The Law, such as it is
Season 3, Episode 4

Larry

This is Larry Lessig. This is the fourth episode of the third season of the podcast, "The Law Such As It Is." In this season, we are focused on the tenure removal of Francesca Gino, the Harvard Business School professor accused of academic misconduct.

The charges against Francesca Gino were grounded in four papers — out of the more than 140 that she had published in the 20 years of her academic career.

These four papers presented what we're going to call "data anomalies," meaning that the publicly posted data doesn't quite make sense as the original data that would have supported the analysis of each paper.

The question is why these anomalies exist. Are they mistakes—made by the people preparing the data to be analyzed—or are they intentional fabrications? And more specifically, in legal terms, was there "clear and convincing evidence" that they *were* intentional fabrications?

Remember where this started: Data Colada — a fantastic group of data scientists, testing the validity of published empirical research — discovered what they called "apparent anomalies" in these four papers. That led them to submit a complaint to the Harvard Business School. That complaint began the two-step process that eventually took Francesca's tenure away.

In the first step, the Harvard Business School investigated the four studies in the four papers in question. They concluded that the anomalies were intentional—that she had committed academic fraud by intentionally manipulating the data in the way the anomalies had revealed. HBS then recommended that the university convene the second step in this process: what's called the "Third Statute" proceeding by the university to revoke her tenure.

We're here, with this podcast, because the university did indeed convene the Third Statute procedure and then found Francesca guilty, making her the first tenured professor in Harvard's history to have her tenure revoked.

But as we covered in the first three episodes, in the first stage of that process, the HBS internal investigation, Francesca was

blocked from gathering evidence to defend herself: HBS forbade her from discussing the charges against her with anyone except the two advisors she was permitted to select. That meant she could not talk to the people who would have prepared the data that she had analyzed to figure out exactly how the anomalies could have been produced. It also meant she was forbidden from hiring a forensic expert who could help her to evaluate the forensic evidence that HBS's own forensic expert was developing to establish her guilt. HBS, in other words, got expert help preparing the prosecution; Francesca was denied expert help in mounting a defense.

And as we're about to see in this episode, that denial of process was fatal to her ability to demonstrate her innocence. Because the mistakes made by the forensic expert hired by the business school led the Investigative Committee of the business school — what we'll call, the "IC" — to conclude that she was guilty. And yet, when those mistakes were identified by the forensic expert Francesca was able eventually to hire, HBS abandoned that evidence, and hired a different expert to craft a whole new case against her.

I make lawyers for a living. Let's just name what this is. This is a pathology of lawyers as advocates, rather than lawyers as justice seekers. The business school had gaggles of lawyers lined up to prosecute Francesca. But apparently, no one had the courage to say —once the evidence that had grounded her initial finding of guilt evaporated—that maybe they should reconsider whether in fact she was guilty. Instead, rather than reconsider, they doubled down.

All that's a bit abstract. Let's begin in this episode with the clearest example of this fatal flaw: what we're going to refer to as allegation number two.

To tell this story, I'm going to have the help of a friend — actually two friends. The first friend is a friend who is expert in data analysis, and who can help unpack the facts behind the charges made against her.

My conversation with this friend is going to be a little bit analysis, a little bit recounting of what actually happened. But the aim is to help you to see just how weak the charges against Francesca are. So weak that certainly no one could conclude fairly—at least with "clear and convincing "evidence," as the rules required—that she is guilty of academic misconduct. And by the time we get to the end of this story, I think you will agree that she is, as I believe and as she fiercely insists, actually innocent.

But for complicated reasons unrelated to this case, my friend is not free to talk freely about the case. So I can't interview them directly for this podcast directly.

That's why I need a second friend, to give voice to the story that my first friend will help tell. That second friend is Ron Suskind. Ron will be the voice of my anonymous first friend.

Ron is a journalist, an author, a documentary-maker, and for the past few years, the president of an AI company, Sourcebase.ai. From 1993 to 2000, he served as senior national affairs writer for the The Wall Street Journal, and in 1995 won the Pulitzer Prize for Feature Writing for a series of articles on inner-city students. His books — including A Hope in the Unseen, The One Percent Doctrine, Confidence Men, and Life, Animated — have been wildly successful best-sellers. I've known Ron for 14 years; he's been a dear friend, with endless advice, some of which I've taken.

Welcome, Ron, and thank you for doing this. Tell us a bit about Sourcebase.ai.

Ron

Well, it's really the AI for media. A couple years ago, I gathered leaders in AI and leading journalists some of whom you may know to build an AI that can essentially be a truth discovery and distribution machine scalable frankly to data pools as large as the Pacific Ocean.

And what it is is an AI that lots of folks in media are using as well as citizens where can essentially have a conversation with enormous pools of document data — evidence, essentially — and discover what is true. We're having a crisis of misinformation, I would submit, in this period. And we would be considered a counter force to that by virtue of being able to use AI to tap into essential fact and evidence instantaneously.

Larry

And it's a rag-based AI.

Ron

It is. It is an LLM rag that means a large language model integrated with retrieval augmented generation — RAG. We're, I think, maybe the most powerfully scalable rag with all sorts of added capacities for precise research and analysis, and that's why it is one that our friends in the media are using as well as folks in government, researchers, and citizens as well.

