Quantifying uncertainty: An obligation or tool of discovery?

Mark E. Harmon

The following is a modified transcript of a talk presented at the December 2016 AGU meeting in San Francisco.

Rather than show you yet another Power Point-based presentation, I will share some of my experiences and thoughts that have lead me to believe that, while we should be quantifying uncertainty in ecosystem analysis, it would be a mistake to think of this solely as an obligation. My thesis is that obligation is not a particularly strong motivation for uncertainty analysis compared to other more compelling ones such as using it as part of our science discovery tool kit. In making this argument I draw upon personal experiences, which admittedly are quite subjective. However, despite being "objective" scientists we need to recognize that our motivations are often quite subjective. And since my talk deals with motivations I think it is quite appropriate to reveal mine and experiences that shaped them.

Before moving on to other kinds of motivations, let us consider obligation. I know that there are many things we should all do. And I also know there are things that I should do. I know, for example, I should eat more fruits and vegetables. I have known this for decades. So how is my plan for eating more fruits and vegetables going? Well, I really should eat more fruit and vegetables. My point is that obligation is not a strong motivator even when enforced with guilt. I say this as someone raised Catholic in New England, a place where rations of obligation, smothered in guilt are daily fare. And that is the problem with obligation, eventually one becomes desensitized to it. To return to my actual diet, I do feel guilty each year as I fill out the dietary preferences in my health insurance questionnaire, but then I also know that this feeling passes in a few days. In contrast, I have no issues in revealing my exercise routine. That is because I vigorously exercise regularly. Why? Because I want to exercise despite the fact it takes time, energy, leads to off-putting odors, sometimes hurts, etc. I exercise because it makes me feel calmer, stronger, and more confident. So perhaps the key to doing something is to want to do something. Which leads me to this question: why should we want to analyze uncertainty?

The first reason for me was curiosity. This stemmed from one of many conversations in which a group shared concerns about the obligation to estimate uncertainty. Most of those conversations concluded that while the estimate should be made, we would not actually make one. However, in one case a large group had just completed an analysis of Canada's terrestrial carbon balance and decided to take the next step and make a rough estimate. Opinions varied greatly. One participant offered that most of the errors would be offsetting and thus uncertainty would be as good as zero. Another, worried about the multiplicative dependencies and complications, offered an estimate of "plus and minus infinity". After careful consideration the consensus was that the uncertainty was plus and minus 10%. Why 10%? I have no idea, but it did make me curious about how to make a more defensible estimate.

My first opportunity to address this curiosity came when I was asked to estimate the carbon balance of an old-growth forest in western Washington state using ground-based measurements. I hypothesized the carbon balance was zero, but realized that the odds of it being exactly zero were also exactly zero.

To test this hypothesis I would need to quantify the uncertainty in the estimate and test if it significantly differed from zero. I realized if I was to do this, even for such a simple system, that the uncertainty analysis could not be done after the fact. Rather it needed to be incorporated at the start so that the components could be quantified and assembled at the end. In addition to hypothesis testing, estimates of uncertainty allowed us to compare against estimates based on eddy covariance. My analysis, which admittedly would be primitive and incomplete by today's standards, indicated the eddy covariance estimate for the year that data were available was outside my range of certainty. Nonetheless, proponents of the eddy covariance insisted theirs was the correct number. Unfortunately, there was no uncertainty estimate for the eddy covariance estimate, in part because it was so large as to be judged "meaningless". However, it really did have meaning: it meant that their estimate overlapped my estimate. Viewed from the perspective of a single number the discussion centered on which of the two numbers was correct and by inference method was correct. Viewed from the perspective of uncertainty, it was a matter of which estimates overlapped. I personally found it more scientifically fruitful to defend a range rather than a single number. By the way as eddy covariance estimates were added in subsequent years, it became clear, that while individual years were above or below zero, the long-term average balance was within the uncertainty that I had originally estimated. And this makes perfect sense in that the ground-based measurements were based on long-term measurements and not annual ones. So the argument, which at the time seemed to center on the "valid" method and number, was really about differences in temporal scale.

As I have continued to work on uncertainty I have the distinct impression it makes one a better thinker and synthesizer. For one thing, as indicated above it changes your perspective in discussions onto a more fruitful path. But it also influences how you view synthesis. Had I been asked 20 years ago what contributed to uncertainty in ecosystem analysis I suppose I would have said "everything". While that's probably the correct answer, it is not very useful. If you asked me now I would list the four major sources: measurements, sampling, model prediction, and model selection. I would also point out the critical importance of how the system is wired together, for example, the degree and sign of correlations, the size of the contributing parts (i.e., high uncertainty of a small part often is swamped out by that of larger parts) and the order in which they contribute. In short I would focus on the system and how it goes together, which is after all the heart of synthesis.

I do not want to give the impression that the only thing one gains by analyzing uncertainty are new skills that allow one to become a super-techno-mensch. I believe that one can gain science insights as well. For example, I had reason to examine the uncertainly effect of some empirically derived parameters for models that predicted the accumulation of live, dead legacy, and dead de novo wood (i.e., the dead wood created by the new forest) after disturbance. During this exercise I found that the uncertainty in the overall carbon balance was not equal over time, that the envelope of responses did not simply parallel the accumulation mean. For two periods, the time when the balance goes from a source to the atmosphere to a sink and the time in which the long-term balance starts to approach zero varied less than 5 years. Interestingly, the time of the maximum sink value could be as early as 25 years and as late as 75 years, a span of 50 years. In exploring the source of this phenomenon I discovered that it was not due to the growth rate of live trees, nor was it due to the amount of legacy wood. Both these influence the timing of the switch from a sink to a source, but not the timing of the peak sink. That variation is due to uncertainty about the timing of the de novo wood accumulation. This makes sense in that this accumulation is dependent on mortality and mortality is highly variable. This leads me to the following observation: we often view uncertainty as a demerit to our knowledge, but I now believe it adds to it, However, to allow it to merit to that knowledge, we may need to shift from the perspective that uncertainty lets us see less (i.e., the result is obscured in uncertainty) to letting us see more (i.e., the full result revealed by uncertainty).

