South by Southwest

Director's remarks as prepared for delivery

March 13, 2023

How Good Do Census Bureau Statistics Need To Be? It Depends...

- Good afternoon, everyone! Welcome to a little talk about how good statistics—and, in reality, the statistical data produced by the U.S. Census Bureau—need to be.

- And you know what? We're going to have a little fun, and hopefully do a little learning in this session.

- I'll talk about topics using some storytelling.

- The first topic is on—of all things—perfection.

- The second focuses on data quality; what it is and what it isn't. It gets to the notion of "data goodness."

- And the last focuses on the "it depends" part of my title.

- But before I begin, I just have to acknowledge two groups of people.

- The first is my family who happens to be here: my wife Della, daughter Clarisa, my granddaughter Renee, and son Emilio. Love you to death and your support means more than anything to me.

- Next is my second *familia*—the photogs of South by Southwest, some of whom are here today. I no longer have that gig but let me tell you—the experience working and learning from you was nothing less than transformative. You helped me to advance my mojo, to nurture and embrace my creativity in all I do, and to just be myself.

- You know, part of my leadership focus as director of the Census Bureau is in large part because of my experiences with you over the years. I tell everyone at Census Bureau they need to bring their whole self to the table to be effective at anything they do—professionally or otherwise.

- OK on with the show! Let's geek out and talk about Census Bureau data!

- First a little background... In case you missed it, the Census Bureau is our nation's largest federal statistical agency.

- We have over 13,000 staff and just under a million square feet of office space at our headquarters, in Suitland, Maryland, besides maintaining six regional offices, a National Processing Center, and a contact center spread across the country.

- We generate trillions upon trillions of data points and statistics each year, both on our wonderfully diverse population and our wonderfully diverse economy.

- We conduct three censuses—population, governments, and economic. We conduct over 130 surveys annually covering our population and businesses.

- On any given day of the year, we have multiple surveys active in the field.

- Hey, I was at a place of business last week and the owner told me they received a letter with my signature on it to complete the economic census. I'm still getting used to folks telling me that.

- But our economic census is still active, by the way. So, if you're a business owner and you got a letter with my name on it to complete the form, please do so asap.

- In any case, our data collections cover a huge array of topics. I won't go through all the topical areas we cover, but let's just say that our footprint includes many important aspects of society.

- To give you just a taste, we cover things like housing, education, families, employment, commuting to work, economic indicators, supply chains, broadband access, emergency planning, and much, much more.

- But at the core of all the data we produce is quality. Our mission is to produce quality data on our nation's people and economy.

- But what exactly does quality mean? That's what I'm here to talk about. And there's no better way to start than by talking about perfection.

- So... What is "perfect?"

- Hey, in our world—the world of statistical data—we'd absolutely love for our data to be flawless. That means no errors, no irregularities... utter perfection!

- But as we all know, perfection is in the eye of the beholder.

- Consider this picture of a snowflake. It's just an ice crystal. I love macrophotography, don't you? Check out how each of the six legs stretches out beautifully. The shape is symmetrical... Its dimensions appear uniform. If that isn't perfect, it's pretty close, huh?

- Well, if you take a closer look, you can easily spot imperfections. Each leg has two sides, and you can see that they don't align perfectly with each other.

- Perhaps if created in a vacuum, this crystal would be perfectly formed, perfectly symmetrical. Alas, things like humidity and temperature affect crystal formation. That's why no two snowflakes are ever alike. Perfection escapes us, in reality...

- So how does this relate to Census Bureau statistical data? Well, let's go there.

- OK, let me start with a profound statement, even though it really isn't if you think about it.

- Now, I don't want to burst any bubbles, but no census ever conducted in our great nation's history has ever been perfect. No survey ever conducted by the Census Bureau has ever been perfect.

- In fact, this is true for all surveys and all data: Imperfections always exist. Yes, we do strive for perfection, and we do advance that ball down the field. But crossing the goal line of perfection is impossible.

- And you know what? That's ok!

- Like the image of our beautiful snowflake, statistical data that are diligently procured and curated can provide much value. But you have to know how to get that value. And value can be had without achieving perfection. That's always been the case.

- So, if the troves of statistical data we produce at the Bureau aren't *perfect*, then how should they be characterized?

- Aha! That brings us to the concept of *data quality*... a topic that I am particularly enamored with.

- Like the notion of perfection, data quality can mean so many things to so many people.

- You know, I like to compare the concept of data quality to music... That seems appropriate for South by Southwest, right? So let's go down this rabbit hole.

