The Law, such as it is
Season 3, Episode 9

Larry

This is Larry Lessig.

Over the course of almost a year now, I've been developing this podcast to unpack the charges made against Francesca Gino in Harvard's determination to remove her tenure as a professor in the Harvard Business School. Francesca is the first Harvard professor in the whole of Harvard's history to have her tenure revoked. As I said at the very start and throughout this podcast, it's my belief that Francesca is innocent of the charges made against her. Not just that Harvard has failed to prove its case or prove its case according to the standard it must prove, but that she did not engage in academic fraud. Period.

It's taken a long time to complete the review, and I apologize for that. This isn't my job. I've got too many jobs. In addition to the actual job I have, which is as a law professor at the Harvard Law School, I've also been struggling mightily to complete a new book. That book is now at the publisher, so I have a couple of weeks to breathe, and that has allowed me to complete the last two allegation episodes and this final episode, or at least final episode for now.

My aim in this episode is to draw together what we've seen across the earlier seven episodes. (The first episode was just an introduction.)

Think of it as a kind of closing argument. I'm going to remind you, if you've actually listened from the start of this podcast, of the arguments of the case. And I want to leave you with a clear sense of the flaws and a clear sense of the basis for my belief that Francesca is actually innocent.

This won't be a short episode. I apologize for that. I've added some breaks. Obviously, you're free to take whatever break you'd like. And if you go to the website, the GinoCase.info, you'll see we've set up an AI that gives you access to the transcript of all of the episodes of this podcast, and a cool little tool that you can use to ask the AI to interrogate the record that we've put in the AI to answer whatever other questions you might have. Okay, so let's go.

What went wrong here? I don't actually believe that Harvard had it out for anyone. Certainly, it didn't have it out for one of its

star academics. So how should we understand how Harvard let things go so wrong? And what should happen now?

As I've thought about this case now for almost three years, I think this whole disaster begins with an understandable panic. The dean of the Harvard Business School was essentially threatened. Data Colada, a group of data scientists who've taken it upon themselves to police academic research and call out data fraud, wrote to the business school in July 2021, and told the business school that they believed Francesca was guilty of academic fraud. And that unless the business school acted, Data Colada would publish their results. And in a not-so-veiled threat, Data Colada sent to HBS now available in the litigation docket, Data Colada indicated that their sources, if not satisfied, would, "take matters into their own hands."

Now this podcast has not been about Data Colada. As I said at the very start, I've been an admirer of Data Colada's work, and I certainly believe it's important that there be an effective check on academic fraud.

But what was striking about Data Colada's threat against the Harvard Business School dean, is that they proceeded contrary to their own stated principles. On their website, they say that this, and this is a quote, "our policy is to share drafts of blog posts that discuss someone else's work with those authors to solicit feedback seven days before the posts go live." I agree with that policy to give their targets a chance to respond first, it's good that they do that. They have smart reasons for doing that. Their work is better for it. But they didn't do that with Francesca. They didn't write to her and tell her the basis for their concerns. They instead went directly to her boss. And apparently the dean, terrified by the idea of nasty words on Twitter, launched an extraordinary process to determine whether, indeed, Francesca was an academic fraud.

And by extraordinary, I mean literally extra-ordinary: beyond the ordinary. Because the business school already had a procedure for addressing questions of academic misconduct. By all indications, that procedure was hammered out by the business school faculty, because that's the standard approach the business school used when introducing new policies affecting faculty. It provided the target of an investigation with certain process rights. It was a standard, regular procedure for determining whether, in fact, someone has committed academic misconduct.

That procedure, however, was apparently not specific or efficient enough for the dean, or at least the dean tweaked by Data Colada. So the extraordinary procedure that he launched changed it. The dean crafted a new way to investigate academic misconduct, and he launched it in this case, subjecting Francesca to its new terms. The dean never brought this procedure to the business school, never asked the faculty to approve it, and never told anybody outside of his own administration that he had devised a new way to determine whether someone was indeed guilty.

But before I describe the new procedure, let me just remind you of our basic timeline and the stages in this process, so that when I'm describing the details below, it will be at least as clear as it can be what exactly I'm talking about. I'm going to divide the description here into two stages. The first stage was conducted by the business school and the second by the university.

In the first stage, the business school ran an investigation to determine whether it believed Francesca had engaged in academic misconduct. That investigation began on October 27 2021 and concluded almost 17 months later, in March of 2023. In June, 2023 Francesca was put on unpaid leave by the business school after the school concluded that it believed that she indeed had committed academic misconduct.

That began stage two. After alerting her of this possibility in July, 2023, in October, then President Claudine Gay initiated a process to determine whether Francesca's tenure should be revoked. That process would be conducted by a faculty committee, seven members, what I'm going to call here the Hearing Committee. And for the next year, the business school and Francesca's lawyers fought about procedures and discovery requests to determine the facts that would be presented to that Hearing Committee. In November 2024, the Hearing Committee conducted a two-day hearing. Two months later in January, it concluded she was guilty. In May 2025 after giving her a chance to appeal that decision, an appeal that I actually wrote, Francesca's tenure was formally revoked by the university.

So it's just about four years from the time Data Colada notified Harvard of its allegations until Harvard did something that it had never done in its 385year history: revoke the tenure of an existing Harvard faculty member.

OK. Every problem in this case began with the business school's special and new procedure for determining academic misconduct.

And it is from this flaw that everything else flows. Consider just two core problems with this new procedure.

First. When you read the description of this new procedure, it's clear that the people crafting it had been inspired by the procedures that are adopted when someone is charged with sexual harassment or sexual misconduct. In that context, the procedure requires total confidentiality. The target is not allowed to discuss the fact that he or she has been charged. There are good reasons for that. If someone has been charged with sexual harassment or sexual assault and begins to discuss it with other people, then the person who has presumptively been harassed or assaulted could suffer retaliation. It's important to protect the victim of harassment or sexual assault, and in order to do that, the person charged with harassment or assault is properly told they must remain silent.

But that procedure makes no sense in the context of a proceeding to determine whether someone has committed academic fraud. Because the person who is alleged to have committed academic fraud is the only person whose privacy is at stake. If they wish to discuss the charge or investigate the charges, risking others will learn that they have been charged with academic misconduct, that's on them. But in this case, Francesca was told she could not speak to anyone. For almost two years, she lived under that gag order. And that gag order, I say, is the first flaw in the procedure.

The second flaw was its restricting Francesca from getting the help she needed to defend herself during this investigation. Because during these 16 months, the business school's investigation committee launched a process to investigate what it determined actually happened. The committee retained a forensic data analyst to help it interpret the data and to determine whether that data showed Francesca had engaged in fraud. Yet Francesca was forbidden from hiring her own forensic data analyst. She was told she was allowed just two people to support her in her defense, and she had already named two people to that position. So while the business school would have expert advice interpreting complicated data files to determine whether there was sufficient evidence of misconduct, Francesca would have no expert support to help her defend against the charge that she had engaged in fraud.

Now, this restriction is just astonishingly absurd. Why limit the defense at all? What purpose does that serve? How does the business school know that two is the appropriate number of people to help her defend herself in the context of a complicated charge of

academic misconduct? And what is the range of skill that those two people are supposed to have? She was advised to get a lawyer and to have a kind of support person, and she followed that advice. No one told her that if she got a support person, she wouldn't be allowed to hire a forensic analyst. Indeed, no one told her that Harvard itself was going to hire a forensic analyst before she had chosen her two.

To be fair, after the fact, the business school denied that it imposed this restriction upon her. I guess that's progress, because it at least shows they get why it is so absurd. But Francesca was absolutely clear that she was forbidden. And certainly having raised the possibility herself, had she not been forbidden, you can be sure she would have hired someone. She didn't because she was not permitted. that is literally what the policy says, and that is what she says the business school staff told her.

Now you might say, why didn't she just ignore that restriction? Why didn't she hire whomever she wanted? Why didn't she talk to whoever she wanted. What were they going to do? Well, what they threatened to do was to terminate her immediately. She had no choice but to follow these rules if she wanted to keep her job while she defended herself. It's not clear to me exactly how they would have had the authority to terminate her, given the procedure that they had launched was not actually the business school's procedure. But that was their threat, and it was effective in bending Francesca to their will. She stayed silent during the whole of the process that the business school launched to determine whether, in fact, she was guilty.

Okay, so flaw two then is this: there should have been no restriction on her getting whatever support she needed to defend herself. The idea that she would be excluded from hiring a data forensic analyst when the business school was hiring a data forensic analyst is absurd. The idea that the Investigation Committee would review the evidence the business school had gathered without Francesca having an expert opportunity to review that evidence and rebut it is absurd. The whole process is absurd.

