The contexts putting your privacy at risk

It’s time for new theories and better training about online privacy. Read this before you click accept all.

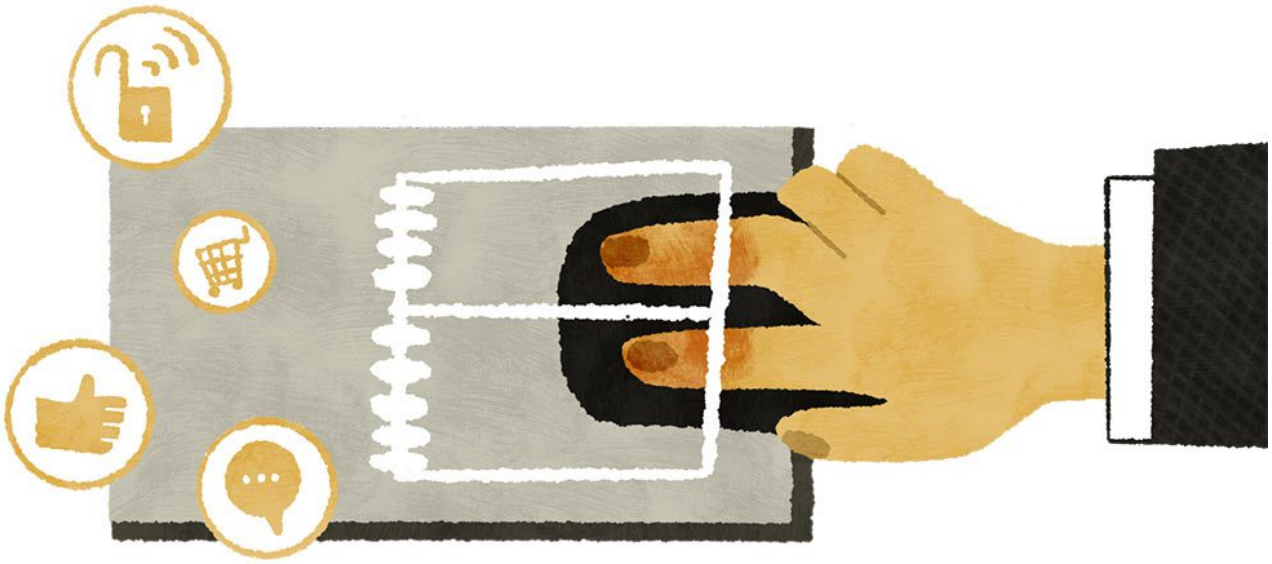

January 1, 2024

By [Tawfiq Alashoor]

It’s the question we face countless times a day: accept, reject or personalize your cookies? Which option you choose says a lot about what you think about privacy.

Our values, beliefs, attitudes and assumptions about privacy have been built up over many decades, since long before today's digital world existed. And although most of us would say we care about privacy, our real-life, day-to-day actions often tell a different story.

This say-one-thing-do-another privacy paradox explains why, in a 2018 study conducted 30 days after General Data Protection Regulation (GDPR) went into effect in Europe, 76% of people browsing the web still clicked "accept all" without a second thought. For most people, the unconscious click is ingrained behavior, detached from any genuine privacy concerns they have about the use of their personal data by third parties.

However, there are signs such behaviors are changing. A 2021 survey found less than 32% of people always accepted cookies when they opened a website, going down to 25% for people aged 45-54. Might people be growing more aware of protecting their privacy online?

Privacy is the focus of my research, specifically in the digital realm, which raises a plethora of issues at the personal, organizational and societal levels.

Traditionally, preserving information privacy has focused on cyberattacks and data security breaches, sensitizing us to the dangers of phishing and the importance of password protection. Yet there is another aspect to privacy beyond information security; it's about the small decisions we make on a daily basis to give our personal data away.

As the number of digital interactions we undertake in the course of a normal day has skyrocketed, we need to enhance our awareness and understanding of the factors at play that influence our privacy decisions — in order that we might make better ones in accordance with our values.

The privacy calculus

The introduction of ChatGPT has turbocharged a technology that has been around for well over a decade: voice assistants. Compared with ChatGPT, voice assistants like Apple's Siri and Amazon's Alexa are "dumb as a rock," Microsoft CEO Satya Nadella told the *Financial Times*. Today's generative AI chatbots can do much more than just cue up a song or turn off the lights on command; instead, they converse with us. But to do so, they need our personal data, and the chatbot's responses grow ever more relevant and accurate the more personal information we feed into it.

During the COVID-19 lockdowns, the use of telemedicine chatbots became indispensable to

help make diagnoses and check up on the health of housebound patients remotely, but it required people to willingly disclose lots of their private medical information via apps. Some might argue this constitutes legitimate self-disclosure — in other words, people have made a privacy calculus, weighing up the costs of giving away their sensitive information in exchange for receiving some greater benefit (healthcare).

