

FTC Putting Disclosures to the Test Workshop
September 15, 2016
Segment 3
Transcript

JAN PAPPALARDO: Welcome back from lunch. I'm Jan Pappalardo. And I'm very pleased to welcome you to the session on Impact on Decision Making and Behavior, where our panelists will discuss studies that evaluate the impact that disclosures have on consumer decision making and behavior. We have a wonderful panel for you today.

But before I introduce them, let me make a small announcement. I understand that there are cards available, if those in attendance would like to ask questions. You can write your question on the card, wave the card in the air, and someone will pick it up. And we will try to see if we can address the question during the discussion session.

So we have a wonderful panel for you today. Our first panelist is Lillian Ablon, an information scientist and Professor at Pardee RAND Graduate School. Next, Idris Adjerid, Assistant Professor at Mendoza College of Business, University of Notre Dame. The next presenter will be Ginger Jin, Director of the Bureau of Economics at the Federal Trade Commission and Professor of Economics at the University of Maryland. And last, but certainly not least, we have Adair Morse, Associate Professor, Haas School of Business, University of California, Berkeley.

LILLIAN ABLON: Thank you very much. So today we've been hearing a lot about disclosures. And I'm going to talk about a study we did to look at consumer attitudes towards one type of disclosure, specifically data breach notifications. I'll give you some motivation for our study, focus on the methods for conducting our research, and then also briefly share some of our findings.

So as a little bit of motivation. Data breaches have become commonplace, and no sector is immune. From banking, to retail, to health care, to the government, millions upon millions of pieces of personal, financial, and health information have been stolen. Often in response to these types of breaches, companies will send out disclosures. They'll send out data breach notifications to the affected consumers.

We wanted to take a look at this practice of sending out data breach notifications to see if they're helping. So a little bit about data breach notifications. All but three states right now require companies to notify people about the loss of personal information. These notifications by state are not uniform, and there's no federal law. The purpose generally is twofold. First, to allow people to take quick action to reduce their risk caused by the data breach. And then to create incentives for companies to improve their data security.

There are studies looking at the cost and the consequences of the companies, but we wanted to focus on the consumer response and the consumer attitudes. Specifically, we looked at three different things. First, the frequency of breach notifications and the type of data lost. Consumer response to the notifications, to the company, and to the company's follow on actions. And then the perceived cost of the breach to consumers. So we'll go through those three findings.

But first, let me tell you about our process. We use something called the American Life Panel survey instrument for this study. The American Life Panel is a nationally representative panel of over 6,000 individuals aged 18 and up, who have agreed to participate in occasional online surveys.

The ALP, the Rand ALP, ensures that it's a nationally representative sample through a number of ways. First, the panel population is recruited through random digit dialing, in order to get the most representative sample of the US population. For a particular study, when we're ready to conduct a study, we draw from the population using stratified, random draws, based on the weights of various demographic groups. Age, race, gender, income, things that we've already collected we're recruiting the population.

So for example, based on past surveys, we knew that males were more responsive than females. So a greater population of females received our particular survey. And then after we get responses, ALP weights the responses in order to ensure that they're nationally representative and weighted against the 2014 current population survey.

ALP is internet based, and for those who don't have internet access, it's provided to them. And so it also allows for real-time pulse of the American public. It also yields a relatively high response rate. ALP surveys generally get between 75% and 80% response rate. So those are some of the pros. It's useful for policy makers to get a pulse of the American public. It's repeatable, nationally representative, and yields a high response rate.

However, responses are based on consumer recall, which we all know are certainly not perfect. Also, consumers may say one thing but act in a different way. For example, consumers may say that they change their pin or password after receiving a notification, but maybe they don't always change their pin or password. And then consumer response and behavior may change over time. We're simply getting a snapshot of the American public.

So our particular survey was conducted the last two weeks of May of last year, 2015. We wanted to get 2,000 people as respondents, so we over-sampled. We received a 78% response rate. We closed our survey on June 1st of last year. And the breach at OPM, the Office of Personnel Management, was disclosed three days later on June 4th, with notifications sent out later in the summer. So again, we were getting a snapshot of the American public. Perhaps sentiment and attitudes towards receiving notifications had changed had they been asked after having received the OPM breach notification.

Let me tell you very quickly about some of our findings. So the first thing we wanted to look at was how often this happens and what kind of data is lost. We found that 26% of our respondents recalled receiving a breach notification in the last 12 months. This extrapolates out to an estimated 64 million Americans. We also asked how many over the population, over their entire lifetime. That was 43% and there were very different differences in demographics. But in the last 12 months there were no differences, no statistically significant differences in demographics.

Of those alerted, over 50% said that they received more than one notification. And this raises an interesting question of breach fatigue. Are consumers becoming so desensitized to receiving so

many notifications, that they ignore the important information and the more they discount the notification or the disclosure. Possibly the case, especially given what we've heard about earlier today.

Respondents learned about the breach most of the time from the company, but certainly quite a bit of the time through other means. And this raises the question of the optimal timing to send out data breach notifications.

The most common types of data compromised were reported to us as credit card information. Although, at least 21% of information is information that is not easily changed. Things like your social security number, your blood type health information.

The next thing we looked at was how consumers responded to the company, to the company's follow on actions, and to the notification itself. Well this was pretty good news for companies. 77% of consumers said they were satisfied with how companies responded. And consumers tended to be pretty loyal. Only 11% said that they stopped doing business with the company. In some respects 11% is high, in some respects it's low.

Often, in addition to sending out notifications, companies will hire the services of a data breach resolution provider to offer things like free credit monitoring or identity theft protection. This is an example of what I received from Anthem. So when offers like this were presented to consumers, 62% said that they accepted these offers. This is much higher than what we heard from other studies. So this may have been an instance where consumer recall was certainly a factor. Other studies had found that it was in the single digits of percentage of people who accepted free credit monitoring.

In addition to accepting these offers for free credit monitoring, we wanted to know if consumers did any additional actions, any additional steps to protect themselves. And here's a place where perhaps they didn't actually-- where consumers said that 51% of them said that they changed their pin or password. Whether or not they actually did, who knows.

And then we also wanted to know what they would want, how they would prefer a company to respond after a data breach. And they were most favorable, or most in favor, of companies notifying consumers immediately, taking measures to prevent future breaches, and offering free credit monitoring or similar services.

The third thing we wanted to know was how consumers estimate the cost of the data breach. We wanted to know the perceived cost, the inconvenience cost, rather than the actual cost. And we found that 32% said that there was a zero dollar cost to them. Of those that reported a non-zero dollar amount, \$500 was the median amount. And 6% put the inconvenience cost at \$10,000 or more. And for this particular group, it tended to be older and those with less than a high school diploma.

So are data breach notifications serving their purpose? First, do they allow people to take quick action, reduce risk? Well on the one hand, no. 44% already knew of the breach before they received notification. On the other hand, 78% said that they took some sort of additional action,

whether it be accepting offers or free credit monitoring, or doing things like changing their pin or password.

Now do they create incentives for companies to improve data security? On the one hand, no. 77% said that they were satisfied with how companies responded, and 89% remained quite loyal to the company. On the other hand, breaches appear to be on the rise, so that might be an incentive for companies to improve their data security.

I spoke very quickly because I wanted to get it all in 10 minutes. But this plus a lot more information is freely available in our report. We also have an infographic. Thank you.

[APPLAUSE]

JAN PAPPALARDO: Idris will be our next speaker.

IDRIS ADJERID: Good afternoon. I'm very happy to be here. Thank you for the opportunity to present to you guys. So today I'm showing a paper on the role of framing, it in terms of how consumers react to privacy disclosures in this case. And this is work with Alessandro Christian [INAUDIBLE]. And we were kind of motivated by, I think, what a lot of the papers today are motivated by is this idea of well, when are disclosures effective?

And we thought that one condition, one important condition, in addition to this idea of usability, and clarity, and attention, is that when you're provided a notice, you react to the objective facts of that notice. If one situation is risky and one situation is not, you change your behavior. And that's predicated on this assumption that consumers are pretty rational and liberative in terms of how they make decisions.

We thought a countervailing scenario is this idea that, well, maybe consumers are impacted by things that are not normative, that are not related to the objective facts of a particular context. In which case, we can say that disclosures are going to be a little less effective. Right? That's predicated on this idea that consumers are pretty predictably sensitive to [INAUDIBLE] irrational. They have susceptibility decision biases and heuristics. So that's some of the motivation for why do some of this work.

And in particular, we were inspired by this idea that privacy context in particular, and other contexts as well, are pretty dynamic. So you have settings where protections are changing fairly rapidly, either over time or between different service providers.

This is a good example. This is actually infographic on what data Facebook collects. This is in 2005. And the further out you go on these cocentric circles, the further-- the more it's available by default. So if you're all the way out for the circle, it's available by default to the entire internet. This is in 2005. In 2014 looks a little different. So you're seeing that much more data is being collected, and by default it's being shared much more broadly.

So when you're thinking about settings where protections are changing fairly frequently, you have two things happening at the same time. One is the objective facts. The objective protections

are changing. The objective data collected is changing. And the other thing that's being manipulated is the relative perception of protection over time between services. So there's a comparative effect. There's also an objective effect.

So that's what we study. And we have this hypothesis that basically consumer disclosure is sensitive-- the sharing of sensitive information is going to be impacted by the relative perception of privacy protection of these disclosures. Even when the objective disclosure is actually identical, the objective facts are identical. And just so I don't leave you in too much suspense, that's pretty much what we find. That these relative perceptions are pretty powerful and predictable effects on consumers behavior.

In terms of implication for policy, we think that's part of the story for limited effectiveness of disclosure in some settings. If we find predictable effects of non-normative factors.

So how do we actually do it? We actually have a range of experiments. We do a few different experiments. I'm presenting one of them. And we recruit individuals on a few different platforms. So the one I'm presenting today is on Amazon Mechanical Turk. We also run it on Prolific Academic, which is another crowd sourcing platform for research. We present this study-- that's a nice little chime. It's very--

[LAUGHTER]

--therapeutic. It's not bothering me at all.

