**Harvard Data Science Review • Special Issue 1: COVID-19:
Unprecedented Challenges and Chances**

A Conversation on COVID-19 With Head of Statistics for BBC News Robert Cuffe

Robert Cuffe¹ Liberty Vittert^{2,3} Xiao-Li Meng^{4,5}

¹Visual Journalism Department, BBC News, Longon, United Kingdom,

²Harvard Data Science Initiative, Harvard University, Cambridge, Massachusetts, United States of America,

³John M. Olin Business School, Washington University in St. Louis, St. Louis, Missouri, United States of America,

⁴Department of Statistics, Faculty of Arts and Sciences, Harvard University, Cambridge, Massachusetts, United States of America,

⁵Harvard Data Science Review, Harvard Data Science Initiative, Harvard University, Cambridge, Massachusetts, United States of America

Published on: Aug 21, 2020

DOI: <https://doi.org/10.1162/99608f92.ca11a017>

License: [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

ABSTRACT

Editor-in-Chief Xiao-Li Meng and Media Feature Editor Liberty Vittert conducted a virtual dialogue on May 18, 2020, with Robert Cuffe, Head of Statistics for *BBC News* for an in-depth conversation about the challenges of communicating accurate statistics concerning COVID-19 to the general public. The three discussed the importance of data accuracy when trying to compare one country's death rate to another, the regional differences in death rates and in how deaths are counted, and the difficulties in estimating the distribution of infection on a regional (country, city, or state) basis.

HDSR includes both an audio recording and written transcript of the interview below. The transcript that appears below has been edited for purposes of grammar and clarity.

Liberty Vittert (LV): Robert Cuffe, the head of statistics for *BBC News*. Robert, thanks for joining us. It seems like you have a lot on your plate right now.

Robert Cuffe (RC): Yeah, I'm not the only one. Work is busy.

LV: Thank you for joining us. I'll just dive right in. We both live in the two countries that seem to have the most coronavirus deaths: the U.S., and the U.K. right behind it. We've heard a lot about different countries and how you compare the deaths between them to see who's doing better and who's doing worse. Can we really do that?

RC: I think everyone's tempted to draw up that league table [comparison table] and to compare countries on the basis of one number, be that the number of deaths, in which case the U.S. doesn't look too great, or the number of deaths per head of population, in which case an enormous country like the U.S. doesn't look quite so bad. Those one-number league tables are not a helpful way to go, because there are so many differences between countries and so many differences in the way that countries count deaths. You really need to understand all those differences before you start to make comparisons.

LV: What are the issues or the differences in data collection between countries when it comes to deaths?

RC: It depends on who you're counting. That's a big one. The metric I prefer to look at for understanding the total death toll is the total number of deaths or the excess mortality. As in, how many more deaths we're seeing this week compared to what we'd expect to see. In previous weeks, that number used to exceed, by far, the number of deaths from the people who tested positive or even the number of people who had COVID-19

mentioned on their death certificate. You had this big excess—this huge, big spike. Some of it was accounted for by diagnosis, but a good chunk of it wasn't. These were in the early days of lockdown, where you're saying, 'These people, they're not the victims of the economic consequences of the lockdown and the impact that it will have on their life chances.' That probably was underdiagnosis. We're starting to see now, in recent weeks, the gap between those two numbers is narrowing, that almost all and sometimes more than all of the excess mortality is COVID registration. You do have that changing pattern over time.

LV: What would be the way that we actually could compare between countries?

RC: The fact that comparisons are difficult doesn't mean that you shouldn't do them. You just need to be aware of some of the limitations. When I say that there are definitional differences, be aware of them. If you've got a signal that is smaller than the likely noise—and you would regard these definitional differences as noise—then don't build your league table. But there are some clear, standout, shining cases where the differences are not just about definitions. There is a reason why South Korea has had such a low death rate. There's a reason why Germany has had such a low death rate. There's a big commonality between the two of them in terms of the vigor with which they've been able to pursue the early testing and tracing.