Larry

Okay, so this is not as sophisticated an effort to tell the truth, and we're not using AI; we're just doing it person-to-person right here.

Ron

Well, let me just say that what you have here though is the kind of thing that, of course, we see as powerful in its evidentiary quality. This test on testimony here, this conversation if you will, with the testimony of your friend who I will play in this in this episode, is the kind of thing, of course, that we put up on Sourcebase.

Larry

We're going to explore how much we can put on to Sourcebase, but right now for this episode, this is how it's going to work.

I asked my first friend, the anonymous friend, to write answers to the questions I would pose in this podcast. They did. And now I'm asking my second friend, Ron, to read those answers. The aim is for Ron to be the voice of my unnamed first friend. The words he will utter are not his words; they are my friend's words. Ron will give my anonymous friend a voice. It's a little bit weird, but Ron's got a pretty good voice, and for a Pulitzer Prize winner, he reads pretty well.

So, thank you, Ron, for playing this role.

You're supposed to say, "Thank you, professor," at this point.

Ron

Thank you, professor. I suppose that's what your friend would have said to you. That's what I'd say. Okay, well, let's go Larry.

Larry

Okay, so let's jump in.

At issue in Allegation 2 is a paper by Francesca and two coauthors, one a professor at Kellogg and one at Columbia, "The Moral Virtue of Authenticity: How Inauthenticity Produces Feelings of Immorality and Impurity."

What was this paper about? What were the authors trying to demonstrate?

Ron

The hypothesis of the paper is that experiencing inauthenticity creates feelings of impurity, which lead to a desire for physical cleansing and greater prosocial behavior as a form of moral compensation. So in simpler terms, when you do something "inauthentic," like acting enthusiastic about a project you secretly dislike, you manifest a desire to cleanse yourself or "make up for it" by doing something good for others.

Larry

This study was conducted partly in the lab (meaning in a physical space where people came into a lab, sat at computers, and answered questions that were presented on the computers) and partly online. Why did they divide the study like that?

Ron

It seems the lab couldn't get as many participants as the researchers wanted for sample size purposes, so they sought more participants online.

Larry

The original charges by Data Collada focused on 20 anomalous entries that came from the online participants. Can you describe what those anomalous entries were, and why they seemed weird.

Ron

The study was supposed to be limited to Harvard College undergraduates. One of the demographic questions was "year in college," which is a question that would make perfect sense for undergraduates. Good answers could be "freshman," "sophomore," "junior," or "senior." Maybe "class of 2015" would be a reasonable answer, or something like that. But Data Colada noticed that in the published data for this study, 20 of the entries had "year in college" set to "Harvard." That made no sense: no Harvard College undergraduate would enter that value, Data Colada reasoned. Certainly, twenty different Harvard College undergraduates wouldn't answer that way.

Larry

Okay, but why did that suggest that Francesca might have been engaged in academic misconduct?

Ron

Well, if real Harvard students didn't enter this data, then these answers must have come from someone else. Then you have to ask *who*? Maybe it was Francesca who provided these answers herself, by just faking data, by just typing it in. As Data Colada explained their thinking, and I quote: "it is difficult to imagine many students independently making this highly idiosyncratic mistake."

Larry

Let's unpack a bit more the particular way that these entries might have been created. First, how were participants recruited? Why would they participate?

Ron

Some were just interested in the research. But some did it for the money. In the physical lab, there was a well-defined process for paying research participants: cash was given to each participant at the end of the study. Online, participants were given Amazon gift cards. The pitch to prospective subjects was along the lines of: "Do this easy survey, and get a $10 gift card."

Larry

And so, how do people learn about these types of surveys?

Ron

Some learn about them from professors at Harvard College announcing them in class. Some learn about them from email lists that are supposed to be limited to Harvard College students, such as mailing lists for undergraduate dormitories.

Larry

So for the online participants, someone would learn that there's an opportunity to earn an Amazon gift card by filling out a survey, and they would follow a link to fill it out. Right?

Ron

Right.

Larry

Now when they logged on to the platform to fill out the survey, the platform (which was served by a company called Qualtrics) would have collected a bunch of data about those people filling out the survey, right?

Ron

Yes, Qualtrics gathered the IP address for the location from which the survey was entered, the "user-agent" of the user's browser (which is meant to report which browser the user is using, like Chrome versus Firefox, and what version, and also what operating system), as well as some other data.

Larry

And so, then, what was the theory about what Francesca must have done to create these 20 anomalous entries in the dataset?

Ron

The simplest theory is that Francesca wasn't getting the results she wanted with her core dataset, so she added these 20 entries to bolster the results.

Larry

But what is the theory? That Francesca doesn't know what "year in college" means? That she's confused about what the form is asking for? Didn't she design the survey to collect the data?

Ron

Right, it's her survey. She knows what the question asks, she knows what a logical answer would be. If she were stuffing the ballot box, so to speak, you'd think she'd be careful enough to submit correct answers — like "sophomore," or "junior" — definitely not "Harvard."

Larry

Now wouldn't you expect that any reasonable investigator would look at this and think it's kind of weird to imagine Francesca is the source of these anomalous entries—again, because if she were the source of these anomalous entries, she at least would have filled out the survey form correctly.