And the bottom line then is that this process that the business school launched to determine whether Francesca was guilty was a catastrophic blunder. It's almost a caricature of the sort of process a business school would invent. As if the question was, how do you most efficiently convict someone, rather than, how do you most effectively determine whether someone is guilty? There is no equivalent in any proper legal context. You can't arrest somebody, lock

them in a cell, have the prosecutor develop a prosecution against them, but then deny them any real opportunity to build their own defense. Yet that is exactly what the business school did here.

Now the reason I'm calling this a catastrophic blunder is not some abstract theory about due process. This isn't some highbrow legal analysis criticizing insufficient process, as if process was something we pursue for processes sake. No, the reason this catastrophic blunder is a catastrophic blunder is that we know for certain that the conclusions the process reached were, in critical ways, completely wrong.

Consider just two. As you'll hear later in this summary, one of the core charges against Francesca, one of the charges that began this whole process, was that she had made up 20 participants in one of her studies. Data Colada pointed to these 20 anomalous entries and suggested she must have added them. The business school committee agreed. But once the business school investigation was over and she was declared guilty, Francesca had her own data analyst look at the data. And once he did, it was conclusively determined that those 20 anomalous entries had been created by a scammer, someone who was simply trying to steal the Amazon gift cards that participants got for participating in the study.

The business school hadn't discovered this. Their forensic experts hadn't discovered this. Data Colada hadn't discovered this. Instead, the business school simply assumed that what Data Colada had alleged, that these 20 entries had been created by Francesca, was true, and therefore that that was evidence of her fraud. And that conclusion grounded their conclusion that she was guilty.

Yet now the Business School and Harvard and I assume Data Colada all concede that this was bullshit. Everybody now understands that these 20 anomalous entries were not created by Francesca, they were created by a scammer. That fact points to this point about process. If the defense had had a chance to examine the evidence before guilt was determined, they would have revealed this error. If the process was as the basics of 101 due process would require, that the defense have a chance to examine the evidence and challenge it, she would not have been found guilty on this charge.

This is the first and clearest example of why process matters and why this catastrophic blunder of due process rigged the procedure against Francesca.

The second way in which we know that this initial flawed investigation produced conclusions we now know are false is even more extraordinary. As I said, to help them build their case against Francesca, the business school retained a forensic expert. That expert produced an over 180-page report that summarized and gathered the data, the evidence that the prosecution believed demonstrated Francesca was guilty. Again, she hadn't had a chance, prior to the business school concluding that she was guilty, to have that evidence reviewed by her own forensic expert. They simply took that evidence as true, as if, because it's by an expert that Harvard surely paid an extraordinary amount of money to produce, the conclusions of that report must be valid.

Yet, as we're going to hear more about later in this episode, after the business school concluded she was guilty based on this report and then proceeded to file charges to have her tenure removed based on this report, and after Francesca's lawyers spent literally millions of dollars defending her against the charges made in this report, Harvard simply abandoned the report. Francesca's experts had produced evidence about a bunch of flaws in the report. Those arguments were obviously powerful arguments, because they led Harvard to withdraw the report. Harvard didn't even believe in the report that had been the foundation to her being found guilty by the business school. Instead, late in the process, it commissioned a whole new forensic report that was released just three months before the hearing was to begin. Now three months might sound like enough time, but the details really matter here. Francesca's lawyers had submitted a response to the tenure revocation process, the Third Statute Complaint on August 1 2024 and Francesca's experts were to submit their reports on September 13. By the time the new expert report showed up, Francesca and her experts had less than 30 days to respond. When you think about experts reading other experts reports, then writing their own reports, and then the lawyers trying to understand those reports and finally explain it in briefs, 30 days is no time at all.

No fair process can presume evidence is true or valid. No fair process can introduce evidence of someone's guilt and not give them a fair opportunity to rebut it. But here, it's clear that a fair opportunity to rebut the report from the original proceeding would either have required that the report be redone or withdrawn. The faculty members on the business school investigative committee said they had relied upon this withdrawn report in concluding that Francesca was guilty. They could not have relied upon it, had she had the

opportunity to show the committee its flaws before the committee concluded she was guilty. Had she had that chance, I cannot believe this committee could have reached the conclusions it did. Had she been given the most basic due process in the first stage of this process, the outcome I am certain would have been different.

< break >

So this new procedure initiated by the dean so that he could say to Data Colada that he was responding as quickly as he could to their charges is the foundational catastrophic blunder in the story of this case. It is the mistake from which everything else follows.

But it's not the only blunder. And the second is just as big.

It's a basic principle of due process that except for some extreme crimes, there's a limit to how old a charge can be. Not because we want to give somebody the right to get away with a crime that they've gotten away with with a very long time, but because there's no fair way to reconstruct the record to be able to determine fairly or correctly or accurately whether, in fact, someone's guilty of a crime.

If the IRS called you and said, "hey, 20 years ago, you reported an expense as a deduction on your 1040 and we believe that was fraudulent, and we are going to charge you with tax fraud," that would be an outrageously unfair charge. Not because you're necessarily innocent, but because no one, 20 years after filing their taxes, has the evidence necessary to prove their innocence. Actually, I do, because I'm obsessive about keeping records, but most sane people don't. And that's why there's such a thing as what we call in the law "statutes of limitations," statutes that say after a certain amount of time a charge is time barred. Again, not because the guilty must go free, but because even the innocent can't show their innocence after the evidence is gone. As Justice Jackson put it,

> "Statutes of limitations are statutes of repose. They protect parties from the burden of defending claims after the evidence has been lost, memories have faded and witnesses have disappeared."

That principle exists in the context of charges of academic fraud as well. Harvard has expressly embraced that principle. The particular formulation that Harvard adopted was developed by the United States Department of Health and Human Services Office of Research Integrity. It states:

> "An allegation about research that is more than six years old cannot be investigated unless the scholar has continued or renewed an incident of alleged research misconduct through the citation, republication, or other use for the potential benefit of the respondent of the research record in question."

That's a very strong principle, with a really important exception. The principle is that no charge of academic fraud relating to work published more than six years before may be investigated. You cannot be forced to defend yourself about charges more than six years old because they're too old. But the exact scope of the exception is critically important. You can't be forced to defend yourself about work that's more than six years old unless the target of the investigation had

> "continued or renewed an incident of alleged research misconduct through the citation, republication or other use for the potential benefit of the respondent of the research record in question."

Now, no one denies that three of the four charges here are more than six years old. Indeed, the oldest of the charges is 14 years old, and the other two, 12 and 11. Thus, under the plain reading of that rule, those charges should not have been investigated unless the exception applies.

Yet they were investigated, and Francesca was forced to bear the costs, literally millions of dollars in legal fees defending herself against these charges, while Harvard has spent, I'm sure, 10 times that amount prosecuting her for these three charges. So how?, or Why? Well, at first, the business school never even defended the idea that they were investigating these three charges, charges that were plainly beyond the six year limit. They didn't say anything about it. They seemed not even to notice it or consider it. And when Harvard finally explained why they thought they could investigate these charges despite the limitation, why they thought the exception applied, they argued only that Gino had cited the old papers within the last six years, and thereby had "used" the old research, and they said that was enough to put them in scope for reexamination under the business school's policy.

But that claim obscures exactly how Francesca had "used" this earlier work. Because all she did was cite the papers. She had listed them on her website. She had cited them from other papers in ways that were not significant, that didn't point directly to the particular parts that were problematic or alleged to be problematic. For example, within the list of citations in the literature review of the section

of a paper these earlier papers were cited. And so the essence of Harvard's argument is that if you even list a paper on your website or cite it in another paper, that's effectively using the research in a way that exposes you to further prosecution and investigation for that research. And that, of course, means that there's no paper that's exempt from investigation, making the rule, the statute of limitations rule, an absurdity.

Now, among lawyers, this should be just a slam dunk argument. Indeed, the Provost of the University is a famous textualist, one of the leading theorists of textualism in the law. I can't believe that he would have reviewed this argument and concluded that it was appropriate to prosecute these three charges, given the way Francesca had actually used the earlier research. It is an obviously absurd interpretation of the rule to believe that merely citing an article exposes you to further prosecution for it.

And indeed, the organization that originally crafted the rule has recognized this absurdity or ambiguity and modified the rule to make it clear that it does not allow for the prosecution of a paper merely because it has been cited. As the rule now states, for the exception to apply, the scholar must cite to the portions of the research record alleged to have been fabricated. A simple citation of the paper is not enough. And under this clarification of the rule, Francesca's citations would not satisfy the exception. Nonetheless, all the way to the very end, Harvard has defended the idea that they can prosecute all four charges when it is absolutely clear, or should be clear, that under the rule, as any reasonable interpreter must interpret it, three-fourths of this prosecution was time barred.

