Decision Making Under Uncertainty
Prof. Natarajan Gautam
Department of Industrial and Systems Engineering
Texas A&M University, USA

Lecture - 22
Safety Stock for Inventories

This recording is for Safety Stock Inventories. Now turns out earlier what we saw was perishable items- things that would perish at the end of the period like a day, a week or a year. Here however, we are going to allow carrying inventory over to the next period. So, period to period you have demand, then you can carry them over. So these are items that do not perish.

(Refer Slide Time: 00:46)

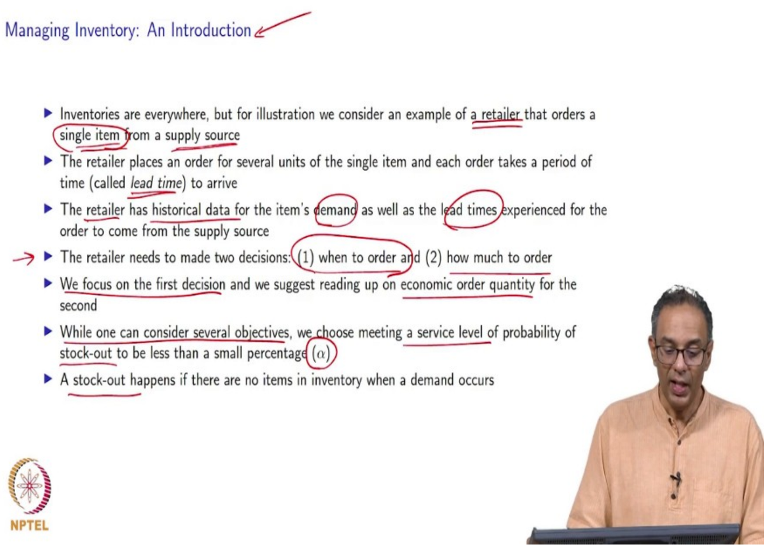

However, in a lot of ways this is somewhat similar; we add another aspect of this called the lead time. Let me just go over this clearly. So, we want to manage inventory; the difference here is we do not have an item that is going to be perishing over time. There is a retailer, and this retailer orders a single item from a supply source.

Again were sticking with a single item, I think these things can be nicely extended to multiple items; similar assumption were going to make in that, we could solve the strategy separately for each item. However, in the multi-item case, there are quite a significant difference in this case because when you order multiple items, you may want to consolidate and order. So, we are going to make an exception here and just say they are only going to focus on a single item.

Then what we did not consider before is the notion of lead time. So, the moment you place an order, it is not going to show up right away, show up in the morning or something like that. Instead it is going to show up after a while, this could be because of the shipping reasons or maybe the product was asked to be made and so on; there are many reasons for why there is a lead time.

Lead time: it is a time from when you place the order and the item arrives at the retailer. The retailer does have some historical data or has some way of figuring out how the demand looks like; and they also know something about the lead time. So, there are two random quantities right now: one random quantity is the demand and the other random quantity is the lead time. This is how much you are going to get when you order, and how long it is going to take. So, both those are assumed to be random and we feel that we have historical data to say a little bit about it.

Now, the retailer has to make two decisions: when to order, so looking at the inventory level, so I am going to draw a picture and actually look at the inventory level and when the inventory level goes low enough, you place an order. So, we will discussed greatly about when to order; we will not discuss how much to order? Now, this is you know usually done in a fairly deterministic fashion. So, how much to order is something that we will not get into, we will touch a little bit of this in the next topic, but for now we will not go there.

So, our focus will be on the first decision, which is when to order, and I would talk, I would say you know- look at economic order quantity for the second topic, which is how much to order? I will touch upon that just a tiny bit. One could look at several objectives, I think several cost related objectives and so on. We are going to pick something rather simplistic, we are going to say: "I want an alpha percentage as the stock out level". So, that means, a stock out happen if there is no items in inventory when the demand comes.

So, when somebody comes and asks you for product, you have nothing in inventory: that is called a stock out, and you want to keep that probability extremely small. So, that is the service level that you are providing or the retailer is providing his customer, and they are saying: well the probability that there is a stock out must be less than alpha, usually alpha is 1 percent or 5 percent. ok.

