

The Law, such as it is

Season 3, Episode 8

Larry

This is Larry Lessig. This is the eighth episode of season three of the podcast, "The Law Such As It Is." In this season, we've been reviewing the decision by Harvard to revoke the tenure of Harvard Business School Professor Francesca Gino.

The first episodes reviewed the procedural nightmare that brought about this result. The last three and a half have considered three of the four allegations of academic misconduct brought against her. This episode considers the fourth, what we're going to refer to as allegation number one. This is the last allegation, and after this episode, there will be at least one more episode that draws together what we've seen across all eight.

Allegation number one is the only allegation within the six-year limitation set by Harvard's rules. Therefore, from my perspective, it's the only allegation that should have been considered by the Business School and the Hearing Committee. That six-year limitation, as you'll remember, expressly says that charges of academic misconduct more than six years old, "may not be investigated." Somebody missed that little requirement when they launched this four-part investigation four and a half years ago, an investigation that investigated three charges that were plainly more than six years old.

I am joined again in this episode by my anonymized friend who we're calling Ava. And Ava will help us to unpack the charges as they relate to allegation number one.

Stay tuned for this one, because there's quite an extraordinary punchline at the very end of this episode. So welcome back, Ava.

Ava

Thank you, Professor.

Larry

Ava, help us understand the context for allegation number one.

Ava

Allegation one involves a paper that Francesca Gino co-authored with two other authors. The paper was titled, "*Why connect? Moral consequences of networking with a promotion or prevention focus.*" It was published in the Journal of Personality and Social Psychology in 2020. In study 3A of the paper, which Francesca led, participants were randomly assigned to one of three conditions: a prevention condition, a promotion condition and a control condition. In the prevention condition, the participants were focused on meeting one's responsibilities. In the promotion condition, they were focused on growth and advancement. In the control condition, they were focused on neither to create such a mindset. Participants initially wrote an essay on hope or aspiration in the promotion condition, a duty or obligation in the prevention condition, or their usual activities for the evening in the control condition.

Larry

Okay, but so what did they do in these three separate conditions?

Ava

In each condition, the participants read a story about a networking event, and they were asked to imagine they were the protagonist making professional connections in the story. The hypothesis of the study was that if you were in the promotion condition, meaning you were focused on dreams and aspirations, you would have a lower sense of impurity and a greater desire to network with other people, whereas, if you were in the prevention condition, focused on obligations, that would produce the opposite results.

Larry

And I'm assuming that this is a published paper, because indeed, that's what the results produced. Is that? Right?

Ava

Yes.

Larry

Okay, so what was the basis of the charge of academic misconduct here?

Ava

Well, here's a study that was produced using Qualtrics. We once again have two files, File A and File B. File A is the downloaded data from Qualtrics, the raw data from the participants' answers. There are differences between File A and File B. In particular, there are 1,066 different values between the original File A and the final File B.

Larry

And so I imagine the charge is that Francesca made those 1,066 changes in order to strengthen the results of her paper.

Ava

Yes, that is the claim.

Larry

And is that claim true?

Ava

No, indeed, of all four of the allegations, I see this as the weakest.

Larry

Let's start with the claim of the Hearing Committee that "all of the alterations are in the direction of the study's hypothesis, and without them, the data would not have supported that hypothesis." Is that claim true?

Ava

Once again, that claim is just false. 39% of the 1,066 changes were of data that was not used to test the hypothesis. 39% of 1,066 is 415. So the question, once again, is why Francesca would change 415 values for no purpose at all?

Larry

Same as it ever was.

Ava

You're right. This has been a theme across the four allegations. The thrust of the charge against Francesca is that there were changes made to the data, and those changes supposedly all supported the hypothesis that established motive. She had opportunity, so therefore she must be guilty. But with all four allegations, the premise is false in all four papers. It's not the case that all the anomalies strengthened the conclusions of the paper.

Larry

In this case, why did the Hearing Committee conclude that Francesca made these changes?

Ava

The committee claims that on a single day, January 24 2020, Francesca opened the original, unaltered Qualtrics dataset sometime in the afternoon, and that she later saved a dataset based on that data, but saved that dataset with 1,066 altered values. The committee thus concluded that in those few hours, she must have made the modifications reflected in File B.

Larry

Okay, but let's be very clear about this. What the evidence shows is just that she opened the Qualtrics dataset that afternoon, right?

Ava

Right.

Larry

When she opened that dataset, for what purpose might she have been opening the dataset?

