JIM PETERSON: I got a call from Tom, cold call, and he says, "Jim, about a year, a year and a half from now, we're going to have this meeting. Can you talk about Bayesian stuff, that stuff you do?" And I said, "Oh, yeah. Sure. A year in advance? No problem. 'Bayesian stuff,' what do you mean?" We have a follow-up from that, and then Tom gave us marching orders and said, "Okay when you're talking about your stuff, talk about it in the context of uncertainty and then how do we reduce it, or what's the important stuff to reduce, or how we figure out what's the important uncertainties from the non-important uncertainties." And I said, "Do we have a few days to do this or what?"

So I'm going to try to do this in half an hour. I'm going to try to go pretty quick, if I can. One thing I want to point out, now that everybody's plugging their book, I want to plug my book with Mike Conroy in the lower right-hand corner on this side. The book basically covers everything I'm going to talk about today. [Slide 1] On my website, I have links to the classes that include in this material, and if you're interested in it, there's actually an e-campus class at OSU on this topic, a lecture and a lab that you can take. It is targeted to a master's level student. All right, let's get started [Slide 2]. Some quick definitions, what is certainty?

An event is considered to be certain if it's going to happen with 100 percent probability. Or if you know what the actual value is, the true value, that's certainty. Anything that falls short of that is "uncertain." So in decision modeling, we use probability distributions to represent uncertainty, at least most sources of uncertainty. Some common forms of uncertainty that we generally need to deal are categorized in terms of reducible uncertainty and irreducible uncertainty [Slide 3]. Irreducible uncertainty includes environmental variability and demographic variability, these are things that additional research will not going to reduce. Hence they are irreducible. For example, you're going to have wet year, dry years, you could study it all you want but you're not going to reduce the water year probability. You might estimate a little bit better, but you're not going to reduce it.

Things that are reducible, this is the thing defined as epistemic uncertainty by Lance (Gunderson) yesterday or linguistic uncertainty. With a lot of efforts and additional study, you can reduce those things. So, as epistemic uncertainty is what I will be discussing. Now, we are generally familiar with statistical uncertainty. We have data, we have standard errors, and we have confidence limits to represent this uncertainty. Observational error is another source of uncertainty. You might not detect every fish that's there or every insect that's there. We also have something defined as structural uncertainty. I'm going to point out source out because it's most often overlooked, and it's also the source of conflict quite often when we're talking about your model versus my model.

[Slide 4] For example, I could make the statement that fish population size is positively related to stream flow, and that's what this figure shows. Now, a positive relationship can take an infinite number of forms. Here, I'm showing three basic forms. I submit to you that based on how I thought the system worked, I would come up with a different stream flow standard. Let's say for example if I thought it was linear, that white line. For every unit increase in stream flow, I'd have an increase in the thing that I cared about. In this instance, fish population size. If the true relation is represented by the dotted green line, there becomes a certain point where I get diminishing returns by increasing stream flows. If the true relationship is the orange lower line, there will be no increase in fish population size if I don't obtain a stream flow greater than this amount where the line begins to curve up. Quite often, these different ideas of how the system works results in gridlock and conflict when we have disagreements about your model versus my model. When I work with decision makers, I generally don't care about which model is correct. I just want to know your model is, that is how you think the system works. We can actually incorporate, multiple models into natural resource decision making, as Lance alluded to yesterday, and we can include them and use them with model probabilities or model weights to combine them and make predictions about the likely effects on management actions.

All right, on to Bayesian inference [Slide 5]. I could create a Bayesian model of a simple linear regression problem. I don't know why I would want to do it, but I could do it. I then could use something called Monte Carlo Markov Chain, i.e., Bayesian, methods to fit the model. It doesn't really make it Bayesian. Probabilistic networks also are often thought of as Bayesian. However, these techniques use conditional probability; that is probability of this event given, some other event. That's not necessarily Bayesian either, though it could be. When I say Bayesian, I'm talking about the use of prior information and combining it with new information or evidence. We do that and we come up with something called a posterior estimate. That posterior estimate is an updated idea of either how the system works, what the value of a particular parameter is, whatever you are interested in.

