1. It's going to have high probability that's close to 1. So you're going to set a higher price. So that makes sense? And expected profits also go up as n goes up.

OK so I said, this is a consequence of the distribution of the n-th order statistic and how it changes is in N increases. OK, so now we know what profit the seller is going to make for posted price as a function of the number of potential buyers. So let's see how he's going to do in the auction.

So we're going to assume what's called an English auction, where the price of the good is going to gradually increase. And potential buyers, you can think about everyone starting with their hand up. Who wants this Picasso for $1? Everyone has their hand up. Who wants it for $2? Everyone has their hands up. Who wants it for $40 million? The hands start going down until there's only one person left. And that buyer's valuation-- sorry, there's only one person left. We're using v sub I to stand for buyer I's valuation for the product.

And then when only one buyer is left, he gets the good at V sub N minus 1, the second highest valuation. Why is that? Why does he get it at the second highest valuation? Yep.

STUDENT: Because until the last person falls away, you might have a higher price at point V.

SARA ELLISON: That's right. The auction's done when the second to last person takes his hand down. So we have no idea what the valuation of the winner is. We just know it's above the second highest valuation. But he gets it at the second highest valuation. So this is the classic open outcry auction that you see on TV shows and things like that.

So here, to compute the expected profits we're going to need the distribution of the N minus first order statistic from the uniform 0,1. So why the N minus first order statistic? Because that's the price at which this thing gets auctioned off.

So if I look up, if I go back to my notes. Actually, I don't think this is in your notes. I'd have to figure it out or look it up in some other textbook or something. I can figure out what the PDF of the N minus first order statistic from the uniform 0, 1 distribution is. And this is what it is.

And so we'll just, to compute-- Oh, yeah. And again, we're assuming this is the distribution of the evaluations. So then, to compute the expected profit in the auction, all we do is we integrate over the support of the distribution, the PDF of the n minus first order statistic. That's what the seller is going to get at auction for this product, times x. I put the x at the end. So it's right before the little dx.

And this is just the formula for calculating expectation from a continuous random variable. Everyone on board with this? So then we just sort of perform the integral. And we end up getting N minus 1 over N plus 1. That's the formula. And so we can make a similar table to the one we saw before, where we have N here. I'm not going to put p on the table here because in an auction, you're not setting a price. It just gets sold at the N minus first order statistic. But we can also put the expected profit in the case of the auction.

And let me put up the last column of the previous table to compare. So if we only have-- why do we get no profit if there's only one potential buyer or one bidder?

STUDENT: [INAUDIBLE]

SARA ELLISON: Yeah, exactly. I mean, the auction's over when there's only one guy left. If you start with one guy, he gets it for free. So anyhow, if there's one potential buyer or one bidder, you do very poorly if you auction it off. And that's compared to expected profit from a posted price of a quarter. If you have 2, then you end up doing a little bit better with the posted price than the auction. But the auction is catching up. And then for N equals 3 and above, you do better with the auction than the posted price. So I think I have-- yeah.

And I should say, by the way, this is not just, it's not that this is true for N equals 2 through 9 and then something crazy happens below. This is general. You can show that for any value of N greater than 2 in this particular model, the auction does do better. So what did this model tell us?

Well, conditional on the assumptions of the model, of course, it told us that the seller is going to do better with an auction when the number of bidders is large. So that might be consistent with intuition you have, or it might not be. I mean, I'm not sure that I had a really strong intuition.

Certainly in both cases of the posted price and the auction, the seller is going to do better if there are more buyers. That seems pretty clear to me. But this model, at least this simple model, shows us that the seller does better with an auction when N is large, large enough.

But here was the sort of amazing thing that you might not have realized when we were going through the example, or going through this model. The seller didn't need to know anything about the distribution of the valuations in the auction. OK so the seller did, in fact, need to know the distribution for the post to set the optimal posted price. But in the auction, the seller is not setting a posted price or not setting a price.

So the seller actually didn't need to know what the underlying distribution of valuations was in the auction model. And furthermore, you could work through this same example using different distributions for the fixed price part of it. So assume that the seller got the distribution wrong and thought that the valuations were uniform on 0 to 1/2 or uniform 0 to 2, when, in fact, they were uniform 0 to 1. And you'll see that the posted price does much worse in that case.