Ava

Yeah, that's the critical first point to realize. When you're producing data from Qualtrics to make an academic study, there could be lots of reasons for opening the original dataset. For example, you might want to check the total number of participants in the study before any exclusions that occurred at later steps that could be in your notes somewhere. It could be in the data you've already downloaded and processed. But if you want to be absolutely certain, you go back to the original source. You open the Qualtrics dataset.

Larry

Okay, but the Hearing Committee assumed that when she opened the Qualtrics dataset, she was opening it for the purpose of beginning the analysis of the data in order to produce her final analysis. Is that right?

Ava

Right. They thought she downloaded the data, turned it into an Excel file, opened the file in Excel, and then worked on the data to get it into a condition that she could then make the analysis using a statistical program called SPSS. And they concluded that when she did that, or while she did that, she altered the data to strengthen her conclusions.

Larry

Okay, but I thought that she had testified repeatedly—and no one challenged this testimony—that this process of data cleaning or preparing the data to be analyzed was the sort of work done by research assistants.

Ava

That's right. Everyone agrees, and it's pretty intuitive that data cleaning was done by research assistants, not Francesca. HBS professors don't clean their own data. So if she

downloaded the raw data and cleaned it up on the 24th, it would have been a first. Or at least a first since the time when she was just becoming an academic, because by this point, she had an army of research assistants helping her with her work.

Larry

Did her research assistants testify that they had, in fact, worked on this data?

Ava

Yes, that's what they said. Her research assistant at the time, Alex Rohe, testified that he had worked to clean up that data before January 24. He specifically recalled working on the data before January 24 and asked Harvard to follow up by checking his email to confirm what he had done before that date.

Larry

Did Harvard follow up on Rohe's suggestion?

Ava

We don't know. There's no evidence that they did. All we know is they didn't introduce any other evidence about his interactions with Francesca, and didn't give her any access to that information. But the key point, as I see it, is that there was uncontested testimony that someone other than Francesca had done the boring, tedious, error prone, worked of cleaning up the data, and had done that work before January 24.

Larry

Which would mean that the file that was opened on January 24 (what we've called File A) would not have actually been the original file from which Francesca would have done the work that led to File B. Is that right?

Ava

Right.

Larry

Did the Hearing Committee have this testimony before it?

Ava

They did.

Larry

Did the Hearing Committee accept that Alex Rohe did, in fact, work on the data prior to January 24?

Ava

No. The Hearing Committee rejected the suggestion of the data that was worked on January 24 was data that had already been prepared by research assistants before that date — Alex Rohe or anyone else.

Larry

Why?

Ava

The committee gave two reasons. First, that there was no evidence that a research assistant was working with the data, “at this late stage” — meaning on January 24. Second, because no research assistant had the opportunity or incentive to alter over 1,000 data points in the direction of the study’s hypothesis.

Larry

But regarding the second reason, we’ve already noted that not all the changes were in the direction of the study’s hypothesis. 415 of the 1,066 anomalies did not strengthen the conclusion of the paper at all.

Ava

Right, so that’s an example of the committee assuming the conclusion.

But focus on the first point: that there was no evidence that a research assistant was working with the data at this late stage.

Larry

Is there evidence that a research assistant was working with the data at this late stage?

Ava

Your question kind of misses the point. While there was no evidence of a research assistant working with the data *on January 24*, there was testimony that a research assistant was working with the data before January 24. That research assistant was Alex Rohe, and that evidence is perfectly consistent with what would have been the normal practice for Francesca working with data: the RAs downloaded the raw data; the RAs cleaned the raw data, such as ensuring that the essays participants submitted were valid; and then the RAs put the data into a form that Francesca could analyze. Francesca testified that an RA would typically give her a file on a thumb drive or sometimes by email.

So while there's no evidence that the RAs were working with the data *on January 24*, they didn't need to work with the data on that date. They could have cleaned the data and prepared it and given it to Francesca sometime before the 24th and then she finally gets around to working with it on the 24th.

Larry

Okay, so this alternative hypothesis that a research assistant had been working with the data to get it into a form that Francesca could work with in the time leading up to January 24; on January 24, Francesca would have taken the data that the research assistant had been working with and then worked with it to produce File B. That alternative hypothesis would imagine that there was another file that Francesca might have been working from or with from January 24. Is there evidence or a record of a file other than what we're calling File A that Francesca might have been working with on January 24 that might already have had the data cleaning done?

Ava

Yes, there is evidence that there was a file that Francesca had on her computer on January 24 that already had the data cleaned and that she would have been working from.

Larry

What was that evidence?