(Refer Slide Time: 04:31)

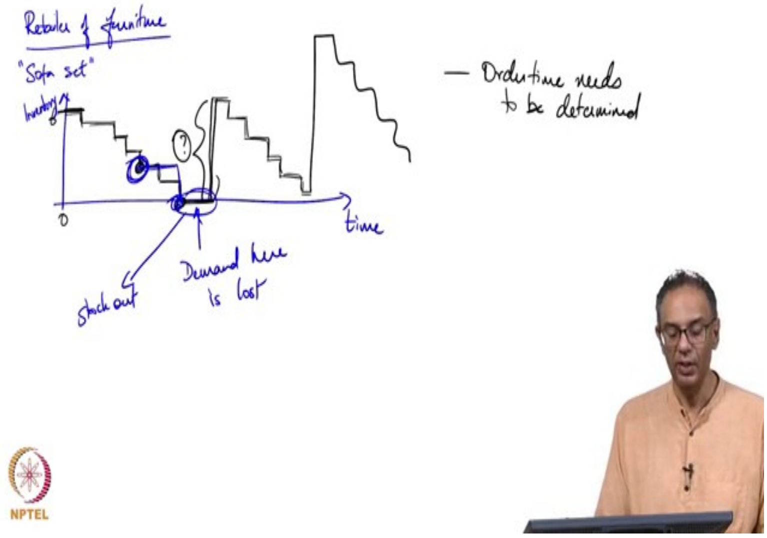

So, very rarely demand comes and you have a stock out (Refer Time: 04:17); so that is what you want to minimize. Let me draw a little picture to explain with an example. So, let us consider this example of a retailer, who is of furniture, there is a special type of furniture, think of a very nice sofa set. This is a very special sofa set that has to be special ordered and it takes a while for the sofa set to come.

So, what the retailer typically does is the following: retail we carry some inventory; so this is inventory and this is time. What the retailer would do is carry some inventory, then when somebody buys, the inventory level goes down by one. It stays like that, another person buys, it goes down by one. Stays like that and another person buys, it goes down by one. Someone else buys, it goes down by one. Someone else buys, it goes down by one. Someone else buys, goes down by one. Every time that happens, it goes down by one and then stays like this.

(Refer Slide Time: 05:28) Then whenever the order comes, it comes from here and then you go like this and then another person orders and then goes here, goes here. Then let us say, right here you place an order and then it will (Refer Time: 05:39) be like this and the inventory level keeps going like this. So, the inventory level is like a step function that is because this is the number of sofa set in the beginning with 1, 2, 3, 4, 5, 6, so there were 6 at the beginning of that period and then beginning of the measurement at time 0 and as time went by, when one person bought it went down by one.

So, every time somebody buys, it goes down by one and whenever a shipment arrived, it goes up. Right. So, that is what happens. Now, turns out the question is: when do you reorder? So, remember that the order time needs to be determined. When should we place an order- should we place an order here? Should we place an order, when it hits 0, and so on? So, those are questions to ask; and how much to order? So, what should this height be? Ok, how much to order is not something that we will look at; we will only look at when do we place an order? That is the situation.

Now in such a situation as described here, if there is a demand, and if it is lost then this situation is called stock out. This is the time when the number of so far in inventory is zero. When you are in that situation, you are what are called stocked out. So, therefore, you have to order sometime early. If you order early enough so that by the time the item comes, you are not down to zero. That is what you want to avoid right, you want to order of early enough.

So that you do not get it to a situation of stock out; we can never completely avoid stock out, but we could reduce the probability to something really small. Like 1 to percent 5 percent; so that is what we want. So, you want to be sure that the probability of stock out is extremely small; that is what were are going to try to optimize.