Now again, for nonlawyers, I don't want to be misunderstood here. I don't think Francesca is guilty of any of the charges made against her. Certainly on the evidence that's presented, you couldn't possibly conclude that she was guilty with clear and convincing evidence. And thus to say that the prosecution should not have happened is not to say that she could be guilty. It is to say that it is simply wrong to allow the prosecution to have occurred, independent of whether it should have concluded that she was guilty or not. The protection of the rule is not against being found guilty. The protection of the rule is against being prosecuted.

And I would say to any academic at Harvard who has done empirical work, this part of the Francesca prosecution should concern you the most. Because are you absolutely confident that the research assistants that you use to produce the data that you then

analyze to produce your research didn't make mistakes? Mistakes that might suggest that you juice the data? Because if you're not absolutely confident, then here's a bit of free advice: don't cite the work anymore. Remove it from your website. Because someday there might be a vindictive prosecutor who comes after you, too.

< break >

Okay, all that's the process stuff, the stuff that we lawyers tend to obsess about. Normal people, most people, reasonable people, are just focused on the bottom line, not the process stuff: is she guilty, or is she not? Did she fabricate her data and commit academic fraud? Or is she actually innocent? Are these anomalies, anomalies that were produced unintentionally or accidentally by research assistants in the process of preparing her data?

In the balance of this episode, I will address that question by addressing each of the four allegations in a way to give you an easy, well, not easy, but easy, relative to reading the whole record yourself, way to understand why with each the prosecution for that allegation fails. This is just one episode summarizing at least four hours of episodes. So, of necessity, my summary is abbreviated, maybe not abbreviated enough, because this is a long episode. But if you want to understand the charge in detail, or the charge with any allegation in detail, go back to the original episode which covered that charge. I will refer you to them as we go through this. But it should be enough, this summary, to get you to see the essence of what failed with each of the four.

But if you look across all four of the allegations, there are two consistent flaws that I want to call out up front, because they'll repeat in each of the four allegations. But they're significant when seen together. Indeed, they are significant in relation to the point I've just made about the burden of proof. Because these two flaws together show, I think, absolutely, why it would not have been possible to conclude with clear and convincing evidence that Francesca was guilty.

The first and most important flaw repeatedly made by the Hearing Committee, the committee that removed her tenure or recommended to remove her tenure in each of the four allegations, is the claim that the data anomalies that had been discovered all tended to support or strengthen the conclusion in Francesca's papers.

That's an extremely important assertion, because it seems to provide motive. Everyone concedes there are anomalies in this data.

If they all support the conclusions of the papers, it would be reasonable to believe that they were introduced by someone who had an interest in supporting or strengthening the conclusions of the paper, and that person would be Francesca.

But with each of the four allegations, that claim is just false. Some of the anomalies supported the conclusions of the paper, some of them did not. And that destroys the presumption about motive, absolutely just destroys it. Indeed, rather than support the presumption that Francesca is responsible for these anomalies, the fact that the anomalies don't all support the paper instead supports the presumption that these are errors. As you'll see in each case, I describe how these errors could have been produced, like what process or procedure could have produced them. But even without a clear understanding of how they could have been produced, the conflicting character of the anomalies negates the presumption they were intentional and intentionally inserted by Francesca. That's the first mistake that is common across these four allegations. The second mistake is the mistake of an inadequate investigation.

Because remember something important about the nature of the work that scholars such as Francesca engage in, or people in her field. In a certain way, these scholars run a kind of academic paper factory. Across the period Francesca was working and the papers she was working on that are the subject of this prosecution, she employed, in total, more than 60 research assistants. The job of those research assistants was to collect and clean the data that Francesca and her coauthors would then analyze. They were the ones working with the data.

Now that doesn't mean they are, of necessity, responsible for any anomaly. But it does mean that an investigation that's trying to figure out the source of an anomaly certainly must consider the possible role of the research assistants, especially when you can't say that all of the anomalies are in the direction of the hypothesis of the paper. Across all four allegations where more than at least 12 research assistants were involved Harvard interviewed just two research assistants.

Just think about what that means.

There are errors in the data, anomalies. Everyone concedes that. Conceptually, we know they could have been produced by the research assistants, or they could have been produced by Francesca. But the investigation just presumed it was not the research

assistants, so it didn't investigate with 10 of the 12 research assistants involved, those research assistants, it didn't interview them. It just concluded, it must have been Francesca who was responsible for the anomalies.

That's not investigation. That's presumption. However sensible that presumption would be if indeed all of the anomalies supported the hypothesis of the paper, in a context where the anomalies don't all support the hypothesis of a paper, that failure to investigate is fatal to the conclusions of the Hearing Committee.

And indeed, I think if you put these two common facts together, there's really nothing more that you have to think about in the context of this determination by the Hearing Committee. When you put together the fact that, in fact, the anomalies do not all support the conclusions of the papers, and that there was no effort to actually determine who between Francesca and the research assistants could have been responsible for the anomalies, these two facts together absolutely negate the ability to reach a conclusion of guilt with clear and convincing evidence.

I would think even if you had motive absolutely established, you still would need to do an investigation that included one of the two possible sources of the anomalies. But here you don't have even unambiguous motive. Here you have ambiguous anomalies, plus a failure to investigate. Those two facts together, should have been enough for someone in this process to step back and conclude that no fair process could convict Francesca of academic misconduct.

< break >

So let's consider, then, each of the four allegations and the problems that are present with each. With each of these allegations, I've used notebook LM to create cheat sheets that might help you visualize the particular errors. You can see the cheat sheets at the website, theginocase.info as well as interrogate the NotebookLM about the case, ask it questions that my really tediously long summaries don't provide the answers to.

This is just an audio podcast, so I'm not going to refer directly to that cheat sheet as we go through this. I will instead go through the allegations in chronological order, going from the oldest to the most current.

You can call up the cheat sheet for each of them, if that would help you follow what I'm going to say. So that means we're going to start with allegation number four.

Allegation number four involves an almost 16-year-old study that Francesca and her coauthors had conducted at the University of North Carolina, which is where Francesca was teaching at the time. The study was completed just as she was moving from North Carolina to Harvard. And the aim of the study was to measure whether the act of signing an honesty pledge would affect the participants truthfulness when reporting their performance on math puzzles. The hypothesis was, if you sign an honesty pledge, that triggers in you a stronger commitment to being honest. And the objective of the study was to measure whether that was in fact true.

Now, 16 years ago, studies like these were not conducted on computers, they were conducted on paper. And what we don't have now are copies of all of the surveys that were submitted by the participants in the study. Instead, we have a file that was sent to Francesca. I'm going to call that file File A which the Hearing Committee assumed was a complete representation of all the data as compiled by the research assistants and then given to Francesca to analyze. They assumed it was the complete representation of the raw data. And the essence of the charge of fraud is that File A is different from File B, the file she used in the analysis that she ultimately produced, and that difference, the committee said, evidenced the fraud.

But the crux of the problem with this conclusion is the presumption that File A represented the raw data as transcribed from the paper surveys. The Hearing Committee believed File A was the original source file. They assumed, in other words, that this was the baseline file. And they pointed to the differences between the baseline file and the file used to evaluate the ultimate paper, and that difference, they said, showed fraud. So everything hangs on the assumption that File A was the complete representation of the raw data. Everything hangs on that being the first stage in this process, because if it was not the raw data, if it was a file that was in process of collecting the raw data, for example, then there's no foundation for saying that File B is a modification of the raw data. If instead, File A is just an interim file along the way between the paper surveys and a final data file representing the paper surveys, there would be no basis, in the law we would say no foundation, for anyone saying anyone had modified anything here. And here there's a plot twist that could have been crafted by Hollywood screenwriters. Because it just turns out that Francesca had the receipts.

As I said at the time, Francesca was working on this project. she was also preparing to move from North Carolina to Harvard. So a

moving company packed up all her boxes in North Carolina and moved them to Harvard, and many of those boxes remained untouched in Harvard's archive. Then when she was removed as a professor from Harvard, another moving company moved those boxes from Harvard to her home. To her garage.

Then when she was preparing to defend herself, in this case, she went through those boxes. And to her extraordinary surprise, she discovered that she, in fact, had the paper receipts for the payments made to the participants in the study. The receipts included participants paid and other studies conducted around the same time, so there are more participants than these in this study at issue. But what these receipts demonstrate was that not all payment records matched with the information in File A, but the payment records did match with the information in File B.

File A was missing participants, which establishes absolutely that File A was not the complete original file, and that instead, File B is the best evidence we have of the completed file. And that means File B doesn't have anomalies within it. File B is the best evidence we have of what the original surveys were.

Now you can imagine how this would be played if it were an episode of Matlock or Reasonable Doubt. And when I read about the discovery of these receipts establishing that File A wasn't actually the representation of the raw data, I thought this charge would just be dropped. The receipts proved File A was not a complete representation of the raw surveys. So there's no foundation to find, at least with clear and convincing evidence, that the data was manipulated or modified by anyone.

