A Principled Framework for Disclosure Avoidance

January 28, 2025

>> Coordinator: Welcome, and thank you for standing by. Currently, all participants are in listen-only mode. Today's webinar is being recorded, and the recording will be posted publicly. If you have any objections, you may disconnect at this time. Now, I'd like to turn it over to your host for today, Shelly Hedrick.

>> Shelly Hedrick: Thank you, Lisa, and welcome, everybody. Thank you for joining today's webinar about the Census Bureau's principled framework for disclosure avoidance. I'm Shelly Hedrick, and I'm from the Communications Directorate. And I'm joined today by two subject matter experts, Michael Hawes, who's a senior statistician for scientific communication in our Research and Methodology Directorate, and Luke Rogers, and he's a senior research scientist in our Population Division.

So, just some housekeeping to start. Throughout this session, you'll be able to post questions in the Q&A window that's in the Webex browser. It's at the bottom right of your screen. Our presenters will be answering what they can in that Q&A window, and they'll also address questions at the end of their presentations. Now, if we can't get to your question today or if you have others, we'll provide some contact information for you to connect with us afterward. Also, we'll be posting these slides, as well as a recording and a transcript, to our website in two to three weeks.

We do have a dedicated webinar series on disclosure avoidance topics that are in the Census Academy section of our census.gov website. And I'll be adding a link to the chat where you can find that, as well as the slides, recordings, and transcript once those become available. So as it happens, today is International Data Privacy Day. So, for all who celebrate, happy Data Privacy Day. And what better way to celebrate than by hearing from our first presenter, Michael Hawes. Michael, go ahead and take it away.

>> Michael Hawes: Great. Thank you, Shelly, and good afternoon, everyone, or to those from points further west, good morning. The Census Bureau has a strong organizational culture of data stewardship. Our ability to safeguard the information provided to us for our censuses and surveys is central to our ability to keep the public's trust. It's also a legal requirement that we must follow. Title 13 of the U.S. Code prohibits us from disclosing the information that's reported by or on behalf of any of our respondents. It also prohibits us from releasing any information that could permit a respondent to be identified. This means using various statistical safeguards to make sure that the products we release don't inadvertently reveal enough information to allow specific individuals and their records to be singled out. We call this process disclosure avoidance.

Now, the challenge we face is that the disclosure risks we're trying to protect against are not static. Much like the ever-growing challenges that organizations face when defending against IT security and hacking threats, statistical agencies have to contend with a

constantly changing environment where existing disclosure threats increase and new risks arise on a continuous basis.

Two of the most vexing of these evolving threats come from the explosion of third-party data sources that can be used as part of an attack on our published statistical products, and the ever-increasing power of computers and complex optimization algorithms that can leverage these data in such an attack.

Maintaining the public's trust and ensuring our legal obligations to protect confidentiality means adapting our disclosure avoidance methods to keep pace with these rising threats. In fact, the Census Bureau has a long history of continuous improvement and innovation of our privacy protections. One recent focus of our efforts has been the strengthening and modernization of our disclosure avoidance methods specifically to address these new 21st-century threats.

Now, it's important to understand that there are many different disclosure avoidance options that agencies can consider, and new techniques and approaches are being developed all the time. Because many of these statistical methods protect against disclosure in very different ways, it can often be quite challenging to compare and contrast them on an even footing.

As the Census Bureau looks ahead to some of the disclosure avoidance decisions we will have to make over the coming years, we realize that we need a scientifically principled way to look at the strengths and weaknesses of these very different disclosure avoidance options in order to properly evaluate between them.

To that end, we established a Census Bureau-wide team to develop a framework of principles that we, as an agency, could use to enable and to guide these evaluations of different disclosure avoidance options. In putting together this Disclosure Avoidance Futures Working Group, or DA Futures for short, we understood that these principles shouldn't be limited in context to any specific census or survey. Rather, they should be a universal set of principles that we can use across the agency, to inform disclosure avoidance decision-making for all of our statistical products.

To that end, we made sure that the working group had expertise representative of the breadth of our information collections and statistical activities and of the myriad important ways that the nation uses our statistical products. Ensuring this breadth of expertise required the active participation of subject-matter experts from our economic, demographic, decennial, communications, and research, and methodology directorates.

This effort culminated in a consensus framework of disclosure avoidance principles that the entire team could comfortably stand behind. We have documented this framework in a working paper that's available on our website entitled Towards a Principled Framework for Disclosure Avoidance. And we're excited to be giving you a high-level walkthrough of these principles today.

At the heart of the disclosure avoidance challenge we face is a set of constraints that we call the triple tradeoff of official statistics. Every statistic that you publish that is derived from a

confidential data source is going to reveal or leak a tiny amount of confidential information. Consequently, the more statistics you publish and the greater the granularity and accuracy of those statistics, the greater the resulting risk of disclosure.