Ron

Right. If I were an investigator, I'd be asking questions here. Francesca wouldn't make this sort of mistake. No one who knows undergraduates, who has attended or spent time at a US college and speaks English even moderately competently, would think these answers make sense.

Larry

So, if you were a careful investigator noticing the anomalies in these anomalous entries, you could have used the Qualtrics database to help you learn something about where these entries came from, right?

Ron

Yes, Qualtrics has great tools for investigating here. Broadly, this is in the category of meta-data, which means data about data, like the last-modified date of a Word file on your hard drive. Qualtrics gathers data like this automatically, as the user submits each survey entry.

Larry

What would this metadata have shown? Or more precisely, what would an investigator have imagined the metadata should have shown?

Ron

The smoking gun here would have been an IP address that was, for example, tied to the IP address at Francesca's home, or maybe her office, or her cell phone, or a friend's house. Whatever, it would link back to a location that Francesca would have been associated with at a time when Francesca would have been there.

Larry

And so if the metadata had shown a place where Francesca was, or reasonably could have been, that would have supported the conclusion that Francesca made these 20 suspicious entries. Correct?

Ron

Right. If the suggestion is that she did it, then the IP addresses should point to her.

Larry

From your understanding of the investigation done by the business school: Did they consult the log files or attempt to determine forensically where or who could have made these entries?

Ron

No. In fact, nowhere close. But there's a bit of a story here.

By default, Qualtrics output files don't include the IP address of the user submitting each survey, nor the user-agent header of the

user's browser. So an unskilled investigator might go to Qualtrics, download the file, see no IP address or user-agent information, shrug, and say it can't be done. "It can't be done" – because the data had to be captured at the time, it can't be reconstructed, you can't *guess* what IP address or user-agent a particular user had sometime in the past. And if it's not in the file, an unskilled investigator might think, then it wasn't captured, and it can't be obtained now.

But in fact, Qualtrics *always* captures this data. You just need to know how to ask for it. Basically, you just have to tell Qualtrics to include this metadata in the download file. I was actually the person who figured this out, not by being super clever, but by just reading the Qualtrics' manual. It seems the HBS investigators didn't read the manual, because their report didn't have that data.

Larry

Was there something else about these 20 entries that Data Colada also flagged?

Ron

Right. Not only did the entries have weird entries for the student's year, they all seemed to have extreme values (like 7's) that strengthened the ultimate conclusion of the study. Not all of them had extreme values. In fact, one row had extreme values in the other direction —1's rather than 7's. But on balance, they did strengthen the result, though not by much. In fact, as Francesca's experts later demonstrated, removing the 20 entries didn't affect the statistical significance of the paper at all. But it is true that except for one row, they all did point in the same direction.

Larry

That is weird. Is there any innocent explanation of how the extreme values might have been produced?

Ron

Turns out there is. People doing a survey in a hurry often choose extreme values. For example, if you do a survey with your keyboard, not mouse, because you realize that's faster, you might tap tab to go to the next question and choose the same number for every question. Or if you are using a mouse, just click click click straight down in a straight line, again yielding the same answer.

So someone just hurrying through the survey could produce these extreme answers because that could be the fastest way to get through?

Ron

Right.

Larry

Ok, so these were the problems with these 20 anomalous entries. The business school investigation of them was incomplete — it didn't know where they came from. But in the first interview they had with Francesca, on February 28, 2022, they presented the entries to Francesca to get her to explain how they could have been produced. Did she, on the spot, have any way to explain them?

Ron

No. On the spot, she didn't have an answer. She did say that the errors were not so obvious that she would certainly have caught them. The "Harvard" rows are strange when considered on their own. But she didn't think she would have noticed them spread among all the others, especially given that this column of data was not even used in the analyses.

Larry

So *she* couldn't explain how they were produced. Did the investigative committee (the "IC") of the business school, forensically, determine whether she had in fact produced them—by tying their submission to her machine, or her location, or anything like that?

Ron

No. Instead, they made this an inference *without* evidence: There were anomalies; she would benefit from the anomalies; therefore, she must have created the anomalies. She, in other words, must have committed academic misconduct – fraud.

Larry

After the conclusions of HBS's IC were made public, Francesca was free to retain her own experts. One of the first things her experts did was to try to unpack what really produced the 20 anomalous entries. Can you describe a little bit about how they would have done that?

Ron

Well, the first step was to use the Qualtrics dashboard to produce a report with the metadata, the IP addresses and user-agent headers, and then download that report and look at those 20 weird rows with "Harvard" as the year.

Larry

And when they did that, what did they conclude?

Ron

Two experts looking at this concluded the same thing. Basically, that there was no way Francesca had done this.

First, the entries all came from a weird, out-of-date browser—a Firefox version 5 browser running on a 2001 Windows XP computer. Almost no one was using Windows XP in 2014. Francesca was always a Mac user. Firefox 5 was released in 2011; by 2014, Firefox was up to, I am not making this up, version 27. From the start, Firefox checked for updates when it was launched, making it unusual for old versions to stay – even before Firefox added an auto-updater (in version 15). Bottom line is this user-agent line doesn't look plausible as any machine that Francesca could have been using. It didn't look like a real user in 2014; it certainly didn't look like Francesca.