But the committee didn't drop the charge, and it didn't conclude she wasn't guilty of the charge. The committee apparently didn't even understand that the evidence established that File A was not a complete representation of the original data. Francesca's lawyers had said there were receipts. They testified there were receipts that showed that File A wasn't complete. They had shown the committee the receipts. But in its final report, the committee stated simply:

"She did not, however, provide those receipts."

This I said in Episode Six was just astonishing. Because at RX 626A and RX 626b in the record, there are copies of the receipts. The committee apparently did not know their own record. Now I don't blame them. Again, this record was more than 2500 pages

long, and the committee was filled with people who have day jobs, and so I don't know the lawyer who was working with the committee and how extensive their commitments were, and I know that Mistakes can be made. But the point is, once you correct this mistake and realize that the best evidence we have of the original data is File B, there is no way you can conclude that Francesca committed academic fraud.

There is certainly no way you can conclude it with clear and convincing evidence. And I mean that modality quite precisely. When you know File A is not a complete representation of the responses from these participants, that it was at most a working file, that was being modified as the responses were being incorporated and checked, you cannot conclude with clear and convincing evidence that File B is evidence of any fraud.

And then this charge got even crazier. As I described much more extensively in Episode Six, essentially, the Hearing Committee believed that after the experiment was finished, Francesca redescribed the experiment from the one that was actually conducted to strengthen the design of the experiment, like she discovered a mistake in the original design, and so just made up a description of the actual experiment to fix that flaw in the original design.

That sounds bad until you realize that the description of the original design of the experiment describes a completely braindead, logically impossible experiment to conduct. Francesca testified that what actually happened was that the original description was just a mistake, and they corrected the description in the final presentation of the work.

And it was a mistake, they said, because as originally described, the experiment could not work, or at least if causation is forward rather than backwards, it could not work. Well, Harvard responded that it wasn't just one time that that description was made. It was made twice in two separate documents, which reinforced, in their view, that the original description was the description of what actually happened. One of these documents was submitted to the IRB for approval from North Carolina University, which does make it sound persuasive, until you realize that it's absolutely clear that the second document is just a copy paste from the first because the very same typo exists in both copies. So this is one document, one document describing an experiment, a brain-dead experiment that is logically impossible to have proved what it was trying to prove, and that description, Francesca said, was just a mistake.

So, there are two theories about what happened here: Number one, a theory that suggests they simply mis described the experiments in that single draft, but conducted an experiment that was eventually described in the published paper accurately. Or, number two, that these supersmart academics planned out a logically impossible experiment and then when it didn't work, surprise, surprise described the experiment in a logically plausible and sensible way.

When you think about those two theories about what could have happened here, no fair fact finder, I believe, could find, with clear and convincing evidence that the second happened over the first. And that conclusion is buttressed by one way in which this allegation is different from most of the others.

As I said at the start, one common theme throughout these allegations is the failure to investigate completely by failing to talk to the research assistants or others that might have been working with the data.

But in this case, they did question the lab manager, and her testimony negated the idea that the experiment was designed in that completely braindead way, surprise, surprise. She would have remembered if it had been that problem. She said she did not.

But more than that, and confirming the outrageous incompleteness of the investigation, the single most important question to have been asked of the lab manager is the question from the earlier part in this description of this allegation, whether File A was actually a complete representation of the participants in the study. But alas, that question was not asked.

So with respect to the allegation that she improperly redescribed the experiment, the actual testimony that in this case negates that assumption, and, logically, it's completely absurd to imagine they initially intended and performed the experiment, because that brain dead experiment could not have been conducted. And so there's no way, with clear and convincing evidence, to conclude that she redescribed the experiment improperly. Instead, she corrected the error in the original description. And with respect to the alleged modifications of File A represented in File B, the evidence shows File A was not the original file.

So we discussed allegation four in Episode Six of the podcast. The cheat sheet on the website nicely summarizes these problems in the allegation. But let's remember this is the allegation where the

paper receipts proved there was no foundation for claiming anything had been modified. That's allegation number four.

< break >

Allegation number three, covered in Episode Seven centers on a 2014 paper entitled "Evil Genius: how dishonesty can lead to greater creativity." The paper included five different experiments, and the one at issue in this allegation was conducted in 2012. So again, we're talking about events that happened 14 years ago.

This allegation, more than perhaps any other, evinces the consequence of the failures to investigate properly or completely. Because here, as with allegation two and one, there was no effort to question the research assistants who would have worked with these data. But as we'll see in a minute here, even the most basic questions to those research assistants could have helped determine which of two very different hypotheses was true. These two hypotheses are the obvious ones. Either A, that in gathering and processing the data, the research assistants made these obvious or plausible mistakes, or hypothesis B, that Francesca had manipulated the data.

No doubt the committee and the business school had doubts about whether the research assistants would have remembered what they did with this particular project or any project. Again, that's why statutes of limitations are such a good idea, just to hark on that point again. But the point is that the questions that would have been asked in this allegation would have been the simplest questions for the research assistants to remember: did you create the data file used in the analysis? Other questions would have been helpful, like, what was your process as you handled data for this study, or which steps were you involved in? It would have been a set of questions, the answers to which would have been remembered, if indeed they remembered anything. And by getting those questions answered, we could have determined whether the thesis that Francesca had manipulated the data here had any plausible basis in fact. Talking to the research assistants seemed particularly key here given that they actually worked on the data for this paper on and off for more than 500 days.

Okay. But let's start with the specifics of the allegation. The objective this study was to determine whether cheating was associated with greater creativity. The basic design was to create conditions where the participants could cheat and then measure the creativity of the cheaters versus the creativity of the people who didn't cheat.

Now measuring creativity, of course, involves a subjective judgment. Everyone concedes that. But the way this allegation was considered, everyone, Francesca's team, Harvard and the Hearing Committee seemed to have believed that determining whether someone was a cheater was an objective judgment, or at least a straightforward judgment. That assumption is understandable because the very first determination the study made was whether someone had lied about guessing a coin flip correctly.

The study asked participants to guess the results of a coin flip. Under Harvard's assumptions about how the study was conducted, that coin flip was rigged. So if someone guessed heads, the report was tails. If someone guessed tails, the report was heads. So, in fact, everyone guessed the coin flip incorrectly, or at least thought they had guessed it incorrectly. And so, the very first question they were asked was whether they had guessed it correctly. If they said they had guessed it correctly, they were lying. And so, the first variable recorded was whether they had lied or not, liars or cheaters. And so from this perspective, it was easy to say that it was an objective determination whether someone was a cheater or not.

But there's something really puzzling about the spreadsheet that collected the data that suggests that there was something more to the question whether someone was actually a cheater. Because the raw data for this study has not been found, but the earliest version of the file that the record contains already includes coding by a research assistant processing data in a way to make the analysis possible. And that coding is pretty weird or strange, if you assume that the only question that they were evaluating was whether someone lied or not. If they said they had guessed it correctly, which, of course, was a lie, the column was marked with a one, and that one meant they had lied. If they said they had guessed it incorrectly, that was the truth, and the column was marked with a zero, meaning they had not lied. But the very next column, right next to that column, was a column marked 'cheated.' And in the first available version of the file, the baseline file that the committee used to evaluate whether there had been improper modifications, the values for that column are produced by a simple formula. The formula said: If the previous column, the column that reported whether someone said that they had guessed it correctly or not was zero, then the value in the cheated column is zero, meaning they didn't cheat. And if the value of the 'reported guessed correctly' column was one, meaning they had lied about guessing it correctly, then in the 'cheated' column, the formula said the value should be one as well. Meaning the

'cheated' column, initially in the first version of this file, produced exactly the same results as the column that said, 'reported guessed correctly.' Or more simply, the 'cheated' column reported the same results as the lying column. And the puzzle this creates is this: Why would the researcher design a spreadsheet with two columns 'reported guessed correctly' and 'cheated,' if they were intended to report exactly the same data? If cheating means lying about the coin flip, the second column is just redundant. Why do you need it if all you're reporting is whether they had lied about the coin flip?

However, if the study defined cheating in a more complicated way, a multidimensional way, based on other factors in their reporting of the respondents' behavior, the presence of two columns would be essential, because it would allow a researcher to record that while a participant may have been honest about the coin flip, they still should be considered a cheater because they had violated other rules on the survey. This distinction is the key to understanding a critical part of this allegation, because the allegation is that with 12 participants, the original file reported them as not lying and therefore not cheating. But in the final version of the file, these 12 are reported as not lying but having cheated. So since the original 12 participants who told the truth about the coin flip were eventually marked as cheaters, the Hearing Committee interpreted that change between the two file versions as fraudulent manipulation of the data, manipulated in order to strengthen the study's finding. But the critical point here is this: That conclusion is only plausible if you assume the two columns were meant to be identical.