Conversely, every statistical technique that you can use to protect confidentiality is going to degrade the resulting statistics in some manner, either by decreasing the availability of the statistics or by decreasing their accuracy. And this isn't a side effect; it's how they protect confidentiality. Therefore, there's an inherent tradeoff that we, as a statistical agency, have to navigate between the degree of protection and the resulting availability and accuracy of the statistics to be released.

At the Census Bureau, these difficult choices are made by the Data Stewardship Executive Policy Committee, or DSEP, who have to decide things like how much disclosure risk is too much or how accurate do these statistics need to be to meet their intended uses. These are difficult decisions, but they aren't made in a vacuum. You'll notice at the center of the triple tradeoff graphic here, we have scientific integrity and stakeholder engagement. These reflect the fact that DSEP makes these difficult choices within the context of the Census Bureau's statistical quality standards, consistent with the best available research and science and, perhaps most importantly, in a manner that's informed by and supported by the invaluable feedback and perspectives of our stakeholders.

At the end of the day, however, this triple tradeoff imposes inescapable constraints on agency statistical products. Any increase along one dimension requires a reduction on one or both of the other dimensions. Increasing the granularity of your published statistics, for example, would require either reducing the accuracy of those statistics or weakening your confidentiality protections, or some combination thereof. Improving the accuracy of the statistics you release would similarly require reducing the availability of those statistics or reducing the confidentiality protections or both. There are many different ways to balance across these three dimensions, but they all have consequences for the resulting statistical product.

With this triple tradeoff in mind, it's important to note that the selection of any particular disclosure avoidance technique, like picking suppression over rounding or picking swapping over differential privacy, does not directly impact agency decision-making within the context of the triple tradeoff. Nearly any disclosure avoidance technique can be applied to implement very different balances along these three dimensions, depending on the implementation parameters selected. In this way, the selection of the particular disclosure avoidance method, like rounding over swapping, should not be confused with decisions about how strong or how weak the implementation parameters of that method should be under the triple tradeoff, like rounding to the nearest three versus rounding to the nearest 1,000.

So, to establish a framework for objectively evaluating between different disclosure avoidance methods, we need a set of overarching principles or characteristics that an ideal applied disclosure avoidance system should meet, while distinguishing those characteristics from any choices relating to the implementation of that system. Now, I'll explain in a moment what I mean by an ideal applied disclosure avoidance system. But first,

I want to reiterate my point from a moment ago, that the exact same disclosure avoidance system can provide a high degree of protection with low accuracy or a minimal degree of protection with high accuracy for any given set of statistical outputs. Finding the proper balance is a choice of the agency that should then be implemented by the system. It's not an inherent feature of the system itself. Both the choice of system and its implementation should be evaluated, but a principled evaluation of a disclosure avoidance system should not hinge on the decisions relating to any particular implementation of that system under the tradeoff.

Okay. So, what did I mean a moment ago when I said the characteristics of an ideal applied disclosure avoidance system? Well, let's start with what is a disclosure avoidance system more generally. Well, simply put, a disclosure avoidance system is a set of one or more statistical methods that transform confidential information about individual data subjects into statistics that describe, estimate, or analyze the characteristics of groups without identifying the data subjects that comprise those groups. Disclosure avoidance systems accomplish this through the application of one or more statistical techniques to reduce disclosure risks in the products being produced.

Now, an applied disclosure avoidance system would be one that performs this function within the operational realities and production cycles of a national statistical office. As such, it acknowledges and is adaptable to the various requirements that stem from the legal, policy, scientific, resource, and stakeholder environments within which it's operating. This can mean being able to adapt to and handle various exogenously determined requirements, such as statistical product design, schedule constraints, exclusions from confidentiality protections, such as the imposition of invariance, and so on.

So, an ideal applied disclosure avoidance system then is one that conforms to a set of overarching principles or features relating to the efficiency, effectiveness, and flexibility of the system as it performs this function, as it transforms confidential information into quality statistics for public release.

So, to reiterate, these principles should reflect characteristics that all disclosure avoidance systems should ideally have, regardless of the specific method being employed or the particular statistical product being considered. And these principles are separate from the implementation decisions that reflect specific choices about the appropriate or desired balance between confidentiality, accuracy, and availability under the triple tradeoff. Those implementation choices will often depend on the characteristics of the data, of the context of the statistical product, and the intended objectives and requirements of the agency and its stakeholders.

Above all, however, it's important to remember that the Census Bureau's mission is to produce high-quality statistics and statistical products. Title 13's confidentiality protections were established in support of that mission. And we must remember that disclosure avoidance is a legal obligation for us and a necessary activity, but it's one that we undertake in tandem with and in support of our primary mission to produce quality statistics.

So, having set the stage with the motivations and contexts for this discussion, let me turn things over to Luke to walk through the characteristics of an ideal applied disclosure avoidance system. Luke.