Ava

The evidence showed that there was a file in Francesca's computer, which was likely the data file with the letter "R" in the name of the file. That letter matches the last initial of the research assistant, Alex Rohe. That file does not survive in the record, I guess meaning someone moved it or renamed it, and it can't be found on her hard drive in the place where it might be expected. But that alternative file could have been the source of the data that was actually analyzed on January 24.

Larry

Okay, so this sounds pretty compelling. Again, remember, the committee's theory is that Francesca downloaded the data, the raw Qualtrics data, on January 24, something she typically never would have done... was then to take that data and to go through the tedious process of cleaning it, to put it into a condition that it could be studied or worked with. And then she worked with it and modified it, tweaked it to get to File B, but this process just literally just never happened with Francesca, because Francesca didn't clean data.

The alternative theory, by contrast, is consistent with her normal practice: taking data cleaned by an RA and working with it. On that theory, she used the file that was named in a way to indicate that it was the data for the study and it had the letter R on it, that would have indicated her research assistant had worked on it. That she opened that file and proceeded to analyze the data in that file. We don't have that file. We don't know whether that file included the 1,066 anomalies or not. And so, once again, as with allegation two and

allegation three and allegation four, we have no foundation for saying that she made any of these changes at all.

Ava

Right, those are the two possible alternatives. But the Hearing Committee rejected the second alternative, which seems to me the more likely alternative, given that she says she never cleaned data herself, and I don't see a lot of HBS faculty cleaning their own data.

Larry

So why did they reject this second alternative?

Ava

The Hearing Committee found that it wasn't likely that the letter R in the file name referred to Alex Rohe because they found that there was another file elsewhere on her computer, that was also a data file also with the letter R on it. See, that other file was created before Alex Rohe was a research assistant. From that, the Hearing Committee inferred that adding the letter "R" to the file was not an indication of the name of the RA who worked on the file.

Larry

Oh, okay, so what's wrong with that argument?

Ava

Well, this again reveals how the committee didn't know their own record. At the time of this other file with the letter R, there was another research assistant working for Francesca whose last name also began with the letter R, Mindy Rock. Yes, R for Rock. Rock was working with Francesca at the time that another file was created, so this file with R actually offers *support* for the hypothesis then Francesca at least sometimes named files with the RA's last name initial.

Larry

Wow, so the Hearing Committee basically rejects the more likely alternative based on an inference that itself was based on a factual mistake.

Ava

Right. They reject the idea that R referred to the research assistant actually working with Francesca at the time because there was another file with the letter R that didn't refer to that same research assistant—apparently not recognizing that other RAs might also have last names that begin with the letter "R."

Larry

Okay, but did Alex Rohe testify that he had been working on data and had sent it or given it to Francesca? Just to be clear about this point again...

Ava

Yes, and this is what we discussed a few minutes ago. That's the craziest thing about this allegation. Alex Rohe — Francesca's main RA at the time — was one of just two research assistants that was actually interviewed during the HBS investigation. He specifically testified that he was working on the data file and implied that he had given her a copy of the data that he worked on. All that he did before January 24. The Hearing Committee thus had evidence from somebody who actually worked on the paper that he had actually produced at least a partial version of the data and had given it to Francesca before January 24.

Larry

And so if he did that, that means that on January 24, when Francesca was deciding to actually complete the analysis of the paper, she could either have taken the data already cleaned by her research assistant and worked from that, or she could have re-downloaded the data, the raw data, and on

her own, cleaned it and prepared it to be analyzed. Is that right?

Ava

Right. They had those alternatives before them — the completely ordinary process that would have suggested the data she worked with was Alex Rohe’s modified version of the data, worked and cleaned version of the data, and the completely extraordinary process that she re-downloaded the data to clean or reclean it herself — they concluded that she did the latter, which seems to me, plainly the less likely of the two. And on the basis of that conclusion, they concluded, supposedly with “clear and convincing evidence”, that she’s committed academic fraud.

Larry

Wow. Seriously, just wow. They had actual evidence from a research assistant that he had worked on the data, and on the basis of a mistaken inference about a naming convention for files produced by RAs, they reject that more probable scenario and instead adopt the less probable scenario on the way to concluding what is even a less probable outcome: that Francesca modified 1,066 values, including 415 values that had no relation to the ultimate strength of the hypothesis as an act of academic fraud.

Okay, but didn’t the hearing committee also conclude that “she analyzed these two files over the course of an afternoon using SPSS running commands in a manner consistent with repeatedly altering the data and then checking whether it improved the results”? Because that sounds pretty bad.

Ava

Yes, the Hearing Committee concluded that.

Larry

Well, that sounds bad. Is it true?

Ava

Sure, it sounds bad, but it's also not true. The finding is speculation refuted by the actual logs created by that SPSS session.