Second, the 20 Harvard submissions came from 20 different IP addresses, mostly in other countries, mostly in data centers, meaning from servers, not the sort of computer Francesca had access to. None of the connections came from any place Francesca might have been.

Third, each entry ends before the next one begins. There was never any overlap, not even by one second. Real users would overlap.

Fourth, each entry specifies a different email address. But none of the email addresses exist anywhere else on the web. That's very unusual: normal users leave digital footprints. These addresses were all from free email services, like Gmail or outlook.com.

Fifth, there were some patterns in the submissions, suggesting the user was using the browser's autofill to copy answers from prior form submissions. The same capitalization repeated in some, then changed to something else in multiple subsequent entries.

Putting all those together, the picture was pretty clear to me. This was one attacker, using a form-filler and "proxies" to bounce traffic through servers in other countries, and using a "proxy switcher" function to quickly change from one proxy to another after each submission. The person was not a genius, because they didn't even try to deeply cover their tracks. But they covered enough to trick the system into saying the submissions were OK.

There was therefore no direct evidence — as opposed to an inference from motive —that this person was Francesca. By far, the better theory was that it was just a scammer trying to get some free Amazon gift cards.

Larry

So once Francesca's experts got to look at the data, they were able to show absolutely that the charge that she had fabricated these entries was just baseless. And indeed, eventually, the hearing committee of the university, — the committee that would determine whether to remove her tenure — agreed. They saw her expert evidence, and they agreed that the 20 anomalous entries were not evidence of fraud or any other misconduct by Francesca. They agreed that this was a product of—as they called it—a "scammer."

For me, this was just one more "holy shit moment" among many in this case.

The whole of the charges made against Francesca by Data Colada with respect to this paper were tied to these 20 anomalous entries.

Data Colada had noticed them. It had flagged them. It had asked the business school to investigate them. All that was completely fair.

But then the business school botched that investigation. Even more astonishingly, it blocked Francesca from conducting her own investigation *before* it concluded that she was guilty. Based on that conclusion of guilt, it asked the university to revoke her tenure. And yet once Francesca had expert support, the kind of forensic support that the business school itself was using to prosecute her, she was able to establish absolutely that she was not responsible for these anomalous entries. Is that a fair summary?

Ron

Yes, the 20 supposedly suspect "Harvard" entries were fully resolved.

Larry

I remember learning all this. Like you, I had not known anything about the prosecution against her before the business school announced publicly that she was guilty. But afterward, I recall learning that the whole basis of her guilt with respect to this charge at least had been blown up. That these anomalous entries had nothing to do with her.

And I remember thinking at that point, "Well, this is all going to be over. The business school will realize they prosecuted her on the basis of a mistake. Maybe they would reflect on their own procedural flaws that helped produce that mistake. Maybe some contrition would have manifested. But certainly, at least with respect to this paper, this process would be over."

Ron

Almost. There was a second finding that the IC made about this paper.

Based on a report by HBS's forensic expert, Maidstone, the IC said, I quote,

> some data in the OSF dataset do not appear in either of the two Qualtrics datasets for this study, that those data strongly support the hypothesized and reported results, and that some data in the two Qualtrics datasets do not appear in the OSF dataset.

That's the end of the quote. OSF stands for Open Science Framework, which is an online platform where researchers can openly share data, methods, and results to make science more transparent. Francesca had deposited the data for this study on a public platform. The IC thought they had identified a gap between the Qualtrics data and public data that could not be explained.

Larry

So, the suggestion was that Francesca must have edited the data to delete observations that weren't helpful to her conclusions, right?

Ron

Right. That was the theory, and that's what the IC concluded.

But here again, it turns out, HBS's original expert's report was just wrong.

After the IC concluded that Francesca was guilty, based on this expert report, the experts changed their minds. In December 2023, 9 months after the business school had declared to the world that Francesca was guilty, the experts issued a new report. That report said, in effect, "oops, we were wrong. There were no missing entries. Every complete response in Qualtrics had a match in the OSF dataset. And every respondent in the OSF data has a match in Qualtrics."

Larry

And so did the IC then update their report, based on the correction that their expert had made?

Ron

No. The IC made no corrections to their conclusions. They had relied on the claim that the data was missing. When Maidstone revealed there was no data missing, they said nothing.

Worse than that, HBS ended up with two parallel processes. Remember the IC finding was from March 2023. In June 2023, HBS filed a complaint with the university, starting the process to remove Francesca's tenure. In December 2023, Maidstone told HBS that rows weren't missing after all. But by that point, the tenure process was already underway. HBS didn't alert the Hearing Committee that the initial Maidstone report was wrong, and Maidstone had made a correction. The prosecution continued without once acknowledging it was based on a mistaken finding by HBS's experts.

Larry

Let's just make the point once again: Here was a second mistake in the business school's "expert" report. The first was its failure to notice that the 20 entries were by a scammer, not by Francesca; the second was its claim that there were entries that had been omitted from the Qualtrics dataset; once it checked its work, the experts discovered this too was wrong. There were no missing entries. Both times, the IC made an error that Francesca could have corrected, if she had had her own expert support *before* HBS concluded she was guilty.

But she didn't, because the rules that the business school imposed upon her forbid her from having expert support before it concluded she was guilty. It's like there's actually a reason why you allow the defendant to provide her own evidence — and test the evidence of the prosecution—before you ask the jury to decide whether she is guilty. It's Due Process 101.