[LAUGHTER]

So this is-- it really wasn't. I mean I'm sorry. I didn't mean to-- I wasn't being sarcastic.

So this study was presented to these participants as a study on ethical behavior. We tell them we're interested in understanding different ethical behaviors. And we manipulated between conditions what was disclosed to them as part of this notice, this disclosure. And we used a few different types of disclosures across experiments.

In this particular experiment, we're using a text disclosure. And then we measured different behavior by asking to share sensitive information with us. And some of it is pretty intrusive. And we asked them to either admit this behavior or not.

So this is our design. It's a little complicated, so I'll just go through it quickly. And I can talk to anyone afterwards, if they're interested. But we had four conditions, and we told participants they were involved in two different studies.

What we would do is, in the second study, participants would have the objectively identical set, identical disclosures. So you're either shown a low protection disclosure or a high protection disclosure. But what's changing is that, whether you're perceiving some kind of increase or decrease in the protection from the first study to the second study. And what's kind of useful about this design is, that in the second round, we can actually evaluate both-- we can disentangle

the effect of the relative change in protection and the effect of the objective difference in protection.

So what did these disclosures actually involve? The high protection disclosure was basically saying that your responses are anonymous. The low protection disclosure just said that you're identified. Whether you look at pornography, that's going to be identified back to your e-mail or not. Fairly straightforward.

We use simple disclosures for a reason. We actually wanted to sidestep this issue of attention, and complexity, and usability. We just said, well, how do people react? Let's assume we've solved all the problems of complex disclosures. How do people behave? These are some of the questions we asked them.

I've had kind of a retrospective moment. You know, my kids are getting a little bit older, and they ask me like, what do you do Dad, like at work?

[LAUGHTER]

I'm a little nervous about telling them I just ask people really uncomfortable questions. And I have all kinds of things about terrible things people have done. So I feel like I need to find a new method at some point. But for now this is working really well.

[LAUGHTER]

We look at the sensitive disclosures. And what we actually find is that this perceived increase in protection actually results in about a 7% increase in the probability of sharing sensitive information. Whereas a decrease, results in almost a 10% change in the probability of sharing information. That's actually holding the objective disclosure completely constant. It's the identical disclosure that you've shown.

In contrast, the objective difference, anonymous identified, actually only results in about a 5% difference. In the second, it actually doesn't have any effect. So we're seeing that pretty powerful effect of these relative factors, and almost no effect or very little effect of the objective change in the disclosure.

Thinking about the methodology a bit. We think there's a few advantages of studying disclosure in this way. The first is that we're able to actually capture behavior by participants. So a lot of privacy research uses hypothetical scenarios. And you think that if you capture actual behavior, that's almost always preferred. Pretty cost-efficient to run.

So I mentioned that we ran several experiments. I'm just presenting one of them. Using Mechanical Turk or other crowd sourcing platforms, plus a survey-based methodology, allows you to run a number of different iterations. So if you have a very nuanced set of hypotheses, it's much easier to run this kind of an approach.

We don't think we use too much deception. And we don't pretend to be some kind of company or anything. We tell them we're researchers. We are generally interested in ethical behavior for a different purpose.

So some deception, not as much as you might have in some other settings. We don't tell them it's about privacy. Which is, I think, some other presenters mention is a bad idea. You don't want to prime them to think about these issues.

False information, lying, is actually a core part of the design. By randomisation we assume that individuals engage in these behaviors at similar levels between conditions. So we're actually looking for people to lie. Do you admit to this behavior? Whereas if you ask people to disclose, for example security numbers, you wouldn't have a ground truth. You wouldn't know whether they're actually telling you the right number or just making something up that's fake.

So in this design, you can get a nice idea of whether they were influenced by these different changes in protection, in terms of their privacy sensitivities. It doesn't require us to develop an app or an add on, things of that nature. I'm not good at that stuff, so this is great for me. So that's always good. But of course not perfect. So there's some things to think about.

This setting is still a little bit stylized. It's not perfectly reflective of how people make decisions online. Like do I share my location information the same ways that I admit to looking at pornography? Probably not. So there might be some disconnect there.

The researcher effect. When they know it's research, people tend to assume that it's a safe place. That we're kind of good actors. And it's difficult to manipulate objective changes and risks in that kind of a setting. So that's can be a bit of a problem.

Longitudinal effects in this design are a little bit difficult, to have studies that see behavior over time. People may not believe you that it's about ethical behavior. They may be onto you. And that's always a risk with some deception. And again this idea of ethical behavior not translated into other settings.

OK. So definitely some limitations to think about. Overall, we think that this is some interesting evidence and supporting this idea that framing the way that disclosures are presented has a powerful effect on behavior. That's a little bit non-normative, which I think has implications.

And it highlights that online experiments can be good when you have this nuanced set of hypotheses, you need to run things multiple times. But it may not be great of external validity. Realism is really core to what you're asking. Or if you want to study people over long periods of time. So that's it for me. Thank you.

[APPLAUSE]

The paper is online if anyone would like to find it.

JAN PAPPALARDO: Thank you. Our next speaker is to Ginger Jin.

GINGER ZHE JIN: How should I do it? [INAUDIBLE]

JAN PAPPALARDO: I think just press that button.

GINGER ZHE JIN: Thank you. Thank you so much for coming to FTC. Before I take the director's role of FTC Bureau of Economics, I have been a Professor of Economics at University of Maryland for 15 years, devoting most of my research time to study informational asymmetries in the market. Especially, quality disclosure and verification. So I'm glad to have the opportunity to share with you the research have done by me and other economists in this area.

Just to be clear, I want to focus on a particular type of information, which is truthful quality disclosure. So economists worry about informational symmetry. For example, sellers may have more information about their product quality, and that information not known to consumers could generate all sorts of inefficiency in the market. So you were thinking truthful quality disclosure should address that problem. Right?

Any economic inefficiency rooted in those information asymmetries should be solved by such information disclosure. And moreover, in reality, we see the information could come from the firms. Could come from a third party, including a government agency or a private agency. Or it could be mandatory or voluntary.

So in this environment the real question we want to ask is, it seems like it should work in theory. Does it really work? Does it really effect people's behavior? And if we're thinking about how would this change consumer's behavior, a lot of experts in this morning have talked about how the disclosed information could catch people's attention, make sure they comprehend the information.

But in terms of real behavior, economists are most time interested in, does this really change their choice of product? Does it really make them sort and match? So I'm going to give you a few examples on that.

On the seller's side, you would wonder when a disclosure regime is available. Do sellers really pick up that option? Do they really disclose information? Did they adjust their price? Did they adjust their quality? Did they actually enter and exit the industry according to the information disclosure regime? And sometimes did they really take the opportunity to game the system?

So in the interest of time, I won't be able to review the whole literature. I just want to give you some examples so that you will have a sense of the positive examples and negative examples.

So talking about positive examples about consumer choice. One study I have done was about restaurant grade, having grade cards. So in Los Angeles County in 1998, the county government had decided to post a letter grade ABC on the front door of the restaurant window. So you can see the example here, C was posted in a corner of the window.

Twelve years later, New York City decided to adopt very similar policies. So on your right hand side is the A actually from New York City. So a lot of studies, including my own, as well as

studies by New York City officials, have shown that this kind of information does work on the consumer side. It raises consumers' perception a great deal. Almost everybody knows the grade card exists. And over 80% of people do consider data when they choose a restaurant.

It also see the revenues become very sensitive to the grade. A grade restaurants do receive higher revenue and C grade do receive less revenue. So to some extent, this works on the consumer side.

Well I could give you another example where it doesn't make any changes at all. I mean this is a famous example cited a lot by the media. I'm not aware of an economist actually documenting no response from this disclosure. But some experts have seen the reason that this does not work because we just don't know how to act upon it. And even the Homeland Security Department itself would admit that this does not communicate anything.

Well if you're thinking that no response is the worse outcome from disclosure, you might be too optimistic. Actually, a study done by Jan and then our staff economist at our Bureau Jim Lecho, has shown that in the context of mortgage loan disclosure-- when I think in 2002, had proposed a disclosure that included a mortgage broker compensation scheme-- and Jan and Jim have done a lab study to show how consumers really understand this information. And how they will make their choice in a hypothetical setting.

What they find is sort of concerning. What they find is that consumers got confused about which loan was less expensive and end up having mistaken choices, which sometimes lead them to choose more expensive loans. It also generates a bias against mortgage brokers and makes them less competitive-- make them into a competitive disadvantage, and therefore, sort of leading to some wrong choices.

So just to give you a very short sum up, that's the consumer response. Could be positive, if they make more informed choice in response to the information. Or they're not responding to the information at all, or sometimes they got a wrong impression and make the wrong choice. So that's the consumer side.

If we switch to the seller side, do they really take up the option of making a disclosure? If we come back to the theory, the theory will predict, well when this option is there, everybody should take it. Why? Because if your quality is above average, you definitely should take it. If your quality is below average, well given everybody above average is taking it, if you're close enough to the average but still below average, you should take it. So that unraveling would go on, until the worst one.

However, in reality that does not work that way. If you come back to the familiar situation of hygiene grade card. This is in Maricopa. They actually try to define grade card pretty similar to what we have seen in the Los Angeles or New York.

And so here you can see that they define A as, you don't have major violations in priority of a priority foundation violations, but you may have some minor core violations. And B is worst

than that. C is even worse. And they disclose everything, actually, online. But the restaurant can choose whether to participate in this or not.

So I'm showing an example that this restaurant was a C. And then she was now participating for a few times and became a B. And exactly how many restaurants actually take up this option? Is it 100% as the theory predicts? It turns out only 57% choose to disclose in this context. And even more surprisingly, not every A chooses to disclose.

In fact, 48% of the restaurants, if they choose to disclose, they would have an A. But they choose not to disclose. Then you'll say wow, this is a complicated situation. There is a lot of other information going on, so this may not be a really clean test of the theory.

So we actually run a similar test in the lab. Just say, OK, you as a sender, you have a secret number 1, 2, 3, 4, 5. You choose to disclose or not. You cannot lie. And then as a receiver, well I either see the disclosed number or I see a blank report. I try to guess what exactly is behind that report if a report is blank. And we give enough economic incentives so that they will mimic the reality.