When examined closely, these logs demonstrated that she ran many commands across four different data sets supporting three different studies. She had them all open throughout the SPSS sessions, running these different commands as she was working across these different studies.

The SPSS logs thus do not demonstrate that she reran the same commands over and over in the way that the Hearing Committee claims to have found. And any duplicate commands were caused by switching between the files she was analyzing; analyzing filtered sets or subsets of the data; closing and reopening SPSS; and by the efficiency of rerunning a command rather than scrolling up in a log to find the results from that command. Again, some of the commands appear similar on the surface, but actually have different controls or different specifications for the statistical model used. Once the command differs even a little, the best interpretation is data exploring – “what result do I get if I analyze the data this way or that way?” or some modification of this, not the data tampering and retesting that the Hearing Committee claimed to have found.

Larry

Okay. But how did the Hearing Committee reach its conclusion in light of this evidence, assuming that they could read SPSS log files?

Ava

Once again enter Professor Freese.

Larry

Oh no, not Freese again.

Ava

Yes. Professor Freese. The Hearing Committee based its conclusion on a statement Professor Freese made with great confidence in his report. From the SPSS log, Freese claimed that it “shows that Gino’s work with the data involved her repeatedly opening Excel data for the study, computing some results that she had already computed, and then reopening the data file.” This behavior, Freese suggested, was very suspicious. It suggested the method by which Francesca would have committed academic fraud on Freese’s account. Francesca would have opened a data file and tested it in SPSS. She would have seen the results, then gone back to the data file and tweaked the data. Then she would have tested it again with SPSS with the same exact command, and if it still needed more tweaking, she would have gone back to the data file and changed some more.

Larry

That’s pretty clever. It’s a pretty good way to fudge just enough to get the results you need. But on this theory, why did she change the extra 415 cells of data that had no effect on the result?

Ava

This theory can’t explain that. But that’s not even the strongest evidence against Freese’s hypothesis.

Larry

What’s better than that?

Ava

The actual SPSS log, showing precisely which commands were run during these sessions. On Freese’s theory, it was the same command being run on the same data over and over again. But the actual SPSS log shows that is just not what happened.

I spent time examining the SPSS log myself. It consists of commands resulting in about 63,000 lines of code. The

interaction between SPSS's menu and log file is subtle. Many SPSS commands dump large blocks of code into the log file. Some analyses, like the mediation analysis used in the study, create hundreds of lines of code in the log from just a couple clicks in the menu. But, okay, you can take out these huge blocks and look at the actual commands that Francesca actually issued.

Inspecting the log carefully, I found that it contains commands from *three* separate SPSS sessions, meaning closing and reopening all of SPSS three times. The three SPSS sessions occur at different points in time, at about two in the afternoon, then about nine at night, and about 10:30. There are 15 file-open commands relating to four different files, and a total of 61 distinct data analysis commands across the four files in three sessions.

On some level, I understand Freese's hypothesis about re-analyzing the same file. But here's the thing: only six of those commands, six of 61, occur more than once for the dataset at issue across the three SPSS sessions – four of them twice, two three times. None more than that. And in one of these instances, Francesca ran the same UNIANOVA twice with identical syntax, one time right after another, *without* reopening the data set in between the two times. If she was changing the file and testing whether the change had the statistical impact she wanted, she would have needed to actually reload the file in order to get the changed version.

Freese does not engage with this fact, which is flatly contrary to his hypothesis of modifying the data in another program and reloading it into SPSS. His theory does not explain why Francesca ran the same command twice in a row without any reloading. The obvious simple hypothesis that she was exploring the data and trying to understand what it showed explains this behavior perfectly. Another natural hypothesis is that she's not some perfect villain whose every command perfectly achieves some evil objective. No, she's more like a normal user, kind of muddling around a program she basically understands, but also a little confused, maybe

distracted, including *innocuously* rerunning the same command. Because really, that's just not a big deal. It's not like we're paying by the command here.

Francesca also ran one UNIANOVA on a key measure only once throughout the afternoon session, again contradicting Freese's theory that she spent time that afternoon changing the underlying data behind that measure.

When Francesca ran different specifications for the same analysis, Freese categorized them as one identical analysis repeated over time. But that's wrong too. If a person were modifying the data in another program and then checking in SPSS whether the modification mattered, the SPSS command the person would run would be exactly the same each time. But Francesca didn't run exactly the same command each time. Instead, she ran slightly different commands, performing different analyses with each command, which again makes sense for someone trying to understand what's in the data, but it's flatly not what you would do if you were testing whether your modifications of the data have achieved your target results.