<break>

Larry

In the tenure proceeding, HBS obviously realized that the scammer story and the missing rows that were no longer missing were fatal to the charges standing behind Allegation 2. They *were*

the entire Allegation 2 charge. HBS could have dropped the charges, or gone back to the IC and asked them to reevaluate the evidence. But instead, as I've described, they doubled down: they sent the paper out to a new expert to see what new problems the new expert might be able to identify.

That new expert then examined the data further and called out a different issue altogether. This charge involved 21 rows of data, some of which were the Harvard entries, and some were the rows Maidstone thought were missing. Specifically, within those rows of data, he found 154 cells that seemed to have been changed between the source data and the data used to conduct the analysis. Based on his analysis, the hearing committee—the committee that would decide whether to revoke Francesca's tenure — concluded that, quote: "all of the alterations are in the direction of the paper's hypothesis."

First, what does that mean?

Ron

If it were true, it would mean that if you examine the changes between the two data files, they all would have made the claims of the paper stronger. Statistically stronger. Again suggesting a motive for Francesca making the changes.

Larry

Well, is that true? Is it true that "all of the alterations are in the direction of the paper's hypothesis"?

Ron

No, it is not true. Not all of the 154 changes strengthen the claims of the paper. At most, half of them do. But half of the changes are in columns that weren't used in the final statistical analysis. So on the theory that Francesca was tinkering with the values to strengthen the conclusions of the paper, you'd have to explain why half of that tinkering was with data that didn't matter to the conclusions of the paper.

Francesca's lawyers made this point at the tenure revocation hearing repeatedly. The Hearing Committee apparently just missed it.

Larry

This is critically important. The whole of the Hearing Committee's argument resolves to this: "there are changes in the data, all of those changes favor the conclusion of the paper; only Francesca

had an interest in making changes to strengthen the conclusion of the paper; so therefore, Francesca must have made the changes." The whole of their argument, in other words, is an inference from motive.

But in racing to their conclusion, the committee missed a pretty important fact about motive: ***not*** all of the changes support the conclusion of the paper. Which raises a pretty obvious question about motive. On the theory that Francesca made these changes intentionally, why would half the changes she made have nothing to do with strengthening the conclusions of the paper?

Ron

Why exactly.

Larry

Okay, I want to flag this flaw, because we're going to see it with each of the four allegations made against Francesca. In each of the four allegations, we're going to see that the hearing committee concluded that all the anomalies that they were talking about, the changes that they had identified, strengthen the conclusions of the papers Francesca ultimately published. In each case, that claim is simply not true.

In every single case, some of the anomalies strengthened the finding, while some had no effect at all.

So in all four papers, the real question about motive should be this: If Francesca really was intentionally manipulating the data to strengthen her conclusions, why would she make changes to the data that both strengthened the results and didn't strengthen the results? Why would she waste her time like that?

But let's return now to the specific changes that were made with this allegation — or that remained, after the scammer charges are gone, and after the missing entries charges are gone.

As I said, the hearing committee said there were 154 changes. There were actually 160, but we'll stay with the HC's number. Half of them strengthened the results—marginally—and half did not. But for the half that did, that sounds really bad, right.

Ron

It's not great. Data values should stay the same from beginning to end. In this case, there was a clear difference between the original submissions and the final data set, and, on the surface, no obvious way to explain it.

Larry

Well, the business school thought it had a way to explain it, right?

Ron

Yes they did. Their hypothesis again was that Francesca had made the changes intentionally, by working through the data file, and changing 154 cells to strengthen the hypothesis of the paper.

Larry

154 cells out of how many? I mean, how many cells could have been changed to strengthen the hypothesis of the paper?

Ron

Well, there were 10,311 cells of data in the spreadsheet. Only some of those cells were relevant to the paper's analyses. 2,455 cells had results that could have been changed to strengthen the hypothesis of the paper.

So that means, if she was seeking to strengthen the hypothesis of the paper, there were 2,455 possible cells that could have been changed, right?

Ron

Right

Larry

But of these 154 changes, how many were changes that actually strengthened the conclusions of the paper?

Ron

80 of 154 strengthened the conclusions of the paper.

Larry

80 of the 154 strengthened the study; 74 did not. Ok, and so as you studied the 154 changes, did it seem like the cells that were changed were just randomly chosen? Did it seem like someone was just working through the spreadsheet, tweaking the results to strengthen the conclusions of the paper?

Ron

No it did not. Because as I looked at the changes, I noticed a pattern. These weren't just randomly selected cells that were different. It didn't look like someone was just moving through the spreadsheet, tampering with the values, asking "what data would be

good in this row to strengthen my hypothesis." Instead, the changes had a pattern. It looked like the data had gotten shuffled, or swapped from one part of the spreadsheet to another.

Larry

So how could that swapping have happened, if not intentionally? I mean, Excel data doesn't just get up and walk around a spreadsheet, does it?