What's the result? We see about a 70% to 80% of time, senders will choose to disclose. This goes up a little higher over time if you do this repeatedly.

However, if you look at people's guess of the real report, of the blank report, they still guess it as somewhere around 2 instead of 1. In theory, they should think only the worst ones do not disclose. But actually, they sort of overcast that. And that actually encourages some senders to hide behind the silence.

You would imagine, well, not only everybody not taking up that option. There could be some price or policy changes. And the price changes are pretty obvious. That's the price will increase if you disclose at a high quality, and would decrease if you were low quality. And if you cannot adjust your price or capacity, there will be a queue for high quality ones.

And then consumers will sort and match. To some extent, that's a good thing. But that also means that not everybody would have their welfare improved because of this information, as they now have to pay a higher price for higher quality.

And again, coming back to the central question, is whether the firms really improve quality if you allow them to disclose their quality. And that's kind of a central question. And we see a positive answer in a restaurant hygiene situation, where not only we see more A restaurants but also we see better health outcomes. Like less hospitalization because of food borne illnesses, and less salmonella conditions, and so forth. So that seems to be a positive outcome.

And I want you to take a minute thinking about how could the negative outcome occur in this situation? It actually occurs in health care as well. In New York and in Pennsylvania in 1991 and 1993, they actually mandate grade report cards about cardiac physician's mortality rate. And you would imagine, OK, if I have a lot of patients and very low mortality rate, I'm a good surgeon. And consumers will come to me and I should sort of receive other benefits. It turns out that the

survey results shows that there is some level of cherry-picking going on. Because now you measure me according to the mortality rate of my patients, I will want to choose a healthier patient to operate on, so that the mortality rate will be lower.

And what's the result of that? That could be the sicker patients have a harder time to find surgeons to have the necessary surgery. And as a result, they may end up in the emergency room and other kind of health care situations more often. And that turns out to be even more costly.

So the lesson I want to emphasize is that even truthful, quality disclosure can be a double-edged sword. On the consumer side, the consumer may take it right and make an informed choice, or they may understand it wrongly and make no response or even mistaken response. On the seller side, not every seller would take up that option if it's voluntary. There could be price changes, and there could be quality changes, and some sellers may even game the system. So I will stop here. Thank you very much.

[APPLAUSE]

JAN PAPPALARDO: Thank you. Adair Morse is our next speaker.

ADAIR MORSE: Great, thanks. OK, so I'm going to talk about a paper. But I've changed the title a little bit. I also want to talk about generally engaging with the private sector and testing. So with that in mind, let me tell you a little bit about the paper first.

So a number of years ago, I did a paper with the national chain of Payday Loan stores. And the idea here is, even if you take a very neutral perspective on Payday Loans-- which I try to, I do research in this area-- and think that the pricing is fair and that it's not predatory in some sense, one has to wonder whether there are some people that have a cognitive limitation or have a bias in their use of Payday Loans. That it's not optimal for some people. And so the idea here is to think about disclosure that de-biases at the point of informing. And this very much reflects the session before lunch.

First, we have to think about what is it that might be causing mistakes at that point. And then how can we target information in a way to de-bias that particular moment of decisions. And so we did this field experiment.

Let me show you the field experiment. We're going to use envelopes. At a Payday Loan store you get your cash in envelopes. And so the idea here, what is the potential problem? Potential problem one is that people may not internalize APR rates. The APR rate is already a mandatory disclosure, but maybe people are not internalizing that because when you go in the Payday Loan store, the pricing is very transparent in dollars.

So what we do then is we make the rates, in some way, very transparent by putting them in comparison to other rates that people may already have. So the Payday Loan rates are high, relative to any other kind of consumer loan. So this makes that very transparent. Here's the APR that you will be subject to.

The second potential problem is that maybe people fail to add up over time, and they focus on a single decision. And so here we express this in dollars, rather than rates. And we do an add up comparison of how to think about the cost of a Payday Loan in comparison to another type of product.

And the results are that the dollar treatment, the second one, is effective not in changing someone's behavior today-- we weren't trying to do that-- but to change their savings behavior between now and the next cycle. We were able to reduce the recurrence of re-borrowing by 10%, which is a large number. The APR, which economists love, doesn't matter at all.

And one of the lessons here I like is that we forget sometimes, in particularly academics we forget in our realm, that the real life situations are not how our theories necessarily predict. So people live month-to-month, and dollars is what matters. And this adding up and forcing them to think about their monthly cycle is really very powerful.

But more generally, I'm going to turn away from the paper now. More generally, the paper advocates for understanding-- and this is very much in the spirit before lunch-- the specific cognitive bias that may lead to suboptimal decision making. And then designing remedies, it could be disclosure or something else, designing remedies that target that particular decision at that particular moment in time. I think that's a very, very interesting and powerful way to think about disclosure.

Now I'm going to change gear a little bit and just talk about this engagement with private industry. So in finance we do a number of field experiments. Now in private industry, you have all kinds of challenges of implementation. In this case, we had to train the clerks to be uniform across customers. You have to deal with randomization that one store is not the same as the other. One day of the week is not the same as another. And that clerks can't be expected to keep track of where you gave each what treatment. And all this kind of thing you have to plan ahead, and you have to be super uptight about it.

Now more generally, whenever I presented this study. When I was writing it, people asked me, why on Earth did the lender agree to do this? You're giving them information that obviously doesn't make the Payday lending industry look good. And this is a very important question because we as researchers want to engage with industry to learn things in a real world setting. Sometimes we like observational data that we could just use, but sometimes the going out into the field and engaging with those products and with disclosure in a real world setting is very powerful.

So why do companies want to do it? Well have companies approaching me on a regular basis, wanting to do something. And their first objective is to show that their product looks great. That as an unbiased researcher, I don't do any of that. They understand that pretty quickly.

But then they have in their mind, oh, well I know what the result is. I know what you're going to find. So I'm OK with doing the research anyway. And then the next thing I say is, what if I don't find that? Oh, then I want to veto it. OK, that's the end of the conversation again. But as a

researcher, that second one takes-- you have to say that upfront. Because otherwise you go nowhere and you get in an argument later on.

So that's the first thing. Is you need to be very clear of the objectives upfront when engaging with the industry. Explain to them the whole benefit of engaging with an unbiased researcher, whether it be an academic or a governmental researcher, is the credibility of that report. If you want to show that the product looks good, find someone that believes in your product and they can write that paper. But it won't get as big a play.

The second thing is it's essential as a researcher to understand the incentives of the company. You have to put yourself in their shoes in order to have a good relationship. Why on earth did the Payday lender agree to do it? I give them a lot of credit. They thought on one sense, maybe people will default less. So maybe it is a profitable situation for us to learn about how they process information.

Second, they had terrible, terrible research coming out. The bias in the research was we're going to show they're terrible. And it's easy to show they're terrible if you're going in with a bias of I'm going to show they're terrible.

And this was the media they were getting. I'm not saying they don't deserve some. I'm trying to be balanced here. But the research, that was what it was saying. So they needed to take a risk to show a more balanced view, that they actually were interested in understanding their people and to what extent are people processing. So they were willing to take that risk, even though they had no veto right. That's the second thing.

And then following up on that. I've talked to other companies since then. And companies, generally, there's many companies that have a genuine desire to satisfy their customers. I've been talking recently until their regime change at American Express. American Express moving into low income, more low income products. And they genuinely want to understand what makes their customer satisfied, so that they can be a better provider of these services.

Another thing American Express has done, also, is that they want to engage with researchers because they want to learn. How should we think about testing? How should we think about understanding our customer? So they actually want to engage in the process, and they don't even really care what the question is. Understanding that as a researcher helps you engage with these companies.

Now don't misconstrue those items from companies that want to figure out how they can exploit behavioral biases. What kind of loan product can we give to make people take out more loans, or to do things to make us more profitable. That's research you don't want to be involved in going down that path. But rather than the positive ones, of actually making products and disclosure help people.

So these are my thoughts on the engagement with companies going out and talking to companies. There are many companies that are actively looking for engagement all the time. And you have

to balance who's looking and what do you want to be engaging with, versus who is looking and where the incentives are going to balance in that.

I want to do one more final thought, which again reflects a lot of the other comments we've already had. It's a little bit different in the last few minutes. So in consumer finance and other fields we're starting to learn about the heterogeneities in people. This is very much like one of the talks before lunch. And how people use products is not the same. People don't use Payday Loans for the same reason people don't do the nutrition use, react to those guidelines the same, and so forth.

When we think about as a policymaker now, taking a policymaker point of view, any changes in disclosure or any changes in any kind of regulation need to understand this heterogeneity. And who a change is going to benefit-- and this is very much what Ginger was saying-- who's going to benefit and who's it going to hurt?

But before we can do that, we've got to understand this heterogeneity and where the different disclosures-- who it forces forward and who it forces backwards. And I think we've done very much too little along these lines of really focusing on going after both the positive and the negative. And I think there's a lot of engagement we need to do going forward to be able to disentangle that. Good, thanks.

JAN PAPPALARDO: Thank you very much.

[APPLAUSE]

JAN PAPPALARDO: If anyone has a question for our panelists, please write it on a piece of paper. And raise your hand and we'll try to get to it. Well thank you, everyone. That was a fabulous session. One question that I have for the panelists, is why did you choose your particular research method. Did you think about alternative methods? And what were the pros and cons of the one that you chose versus others that you considered?

LILLIAN ABLON: Well I'll start. We really wanted to know what the American population, what consumers felt. What their attitudes were towards data breach notifications. And so in order to get that the American Life Panel is already set up. It already has all the demographic information of the population of the participants. And it was just kind of a well-oiled machine for us to tap into. It's been around for over a decade, and there have been hundreds of studies that have been done with the ALP. So it seemed pretty obvious to us just because it was right there and accessible for us.