So, here's a real quick example. I can estimate the probability that, say, a fish species occurs in a stream given I didn't collect it while sampling. I have a couple of papers on that with my colleague Peter Bailey. Now, I was teased a few years ago by my colleagues who happened upon a journal of crypto-zoology where somebody used our techniques to estimate the probability that Sasquatch occurs given that people haven't seen it. So these techniques can be abused just like any other quantitative technique. The beauty of Bayesian techniques is that we could use that posterior estimate as a prior estimate and actually learn iteratively. That's the whole idea-- pretty simple huh?

When I teach these concepts, what I try to get in people's heads, is that this whole idea of prior information, prior information could be in the form of data. It could be in the form of reports. It could also be in the form of institutional memory in people's heads. We don't need to reinvent the wheel every time if we can use Bayesian techniques to take the information we have, combine it with the data that we need, and come up with some revised estimate.

I'm going to give you an example of how we might combine prior information with data [Slide 6]. On the left-hand side, I have this thing called the imprecise prior and we have on the Y-axis, the probability of redd [scour]. We're in the Pacific Northwest and biologists in the Central Valley are concerned with redds scour. They want to know what is the probability of redd scour given some discharge value. Here, I have an imprecise prior on the left-hand side, so that's mean (box) and 95 percent confidence limits (lines). On the right-hand side, I have a precise prior. There is the data, it's the same data on both sides, it's the observed redd scour events that occurred I think in three out of seven years. I can actually combine those two information sources and come up with a posterior estimate. When I have an imprecise prior and I have a sufficient amount of data, notice that the posterior more closely resembles the data. That's because the evidence is stronger for the data. This is not mumbo jumbo or magic. It's just mathematics. On the right-hand side, there is the posterior distribution when I combine a precise prior with data. Notice that the information is shifted up. That posterior looks a little bit more like that prior because it has more information. Now, people are sometimes bothered by this and think it (Bayesian) is magic or voodoo or something like that. It's just mathematics.

[Slide 7] If the prior is more precise than the data, the posterior is going to come close to resembling the prior. If the prior is less precise than the data, then the posterior is going to more or less look like the sample data. This isn't controversial, just mathematics. Sample size plays a role as well. The greater number of samples, the greater influence the prior and the posterior is going to have. And in fact, Bayesian scientists like to talk about prior sample size that's contained in the mean and the variance of, say, a normal distribution. You can actually estimate something known as prior sample size using the mean and variance.

So, more information, greater influence; less information, lower influence. Makes sense, right? That's Bayesian an inference. Now, how do we use this anyway? How do we use this for flow management or decision making? [Slide 8] Well, first we need to get the priors. How do we do that? Previous studies and published reports are ideal- but only if the reports contain useful information. Now, what I find really kind of frustrating is that not nearly enough agencies or other folks publish their data so that getting this information from previous studies can be very difficult for meta-analyses. If we (scientists and managers) actually report things in certain ways that people could use it in the future, such as means, variances, effect sizes, we would honestly be a lot better off. Otherwise, it is a loss of information and waste of scare management resources. It really is. We can also use elicitation to obtain expert judgment. Robin talked a bit about this during lunch. There are ways to go about eliciting values from folks. You could even get measures of uncertainty. You use something called a four-step process to obtain estimates that is used as a starting point, a prior. Alternatively, people sometimes use this thing known as diffuse or non-informative priors. I'm not sure why you would do that for management decision making. If you had some information, I think you'd want to use that information.