And why would you have two identical columns in the data? But if you think the two columns were meant to report two different variables, then the difference between the two columns wouldn't even suggest that anybody had improperly changed the values. Instead, it would indicate that somebody had properly determined that these 12 were cheaters, even though they hadn't lied about guessing the coin flip correctly. Okay, so how could someone be a cheater, except by lying about whether they had guessed the coin flip correctly?

Turns out, lots of ways. For example, the survey instructed the participants not to consult information beyond what they knew. That means, if they had gone to the web and looked up answers to the survey questions, that would be cheating. You might say, "What answers could you look up as you're trying to fill out a psychological survey designed to measure creativity." But it turns out, one of the

measures of creativity in this survey was a pretty common measure of creativity. Lots of psychological studies used the same task to measure creativity. So if you wanted to cheat in answering these questions, you could have opened another web browser or tab and done a Google search on this type of question, discovered any number of examples providing answers to these questions on the web, and then just copied those answers into the survey. And if you did something like that, that would be cheating. And if we imagine these two columns were reporting two different things, column one, reporting whether you lied, and column two, beyond lying, reporting whether you did something else to indicate that you were a cheater, then we can understand why the first version of this file could have marked these 12 as not cheaters, but the subsequent processing of the file could have revealed them as cheaters.

Okay. But why, you might ask, would anybody want to cheat? It's just a psychological survey. I mean, it's not going to mean you don't get into law school. So what's the interest in cheating? What do you gain?

Well, it turns out that the people taking this survey were hired on a platform that Amazon runs called Mechanical Turk. And those people, people who work on Mechanical Turk, have an interest in having a good reputation as being reliable or informed or just good participants in a survey. So those participants would have an interest in getting good grades, because that would help them get hired on other jobs. And thus, they very much would have an interest in getting the right answer or getting as good an answer as they could. They would have an interest, therefore, in cheating. If these two columns were meant to measure two different things, the first column measuring whether they lied about guessing the coin flip correctly, and the second, adding to that judgment other evidence of cheating, then these participants certainly would have had an incentive to cheat, even if they hadn't lied.

Okay, all that's a little bit abstract. Let's get very specific. So after guessing the outcome of a coin flip, the participants in this study were supposed to do two tasks designed to measure their creativity, these two subjective tasks.

One task is called the remote association test, or let's call it the RAT task. The way that worked was that participants were given three words and then they were supposed to report on a word that was associated with those three words. So if the three words were something like 'blank, white and line,' then the associated word

would be 'paper.' The RAT task is a common task in psychological studies, and so it's easy to find many examples of the word associations on the web.

The second creativity task was called the Usage Task, and here, the participants were given a word and asked to list all the possible uses of the thing the word identified in the real world. So if the word was say, newspaper, somebody could answer the question by 'how do you use a newspaper?' by saying you read a newspaper, it's a perfectly fine answer. It's just not very creative. But if you said you could use the newspaper to make a paper airplane or wrap fish or as a substitute for wallpaper, those are also possible uses of a newspaper, but they're much more creative. And the objective was to measure creativity based on how creative someone was in identifying the uses for the thing named by the word.

But while you could look up the answers to the RAT task on the web, with this task, you really couldn't look up answers on the web. You couldn't find answers to the usage task on the web. So if you're trying to cheat, you could cheat with the RAT task. But cheating wouldn't change your ability to do well with the usage task.

Okay, here's where it gets really interesting. Because what's interesting about these 12 anomalous entries, these entries of people who didn't lie but were marked as cheaters in the final version of the file, even though in the initial version they were reported as both not lying and not cheating, the interesting thing about these 12 anomalous entries is that they were all really good at answering the RAT task, but just average in answering the usage task. That meant for the sort of thing they could have cheated on, their grades were very good. But for the sort of thing they couldn't cheat on, their grades were just average.

Now you might say, well, "If they're cheating, doesn't it take longer for them to perform the work than the others who are not cheating?" And the answer is yes, it would take longer. And indeed, for these 12, the average time to complete their survey was longer, statistically significantly longer than the average time that it took the others to perform the survey. So again, a pattern consistent with a cheater.

Now in the full episode, I go through some of the other indicators, all of them suggesting that these 12 were cheaters. But the only important point to see here is when you look at the actual answers they give across a range of questions, it would have been perfectly

reasonable for a research assistant to conclude that these 12, though they hadn't lied about guessing the coin flip correctly, could still have been properly considered cheaters.

Now, again, we don't know exactly what the research assistant who was preparing the data here was told to do or actually did, but here is another clear example of where even a simple investigation could have disambiguated between two very different hypotheses.

Hypothesis number one that the research assistant was doing exactly what I said. An investigator could have asked the research assistant, "Do you remember what you were doing with respect to this column marked 'cheated'?" If indeed they were instructed to make a holistic determination of whether the participants had cheated, they would remember that. They would say that. Then we would have absolute evidence that these changes were not anomalies produced by Francesca to strengthen the conclusions of her paper. We would have absolute evidence that they were instead the research assistant doing the work they were hired to do, preparing the data to be analyzed.

But if the research assistant had said "No, all we did was to code the column marked 'cheated' to be the same as the previous column," then you would have pretty good evidence that Francesca had modified the data, or somebody had in order to strengthen the conclusions. So the point is obvious: A simple investigation could have resolved this completely.

But even without that obvious question being asked, the point is that the strong support for why these 12 could have been marked as cheating, even though they did not lie, means that there's no way you could conclude with clear and convincing evidence that they're being marked as cheaters means Francesca had modified the data to strengthen the conclusions of her paper.

Okay. Those weren't the only problems that the Hearing Committee found in the data in allegation number three. I go through the other problems in tedious detail in Episode Seven. But I think we can summarize that tedious explication quite simply. Yes, we could identify anomalies in the data, and for some of those anomalies, we have the evidence of how those anomalies were produced.

For example, with the second subjective task, the Usage task, we can show that the research assistant copied data from the wrong location in a spreadsheet and that produced the error. But with this

error, we can be confident it was a mistake made by the research assistants. Shouldn't be any ambiguity of that. And why can we be confident that it was a mistake made by the research assistants? How do we know that copying from the wrong location was unintentional rather than intentional? Because the result of that copying was to weaken the conclusions of the paper. Francesca had no reason to modify data to weaken her study.

So with respect to that error, we can say it's error because there's no reason to believe Francesca had intentionally copied this data that weakened her study. Now, not all of the mistaken copied data weakened the study. Some of the mistakes strengthen the paper. But you know, that's the nature of errors: they go both ways.

And the critical point is this, having established that there were errors weakening the paper, that destroys the inference that these differences were intentional differences designed to strengthen the outcome of the paper, because they were the same sort of errors made in the same context.

Now, as I noted in the episode, it's not possible to explain all of the errors for the participants. One in particular, the RAT scores for four participants, didn't add up each of the answers in each of the 17 questions for four of them, they just didn't add up. Francesca can't account for that error, neither can I see how that might have occurred.

But with respect to this problem, and with respect to an allegation that there was some effort to cover up problems with the data when the data was shared with one of the members of the Data Colada team, the critical point again, is this: Any uncertainty here could have been resolved by talking to the research assistants. A simple interview could have distinguished between the scenario where it was plausible that the research assistants made the mistakes and the scenario where it seems clear that Francesca must have been the person responsible for the mistakes or the errors, and if Francesca then clear it would have been data fraud.

Okay, drawing the allegation together. It is clear that some of the anomalies in this allegation were just errors caused by research assistants, because they were plainly weakening the conclusions of the paper. And I said that fact sweeps in the other errors as well, even some that strengthen the paper, because they were all errors of the same sort. Their source is similar, and it is likely to be the same sort of error. With the 12 participants, if indeed the two columns

were to represent two different variables, the changes aren't errors. They're not anomalies. They are the research assistant doing his or her job. The evidence shows why these participants could have been marked as cheaters, even though they were not liars. The only way to prove with clear and convincing evidence the contrary assumption, the assumption that this was fraud, was to interview the research assistants. Yet again, something Harvard couldn't be bothered to do.

The bottom line for this allegation is this: the profound weakness in this allegation is caused by the failure to do the most basic investigation by interviewing the research assistants who had worked with the data. That weakness destroys the ability to conclude with clear and convincing evidence that Francesca had modified his data.

That's allegation three. Let's take a break.

< break >

Okay, now we get to allegation number two, covered in Episodes Four and a brief Episode Five. I've already described a little bit of the scandal around this allegation at the start of this episode. This is the allegation involving the 20 scammer entries. So allegation two involved a paper titled "The Moral Virtue of Authenticity," and a study conducted in 2014 that was reported in that paper. I described the paper in episode four, but that detail of what the paper tried to do is not important here.