- Think about music and what you like about it. Some folks are attracted to the beat. Others dig the melody. Some like the instrumentation. Still others find the lyrics captivating. And others appreciate the *composition*—how all the elements fit together to create a true piece of art.

- As for me, I've been a live music photographer. I've pretty much seen it all: the rappers, hip hop, hard rock and metal, indie, blues and jazz, country, K-pop, folk, classical... You name it. And in their own way, each is magnificent.

- But to really appreciate music, it needs to speak to you. And music can speak to you differently depending on your needs. If you're happy and want to party, there's s music for that. If you are melancholy, there are songs that provide you comfort. And if you just need some background music while working or having a meal, there's a playlist for that.
- The point is, depending on your needs, there will be different combinations of melody, beat, lyrics, composition that fit the bill... right?
- Well, guess what? The concept of combining elements to meet your needs applies exactly to data quality, as well! Let me show you how and why.
- There's a framework that the Federal Committee on Statistical Methodology published in 2020. It's appropriately titled of all things, "A framework for data quality." It lays out three domains to data quality:
  ◦ Utility
  ◦ Objectivity
  ◦ Integrity
- ...And the interpretations of these three are as you would expect. I'll discuss each.
- *Utility* relates to how well the data address one's needs—how useful is it for what you are trying to do? Obviously, needs can change from one person to the next, or from one problem you are addressing to the next.
- But you know what? Utility isn't uniform; it varies by your purpose or objective.
- So, we have some land in the Hill Country... it's pretty rocky and hard to get around out there. I wouldn't take a Lamborghini out there because it would be useless. I'd rather have my four-wheel drive pickup that's made for off-road travel. A pickup fits my purpose of navigating the Hill Country terrain.
- But if I needed speed on a smooth highway... hey, the sportscar would be sweet and meet my needs.
- The same principle applies to data utility. Some of our surveys provide great data at the national and state levels. But they aren't useful for small geographies such as rural counties or neighborhoods in cities.
- Our 2020 Post-Enumeration Survey is a great example of that. It assessed the quality of the 2020 Census at the national and state levels. The survey estimated who was overcounted and who was undercounted. But it wasn't designed for substate or city estimates. The sample size is too small for that.
- On the other hand, our flagship American Community Survey, which we call the ACS, is the opposite.
- Using our 5-year ACS data, you can get estimates down to the census tract and block group levels on a variety of socioeconomic topics such as poverty, disability, education levels, home ownership, household income, and so on. That's why the ACS is used in most of our policymaking in this country.
- Now let's turn to the second dimension of data quality is *objectivity*. It reflects accuracy, reliability, and error structure. Many folks—especially us statisticians—like to focus on this dimension a lot when thinking about data quality.
- We're all familiar with margins of error in polling, right? Well, that's one measure of objectivity. It communicates the natural variation of a statistic that is based on a sample of a population instead of a census.
- We statisticians take comfort in being able to control the size of that variation by how large of a sample we use and by how we design a survey sample.

- And then there is the nefarious influence of *bias*. Bias refers to systematic error or shift in a measure that you often don't know about.

- To give you an idea of bias, let's think about here in Texas, where we love our SUVs and trucks. Inevitably, our vehicles get out of alignment. At least that's what the mechanics tell us.

- Now, if you take your hand off the steering wheel while driving you should keep going straight, right? But if your alignment is off, you'll drift to the left or to the right, and you have to make constant steering correction just to keep straight on the road. When that happens, there's a bias in your tire alignment, a systematic misdirection of one or more tires.

- Unfortunately, when it comes to statistical data, you don't have the luxury of taking your hand off the steering wheel to see if you drift to the left or right. We know those imperfections exist in our data, and we do our best to minimize them.

- We do that with a lot of research. And then we develop ways to anticipate and correct for bias in our data.

- As you might expect, there are many sources of bias. We do many censuses and surveys, and not everyone participates. The participants may be and likely are different from those that do not participate. That's nonresponse bias.

- And even among those who participate, some segments of the population may not fully understand the questions being asked, or misinterpret them, or deliberately provide untruthful answers. That kind of error is called response bias.

- The last of the three elements of data quality is *integrity*. Now, the concept of integrity reflects the scientific rigor used in gathering and processing the data. How thoroughly were the design and the methods thought through?

- If I do a survey of adults in Texas only in English only, it's going to be impossible to make conclusions about all Texans because over one in three speak a language other than English at home.

- And if that survey of Texans is an online survey? Well, over one in eight households in Texas doesn't have broadband access, so that would have an impact on the study's scientific integrity.

- By the way, integrity also includes as well as protections against manipulation, influence, and of course, unauthorized access. So, there is a bit to unpack in terms of this aspect of data quality.