>> Luke Rogers: Thank you, Michael. And good afternoon, everybody. As Michael discussed, the DA Futures Working Group wanted to develop a scientifically principled way to look at the strengths and weaknesses of different disclosure avoidance options in order to properly evaluate between them.

After substantial discussion and debate, the working group converged on a set of seven principles that reflect the inherent characteristics of any applied disclosure avoidance system that should ideally have. I want to stress, however, that these are just ideals. It is unlikely that any disclosure avoidance system could perfectly meet all of these principles. Rather, the principles offer us a set of ideal yardsticks against which to evaluate and compare the different disclosure avoidance options.

So, for the first principle, an ideal applied disclosure avoidance system, or DAS, should be able to support a meaningful assessment of disclosure risk. Put more simply, in order to protect something, you need to know where it might be vulnerable and how vulnerable it is. With that in mind, a DAS should be able to assess the expected disclosure risk of a statistical product release against known confidentiality attacks, and it should also be able to predict the disclosure risks of those statistical releases against likely future attacks.

As you probably know, with the publication of any statistical product derived from confidential data, some disclosure risks will be unavoidable. But we must be able to understand and evaluate the natures of those risks if we are going to protect against those attacks and if we could put our subjects at risk.

For the second principle, a DAS should be able to assess the impact of confidentiality protections on the availability and accuracy of the resulting statistics. Under the triple tradeoff, any disclosure avoidance method that you select will have an impact on the availability and/or accuracy of the data. How substantial this will be will, of course, depend on implementation parameters selected. An ideal DAS will be able to assess and predict the impact before the parameters have been selected and the predictions have been applied.

For our third principle, an ideal DAS should be able to target protections for vulnerable data subjects. Disclosure risks aren't evenly distributed across underlying populations. Some data subjects' records are more at risk of re-identification than others, particularly when their characteristics differ substantially from those around them. The more unique a record is within a population being described, the easier it is often to re-identify that record. Consequently, an ideal DAS should allow us to target protections where they are most needed, rather than applying uniform protections for all records. This could result in overprotection for much of the data and an unnecessary degradation of data accuracy and availability as a result.

>> Michael Hawes: So, Luke, this is a really good point, and I'd like to offer an example here if I can. One very commonly used disclosure avoidance technique is to round your published statistics. Because rounding techniques are typically applied in a consistent

manner across an entire variable or through an entire table, the uniformity of rounding methods may result in overprotection of some data subjects' information and under protection of those who may be more vulnerable to re-identification. In the business data context, for example, rounding annual research and development investments to the nearest $100,000 may provide ample protection for small firms, but that same degree of protection might be insufficient for a Fortune 500 company. If we had to pick a uniform rounding rule sufficient to protect the largest companies, it might overly coarsen the resulting statistics, making them unfit for their intended uses. By contrast, suppression methods that rely on like P percent rules, for example, can provide variable degrees of protection depending on the relative risk of identifying a particular firm in the statistic.

>> Luke Rogers: Yeah. Thank you, Michael. Moving on to principle four. It should be able to target accuracy for your statistical products. Under the triple tradeoff, disclosure avoidance involves implementation decisions about the desired balancing, that balance between confidentiality, availability, and accuracy of the released statistics. And more specifically, about the relative prioritization of both the use cases those statistics are intended to support. An ideal DAS should give decision-makers the most flexibility to achieve and implement agreed-upon balance that reflects those priorities, targeting accuracy for those statistics and use cases that reflect the greatest value for the nation.

>> Michael Hawes: So, Luke, this really gets at the heart of our mission as a statistical agency. We don't publish statistics just for the sake of releasing statistics. We collect and publish our products because the country needs quality statistics to support specific and important use cases. But different use cases have varying accuracy and precision requirements. And this has consequences for disclosure avoidance systems' ability to target accuracy and precision to where it's most needed.

For example, like in the public health context, analysis of rare or highly infectious diseases typically requires high precision, sorry, high degrees of precision for what might be very small counts or very small geographic areas. A suppression-based method, such as one that redacts any count below 10, would make it impossible to know if a particular area had zero cases or nine cases of the disease, which, depending on the severity, could reflect two very different statistical outcomes. Instead, a potentially more useful disclosure avoidance method that would permit achieving accuracy targets in this context might be reporting the data ranges, reporting the data in ranges, such as like zero, one through five, or six through 10 cases, and so on. So a system's ability to target accuracy and precision in support of different use cases is really important to ensuring the resulting usefulness and value of those statistics.

>> Luke Rogers: Thank you, Michael. That's really interesting. And moving on to principle five. An ideal DAS should be able to track cumulative disclosure risks over time. We know that every independent statistic you release from the same confidential data source contributes to the overall disclosure risk. But the cumulative disclosure risk across multiple data releases is not necessarily a linear function. The combined disclosure risk of two products can often be greater than the sum of the individual releases. As a result, the disclosure risk of multiple statistical products using the same data source needs to be considered holistically. Because we often release many statistical products in sequence that

are derived from the same underlying confidential data, an ideal DAS will need to assess the combined disclosure risk of multiple related statistical releases over time. Do you happen to have any good examples of how different DA mechanisms might fare against this principle, Michael?