So the important point here is that there's a large reliance when you're setting a posted price on knowing the distribution of valuations. That reliance isn't there with the auction. And then the third thing to point out, this model doesn't have any transactions cost in it. So we talked about how transactions costs might be important. And you might not want to auction off things that are kind of low value items just because it's a hassle, and selling it at posted price is easier. This model doesn't have that in it at all. And so you have to keep that in mind when you're interpreting the results. OK, yep.

STUDENT: Isn't there like a transaction cost to gather all the people who do the auction?

SARA ELLISON: Yeah, absolutely. And that's, like I said, none of those transactions costs are in this model. So you could write another model that had transactions costs that could be a function of the number of bidders or something like that. And the transaction costs could-- I mean, you could think of various ways to model it.

STUDENT: If it's like, I don't know, if it's like wealthy people, then you have to [INAUDIBLE]

SARA ELLISON: That's right. Yeah, if you're running an auction at Sotheby's, you can't serve chicken fingers or something. Yeah, that's right. You have to have a nice auction house and yeah, exactly.

STUDENT: Is there a model to tell you how much you need to spend to make an auction happen? Because this is the optimal spot for a fixed price, gap is [INAUDIBLE] for the auction. [INAUDIBLE]

SARA ELLISON: That's right. You could think of, yeah, have that interpretation of the gap, as well.

STUDENT: So I don't understand the part where we keep doing this calculation, we don't need to know about an auction distribution because in N minus 1 first statistics, we are assuming that we know the--

SARA ELLISON: We needed to know. But the seller didn't need to know because the seller is not setting a price. What is the seller doing? Seller's just auctioning it off. So we needed to know in order to calculate the expected profit. But, yeah. But he didn't need to know.

STUDENT: I thought the seller's doing these computations to see whether auctions that are option or not.

SARA ELLISON: If the seller wants to do this calculation, then the seller would have to know in both cases, absolutely. Yep, yep, you're right. Now I'm glad for that clarification.

ESTHER DUFLO: You choose but if he gets it wrong, just maybe he has chosen the wrong mechanism, but at least he hasn't chosen the wrong price. You just need to know to choose the scope of error less than you actually-- you have to go very specifically. And you can show that, [INAUDIBLE] distribution was uniform from 0 to 2.

You think it's from 0 to 1, you're going to do your comparison under the same misleading model. You're still probably going to-- it's still possibly going to tell you when you're going to [INAUDIBLE] auction versus fixed price. The number of the N is going to change, the critical N. But that's, for most Ns, you're still going to get the right choice. But then if you go for a fixed price, you'll get the price completely wrong. Whereas, if it turns out to be for an option, there's an issue with price [INAUDIBLE] itself.

SARA ELLISON: Yeah, exactly. Was there another question? Yeah.

STUDENT: Quick clarification, as well. So the--

SARA ELLISON: I can go back, if--

STUDENT: Yeah, if you don't mind going back to the equation for option. Is how universal is that model?

SARA ELLISON: So in order to compute the expected profit from the auction, I had to make an assumption about the distribution of valuations. And keep in mind, I also made an assumption that those valuations were independent. So there are lots of cases in an auction model where you wouldn't want to-- where you might be hesitant or might not be particularly confident in the distribution of the valuations. But in particular, you wouldn't want to make the assumption that the valuations were independent.

STUDENT: That's not the most common, at least in reality.

SARA ELLISON: In some auctions that's a reasonable assumption. And in other auctions, it's just not a reasonable assumption. So anyhow, I would say this is a somewhat general model, but not completely general. It's not going to be a good-- and the other thing is that we ruled out things like, we ruled out being able to bargain over the posted price. We ruled that out completely. We ruled out setting a reserve price in the auction and not selling if you're below that reserve price. So there were certain things we ruled out just to simplify the model. Yep.

STUDENT: English auction and the Dutch auction...

SARA ELLISON: So there's actually sort of a very famous theorem in auction theory that says that in some kind of limited sense, all standard auction mechanisms are equivalent. And so that doesn't mean they're equivalent in all ways. But it means the expected revenue that the seller gets from a whole big set of standard auctions, regardless of what mechanism. And the English and the Dutch auction are in that set. The expected revenue is going to be the same

So yeah, that's called, not surprisingly, the revenue equivalence theorem. So, yes.