(Refer Slide Time: 07:54)

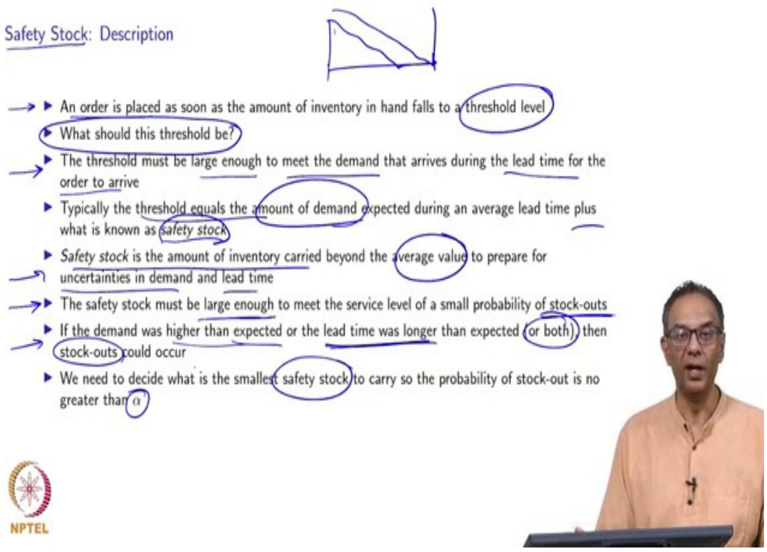

Now, let us fold your head, let us look at a few more details; so, some more description. So, there is a notion of safety stock. Safety stock is something that we defined a little while ago, but we are going to say more precisely what safety stock is. So, as soon as the amount of inventory falls to a threshold level, keep placing order (Refer Time: 08:17). As soon as the inventory level fall with that value we depicted (Refer Time: 08:23), you place an order.

But, what should that value be? So, that is the question we are trying to answer. When should we place an order? Turns out, we want the following feature: the threshold must be large enough to meet the demand that comes during the lead times (Refer Time: 08:42). That means, the threshold should be high enough that by the time the order comes in, you are not hitting zero. So, your demand does not fall off to zero by the time the order comes in, you do not want that. You want to order it early enough, so that when it comes you are not at zero. So, that would be a good thing to do. So, we will draw another picture later to explain that.

Typically, what happens is the threshold itself equal how much demand you want plus what is called safety stock. So, how much demand not you what you want, how much demand you expect to get? So, during an average lead time, let us say the average lead time is two weeks. Two weeks, whatever my demand is going to be, that much I must carry.

So, you should order when at least reaches that. You should place an order when that happens. But, you should have in hand a little bit more than that, that quantity is called safety stock. I will give you another picture later and explain this more clearly.

So, safety stock is the amount of inventory that you carry, which is beyond the average to prepare for uncertainties in demand and lead time. So, there are two uncertainties, the two random quantities: one is the demand and the second one is the lead time and we want to be sure that we are well prepared to take care of those uncertainties.

So, on average you want to have a certain amount and that is called the standard quantity and then above that you want to have what is called as safety stock. So, the safety stock must be large enough, so that you have a small probability of stock out. So, if you go on increasing the size of safety stock then the probability stock out go on reducing, but it is also expensive to carry too much inventory. So, at some point you are saying: you go on increasing and as soon as you reach that alpha value, you will stop.

So, if the demand was higher than expected or the lead time was longer than expected, then you could get a stock out and that is when your safety stock is really going to be useful. When the demand is larger than what you expected or the lead time was a bit longer, or both together is even worse, then you would could get a stock out.

So, you want to stock enough i.e., safety stock of a decent value, so that the probability of stock out is low. In particular, we want the smallest safety stock so that the probability of stock out is greater than alpha. So, you want to find: what is the smallest value to carry, so that the probability of stock out is no larger than alpha. So, let me take another example, which is near and dear to a lot of people, especially in places where there is not much water.