Ron

At first I had no idea. The values in a row should stay in that row. Row 1 data can't just appear in row 9, and vice versa. But when I looked deeper into the mechanics of Excel, I was kind of surprised to learn about a feature that Excel calls "Cut and Insert." That's not a terribly helpful name. If I were naming it, I'd call it "Swap." But basically, the way it works is this: if you highlight a range of cells, press the shift key, and then drag the highlighted range, Excel will take the selection you started with, and put it where you drop it. That part seems intuitive, like a drag or move. Much less intuitive is this: Excel will also take the block of cells where you dropped it, and move them to the place you started. It essentially *swaps* the source range with the destination range. The two sets of values trade places.

Larry

Wait. So if you highlight a 2x2 block of cells — A1 to B2 — and then shift-drag that block to C8, the values in A1 to B2 will now be in C8 to D9. But what you're saying is that whatever was in C8 to D9 originally, is now in A1 to B2?

Ron

Exactly. The two blocks are swapped. Regardless of how many cells you've highlighted to move, a block of the same size gets swapped with the block you're dragging.

Larry

Why would Excel even have a feature like this?

Ron

Well, Excel is kind of a Swiss army knife of data tools. What features doesn't it have? Actually, I have wanted this feature at a couple points in my time working with Excel, but I didn't know what it was called or what its shortcut key was. And since I learned about it, I have actually used it a few times. So I think it *is* an appropriate feature to include in Excel.

But it's a dangerous feature. Excel doesn't pop open a dialog box to say "You are swapping 3 cells with 3 other cells," or 300 cells with 300 other cells, as the case may be. Excel just does it. Excel's online help does warn that "Excel doesn't warn you." But if you haven't been warned, you're not likely to notice the swap. Instead, if you're selecting a range of 8 rows by 9 columns to look at as you check data for problems such as inattentive users, you might not notice that a block of 8 rows by 9 columns has been swapped with another block. Especially when all the data kind of looks the same, all survey data, in this case all single-digit numbers from 1 to 7.

Larry

Ok, tie all this together for me. Why does this swapping function matter to the story about these 154 cells?

Ron

Well, this feature actually could explain a lot about the 154 anomalous data entries.

One area of anomalies is a block of cells 9 columns wide by 8 rows deep. When I looked at the data, I realized this 9x8 range had been swapped with a different 9x8 range. And when I realized that, this swapping function became a suspect. When I realized that this range had been swapped, it became plausible that this hidden Excel feature was actually responsible for the change. If the rows and columns had been arranged in a particular way — and actually, very similar to the way they were arranged in one version of the files found on Francesca's hard drive–then the 9 columns and 8 rows that I had found were swapped could have been unintentionally swapped by this hidden Excel feature.

And that single unintentional swap would then account for 94% of the changes observed in the file. Of these changes, 44% were changes irrelevant to the study hypotheses. And another 11% were changes to a measure that was the average of other measures affected by changes, so it would have automatically changed as a result.

Larry

So one inadvertent (or unexpected) swap, caused by shift-dragging a block of cells rather than just selecting a block of cells, could explain 94% of the anomalous entries?

Ron

Yes, it could.

Larry

What about the other 6%?

Ron

They can't be explained by the hidden swap feature. But they could have been produced by another hidden Excel feature. This command is Copy and Replace, Control-Drag. This copies cells, meaning creates a second instance of the same data, but in a new place, overwriting what's there, again without a warning message. There are two other blocks of cells, for which Copy and Replace would fully explain the observed changes. And then, yes, every change from the original file to the posted file – would all be totally explained.

Larry

So one hidden feature of Excel — the swap feature — could explain 94% of the changes. The other hidden feature — this copy and replace, control-drag feature done two times, could explain the rest. And both of these unexpected changes would happen simply because a shift key, or a control key, was inadvertently depressed as the person working with the file was moving data to line up the columns to be analyzed.

Ron

Right. The three inadvertent edits together would explain each of the 154 changes, and thus create a plausible explanation for how the changes actually happened. Meaning, whoever was "cleaning" the data file could have been a bit sloppy with how he or she was using the keyboard and mouse, inadvertently producing *exactly* these changes. Not by randomly tweaking the data to strengthen the results — these changes were not randomly distributed — but by these hidden Excel features that are, as you inadvertently trigger them, easy to overlook.

Given this pattern in the changes, this alternative explanation seemed wildly more plausible to me than the idea that Francesca was simply tinkering with the data to produce a stronger paper.

Larry

But why would anyone be making any changes to the data at all? Why don't you just move from the surveys — in lab and online — to the data that gets analyzed?

Ron

In this case, there were both lab participants and online participants. The data from the two groups arrived in two different files. Those files had to be merged. That merging sounds easy, but the columns don't line up because the two surveys had different fields due to different setups. So an RA would have had to copy the rows into a single file, rearrange columns, and get it all lined up.

Then there's the process of data cleaning. This is empirical work with data from survey respondents. Survey respondents are messy, to put it politely. RAs look for incomplete answers, people who did not take the study seriously (by writing bogus essays, for instance) or people who were obviously confused or didn't even finish the task. Those would be flagged and possibly excluded. It's a lot of data massaging. We could imagine a super-RA doing it all perfectly. But it's also easy to imagine an RA making a mistake.

Larry

And again, you said RA. You don't believe Francesca was doing this massaging, right?

Ron

I don't know, I wasn't there. But a HBS professor getting columns in a spreadsheet lined up? And merging files? It seems way above her pay grade. More like the work RAs do.