IDRIS ADJERID: Yes, I think I mentioned some of these things. But I think, clearly, like a field experiment is kind of the holy grail a lot of times when we think about good causal research. But sometimes it's not possible. And then you're thinking about, well, what are the alternatives to that? And I think that crowd-sourcing sampling and those kind of approaches we talked about a lot in last session.

And I'd like just I think one of the things, is that there's actually a decent amount of tacit, unobserved knowledge around these platforms. It's not as trivial to recruit people as you would think. I feel one thing is to keep in mind some of the pitfalls and some of the good practices around these platforms is something that we want to highlight more, in terms of using this platform to do the research.

But the advantages are what I mentioned before. The replicability, the ability to test different nuanced hypotheses, and things of that nature.

GINGER ZHE JIN: Well as an economist, I'm mostly interested in economic incentive. And especially how people will use those incentives in a different regime. My own dissertation was about disclosure, but in a completely different setting was about an HMO insurance. And there I have seen the literature showing that what people say is very different from what people do. And that's why we actually engage in very long and many years of data to try to figure out exactly what people do.

And I want to echo with Adair that even in the observational study we have done, was a Los Angeles County we actually engaged was Los Angeles County very early on. Because they are very interested to know what the impact of their policy. And actually they have been very instrumental to help us to find out, for example, sales revenue data from another part of the county government, so that we can document how restaurant revenue responds to the regulation.

ADAIR MORSE: In my setting, or rather in my field, there's a kind of cottage industry in household finance research. Let's go around and document all the mistakes people make. And that's fun in some ways. That we can document mistakes. And these things get published easily.

My view is this the stress, and in particular a lot of the consumers, a lot of the people I deal with are constrained. And they're stressed out. And the stress of that situation makes field experiments a little bit unique in being able to get to real decision making. And that is not true for all kinds of studies.

But the clear advantage of being able to do a field experiment, and are people actually cutting back on their Starbucks, whatever consumption, just day-by-day, to be able to pay back their loan. That's something that I think is harder to replicate so much in observational. Most of my work is observational work, but in this case we were able to get to that point. And it's nice to be able to see where there are mistakes and we can help those mistakes. But also where it is that people just are not-- it's not all mistakes. And so there was a balance here. That's the one.

And then the flip-side of the question was what I would do differently. Field experiments there's no going back. So you've got to really design that disclosure right. And I think the kickback we got on that paper is that you can't disentangle. So when I should do the dollar treatment, you can't disentangle it's in dollars versus the comparison. So what is really the behavioral bias? And that I would do differently. To have one without the comparison in dollars. Simple things like that.

JAN PAPPALARDO: I have a question from the audience for Idris. What does framing a privacy policy mean?

IDRIS ADJERID: So framing has a pretty broad literature in psychology and behavioral economics. And I think the plain language version of it is that, can we change the way that we present some pieces of information that doesn't impact the objective features of that information and somehow changes the conception of it. What is the goal of that context? What are the things that are primed? What are the considerations?

So in privacy settings, we have competing considerations. I have some benefit from disclosure of information, but at the same time I have some risk. Can I change the framing of a message that either highlights the benefits or the risk for example? So the subtle changes in language and presentation that don't affect the objective facts but can affect behavior.

JAN PAPPALARDO: That's it. I have a follow up to that. There was talk earlier today about relative information mattering to people and making things more concrete for them. Framing matters. To what extent can that be used to lead consumers closer to the truth versus further away from the truth? Is there potential problems with manipulation?

IDRIS ADJERID: I think there is a risk. And I think it's like any tool, it's a double-edged sword. So like Adair's work on de-biasing people. You can either work to amplify the rational pieces of how we make decision making and do away with the biases. Or if you're more paternalistic, you can say I think consumers should act in a particular way. I'm going to exploit their decision biases to nudge them in that direction. That's like a big controversy in terms of what way to do it.

But I think you can also see the flip side of that pretty easily. That you know firms that have a lot of incentive to get you to share your personal information with them might use the same kind of insights in a way to nudge you towards giving them information. So it's a tool. And it can be used depending on the motives by policymakers or firms. A lot of factors, I think, influence how people use those tools.

GINGER ZHE JIN: I want to add the tension is not only the appearance of the information, but also on the content of the information. For example, in nursing homes there has been disclosure on nursing home quality. There's so many dimensions of nursing home quality. The government is in a very tough spot. Decide, OK, what to disclose and what not to disclose. If you only disclose, say, staff ratio then maybe you see quality improvement on staff ratio. But decline of quality and other dimensions that are not reported.

So that's a constant tension, that how would you make it informative and simple, but also to some extent, comprehensive and address the need of information.

JAN PAPPALARDO: So if you could have any procedure to extend your research, what would you do? Any suggestions that we'd like to do going forward?

LILLIAN ABLON: For us, one of the-- to keep going back to the ALP. One of the things is that it can be repeatable. And so we were hoping to have kind of a longitudinal study, to be able to go back and see what consumers think about notifications.

Certainly, we would care about or we would want to get more details into what sector sent out the notification. To see if there's attitudes or sentiment towards the retail sector versus the health sector, for example. Certainly, we'd want to get a little more insight into things like why certain population put an inconvenience cost of \$10,000 from having received notification. So kind of digging into some of the questions and getting a little more nuanced.

We're also interested in-- notifications can come in many forms. Text message, an e-mail, a paper letter. And being able to see how that affected people's attitudes would certainly be of interest to us.

ADAIR MORSE: So in consumer finance, the buckets of products are very different. So a credit card is very different from an installment loan, which the platforms now like Lending Club and those are all mostly installment loans. It's very different from other kinds of lending products. And it's just very discrete.

And we have no idea, really, what is the optimal form of financing where people optimize for themselves on consumption, and borrowing, and total utility. We have no idea what kind of product optimizes that a lender is able to offer profitably and a borrower borrows. It seems like an obvious question. But we don't know how people process information and how that information processing effects their being able to manage month-to-month with their finances.

And generally, just the landscape. We're just starting. And yet we have student loans of over \$2 million. We have about \$1 trillion in credit card debts. And we have platforms starting up. And we just have no sense of all this, and we're just starting to go after disclosure. There's a lot of questions we need to start to get into the products.

JAN PAPPALARDO: Well one other question I was wondering about, is this session has focused on the effects of disclosure on behavior. And if we do not test consumer comprehension of the disclosure at the point a decision is made, how do we know whether the disclosure brought them closer to what they would do under truthful, accurate information versus not. Any thoughts on that?

GINGER ZHE JIN: That's a tough identification question. And we have encountered in some situations. I mean one difficulty is that we have to know what other information consumers already know before receiving this information. That's sort of hard to get, especially if the disclosure is already done. So I guess that would depend on the context. You would try to find a control group who see exactly all the other information, but not this piece of the information.

ADAIR MORSE: I think the one thing I have learned is, that we don't measure, is how things stick. So the nutrition stuff from earlier. Do people remember this? And one thing is affecting [INAUDIBLE] at the moment, and the other thing is affecting, does that disclosure go forward? In my area, there's a whole literature on financial literacy, which is not disclosure impact, and it's in contrast to disclosure.

And literacy is not hugely effective is sort of the consensus in changing long-term, making the optimal decisions. Going forward now there's a debate on that. But I think we really need to push

on whether we can do better on disclosure or get the literacy, and whether it's nutrition literacy or any other kind of literacy, such that it's not overwhelming, so that people actually process what they need to retain. Retaining literacy, retaining disclosure I think is something that we should be targeting.

GINGER ZHE JIN: And to follow up on that, I would say disclosure is actually evolving over time even in the context of restaurant hygiene card, something that looked successful 10 years ago. I mean, today it may look not informative at all. Because 80% or even 90% of restaurant have A.

So when everybody has A, what's the information content of that? Then consumer will try to demand more details beyond the letter grade. So I think we, either as a policy maker or as a researcher, need to be really aware of sort of changing landscape, that consumer's need and consumer's information source could be very different from day to day, from year to year.

JAN PAPPALARDO: Other thoughts on that? Any other thoughts on that from the panelists?

LILLIAN ABLON: Well, just kind of to echo what had been said, specifically the language of data breach notifications is written at a higher education level. And with millions of people having their data breached and having their data stolen, one might think that it's imperative to have something at a lower reading comprehension level. And there are some efforts to do that.

Although, it's interesting to hear that the financial literature has found that it's not necessarily helpful. So hopefully, those efforts will pay out for data breach notifications. But maybe they'll find out that it doesn't really help.

JAN PAPPALARDO: One other thing I was wondering about as you're talking about people having different preferences, heterogeneous people, much of our work is done with paper and pencil disclosures. What do you see as the future? Is it possible to have disclosures that are more targeted to individuals, to different sub-populations, or even to the person themselves? Thoughts on that?

IDRIS ADJERID: Yeah. I think we've actually done some work that's motivated by the kind of heterogeneity, not necessarily in disclosures of privacy policies, but also in the control mechanisms that people now. How do you limit the sharing of information? Things like that.

And I think it's kind of similar to the discussion about being a double-edged sword. That introduces opportunity to kind of manipulate in subtle ways the nature of these mechanisms that can impact behavior in kind of important ways. And I think in terms of future research, I think this is where privacy [INAUDIBLE], is how do we design disclosures that are effective and control mechanisms that are effective, whether it be the language, whether it be these subtle changes in how we present them to induce behavior that we think is good for consumers.

ADAIR MORSE: One thought that may be relevant, and it's not exactly disclosure, but something that is effective where the financial literacy is not effective is role playing and the idea of disclosure being a role playing exercise, which could be heterogeneous. When these financial

training courses are done more in a role playing setting rather than let's do a lecture-- it's too soon to make this conclusion-- but it's generally more effective. And so the idea of heterogeneous disclosure and role playing combined may be something that could offer opportunities going forward.

JAN PAPPALARDO: Well, with that, we're out of time. I would like to thank our panelists for a wonderful session.

[APPLAUSE]

HAMPTON NEWSOME: Good afternoon, everyone. I'm Hampton Newsome, an attorney at FTC's Bureau of Consumer Protection. And now, we're onto the case studies panel. And we will be highlighting several different disclosure studies. It'll provide these experts an opportunity to discuss the goals and designs and results of their work.

