"You do understand that people don't trust technology?": Explaining Trusted Execution Environments to Non-Experts

McKenna McCall*†, Carolina Carreira*†‡, Miguel Flores†, and Lorrie Faith Cranor†
†*Carnegie Mellon University*
‡*IST University of Lisbon, INESC-ID*
{*mckennak, cdacunha, miguelfl, lorrie*}*@andrew.cmu.edu*

***Abstract*—Trusted Execution Environments (TEEs) protect confidentiality and integrity of trusted applications by creating an isolated environment for executing code. Prior work has shown that users may feel more comfortable sharing data when they know it will be protected by a TEE—especially if they understand what a TEE is. In this study, we evaluated text-based explanations introducing TEEs to non-experts. We analyzed existing TEE explanations to develop candidate explanations and evaluated them via vignette scenarios with 966 crowdworkers. The explanations that enhanced understanding most were *non-technical* ones that highlighted specific threats that can be *prevented* by a TEE. Surprisingly, even the explanations that enhanced understanding had little effect on willingness to use the TEE-enhanced technology. These results provide insights into ways to communicate technical security concepts more effectively but also suggest that explaining security technology might not be enough to address users' privacy concerns.**

1. Introduction

While the demand for users to share their data grows, consumers are expressing concern and confusion about how their information is used [1]. Confidential computing [2] seeks to protect users by restricting computations on sensitive data to Trusted Execution Environments (TEEs). These environments guarantee the authenticity of the executed code, the integrity of the runtime states, and the confidentiality of the code and data [3]. Confidential computing and TEEs have applications in AI and machine learning [4], [5], [6], IoT [7], and blockchain smart contracts [8].

TEEs are not only being explored for their technical strengths. Prior work also suggests that when users are made aware of cloud-based TEEs in home IoT devices, they may feel more comfortable with their data being collected [9]—especially when they understand what a TEE is. However, this work did not investigate how technologists should explain TEEs to end-users, which is itself a challenge. Unlike some security concepts like passwords, most people are not familiar with TEEs, or even the technologies they rely on. While fully explaining TEEs to a non-technical audience may not be feasible, understanding all of the technical details may not be necessary to understand enough of the security benefits they offer to help users feel secure. Indeed, the original study investigating the impact of TEEs on comfort evaluated understanding based on just three high-level TEE concepts [9].

In this study, we evaluate strategies for explaining TEEs to enhance both *understanding* of the capabilities of a TEE as well as *comfort* using TEE-enhanced technologies. Ideally, an explanation would be nuanced enough to communicate what guarantees TEEs offer—without over-promising or being overly pessimistic about risks, which could discourage people from using technology. We based our explanations on common themes we found in existing TEE explanations from technical websites, forums, research papers, and popular media. We evaluated candidate explanations through a series of True/False questions via two online surveys of 966 Prolific crowd workers. We used vignettes in our surveys to evaluate explanations across different scenarios where TEEs might be used. We include both home IoT scenarios as well as AI and medical research applications, as they have also been identified as potential use cases for TEEs [10].

Our first survey addresses two research questions:

- **RQ1:** Which explanations improve TEE *understandability* for non-experts? Is there a best overall explanation or do different scenarios benefit from different explanations?
- **RQ2:** Which explanations enhance *willingness* to use the TEE-enhanced technology? Which ones promote the feeling that data will be *safe*?

Based on the results of our first survey, we developed an FAQ to supplement our explanations by answering real questions asked by our participants. We also asked follow-up questions to better understand what contributes to the perception of safety. In the second survey, we answer the following research questions:

- **RQ3:** Does an FAQ further improve understandability? Does it increase willingness to use TEE-enhanced technology or the feeling of safety?
- **RQ4:** Which aspects of TEE scenarios contribute to the belief that data would be safe/unsafe?

To the best of our knowledge, ours is the first study

* *The authors contributed to this work equally.*

to investigate strategies for explaining TEEs in a way that is accessible to non-experts. While we found that many existing explanations use technical jargon and focus on broad security guarantees offered by the TEE (e.g., attestation, confidentiality, and integrity), what performed best in our experiments were *non-technical* explanations that highlighted specific attacks *prevented* by a TEE. We also found that people generally answered comprehension questions correctly when we provided information directly in our explanations or FAQs, but struggled to answer questions that required them to make inferences based on our explanations.

Surprisingly, in contrast with prior work [9], we found that our explanations had little effect on willingness to use TEE-enhanced technology or feelings of safety. We believe that this is due to methodological differences between ours and the previous study (namely that we focus on high-level feelings of comfort and safety, while they focused on specific data-sharing conditions) and, importantly, our observation that TEEs cannot address all of the privacy concerns raised by our participants. These results provide insights into ways to communicate technical security concepts more effectively but also suggest that explaining security technology, while useful for improving transparency, might not be enough to address users' privacy concerns.

The rest of the paper is organized as follows: Section 2 covers background and related work; Section 3 explains our methods for collecting and analyzing existing TEE explanations to identify themes for testing via our surveys; Section 4 describes our survey methods; Sections 5 and 6 present the results of our two surveys; Section 7 includes additional discussion; finally, Section 8 concludes.

2. Background and Related Work

In this section, we describe background on TEEs and where they are used. Next, we outline some related work on the importance of communicating with users about security. Finally, we summarize work that has attempted to explain technical security concepts to end-users.

2.1. Trusted Execution Environments

TEEs are combinations of several security processes, including hardware security extensions, cryptographic modules, secure distributed systems protocols, and more. According to Sabt et al. [3], a TEE can be defined as a tamper-resistant processing environment that guarantees the *confidentiality* and *integrity* of the executed code and data (preventing unauthorized reading and modification, respectively). A TEE also provides remote *attestation*, a process where the TEE proves to a remote verifier (such as a server) that it is operating securely and that the integrity of its code and data has not been compromised.

TEEs can be found in many Android phones. Authentication in Android is typically handled by code residing in a TEE based on ARM TrustZone [11], a set of security extensions that enable ARM processors to run in two distinct modes—secure and non-secure. Most Androids also use TEEs to process mobile payments, secure banking, device reset protection, and detect malware [12].

Another type of TEE, Intel SGX, allows applications to create protected areas in memory in some Intel CPUs [13] and also has applications in smart home devices. Ayoade et al. [7] propose using Intel SGX for decentralized data management in smart home applications. There are several other TEE technologies in the realm of confidential computing that target cloud computing, such as AMD SEV-SNP [14], Intel TDX [15], and ARM CCA [16].

