

# If It's Worth Doing, It's Worth Doing With Made-Up Statistics

*Posted on May 2, 2013 by Scott Alexander*



I do not believe that the utility weights I worked on last week – the ones that say living in North Korea is 37% as good as living in the First World – are objectively correct or correspond to any sort of natural category. So why do I find them so interesting?

A few weeks ago I got to go to a free CFAR tutorial (you can hear about these kinds of things by [signing up for their newsletter](#)). During this particular tutorial, Julia tried to explain Bayes' Theorem to some, er, rationality virgins. I record a heavily-edited-to-avoid-recognizable-details memory of the conversation below:

**Julia:** So let's try an example. Suppose there's a five percent chance per month your computer breaks down. In that case...

**Student:** Whoa. Hold on here. That's not the chance my computer will break down.

**Julia:** No? Well, what do you think the chance is?

**Student:** Who knows? It might happen, or it might not.

**Julia:** Right, but can you turn that into a number?

**Student:** No. I have no idea whether my computer will break.

I'd be making the number up.

**Julia:** Well, in a sense, yes. But you'd be communicating some information. A 1% chance your computer will break down is very different from a 99% chance.

**Student:** I don't know the future. Why do you want to me to pretend I do?

**Julia:** (*who is heroically nice and patient*) Okay, let's back up. Suppose you buy a sandwich. Is the sandwich probably poisoned, or probably not poisoned?

**Student:** Exactly which sandwich are we talking about here?

In the context of a lesson on probability, this is a problem I think most people would be able to avoid. But the student's attitude, the one that rejects hokey quantification of things we don't actually know how to quantify, is a pretty common one. And it informs a lot of the objections to utilitarianism – the problem of quantifying exactly how bad North Korea shares some of the pitfalls of quantifying exactly how likely your computer is to break (for example, “we are kind of making this number up” is a pitfall).

The explanation that Julia and I tried to give the other student was that imperfect information still beats zero information. Even if the number “five percent” was made up (suppose that this is a new kind of computer being used in a new way that cannot be easily compared to longevity data for previous computers) it encodes our knowledge that computers are unlikely to break in any given month. Even if we are wrong by a very large amount (let's say we're off by a factor of four and the real number is 20%), if the insight we encoded into the number is sane we're still doing better than giving

no information at all (maybe model this as a random number generator which chooses anything from 0 – 100?)

This is part of why I respect utilitarianism. Sure, the actual badness of North Korea may not be exactly 37%. But it's probably not twice as good as living in the First World. Or even 90% as good. But it's probably not two hundred times worse than death either. There is definitely nonzero information transfer going on here.

But the typical opponents of utilitarianism have a much stronger point than the guy at the CFAR class. They're not arguing that utilitarianism fails to outperform zero information, they're arguing that it fails to outperform our natural intuitive ways of looking at things, the one where you just think "North Korea? Sounds awful. The people there deserve our sympathy."

Remember the [Bayes mammogram problem](#) ? The correct answer is 7.8%; most doctors (and others) intuitively feel like the answer should be about 80%. So doctors – who are specifically trained in having good intuitive judgment about diseases – are wrong by an order of magnitude. And it "only" being

one

order of magnitude is not to the doctors' credit: by changing the numbers in the problem we can make doctors' answers as wrong as we want.

So the doctors probably would be better off explicitly doing the Bayesian calculation. But suppose some doctor's internet is down (you have NO IDEA how much doctors secretly rely on the Internet) and she can't remember the prevalence of breast cancer. If the doctor thinks her guess will be off by less than an order of magnitude, then making up a number and plugging it into Bayes will be more accurate than just using a gut feeling about how likely the test is to work. Even making up numbers based on basic knowledge like "Most women do not have breast cancer at any given time" might be enough to make Bayes Theorem outperform intuitive decision-making in many cases.

And a *lot* of intuitive decisions are off by way more than the make-up-numbers ability is likely to be off by. Remember [that scope insensitivity experiment](#) where people were willing to spend about the same amount of money to save 2,000 birds as 200,000 birds? And the experiment where people are willing to work harder to save one impoverished child than fifty impoverished children? And the one where judges give criminals several times more severe punishments on average just before they eat lunch than just after they eat lunch?

And it's not just neutral biases. We've all seen people who approve wars under Republican presidents but are *horrified* by the injustice and atrocity of wars under Democratic presidents, even if it's just the same war that carried over to a different administration. If we forced them to stick a number on the amount of suffering caused by war before they knew what the question was going to be, that's a bit harder.

Thus is it written: “It’s easy to lie with statistics, but it’s easier to lie without them.”

Some things work okay on System 1 reasoning. Other things work badly. Really really badly. Factor of a hundred badly, if you count the bird experiment.

It’s hard to make a mistake in calculating the utility of living in North Korea that’s off by a factor of *a hundred*. It’s hard to come up with values that make a war suddenly become okay/abominable when the President changes parties.

Even if your data is completely made up, the way the 5% chance of breaking your computer was made up, the fact that you can apply normal non-made-up arithmetic to these made-up numbers will mean that you will very often *still* be less wrong than if you had used your considered and thoughtful and phronetic opinion.

On the other hand, it’s pretty easy to accidentally Pascal’s Mug yourself into giving everything you own to a crazy cult, which System 1 is good at avoiding. So it’s nice to have data from both systems.

In cases where we really don’t know what we’re doing, like utilitarianism, one can still make System 1 decisions, but making them with the System 2 data in front of you can change your mind. Like “Yes, do whatever you want here, just be aware that X causes two thousand people to die and Y causes twenty people an amount of

pain which, in experiments, was rated about as bad as a stubbed toe”.

And cases where we don't really know what we're doing have a wonderful habit of developing into cases where we *do* know what we're doing. Like in medicine, people started out with “doctors' clinical judgment obviously trumps everything, but just in case some doctors forgot to order clinical judgment, let's make some toy algorithms”. And then people got better and better at crunching numbers and now there are cases where doctors [should never](#) use their clinical judgment under any circumstances. I can't find the article right now, but there are even cases where doctors armed with clinical algorithms consistently do worse than clinical algorithms without doctors. So it looks like at some point the diagnostic algorithm people figured out what they were doing.

I generally support applying made-up models to pretty much any problem possible, just to notice where our intuitions are going wrong and to get a second opinion from a process that has no common sense but is also lacks systematic bias (or else has unpredictable, different systematic bias).

This is why I'm disappointed that no one has ever tried expanding the QALY concept to things outside health care before. It's not that I think it will work. It's that I think it will fail to work in a different way than our naive opinions fail to work, and we might learn something from it.

**EDIT:** Edited to include some examples from the comments. I also really like ciphergoth's quote: "Sometimes pulling numbers out of your arse and using them to make a decision is better than pulling a decision out of your arse."