Some commonly used Bayesian tools [Slide 9]. The most popular one is, MCMC. If you ever hear of the thing MCMC, it just stands for Monte Carlo Markov Chain Methods. It is highly versatile and can integrate multiple data types and multiple data sources. I could take data from different studies and combine them in a single framework to estimate a common parameter or actually parameterize an entire process based model. It's really a natural fit for decision making and decision modeling. In fact, MCMC defines a joint probability distribution. So the previous talk we just saw, I would probably use MCMC to come up with a joint probability distribution that incorporates all the sources of uncertainty that were identified, and obtain a single posterior distribution. We also have these things called probabilistic networks, aka Bayesian belief networks, directed a-cyclic graphs, and influence diagrams. These methods represent something known as a joint probability distribution. They're very, very efficient at propagating uncertainty from multiple sources. For example, if you think about it this way, uncertainty in some flow metric plus the uncertainty in -- what's the physical habitat look like at a given flow, and maybe the uncertainty of the ecological response to the flow and habitat? So, probabilistic networks tend to propagate uncertainty quite efficiently, and they do a very good job with it. Now, the problem is that the as you increase the number of parameters in a probabilistic network, the uncertainty propagates and gets bigger and bigger and bigger and bigger to the point where it forces you to simplify your models. Otherwise, you can end up with a model that's basically useless for making a decision because the estimates of management effects will be very uncertain.

I'm going to give you an example of how you would go about using these things [Slide 10]. This one is from something I worked on with folks in the Southeast when I was with the Georgia Coop Unit, something known as the Southeast Resource Assessment Project. And the idea was to evaluate the climate change effects on stream flow and temperature and how it is eventually going to affect the aquatic biota. This is the Apalachicola/Chattahoochee/Flint Basin, the source of tri-state water wars and this was our study area right here. All of the habitat and fish population data were from the southern portion of the basin, from the Flint River Basin, area right around here. We needed to expand our models to species that occurred in the upper section of the watershed that contained cool and cold water fish species. We had very sparse data available for this area, but it was primarily from adjacent basins in the Blue Ridge physiographic province, which is pretty much where the upper Chattahoochee is located. Our task was to develop tools for estimating the response of the cool water biota to changes in the hydraulic regimes and the temperature regimes.

I recently published a paper with Collin Shea, and we used something known as a Bayesian State Space Model, to estimate the response of fish populations to flow variability and some other flow metrics. [Slide 11] Here on the Y-axis, is the probability of reproductive success and on the X-axis is flow variability during spawning and rearing time period. What we did was model the probability of success for individual fish species as a function of species traits. In this instance, it happens to be broadcast spawning species. This dotted line is the average response. This white area you see on either side right there is the predictable variation in the response from species to species. So if I were to look inside this white area, you would actually see individual lines for all those species that we included in the analysis and for broadcast spawners, it was something like 15 species that we had data on.

[Slide 12] Now, this information right here represents how much variability we expect in response from species to species given that they are broadcast spawners. And we can use that information. We had limited data on rainbow trout that we collected as part of a study with the Coweeta LTER. We used that light-shaded area as our prior, and our posterior estimate with 95 percent confidence limits are shown right there in that dark line with that little bit of orange or shaded line. We could then use that in our models to estimate the response to rainbow trout populations to changing flows and thermal regimes.

So, how is this useful for water resource decision making? [Slide 13] Well, I'm going to assume that we have objectives and decision alternatives. What we're going to do is we use this information to build decision models. And in fact, I would probably start decision modeling with existing information. I wouldn't go out and do a study right off the bat. I'd start with expert opinion, meta analysis, and expert judgment. Data would come in later.

{Slide 14] So the common question I often get is, whoa, whoa, whoa, you're not using data yet? Won't the assumptions or any of these Bayesian mumbo jumbo stuff affect decision making? How is it going to affect the decision? You have all this uncertainty in the decision model. And the answer is yeah, sure, maybe, probably, but what we really want to do is something known as a sensitivity analysis. So going back to the task that Tom set out for me-- how do we identify what are the important components, the key uncertainties, what can we do about them? The answer is sensitivity analysis.

Now, the sensitivity analysis I am referring to isn't just simple perturbation analysis. In decision sciences, we focus sensitivity analysis on one key aspect: what would we do differently if we knew X? What would we do differently, not how much does it change habitat or how much does it change the fish population. But given a set of decision alternatives, if I knew X with perfect certainty, would I make a different decision than if I didn't know X with perfect certainty? I guarantee you the type of sensitivity analysis focuses you on decision making. You also use that to prioritize research and monitoring. That's what you use that for. So it focuses decision making on what we need to know and how much (we need to know) is enough?