STUDENT: I'm just wondering under what conditions can we assume independence? Because it seems to me that the prices that one person might offer will be conditional on what the other [INAUDIBLE] are.

SARA ELLISON: Yeah, no. It's a good question. So the question is, under what circumstances does independence make sense? So let's suppose I'm auctioning off a bottle of wine. And it's a bottle of wine that the person who buys it is going to drink, so not likely to resell it, not likely to do anything else with it but just drink it. Well, then we can think, the valuation that those buyers are going to get from that bottle of wine can reasonably be assumed to be independent.

STUDENT: [INAUDIBLE] certain assumption what the price is.

SARA ELLISON: That's right. And the enjoyment you get from drinking that bottle of wine doesn't really have anything to do with the enjoyment I would get from drinking that bottle of wine.

STUDENT: If we had envelopes, for example, for the price, then they are independent. I don't know what other people are writing. And I just--

SARA ELLISON: Well, no, no, no. So that's-- you're talking about the dependence in bidding as opposed to valuations. And I kind of glossed over the distinction between a bid and a valuation here. That's one reason why I use the English auction mechanism. Because, in fact, the optimal thing-- we didn't talk about what the optimal thing for bidders to do. But in the English auction, the optimal thing for them to do is to bid their valuations. In some auctions, the optimal thing to do is not to bid your valuation. So there is a distinction between a bid and a valuation. Here, we didn't have to worry about it because the mechanism I chose.

Did you have a question?

STUDENT: No, I just think if you have a large enough, N you can sort of [INAUDIBLE] three people on there.

SARA ELLISON: Well, no, not necessarily. Not necessarily. I mean, if, for instance-- So let's think of a case where the government is auctioning off the rights to cut down timber in a tract of land. So they have public lands. And they're auctioning. And there's sort of a bunch of firms who are going to go and sort of cut down timber and mill it and sell it.

So for instance, my valuation-- I mean this is kind of a more complicated model. But my valuation should depend on everyone else's valuations if I'm one of the potential bidders there, because there's information in all of their valuations about how much timber is actually can be cut down and milled in what it's market value is. This is sort of a complicated, there's a complicated question of how to estimate the valuation. And so your valuation is going to be relevant information for me, as well, in a case like that.

OK, so this just gives you a little taste of auction theory. I happen to like auctions and I hope you thought it was interesting. But there's lots of other interesting and sometimes complicated economic issues having to do with auctions and auction theory. But it's also a nice illustration to show you that order statistics are important in the real world, too. So let's see. I still have a couple other things to say about auctions that will take us to the end of the lecture. So as I mentioned earlier, eBay, which was founded in 1995 as an online auction site exclusively, now only has about 15% of its listings in auctions.

So the vast majority of its listings are Buy It Now listings. There's also some kind of best offer listings which are kind of a hybrid. But the number of pure, the percentage of pure auction listings is down to about 15%, maybe lower by now, 15% by the last time I saw the data.

OK, so here's a graph illustrating that point. So this is taken from an NBER working paper, sort of mentioned up above. And what they did is looked at-- so the darker line is the share of active listings on eBay that are listed in auction format. And you can see as recently as January of 2003, that was almost 100% of active listings on eBay were auctions. And there have been some sort of changes in the rules, the listing rules in eBay over the years. But there has also been this sort of gradual decline over this period of auction listings.

So now we're down to, like I said, about 15%, and maybe even lower, of active listings or auction listings.

STUDENT: Was January of '09 where they kind of reevaluated their structure?

SARA ELLISON: So I should have looked that up to double check before I came in today. There was some kind of yeah, there was some kind of a rule change that occurred there. Maybe it was sort of a difference in the fees that they charge for different types of listings. And I knew that at one point and I forgot to look it up.

STUDENT: [INAUDIBLE]

SARA ELLISON: Pardon me? Yeah, although, I mean, this is-- keep in mind, this is probably a decision that the macroeconomic variables are not going to have a huge effect on. You're just a seller and you're like, auction or posted price. So this isn't not-- this is just the fraction of the listings that are auctioned versus posted price. So it's not like the number of listings plummeted. It's the fraction that was auctions versus posted price plummeted. So I said that.

STUDENT: [INAUDIBLE] the sharp drop in fall 2008 coincides with the decision in September '08 to allow good until canceled [INAUDIBLE].