>> Michael Hawes: That's a good question. Assessing cumulative disclosure risks across independent products can be tricky. Swapping or rounding mechanisms, for example, typically provide very limited ability to predict the cumulative disclosure risk across multiple releases. By contrast, linear program cell suppression methods and certain forms of noise injection can more readily account for cumulative disclosure risk across multiple releases.

>> Luke Rogers: Thanks. So for principle six, an ideal DAS should be transparent. As a statistical agency, we have the professional and ethical obligation to be transparent about our methods and about the known limitations of the statistics we produce. Because under the triple tradeoff, any application of disclosure avoidance will have unavoidable consequences for the resulting usability of the data, an ideal applied DAS should be explainable to the data users. It should allow for uncertainty that is being introduced to protect confidentiality to be effectively assessed. It should allow for information or tools to be provided to data users to enable them to factor the uncertainty being added into statistical analyses. It should permit external assessment of the degree and scope of confidentiality protection. It should permit external assessment of the efficiency of the system's tuning. In short, an ideal DAS will allow us to explain the impact the system has on the utility of the statistics we release to policymakers and data users. Now, Michael, transparency isn't something we have traditionally thought much about in the context of disclosure avoidance. How might different DA methods fare against this principle?

>> Michael Hawes: So that's a really good point. The different components of transparency that you mentioned can be challenging for many disclosure avoidance methods. Cell suppression and certain forms of rounding are relatively intuitive to non-technical audiences and permit transparent interpretation of their impacts on data usability. But other methods, like swapping, may be easy to explain from a methodological standpoint, but their impacts on data accuracy will often be opaque to data users. Still, other methods like formal privacy may permit greater visibility into the impacts on the usability of the statistics but may be harder to understand conceptually. In short, robust transparency is going to be a challenge for many disclosure avoidance methods, either because of their complexity or because their impact on accuracy can't be readily assessed or communicated to data users. But this is certainly an important principle, and it's one that many disclosure avoidance methods struggle with in one way or another.

>> Luke Rogers: Great. Thank you. So lastly, our seventh principle. An ideal DAS should be feasible. As a statistical agency, we operate with an environment of defined research and production schedules and often of significant budgetary and human resource constraints. An ideal DAS will need to be able to effectively and efficiently operate within these limits.

So, we have identified seven characteristics of an ideal applied disclosure avoidance system that we can use as yardsticks to evaluate the strengths and weaknesses of different

candidate disclosure avoidance system options. To that end, it is important to stress that these principles are intended to be a tool for agencies to use in their evaluations. They are not intended to prescribe the use of one particular DAS over another. In fact, there is often going to be a tension between these principles. For example, the sophistication of an approach that can improve risk assessment and tuning can also make transparency or feasibility more challenging.

Similarly, not all the principles are necessarily equally important. For a variety of reasons, an agency may place a higher value on transparency or on targeting protections for vulnerable data subjects because they may consider one feature to be more critical for their organization.

Above all, these principles are intended to help differentiate decision-making about the selection of a system that is a particular SDL technique or algorithmic framework from the implementation choices about the desired balance of the triple tradeoff.

But the principles only give us half of what we need. They help us guide us to the selection of the DAS to use, but they don't help us answer the difficult questions about the implementation choices under the triple tradeoff. Determining those implementation choices requires answering some tough questions about the desired balance under the triple tradeoff, independent of the selection of which disclosure avoidance mechanism to use. Questions like how protected is protected enough? Which use cases require the greatest accuracy? Should any data elements be excluded from protection? And what level of granularity is necessary to support adequate coverage across data user groups?

To help in answering these questions, the DA Futures Working Group is currently developing a parallel framework to inform discussions and decision-making about these implementation choices and achieving the proper balance under that triple tradeoff. We look forward to sharing that framework with you in the near future.

Now, before we turn to answering some of your questions, I want to end on a point that Michael brought up earlier in today's presentation. Decision-making about disclosure avoidance is difficult. In addition to making decisions about which disclosure avoidance mechanism to use and evaluating those options under the principles we laid out today, the triple tradeoff requires us to make difficult choices about the overall confidentiality protection, availability, and accuracy of our statistical products. At the Census Bureau, those decisions are made by DSEP. But as Michael pointed out, they are not made in a vacuum. Each of these difficult decisions is made based on the best available research and science, guided by the Census Bureau's commitment to scientific integrity, our adherence to our statistical quality standards, and, most critically, informed by the feedback and perspectives of our stakeholders. So now, let's turn things back over to Shelly and take some of your questions. Thank you.