We've got four presenters. First, Collin Campbell, a professor in the Department of Marketing and Entrepreneurship at Kent State University. He'll talk about his work on native advertising recognition. Then we'll have Sarah Farnsworth, who's Vice President of Scientific Affairs with PEGUS Research. And she'll address disclosure testing for drug labeling.

And then that'll be followed by Manoj Hastak, who's a Professor at American University. And he'll cover research delving into the impacts of qualifying disclosures. And finally, we'll have Heidi Johnson from the Consumer Financial Protection Bureau. She'll discuss some lab research that the agencies discuss as conducted related to disclosures. So with that, I'll step aside and invite Collin up.

COLIN CAMPBELL: Thank you so much for the kind introduction, Hampton. It's a pleasure to be here and present some work that I've conducted with myself and the co-author Larry Marks at Kent State. And what we're looking at here is native ad recognition in a social media context.

And we've seen there's been a tremendous shift of ad dollars toward native advertising. And the majority of this seems to be in a mobile context. So a lot of money is shifting toward this medium. And in many cases, we're seeing this, again, being viewed on mobile devices, which are in some cases a different context.

Here, I just pulled some examples from my own sort of Facebook feed. I've got two posts that are from friends and one in the middle that's an ad. And you'll notice, like most native advertising, the ads blend in. So they have the same style and context as non-advertising content, except for the disclosure notice at the top, the suggested post indicated there.

So it's in some cases quite challenging for consumers to recognize this. And there hasn't been research looking at it in a social media context yet. There has been some really important work taking a look at it in terms of articles, native ad articles, like in The New York Times , things like that, which is very useful.

But what we're seeing as possibly a difference is that when people go to social media, they're not looking at this content in the same way as they might be in other mediums. So if you're looking at a newspaper article, you're much more involved, hopefully, when you're reading through that than if you're looking at your social media site. People typically go to these sites for fun entertainment, sort of a light context.

We know from research that's been done previously that when people do recognize something as advertising, they generally proceed with much more critical processing. And they're much more critical of what it is they see. So one of the big concerns is if people don't recognize this content is advertising, there might be some potential harm that occurs to them. What we really want to get into later on is taking a look and seeing that if we vary perceived harm or things like that, what kind of effects might exist there?

But the first question is really taking a look and seeing do people actually look at these disclosures. We've seen not a whole lot of outcry from consumers with advertising on social media, unlike what we've seen with articles. And we've also seen industry suggesting that consumers tend to like not very overt disclosures. So we're taking a look in this context to see are people actually paying attention to these disclosures or is it that they're just not actually noticing them at all.

So in our first study, what take a look at is very two different things. One is ad position, where the ad actually is placed. And then secondly, it's brand familiarity. The reason we take a look at ad position is we want to see if one of the cues people might use to recognize advertising, where it's placed on a site, if that's gone, if that affects the recognition.

The second thing we look at is brand familiarity, if it's a well-known brand or an unknown brand. So we had a nice little 2 by 2 design here. We used Amazon MTurk. We had high familiarity, low familiarity. And then we had whether it was placed in the stream, so in the main part of your Facebook feed, or on the side bar.

In all cases in this study, study one, we use suggested posts, which is what Facebook users for its disclosure tag. These are examples of the stimuli we used. And here, we've got the Ford ad in the center on the left. And on the right-hand side, we have it on the side bar. And we pre-tested these to be equivalent. There's no real difference in recognition or ability to pick up on things between these two different sites in terms of other attributes.

And what's really interesting here is that we see on the left-hand side, in the stream condition, with an unfamiliar brand, people have lower ad recognition. And what's important to note there, in that blue bar on the left-hand side, all they had to recognize it as an ad was the disclosure itself. In the other bars, they have additional cues.

They had the location of the ad. And they also potentially have whether or not it's a familiar or unfamiliar brand as well. So we're seeing here that additional cues beyond just the disclosure is needed to really boost ad recognition. Second thing to note as well, too, this is measured on a seven point Likert scale. Ad recognition is pretty high generally, which is kind of cool as well, too.

In a second study, what we do is we go a bit deeper. And we stay here with brand familiarity. We have low and high familiarity. But we also introduced image professionalism, so a professionally produced ad or more low grade image that might come from a consumer.

We also look across five different disclosure conditions, so four different disclosures as well as a no disclosure condition as well, too. So a bit more of a complex design, sort of what Craig said not to do this morning. But it's still kind of cool to look at.

And we also looked here at ads alone. So in the first study, we were looking at a mock up of your Facebook feed. And in this case, you're looking at just the ad.

And the reason we did this is because, one, that's how people see it in their mobile devices. And two, we're thinking maybe people are being distracted by other content on the page. So if we really make it clear, maybe we're going to see more of an effect of disclosure.

This is a really interesting site as well, too, Olapic. It's a company that actually advertises and sells user generated content to be used on advertising. So they'll actually go out and connect. If you're a company, they'll go and find content created by consumers.

And they'll use that as a basis for ad material. And they claim that there's a better response to user generated content as opposed to professional images. This is one of the reasons we're looking at professionalism of the image as part of the study.

These are the two different stimuli we used. We had 20 different conditions, but these are just representative of that. So we have a professional photo on the left-hand side. On the right-hand side, we have the unprofessional photo or less professional photo. And then we used Dunkin' Donuts and Dough Boy's Donuts. And we had different disclosure text across the top.

In terms of taking a look at a three way interaction when it comes to analysis, it can get a bit tricky. But I've been trying to make it hopefully more approachable. This is taking a look across the five different disclosure conditions. Each particular column is one different analysis.

What we're seeing here is that across these four different disclosure conditions, as well as the no disclosure condition on the right, we are seeing effects of these other criteria, so the professionalism of the image or also taking a look at the familiarity of the brand. What's interesting to note here is it's sponsored post behaves similarly to no disclosure as well. And we're seeing the same pattern across those two conditions where you need to have those two cues coming together in order to get a boost to ad recognition.

So really interesting to take a look at that. One of the other important questions you want to take a look at was slicing the data a little bit differently. What we wanted to see here was looking across all those five disclosure conditions, the four disclosures and the control, is there that added effect of having additional cues? So what we wanted to take a look at here was if there's no other cues, what would happen? And then what begins to happen if you have additional cues added in?

So here, we can take a look. These are the results for ad recognition looking across for the unfamiliar and unprofessional conditions for the five different disclosure conditions. And there's no significant difference across these. So this means that when there's just a disclosure with no other cue, there's no real difference caused by that disclosure.

What we take a look at next is if we add in a professional photo. What we see here is the red arrow, the second from the far right, promoted by a brand that was significantly different than three with white arrows. So, sponsored post, suggested post, and also no disclosure. So in that one condition, that disclosure did boost ad recognition in conjunction with a professional photo.

Likewise, in the next condition, we took a look at an unprofessional photo, but a familiar brand. And again, we see a similar pattern happening here. So the promoted by brand boosts ad recognition in conjunction with that familiar brand. So that additional cue is working to boost recognition.

What we do notice as well, though, is that there's a threshold effect. When you've got both of those happening together, then the disclosure itself isn't actually all that extra-- doesn't provide much marginal benefit, so to speak. So here, we can see the red bars. Those are not significantly different. But you will notice they are all a bit higher.

So, very cool to see here. So taking away from this, the big sort of like takeaways are the fact that we need multiple cues in a social media environment. We think people are not really processing all that deeply when it comes to this content. And that you need more than just a disclosure, potentially, to take a look at trying to boost ad recognition.

At the same time, though, we do have very high ad recognition rates. We had a seven point Likert scale in the first study, nine point in the second one. And they were well above average in both cases. So it's a very good sign there.

The promoted by brand was the one that did work. Again, we're not seeing deep processing. And there may need to be different forms of disclosure that are being used to actually take a look at this and trying to boost ad recognition.

Important notes in terms of the FTC's guidance on disclosures, the language we find that works is different. But it's important to note that within their disclosure guidelines, they do recommend taking a look at the totality of what's being conveyed. And our research really supports that.

So it really shows that it's not just the disclosure itself. It's other cues, like the familiar to the brand, the location of it, or the professionalism of the photo that might hint at boosting ad recognition. Thank you very much.

[APPLAUSE]

SARAH FARNSWORTH: Well, it's a pleasure to be here with you today. My presentation will be a little different, as we don't directly test product disclosures. However, we do routinely test

consumer comprehension of over-the-counter product labels. And so we thought our model could be useful if disclosures were ever systematically tested.

So for a bit of background, as we've been discussing today and just before lunch about comprehension of disclosures, an important factor that impacts the effectiveness of disclosures is whether or not consumers can actually comprehend them. So as I mentioned, we do propose that the model currently used to test over-the-counter drug product labeling may be relevant to evaluating other product information, such as disclosures. So today, I'll discuss some studies used in the approval of an over-the-counter medicine that was submitted to the division of non-prescription drug products at the FDA as an illustration of this model.

I first wanted to give you a brief background to PEGUS research, the organization I'm representing, and the type of work that we do. So we partner with pharmaceutical companies and consumer health care product companies to develop and optimize the actual packaging for an over-the-counter medicine or product. So we work with them to help develop and refine the language on the drug facts label and also the formatting and content of the overall package.

Once that's developed, we design and implement various studies to test the over-the-counter label. The first type would be a study to determine if consumers can adequately comprehend the information on the label. And those are label comprehension studies.

And if consumers can then in turn use that information on the label once they comprehend it, but actually use that information to apply it to themselves to make an accurate decision about whether or not that product is appropriate for them, and those are called self-selection studies. And lastly, we also conduct studies to determine if consumers can use the product safely in an over-the-counter setting, a simulated over-the-counter setting, by following the label directions on the package. Then the results of these studies are then submitted to the FDA for review on approval of new over-the-counter products.

Today, I will be focusing on the comprehension piece of this overall development program. So the regulations state that over-the-counter labels must be likely to be read and understood by the ordinary individual, including individuals of low comprehension under customary conditions of purchase and use. So our label comprehension studies are designed to demonstrate consumer comprehension of the OTC label, namely the directions and warnings in that content.