So again: the logs showed that there were different commands runs, meaning that Francesca was not changing the data in a separate program and testing whether that improved the results. I think Freese knows how to read the SPSS logs, but he came into this looking to find evidence that Francesca was guilty, and his whole report leads down that path, not engaging with the evidence within these very same logs, that pretty well points in the opposite direction.

I think we all know that busy researchers and important professors get interrupted. Francesca's calendar for the day shows she had an office meeting and a child to pick up, but within the general period of the timestamped SPSS afternoon analysis. Freese assumed that all data analysis would be done in one continuous session with no reloading, but that's really not consistent with what the calendar says she did that day. Now, the SPSS log does not record the time of individual

commands, just the start of each session and the order of commands within each session. Still, the logs show scattered work during the business day, consistent with a person having meetings, forgetting what she was doing, coming back to it and redoing some work she had already done. And then late at night, surely, after the kids were asleep, the logs show somewhat more purposeful work overall, in an order matching what's in the paper.

Francesca's innocent explanation thus matches what you'd expect from a busy professor and someone with young kids. Her youngest was two and a half months old at the time. Freese imagines an evil genius tweaking data— well, there's just not support for that in the SPSS logs.

Larry

So it sounds like both you and Francesca's forensic expert disagree with what Freese and the Hearing Committee concluded after examining the SPSS log.

Ava

Yes, the Hearing Committee's conclusion from this data is refuted by an informed interpretation of the SPSS log file.

Larry

Okay, so pulling this together. Basically, we have two experts: Harvard's data expert, and Francesca's forensic expert. Harvard's data expert, Professor Freese, said she downloaded the data and cleaned it and then made the changes that produced File B. Francesca's expert said she worked from a file that had been created by somebody else, and so we don't know whether the differences between the raw data and the final product are the product of her modifying it, or whether the anomalies were created by the research assistants.

Ava

Well, no, it's even more outrageous than this. Because this was a rare instance when both the Business School's forensic expert and Francesca's forensic expert actually agreed.

Larry

Wait, I thought you said Freese was Harvard's expert.

Ava

Yes, Freese was the data and statistics expert who came into the case at the very last minute to support removing her tenure. But there was a second expert who HBS brought into her tenure case: a forensic expert.

Both HBS forensic expert and Francesca's forensic expert testified at the tenure revocation hearing that they believed Francesca did *not* work from data downloaded from Qualtrics on January 24. Both experts believe that Francesca worked from data that she had copied and pasted from some other file, specifically not, as the Hearing Committee concluded, from the Qualtrics data.

Thus *both* experts are effectively agreeing that the real file that she was working with beginning on January 24, was *not* the one downloaded from Qualtrics, but some file created before January 24. And again, Francesca argues that that file would have been created by an RA, most likely Alex Rohe.

Larry

So, what basis did the Hearing Committee have for concluding that she downloaded the data on January 24, cleaned it herself, modified it, and with that modification, eventually producing File B?

Ava

No evidence. Literally none. It was a fabrication based on Freese's hypotheses about what must have happened. Hypotheses driven by their conclusion that she was guilty. She must be guilty, that this is the cleanest way to explain how she was guilty. She started from a clean, raw data file, did what she never did in the process of working with data, namely clean the data herself, ignoring the work she had hired a research assistant to do, work which the research assistant testified he had done prior to January 24, and then she

manipulated that data to produce the anomalies that are reflected in File B.

Larry

Okay, so I take it you would then say that there's not clear and convincing evidence that she modified the data to produce File B.

Ava

There is not clear and convincing evidence. There's not even more likely than not evidence. There's pretty close to no evidence at all. There's no evidence in this case to support the conclusion that she downloaded the data on January 24 and manipulated that data herself to produce the anomalies that are reflected in File B.

Larry

But if she didn't manipulate the data, what explains the 1,066 differences?

Ava

I spent some time exploring the data to try to answer that question. What I discovered is that there are good reasons to think these were errors, not falsifications. Even without access to all the evidence we would ideally want, 62% of the changes can be explained by just one of three known types of data processing errors.

This one error is an Excel "all-to-one" error. It happens when a user copies and pastes, but selects a range rather than a single cell, arranges like group of cells when giving the Paste command. If you paste a single cell into a range, the Paste command repeats the cell's value across all of the cells of the selected destination range. If the RA had committed that error, then 62% of the changes would be accounted for by that alone. You know, it's sort of a weird thing to do. If I wanted to paste a cell into a cell, I select the one source, copy, select the one destination, paste. But if someone for some reason has a different habit, like has a habit of selecting a group of cells

in the supposed destination and does that over and over, yeah, you could change a lot of data. You could override a lot of data that way, maybe without even realizing it.