Instead, what's important is the allegation of fraud. And The core of this allegation is what I've already described. This was the claim that there were 20 fraudulent entries created in the survey which the business school assumed Francesca must have entered herself, even though what made them seem anomalous was the error in answering the most obvious questions that anybody knows who knows anything about university life would have answered correctly. The question was, "what year were you in school?" And rather than listing first year or freshmen, the answers were Harvard. Now as a professor at Harvard and as a person who spent time on college campuses with college students, that's not the sort of error Francesca would have made, and that seems pretty obvious to me. But apparently it wasn't obvious to Data Colada or to the business school or the business school's data expert.

They looked at these anomalous entries and they concluded Francesca must have entered them. Now as I described at the

beginning of this episode, when Francesca was able to hire the experts who could analyze the data using the metadata from Qualtrics, those experts used the IT forensics to demonstrate absolutely these 20 entries were scammers that had nothing to do with Francesca. And ultimately, the Hearing Committee agreed this was not Francesca.

She was innocent of this charge. Then in addition, there was a claim by the business school that there was a gap between the data in the study and the data it was published. Data that was published to something called the Open Science Framework, suggesting that Francesca had made up the observations that were in the published version because they were helpful to her conclusion. And that charge also helped convince the Business school Investigation Committee that Francesca was a fraud. She'd made up these data. But nine months later, the expert who had made those allegations admitted that they were completely wrong and that there were no missing entries.

They failed to make that notice transparent to the Hearing Committee, but they admitted that this charge was wrong. So that's two for two. And you might have thought that having recognized that the basis for the charge had been removed once Francesca had had the ability to review the forensic evidence, you might have thought the business school would have withdrawn the charge at that point. But no, that's not what happened.

Instead, the business school hired a new expert. And the new expert went through the same data. And the new expert agreed that the 20 anomalous entries came from a scammer. But they then focused on 154 altered cells, anomalous data, and 154 cells. And between the source data and the final file, those 154 entries, this new expert suggested, suggested that Francesca had engaged in academic fraud.

Okay. Now here again, the Hearing Committee made its false claim that all of the changes strengthened the paper's hypothesis, thereby suggesting motive. But in reality, only about 80 of the 154 changes strengthened the results. The rest of the changes affected variables that were totally irrelevant to the hypothesis of the paper.

But still, there was the question. Still there is the question: how did these changes come about?

And here the point of episode four was quite important. The theory of the Hearing Committee concluding that Francesca had

engaged in fraud by altering these 154 cells to strengthen the paper, even though 44% of the 154 did not strengthen her argument at all. The assumption of the Hearing Committee was that she would have moved through the cells and made the changes to strengthen the results that would be produced by the data.

But Francesca's experts discovered a pattern in the changes, a pattern that would be explained by the RA working on the data, making a very particular kind of Excel error. when manipulating the data. That error is replicable. We can take the original file, make that error and produce the resulting file exactly. And so the question I pressed in Episode Four was this: given two alternatives, alternative A, that an Excel error explains the 154 differences, and alternative B, that Francesca moving through the file randomly made these 154 changes or differences that were observed, some benefiting her, some not. Which of the two is more likely?

And it wasn't just these two points that I was trying to get you to think about, which of these two is more likely. Instead, in episode four, I tried to make a sharper point. What I said was, what are the chances that B would produce alterations with the same structure, the exact same alterations as were reported from assumption A? Or again, when you have one way to account for the changes, through a replicable error made in manipulating the Excel file, what are the chances that a random set of changes would replicate those changes precisely?

Or again, or put more strongly: There are two hypotheses here, an innocent hypothesis and a guilty hypothesis. The innocent hypothesis is that this Excel error produced the changes. The guilty hypothesis is that these changes were the result of intentional manipulation. But, it just so happens that the random but intentional manipulation produced the same 154 changes as the innocent hypothesis does. When putting these together, my point was simple: What are the chances that the guilty hypothesis is true, given it produced exactly the same results as the innocent hypothesis? If Francesca wanted to strengthen the results by manipulating the data, she had 2455 observations in that spreadsheet to select among. 2455 cells that she could have changed. So what are the chances that she would have happened to select the same 154 changes that the innocent explanation identifies? I mean, that's a probability question, and the math to do that probability question is way above my pay grade. But even as a lawyer, I can tell you the chances are obviously incredibly small. A person randomly editing cells would edit them, yes,

randomly. But these cells have a pattern in rectangles more consistent with the kind of drag errors for which Excel is, for better or worse, notorious.

Okay. Finally, the committee was convinced that these changes had been made by Francesca, because they were all made on a day when RAs were unlikely to be working, Thanksgiving 2014, 12 years ago.

But as Episode Five demonstrated, very short episode, but one, I thought I had to add, because someone raised this question to me, you can only reach that conclusion, the conclusion that it all happened on one day by willfully ignoring the evidence provided by Francesca's lawyers and experts that there were at least five earlier files in the archives leading up to the file worked on on Thanksgiving 2014. Because, in fact, the file that was worked on on that day in 2014 was a file that had been worked on over many months. Only by ignoring these other inconvenient files that can the committee assert that she must have made those changes in a single 24-hour holiday window.

That's allegation two. And to summarize, and use your cheat sheet, if it's going to help you see what I'm summarizing here, I'm going to call that strike three. Strike one was the bogus claim about the 20 participants. All concede that was the work of scammers, not Francesca. Strike two was the allegedly made-up data. Oops, Harvard's first expert confessed no data was actually made-up. And now strike three is this, 154 alterations, 44% of which have nothing to do with the conclusions of the paper, all explained by a plausible Excel error with an effectively 0% chance that the same pattern of changes would have been produced by simply tweaking 3% of the potentially conclusion strengthening variables. Let's call this allegation the scammer allegation. This allegation two, like four and three fails.

Okay. Finally, allegation one. Now notice again, because I won't let you forget this. This is the only allegation made within the six-year period of the rules. The only allegation that should have been investigated under Harvard's own rules. The only allegation that Francesca should have had to defend against.

The paper at issue in this allegation was published in the Journal of Personality and Social Psychology in 2020. It was titled, "Why Connect? The Moral Consequences of Networking with a Promotion or Prevention Focus." Again, the details of the paper aren't

important, though we explain them in Episode Eight, in some detail. What is important is the nature of the fraud charge.

The charge of fraud involved 1066 anomalous entries that the Hearing Committee concluded Francesca must have made, because they said they all strengthened the study, and so they all evinced motive. Okay, but surprise, surprise and again, as with each of the other allegations, that is simply not true, 415 of the 1066 anomalies, 38% had nothing to do with the hypothesis of the paper. So once again, this was not a case of clear motive. It was a case of mixed evidence, which destroys the presumption of motive. Almost 40% of the changes had no impact on the study's hypothesis. On the assumption Francesca was making these changes, why would she have gone through the trouble of making these irrelevant changes.

But what got the Hearing Committee going here was that they thought they had a slam dunk case demonstrating Francesca had made the 1066 alterations. And the evidence for their slam dunk case was this. That on the day Francesca shared the written-up results with her coauthors, earlier in that day, she had downloaded the raw data for the file used in the analysis of those results from Qualtrics. And because she had accessed the raw data on that day, but then shared results based on a file that differed from that raw data, differed because of these 1066 alterations, they concluded she must have been the one making those 1066 alterations. And then, even more deviously, they alleged she had spent most of the day

> "using SPSS running commands in a manner consistent with repeatedly altering the data and then checking whether it improved the results."

Okay, that's an important quote. Let me say it again. They said she had spent the day,

> "using SPSS running commands in a manner consistent with repeatedly altering the data and then checking whether it improved the results."

Now that's a pretty strong charge. And someone making that charge is implying that they looked at the only evidence that could support that charge, the log for the SPSS session, and had drawn from that evidence conclusive evidence to support that claim. Or again, the Hearing Committee is saying that the SPSS logs support the claim that Francesca basically spent the day obsessively tweaking the data to juice the results of her study.

Now not to give away the punch line here, but that claim is utter bullshit. But to see just how utter that bullshit is, we need to unpack a little bit more what the Hearing Committee claimed.

The first bit of weirdness in the Hearing Committee's claim is the suggestion that Francesca had done something with this data that she never did. That she had started from the raw data set and cleaned the data to prepare it to be analyzed herself. As I've said, as I've explained throughout the course of these episodes, that work was the work done by research assistants. Yet the Hearing Committee's claim is that on this one occasion, Francesca demoted herself to the level of a research assistant, and spent the hours it would have taken to prepare the data to be analyzed, and then spent the hours it would have taken to repeatedly modify the data, to juice the results, to produce the paper. Weird.