So these studies utilize one on one standardized interviews, so face to face interviews with a general population of consumers to collect data on comprehension of the key warnings and directions. It is important to know the FDA issued a guidance for industry on the conduct of these studies back in 2010 to help standardize how these studies are conducted. So overall, these studies are conducted in accordance with that FDA guidance, of course, and accordance with Code of Federal Regulations and good clinical practices.

So like I mentioned, the Code of Federal Regulations governs the content and structure of drug product labeling. However, the wording and other package elements should be developed and optimized through a series of iterative, qualitative, and quantitative comprehension studies. So

before getting to that point, not only the wording and the language and how it's presented to consumers, but also the formatting of the label.

Earlier today, we talked about elements such as presenting things in a structured, organized way. Luckily, the drug facts label is stimuli that's already pretty organized and structured in terms of sections. But things like how much white space is used, the bullet points, how they're utilized, how bolding is utilized, everything like that is developed in early qualitative and pilot studies.

And there's quite a bit of content on a drug facts label. So testing is prioritized on the messages with the greatest clinical consequence. If a consumer fails to understand them, and then thus logically fails to follow that direction. So the endpoints of a study are prioritized based on the clinical implications of a consumer failing to understand that direction.

And those with the greatest clinical implications are assigned a target performance threshold. The greater the clinical implication, normally the higher the target. So for example, something that would preclude a consumer from use of a product that would be dangerous for them to use, that warning would typically be in the Do Not Use section of the label. Those are frequently determined to be primary endpoints with a threshold of, say, 90% or higher. Whereas, something that doesn't quite have as great a clinical implication if it's not understood or followed, sometimes is assigned a lower target performance threshold, such as 80% to 85% for example.

So in sum, the overall model that's followed, it's pretty simple. But it's useful. And it works. So first, we create the content or the stimuli in the drug facts label that participants are responding to.

We conduct early qualitative work with consumers. Qualitative work isn't always done, but we find it very helpful. Because we can first sit down with consumers in a one on one setting and talk to them about the concepts on the label, talk to them about the drug and the indication that it's intended to treat to kind of get in their own words how they think about the issues, how they think about the product.

And then we identify key messages and end points and the associated target performance standards. And through pilot testing, we also develop the comprehension questions, how we actually assess comprehension and the associated scoring criteria. A pilot test is conducted. And then further refinement of the data collection instrument, the label, and other elements-- there's usually multiple rounds of refinement. So it's an iterative process.

So then you test again. And ultimately, the culmination is a large pivotal study intended for FDA submission, where you evaluate your final results against the targets. Oops. Got a little ahead of myself.

So in the time that I have left, I wanted to just briefly present a case study for you for Nasacort Allergy 24 hour. You may have seen this on product shelves or ads for this. It was a first-in-class switch for a prescription to over-the-counter switch. It was a first-in-class.

So a series of label comprehension studies were conducted into 2011 and 2012 to develop labeling that was adequately understood by consumers, a general population of consumers. As I indicated earlier, we did the first conduct qualitative and pilot studies to refine the methodology and data collection instrument as well as the stimuli. Then typically, the final study protocol that describes the study design, the statistical analysis plan, and other elements is summarized in a study protocol and typically submitted to the FDA for review and comment prior to fielding.

So these pivotal studies for FDA submission were conducted in two phases. Both the outside package, in this case, and the package insert inside the product carton were tested for consumer comprehension. And in this case, participants were randomized to view one stimuli or the other to limit any influence or bias.

It was a large sample. 886 consumers viewed the outer carton. And 734 reviewed the package insert. Very few exclusion criteria are applied in these studies. We want to ensure it's a general population of consumers rather than only those who have the condition that the drug is intended to treat.

And this is because when a product is over-the-counter, one may not have the condition at that time, but could develop it down the road. Or someone could be purchasing the product for someone else in their household that does suffer from that condition. And in this case, participants down to age 16 were recruited at 15 mall sites spread across the country to ensure both demographic and geographic diversity and a representative sample.

So I'm running out of time. So I'll zoom through this, sorry. So we collect data, like I mentioned, in one on one interviews. They're standardized interviews at all sites. And we enter data real time in electronic data capture system. And then each key message or end point has an associated question or questions to assess the comprehension.

We also use a pre-specified answer key rather than determining what's correct or acceptable after the study's conducted. That's determined in advance of study conduct. And a keynote about sample composition in this study design in these types of studies is that the FDA typically requests that around 20% to 30% of participants qualify as low literacy.

And in these studies, we used the rapid estimate of adult literacy in medicine. And that's a validated instrument, too. It's a brief survey. It's really brief to administer.

It's about a list of 66 words that are read out loud. And their ability to correctly pronounce those words is used as a measure for literacy. And in this case, approximately 30% of the sample qualified is low literacy.

So briefly, this is the drug facts label tested. We used scenario questions. They're kind of like a story problem. Because they require a higher degree of assimilating package information into made up real life situations.

The questions are read out loud. They're open-ended and not leading. And there's also silent pre-coded answer alternatives for ease of data collection.

So I'm out of time. So I'll briefly show you an example. So on this label one of the important end points tested was children under 2 years of age do not use. So this is an example of a scenario question.

You can see it's an OK or not OK type of question. So that's a 50-50 chance of getting it right. So in those circumstances we ask participants to explain their answer. And then you can see an example of a pre-specified answer key at the bottom. Sometimes acceptable answers are also accepted, because people tend to speak in abbreviated ways.

So when there's a list of components to a response, such as this example, we accept other options as acceptable. So in this case, for this question of the allergic reaction, a correct response was stop use and ask a doctor or seek medical help right away. But if someone just mentions stop use, ask a doctor, or seek medical help, that was determined that it basically demonstrated an acceptable level of comprehension, showing that they saw the warning and they understood the general message.

I'm sorry. I went over. But that's kind of a general overview of the models we use to test over-the-counter package information. And I can talk more about the implications of how that might relate to product disclosures. But hopefully, you can see that drug package information it's a lot of information just like product disclosures.

And we believe that a similar model with face to face or other methodology, but with standardized interviews and structured data collection, we could test product disclosures if that ever became required to submit that data to the FTC. Thank you. Sorry about that.

[APPLAUSE]

HAMPTON NEWSOME: Thank you, Sarah. Now, we have Manoj Hastak.

MANOJ HASTAK: I could continue for you if that's-- OK.

SARAH FARNSWORTH: You can finish my presentation for me.

MANOJ HASTAK: Oops. I went a little too fast there. Thank you, Hampton. And thank you to the FTC as well for inviting me to speak here and for a really good conference so far, really enjoying it.

So let me jump right in since I have less than 10 minutes here. And my goal is to talk about three studies. So I'm going to try and do Colin one better. Although, I have a feeling I might get cut off at the end there.

So quick disclosures about myself, I have a long standing association with the FTC. I've been an in-house consultant and advisor for them for a number of years. And so I've participated in or led a number of FTC studies on a variety of topics, including disclosures.

And so today, I'll talk about a couple of the FTC studies. And I hope to also talk about a non-FTC project that I'm working with, a colleague that's linked to some of the FTC work. So my talk is going to focus on three methodological issues. And to the extent I can, I'm going to talk less about the results and focus more on the methodological issues.

The first one is assessing ad communication versus believability in communication and disclosure studies. By ad communication, I mean measures that try to assess what consumers think the ad stated or implied. What are people taking away from the ad?

But believability focuses, as the word suggests, on whether consumers believe those claims in the ad. And often, one or the other is assessed in our research, but not both. So that's kind of one issue.

The other is probing for consumer interpretation of disclosure intent. So again, often we focus on issues of consumer attention, to disclosures or comprehension, or even behavioral impact. But the less often, we focus on what consumers think the purpose of the disclosure is. And so it's an interesting issue as to whether people differ in those assumptions and expectations and what effect that has on the efficacy of disclosures.

And finally, you've seen a lot of eye-tracking data in this conference. It's really a growing tool, certainly in my field, to assess consumer attention. And so I'll present some data on that as well.

So my first study is an FTC study. And by the way, this is available on the FTC website. There's a report that I've done that's on there. It's on "up to" claims in advertising.

These claims are pretty ubiquitous. You see them, for example, in weight loss claims, lose up to 50% or weight loss by up to 50%. You get them in ads touting earning claims, make up to \$500 or earn up to \$2,000 a day, internet speeds up to five times what you have achieved before. So, very, very common kinds of claims.

And the study was trying to understand whether people saw the "up to" part of the up to claims as a disclaimer in and of itself. In other words, did that moderate consumer expectations. And so we did a mall intercept study with a product that was a replacement windows product, Bristol Windows. We had three ad treatments that I'll show you and decent cell sizes.

So here is the first sort of our original ad, very simple and sparse ad. And we wanted to use something fairly simple in order to reduce noise in terms of consumer communication and comprehension. And you see the "up to" claim there, save up to 47% on heating and cooling bills.

Here is our first comparison ad, the cleansed ad, in which we took the "up to" language out of the ad. The rest of the ad is identical to the original ad. So again, trying to see whether there's any difference in consumer takeaway and believability in the "up to" versus the non-"up to" condition.

And then here is our disclosure ad. We put in a big fat disclosure at the bottom, clear and conspicuous as others have talked about this issue. And here is a better view of it. We basically said the average user of this product saves about 25% on heating and cooling bills.

I've crammed a lot of results into this one slide. So let me just take a moment with this. And I'll go row by row. So the first rule basically is asking people based on what the ad said or suggested, how many Bristol Windows users can expect to save about 47%?

And we had appropriate filters making sure people actually took that 47% number before we asked them this question. And then you have the three columns comparing responses to the three ad treatments. And what you find is that there are no differences. There is no difference in take away between people looking at an "up to" versus a non-"up to" claim.

So the "up to" qualifier doesn't do anything. And then quite surprisingly, the big disclosure doesn't seem to do anything either. It actually keeps people at the same level.

The second row gets at the believability measure. After people had been asked the question in row one, we kind of said, now we're going to ask you about your personal opinion of Bristol Windows, and then said the same question. In your personal opinion, how many users of this product will save about 47%?