I also examined how other errors can explain some of the anomalies. Errors like a flip error, where values coded on a 1–7 scale are incorrectly reversed, like seven to one during recoding. That’s an easy error for RAS to make, and that error seems to be present in this data too.

Or, the study in Qualtrics may have initially assigned arbitrary numeric coding to Likert responses.

Larry

The Likert responses?

Ava

Yes, these are the responses that measure the intensity of feelings like, strongly agree to strongly disagree. These labels make sense to people, but the data analysis requires sequential numbers. That means the data would have to be recoded in Excel or Qualtrics, depending on when the RA realized that numbers are needed and the numbers could be messed up.

Larry

Messed up?

Ava

Yes, because apparently Qualtrics sometimes converts the logically-ordered English labels into out of order numbers. You would think like, strongly agreed was one. Slightly agreed was two. Neutral could be three. But apparently sometimes Qualtrics changed those numbers so they’re not in order, and at that point you have to put them back in order. It’s strange. I don’t know why Qualtrics would mess that up, but I read multiple online threads of other users complaining about it. So we need some information from Qualtrics to determine what role this error might have played and whether this error was likely to have occurred in this dataset.

These are the kinds of mistakes known to occur when cleaning and formatting large datasets. You know, stuff happens. Computers are complicated, multiple systems connecting... And as I say, just the first of these errors, the all-to-one, covers 62% of the data anomalies.

Larry

Still, though, that's a lot of speculation.

Ava

It is speculation. But the point is, it's a way of understanding how the results we see could have been produced while the data was being prepared by the RAs rather than assuming, and they did assume, that Francesca was a fraud.

Larry

But it'd be really great, wouldn't it, if we had some way actually to see what files she had opened or what files she had used on January 24. I mean, it would be really great if there was a spy that had been sitting in the room with her that could tell us what exactly she did on January 24. Because if there'd been that spy, that spy could tell us whether it was this lost file from Alex Rohe that she had opened on January 24 or whether she had, as the Hearing Committee found, downloaded the data herself, cleaned that data and then done that manipulation of the data in order to produce File B.

Ava

What do you mean, we should have had that spy?

Larry

I mean, it would be really great to have a spy in the room like that, right?

Ava

Okay, but you know, we did have sort of a spy. There was a spy sitting in the room watching which files were opened on her computer. And that spy such as it is, it's the computer's

operating system. The computer itself recorded which file she opened that day.

Larry

What do you mean by that?

Ava

Well, the log files of the computer of the operating system would have recorded, for example, that she inserted a thumb drive and opened a certain file from that thumb drive, or if she downloaded data at a certain time from the Internet.

Larry

Wait... so the log files of the operating system of the computer she was working on would have captured and reported what actually happened that day. Did Francesca or her experts look at that log file?

Ava

Well, her team tried to get them. They asked Harvard to supply the log files so they could see exactly what happened on that day, exactly which files were opened.

Larry

And what did Harvard say to their request?

Ava

Harvard said, we don't have them.

Larry

Why didn't Harvard have them?

Ava

It appears Harvard's investigators made maybe the most basic forensic mistake that could have been made in any forensic study involving a computer.

Larry

What mistake was that?

Ava

Well, hold on to your chair here, because this is really going to blow your mind.

When the Business School opened the investigation in October 2021, the first thing the Research Integrity Officer did was to seize Francesca's computer. Francesca asked them whether they had taken an image of the computer at the time they seized it, and in an email dated November 4, 2021 the RIO told Francesca that he had made such a copy of her hard drive: "All the electronic files included in the inventory are *forensic copies*, and the original sources remain available to you."

Forensic copies. But 26 months later, when Francesca's forensic expert received the evidence from HBS, he discovered that, in fact, no forensic image of Francesca's computer had been made.

Larry

What do you mean by an image? A picture?

Ava

Not exactly, but sort of. A full copy of the entire hard drive. It's Forensics 101 that when you're doing a forensic analysis involving a computer, you basically take a snapshot of the hard disk of the computer. Not just the files that normal users access, but every single file on the computer, including system files and log files. And you do that at the start, immediately, because any delay could cause logs to be overwritten.

So if they had taken a snapshot, a full forensic image when the computer was seized, it could have been possible to determine exactly what happened on the day when the files were on that computer. You could have determined whether a thumb drive had been inserted into the computer, whether a file had been read from the thumb drive, which was Francesca's theory of what happened, or whether there was no thumb drive inserted that day, and instead the data that

produced the file that actually became File B was downloaded from the Internet, as Harvard alleged.