Confidential computing is particularly relevant in the medical domain, because patient data, such as data from clinical trials, has some of the strongest legal protections in the US [17], [18]. Data aggregation enabled by confidential computing gives healthcare providers the ability to improve patient or research outcomes while safeguarding patient privacy [19]. TEEs could be used to protect machine learning for medical applications and ensure compliance with medical regulations [5].

2.2. Importance of Understanding Security Technology Basics

While technical expertise may be required to understand the details of security technology, even a basic understanding can help users make informed decisions and better protect themselves, while misconceptions and poor usability can lead to worse outcomes. An early example of opaque security technology hindering users can be found in the seminal paper by Whitten et al., which provides empirical evidence that users who lack understanding about how public key encryption works behave in ways that undermine the security and privacy of their encrypted email. The authors conclude that an unusable or incomprehensible security mechanism will not be used effectively and thus not provide security [20].

Misunderstanding the functionality and limitations of security tools can also lead users to develop a false sense of security [21], [22]. This overconfidence may lead them to engage in riskier behaviors under the mistaken belief that they are protected. For instance, Bravo-Lillo et al. [23] showed that misconceptions about web browser security warnings can give users an illusion of safety. In addition, interview [24] and survey studies [25] have investigated users' misconceptions about how attackers steal passwords, finding that misconceptions led users to believe vulnerable passwords were secure.

Misconceptions about security tools and design choices can also hinder their adoption [26]. Users may also face usability challenges with cookie banners due to their design [27]. Similarly, users struggle to comprehend iOS privacy labels because of jargon and unfamiliar terminology [28], [29]. On the other hand, informing users about security technology can have a positive impact. For example, Furnell et al. [30] find that more information can motivate users to choose strong passwords. When security- and privacy-enhancing technologies are mentioned to users

as part of the consent process, users need a basic understanding of what protections these tools can and cannot offer if they are to make an informed decision. The European Union's GDPR [31] requires that organizations must provide clear and accessible information to ensure users understand how their data is used. Similarly, in the US, HIPAA [18] mandates that healthcare providers give patients clear information about their privacy rights and how their medical information is shared.

While most of the work in this space focused on technologies users employ to protect themselves, explaining TEEs is a fundamentally different task since they are hidden from users. Our goal in explaining TEEs is more to improve transparency and comfort than to change user behavior.

2.3. Explaining Technical Concepts

Our study builds upon prior work on short explanations to communicate technical concepts and evaluate them using online surveys [32], [33], [34]. Prior work on explaining security concepts mostly focused on perceived security [34], [35] and did not address comprehension or focus on non-TEE related contexts [32], [33], [33], [36].

Several research studies have proposed and tested explanations of other technical security concepts with end-users. Research on formal verification has also emphasized the importance of communicating technical concepts to non-technical audiences and identified it as a priority and a challenge for future work [37].

Xiong et al. [38] attempted to explain differential privacy with experiments to investigate the effects of different communication approaches. They found that, despite the positive effect of the explanations, participants struggled with understanding some of the more technical jargon. Karegar et al. [39] studied a possible solution by addressing the impact of metaphor-based explanations of differential privacy. They found that metaphor explanations can help understanding but can also lead to misconceptions. Cummings et al. [40] attempted to design better explanations about differential privacy but highlighted the difficulty of crafting explanations that satisfy user interest and preserve the integrity of the technical content.

Similar to our work, Akgul et al. [33] investigated whether text-based explanations improve users' mental models of encryption. They found that changing pre-existing mental models can be challenging, but educational interventions can work. Interestingly, they concluded that their explanations may have slightly oversold the capabilities of encryption. Shen et al. investigated users' understanding of smartphone permissions and observed that short explanations within user interfaces led to better comprehension. The authors found that adding information to permissions dialogues made it more clear to users how their choice affected the way that their location would be tracked [32]. When it comes to explaining encryption, not all authors agree. Distler et al. [34] attempted to explain encryption and concluded that explaining encryption does not necessarily maximize perceived security. They focused primarily on the feeling of security and did not study users' comprehension of encryption. This is also the case for Stransky et al. [35], but their results suggest that using text disclosures about encryption makes users feel more secure seem more effective than iconography.

OS and browser security warnings are designed to provide actionable security information to non-technical users [36], [41]. Wu et at. [42] show that warning notifications in Signal can improve comprehension of the purpose of security mechanisms and promote favorable privacy outcomes. Well-designed password meters can be an effective communication tool to inform users about their password complexity and are a good way to provide actionable feedback about password strength [43]. Privacy and security "nutrition" labels are designed to provide succinct information to users that can inform their decision-making [44], [45].

To our knowledge, the only other attempt to communicate about TEEs to end-users is from Musale et al. [9], who investigated the impact of TEEs on data-sharing preferences. They also looked at the impact of understanding TEEs, finding that people who understood TEEs were more likely to be comfortable sharing their data. For example, they found that participants who understood TEEs were significantly more comfortable with their data being collected if they were "notified" of the data collection than those who did not understand TEEs. To assess TEE comprehension, the authors asked three True/False questions about secure storage, secure computing, and remote attestation. In one question, they ask whether the statement "non-authorized persons can modify/change the nature of the algorithm being used or gain access to the image database" is true or false. While their study focused on understanding the impact of TEEs on existing privacy norms, ours focuses on how to effectively explain TEEs. For this reason, we conduct a broader assessment of 10-12 questions that address different aspects of a TEE. We also focus on high-level feelings of comfort and safety instead of specific data-sharing conditions. While their work did attempt to explain TEEs to their participants, the goals and methodologies of that study were fundamentally different from ours.

3. Developing Candidate TEE Explanations

In this section, we describe our approach for developing candidate TEE explanations, which is based on a technique from prior work on differential privacy [40] and other guidelines for writing effective explanations [34], [46]. First, we describe how we analyzed existing TEE explanations from technical websites, forums, research papers, and popular media to identify common themes. Next, we explain how we used these themes to develop explanations for evaluating in our study.