LV: I have a question for both Xiao-Li, as my favorite statistician, and Robert.

Xiao-Li Meng (XLM): Thank you.

LV: There's an [article](#) in *The Washington Post* in which a leading [epidemiologist](#) at Boston University said, "State and local leaders should be studying the estimates of excess deaths in their communities and basing consequential decisions about reopening businesses and social activities on those figures." To me, that presents to you that excess mortality is a good way to compare. But to actually make decisions based on it—I see a lot of problems with that in terms of what you're attributing deaths to, what their mental health issues are, and all the different issues that would come with this. What do you both think as statisticians? Is excess mortality a number you can just use to compare? Or is it a number that you could actually use to make decisions?

RC: Excess mortality is enormously helpful for understanding how bad things are. It's certainly the number that we're trying to minimize, or that government action is trying to minimize. But, in and of itself, it doesn't help you to decide between the things you're trying to balance. The extra deaths we're seeing are above and beyond what we'd expect to see. They're a combination of not only the deaths that are attributable directly to the virus, or the biological deaths, where people are infected and died. Excess deaths include the people who've died directly biologically after infection from the virus, but they also include the people who have died because of the strain the virus puts on our society, on our emergency rooms, and on our intensive care units. It also will eventually contain the victims of the measures that we take to try and control the virus—the people who have a tough time while locked down or the people who become unemployed as a result of the economic consequences. That headline number, because it includes the two things we're trying to balance off against each

other, doesn't necessarily give us guidance. You need to go into the details of the causes of the deaths before you can understand what policy decisions you can make in order to trade off these two things.

XML: I want to say that I cannot agree more and I cannot disagree less. For me, I have two reactions. One is, as Robert just said, if you make a decision, it's all about long-term consequences. Looking at the excess number obviously at this moment is one important piece of information. But what you really want to know about is that if you make this decision, what would be the consequence? What would be the actual deaths later? What would be the quality of life? What would be many, many other things? I think that the excess number should definitely be a part of it, but to decide purely on it would be quite a mistake.

I also want to pick up on one more thing, which as statisticians, we all understand, but in the general public people may or may not understand why it's important: Your decisions shouldn't be just based on your own number. You should at least base them on your neighbors' numbers, and you should try to borrow more information. Such as with Bayesian shrinkage estimations, obviously you need a lot more data. Given the quality of the data, trying to decide just based on your own numbers is probably not really a good idea, statistically speaking. I know that's a hard point to get across: why other people's numbers matter. It's all a part of a pattern for us to estimate it better.

LV: Excess mortality is not the only number you look at when you're trying to understand when shutdown could happen or how well countries are doing. One thing we've seen a lot of, especially in the U.S., is that because there are multiple states, the different states have very different rules and are making very different decisions. We've seen almost this social experiment in how different people are reacting to different orders, different laws, and different policies. In Georgia, for example, they decided to open up a while ago and were the first to do it. You would see all these news stories a couple of days after they opened up, saying, 'Oh, there's this huge spike in the number of cases two days after they opened, horrors!' or, 'There are so many more deaths three days after they opened up.' From my understanding, that isn't really a fair way to measure it either. Do you have a time period that you would think to look at?

RC: If you're trying to understand the immediate consequences in terms of the direct effect on the virus after opening up, it's probably going to take you a couple of weeks before that starts to filter through, because it takes a while for somebody to get infected and then to develop symptoms. That's five days already, and then maybe another week or even longer before they start to get sick enough to need to go to the hospital. They might spend a couple of days in the hospital as well. Then, of course, if things go wrong or they don't recover, they do, in fact, die. It will take a while for that death to be recorded. You're talking about three or four weeks before you start to see the effect of a change in lockdown show up in the figures. You get some interim measures along the way—like the number of people who are reporting symptoms—if you've got good surveillance like population-wide surveys where you're randomly sampling people and asking them to swab themselves. Then you can get much faster information, but that's a huge undertaking. It does take a while before you start to see the effects of changes to lockdown show up into those clinical measures. If we're

moving on to the economic component, that takes even longer. That's a long old job, before you start to understand that stuff.