Larry

So why didn't Harvard have that image? Did they lose it?

Ava

No, the Keystone cops that ran this investigation never took an image of the computer's hard drive.

Larry

What do you mean, they never took an image of the computer's hard drive?

Ava

I mean exactly that they never took an image of the hard drive. What they did instead is they had a technician copy some files from Francesca's computer onto a separate drive. Not all the files, just the particular files that the RIO thought might have been involved in the research at issue. They made no effort to copy everything. They didn't use imaging software.

On some level, well, I guess I sympathized with the technician. What training did he have about what it meant to take a "image"? And did he have the right tools to do that? But HBS is a big and wealthy place. Boston has great tech talent. If they wanted this investigation done properly, they could have found someone to do it properly. But... they didn't. By the time this blunder was discovered — all these log files, of course, had been replaced with new log files, as the operating system, and its standard practice, goes out with the old and in with the new. Don't want to keep gigabytes of logs around forever, so it really does make sense for the operating system to get rid of the old logs. But there was no way to go back to see exactly what happened on that day, because by the time Francesca's lawyers discovered Harvard's mistake, the operating system had erased the log files.

Larry

Okay, so this is a little mind boggling. You're saying that if the business school had conducted their investigation in the way that any forensic investigation of technology would have been conducted, if they had gone into her office and taken an image of the entire computer disk at the time the investigation started, the image could have the evidence necessary to distinguish between two fundamentally different hypotheses:

Hypothesis 1: That she had taken data given to her by her research assistant, data that had already been cleaned and prepared for her to analyze, and then analyze that data, producing File B, or

Hypothesis 2: That she had downloaded the data from Qualtrics, and cleaned that data and prepared and therefore had done the manipulation herself.

That evidence could have been preserved in the image taken of the computer disk, but Harvard didn't have that evidence because Harvard's inept investigation had not understood that when you're investigating a computer incident, you take an image of the entire computer disk. You don't copy a few folders of files.

Ava

That's all right. HBS failed the most basic test of computer investigations. And that failure denies both HBS and Francesca the evidence to distinguish between these two hypotheses. Now, my view is that these two hypotheses are, on their face, so different in probability that you don't exactly need that forensic evidence... but still, that evidence would have resolved this question with certainty.

Larry

And we don't have that evidence again because of Harvard's own basic ineptitude.

Ava

That's right. The incompetence of the HBS investigation denied Francesca the chance to *prove* what seems clear, even without this evidence: that whatever explains the gap between the Qualtrics data and the file that she used in the analysis, the explanation is *not* that she intentionally modified the data. Whatever modifications, whatever anomalies, whatever reasons there were for the differences between the raw data and File B, the evidence Harvard claimed to have preserved but didn't actually preserve, would surely have been helpful in her effort to prove that this was not research misconduct.

Larry

Okay, so maybe if you're a prosecuting University and you realize that your own complete incompetence has denied the person whose career you are destroying the opportunity to establish your innocence, maybe at that point you say to yourself, "Oops, sorry, we screwed up. And because we screwed up, the evidence in this case is incomplete. And it's incomplete in a particular way, in a way that fundamentally weakens the defendant's opportunity for a defense that's on us, and we therefore refuse to proceed with the prosecution on the basis of this allegation." You would think that's what they would say. So is that what they, in fact, said?

Ava

No, they did the opposite of owning that incompetence.

In fact, what they did is even worse. At the tenure revocation hearing, in November 2024, the RIO blamed Francesca for this mistake. He claimed that he asked Francesca what files to copy because he was worried about her privacy. Yet, the subset of files that he actually copied included a folder called 'Personal,' which included her tax returns. If he was actually concerned about her privacy, he certainly would not have copied those files. And anyway, he told her in an email he had made a "forensic copy." Forensic. That's a technical term with a technical meaning. Copying some files from some folders is kind of the opposite of forensic.

Larry

Okay, wow. So HBS incompetence not only destroyed this evidence, but then they blamed her for it.

Okay, so there's one final bit to this charge related to allegation number one. The committee found that Francesca had "already drafted a description of the study results, suggesting she had altered the data to conform to her pre written draft."

Ava

Yeah, this bit of the charge is even more bizarre. The claim is that the Hearing Committee found a file on her computer that was drafted before the data analysis had been completed — before January 24 — that described the results that the analysis ultimately supported without the details, such as the means, stats or p-values. That supported the claim, the committee found, that she must have modified the data to conform it to this description written in advance.

Larry

Okay, so had the Business School made this charge against her originally?