I’ll give an example of how you go about doing that [Slide 16]. We worked on this project in Georgia. The team included several people in this room. Jonathan was there, Tom was there, and Brian was there. The project was called water availability for ecological needs in the ACF Basin, we built a spatially explicit model based on stream segments. And if John is still here, we used channel types or channel classifications as one of our state variables. We had flow models, habitat models, and fish meta-population models for 43 fish species. We represented statistical uncertainty with errors in the flow and habitat models and structural uncertainty about the population dynamics using alternative fish population meta-demographic models.

I can go through and show you all the fancy statistical approaches to evaluating how decision is affected or how uncertainty affected decisions but let’s keep it simple. [Slide 17] Here, I have alternative local extinction models, i.e., extinction of individual species in stream segments. One model estimates extinction as a function of chronic flows and the other as a function of acute low flows. So on the X-axis, Kenneth Odum estimated changes in our discharge statistics as a function of daily water withdrawal, a million gallons per day. On our Y-axis is the loss of fish species richness as a function of water withdrawal. So you’ll look at this green line, this assumes that is chronic flows are running the system and driving the local extinction of fish species. This is our estimate of the change in species richness with daily water withdrawal (green line). However, under an alternative model – and that, is acute low flows largely drive the extinction dynamics of the system – you see a very different picture right here (blue line). So if I thought this is how the system worked, you can’t take much water out of this system, can you? If I thought chronic flows are how the system works, you can probably take a little more out and not really have much of a measurable effect on fishes. So I would call this a key uncertainty. Now, how might I manage the system? I might create something known as a composite model or an averaged model. I might have equal faith in both models, and I might manage based on that dotted line (the averaged effect).

Now, reducing uncertainty, this is the important. [Slide 18] I have uncertainty and it has a big effect. This is a key uncertainty. There are a few other key uncertainties. I can guarantee you if you go and do a sensitivity analysis that you will have, generally speaking, even if you have a 60-parameter model, less than half a dozen will be the things that largely drive a decision. Now, I’ve put together some really complicated models for the California Central Valley Project, for white sturgeon, green sturgeon, and Chinook salmon. Each one contained over 60 parameters, and for each one less than 10 parameters came out as those that largely drove the decision. Focusing on all of the parameters in the models is unnecessary if 10 are running the show or six are running the show, those are the ones I want to focus on. Those are the one I want to monitor; those are the ones that I want to reduce uncertainty. They make a difference. So we can do new studies or experiments, or we could do adaptive management. We can do monitoring, and we can use that information to revise and update our beliefs of how the system works.

[Slide 19] We also want to know though is not just to identify those components but what will it give me, what will I gain by doing future study or by doing additional monitoring. How much information or data is needed? How much do I need to reduce this uncertainty, which Tom has asked me to talk about? And how much is enough? How much of this additional study is enough?

[Slide 20] Again, we can draw from the decision sciences and we can use their techniques to estimate these things, to look at the tradeoffs. One is called estimating the value of information. It's the expected value of a decision: what are you going to get from that decision if you know everything or a particular component with perfect certainty. It could be done on model parameters, or it can be done on model inputs, it could be done on system state. What I mean by system state is what's the current amount of habitat or what's the current flow regime? Now, we express this value of information in the currency that is valued by the decision maker. If it's happiness points, it would be happiness points you're going to gain. If it's money, it'd be money, if it's water, it would be water. If it is fish species persistence, or abundance, or whatever, that's the currency that's going to be estimated. So for example, by reducing this uncertainty, you're going to get X (a value) more fish.