XLM: As you said, the reporting is typically delayed and depends on what kind of surveillance system you have. Years ago, I actually did a study on the CDC reporting delays on the AIDS epidemic. What we found, which statisticians and others would expect, is that the amount of reporting delay was related to the kinds of health care conditions in that area or in the system. There is a confounding factor there. If everybody delays by four weeks, we know how to make an adjustment. It makes things even more complicated if the delay itself is a function of all these complicated factors. I wonder if, in this current pandemic, you guys have noticed such issues in how you are trying to report this with these kinds of very nuanced conditions, or you do say that it's just too complicated to communicate?

RC: Thankfully, it's not my job to try to do this stuff from scratch because of all of the different factors that you're trying to come by. You're absolutely right. You can make a simplifying assumption and say it takes five days from infection to symptoms. But that's not the case for everyone. Of course, it's not. It takes so many days for reporting to happen. Of course, that's not the case. You've got distribution on everything. My job, thankfully, is not to try to work that out. That's part of the jobs of far smarter people than me. But that is being taken into account in some of the estimations. You do get a rough sense when you look at the outputs from the models of whether or not they're taking regional differences into account. Those are only one part of those kinds of statistical elements of measuring. If you're trying to build a model for what's happening to infections based purely on deaths, well, the infection rate is not the only thing that determines who dies. There are regional differences in death rates and they are driven by ethnicity, socioeconomic status, and lots of other things as well. Unless you account for those pieces as well, you're going to make mistakes in your estimation of the distribution of infection. I think that's something that infectious disease modelers need to take into account as well. Is the model applying the data well? Which pieces of data? It's not just the data that you use to build up the model in the first place, it's how you check in with the real world as well afterward.

XLM: On that note, Liberty, if you don't mind, I want to follow up with another question, because I'm inspired by what Rob just said. By the way, I love what you just said about how you've got distribution on everything. I'm going to steal that quote. That's a great quote about everything has variability. But the question I have is—it's an obvious one as a statistician—you see these numbers being reported all over the place, but people tend to report just one number. You don't report an uncertainty assessment, which for statisticians is everything. I absolutely hated when people just gave me one number to ask me to make a decision during those days when I was dean. I would say, "Give me two numbers," but nobody gave me two numbers. I understand that when you're trying to report, say a confidence interval, usually, the ironic thing is when you report a confidence interval, you project to the general public that you don't have much confidence in what you're reporting because you give two numbers. How do you deal with that issue? How do you convey to the general public

that these numbers are useful, but they come with such a grain of salt because of all the complications? What are the things you do to make that happen?

RC: When you said you only wanted to hear two numbers, was that the top or the bottom end of the confidence interval?

XLM: I actually shouldn't have said just two numbers. [Laugh] I should have said I wanted to see the whole histogram.

RC: When you're telling a story to an audience, you need to make decisions about what they will read and what is useful for them to hear. It comes back to that question about the signal and the noise. It's important to tell people about the noise if it changes a story.

If unemployment has gone down by about 15,000 people and the margin of error on that estimation in the country you're looking at is plus or minus 80,000, that's a very different story than if the margin of error is plus or minus two or three thousand. In lot of cases where there is a really big change, you don't talk about the margin of error. Who needs three numbers? The statisticians may love it, but you're dealing with the lay audience, most of whom are not, let's face it, nerds. You want to provide stuff that they will find interesting because they don't have to read your information. They could be reading stuff about "Love Island" or anything else. Unless you make it worth reading, people won't bother.