- But at its core, integrity is about how well the research structure is built. If the roof of your new house isn't built well, it doesn't matter how good the rest of the house is, you're going to eventually get water damage from rain. If your windows aren't sealed properly, you'll have problems there too, regardless of how well everything else is built.

- OK, so those are our three data quality elements at a glance—utility, objectivity, and integrity. Analysts, policymakers and the public seem to favor one of these dimensions over the other.

- Now let's get back to the framework. I'd like to home in one of these dimensions—the first one, utility.

- Under our FCSM data quality framework, utility features the following dimensions:
  - Relevance
  - Accessibility
  - Timeliness
  - Punctuality
  - Granularity

- *Relevance* is defined pretty much as the word suggests. It's the extent to which data meet a user's needs.

- Next is *accessibility*. It focuses on the ease of acquiring statistical data and associated documentation, as well as the ability to understand the products, themselves.

- For example, our American Community Survey public use microdata is relatively easy to access. You can just download it from the web. But it takes a bit of processing knowledge and statistical training to use the data.

- Then there are data visualization tools like the Community Resilience Estimates, the Census Business Builder, and My Community Explorer. These offer easy access to data using geographic data visualizations that can facilitate community needs assessment and economic development planning.

- The next dimension is *timeliness*. It simply refers to the time lag between the date of observations and the public release of the data.

- The pandemic made us very aware of this dimension. For instance, our nation needed contemporaneous data during the pandemic, and two high-frequency surveys were developed to meet that need—the Household Pulse Survey and Small Business Pulse Survey.

- Next, there's *punctuality*. While related to timeliness, it specifically focuses on the ability to stick with the official data release schedule. For instance, punctuality was a key concern in the delay of the apportionment counts and our redistricting data release.

- And finally, *granularity* is all about disaggregation—the amount of detail baked into data items like time, geography, and sociodemographic characteristics such as age, income, industry codes for businesses, and so on.

- But today, I'd like to focus on just one dimension of data utility—*relevance*.

- Now, like most things in life, relevance can mean many different things to different people. On top of that, relevance is situational to what you are using the data for.

- For instance, last year I addressed the Alaska Federation of Nations, an annual gathering of all tribes in that state. Some tribal leaders reported their need for better data for their governance, specifically for their planning in housing, education, and health.

- They spoke about how the employment and earnings questions don't really capture the sustenance work performed in remote villages. Asking for monthly income or hours worked is almost meaningless when you live in a remote village and need to fish and hunt to sustain your village residents. The absence of meaningful data clearly reflects a cultural relevance gap *from the tribe's perspectives*.

- Now take macro-economists. From their perspectives, the failure to capture data on sustenance work in remote Alaska may no bearing on their national economic analysis.

- Now let's look at this a bit more broadly.

- For about the past decade, I've spoken about a renaissance that is unfolding before us, as a result of technological advancement and globalization. It has fundamentally changed society and our culture.

- We're now data-driven. We increasingly rely on immediate, easy access to information. In fact, we expect to be "catered to" by way of tailoring algorithms based on our internet use.

- This renaissance also features cultural changes. With the ever-increasing use of social media, connectivity between families and friends, a new wave of virtual communities has sprung up. More and more people have smart phone access, and they use social media apps to connect with each other.

- Technological advances in genealogical research and DNA testing have amplified interest in people knowing who they are racially, ethnically, and culturally. Our nation's population is becoming more diverse, including mixed race, mixed-ethnicity, and mixed race-ethnicity. We recognize the diversity within racial and ethnic groups.

- Sexual orientation and gender identity is recognized as an important part of who we are as a nation.

- And then there's our economy. As it evolves, so do aspects of work, such as telework, remote work, as well as the classic office work and manufacturing work locations. New industries are emerging, for instance the cannabis industry with a slew of new products.

- And even the term "work" has taken a more complex meaning. Sure, we still have folks working regular hours, with salaried monthly and biweekly paychecks. But a sizeable chunk of our workforce can only make ends meet by working multiple jobs, or doing gig work, or some combination.

- Here in Austin, many of my photog family often hold multiple jobs besides passionately pursuing their gig work as a photographer. The same is true of musicians and other artists. Several of my media interviews this year have been with freelance journalists—gig workers.

- This is all important context when thinking about data quality in our federal statistical system. We pride ourselves in gold-standard data collections like the American Community Survey, the Current Population Survey, our decennial census, and the many other data collections we conduct.

- But as society evolves—and does so rapidly—the relevance of data items we collect will be affected over time, and not necessarily in a good way. And that will affect their utility, a key dimension of data quality.