>> Shelly Hedrick: Hi. Thank you, Luke. We're going to now turn to your questions. And if you've been using the Q&A feature to ask your question, thank you. Our team has been trying to answer those along the way. And if we haven't answered your question, we may try to answer it out loud. Michael and Luke, we did have a few questions. Let me go ahead and start with the first one. The question is, you say a disclosure avoidance system protects

the data subjects within the datasets while providing statistics and insights about the groups those subjects belong to. Have you considered aspects of group privacy, that is the need or desire to, at times, protect some of those groups?

>> Michael Hawes: So it's a really interesting question. It's one that I think disclosure avoidance practitioners have increasingly been turning their attention to in some degree or another. I think if you look at the various legal frameworks that statistical agencies operate under that obligate us to protect confidentiality, those are framed very specifically in terms of protecting individual confidentiality, protecting the confidentiality of particular data subjects, whether those are individual people or firms or establishments or organizations, protecting the response information about individual data subjects.

It doesn't say anything about protecting necessarily, about protecting group information. But I think we're increasingly finding that there are aspects of group privacy that merit consideration, particularly when you consider that information about one member of a group can be, well, can contribute to disclosure about other members of the group or situations where there is a particular risk of harm relating to kind of a group's information versus an individual. So this is something that hasn't traditionally been a focus of disclosure avoidance, but it's something that practitioners are increasingly looking at, and it's certainly something that merits consideration down the road. So I'm looking to see if we have other questions here to answer.

>> Shelly Hedrick: Yep. And sorry. I am looking as well. Just a second. I think we do. Oh, here's one. Are these principles listed in order of importance or perceived importance?

>> Michael Hawes: I can take that one, Luke, if you want me to. No, these are not listed in order of importance. And, in fact, as Luke mentioned in his remarks, these principles can often be in tension or in conflict with each other. And the relative importance of one principle over another may vary agency by agency or context by context. There may be situations where transparency is most important. There may be situations where the ability to target protections or the ability to target accuracy are most important for the agency. And so, those contexts in which agencies are operating are often going to influence the relative importance of one or more principles over others. So these are certainly not all equal, and they're certainly not presented in a particular kind of order of importance. Those are really going to vary agency by agency and context by context.

>> Shelly Hedrick: Going through the list here. Hang on just a sec. Just a second.

>> Michael Hawes: Actually, there's one here that I noticed, if I can pick this one up. So there's a question. I would imagine that different disclosure avoidance techniques can provide different levels of confidentiality for the same degree of accuracy and availability. I guess that some techniques are more efficient. How is that incorporated into the framework?

So that's an excellent question, and I think that really gets at the efficiency aspect of these. I think the more efficient a system is, the better it's able to, for example, meet principles three and four, the ability to effectively and efficiently target protection, the ability to effectively and efficiently target accuracy. You're going to get greater degrees of protection

for the same degree of accuracy and availability as if you had a less efficient technique at your disposal. So to kind of use an example that came up earlier in the presentation, a lot of rounded techniques, for example, are applied uniformly across data, which really limits your ability to target for protection or to target accuracy.

So in those cases, what you would find is to achieve some kind of exogenously determined degree of protection, you would have substantially less data accuracy than if you adopted another mechanism that was more efficient along those dimensions that really allowed you to more effectively target protection and accuracy. And in those cases, you would get, for the same degree of protection, you would get much more usable data as a result. So that really gets at kind of the efficiency of principles three and four in particular. So great question.

>> Shelly Hedrick: Great. Thanks, Michael. I'm sorry. Were you done?

>> Michael Hawes: Yeah. With that question, yes, I was. Yeah.

>> Shelly Hedrick: Okay. And again, folks, I've put some links in the Q&A. So, I put a corrected link in after the first. So those should be working. So take a look at those. That's a link to the paper principles that we've put together, as well as a link to the webinar series for this particular webinar, where we'll be posting the slides, the transcript, and the recording. We do have another question here, Michael and Luke. It's thanks for the great presentation. I'm wondering if you can share more about your process for developing and refining these principles. Were there any principles that you considered but did not make it into the final set?

>> Luke Rogers: So I can maybe start to speak to this one and Michael, please jump in. So as a group, I think one of the first things we did was we talked a lot about how the 2020 Census went. And we started kind of with this. Then we kind of moved into like a core set of principles of just what would a ideal applied disclosure avoidance system look like. So forget about the 2020 Census, you know, now that we've kind of recapped lessons learned, but then knowing what we know now, what would an ideal applied disclosure avoidance system look like and just with all options on the table? And we had all these people with different perspectives chiming in. And we just started kind of brainstorming.

And then it was a while ago that we started this process. And it took a lot of meetings and a lot of conversations. But I think as we started documenting and kind of whittling in, like honing in on specific items, we started having certain debates about phrasings and just general conversations about does that fit into the triple tradeoff. Is that something about implementation choices, or is this something more about the principles? And I think that's where we really started to kind of get this clear delineation between those two concepts of there's a difference between choosing a DAS and developing a set of principles about the actual disclosure avoidance system and then setting aside the implementations choices about how to tune that DAS.