In the *BBC*, we're very careful to make sure that we're putting out factual information. We double source material. Our information is not wrong, but that doesn't mean we need to show all of the uncertainty all of the time in every story that we do. We're not parliamentary researchers for MPs (Members of Parliament). They have people who do that stuff for them. If you have advisers who should be doing that work, it's not necessarily our job to do it. Of course, we don't want to make mistakes, but I think it's a different question. We have different audiences that we're aiming at.

LV: When you have this job, to tell stories and to also decide, in a sense, what people should hear, there are really tough conversations that are surrounding coronavirus. This whole idea of economic shutdown versus lives lost and how you make those measurements is an unbelievably emotional and difficult conversation, but it's something that governments do all the time. The Department of Transportation, for example, in the U.S., actually has a number of how much a life is worth when they're deciding about the cost of implementing new safety regulations. Life insurance companies do it all the time. This is something that happens every day. While it's a really tough conversation, we still have to have it, or at least I think so. My question is, how do you all have that conversation of saying, "This is how many lives we'll save, we think, by doing this lockdown or these very difficult economic restrictions, this is how much money will be lost, and this is how we're deciding how long to do it for"?

RC: I think it's even more complicated than trading off the cost of the lockdown against the individual lives that you save. The first problem is the number of lives that you save is counterfactual. You just have no idea what the number actually is. You know how many people are dying. But you don't know how many people would have died if you hadn't locked down.

LV: Exactly.

RC: There's much that we don't know about this virus and about what's going to happen to it over the next few months. We're still learning from many other countries. You see a very different response in Sweden where the lockdown has been societally mediated. It hasn't really been run by government dictates. But before you get into those comparisons, as we were talking about at the start, you need to be aware of the differences between countries, not just in how they count deaths, but also in the hand that they've been dealt. New York is an incredibly densely populated area that's just been hit so hard by the virus. We've seen that in London. In the U.K., that's the hardest hit part of the country, where you've got these massive transient populations. All of those kinds of population density factors and the risk of dying factors—all those things contribute. There isn't a model for that right now that I'm aware of. What you're actually doing is you're incorporating not just understanding of the virus, but you've got to project forward and say, 'The cost of this lockdown is not just the financial cost, but also the damage it's going to do to the economy in the long run and the damage that it will do to individuals in the long run and to their health.' On top of that, you've got to add the cost of the measures that you take to mitigate the lockdown.

In the U.K., the government has implemented a scheme whereby it pays about 80 percent of the salary of employees up to a certain amount so that employers don't lay them off. They can keep them in the job and they can hope to get through until the lockdown comes off. The hope of doing that is you don't have too much scarring of the economy. You're able to bounce back because the people are still able to work and they're not unemployed. That scheme has a cost as well, and hopefully, that is a benefit in the long term. By the time you start to put all those things together, then my head is hurting. It's an enormously complicated problem. I hope that the people who are making these decisions are flying on the basis of the best models that they can conduct. But there's no perfect model for this.

LV: We're talking about models now. What do we have a model for? What don't we have a model for? There are multiple models from all different sources and all different institutions and groups saying how many people would die if we didn't do anything and how many people will die because of the shutdown. There are enormous differences in these models and there is enormous criticism on all sides about these models. How do you decide what models to report? How do you decide what models not to report? Have there been things that you feel that you would have done differently over the past couple of months? Are there things that you would rather have done based upon all these different types of models?

RC: I think we've become increasingly skeptical of models that are basically statistical models— models that just look at the patterns in the data and are not based in a mechanistic understanding of what's going on with the virus. No model is perfect and the models that try to piece together all the different facts we know about how viruses work and apply them to the coronavirus are like trying to put together a jigsaw with lots of pieces from different boxes. They're really difficult and they're very fragile, or I've been told that they're very fragile by people who know more about it than I do. But they still seem to be more satisfying than the people who just look at the second derivative (the number of deaths that we're seeing in a country) and then try and apply that globally as some kind of law, or the models that just use a 'Here's the date of lockdown and all these 18 different countries' without any understanding of how lockdown was implemented and what else had happened. Or 'Here was a pattern of deaths they were following up until that point. Let's just project it forward.' Those kinds of projections we've become less and less trusting of. We've seen a few of those come out and those projections have tended to be revised pretty quickly within a couple of days. That doesn't mean that we're saying all the models that led to all the decisions that have been made so far are perfect. But there seems to be more underneath them.