3.1. Identifying Existing TEE Explanations

We conducted a Google search using the term "Trusted Execution Environment" and restricted results to the last five years. The first five pages of results included 42 unique

Code	Description	Frequency
Reputation	Leverages pre-existing trust/reputation of recognizable companies	2
Verified	Application running in the TEE is verified	2
Attestation	Process to check that the software supporting the TEE is the code we expect	4
Trust	Explanation mentions the word "trust"	5
Unsubstantial	Generic/un-detailed description	8
Threat	TEE protects against untrusted OS/peripherals	8
Techniques	Describes particular TEE (e.g., Intel SGX, Arm TrustZone)	10
Cryptography	Mentions cryptographic concepts	10
Technical	Explanation uses technical terminology (e.g., "confidentiality," "attestation")	11
Integrity	TEE prevents unauthorized modification	16
Prevents	TEE prevents some undesirable behavior	17
Secrecy	TEE prevents unauthorized access	21
Isolation	TEE ensures isolation from the rest of the system	23
Hardware	Mentions that a TEE is *hardware*-supported	23

TABLE 1: Codebook for explanations found in the wild and how frequently each code was identified in the explanations. Each explanation could have up to 7 different codes.

URLs that had 32 TEE explanations. These results came from diverse sources, mostly aimed at an audience of technical experts. We obtained eight additional explanations through searches targeting well-known, general audience platforms like the New York Times, Medium, and Forbes.

We removed 12 sources from the initial 50 that did not include substantive TEE explanations (i.e., the source mentioned TEEs, but did not provide any explanation about what they are). We removed two others because their explanations were incorrect or misleading. We analyzed the explanations from the 36 remaining sources to identify themes to test in our experiments. 19 sources came from technology-focused websites from companies that provide TEEs (e.g. Intel, NVIDIA, AWS, Google Cloud). Media sources, including general audience magazines and news websites, accounted for 6 explanations. The remainder came from a mix of scientific publications, forums, social media websites, governmental websites, and Wikipedia (see supplementary materials [47] for all explanations and sources). This diversity of sources ensured a broad spectrum of explanations to reflect the variety of information available to the public.

Two authors independently reviewed the explanations to identify themes and assigned a code for each theme. The number of codes per explanation largely depended on the size and complexity of the text. We ended up assigning eight codes to the most complex explanation. The coding process began with a few initial codes based on our prior knowledge of TEEs. Codes were added based on themes that emerged during the analysis. After reviewing all explanations, the coders discussed the themes to develop a shared codebook. They repeated the process of reviewing explanations, coding, discussing all disagreements, and refining the codebook twice more until they reached 100% agreement. The final codebook has 14 codes. Here we list the codes that appeared in at least five explanations with the number of explanations in which they appeared in parentheses: *Isolation* (23), *Hardware* (23), *Confidentiality* (21), *Prevents* (17), *Integrity* (16), *Technical* (11), *Cryptography* (10). *Techniques* (10), *Threat* (8), *Unsubstantial* (8), and *Trust* (5). The complete codebook, including a description of each code, can be found in Table 1.

3.2. Designing Candidate TEE Explanations

We developed new explanations that used key themes found in existing explanations and iterated on their wording through pilot testing. We designed our explanations to be composable so that we could separately test each component in controlled experiments. We identified *Confidentiality*, *Isolation*, and *Integrity* as themes that seemed fundamental to a TEE explanation and should be included in every candidate explanation using either *Technical* or *Non-technical* language. In addition, we identified *Hardware*, *Trust*, and *Prevents* as themes that might aid understanding. We decided to test explanations that included *Hardware*, *Trust*, or an *Unsubstantial* explanation, as well as explanations that either explained what a TEE *Prevents* or includes *No Prevents* clause. In order to keep the number of treatments in our survey manageable, we did not evaluate the less common themes.

As shown in Figure 1, the structure of each explanation is: (1) a high-level sentence introducing the concept of a TEE as a security mechanism (one of the following themes: *Hardware*, *Trust*, or *Unsubstantial*), followed by (2) a sentence introducing the concepts of isolation, confidentiality, and integrity in either technical or non-technical language (one of the following themes: *Technical* or *Non-technical*), and, only for some explanations, (3) a third sentence introducing a specific threat that a TEE can prevent (theme: *Prevents* or *No Prevents*). Our candidate TEE explanations are the set of all 12 possible combinations of themes that follow the structure above. Complete TEE explanations are shown in Appendix A.

4. Survey Methods

To evaluate our candidate TEE explanations, we conducted two online surveys. The initial survey focused on evaluating our explanations (**RQ1-2**), while Survey 2 tested some follow-up research questions (**RQ3-4**) based on the results of Survey 1. In this section, we describe the methods we used to conduct and analyze data from both surveys.

4.1. Survey 1: Evaluating TEE Explanations

The purpose of Survey 1 was to evaluate our candidate TEE explanations to identify which themes, adopted from existing explanations, are best at enhancing understanding (**RQ1**), willingness to use TEE-enhanced technology, and feelings of safety (**RQ2**). We constructed four scenarios to use in our surveys. Each scenario describes a situation where personally identifiable data is collected for some purpose. The data collected in each scenario is the same, but the *setting* and *purpose* of collection depends on the scenario.

We choose medical research and smart home settings because they are both promising TEE use cases [5], [7],

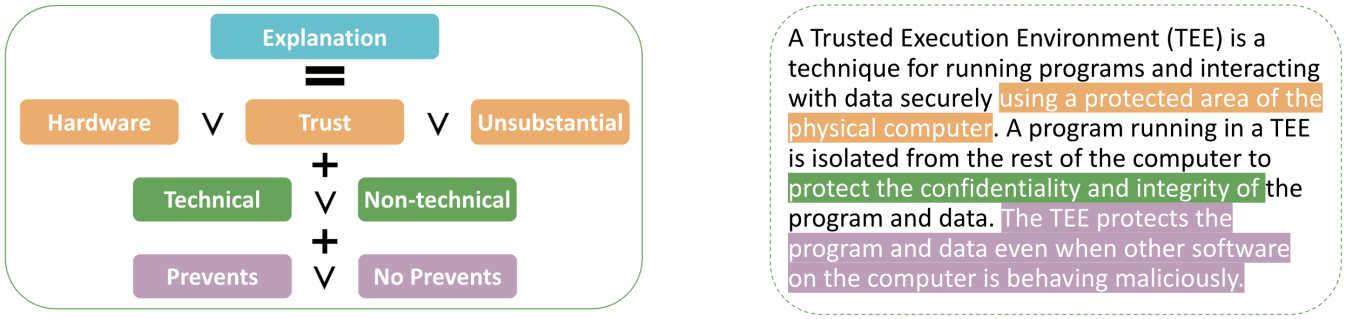

Figure 1: Diagram illustrating the design of the initial TEE explanations and an example candidate TEE explanation.