SARA ELLISON: Oh, great. Thank you. OK, and here's another graph that illustrates kind of a similar phenomenon using different data. So from the same paper. And what they did is they used Google Trends data. Has anyone used Google Trends data? What did you use Google Trends data for?

STUDENT: Different companies to see how they change [INAUDIBLE] in popularity [INAUDIBLE].

SARA ELLISON: OK, so for those of you who are not familiar with Google Trends, Google Trends is a sort of a free facility that Google makes available to anyone who wants to use it. And you can go on to Google Trends and download data on the number of searches that have been done on particular search terms. And you can download the data over time, which is what they did here. So they were looking at the search term online auction versus a search term online prices. And they saw these sort of trends over time changing.

You can also use Google Trends to look at Google searches geographically. So you can ask them to give you data on Google searches sort of by region or by state or by city. And so this is actually sort of an aside, but Google Trends can be a very useful tool for sort of gathering data for various purposes. So it basically just tells you how popular search terms on Google are.

So what they did is they downloaded Google Trends data. And I should say also that Google Trends doesn't give you the raw number of searches. What they do is they give you some kind of, I don't know, a normalized or a search volume index, I guess they call it. So what they did here is they got Google Trends data on two search terms, online auctions and online prices. And you can see that there has been this sort of decrease from the index high of around 90 to now down to around 20 for the search term online auctions. The search term online prices has bounced around quite a bit but has maintained something like the same Google search volume.

This is an indication that people are just-- they're as interested in online prices as they have been for a long time. They're getting less interested in online auctions.

STUDENT: I would just be interested to see if we just aggregated data by the type of goods how that [INAUDIBLE].

SARA ELLISON: Well, strange that you mention that. So what's going on? So we have a theory that we developed just a few minutes ago, suggesting what a seller might prefer to sell an item with a posted price versus an auction. And that theory is useful for thinking about information asymmetries. But it was an incomplete theory in the sense that the model didn't include anything about transactions costs. And furthermore, these eBay graphs suggest that something has been changing over time. But that could be, our theory could suggest what has been changing over time that that's driving these trends.

So three broad hypotheses-- and I should say, these are more or less discussed in the NBER working paper that I mentioned, although actually, they don't discuss three very much, which I think is kind of an important hypothesis. So one hypothesis is that there has been a compositional shift of sellers and types of products on eBay.

So it used to be that eBay was mostly used for memorabilia, collectibles, antiques, estate jewelry, things like that. And now it's being used mostly for iPhones or Xboxes or whatever. OK and those products, the sort of new mix of products, is not a bunch of products where the auction mechanism makes sense.

Basically, could be that consumer tastes have changed. Online auctions are just not as fun as they used to be. So in the late 1990s, everyone was just like, hey, I'm going to go on eBay and I'm going to buy something on auction. And then the novelty wears off and then you're like, oh, the transactions costs are killing me. I'm sick of this. So that's certainly my-- I can tell you from my experience, that's what happened.

And then also, as was pointed out earlier, the price discovery benefits of auctions have declined over time. So online search has made it easier to find prices for comparable items. But eBay itself might have killed the market for online auctions because eBay has created thick, national markets for lots of things that didn't exist before. Those thick, national markets provide information on price. And so now maybe, you just go on eBay and you see, oh, what did this liberty head silver, or liberty head gold dollar from 1863 sell for on eBay?

OK, well, I got it. You know, I have 20 auctions, bid on 20 auctions, so now I know how to set my price. So eBay may have been victimized by its own success in some sense. So the paper does a little bit of this. And let me just mention briefly-- I won't go into any detail.

But what they can do, what they've done in the paper is they can decompose the shift over time from auction to posted price, look, exactly as you were suggesting, look at the different product categories, and look also at the experience of the typical seller. Because that lets you get at some of these informational questions. Like if the sellers that are selling on eBay are just much more experienced than the sellers that were on eBay 15 years ago, then maybe they don't need the auction mechanism to tell them about the price distribution.

And what they found, actually, was that these explanations only account for a fairly small fraction of the shift. And instead, the returns to sellers using auctions have diminished. And if you're interested, you can check that paper. But it's not a hard paper to read. And it's sort of a nice illustration of how you can take a data set and decompose it to answer different economic questions. OK, we'll call it a day.