LV: Right now, we have leaders of multiple different countries saying that we will have a vaccine most likely by the fall. I also see these news stories that say we have never approved a vaccine for a coronavirus. How do we determine whether that means they make the approval for a vaccine less stringent or more stringent? How much hope should we really have as the general public for a vaccine, statistically speaking?

XML: I can tell you, based on the articles that we have received (some of which have already come out, others that will come out later) that I think all experts agree there will be at least one more year and possibly 18 months until we get a vaccine. That's the very conservative estimate. I don't know how others will think the vaccine will be available by the fall. It's going to be a long process, partly because you do have to test them with clinical trials, which is obviously time-consuming. We want to make sure that these vaccines come out safe and effective. You may remember Jeremy, the previous guest, saying something about how it would be disastrous if the vaccine turns out not to be ineffective. It's a much worse outcome than some safety trade-off. To think about the release as happening this fall is being seemingly unrealistically optimistic based on what we have seen. The other part of the question you have is, "Can we speed up the approval process?" One of the articles was talking about how the FDA (the Food and Drug Administration in the U.S.) can think differently about the trade-off between a false positive and a false negative because that is usually how they decide those things. During the pandemic, you can make a different trade-off. There are suggestions that things can be speed up, but with theoretically understood trade-offs, and how to control these trade-offs. We can't speed it up just because we want the vaccine. You do have to make sure there is a scientific process involved.

RC: Hopefully, it'll get harder and harder to research a vaccine as incidence is falling. It takes longer and longer to test a vaccine with fewer cases, so that would be great news if that vaccine never arrived because even though it was effective, we didn't have enough people to test it on and we never got a signal.

Unfortunately, I suspect that's not going to be the case. What we know in the U.K., where the government has started to do random sampling of the population to see how many people have the virus, is that about one in 400 people in England has the virus at the moment. That's going down at the moment. There are fewer and fewer people who have the virus every week. If you think about how many people you need to study in a vaccine trial in order to get a good chunk of people who have the infection in the control group and a good chunk of people who would have had the infection but didn't get it because they were vaccinated, that means you need to run a really large trial and it becomes harder and harder to do that as the virus comes down and down and down. That's the hope. The hope will be, of course, that the virus gets so far down that we're never able to conduct a vaccine trial. But I think that's very unlikely.

XML: That itself is a part of the difficulty with the clinical trials. That's the whole process.

LV: Robert, thank you so much for talking to us. Thanks for all the good work you're doing at the *BBC*, really helping statistics be communicated correctly.

RC: Thanks very much. Pleasure talking to you guys.

LV: I suppose this is almost a case of no data is good news, in that if there isn't a vaccine able to be made, it's because the cases have gone down so much. But unfortunately, it doesn't seem like that will be the situation here.

XML: Unfortunately, I have to agree with you on that. On the other hand, it's a very interesting way of thinking about data. It shows, again, the vast scope of the landscape of the data. And this is a production of the *Harvard Data Science Review*. Thank you for listening.

Disclosure Statement

Robert Cuffe, Liberty Vittert, and Xiao-Li Meng have no financial or non-financial disclosures to share for this article.

©2020 Robert Cuffe, Liberty Vittert, and Xiao-Li Meng. This interview is licensed under a Creative Commons Attribution (CC BY 4.0) [International license](https://creativecommons.org/licenses/by/4.0/), except where otherwise indicated with respect to particular material included in the interview.