Indeed, this pact of reciprocity informs a line of research related to data disclosure. Chatbots imitate the social norm of small talk, whereby the more my conversation partner discloses information about him or herself, the more I feel compelled to reciprocate and disclose information about myself. This builds trust between parties, and an interpersonal relationship is developed.

Data-seeking companies try to replicate this dynamic by having chatbots engage in reciprocal conversations. But there’s a fine line between encouraging people to reveal more of their information in ways that build trust for some positive purpose (like better healthcare) and manipulating people to part with their private information, which leaves them feeling exploited, vulnerable or used, and which, in turn, damages the overall utility of the system if it is no longer able to deliver trustworthy results.

In research with colleagues, I have tested the environments and conditions that influence users giving their private information. Here I summarize some key findings.

The anthropomorphism effect

Anthropomorphism is when people ascribe humanlike characteristics to a nonhuman entity. By this, we're not talking about the mere fact of a chatbot having a humanlike voice; rather, we wanted to study the effect that a chatbot with a humanlike voice might have on a user in terms of how they interacted with the technology and the level of trust they felt afterward.

For our study, we developed a chatbot called Amanda and ran two experiments: one using text-based chat and one voice-based. We asked participants to join a conversation with Amanda, starting by introducing themselves, then engaging in reciprocal questions and answers.

In the treatment group, Amanda would disclose some information like this: "You know what makes me furious? Sometimes people mispronounce words, or even worse, speak quietly and expect me to understand what they say. What are some of the things that make you furious?" And when the person disclosed what annoyed them, Amanda would utter an emotional reaction like, "Tell me about it!" designed to elicit high intimacy, whereas in the control group, Amanda's utterances were kept procedural: "Alright, let me ask you the next question." After the interaction, participants filled out a questionnaire to test subsequent

trust in conversational agents.

In both experiments, we found that when the chatbot engaged in reciprocal self-disclosure with more prototypically human characteristics, these anthropomorphism features led users to ascribe more trustworthiness to the system, both cognition-based (i.e., a rational decision to trust the entity based on its perceived competence and superior thinking abilities) as well as affect-based (i.e., a decision to trust based on a perceived emotional bond or feelings of empathy or rapport).

Knowing this can be helpful to chatbot developers to improve the conversational components of their apps, specifically those where reciprocal self-disclosure and trust are central to the delivery of critical functions, like healthcare.

However, it is also important to make users aware of how cybercriminals, for instance, could exploit these same features to manipulate people into making disclosure decisions inconsistent with their privacy preferences.

On the one hand, anthropomorphizing chatbots to be more closely aligned with a user's own gender, ethnicity, accent, etc., helps in making them more relatable and delivers a more personalized, user-friendly experience.

On the other hand, using anthropomorphism for manipulative self-disclosure — lulling users into a false sense of familiarity so that they give too much of themselves away — lacks transparency and raises a host of ethical issues related to privacy.

Too tired and too happy to care

In a separate, award-winning study, my co-authors and I conducted another series of experiments, this time to test the effect of various contextual factors on users' willingness to share their personal information online.

First, we wanted to check the effect of being mentally overloaded or tired on information sharing, given that people tend to let their guard down when they are cognitively depleted. So, we gathered participants, controlling beforehand for their privacy concerns and their proclivity for sharing personal information. We presented them with a cover story about a mobile health app, explaining that their input was vital for the development of this app.

Then, we made them complete a series of timed reading and writing tasks designed to

deplete them. After the tasks, we asked them to report their mood, using adjectives ranging from happy and content to fed up and grouchy. And based on this, we measured their actual disclosure behavior.

We repeated this experiment twice, once to manipulate users' mood state, and again to manipulate both depletion and mood states.

These experiments helped tease out the contextual factors that affect when people are more or less likely to act in accordance with their privacy concerns. The privacy paradox is real, albeit conditional on a combination of external (cognitive) and internal (mood) factors.

It's not just that people who are cognitively depleted are more susceptible to oversharing despite their stated privacy concerns, but if they are cognitively depleted *and* in a happy mood, their propensity to act *opposite* to their own stated privacy concerns goes up.

Think of your own behavior on a Friday night after an exhausting week at work. TGIF! You pull out your phone and start cheerfully interacting on social media. This is when you might mindlessly click "accept all" and unwind with a fun online poll in which you freely let flow a veritable goldmine of information on your likes, dislikes, personality traits and political opinions, when under different conditions — say, mentally alert and grumpy — you would never do so.

(Note: This is not to say that negative affect is preferable. Other research shows that negatively primed or frustrated users can be more prone to violate online security policies. As I will explain later, the context matters.)