And indeed, that's precisely what Francesca told HBS: As she told them, the data merging and cleaning was not done by her, it was done by RAs. It was the *analysis* that was, and should be, done by professors. That's the norm in her field. That's what she said happened here.

Larry

Ok, so what you're saying comes down to this: Look at the pattern of the changes, and recognize first that there is a pattern. These aren't the sort of changes that would be produced by simply moving through the file, and tweaking random cells to strengthen the result. Something more general happened to produce the anomalies, and you identified what those more general steps could have been. In preparing the data, in a process that Francesca had testified included the ordinary steps to prepare the data, these hidden Excel functions could have produced these anomalous results precisely.

Ron

That's right.

Larry

That's your read. There is one more way to see the weakness in the HC's conclusion, at least it seems so to me.

The hearing committee considered two fundamentally different theories about what happened here. In one, Francesca moved through the data, intentionally making changes in 154 cells, with 80 of those changes actually strengthening the result. In the other, those changes were inadvertently produced in the process of cleaning the data, by triggering these hidden Excel features.

Ron

Yes, these are the two possible theories.

Larry

The hearing committee looked at these two theories. They concluded that there was "clear and convincing evidence" that the first theory, that Francesca had tinkered with the data to produce these 154 changes, was true.

Ron

That's what they concluded. They weren't super explicit in discussing the probabilities of each theory, but yes, they said the explanation was intentional manipulation, and so they didn't credit the hidden Excel features theory.

Larry

But then here's where—when I was studying this carefully in the process of writing the final appeal—it struck me there was another way to see why it's option two that must be true rather than option one. And here's the argument that I made.

Francesca and her experts had demonstrated a procedure — a series of steps which included these two hidden Excel features unintentionally triggered — that would have produced the 154 anomalies. Those steps are replicable; the error is replicable. Right now, we could take the original file, perform the steps constituting this procedure, and the resulting data file would have the same 154 anomalies, some strengthening the results, some not. So it's not random or magical, it's a deterministic process that gives us the anomalies that were being alleged to have constituted fraud. But on

this account, they were not fraud. They were inadvertent mistakes by whoever was preparing the data file for analysis.

On the HC's account, the changes were not inadvertent. They were intentional. On the HC's account, Francesca had worked her way through the data file, making 154 changes to the data. 80 of those changes strengthened the conclusions of the paper. Those 80 were among 2,455 possible cells that could have been changed to strengthen the conclusion of the paper. But regardless of the proportion that strengthened the paper, the claim is that she was simply moving through the file, tweaking the data by changing the results in 154 cells.

But the weird thing if the HC was correct is this: What are the chances that Francesca's random changes of 154 cells — out of 10,311 cells — would be identical to the changes produced through the plausible procedure that included these hidden Excel features?

I mean, think about it. There were 154 changes spread throughout the spreadsheet. The theory that she made those changes intentionally suggests an extraordinary coincidence: that in randomly selecting 154 changes — 80 of which strengthened the result (among the 2,455 that could have strengthened the results) — Francesca just happened to make precisely the same changes that would have been produced by the plausible sequence of steps, including the hidden Excel features, that she and her experts had identified.

How likely is that? How likely is it that Francesca would have randomly selected the very same cells in the very same pattern to produce the very same anomalies? What are the chances that she would have picked 154 cells exactly?

Ron

That is a vanishingly small probability. If Francesca were randomly modifying the file, she wouldn't happen to modify it in groups or bunches like this.

<break>

Larry

Now of course, to confirm that this alternative account — that this pattern of changes was made inadvertently, by RAs working with the data — was actually true, or even plausible, you would

have to speak to the RAs. To confirm this alternative theory, in other words, someone would have to ask the RAs who actually did the work: "Could this have happened? Did you notice something like this ever happening?"

Ron

For sure. If I were running the investigation, I would want to talk to the RAs to see what they say about how they used Excel, and check whether that lines up with any of the patterns that seem to explain what happened to the data.

Larry

And this then brings us to the second consistent theme that we will see throughout each of these four alleged cases of academic misconduct: Did they interview the RAs who were charged with cleaning the data, and producing the files that would be analyzed by the academics writing the studies?

Ron

No, they did not. The business school didn't interview the people who actually manipulated the data to produce the files that Francesca analyzed for this study.

Larry

Why? Why was there no investigation involving the RAs?

Ron

I'd say the IC thought they felt they didn't need to ask the RAs. It's the inference from motive all over again: No RA would have an incentive to manipulate data, and so if you believe the changes were intentional, there was no reason to believe they were made by the RAs.

Larry

But that's assuming the conclusion. Because, obviously, these anomalies could have been produced intentionally, or they could have been produced inadvertently, by mistake.

Ron

Right.

Larry

Why didn't Francesca interview the RAs herself?

Ron

Well, remember, when all this started, she was forbidden from talking to anyone except her two advisors. That gag order was lifted only after HBS had called her a fraud. Then she was gagged again during the tenure process. Her tenure lawyers did reach out to RAs as part of that process. But at that point, given the media frenzy, no one was eager to be pulled into this mess. Most RAs did not even reply to the request for an interview. One RA had her lawyer respond to convey she was unwilling to engage.