The second bit of weirdness in the Hearing Committee's claim is that it is flatly contradicted by Harvard's own expert. The Hearing Committee asserted that the data for the file that Francesca turned over on January 24 must have been downloaded from Qualtrics on January 24. But Harvard's own forensic expert testified that he didn't believe that the source for the file that was shared on January 24 had been downloaded from Qualtrics on January 24. He believed the data file existed before January 24. Inexplicably, the Hearing Committee ignored Harvard's own forensic expert. And he was not the only one who said that the data file that was worked on on January 24 had been created prior to January 24. One of Francesca's RAs had testified, and again, this was a rare case, when the RAs had been interviewed, that he had worked on a version of the file before January 24 and he had given her a copy of that file before January 24. The record doesn't include the file, or didn't include the file anymore. But it did include a directory that did name a file that was named in a way to indicate that it was provided by this research assistant, Alex Rohe. The file had the initial 'R' added to the name. It was a data file, that sort of file Francesca would have used or taken from the research to use to analyze the data.

But the committee rejected this testimony. Why? Well, because it turns out there was another file on the computer that also had this mysterious letter R in its file name. But when that file was created, Alex Rohe was not a research assistant for Francesca. And so the Hearing Committee deduced that the letter R was not a convention used to refer to the name of the research assistant. QED.

But what the Hearing Committee missed is that at the time this other file was created, there was a different research assistant working for Francesca, Mindy Rock. Notice her initial is also R. And so the second file actually supports the suggestion that there was a convention to name data files with the initial of the RA working on that file.

Nonetheless, the Hearing Committee was not to be discouraged. They were convinced that the data file that Francesca turned over to her co-authors on January 24 had originated as the data she had downloaded from Qualtrics on January 24.

Okay, but the crux of the allegation made against Francesca, the most important part of that allegation is the claim by the Hearing Committee that she spent the afternoon and evening of January 24 working with these data files,

> "using SPSS, running commands in a manner consistent with repeatedly altering the data and then checking whether it improves the results."

That's a pretty damning conclusion. The suggestion was she just sat there, repeatedly running the same commands over and over again, tweaking the data to strengthen the results of the paper. And as I noted in Episode Eight, that sounds pretty bad. But as my collaborator, the Anonymized Ava, observed, it was also complete and utter bullshit.

Because Ava, leveraging her powerful data skills, actually took the log files from those SPSS sessions and examined them. The actual commands were interspersed with about 63,000 lines of code automatically generated by SPSS as Francesca used SPSS menu. But if you strip away the junk, it's possible to see what Francesca was actually doing on SPSS that afternoon and evening.

And what did Ava discover? It wasn't that Francesca was running the same commands on the same data set over and over and over, tweaking the data set to juice the results. Instead, what the actual logs showed was that of the 61 distinct commands that Francesca performed that day only six, six out of 61, occurred more than once for the data set at issue. Four of them occurred twice, and two of them occurred three times, none more than that. That means the claim that she spent the whole day perpetually running the same commands over and over and over again to juice the data was utter bullshit. It just did not happen. The objective evidence from the log file shows it just did not happen.

Now here again, some of the 1066 changes that were identified did not strengthen the results of the paper. And again, as I've said throughout the episode, where there are changes that both strengthen and don't strengthen the results of the paper, that should negate the inference of motive. Because there would be no reason to produce changes that don't strengthen the results of the paper if you are intentionally trying to fabricate results.

So where did the anomalies come from? Well, here again, Francesca offered testimony about how a common kind of Excel error could have produced those anomalies. And again, that evidence was ignored.

But for me, the real kicker with this allegation comes at the very end. Because it's conceivable we could have known exactly what happened on that computer that day. If Harvard had done its job competently, we could have known exactly from where Francesca got the data that she worked on that day. We could have determined exactly when the file that she worked on was opened and from where it came.

And we could have done that because computers are spies. Built into the operating system of a computer is a persistent surveillance mechanism. That mechanism gathers data about what you do on your computer and records it in computer logs. And that surveillance mechanism reports, for example, or could report that a thumb drive was inserted at 12:41, that it would have reported whatever identifying information there was about that thumb drive. It could have reported that a file was opened from that thumb drive at 12:42, it could have reported that the thumb drive was removed at 1:12. That surveillance mechanism reports that information and stores it in log files. And a forensic expert can read those log files and determine precisely what happened with that computer on that day.

And so, in this case, a forensic expert could have been able to read these log files and determine precisely whether Francesca did as Francesca said she did, or whether Francesca did as Harvard alleged she did. That data, in principle, was available.

But as it turned out, because of the keystone cops investigation that Harvard launched against Francesca, that data was destroyed.

Why was it destroyed?

This is really the most extraordinary part to me. When Harvard notified Francesca she'd been charged with academic misconduct, it

seized her computers. They told her that they had taken, "a forensic copy" of the "hard drives" of the computers.

Now what that means for anybody who knows anything about conducting an investigation involving computers is that they basically take an image of the computer. I don't mean a photograph. I mean they take an image of the state of everything on the computer at the time the investigation begins. An exact copy of the entire hard drive of the computer. And sometimes even more, but at least the entire hard drive of the computer, so that they're able to perfectly reconstruct the entire state of the computer at the time the charge of misconduct is made.

It turns out Harvard didn't do that. Instead, they just sent a technician down to Francesca's office and had him copy off of the computer certain files they thought would be relevant to their investigation. Which means all of the log files were vulnerable to being overwritten depending on when the computer was fired up again, because the log files have an automatic routine that deletes them after a certain amount of time.

So When Francesca was finally able to hire a forensic expert, and that forensic expert was able to do the work necessary to defend Francesca, they made a request to Harvard to produce the copy of the hard drive. The forensic copy of the hard drive. Harvard reported they didn't have one. They didn't have one because all they had were copies of files taken from the computer. And by this time, the log files were all gone.

The files that could have shown, absolutely, Francesca's innocence were gone because of the inept investigation Harvard had launched against Francesca.

Now Ava commented at this part of Episode Eight that she understood and felt sorry for the technician. What did he know about computer forensics? But Harvard's a big institution. The business school is an especially rich institution. And the idea that they would run a third-rate investigation of the data forensics necessary to establish whether a tenured member of their faculty had committed academic misconduct is itself misconduct. The simplest evidence Francesca could have pointed to to objectively establish her innocence, oops, Harvard destroyed through its own incompetence.

Okay, that's allegation one.

The committee fabricated a scenario that Francesca was doing something Francesca never did: preparing her own data. Despite

Harvard's own expert testifying, that's not what happened; despite Francesca's testifying, that's not what happened; despite Francesca's RA saying he had prepared the data files in advance of this one magical day, so that's not what happened; despite the fact that the objective evidence from the SPSS data log shows absolutely without doubt that the scenario of manipulation that the Hearing Committee concluded Francesca had engaged did not happen.

Nonetheless, the committee concluded that there was, "clear and convincing" evidence that she had done what there was literally no direct evidence she had done. And then, to add insult to injury, the incompetence of Harvard's own technology staff meant she had no opportunity to use simple direct evidence to prove what wasn't her burden to prove, but to prove her innocence.

Let's call this keystone cops allegation number one. And that's the last allegation against Francesca. Let's take a break.

< break >

In late 2024, *the Atlantic Magazine* published an extraordinary piece titled "The Fraudulent Science of Success." The piece began with an account of Francesca's case and relied heavily on people like Nick Brown and James Heathers, who were in the business of analyzing data sets, looking for anomalies. The piece then pivoted to talk about Juliana Schroeder, who was one of Gino's coauthors. And after Francesca's story broke, Schroeder began to audit their joint work. They launched a project called the Many CoAuthors Project, which gathered all the joint work and tried to isolate errors in the work. Indeed, the project discovered many errors. None of the errors it discovered were tied to Francesca. Critically, the piece then pivots to Schroeder's own research, and it reveals that in Schroeder's own research, there were problems that were uncovered during the self audit. The author of *The Atlantic* piece presses Schroeder to account for these errors. And after confessing that she was, "deeply ashamed" of the errors discovered, the article describes her response as follows.

> Still, she said that the source of the error wasn't her. Her research assistants on the project may have caused the problem. Schroeder wonders if they got confused. She said the two RAs both undergraduates, had recruited the women at the gym, and that the scene there was chaotic, sometimes multiple people coming up on them at once, and the undergrads may have had to make some changes on the fly, adjusting which participants were being put into which group for the study. Many things went wrong from there, Schroeder said. One or both RAs may have gotten ruffled as they tried to paper over inconsistencies in

> their recordkeeping. They both knew what the experiment was meant to show and how the data ought to look. So it's possible that they peaked a little at the data and reassigned the numbers in a way that seemed correct.

But *the Atlantic* author wasn't satisfied with that account, and two pages later, he continues,

> I spoke with Schroeder for the last time on the day before Halloween. She was notably composed when I confronted her with the possibility that she had engaged in data tampering herself. She repeated what she told me many months before that she definitely did not go in and change the numbers in her study, and she rejected the idea that her self-audits had been strategic, that she had used them to draw attention away from her own wrongdoings.

And then in the final pages of the essay, the author struggles with the ambiguity and with the increasing recognition of many in the field, that the field is marred with many similar problems, many examples of data just not adding up.