So let's go back to these alternative extinction models. I usually estimate what's the gain in terms of the biological outcome. [Slide 21] Here I'm going to estimate what's the gain in say, water yield. What additional water can you take out of the system? I'm going to assume a constraint that species loss has to be less than or equal to 5 percent of the total species. That might be an acceptable loss, or I can combine objectives (e.g., species and water) with utility function but I don't want to go there for this example. So, getting back to daily water withdrawal. If I assume that the fish local extinction dynamics is chronic flows, I can take 3.5 million gallons per day. If it is, acute low flows I can only take 1.37 per day. When I weight those by the belief in those models, I get 2.44 mgd. For the composite estimate, I get 1.83. Here's the math, the value of perfect information. If I resolve this uncertainty, I get an extra 610,000 gallons per day. That might be pretty good if you're a decision maker.

[Slide 22] Now, not all information is perfect, right? We have sample data. We've got sample variance to deal with. So what we need to do in incorporate this uncertainty to estimate something known as the value of imperfect information. We do that using some modeling or some simulations to figure out what's the error likely to be if we conduct that study.

[Slide 23] So in this instance, I did some simulations – don't need to worry about the details. If I sample on two occasions, I'm going to get an error of about 35 percent in species richness estimate, and then I use a little Bayesian statics to estimate the true species richness, given I estimated 25. It's 25 plus or minus 4. I plug this estimate into my model, do some magic Bayesian stuff, the value of imperfect information if I take two samples, collect two samples, is 0.26 million gallons per day. I could do the same thing for four, five, six, seven, eight sample occasions. (I can tell you right now that it asymptotes out at four sample occasions.) I could put more effort and get a better estimate of species richness and the value of information there is a little bit higher 0.49. So I can actually say "You know what, I'm not going to gain anything beyond four samples." You could do this for any model parameter. You can do it for whatever you want. You could identify those key uncertainties; you can say what it going to take to reduce uncertainty. It's just math, nothing magic, and we use Bayesian techniques to update the probabilities.

[Slide 24] Reducing uncertainty, monitoring. So Mary Freeman, through Jonathan Kennan and the USGS Water Smart Program, had a graduate student Rachel Katz, who's shown here, go about and do some monitoring in the Basin to reduce some of these uncertainties in the model. So Rachel did a bunch of sites, 21 sites, and let's see right here, Potato Creek and Ichawaynochawy Creek. Samples were collected in 2011, 2012 and also 2013. I'm only going to include 2011 and 2012.

[Slide 25] So, here are the updated model probabilities for the acute flow models. We're using Bayes' rule to update these probabilities through time. Our prior is 0.5. We pull in the data that Rachel collected, and we come up with a posterior. I am showing you two species that really kind of behaved really nicely – It doesn't always work this well – mosquito fish and spotted sucker. Look, I have fish in this picture. So, prior 0.5, 0.5. When we look at the mosquito fish data, we evaluate what happened compared to what we estimated. We update that probability. Now we're having a little more faith that the probability that the acute flows are what drives a mosquito fish extinction dynamics is a little bit higher. But if I look for spotted suckers, there's less evidence for it. I think it's more chronic flows, not acute flows that affect suckers. I go through it again, I take the posteriors (yellow bars) as a prior, I compare data to it, I get a posterior. Same thing for spotted sucker. Key point: I'm iteratively updating through time. This adaptive management approach is more of the Australian/decision theoretic school, where you update things through time. Either way, you still can update knowledge using Bayes' rule.

[Slide 26] Here are the estimates for the "averaged" model and how much they changed with the posterior weights for all the species adjusted (Blue line) versus the original estimates (green line). So we're improving the model estimates through time.

[Slide 27] To wrap things up, I believe that Bayesian approaches are vastly underutilized for water resource management. There's a recent estimate that about 15 percent of ecologists use Bayesian statistics or Bayesian approaches. I think it's underutilized, particularly for managers. There's some cost savings by leveraging existing information, and I really don't think we need to conduct a study, intensive study on every site that we're going to be evaluating stream flows. So with Bayes we can leverage information, that's the whole idea. We can use Bayesian techniques to propagate uncertainty and also to figure out what uncertainties are important. We also can use Bayes to update information and learn through time. We can narrow those confidence bands, and we can figure out which uncertainties are important. It's a real natural fit with the decision modeling, because we can do all these things that I just talked about. And that's sort of the key, at least the way I see it. I'm a little preachy about this, but that's why I wrote a book.