As bullish as I am on quantifying uncertainty, I do have some concerns. I think we are rapidly passing into a phase in which uncertainty will be acknowledged and quantified. If used properly I think this will improve ecosystem science and thinking. Can uncertainty be used improperly? Unfortunately the answer is yes. First, it can be weaponized. Weaponized uncertainty flourishes when true nature of modern science is misunderstood. Of course the goal of modern science is to create a mechanistic understanding, but that knowledge for many scientific fields such as ecology is inherently imperfect and uncertain. So while I understand that we should be certain about the science, the fact is that much of what we know now and will know in the future will remain uncertain at some level. This misunderstanding of uncertainty and science (the certain) can be used as a weapon. We see this in today's climate debate: action is not possible until the effects are certain. Of course since there is always some form of uncertainty in effects it follows that action is never possible. While I understand some will always use weaponized uncertainty as a policy strategy, we can dull its effect to some extent by not attaching value too uncertainty. That is, while it is good to have low to no uncertainty, the level of uncertainty is what it is and it is not good or bad per se. Otherwise systems with inherently high natural variation will be worse than ones with low natural variation. If this notion seems far-fetched a colleague once explained to me how ecological systems are inherently "bad". The reasoning was that they have high natural variability. I am not sure what good versus bad has to do with natural variability; to me natural variability makes systems interesting. I suspect he was really trying to convince me his science was good via weaponized uncertainty.

It is not only the policy arena that can be limited by misuse of uncertainty; it can impact science. I recently experienced weaponized uncertainty in a discussion about Oregon's forest carbon stores. One contributor felt that belowground live stores should not be included because of the high degree of uncertainty in the model used to predict these stores. The model is quite well known, a simple ratio of above- and belowground live stores, but that ratio ranges between 15 and 30% leading to a two-fold range in live belowground stores. From this perspective exclusion of this pool seems like a rational thing to do. However, consider the following: while the parameters for the models predicting aboveground stores are well characterized, there is a range of models to select among, and for some well sampled species the variation of model predictions for a given diameter approaches a factor of two. Therefore the plan to exclude live belowground carbon from the analysis does not makes sense because it means a pool that we know exists and has pretty much the same overall uncertainty as one we "trust" (live aboveground) will be neglected. To me this is the consequence of seeing uncertainty as a "problem" rather than just a metric of the current level of understanding. If uncertainty is the problem, then one way to get rid of the problem is to get rid of the culprit, in this case live belowground parts. In my opinion we would be far better off just reporting the uncertainty and then take future steps to actually versus artificially reduce it if we judge the current level to be too high.

Another concern is that uncertainty estimation can become an undue burden. This would have the effect of slowing down and in some cases obscuring the underlying big picture science. No doubt that the former was, in part, an early motivation for not analyzing uncertainty—it would take too much time. My concern is mirrored in some of what I see in terms of the current use of statistics in ecology. I freely admit I am past my sell-by date, but many ecological analyses in my opinion have gone into a zone of statistical overkill. For example, why is it more important to publish a table of F statistics than a table summarizing the actual data? Is not the data part of the story? Or is the story just the analysis? Do not get me wrong: I am not opposed to statistics, but I do have to wonder if the balance between analyzing data completely with statistical rigor versus understanding processes and controls swung too far toward the former. To avoid uncertainty overkill I am advocating the view that quantifying each and every aspect of uncertainty of an analysis is not always necessary. Instead, the focus should be on the key sources for given systems. Therefore an important investment will be to determine the general magnitude of sources for various systems as an initial guide to where the emphasis should be placed. If a source is relatively small or well characterized, then it can be addressed without expending undue effort. If an uncertainty source is large, then it should be addressed explicitly. The point is that we need to be efficient and effective, recognizing there are costs as well as benefits.

My final concern is about hammers. It is related to the old adage "if you only have a hammer, everything looks like a nail." Key tools for uncertainty clearly include statistics and probability. I do not see how uncertainty estimates could not employ these tools, but there are aspects of uncertainty, particularly that related to model selection (or as it sometimes described structural uncertainty) that cannot be easily addressed with these tools. And since these aspects of uncertainty cannot be addressed with the available tools, they currently tend to be ignored (hence the hammer analogy). Which is very unfortunate because the heart of science is the discrimination of different models (or as we usually refer to them hypotheses.) My overall concern is that if we do not address model selection we are not harnessing the full power of science. To address this shortcoming we need to find and use the appropriate analysis tools. Moreover, we need to change our perspective from only running the "best" model to running all the viable alternative models. If the uncertainty caused by model selection is unacceptably high, then either a more generalized model needs to be developed or a method to discriminate between the models needs to be devised. I acknowledge that this recommendation is collides with my previous concern; however, model selection is crucial to scientific advancement and potentially has high enough stakes to be worth the effort.

Let me conclude by noting that there are many reasons for quantifying uncertainty in ecosystem analyzes. I agree with the principle that we have some obligation to provide these estimates, but there are other more compelling motivations. These include: 1) learning more about estimates because the variation in estimates is important information; 2) understanding how systems are put together which increases our synthetic understanding; 3) as a tool for more rigorous hypothesis testing; 4) as a way to compare estimates derived with different methods or those derived for different places, and times; 5) and as a way to harness the scientific method more completely. I believe that these motivations all fall under the rubric of discovery. What could be a better motivation for scientists?