And that's why in the presentation, they're very clearly, and we kind of harp on it quite a bit, we try to clearly make them distinctly different ideas. And that was a big part of that early discussion, was just, you know, really teasing those two ideas apart. Because when we

first started, those were really heavily entangled. And it wasn't until we started coming up with these, the seven principles, what became the seven principles, that we were able to tease out that idea of the triple tradeoff and as a separate framework. So, I think that doesn't really answer your question about the deliberations that happened. But I think it kind of gets to the process of it was a lot of going back and forth and trying to disentangle different concepts that were even, you know, I think for us at times, hard to grasp and, hopefully, the fruit of those different conversations and deliberations that will help you all and help everyone else that's dealing with the similar problems of confidentiality and protecting confidentiality in the future.

>> Michael Hawes: Yeah, I'll reiterate everything that Luke just said. It was very much a process of kind of arriving at these, and it involved substantial deliberation and kind of consideration of a lot of different viewpoints and perspectives. And it's really a testament to both the expertise and the passion of the members of the working group that I think we came to the final set that we did, and that it's a set that I know I strongly stand behind and feel we achieved something important here.

So the specific question about like were there principles that were originally considered that didn't get included or kind of how did it come about, I think it's not so much that we considered some and got rid of them as more as we refined what we meant by the principles that we were starting to explore, particularly as we were considering the context in which they could be used. So as Luke mentioned, like, this effort originally kind of started as part of our kind of thinking about our experience with disclosure avoidance for the 2020 Census. But we didn't want to focus our discussions just on census data.

We realized that if we were going to go through this exercise of coming up with a principled framework for evaluating disclosure avoidance systems, we wanted it to be something that applied to all of our disclosure avoidance decisions across the scope of statistical products that we produce from the econ side of the house to decennial to the demographic directorate.

And doing that made us realize like, okay, so we need to think broadly about what some of these principles mean. And that required kind of really thinking about what we meant by certain terms. And we had actually very robust discussions about individual words in some of these that were very enlightening for many of us, I think. And we often had to stop because we also didn't want to stack the deck, so to speak, in favor or against any particular method because all of these statistical methods to protect confidentiality are different tools in our toolbox. And they're all better or worse at what they do in different contexts. And so, we wanted to make sure that these principles reflected that and reflected that there are going to be different situations where different disclosure avoidance approaches are appropriate. And we wanted the principles to help guide us in understanding when that was.

And so, as we kind of thought about these questions and thought about these characteristics of an ideal system in the context of different types of information collections, in the context of different types of data, it really helped us kind of hone in on

what we thought and arrived at as a set of principles that really we hope are kind of universal in this regard.

>> Shelly Hedrick: Okay. Okay, great. We had another question here. This one's about engagement. How do you select the stakeholders for stakeholder engagement? Do you include the general public, or are they just internal stakeholders?

>> Michael Hawes: So that's an excellent question. We have many, many diverse stakeholders. And we have internal stakeholders, many of whom are represented on the project team here. And we have even more types of external stakeholders. And they extend, excuse me, they extend beyond those that you might normally think of as an agency's kind of key stakeholders, insofar as there are lots of different groups and individuals and organizations that are kind of key players in what we do as a statistical agency, from appropriators and policymakers at the federal level, at the state level, your decision-makers, your state demographers, your city planners at the local level, your tribal leaders. You've got privacy advocacy groups. You've got civil society groups advocating for better data about particular sub-portions of the population. You've got those advocating for better data on rural communities, those for advocating for better data about places and towns. There's lots of different perspectives here and lots of different kind of groups that reflect and that kind of advocate for those perspectives.

Now, we have formal mechanisms for stakeholder engagement. In the past, we've done things like federal register notices, where we're just putting out calls to the public. We have formal public channels for feedback, like our federal advisory committees, like the Scientific Advisory Committee, for example, or the 2030 Census Advisory Committee. Those are two great examples that have given lots of kind of input on these sorts of questions. We work with other third-party groups, like the Committee on National Statistics at the National Academies of Science, Engineering, and Medicine. We get lots of feedback from stakeholder groups that represent different kind of subsets of data users and more. And so, so there's no comprehensive list. We're always looking for additional ways to effectively kind of solicit and be able to ingest feedback from these groups. But we recognize that there's no one way of collecting that feedback. And we, as an agency, need to be diligent about how we are doing that in an effective way.

>> Shelly Hedrick: Thanks, Michael. Another question. How does this impact the ACS and Census 2030? Will Census 2030 be using similar methods to Census 2020, or are these being re-evaluated?

>> Michael Hawes: So, these principles, as I mentioned before and as Luke mentioned in his remarks, these principles are designed to be kind of universal for us. We wanted to come up with this set of kind of evaluation criteria to use for all of our disclosure avoidance decision-making. So I guess the short answer to that question is, future decisions about disclosure avoidance for the ACS, future decisions about disclosure avoidance for 2030, future decisions about disclosure avoidance for our other demographic surveys, like CPS or SIPP, or about our economic data collections, like the Economic Census or some of our business statistics, these will all be kind of evaluated in the context of these principles, we

hope, because we sought to come up with this principled framework to help us approach disclosure avoidance in this kind of scientifically principled and rigorous way.