(Refer Slide Time: 11:40)

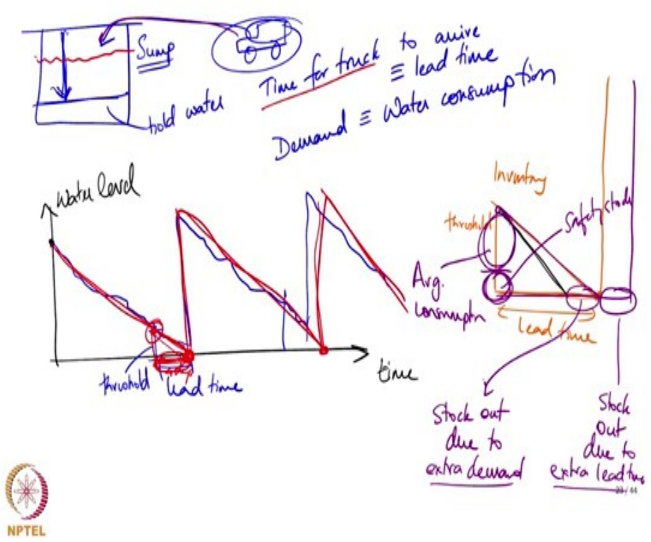

So, think of another special inventory; this is not a retailer example. So, let us say this is a sump, it is sump to hold water. So, a water sump holds water. Then there is a water truck that comes and fills the sump. So, what do we normally do? Well the sump is like an inventory, the sump level keeps going down as people use it. So, let us say this is for a lit colony of individuals' houses and then the water levels in the sump keeps going down and at some point you will call and have the truck come and deliver.

So, usually the truck does not come right away; it never does. They will tell you: we will be there in about four days ma'am, something like that. So, they will give you a lead time. So, the time for the truck to arrive is called lead time.

And your demand is basically the water consumption. So, if I were to plot a graph of the water consumption level over time, the inventory level, water level over time. What will happen is, let us say we had a refilling right here, then the water level might kind of go down like this, go like this and then you would order.

So, this is where your reorder point is. Your order here and then it will come in a few days. So, this is when it comes. As soon as it comes, you would go up, so somewhere here. The water level goes up and down like this and then you probably reorder here. And then it will show up right here, and then you would fill it and keep going down.

Now, I am going to say a few things. First thing that I do want to say- so, this is the lead time, so this is the threshold. Now, I have not said anything about safety stock, I am going to come to that in a second. Now the first question is what is the lead time? Now many times people consider things like water consumption to be somewhat deterministic, they do not consider that as terribly random.

So, many times people just model this as a straight line and then order goes up like this in a straight line or it goes up and so on. If it is a straight line and if there is no uncertainty; so, this is the first example there is no uncertainty, then you do not have to be in the situation of the dreaded stock out. So, imagine not having water- that is the worst thing that will happen. You want to be sure that you never run out of water right. So, among all our needs- things like air, water; these are tremendous need and you have to make sure that that never happens. We never want to have stock outs.

Now, that is easy to avoid if everything was deterministic. So, what happens is: let us say they consumption is somewhat deterministic; it is not a terribly bad example because you know most of us use the same amount of water, for bathing, for cooking and for using the bathroom and so on. So, it is not that big a deal. However, I will talk about the random things (Refer Time: 15:21) very soon. Now, when would you reorder? Well it is easy because you just look and then see: ok, I want to reorder so that during that time I am surely going to have positive quantity.

So, I know when I am going to become empty, if my lead time is also deterministic, which is definitely not the case in real life. However, pretend like the time for the truck is fixed. If the person says- I will come in 4 days, he comes exactly 96 hours nothing less, then what happens is- you could order nothing more or nothing less. So, what happens is you could order just before, so just exactly 4 days in advance you order, it will come then they would fill it up.

Now, how much you fill in water? Again it is not that big a deal because either you completely empty the truck which is an option or you completely fill up the sump. These are two frequently used options, but you know in this day and age and water is expensive. You could think of filling up somewhere in the middle. There is no particular reason to believe that it actually has to be filled up or the truck has to be emptied, I mean it all depends on the price of the commodity.

So, let us just not worry about; anyway that is not a decision that we are going to be taking. So, let us not worry about that right now. So, our situation is how much in advance do you order? That is when we look at a threshold level; now let me zoom into that threshold level picture and then say the following. So, if you look at the threshold level, so now let us say: if I were to be consuming at a deterministic non-random fashion; this is the threshold inventory level.

This is the lead time, I will plan so that my threshold is exactly equal to the demand during the lead time. The demand during the lead time is exactly equal to the threshold and thereby I will come here and then my new stock will show up here; go up. So that is what will happen.