[14], [15], [16], [19]. We chose not to use a smartphone scenario, despite it being another TEE use case [11], [12] because we wanted to focus on emerging TEE applications. The fact that most people already have smartphones [48] could also have biased our results. In the medical research setting, we ask participants to imagine there is a medical research study that involves collecting personal information if they choose to participate. In the smart home setting, we ask them to imagine shopping for a smart device that will collect personal information about them if they choose to purchase it.

We also have two variations of each scenario, one where the purpose of the data collection is to develop technology involving AI and one not involving AI. We included AI in our scenarios because the adoption of AI has been growing in both medical research [49] and the smart home context (e.g., Google Home [50] and Alexa [51]) and there is evidence that people are wary of AI [1], which could factor into their willingness to use the technology.

Each participant receives one medical research scenario and one smart home scenario randomly. They are also randomly assigned the technology with AI or without AI in each scenario they receive. For example, one participant may receive the "medical research with AI" scenario followed by the "smart home with AI" scenario while another may receive the "smart home without AI" scenario followed by the "medical research with AI" scenario. For each scenario, participants are told that the data is stored in the cloud and protected by a TEE. The complete scenario text for all four scenarios is shown in the supplementary materials [47]. The scenario text is followed by a random candidate TEE explanation (from the set of 12 explanations), and they receive the same explanation for both scenarios.

In the first part of the survey, we introduce the scenario and confirm participants are paying attention by asking them to select the purpose of the medical research study or what device they're shopping for. If they answer incorrectly, they are asked to re-read the scenario text and try again.[1]

Next, we asked participants to rate their willingness to participate in the medical research study (for the medical research scenarios) or willingness to purchase the smart home device (for the smart home device scenarios), and how safe they believe their data would be, each on a 3-point Likert scale. We then evaluated comprehension via 10 True/False questions and allowed participants to ask us any lingering questions they had about TEEs. We ended the survey by collecting demographic data: age, gender, highest education level, experience/education in computing, as well as prior experience with medical research and smart devices. The complete survey instrument can be found in the supplementary materials [47].

We solicited feedback from TEE experts outside of our team on the technical accuracy of our explanations, scenarios, and comprehension questions. We refined the survey questions through multiple pilots.

4.2. Survey 2: FAQs and Understanding Aspects of Feeling Safe

Survey 2 was similar to Survey 1 other than the introduction of an FAQ to answer some of the most frequently asked questions participants had in Survey 1 (**RQ3**) and asked additional questions to understand which *aspects* of the scenarios led to the belief that data shared with the TEE-enhanced technology would be safe or unsafe (**RQ4**).

Participants were randomly assigned to one of three FAQ conditions: one where they were *Shown* the FAQ on its own page after the first scenario was introduced (which they could not click past for 60 seconds) *and* as expandable text on subsequent pages; one where the FAQ was *Hidden* by default and *only* offered as expandable text; and one where they were not given an FAQ (*None* condition). To keep the number of survey conditions reasonably small, we did not re-test all of the TEE explanations from Survey 1. Since the *Technical* and *No Prevents* themes generally led to worse comprehension scores, we used the *Non-Technical* and *Prevents* themes in all of the explanations. Thus, we had 3 explanation conditions (*Hardware*, *Trust*, and *Unsubstantial*), plus we added a fourth no-explanation condition (*None* condition) to serve as a baseline for the questions about aspects of safety.

Because some participants mentioned in Survey 1 that they do not believe data could ever be "Completely safe,"

1. Eight participants answered incorrectly the first time, but four succeeded after we gave them a second chance at the attention check.

when we asked participants how safe they believe their data would be in Survey 2, we used a 4-point Likert scale, adding "Mostly safe" to the 3-point scale (Completely safe, Somewhat safe, Not at all safe) from Survey 1. We also asked participants to rate how much different aspects of the scenarios contributed to the belief that their data would be safe/unsafe on a 5-point Likert scale and to expand on "anything else" that contributes to those feelings in a free-response field. Finally, we added 2 True/False questions about the topics covered in the FAQ.

Constructing the FAQ. Our FAQ is based on the questions participants asked in Survey 1.

Our FAQ answers three questions:

1) How do TEEs work?
2) How do we know the TEE is working correctly?
3) How are TEEs used in real life?

The answer to the first question includes additional technical details about how TEEs work, specifically Arm TrustZone [52] and Intel SGX [13]. To answer the second question, we described attestation and mentioned that researchers are continuing to develop ways to ensure the applications running in the TEE work as expected. Finally, we used authentication in Android [11], [53] as an example of a real TEE use case in the answer to the third question. The complete FAQ text may be found in the supplementary material [47]. We also provided links to the resources cited in this paragraph (plus general information about confidential computing [54]) at the end of the survey.

4.3. Recruitment

We used the same recruitment process for both surveys. We recruited 469 Prolific participants for Survey 1 and 501 for the second using quotas [55] to ensure approximately equal numbers of men and women.[2] People who participated in Survey 1 were not allowed to participate in Survey 2. Our participants are adults located in the US who are fluent in English. We paid participants $2.50 for Survey 1 (median completion time approx. 10 minutes) and $2.75 for Survey 2 (median completion time approx. 13 minutes).

We reviewed results for low-effort or nonsensical free-text responses (none in either survey) and removed responses for participants who failed both attention checks (none in the Survey 1 and 4 in Survey 2). We were left with 469 responses for Survey 1 and 497 for Survey 2.

Table 2 shows the participant demographics, which are similar for both surveys. Participants were balanced across gender, generally young (73.4% under 45 in Survey 1 and 73.8% in the second), and college-educated (63.3% in Survey 1 and 61.6% in the second). Few participants were familiar with TEEs before taking our survey (around 7% for both surveys) or have a career or formal education in a computing field (16.4% in Survey 1 and 21.5% in the second). Our participants tend to have some experience with smart home devices (81.4% in Survey 1 and 86.3% in the second) but not with medical research/work in the medical field (23.2% and 28.6%).

2. Sample sizes were determined by a rule-of-thumb estimate for the logistic regressions we planned for our analysis [56].