>> Shelly Hedrick: Great. A similar question or different question. Have you considered principles that acknowledge existing error in data as a starting point of disclosure avoidance?

>> Michael Hawes: So that's an interesting question. I mean, there are numerous sources of error in any information collection. You've got measurement error. You've got coverage error. You've got operational error. And when publishing statistics that are derived from a confidential source, you've got disclosure avoidance error. All of these impact the usability of statistics in different ways. In fact, if you read the paper -- We have a link to the paper in the slides. It's available on our website. If you read the paper, we do have a discussion about what it means to be assessing accuracy and targeting accuracy as part of the disclosure avoidance process in the context of some of these other sources of error. And so, I would point you to that.

I think it is certainly something to consider. It's knowing the relative magnitude of disclosure avoidance error, for example, in the context of some of these other sources of error, helps inform some of the decisions about, well, okay, how much disclosure risk should you be willing to accept to achieve greater degrees of accuracy. I think if the underlying data are already highly reliable, highly accurate with very small amounts of other sources of error, that's going to be a different calculation about how much risk is acceptable from a disclosure standpoint than if the underlying data are already very subject to other sources of error. Then perhaps accepting large amounts of disclosure risks to achieve small amounts of additional accuracy may not be as advantageous from a cost-benefit standpoint. So, certainly something to consider. I don't think there's hard and fast rules there, but it certainly can inform the decision-making.

>> Shelly Hedrick: Okay. Hey, another question. Were other statistical agencies besides the Census involved in the principles' discussion?

>> Michael Hawes: That's a great question. In a formal capacity, no. This was the product of our Census Bureau working group here, our DA Futures Working Group. We have, however, benefited from some really valuable feedback from a variety of sources as we've been kind of revising these principles. We first presented an earlier version of these principles at an NBER workshop on privacy, on statistical privacy, back in May of last year. We got some really useful feedback from the attendees of that workshop, many of whom were from other statistical agencies and from kind of partner organizations in the private sector.

We also got some excellent feedback from some anonymous reviewers at the Harvard Data Science Review, where this paper is currently under review. We've presented versions of this with several of our kind of key stakeholder groups, including the Federal State Cooperative for Population Estimates, and we got some excellent feedback from them, and at the last summer's meeting of the American Statistical Association, and got some excellent feedback from those in attendance there. So in terms of kind of formal partnership with other statistical agencies in drafting these, not really. But in terms of the

feedback that we've received from many sources, many of which reflect members of those other statistical agencies, yes, it's been really helpful.

>> Shelly Hedrick: Very good. Another question is, have you analyzed or compared the NMF, which is the noisy measurement file, and the data to best understand the effects of DP, meaning differential privacy, versus the decisions made? In other words, whole, integers, et cetera. Apologies if I've oversimplified what can be done, but curious to know what that might have looked like and how these seven principles apply to what happened in Census 2020.

>> Michael Hawes: So there's a lot packed into that question, and I think it merits attention, and I think it merits more attention than I'm able to give it in the context of the short amount of time that we have left with this webinar. So I don't want to make light of the issues that that question represents because I think they are certainly ones that merit discussion. But I can see who submitted that, and I'm going to beg this person's forgiveness for seeing if I can try to answer it in a separate forum where we can do that topic more justice.

>> Shelly Hedrick: Very good. Another question. Do you imagine that the weighting of the different principles might differ for different statistical products. And so, they might use different DASs? Is consistency of DAS something to aim for? And perhaps that contributes to principle six, which is transparency.

>> Michael Hawes: So that's a great question, and, and I think I touched on this in one of my responses earlier. Absolutely. I think the relative importance of one principle over another can absolutely vary depending on the particular statistical product under consideration. I think there are certain cases, certain types of data products where transparency might be paramount. I think there are others. I can particularly think of some very sensitive surveys that we collect here at the Census Bureau or that other statistical agencies, like National Center for Health Statistics, for example, collects where like the the primary objective might be making sure that you can properly target protections if you're talking about really sensitive data subjects or data elements that could cause irreparable harm to your survey participants. So I think, yeah, based on the agency's context, based on the context of the particular statistical product being considered, I think you can absolutely have different kind of relative prioritizations of these principles. And that's certainly something that an agency should be aware of and should certainly consider. There's no one set of kind of priorities that should be followed. These are a set of yardsticks to use, and which yardstick is going to be most important will vary depending on the situation.

>> Shelly Hedrick: Okay. Another question. Did you do any practice decisions on use cases examples or examples based on the framework as it was being developed? And if so, what kind of scenarios did you consider?