Now, let us think about: what if my demand was a little bit higher. My demand was a little bit higher, I would be here right. Even if that lead time is fixed constant that exact four days; if my demand was a bit higher, I would hit here and I would be in a stock out situation. Because my demand is higher than how much water I have during that lead time. That is a big problem and this could happen because maybe somebody had extra guests in their house and suddenly they used up a little bit more water than they normally do. The other thing that could happen is- the lead time itself could be a bit longer.

The other thing that could happen is- the lead time itself could be a bit longer. So, that is that is another possibility. It is entirely possible that the lead time itself could be a bit longer. So, lead time itself could go here, so even if you had a deterministic consumption, the order arrives only here, and there would be stock out here out here. So, these two are situations for stock out due to extra lead time, this is stock out due to extra demand. Both are problematic: you had more demand or more lead time. Well, even worse is both.

So (Refer Slide Time: 18:58), we could have both: extra demand and extra lead time; so you want to have enough so that your threshold must be such that, this would be roughly the average consumption on an average lead time. So, I take the average lead time, compute what is the average amount I would consume. And then, I would keep little bit here- this is the safety stock. And the safety stock is there so that if my lead time or my demand is a little bit larger than what we had expected than the expected value or both are larger than the expected value, then I want to be sure, I have enough.

So, this one essentially does that. So, I want to keep enough safety stock beyond what is the average consumption level so that when the amount of inventory reaches that level, you would make a call. So, in the water situation as soon as the water level reaches this particular value, you would go ahead and make the phone call and then say: “hey, can I go ahead and can you guys go ahead and bring your water”. And the guy will say: “well will come anywhere between 3 days and 5 days”, and whenever the person shows up after the fifth day, we could be in trouble and this is exactly why we need enough of a safety stock.

Now, also notice in the water situation if we carried too much safety stock, then when the water person comes or when the truck comes, they may not be able to fill a whole bunch because all the water that they bring may not fill yourself. So, there are other concerns one could think about. Now, the question is: how do you quantitatively figure out- what is the safety stock? So, that is going to be our next thing.

(Refer Slide Time: 20:40)

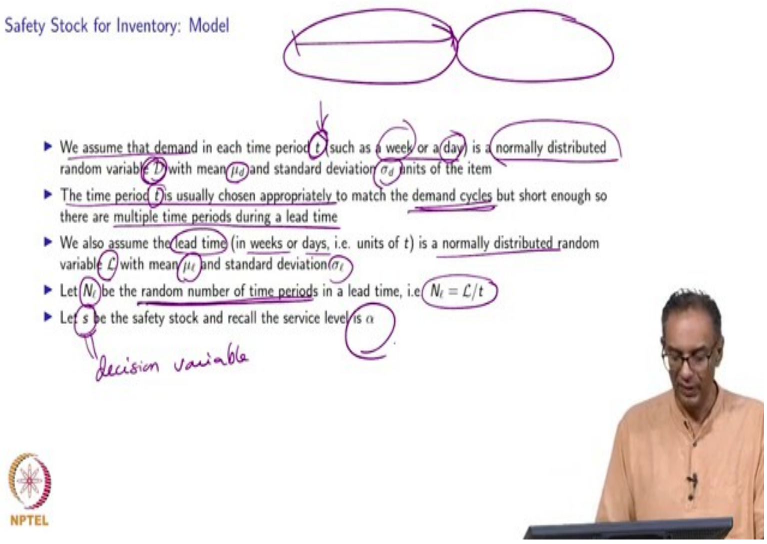

So, we assume that the demand like I said before is normally distributed. So, we straightaway pick a distribution. We could work with other distribution; I just want to let you know that that is not impossible, but there are some other assumptions you would need. We will not get into that, we keep life a little bit easy and say we have a normally distributed demand.

Now, I am using the letter *D* in script to mean random variable. So far we never used a letter *D* to be the random variable, we usually use x, y or z. The reason I am using *D* that the letter *D* stands for demand. It is easy to remember which letter is what; μ_D is a mean demand, σ_D is the standard deviation of the demand. So, these values are given to us- the mean or the standard deviation of the normally distributed demand.