	Initial Survey		**Follow-up**	
	n	%	*n*	%
Gender				
Male	229	48.8%	245	49.3%
Female	228	48.6%	242	48.7%
Non-binary / third gender	11	2.3%	8	1.6%
Prefer not to say	1	0.2%	2	0.4%
Age				
18-24	74	15.8%	82	16.5%
25-34	151	32.2%	164	33.0%
35-44	119	25.4%	121	24.3%
45-54	63	13.4%	81	16.3%
55+	62	13.2%	48	9.7%
Prefer not to say	0	0.0%	1	0.2%
Highest Education Achieved				
Less than high school	9	1.9%	5	1.0%
High school or equivalent	161	34.3%	184	37.0%
Bachelor or associate degree	207	44.1%	225	45.3%
Graduate degree	90	19.2%	81	16.3%
Prefer not to say	2	0.4%	2	0.4%
Familiar with TEEs?				
Yes	35	7.5%	36	7.2%
No	434	92.5%	461	92.8%
Experience in Computing?				
Yes	95	20.3%	107	21.5%
No	374	79.7%	390	78.5%
Experience With Smart Homes?				
Yes	382	81.4%	429	86.3%
No	87	18.6%	68	13.7%
Experience with Medical Research/Work?				
Yes	113	24.1%	142	28.6%
No	356	75.9%	355	71.4%
Total	469	100%	497	100%

TABLE 2: Demographics of participants for both surveys.

4.4. Qualitative data analysis

Our study has two open-ended questions. In the first open-ended question, we ask participants if they have any questions about TEEs (this question is in both surveys). In the second, we ask about aspects of the scenario that contribute to their belief that their data would be safe or unsafe (this question is only in Survey 2). In this section, we describe how we analyzed these questions. Codebooks, including descriptions of the codes, can be found in supplementary materials [47].

Questions about TEEs. The questions asked by participants in Survey 1 were coded by two of the authors. Initially, one of the coders reviewed the participants' answers, constructed the codebook using thematic coding, and trained the other coder on the codebook. Next, the coders independently coded all responses, met to discuss disagreements and update the codebook, and then re-coded the answers again.

This process was repeated two times until all disagreements had been resolved. We started with the same codebook for Survey 2 and involved a third author as a coder. The same initial coder reviewed responses and trained the other coders on the codebook. Then, all three coders independently coded all answers, meeting to resolve differences and update the codebook. This process was repeated twice until all 100% agreement was reached.

Aspects contributing to feeling data is safe or unsafe. The second open-ended question about aspects of safety was analyzed much like the first. One coder started by reading through all the answers, developing a codebook, and training the other two coders. After, each coder independently coded all responses, meeting to resolve differences and update the codebook. This process was repeated twice until all disagreements were resolved.

4.5. Quantitative data analysis

We performed logistic and ordinal logistic regressions as well as Mann-Whitney U and Wilcoxon signed-rank tests.

We kept our predictors consistent across all regressions, except we added the FAQ as a predictor for Survey 2. We choose predictors that allow us to explore the relationship between our explanations and the outcome variables (e.g., whether they answered a True/False question correctly). Our predictors are the *explanation* shown, *computer science experience*, *medical or IoT experience*, and *FAQ* condition. To tailor our analyses to the scenarios, we use medical experience as a predictor for medical scenarios and smart home experience for smart home scenarios. We selected our models and planned our analyses in advance to limit Type I error rates associated with running multiple tests.

The explanation in Survey 1 consisted of all possible variations and combinations of explanations between the first three sentences (a total of 12 different explanations). The baseline for these predictors is *Unsubstantial*, *Technical*, and *No Prevents* for each of the explanation sentences. In Survey 2, we had four possible explanations shown, and the baseline is *None* (no explanation). The baseline for medical and smart home experience is False (no experience).

Comprehension questions. We assess user understanding based on a set of True/False questions (10 in Survey 1 and 12 in Survey 2). To analyze these binary outcomes, we performed logistic regressions for each comprehension question in each scenario (e.g., 10 questions $\times$ 2 scenarios $=$ 20 models for Survey 1). Because we used different models for medical and home IoT scenarios, each participant was in the data set exactly one time, so we do not need to account for repeated measures in our model. The coefficients represent the log odds of the outcome occurring for a one-unit increase in the predictor. A coefficient greater than zero indicates that the predictor increases the log odds of the outcome variable to 1 (a correct answer). Conversely, a coefficient less than zero would indicate a negative impact. Additionally, we assessed the significance of each predictor by looking at the p-value with a significance level of 0.05.

Safety and willingness questions. We also asked participants about their perceptions of safety and willingness to engage with our scenario. These questions did not have binary answers. Instead, participants answered using a 3-to-5-point Likert scale. To ensure consistency across models, we binned all Likert scale data used in statistical analysis into 3 levels. For the willingness to engage with our scenario, there was no need to re-bin, we had "Would not" (baseline), "Maybe would," and, "Definitely would." For the safety perceptions, we binned all answers into "Not at all safe" (baseline), "Somewhat safe," and "Safe" (from binning "Mostly safe" with "Completely safe.").

To analyze this data, we used an ordinal logistic regression. We conducted one regression per question, per scenario, using the same predictors as the comprehension questions. Because we used different models for medical and home IoT scenarios, each participant was in the data set exactly one time, so we do not need to account for repeated measures in our models. We conducted Brant tests to ensure the proportional odds assumption for all predictors. The results indicated that the proportional odds assumption holds for all predictors ($p = 0.05$). Interpreting ordinal logistic regressions is similar to binary logistic regressions. The difference is that a coefficient greater than zero indicates that the predictor increases the log odds of the outcome variable reaching or exceeding a higher category when compared to the baseline—for example, a positive coefficient for our willingness questions would indicate the participant is more willing to use the technology, while a negative coefficient indicates they are less willing. We keep the same significance level of 0.05.

We also wanted to understand the aspects of the scenario that affect perceptions of safety with vs. without information about TEEs. For this, we used the Mann-Whitney U test, a nonparametric test that allows for the comparison of median ranks between two independent groups, even with non-normal data. To compare the average score for questions **Q1**-**Q5** to the average score for questions **Q6**-**Q10** we use a Wilcoxon signed rank, a nonparametric test to compare the median of the differences between two groups. We examine each scenario separately and use a significance level of 0.05.

4.6. Limitations

We chose an IoT and Medical scenario to reflect real-life situations where our participants could make choices about the use of a TEE-enhanced technology. However, these scenarios, while designed to be realistic, might not fully capture the complexity of a real-world context and may not be representative of the entire range of contexts where users may need some understanding of TEEs. The study relies on self-reported data, which may be affected by social desirability bias or participants' willingness to disclose their true thoughts and feelings. We tried to mitigate this limitation by ensuring the confidentiality of the participants. We also checked to make sure participants had read and understood the scenarios we were asking about. Moreover, within the online crowdsourcing platforms available, Prolific

seems to be one of the most reliable [57], [58]. Our sample of participants is skewed young and may not represent the larger population.