Letting your guard down

The dangers of making undesirable privacy disclosures when being in too relaxed a state of mind was highlighted in another study of mine, this time using stress-relief breathing techniques. With colleagues from Copenhagen Business School, we had participants do typical breathing exercises, then had them answer questions related to demographics, health, lifestyle and other personal information.

Our findings suggest that when individuals are relaxed, they're more likely to reveal personal information compared with those in a neutral state.

The point of all these studies is not to teach data-gatherers how to trigger the mental and

emotional states that would foster users yielding more of their information against their better judgment. Rather, organizations and app developers should consider incorporating responsible guidelines and features, such as pop-up alerts, to promote mindful privacy practices, so that users' disclosure choices actually match their values.

Why we need privacy training

What my research shows is that individuals need to appreciate the impact of psychological and environmental factors on privacy decisions, understanding how they could end up making choices inconsistent with their preferences if they're not careful. People must engage their brains and apply more effort when processing information and interacting online — admittedly, easier said than done.

Companies also bear responsibility for balancing the economic possibilities they might gain from leveraging people's data with those same people's desire (and right) for privacy. Increasingly, we are seeing websites changing their data collection methods by offering a variety of privacy awareness pop-ups, enabling users to make easy, and at times quite granular, privacy choices before they can continue browsing.

(Some of these privacy options are so exhaustive that users may give up and exit, or just surrender and "accept all" as a consequence of mental overload again.)

Such moves are signs that privacy regulation is having an effect on companies as well as on individuals. The end result should be better privacy protection over time, even if the reason companies and individuals are choosing privacy settings is because they have to, not out of any deep conviction for privacy standards and awareness.

SETA vs. PETA

Despite some regulatory progress, policymakers still have a long way to go. GDPR in Europe, and similar consumer privacy acts in a few U.S. states, represent important steps toward instilling the notion of privacy protection in the public.

One concern I have, though, is that, when it comes to information protection, the emphasis is on Security Education, Training and Awareness (SETA) — that is to say, information security is confined to technical protection mechanisms. So, we are taught to exercise vigilance regarding our passwords, software updates, two-factor authentication, cyber threats and so

on.

However, there's another level of information protection no less deserving of our attention, which is why I argue we need to introduce complementary Privacy Education, Training and Awareness, or PETA, in addition to SETA.

Consider another study in which we ran a classic SETA workshop. The idea behind SETA courses is that once individuals learn about the risks, they will be less susceptible to making the human errors that account for many data security breaches, such as clicking on a link in a phishing email.

We divided participants into two groups: in one we gave them a privacy awareness nudge, and in the other we didn't. Everyone completed three surveys: one immediately before the workshop, one immediately after, and the last one two weeks later. Survey questions included things like, "Do you use the same password for multiple websites? Do you include your birthday when creating a password? Do you share some of your personal details on social media networks?"

The control group could give yes/no answers or skip any question, while the treatment group could answer yes/no or choose a third option of "I prefer not to provide this information," making people explicitly aware of their ability to withhold private information.

Participants who received a privacy nudge shared less personal information compared with those who didn't receive any nudge.

However, we did not expect to find that, for both groups, whether they had received a privacy nudge or not, even after learning about security risks and protections on the SETA course, they still admitted to disclosing sensitive information to a high degree.

This is an alarming finding, which reinforces my call for PETA.

How would a PETA program be different?

- First, it should start by explaining the distinctions between privacy and security.
- Next, learners need to be trained on the psychology and behavioral economics of privacy decisions, related to some of the issues touched on in this article.
- Finally, there should be thorough discussions of privacy theories and the importance of privacy regulations and best practices.

Crucially, there needs to be thinking on two levels, which the Nobel Prize winner Daniel Kahneman described in his book, *Thinking, Fast and Slow*, as System 1 and System 2 modes of thought.

Most thinking is automatic, intuitive and less effortful (System 1), like clicking "accept all" without thinking through the ramifications of this privacy decision. The goal of PETA would be to accustom people to slower, deliberate, more effortful thinking (System 2), like pausing for a few moments to process information related to a privacy decision, allowing sufficient time to weigh up the potential pros and cons of allowing third parties to track you online.

If human error is still the biggest cause of cybersecurity breaches, then training humans to resist privacy fatigue and equipping them with mental models for informed privacy decision-making seems like a sensible way forward.

The general relativity of privacy

In 2023, an Australian doctor, Graeme Siggs, was watching TV and happened to see an interview with a professor, Dan Angus. The doctor noticed a brown spot on the professor's cheek and, being a skin cancer specialist, took a keen interest. He went online and searched up old photos of the professor, observing how the spot had grown bigger over time. He contacted the professor and urged him to get it checked. It turned out to be melanoma and, thanks to the doctor's timely online intervention, Prof. Angus got life-saving surgery.