Larry

That's understandable, given the firestorm. But it just emphasizes again the basic flaw in this whole process: If you're going to charge someone with academic fraud — and thereby, effectively end their career — you got to do the work *before* you announce their supposed guilt. It would have taken some work to reach out to the RAs and talk to them, no doubt. But that work is only feasible *before* guilt was declared. Especially in a case like this, where *everything* is an inference from motive. Maybe it wouldn't be necessary if the evidence was strong. Maybe if you had smoking gun evidence of manipulation—like if the IP addresses for the scammer entries were actually linked to Francesca's home office, or the Amazon gift cards were used to buy her a new computer. But if your case is based on inference alone — she had a motive, so she must have done it—it is not unfair to insist that you do some work to provide some direct evidence to back up that inference. Especially if the standard of proof is "clear and convincing evidence."

Ron

That certainly seems right to me.

Larry

Okay, so let's pull this all together. For Allegation #2, here is what we've got.

Data Colada noticed 20 rows in the dataset supporting the claim that the paper behind Allegation 2 had anomalies within it. They're in the business of flagging anomalies, and so that's what they did. We should be happy there are people like this in the world. People like this flagging anomalies in academic research helps keep research honest.

It was then HBS's job to investigate whether those anomalies were actually evidence of fraud.

But as we've seen, that process was fundamentally unfair — because it effectively locked the defendant in a room and denied her any effective support to build a case to defend herself. She was not allowed to question any of the witnesses, any of the people who actually worked on the data files at least before her guilt was declared; she was not allowed to hire a forensic expert to examine and rebut the forensic expert that the HBS was relying on to conclude that she was guilty.

The investigative committee of the business school concluded she was guilty. With respect to this paper in particular, the business school concluded that she had manipulated the data, by creating these 20 rows of anomalies, to strengthen her paper—and while she was at it, to earn some extra income from Amazon gift cards.

But then that whole theory collapsed. That wasn't Data Colada's fault. Data Colada had no way to know the theory would collapse. They didn't have access to the information that would have shown there was no basis for believing she had manipulated this data. They didn't have access to the metadata that would demonstrate these were just scammer entries.

But the business school did. Yet the completely inept investigation by the business school did not discover this before it concluded that she was guilty. Specifically, it did not figure out how to get the key metadata from Qualtrics to show she did not create these weird entries. Instead, the business school just concluded on the basis of this flawed evidence that she was a fraud.

And then, once the flaw was revealed, once the mistake in their evidence was shown to them, rather than retreating, they doubled down. They sent the paper out to a new expert to search for more evidence of misconduct. The new expert discovered 154 entries in the dataset that were anomalies. On their theory of the case, Francesca had gone through the dataset and manipulated the data to produce a stronger result. They had evidence of this to support that theory, because they had two files that were different. Different because of these 154 different entries.

The hearing committee that revoked her tenure relied on these 154 anomalies to find her guilty. It first stated that the anomalies were all in the direction supporting the paper's hypotheses. That was false. They were not. Instead, the anomalies were a weird mix: some supported the hypothesis, some did not. To assume that Francesca had produced these anomalies intentionally, you'd have to explain why she manipulated parts of the data that just didn't matter.

And as to the data that did matter, you'd have to explain the patterns in the manipulation. On the HC's theory, she modified the data to strengthen her results, by altering 80 cells among 2,455 that could have been altered. But Francesca presented evidence of how those very same anomalies could have been produced inadvertently by the RAs preparing the file. She showed how the RAs could have followed a procedure that could have triggered these hidden Excel features, and produce exactly the pattern of changes that were identified. And so, as I argued in the final appeal, given this completely plausible way that these 154 anomalies precisely could have been produced inadvertently — in a process that was completely replicable or deterministic — what is the chance that the very same anomalies — in the same pattern and the same order — were in fact produced by Francesca randomly going through the file and tweaking 154 cells?

What are the chances? Exactly zero.

And yet, the HC concluded that there was clear and convincing evidence that Francesca had made these 154 manipulations intentionally. Without investigating the alternative, without interviewing the RAs who actually prepared the data, without acknowledging that in fact *half* of the anomalies don't even matter to the result, the committee concluded not that it's more likely than not that Francesca had committed fraud, but that her fraud had been shown by "clear and convincing evidence."

There's a technical legal term to describe all this: bullshit.

That's the end of our analysis of Allegation #2. Thanks to Ron, for giving voice to my anonymous friend. Thanks to my anonymous friend, who has me the chance to voice the claim that my university has committed a fundamental injustice against one of its most promising scholars.

And to those of you doing the work to understand this injustice — you, the listener — thanks for listening. Stay tuned for the next episode, when we'll turn to another of the papers she was charged with fabricating — with even less evidence than the one we've just reviewed.

One final note about Data Colada: As I said in the first episode, I have invited Data Colada to participate in these four episodes about the particular charges. I emailed each of the primary members of that collective. I have received no response from any of them.

This has been the fourth episode of this season's *The Law Such As It Is*. These podcasts are produced by me, Larry Lessig, and produced technically by Josh Elstro of Elstro Productions. You can follow this case and find the documents related to this case on our website. That website is theginocase.info. At that website, you can also find a way to sign up for our Substack, which will echo the same material on that different platform.

This is Larry Lessig. Thanks for listening.