4.7. Ethical Considerations

The surveys and consent forms were approved by the IRB at the authors' institution(s). The only personally identifiable data collected were Prolific IDs for recruiting and paying participants and IP addresses for bot detection.

5. Survey 1 Results

In this section, we summarize the results of Survey 1 evaluating candidate TEE explanations.

5.1. RQ1: Factors Influencing Comprehension

To measure our participants' comprehension of TEE concepts, we asked them a series of True/False questions. Each question has one correct answer. We evaluated their responses for correctness and summarize the scores in Figure 2. The full regression table can be found in Table 7, and scores per scenario and TEE explanation can be found in the supplementary materials [47].

Participants are less likely to recognize TEE limitations. Looking at the overall trend in comprehension scores across all conditions, shown in Figure 2, participants are more likely to correctly answer questions about the features of TEEs (**Q1** through **Q5**, 86.5% correct overall) than about the limitations (**Q6** through **Q10**, 71.5% overall). We also tested this hypothesis with a Wilcoxon signed-rank test comparing the average score for questions about limitations, to the average score for questions about features in Survey 2. We found a significant difference between these two groups for the two smart home scenarios and for the medical scenario without AI ($p - values$ for these three scenarios ranged between $4.378e - 09$ and 0.0044, see Table 6).

TEE *Prevents* and *Non-technical* explanations can improve comprehension. As discussed in Section 4.5, we ran 40 regression models to predict the relationship between each explanation factor and experience on each of the four scenarios and 10 comprehension questions. We found that some explanations are better at describing TEE concepts than others. The *Non-technical* explanation is especially good at explaining who is (or is not) allowed to access the data in the TEE (e.g., **Q2** and **Q5**) while the *Prevents* explanation is best at explaining that the TEE protects against malicious software on the rest of the computer (e.g., **Q3** and **Q4**). The effects for the *Non-technical* explanation hold across all scenarios for **Q2** (significant for the medical scenario without AI) and all medical scenarios for **Q5**. Similarly, the effects for the *Prevents* explanation hold across three of the four scenarios for **Q3** (significant for the medical scenario without AI) and across all scenarios for **Q4** (significant for both medical scenarios and the smart home scenario with AI).

5.2. RQ2: Factors Influencing Willingness and Feeling of Safety

While the *type of technology* we described in the scenario seems to have an effect on participant willingness to use TEE-enhanced technology and belief that TEE-enhanced technology will keep their data safe, this does not seem to be the case for the *TEE explanation*. Most explanation predictors are not significant in our regression models, except for the *Non-Technical* explanation, which seems to make participants feel significantly safer in the smart home device scenario without AI (see Table 3).

TEE explanations seem to have little effect on willingness and feeling of safety. Regardless of TEE explanation, our results (shown in Table 3) suggest that participants were nearly equally willing to engage with the TEE-enabled technology in our scenarios. 20%-22.4% said they were "definitely willing" and 50.4%-55.7% were "maybe willing" across all scenarios. Similarly, participants seem to believe their data would be nearly equally safe regardless of how the TEE was explained. 24.0%-28.3% said it would be "completely safe" and 62.3-66.9% said it would be "somewhat" safe across all scenarios.

5.3. Questions from our participants

The survey had two opportunities for participants to ask us questions about TEEs. We received 310 questions from 252 participants. In this section, we describe the most common types of questions asked. Note that participants may have written multiple questions in a single response, with each question potentially having one or more theme. Our codebook describing themes and the frequency of each theme can be found in the supplementary materials [47].

Our participants had many questions that our TEE explanations did not answer. The most common questions were about TEEs, but there were many other questions about the scenario where the TEE is used, potential risks they might encounter, and what guarantees there are that the TEE will function as described. We also received some comments indicating that participants lack trust in TEEs and data privacy, in general. Participants asked more questions after the medical scenarios, and these questions were more likely to focus on TEEs and potential risks than in the smart home scenarios.

Attributing quotes to participants. When attributing quotes, we report the participant ID, which survey the quote is from, and the treatment they were assigned. For Survey 1, each participant receives three letters, corresponding to the three sentences in the TEE explanation they received: (H)ardware, (T)rust, or (U)nsubstatial; (T)echnical or (N)on-technical; and (P)revents, or (N)o Prevents. For example, (P30S1-HTN) means participant #30 in Survey 1, who received the *Hardware*, *Technical*, *No Prevents* explanation.

Questions about TEEs. The most common questions we received were about TEEs (143 responses). 49 of these questions asked for more technical details generally: "How

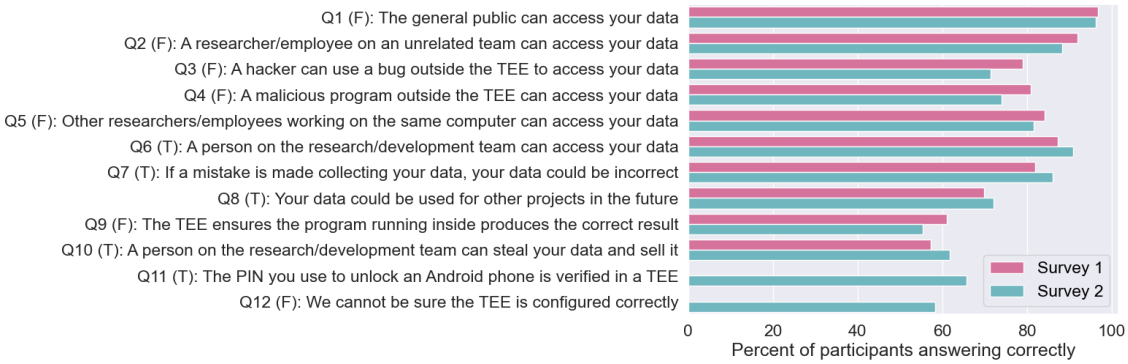

Figure 2: Overall scores for the True/False comprehension questions in both surveys. **Q1**-**Q5** are about features of TEEs, **Q6**-**Q10** are about limitations of TEEs, and **Q11** and **Q12** are questions that can be answered based on information in the FAQ and appear in Survey 2 only. The correct answer for each question is shown in parentheses by the question number.

exactly does a TEE work?" (P55S1-TTP). 21 participants wanted more information about how the TEE creates an isolated environment: "How exactly is a TEE isolated from the rest of the computer?" (P146S1-HNP). 15 participants wanted more implementation details: "Is there a second set of RAM with an independent CPU or something?" (P379S1-HTP). In addition, 15 participants wanted more information about what else the machine is capable of: "...I presume that means that the [researchers] cannot use any other programs at the same time?" (P48S1-UTN). It is noteworthy that only 13 participants asked questions that should have already been answered by the scenario text or the TEE explanation they received. Since this represents relatively few of the responses we received (less than 5%), it suggests that most participants were paying attention to our survey.