And what we find is there are no differences across, again. But interestingly, there are no differences between the two rows either, which kind of surprised us. Because we thought there would be some discounting. The ad said 47%, but I don't believe it. It's probably less.

And then we asked the same thing a slightly different way. In your personal opinion, what would be the average savings that people would achieve? And again, we don't see a lot of discounting and no differences between the treatments. All right.

So again, I'm not going to focus too much on the study results. I'll talk a little bit about these different measures that we are using here. But let me follow up first with our second study which is the one I'm doing with a colleague at American University where we want to replicate the FTC study and also extend it by using some additional measures. So we used both concurrent think aloud protocols and eye-tracking measures. And what I'm presenting here is really data from two sort of related studies, if you will.

So the ad was tailored more to the student population, because we use the American University Behavioral Lab in the business school. But it's very similar to the FTC ads. Charge your phone 43% faster-- I think the previous one, sorry-- versus charge your phone up to 43% faster versus a disclosure condition that says basically that the average person will charge 17% faster.

And looking at the results, there are some similarities. We replicated some aspects of what we found with the FTC study. But other results are a little bit different.

So again, very quickly what we are finding is the "up to" qualifier doesn't seem to do anything. Adding the "up to" doesn't really give people a red flag, if you will. But the disclosure here was

effective. If you just look at the first row, you drop people's take away from the 70s and 80s to about 37%.

The other interesting thing we found was that while the disclosure was statistically significantly affective in the first row, if you will, which is the ad communication measures, it wasn't effective in the second row. Everybody's engaging in discounting when we asked for their personal opinion. But the disclosure doesn't drop that number much further, so some interesting results which I'm happy to talk about in more detail.

The eye-tracking data was interesting as well. I'm showing you two measures here. The disclosure dwell count is how many times people actually looked at the disclosure. And you can see that there is no difference between the "up to" and the non-"up to" conditions. But the disclosure condition people look at that disclosure area, that area of interest, a lot more.

Let me just quickly say going back for a moment. If you look at the disclosure ad, the area of interest really was not just the disclosure, but that entire area including the charge your phone up to 43% faster. So what it's really showing is that people are getting to that area a lot more. Sorry, I'm going the wrong way here. Wasting time as usual, sorry. And I've gone too far.

So people are fixating on that disclosure a lot more. And disclosure dwell duration in seconds shows the same thing. People are looking at the disclosure. Interestingly, what we found is that while people tend to look at that disclosure area when there is a disclosure there, it does not change how much time people spend looking at other areas of interest in the ad, the ones that a marketer would be interested in.

So at least our initial data are suggesting that the disclosure is not interfering with attention to other information. I probably don't have time for the third study. But I just want to show you quickly sort of the punch line with the third study. And I won't describe it in detail.

We kind of tried to assess why people thought the disclosure was there in this study. And what we find is sort of an even split between people saying that it's there to protect the company-- so people said things like they're trying to cover their ass, for example-- versus this is useful information for consumers. They're trying to tell us something interesting. And again, I'm happy to talk a little bit more about that.

So just to sort of round things off for a methodological perspective, I think what I would sort emphasize is the value of multiple measures. A lot of people have mentioned this, but having multiple measures provides you with much more interesting insights into sort of the rich impact the disclosures may have on consumers.

Probing for disclosure or comprehension as well as intent, I think intent data is useful. Somebody in the last presentation, a couple of people, talked about heterogeneity of consumer responses and preferences. This is an interesting dimension of heterogeneity, that different people come in with different priors about why disclosures exist and how the effectiveness of disclosures would be useful.

And finally, the value of replication. We replicated portions of the FTC study. At least in my field, this is rarely done. Although, it's beginning to sort of-- one of the journals has now started encouraging replication and extension. So that's kind of a plug for that. Thank you very much.

[APPLAUSE]

HAMPTON NEWSOME: Heidi Johnson.

HEIDI JOHNSON: All right. So I'm going to be discussing findings from a study that several of our researchers did in a laboratory setting. But I want to open by talking more broadly about CFPB's interest in research on disclosure. I think it's probably pretty obvious as a regulatory agency this is a lot of what we do. It's a really large share of consumer protection regulations included disclosure component.

And so actually the Office of Research at CFPB has decided this is something really important for us to kind of strategically invest in and focus on. So we're going to be trying to build up a greater knowledge base about this over the course of the next few years. And I wanted to share some of the main priority areas we've identified that we're hoping to learn more about.

So the first one is efficacy. This is really focusing on how different dimensions of a disclosure might influence outcomes for consumers. We're also really interested in questions around methodology, I think, matching up with a lot of the discussion we've had here today. So, trying to figure out when is the right time to use which different methodology, what kind of outcome measures should we be thinking of, and how do we study that.

And then the last component, focusing on market effects, trying to think holistically about the effect of a disclosure intervention on the marketplace. Not just how do consumers respond, but also adding in considerations about how firms might respond to that and the interaction between the two. So the efficacy component is worth drilling down on a little bit further.

In the course of developing disclosures at CFPB, a lot of what we do is looking at, OK, based on this content that we're developing, the words that we put on the page, how do consumers understand that content? And we're interested in really expanding a lot further in that and looking at these other elements that are kind of lessor explored. So the first is thinking about the context in which a disclosure is provided.

Financial transactions happen in so many different contexts. And consumers are receiving information in so many different contexts from the mortgage closing table, to account opening at the branch of a financial institution. Or consumers might receive information in the mail at home and be reviewing it on their own. And so we're hypothesizing that where those decisions are happening might really impact what people take away and how they think about a disclosure. So we want to learn more about that.

We're also really interested in thinking about different outcome measures, so not just understanding. But what makes people pay attention to a disclosure? And do they incorporate that into their decision making?

And we're going to tackle these questions through a wide variety of methodologies. In the course of rule making when we're developing a disclosure, we'll do a lot of research through qualitative methodologies, like user testing, cognitive interviews. We also are able to partner very frequently with private companies and run field trials, which is always really exciting to see what's happening out there in the real world.

And then another component of our research is more foundational. By which, I mean not really in the context of a specific disclosure-- so what do we know about this particular one-- but more generalizable lessons that we can glean and apply to future disclosures. And so it's in that latter category that the study I'll be sharing falls into.

So this lab study focused on context and attention. We were really interested in using the lab for this, because the lab is a great opportunity to isolate effects. You have control over the environment. And you can really figure out what's affecting what here.

So we think of this as a good way to kind of glean those more generalizable lessons for future disclosures. The primary question here was what affects attention. And we were looking at two factors.

The first was the design of the form, so the arrangement of information on the page. And the second was the context. So were people reading this in isolation on their own? Or were they in the presence of someone else, in this case, a researcher?

So this study was conducted by Alycia Chin and Dustin Beckett of the CFPB. It was actually conducted at the end of another study that they were conducting at Gettysburg College with 192 students. So in the course of doing this other study, we would always be providing participants with information about their privacy rights, and how they're going to get paid, and other features of the study. And so our researchers kind of capitalized upon that opportunity to provide that information under a few different conditions.

The way we measured whether they were paying attention is this information here that was embedded in that disclosure. So I'll summarize really quickly. Essentially, this just offered participants the opportunity to learn about studies paying a rate of \$35 an hour if they initialed somewhere on the form.

And that was the information that was chosen to put in here, because we would expect a dominant response to that. This was presented to people who were already in the lab. So we know they're kind of generally willing to participate in these studies.

It's not really signing them up for any particular time that they have to show up to. This is just asking do you want to be notified. So that's kind of pretty easy to agree to. And then \$35 for college students actually is pretty persuasiveness.

This is more than double the amount that they were typically paid for studies. So we thought if someone actually sees this information, they're pretty likely to put their initials down on the page and sign up for it. So that's what we were looking for, did they initial.

So the first factor that we were looking at is, as I said, the design. So you don't need to worry about trying to read what looks like fine print here. But essentially, this is to show you that it was the same information on both just arranged in different places.

And that difference was this signature box. So on one form, it was placed at the top. On the other, it was placed at the bottom. And the information there was asking participants to affirm that they had carefully read the information, and then signed there.

And this was inspired in part by work done by [INAUDIBLE] in 2012, which showed that when people were filling out a form where they were asked to report information, the accuracy was much higher when people signed the form at the top. And so we thought that might be kind of interesting to apply to engagement with the disclosure. If you have to sign at the top, are you more likely to read through to the bottom? And so that area circled in blue is the information about studies at \$35 an hour. So that's what we were trying to see if people actually noticed.

OK. The environment is the other factor that's randomized. So this is the setting that people are in. This is the lab at Gettysburg College. The black dots are the participants at different workstations. The orange dots are our researchers, one at the front of the room, one in a room off to the side observing the session and providing payment at the end.

And in the first condition, participants were provided with this disclosure after the end of the study while they were sitting at their workstations. So they were asked to review it. They were kind of left on their own to do that. And then one by one, they were called up to the payment window to return that disclosure and receive their payment.

In the other condition, they actually weren't provided that disclosure until they got to the point of being provided with the payment. So they just sat quietly at their desks, wait be called up. And then when they were face to face with the researcher, they were given the disclosure and asked to review it.

So our results were you can see on the y-axis here-- I'm not sure if you're actually able to read that. But I'll walk you through it. The y-axis is the proportion who initialed the form, which is our measure of attention.

We saw that when it comes to the design of the form, more people did initial it when they signed it up at the top. But the difference wasn't statistically significant from when they signed at the bottom. So that might be an area to explore more with more power.

However, on the right-hand side on the graph you see there, we saw a big impact on attention when people were looking at it alone versus face to face with the researcher. So many more people initialed it when they reviewed it on their own. So this was pretty interesting for us.

We found that in this particular case, the form design didn't influence attention. We don't give up on design. There's probably a lot more ways that we could have constructed it that might have had an influence. But in this particular case, it looks like arranging the signature box may not be as promising.