- In fact, we've already seen that with our race and ethnicity collection and reporting standards. We continue to use the 1997 standards. Fortunately, through the leadership of our new Chief Statistician Karin Orvis, OMB is in the process of reviewing and revising those standards.

- There's also the issue of collecting data on those involved in the ever-growing gig economy. Many Americans supplement their employment income with gig work. Gigs are sometimes referred to as contingent jobs or alternative work arrangements.

- Both the American Community Survey and Current Population Survey capture some data on these, but only when they are the main sources of work. So, we have a recognized data gap for everyone who uses a combination of jobs and gig work to make ends meet as well as for conventional workers who use gig work to supplement their income.

- Now to be clear, the federal statistical agencies do adapt to our changing society. The Contingent Worker Supplement to the CPS has been in place since 1995. Questions on computer and internet use first appeared in the ACS in 2013. Hey, the question on whether a house had a flushing toilet was removed from the ACS in 2016. And, of course, the 1997 standards for collecting and reporting race and ethnicity are currently being revised.

- So, change does occur. The issue is the extent to which societal change is outpacing revisions in our data collections. It seems to me like we are getting further behind.

- There are, however, shining examples of quick adaptation. It relates to the data quality dimension of timeliness, although one could argue that it overlaps with relevance.

- Early in the lockdown period of the pandemic, the federal statistical system realized there were no contemporaneous data available to tell the story of how households in our nation were faring. The ACS would be of little use, due to the 9-month lag between the last data of collection and the publication.

- So a collaboration was formed with other federal statistical agencies to create the Household Pulse Survey. It featured quick turnaround national surveys, focusing on how the pandemic affected households socially and economically. The first survey was launched April 23, 2020, and the survey program continues to this day.

- As an online survey with a low response rate, the quality threshold of the Household Pulse Survey is well below that of our flagship surveys. In fact, the product was labeled experimental.

- Yet the information still provided a valuable, contemporaneous glimpse into how the nation was dealing with the pandemic. Policymakers and the public loved it. It was the right data product for that time. It aligned with the public's appetite for contemporaneous data during a profound and vulnerable period in our country's history.
- I'll note, by the way, that there's a similar sister survey of establishments called the Small Business Pulse Survey that served the same purpose, but for businesses.
- Ok, I've talked about data quality in the context of data utility, specifically the dimension of data relevance maybe with a little timeliness sprinkled in. So where does this leave us? How good do Census Bureau statistics need to be?
- First, let's recognize that data can be collected that are highly accurate, reliable, and demonstrate little bias. Yet they can have little relevance, or their relevance can diminish quickly with time.
- And the risks rise when an ongoing assessment of the content we collect is not baked into our statistical programs. For instance, we always begin planning the next decennial as data collection for the current one ends. I suggest that we adopt that model for all the data we collect, be it from surveys or even from administrative data sources.
- If you hadn't yet realized, this is also a matter of data equity. Collection of job and income data is easiest for those fortunate enough to enjoy conventional jobs paying a monthly salary or biweekly paycheck. But among underserved vulnerable populations, it's difficult to provide accurate employment and income responses for those involved in episodic work, gig work, and multiple part-time jobs. The cognitive response burden is considerably higher for these folks.
- Unless we're willing to explore better ways to capture data from these segments of society, an inequity will exist in data quality across various subpopulations associated with those who are most affected by inequities. That's why we need to continuously assess the relevance of the data elements we currently collect, even when they have recently been revised.
- But let me directly answer the question posed in my session title: How good do Census Bureau statistics need to be?
- The answer is truly that "it depends." It depends on the specific questions that society needs to answer.
- The U.S. Constitution calls for a decennial census for the distinct purpose of congressional apportionment. Those 50 population counts—one per state—must be as accurate as possible.
- More generally, to meet their objectives, our many survey programs each require a different balance between our data quality dimensions—utility, objectivity, and integrity. Flagship surveys like the ACS require a heavier reliance on objectivity and integrity and less so on utility… specifically on timeliness.
- On the other hand, our high-frequency survey programs—like the Household Pulse Survey—sacrifice objectivity to bolster timeliness.
- Both contribute to our knowledge base and both benefit the public. They just do so in different ways to address different data needs.
- Just like our music metaphors, there are different combinations of melody, lyrics, and beat that meet one's immediate personal needs.
- So, to all my Census Bureau data nerds out there, I implore you: The next time you need a problem or issue addressed with data, think about the combination of data dimensions that will best fit your needs.
- And to my photog family: if y'all are out hearing music tonight, think about what mix of musical elements best fits your personal needs of the moment… And then find that band to that venue.
- Thank you everyone.