Questions about the scenarios. We received 91 questions about our scenarios. 41 asked about the data involved, including what data is collected, data retention policies, and how/whether the data is anonymized: "What happens to my data when I no longer wish for it to be stored" (P149S1-TNN). 25 questions were about the people on the research/development team: "How many people have access to the TEE, what are their qualifications...." (P12S1-HTN).

Questions about risks. 74 participants asked about potential risks. The most common risk, mentioned by 27 participants, was hackers: "I get that the program is safe from other possibly malicious programs, but what about hackers" (P47S1-TTP). 23 participants were concerned about people behaving maliciously, including the people in the scenario with legitimate access to their data: "What kind of process ensures that the researchers will not share my data?" (P181S1-HTN).

Questions about guarantees or real-world uses. 51 participants wanted to know how they could be sure the TEE would work: "...how [is] it guaranteed that it can't be accessed?" (P87S1-TNN). 21 had questions about real TEEs, including 11 who asked whether they had ever been involved in a breach: "I would like to know if there have been cases in the past where TEEs have been hacked" (P204S1-HTP). 4 wondered whether TEEs are actually real: "Is it a real thing? Or a hypothetical idea just for the study?" (P127S1-TTP).

Other concerns. 43 participants did not ask questions, instead using the space to share opinions. 10 commented on the scenario "...I was biased about this to begin with. I don't trust these devices" (P62S1-TNP). 16 wrote about technology: "You do understand that people don't trust technology?" (P389S1-UTN). 10 people mentioned that they don't trust TEEs: "I don't trust my information will be secure, especially with the words 'trusted environment' " (P243S1-HNN).

6. Survey 2 Results

Our first survey shows that TEE explanations might be effective at communicating TEE concepts, especially when they are *Non-technical*, mention specific threats a TEE can *Prevent*, and do not expect people to infer new things. However, our explanations had little effect on willingness to use TEE-enhanced technology or the belief that the TEE will keep data safe. Moreover, the questions we collected suggest that a potential reason that TEE explanations have little effect may be that our participants' primary data privacy concerns are beyond the capabilities of TEEs. In this section, we describe our follow-up survey, Survey 2, where we introduce an FAQ based on the questions asked by participants in our first survey and additional questions related to their belief that the TEE-enhanced technology will keep their data safe.

6.1. RQ3: Effect of FAQ

We summarize the scores for each question in Figure 2. The full regression table can be found in Table 8. Having an FAQ seems to have helped participants answer questions about TEE features correctly, but it also seems to have made them less likely to answer questions about TEE limitations correctly. The FAQ had little effect on participants' willingness to adopt TEE-enhanced technology. Still, it did tend to

make people feel more confident that their data would be safe when protected by a TEE.

Participants did interact with the FAQ. To determine whether people were reading the FAQ, we added two comprehension questions. **Q11** asked about real-world use of TEEs ("How are TEEs used in real life?" in the FAQ) and **Q12** asks whether we can know that a TEE is configured correctly ("How do we know the TEE is working correctly?"). In both cases, participants were significantly more likely to answer the question correctly if they had an FAQ than if they didn't (Table 8). 77% of participants who received a *Hidden* FAQ expanded the questions at least once.

The FAQ has a mixed effect on comprehension. Participants with an FAQ were more likely to correctly answer questions about TEE features (**Q1**-**Q5**) or the FAQ-specific questions (**Q11**-**Q12**). The difference between the *Shown* FAQ and *None* FAQ condition was statistically significant for **Q2**, **Q3**, and **Q5**, for different scenarios. On the other hand, the *Hidden* FAQ condition was significantly better than the *None* FAQ condition for **Q3** in only the medical scenario without AI. Meanwhile, both *Shown* and *Hidden* FAQ conditions were better than the *None* FAQ for **Q11** and **Q12** (statistically significant for all scenarios). When the FAQ helped participants answer the question correctly, we found similar results for both types of FAQ presentations, except for the FAQ-specific questions, where the *Shown* FAQ was better than the *Hidden* one.

Interestingly, having an FAQ made it more likely that participants would answer questions about TEE limitations (**Q6**-**Q10**) *incorrectly* than if they didn't have an FAQ at all. The difference between the *Shown* FAQ condition and *None* FAQ condition was statistically significant for **Q8** and **Q10** for a few (but not all) scenarios. The *Hidden* FAQ condition was significantly worse than the *None* FAQ condition for **Q6**, **Q8**, and **Q10**, also for a few (but not all) scenarios.

Having an FAQ or explanation seems to have little effect on willingness to use technology. Similar to the findings in Survey 1, where we saw almost no impact from different explanations, in Survey 2, we see that explanations and FAQ do not seem to significantly affect participants' willingness to use TEE-enabled technology (see Table 3).

Having an FAQ or explanation seem to make people more confident their data will be safe. More participants receiving an FAQ believed their data would be completely or mostly safe (75% for *Shown* and 71.4% for *Hidden*) than those who did not receive an FAQ (64.1%). We observed a similar trend between the participants receiving a TEE explanation and those not receiving one (70-74% depending on the explanation vs. 63.2% for no explanation). These differences are significant only for the *Shown* FAQ condition in the smart home scenario without AI (see Table 3).

6.2. RQ4: Aspects Contributing to Safety

For each scenario, we asked participants which aspects of the scenario contribute to their belief that their data would be safe (or unsafe), including: the use of a TEE, that a hospital (or company, depending on the scenario) is collecting the data, the people on the team, what data is collected, and the purpose of the data collection. We also asked if any other aspects of the scenario not already mentioned contributed to their belief that their data would be safe or unsafe. Overall, we found that providing information about TEEs (by giving them an explanation *or* an FAQ) seems to make people more confident that the TEE would keep their data safe. There were many other aspects of the scenario that people were concerned about that TEE explanations and FAQs did not address. The results in this section are supported by a Wilcoxon signed-rank test, see Table 5 for details.