Now, this *t* is a little bit complicated to think about, the reason for that you will see in a second. So, the demand itself is the number of items and the time period *t* is a week or a day. So, the unit of demand is how many? So, you have to fix something. So, let us fix the time period *t*, let us say that is a week or a day. So, that depends on the applications and some applications make sense to call *t* a day, in sometimes you know you can call it a week and that depends on how long it takes to ship some item.

Now, this should be chosen appropriately so that the demand cycles can be matched accordingly. So, let me just say a little bit about this. What I want to say is that: if you think about the demands, the time period *t*, and the demand does not have to be IID in each day. Let us say time period t is one week. The demand does not have to be IID, you know on a Monday, on a Tuesday, on Wednesday and so on. Maybe the weekend the demand is higher, in a weekday smaller and things like that; It could go anything. However, from week to week, the demand must be identically distributed, we have that requirement. So, we pick a period t appropriately so that the demand cycles are nicely modelable.

Sometimes you know you could have a demand for a month, sometimes for a year, you want to pick a quantity that is reasonable enough so that you have the other thing being multiple time periods during a lead time. So, multiple time periods need to be there during lead time. So, picking one year is probably too large. So, one day or one week- it is probably most often chosen. I mean maybe even one hour, two hours are possible and looking at much smaller time unit. So, depends on how far is the delivery location.

So, lead time again could be in weeks or days could be months, it depends on shipping from another country could be month. We assume that is also normally distributed and it is using the letter *L*. *L* for lead time, mean is μ_L, and started deviation σ_L. So, the lead time now this has units of time, *D* is dimensionless, *D* is a number: how many I get as demand and *i* has this same dimension as the lead time. So, we try to pick and that is the thing that we have to be a little bit careful about. So, I am going to use a letter N_l to be the number of time periods inside a lead time. So, $N_l = L/t$

So, whatever lead time divided by t then it just tells me the number because when we want to compare two things, it is better that we compare two things that do not have units and therefore, we say: N_l is the number of time periods that is sitting inside a lead time. And s is what we want. So, this is our decision variable- what should be the safety stock? Remember, we have a non-perishable item. So, even if our demand comes in and we have items left, we could use that in the next period. You can keep using it and add infinitum.

So, there is no issue and remember that the service level is α. So, you want to be sure that the safety stock is just small enough, so that the service level of α can be reached.

(Refer Slide Time: 25:02)

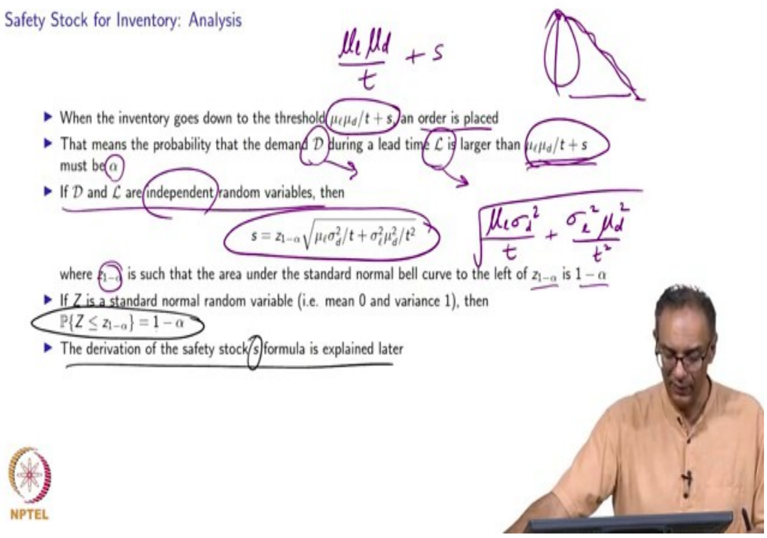

Now, we go into the analysis; I am not going to give you the proof, which we will do in the next lecture. However, I just want to say a few things. So, when the inventory level goes to a threshold of this quantity. So, when the inventory level reaches this threshold of $\frac{\mu_l \mu_d}{t} + s$.