Question: was this a violation of Prof. Angus's privacy? After all, he hadn't consented to having his online photos used by a stranger in this way.

This real-life story illustrates a new theory I am developing called the "general relativity of privacy." Putting a spin on Einstein's "general theory of relativity" — which posits that all states are relative, according to a space-time continuum — I suggest that all privacy is relative, according to a *contextime* continuum.

Think about it: While it's true that space, such as geolocation, does identify us in the digital world, the most significant shaper of our online presence in relation to privacy is *context*.

The surfacing of Prof. Angus's old photos, even if it were in violation of his privacy preferences, was beneficial within that specific context, and he acknowledged as much when he wrote to Dr. Siggs, thanking him for saving his life: "I've got two young girls, a wife who loves me, a family that loves me. Graeme, when I come to Adelaide next, I owe you a beer

and a hug."

Conversely, the surfacing of old off-color jokes by U.S. comedian Kevin Hart, which may not have offended many people in 2009, were no longer considered funny in 2019, and Hart had to apologize and withdraw from hosting the Oscars that year. The context was decidedly different.

Social norms and privacy laws change. Moreover, in the digital world, time is a timeless construct — yesterday's photos, videos and posts are perennially present, but our perception of them changes according to the context.

Consider another scenario: a digital platform uses AI to generate personalized ads. The success and reception of these ads depend heavily on their context and timing.

For example, an ad for fitness gear might be well received if it were something you had just been searching for online, but the same ad might feel intrusive if it popped up during a private conversation you were having on a messaging app with a friend.

Given my findings so far on the dynamic nature of personal data's meaning and value, which shift according to situational factors (e.g., cognition, emotion, ethics, etc.), I believe the "general relativity of privacy" represents an important new direction for future research.

Rethinking the whole privacy concept

In the meantime, all of us, in our private lives and in our organizations, should be taking time to think more deeply and seriously about the whole privacy concept.

For managers of products and services, particularly those whose business models rely on chatbots and AI, they should prioritize privacy measures and procedures, paying special attention to *contextime* in the development process. Forward-thinking managers and CIOs should adopt a privacy-by-design mindset.

Ask yourself: Do your pop-up alerts genuinely help users think through their privacy choices and make decisions aligned with their values? Or are you exhausting or manipulating users into unethically surrendering more of their data?

Instead of wasting annual budgets on ineffective SETA programs, organizations should invest in better PETA programs that equip people to make more effective privacy calculi. This could help customer service departments, for example, improve information requests for

authentication, reducing identity impersonation and avoiding cybersecurity threats and attacks.

Privacy and personal data are destined to become even more valuable currency, eventually serving as the basis of a "trust score." Apple has been doing this for several years now, essentially "scoring" users' browsing, purchases, searches and downloads based on a historical record of their digital activities, ostensibly to prevent fraud and other malicious activity.

We are on the cusp of a revolutionary new digital privacy age. What will you do the next time you're asked to accept cookies?

SOURCES: "My name is Alexa. What's your name? The impact of reciprocal self-disclosure on post-interaction trust in conversational agents" by Kambiz Saffarizadeh, Mark Keil, Maheshwar Boodraj & Tawfiq Alashoor. *Journal of the Association for Information Systems* (2023).

"Too tired and in too good of a mood to worry about privacy: explaining the privacy paradox through the lens of effort level in information processing" by Tawfiq Alashoor, Mark Keil, H. Jeff Smith & Allen R. McConnell. *Information Systems Research* (2022). This paper won a 2024 Best Published Paper Award from the Communication, Technology & Organization (CTO) Division of the Academy of Management (AOM).

"Take a deep breath and tell me all about it: an experimental study on the effect of breathing on privacy decisions" by Tawfiq Alashoor, Andreas Blicher & Rob Gleasure. 15th Anniversary NeuroIS Retreat (2023).

"An online randomized field experiment on the importance of Privacy Education, Training & Awareness (PETA)" by Tawfiq Alashoor et al. 3rd International Conference on Computing and Information Technology (2023).

"It is seriously time for Privacy Education, Training & Awareness (PETA) programs" by Tawfiq Alashoor, chapter in the forthcoming book *Effective methods for teaching business-related topics during and post crisis*.

"Mind your digital grave: how digital traces mutate into digital zombies" by Mazen Shawosh, Tawfiq Alashoor & Nicholas Berente. TREO (Technology, Research, Education & Opinion) Talks, held in conjunction with the 43rd International Conference on Information Systems

(2022).

A version of this article is published in IESE Business School Insight magazine (Jan.-April 2024).

https://www.youtube.com/embed/0zOH30t_8Tw

Array

Tawfiq Alashoor

Assistant professor in the Department of Operations, Information & Technology at IESE. His research focuses on privacy and cybersecurity decision-making in emerging AI-supported technologies, including apps, conversational agents and robots.