Ava

No, they didn't. They didn't make this charge because the practice that the Hearing Committee flagged is completely normal within the field that Francesca was working in.

In co-authored papers, in describing and designing the study, the authors often draft a version of the conclusions to decide whether the paper is actually worth pursuing. That's especially true when a paper might include multiple studies and you have to decide which of them, which combination, is worth prioritizing. Frankly, it's just prudent. There's no point in conducting a study that, if successful, would yield a finding that gets a shrug from the field. So, the fact that there was a draft describing the conclusions that eventually were produced is a complete nothing burger.

Some people like to dump on Francesca's field. It's not necessarily the very fanciest in methodology, but in writing up the finding in advance, they're really not alone. For example, Amazon trains its employees when designing a new product or service to write the press release for that product or service before they write the first line of code. If the press release is lame or hard to write, probably the product will be a flop. For Francesca and co-authors, the finding *is* the product.

Larry

So the committee drew its conclusion that this early draft was evidence of fraud based on a completely undeveloped record—undeveloped because no one had ever suggested that there was anything wrong with his practice. And had they done so in the Business School, at the stage of the investigation in the Business School, it would have been dropped immediately: Francesca could have pointed to 50 other scholars who similarly write their conclusion up before they do the research to check whether a study is worth pursuing. Instead, it was the Hearing Committee's wild speculation that they had found a smoking gun, and they then deployed this as yet another way to support a conclusion that they had obviously started with: that Francesca was guilty.

So that's allegation one, right?

Ava

Right.

Larry

Okay, let's just pull it together. The essence of allegation one is that there are two files, File A and File B. File B is different from File A. That difference is captured in the 1,066 anomalies.

The Hearing Committee believed that those 1,066 anomalies were produced by Francesca. They found that all of those anomalies strengthened the results, and therefore that was the evidence of fraud. On that basis, they declared her guilty, but

as with the earlier allegations, the claim that the anomalies all strengthened the results is just not true. Harvard says it's true, but if you look at the actual changes and the statistics the paper reports, it is not true. This undermines their motive theory absolutely.

Also, as with the earlier allegations, the foundation for the claim — that File A is the file Francesca worked from — is equally wrong. Not only would it have been extremely improbable that Francesca would have downloaded the Qualtrics data and cleaned it herself, there was testimony that a researcher had produced the data that was analyzed by Francesca on January 24 and given it to her before January 24 in the standard way in which research assistants shared data with Francesca at that time, meaning on a thumb drive.

If the data from her research assistant, Alex Rohe, had been the data that she had worked from on January 24, then to establish she had committed fraud, we would have to compare that data, that file to File B, and any anomalies between these two might be the foundation for a claim of academic fraud.

But we don't have Alex Rohe's file. We only have a reference to it on the computer, so there's no foundation for the claim that she had modified any data at all.

And worse, the incompetence of the university's investigation likely destroyed the evidence that could have established the source of the data that she began working with on January 24. If Harvard had done what the Research Integrity Officer said had been done, if he'd actually produced a forensic copy of the computer at the time the investigation started, the log files on that copy could well have revealed what files were opened from what source on January 24, resolving finally, whether it was a thumb drive that had provided the data or data downloaded from Qualtrics on the internet.

We don't have that evidence. What we have is:

1. The ordinary practice that she would never have cleaned data herself.

2. Both forensic experts agreeing that the source of the data she analyzed on January 24 was copied from some other file. And
3. Her own research assistant testifying that he had provided data, a file, to her that he had already done at least partially the cleaning of the Qualtrics data.

Add those together, and again, it's not just that you don't have clear and convincing evidence of her guilt. It's not just that you don't have more likely than not evidence of her guilt. Add those facts together, and you have clear and convincing evidence that there is no foundation for finding that she made any modifications at all.

An integrity-driven prosecution of Francesca would have confessed this error and dropped the prosecution. But that's not what Harvard did. Harvard tripled down.

That's the end of allegation number one, which means that's the end of all four allegations that we're covering in this season of the podcast.

In the next episode, I'm going to pull all this together so you can be reminded of where we've been and have a concise way to understand the essence of the defense, the defense to establish that in fact, Francesca Gino did not commit academic fraud.

But I want to end by thanking my anonymous friend. Anonymously. No one would believe just how much work this friend has done to help understand the facts in this case and unpack the arguments. Effectively, all of that work done pro bono, which for the non-lawyers on the line means all of that work done for free. It's rare to find such talent in general. It's especially rare to find such talent willing to devote their talent simply to demonstrate a simple truth that Francesca Gino is innocent.

Stay tuned for the final episode. This is Larry Lessig.