Now, remember μ_l: lead time, is a time unit, t is a time units; so, time units cancel away, μ_d is a number, s is an inventory number. So, both the unit wise we are good and let us say you place an order as soon as the inventory level reaches that quantity. So, what does that mean?

That means, the probability that is larger than $\frac{\mu_l \mu_d}{t} + s$ needs to be at most α.

So, the probability that the demand during this lead time, as the demand *D*, there is lead time *L*, is bigger than this threshold quantity: that is when you order. When we reach the threshold, you want order so that the threshold is high enough that the probability of stock out is extremely small: alpha.

Now, we are going to make an assumption not only that *L* and *D* are normally distributed, we are also going to assume that they are independent random variable. If they are not independent, there is another way to solve this problem. And there are some books and articles that talk about what should be the safety stock levels if they are not independent. It is not a terribly bad idea to assume that they are indeed independent.

So, I will go ahead and do that and then we can prove, which I will do in the next set of slide is

$$s = z_{1-\alpha}\sqrt{\mu_l \sigma_d^2/t + \sigma_l^2 \mu_d^2/t^2}$$

Again lead time has units of time, σ_d^2 does not have any units. Likewise, σ_l^2 has units of t squared, μ_d^2 does not have any unit. So, this thing is the unit of how much demand we have, z is our standard normal bell curve so that the area to the left of $z_{1-\alpha} = 1-\alpha$. I am going to do a little picture to actually show you that.

(Refer Slide Time: 27:47)

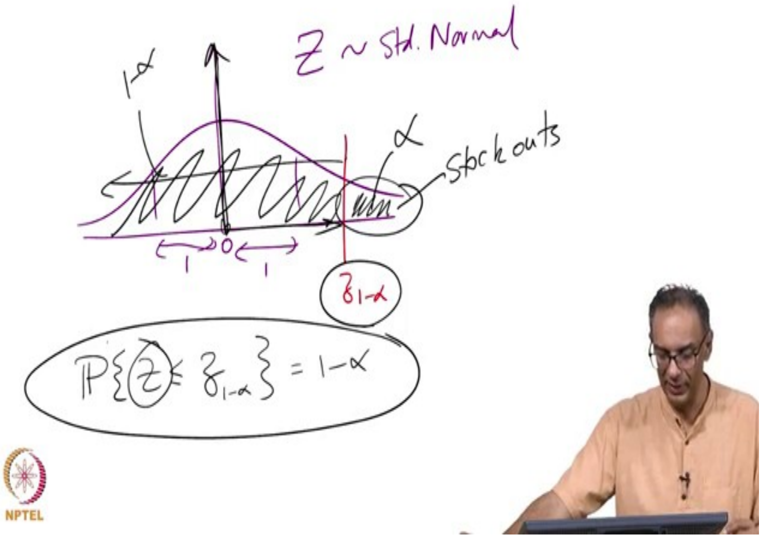

This is the standard normal curve, where you see one and zero. So, this is 1, this is 0, this is 1. This is z standard normal. Now, area to the left of red area ($z_{1-\alpha}$¿ under that curve is $1-\alpha$.

So, what is the value of z so that the area to the left of z is $1-\alpha$? That is the same as asking what is the probability that z is less than or equal to this is $1-\alpha$? The area to the left of $z_{1-\alpha}$ should be $1-\alpha$. So, we want to compute that. Why are we doing that? Well think about this and we are going to explain very briefly. If I exactly add the mean and there is no uncertainty then this is where we would be. This tells me how much away from the mean I want to do. So, I want to be large enough. So, that the probability that I am here is α. So, this is where you get stuck out. So, this is corresponding to stock out.

So, I want to pick a safety stock that is large enough so that my probability of stock out is α. So, that is basically what we want. We want to be sure that we are here. The probability that the random variable z is less than or equal to $z_{1-\alpha}$. We will see this result, and then derive it a little bit later in the next topic.

Now, I do want to summarize very quickly by saying that we have seen the normal distribution and this is something that we did especially when we have a table. On the other hand, if you do not have a table, you just want to use the software. This will work as well. We will talk about how to derive those *s*: the safety stock a little bit later in the next item. So, we will stop here for now and then we will continue with this next.

Thank you.