But what's striking to me when I read stories like this is how vigorously people hold on to the least likely explanation for the problems they've identified. As I've said throughout these podcasts, this is a weird field of science where the work is almost factory produced, where there are many people involved in the production of the science, and where, frankly, not much actually hangs on the results. This is not like Boeing screwing around with safety standards. No plane is going to fall out of the sky if there's an error copying data in an excel sheet. So, it's not hard for me to understand how people would not take seriously enough the need for precision in the work that they do as they prepare these data. Especially when you understand who's preparing the data, undergraduate research assistants, people with a million other things going on, people for whom absolute accuracy may not seem as important as answering that next snap.

I don't say any of this to excuse data fraud. I don't say any of this to suggest that data fraud doesn't exist. But I do say it to suggest that anyone raising a question about the work of someone else in this field, or I'd say any field, needs to start with the obvious and only with compelling evidence move to the unlikely. They need to start with the presumption that this is error and only with the most compelling evidence move to the conclusion that this is data fraud. Or, as the Atlantic author put it, data tampering.

Because between error and evil, the vast majority of cases will be error, and in that world, the real evil is attributing to others an intent which you don't have the evidence to support. That's what I believe happened in this case.

It was totally appropriate for Dana Colada to raise questions. Those questions should have triggered a careful investigation by the Harvard Business School. They didn't. They triggered an embarrassment. A plain denial of fair process, denying Francesca an opportunity to understand, to explain or to defend the basis for the charges against her. And then once the bureaucracy got going defending its decisions, the machine could not stop. Until it did stop, with the judgment from the Hearing Committee that ignored the failure to investigate the most plausible alternative accounts for the anomalies that they reflected upon and who repeatedly asserted what was plainly false, that all of the anomalies strengthened the conclusions of the paper.

Of all the simple ways to understand what went wrong here, here is the simplest. When you have an allegation of data manipulation or data fraud, and in fact, the anomalies are ambiguous, as they were in this case, that should end the investigation until there is clear evidence that the target of the investigation did the manipulation. Evidence, such as the evidence that was lost with Allegation One when the university failed to preserve a forensic image of the computer, the investigation should end.

Now you might say, "Well, what if an evil genius recognizes that's the rule? What if they create changes that don't strengthen the paper, while they also create changes that do strengthen the paper?" I agree, that's a real problem. And I agree we need to take special steps to investigate and prosecute that problem. But no one alleges in this case that Francesca was that sort of evil genius, and even if they did, there is no evidence in this case to show, with clear and convincing evidence, that that's what she was, or that's what she did. The discipline that this process demanded was for someone to ask cleareyed and directly, could the evidence here meet the standard? No one asked and answered that question.

At some point, serious people in this process must have recognized the embarrassment that was the business School's process. They did nothing to stop it. At some point, serious people realized that the foundation for the judgment made by the investigation committee of the business school had crumbled. They did nothing to stop the process. They just hired new experts to find new reasons

to continue the prosecution. At some point, serious people realized the fundamental unfairness in forcing Francesca's lawyers after spending a year preparing for one case to remake their case within 30 days to defend against a slew of new charges by a new expert, but they did nothing to stop it. And at some point, whether the President of this University or someone in his administration must have looked back at the work that had been done and seen the problems and recognized the injustice, but they too, did nothing to stop it.

And to this day, an army of lawyers defending and rationalizing their work continues that work, they continue to press. Indeed, press even further, as I've learned, they've now tried to force me to turn over my emails claiming that I can be deposed because I would have evidence to answer the question whether Francesca was guilty.

"What do you expect?" a friend said to me. "You've become a traitor." Really? A traitor? A traitor to an institution whose brand is "Veritas" — truth.

What I see in truth is four allegations raised against Francesca found to be true by a Hearing Committee, but that no fair application of the standard that that Hearing Committee was to apply clear and convincing evidence could find those allegations were true. And to the extent I've succeeded in raising in your head questions about any of them, we should recognize that there is a question only because Harvard failed to conduct a proper and adequate investigation.

I'm here today, this case is here today because its target, Francesca Gino, refused to shut up and go away. When the business school concluded she was guilty, they tried to get her to quietly resign, as they and the university have done with many other professors in the history of Harvard who turned out to embarrass Harvard because of their misbehavior.

But Francesca refused, because Francesca did not commit academic misconduct. She fought back. She fought back with all of the resources she and her family had. Even then, that wasn't enough, and it required the help of others, including the help of Bill Ackman, who, after careful study of these allegations, concluded, as I did, that the allegations are false.

So where should Harvard go from here?

Well, first: I'm not under any illusion that a podcast is going to remake the world. I'm pretty sure that Harvard will continue to ignore what we've said here, what the story reveals, and just get on with the many other fights that it has, some of which I think are the

most noble fights in America today. There's litigation happening now. I'm not sure what will happen in that litigation. I'm not sure what could happen in that litigation.

But I still think Harvard should step up and do the right thing. At a minimum, the President of Harvard should say it was wrong to investigate and prosecute Francesca for at least three of the four charges. And if they want to continue to insist on the correctness of the prosecution of the fourth charge, allegation number one, then they should have that allegation evaluated by an independent fact finder, an actual fact finder, an actual person who understands how to evaluate evidence against a standard. Maybe a former judge. Indeed, there are former judges on the Harvard Law School faculty. They could listen to the evidence and make an evaluation of whether there is any possibility of finding with clear and convincing evidence that for the one allegation, allegation number one, Francesca Gino had committed academic misconduct.

And when that process concludes that there is absolutely no chance that a properly applied, clear and convincing standard could have concluded that Francesca is guilty, then the university needs the courage to admit it, to admit they got this one wrong and to correct the mistake they created.

God knows the error was understandable. I wrote the appeal to the Hearing Committee's judgment, and we filed it just at the moment the University was under its most extreme pressure from President Trump and his administration. I know they were distracted. I know a lot was going on. I'm not saying anybody was malicious or evil.

They were just not courageous. They were just not committed to Veritas. The bureaucracy got going, its inertia was unstoppable. No one was going to have the courage to stand up and say, this is not right. It was easier for everyone to just go along, easier for everyone to just let it happen.

When I was a student in college, I spent a summer in Eastern Europe and the Soviet Union. It was 1982. On a very hot day, I was in a bus crossing between Romania and Bulgaria. The line at the customs checkpoint was endless. I think it took us an hour to get across that border. I spent the time looking out the window. This was 1982, there was no TikTok or Instagram, but the images out the window were extraordinary. Nonetheless, it was Eastern Europe in the 1980s. The cars were not BMWs, yet they were beautiful and

elaborately filled with all sorts of things. Boxes tied to car roofs. Trunks not quite closed, but fastened to stay shut.

A couple cars ahead of us, there was a truck. The truck was carrying cages of geese. At one point, the truck hit a bump, and the gates of the truck opened, and a bunch of the cages fell onto the road. Geese were everywhere, just as many racing into the traffic as off the road. But rather than slowing the slow roll of cars and trucks, the rolling continued. Truck after car after bus just rolled over the geese. They squealed as they were crushed. But no one could be bothered to stop to clear the way to fix the problem.

I've thought about that scene many times across my life. I've thought about it here as I've tried to help Francesca tell her story. I've so much wanted to see someone stop, someone have the courage to stop. Someone break the rhythm. Someone have the courage to do what's right.

But courage doesn't mark us at this time, in this place, even at this university. And for that, I am very sorry. But that sorrow is but a fraction of the sorrow I have when I think about the injustice that Francesca Gino has suffered. We should be better than this, especially, *we* should be better than this.

< break >

I think this is the last episode of this podcast. It's the last for now. I am extraordinarily grateful to Ava for her incredible work across the course of this investigation, helping us understand the facts and discovering many of the reasons why we can conclude with confidence that Francesca is innocent. I'm grateful to Francesca for letting me tell this story. I don't know how anyone goes through what she has gone through, certainly not with the poise and integrity and courage that she has shown. There have been moments of weeping. You've heard some of them on this podcast. But I see only strength. Maybe it's easier to be strong when you know your rights, but maybe it's harder. Maybe knowing your rights and yet facing what she's facing is the most depressing reality, I hope, I don't know.

And finally, I'm grateful to you, listeners. Many of you have written and some of you have offered support. We have no clear ask of you, other than to spread the word to the extent you think there's word to spread. I hope some will help reform the reality of the process at Harvard and if not just at Harvard elsewhere too, because some of what I've heard is parallel stories about this kind of problem everywhere across the world. We can't let that happen without

fighting to resist it. And if you can help in that elsewhere or here, please do.

Thanks to Josh Elstro of Elstro production for working through these difficult podcasts and the corrections and the edits necessary to make sure that they are presented as professionally as he presents them.

This is Larry Lessig.

Thanks for listening.