>> Michael Hawes: So that is a fantastic question, and it really gets at the heart of much of what my role was during some of the discussions that we had. I frequently had to put myself in the context of devil's advocate, making sure that as we were coming up with nuances, as we were trying to settle in on kind of the, the precision of the principles as we

laid them out, that we were thinking about the ramifications that these principles would have in very different contexts.

So I've been involved in disclosure avoidance within the federal government for the better part of a decade and a half. And I've done these kinds of things in a lot of different contexts. And so, I found myself thinking back about, okay, how might these principles have been interpreted, or the draft principles at the time be interpreted in the context of decisions about disclosure avoidance for the civil rights data collection that the U.S. Department of Education publishes, or in the context of some of our economic surveys, or in the context of health data that I've been asked to opine on at one time or another and looking at how these might be interpreted or more or less informative as principles when thinking about one type of disclosure avoidance approach over another. And so, there was a lot of these kinds of exercises, these hypotheticals, or kind of walk-through case studies, how might these have been used in some of those decision-making. And it was very informative, and it helped us kind of hone in on, I think, what is very precise language that can apply no matter the context.

>> Shelly Hedrick: Okay. Another question. We still have a few more minutes, so I think we can get a couple more in. The question is, do you have thoughts on how principle five, to track cumulative disclosure over time, might be relevant to releasing results of analysis conducted by external researchers and FSRDCs? It seems like it might be important.

>> Michael Hawes: It is very important. And I thank you for that question. Tracking cumulative risks over different products and over particularly different types of uses of the statistics, I think is particularly vexing in some situations more than others. I think it's a very different animal, so to speak, to think about tracking cumulative disclosure risks when you're talking exclusively about releasing sets of tables, like sets of statistical tables, which is kind of the traditional form of a statistical agency's data products.

But I think bringing up kind of researchers in the Federal Statistical Research Data Centers is a really interesting point because, okay, if you've historically been focused on tracking cumulative risks over data tables, how do you then incorporate the incremental, if marginal added disclosure risks of an extensive set of model parameters or regression coefficients? Those may have a very different kind of disclosure risk profile to them and are going to contribute to the overall disclosure risks in very different ways.

And so, having kind of unified frameworks that allow you to really effectively track disclosure risks across those different types of things can be very challenging. And I think to my knowledge, there is no disclosure avoidance method that is foolproof for any type of disclosure avoidance, or sorry, for any type of statistical output that's foolproof in this kind of tracking. Some do a better job than others, certainly. I know formal privacy has gotten a lot of attention in this regard, but there are some types of analyses, some types of model parameters where there's no effective way for formal privacy to evaluate the disclosure risk or to evaluate it in kind of tangible ways. So this is certainly a work in progress and certainly something that is going to require a lot more research over the coming years to fully answer.

>> Shelly Hedrick: Okay. And I think we just have time for one more question. This is, to what degree does the application of a DAS require an understanding of how the data are to be used after publication? How are those uses evaluated for 2020 and how will they be evaluated for 2030?

>> Michael Hawes: So I think we don't have time to answer the second part of that, about how they were used in 2020 or how they'll be used for 2030 or for kind of any other kind of concrete case. I think that's a much more extensive discussion. But I think to the first part of the question, which I think is a really, really useful one, I would say that is exactly what principle four is all about. The ability to target accuracy as part of your ability to produce quality statistics is essential for a disclosure system. Our goal is to produce quality statistics. Doing that means knowing how the statistics are going to be used so that you can maximize for any given level of protection, so you can maximize the usability of the statistics you're releasing. And usability here is, is both of the other dimensions. It's the accuracy dimension, certainly, and it's also the availability dimension. Are you producing all the statistics that are necessary to support those really important use cases? And are you releasing those statistics with a degree of accuracy and precision sufficient to support those uses in an effective way? So in order to even be measurable against principle number four, you have to be able to understand how the data are going to be used. Luke, I know this is something near and dear to your heart. Do you want to add to that?

>> Luke Rogers: I was going to say in addition to and separate from the principles, the understanding the use cases from the stakeholders particularly is critical to tuning that, the triple trade-off. So it's a big part of that component that we'll be talking about in the future. So stay tuned for our work when we start talking more about that publicly.

>> Shelly Hedrick: Great. Okay. Well, then that wraps up today's webinar. But I want to say that there were a couple of questions we couldn't get to. We do have an inbox, and that is 2030DAS, 2030DAS@census.gov. So, please don't hesitate to send us your questions or comments to that inbox. And, you know, as I mentioned earlier, we will be having the slides, the recording, and the transcript posted in that Census Academy page in the next two to three weeks. Both of the links that I posted a little earlier should be working. I originally posted a bad one, but they both should be good at this point. You might have to click on, you know, continue to an external website to be able to access them since we're coming in from, you know, the Webex platform, but they should be working. Thanks to everyone who joined us. Thank you for your excellent questions, and thank you our presenters, Michael and Luke. Have a great rest of your day.