But it was really interesting for us to learn that the context did have that big effect from 5% of people initialing to 35% when they reviewed it on their own. And so that's something that we would want to think about in developing future disclosures. It also kind of makes you question what are the other opportunities to further increase attention. If we could see that big a swing from this one little change, from 5% to 35%, what are those other factors that maybe we haven't thought about yet that could also boost it up further?

And also another big takeaway for us is in thinking about outcome measures. So we realized at some point in some however small way, everyone paid attention to this disclosure. Because we had 100% rate of signing their name. Everyone provided their signature.

But some people didn't go much further beyond that. And so if we had set up our outcome measures to look at how many people signed it, we would have concluded something very different. So that's something that we're going to keep in mind for future studies is figuring out how to construct the outcome measures and how to capture what we're really looking for there. And I'll wrap up. Thanks.

[APPLAUSE]

HAMPTON NEWSOME: OK. As you probably noticed, one of the themes of our discussion today has been the methodologies used with testing. And with this panel, we've got four very interesting studies. So they involve mall intercept, lab studies, and also an online study. And so I'd like for us to kind of discuss some of those issues.

And Manoj, I'm going to pick on you first. With your "up to" studies, you were studying fairly similar claims. But on the first one was small intercept. And on your follow up, you chose to use a lab study. And I was wondering if you could get into what went into that decision and the benefits and drawbacks you see from the lab study.

MANOJ HASTAK: Sure. So the FTC study, as I said, produced some surprising results. And so I felt like there was some value to try to replicate it. But also some when I presented that study early on, there were also some comments from the audience.

And one of the thoughts was is there a lot of noise in mall intercept studies. People who have done those kinds of studies would realize that it's a more noisy environment than you will get in a pure lab, for example. And there is more heterogeneity than you might find in a student sample. So that was one issue.

Cost was an issue as well. I know other panelists talked about it. A mall intercept costs quite a bit of money. And as an academic, that was kind of difficult to do.

But our goal was really to do what I would call a constructive replication. We didn't want a literal replication of the FTC study. So we chose a different product, which is more relevant to students obviously.

We changed the methodology a little bit. We had student subjects using a behavioral lab. And so the fact that some parts of the findings are replicating between the FTC and the AU study, if you will, I think is quite encouraging. Because it kind of tells you that this effect, or lack thereof, may be more robust.

But certain other things didn't replicate. And so that does leave a lot of open questions. Because it raises questions about whether this is because of different methodology or are there really maybe different situations where the effects of "up to" claims and disclaimers may be different. So it sort of tries to sort of solidify our understanding of some issues while leaving some questions open for future research.

HAMPTON NEWSOME: OK. And Colin, yours was an online study, which makes sense. Because you were looking at social media ads essentially. But was there anything in particular? Did you consider any other approach?

And also just generally in your work, have you done a lot with online studies? And what benefits do you see? And what drawbacks do you see?

COLIN CAMPBELL: Yeah, for sure. So we used Amazon MTurk. As we discussed quite a lot today, a lot of the academics are using that, which is fantastic. I use that as well, too.

And some of the issues we've identified previously as well. We definitely use attention check questions and looked at time people spent on the surveys, things like that. Big benefit of MTurk, of course, you can launch a study and have results back in a matter of hours in some cases, which is fantastic. So, spending less time getting ready to run it or running it in labs, things like that.

I've also used student samples, of course. And also panels as well, consumer panels, which are more expensive. And at least in the studies I've run, I've not seen major differences between them, basically no differences between them, which is quite interesting.

It could be contact specific. I'm not sure. But at least from what I've seen, the consumer panels have not been any different from what I've run with MTurk. In some cases, the exact same study hoping for differences and not getting them.

One of the criticisms we've received with this study is the fact that people are looking out the stimuli artificially longer than they might in a real world environment. So we're actually running some stuff coming up. We're going to actually have a video of a scrolling feed where, like, what you would look at if you were going through your own social media, which is probably more realistic. And we're curious to see what happens with that. So there's always advantages or disadvantages to any particular approach.

HAMPTON NEWSOME: OK. And Sarah, yours is a mall intercept, which makes sense. Because generally, at least, you know, when I buy OTC drugs, you're in the store looking at the label. But have you considered using other approaches? Or are other approaches used for this? And are they effective, less effective, than the mall intercept?

SARAH FARNSWORTH: Sure. So yeah, we did mail intercept only for a long time. And really the reason we chose mall intercept was because it allows you to access a wide variety of consumers in a short amount of time. So there's mall research sites all across the country. So you can get a great sample size and geographic and demographic diversity quickly.

However, we're starting to explore other models where we pre-recruit using some market research databases and other panels to pharmacy sites. So it's a more naturalistic environment where someone might go to get their OTC medication. So the environment is a little different.

And also, we think that the pre-recruit and scheduled appointments may actually speak to the cognitive model that Michael discussed this morning in terms of attention, motivation, attitudes, and beliefs. We were tending to observe that when you pull someone off a mall floor, you're interrupting their outing to go shopping. They're not quite as motivated or attentive. And so that in turn can influence the accuracy of your comprehension assessment.

We're wanting to make sure the content of the label is clear and easy to understand. And if you have so many confounding variables, like motivation and intention, then sometimes you're not getting the best assessment. So we kind of weigh those pros and cons when choosing a model. We're also piloting out some online work as well.

We'd still choose face to face research. Because we think there's a lot of value in that. We don't feel for this type of research the self-administered questionnaires are the best way to conduct the interview and gather that data. But we're looking into ways to still utilizing the face to face personal interviewing style, but in an online environment.

HAMPTON NEWSOME: OK. Does anyone else want to weigh in on these issues, the benefits of online approaches? We talked earlier about representative samples. But there are also benefits in terms of using graphics and cost and that kind of thing. Any thoughts?

MANOJ HASTAK: Well, I'll share one thought. I mean, to me as I've sort of looked at these methodologies over time, the online panels have become a lot more sophisticated in what you can do and what kind of data you could collect. You know, so you could do-- I mean, I've seen packaging studies where people can virtually rotate the package and zoom in and do all kinds of things.

It clearly doesn't simulate how people really in reality interact with these products. That part is true. But the fact that you're able to do a lot more and the cost factor-- from everything I hear, the costs are significantly less than doing mall intercepts-- is kind of moving a lot of the research-- including litigation research, which really is usually the last barrier to cross. When all the i's have been dotted and the t's have been crossed is when I think litigators and judges accept this kind of research. So I think the movement is in that direction. And more and more research will be conducted online would be my guess.

HAMPTON NEWSOME: OK. Just to move into kind of some of the substance of your studies, Heidi, the one thing that strikes me, that result where the people were more likely to initial the form when they were alone in the room, right? Am I reading them right?

HEIDI JOHNSON: That's right, yeah.

HAMPTON NEWSOME: And when someone is, I guess, looming over their shoulder or whatever, they blow it off or something. What's going on there? Do you have any speculation about--

HEIDI JOHNSON: It's a great question. And it actually reminds me of the conversation we had before the panel, where you hypothesized it might move in the other direction. So I think certainly we could come up with a lot of ideas for how people are influenced and what kind of cues they pick up on to make their decisions about how much attention to spend on something.

In this particular case, our researchers have a few ideas that they're actually going to be exploring over the course of the next year. So we're really excited to be able to study it a little more and disentangle those mechanisms. One hypothesis is that there's a difference between the time that was allotted between these two different conditions. So maybe if you're kind of on your own with not much else to look at, because they were in a lab setting, they're going to look at it. Because there's not much else to do, and they have more time to do that.

Another possibility is that there is some element of social pressure. So maybe when they were face to face with the researcher, they might have been responding to some kind of sense that there wasn't an expectation that they would review it carefully or that they were using someone else's time. So they're going to be conducting more studies to kind of try to see which of those two might be at play and if there's anything else going on there.

HAMPTON NEWSOME: OK. And Colin, just briefly, kind of in nutshell, what does your crystal ball say about native advertising over the next couple of years? It's something you've spent a lot of time on. And how do you see things evolving with research and also just what's happening in the market?

COLIN CAMPBELL: I wish I had a crystal ball. That would be amazing. The one thing that I am definitely seeing and hearing from talking with the industry as well is the shift worker programmatic. So they're getting much more nuanced in how they're placing native ads. And they're trying to get much more sophisticated with the way that they're doing that.

So currently, the whole essence of native is that it blends into the environment in which it's in. And in many ways, it's very boutique to do that. Because you have to understand the particular page or site that you're on. And the challenge is how to do that in an automated way.

Well, they seem to be making systems to do that. So I think we're going to see a big shift toward having more of this content out there as we move forward. Because people also seem to respond to it more positively as well, too. So I'm expecting to see more of it and probably more sophisticated forms of it coming out as well.

HAMPTON NEWSOME: OK and Sarah, with the drug labels, what happens when it fails the study? And what's the metric that used to say, well, this isn't failing? And also, what are the

things that consumers-- is there a particular part of the drug label that consumers have particular trouble with that always kind of pegs problems?

SARAH FARNSWORTH: So when an end point doesn't meet the target performance threshold, we kind of evaluate how far off it is, factors that could have played into that not quite meeting the target performance threshold. And that's one reason why it's an iterative process. So hopefully, by the time you get to the culmination of the pivotal large study intended for FDA submission and OTC approval review and approval, the goal is to have refined your stimuli and your questions well enough that you can optimize comprehension.

And sometimes, we've been in a position where we feel like we've done everything we can on a message. And it just might not reach the threshold that FDA would like to see based on the clinical implications. So sometimes, you just kind of have to let the cards lie and go forward.

But in terms of areas on the drug facts label, it really depends on the indication. But in my experience, it seems that consumers really pay attention to the Do Not Use section, the direction section for actually how to use the product. But for whatever reason, it seems that ask a doctor or a pharmacist before use doesn't feel quite as concrete.

It's more of a conditional direction. So it's not telling you that this is OK or not OK for you to use. It's more of a conditional direction. So consumers tend to struggle a little bit more of that in my experience.

OK. We have about two minutes. Does anybody have anything they want to add, a cautionary tale maybe from somebody else's research or something like that? OK. Well, let's wrap it up there, so we can get on to the next panel. Thank you so much.

[APPLAUSE]