Attributing quotes to participants. We attribute quotes using a similar strategy as Survey 1, except that, here, the treatment is represented using two letters. The first letter is the TEE explanation: (H)ardware, (T)rust, (U)nsubstantial, or (X) for no explanation. The second letter is the FAQ condition: (H)idden, (S)hown, or (X) for no FAQ. For example, (P68S2-HX) is participant #68 in Survey 2, who received the *Hardware* explanation and the *None* FAQ condition.

Explaining TEEs seems to make people more confident the TEE will keep their data safe. In the group that had access to information about TEEs, 80.3% said the use of a TEE made them feel their data was definitely or somewhat safe, while only 52.3% of those who had no access to information about TEEs said the use of a TEE made them feel their data was definitely or somewhat safe ($p - values$ between 0.0002 and 0.0026 depending on the scenario).

TEE information seems to have little effect on other aspects. For aspects other than the use of a TEE, providing a TEE explanation or FAQ seems to make little difference to our participants' feelings of safety ($p > 0.05$ except for the people involved in the medical scenario without AI with $p = 0.0413$). For example, 61.6% of participants reported feeling definitely or somewhat safe about the purpose of the data collection when they had information about the TEE vs. 57.8% without information. For the other aspects, providing information about the TEE made people somewhat *less* sure their data would be safe. The place where information made the biggest difference was when we asked about the people involved in the scenario. Here, 50.8% of participants with information about the TEE reported the people made them feel their data would be definitely or somewhat safe, while 54.4% without information about the TEE said the same.

Other aspects of the scenario mentioned by participants. Here, we describe some of the most common aspects, not already discussed above (many participants used this opportunity to expand on their previous answers). We received 660 responses total from 382 participants, where each response might mention one or more aspects of the scenario. 392 responses mentioned at least one aspect contributing to the feeling their data would be unsafe and 249 responses mentioned aspects contributing to the feeling their data would be safe.

Some aspects of the scenario participants reported contributing to the feeling that data would be unsafe in-

	Survey 1 Medical Scenario Without AI		Survey 1 Smart Home Scenario Without AI		Survey 1 Medical Scenario With AI		Survey 1 Smart Home Scenario With AI	
Variable	**Willingness**	**Safety**	**Willingness**	**Safety**	**Willingness**	**Safety**	**Willingness**	**Safety**
Expln sentence 1 [Baseline = Unsubstantial]								
Hardware	0.33	−0.03	0.32	−0.05	−0.01	0.02	0.11	0.07
Trust	0.15	0.14	0.36	0.08	−0.33	−0.30	−0.12	−0.18
Expln sentence 2 [Baseline = Technical]								
Non-Technical	−0.25	−0.09	0.34	0.87**	0.18	0.13	−0.07	−0.25
Expln sentence 3 [Baseline = No Prevents]								
Prevents	0.27	0.23	0.30	0.03	−0.17	−0.18	−0.30	0.41
Medical/Smart home exp	0.43	−0.15	1.18***	−0.18	0.14	−0.23	0.85**	0.59
CS Experience	0.01	0.02	0.11	0.75*	0.31	0.25	0.13	0.12

	Survey 2 Medical Scenario Without AI		Survey 2 Smart Home Scenario Without AI		Survey 2 Medical Scenario With AI		Survey 2 Smart Home Scenario With AI	
Variable	**Willingness**	**Safety**	**Willingness**	**Safety**	**Willingness**	**Safety**	**Willingness**	**Safety**
TEE Explanation [Baseline = None]								
Unsubstantial	−0.31	0.63	−0.21	−0.36	−0.01	0.08	0.45	0.95*
Hardware	0.11	0.31	0.46	0.65	0.19	0.68	−0.01	0.31
Trust	0.02	0.72	−0.14	0.08	−0.37	0.44	0.04	0.45
FAQ [Baseline = None]								
Hidden	−0.09	0.28	0.53	0.52	−0.28	0.18	−0.28	0.45
Shown	−0.01	0.56	0.32	0.87*	0.42	0.42	0.24	0.54
Medical/Smart home exp	−0.05	−0.22	1.83***	1.16**	−0.01	−0.01	1.24***	0.52
CS Experience	0.08	0.16	0.13	0.44	0.21	0.14	−0.16	−0.36

TABLE 3: Regression table for questions about willingness to use technology and belief that the TEE will keep data safe in **Survey 1 and 2** respectively. The first column is how willing the participant would be to use the TEE-enhanced technology, the second column is the belief that their data will be safe. There is one ordinal logistic regression model for each question in each scenario (24 models total). The numbers in this table are the log-odds coefficients for each predictor, with the baseline explanations used in each model noted in *italics*. Statistical significance is noted with asterisks and shaded cells: blue for positive coefficients and orange for negative.

clude: prior experiences with/knowledge of breaches (42 responses), the future use clause in the scenario text (28 responses), the belief that some of the data was being collected unnecessarily (24 responses), the use of AI in the scenario (19 responses), the risk that there could be a bug in the TEE code (14 responses), and the risk that their data would be sold (12 responses). 5 people were concerned that we mentioned future research: "The fact that it says researchers are looking for new ways to verify the program is working correctly. That makes me a little hesitant. Sounds like there are still bugs…" (P144S2-TS). 3 people seemed suspicious about being told to "trust" the technology: "Comes across a bit like: 'Yeah trust me bro your medical records are totally safe bro, trust me, bro there's an acronym. You like acronyms right man?' " (P237S2-TX).

Most aspects participants reported contributing to the feeling that their data would be safe were repeated from previous questions. The most common new aspect contributing to the feeling of safety is the perception that the data being collected is not interesting enough to an attacker anyway (25 responses).

6.3. More questions from our participants

Similar to the first survey, we gave participants two opportunities to ask us questions they have about TEEs and received 267 responses. In this section, we summarize the most common questions we received, following the same structure as Section 5.3, and how questions differed between participants who did and did not have access to an FAQ. We began with the same codebook as in the first survey, with only a few additional codes emerging during the analysis. The codebook describing all of the themes and how frequently they occurred may be found in the supplementary materials [47].

Giving participants an FAQ made them less likely to ask questions about TEEs or the scenario, which were the most common kinds of questions we received overall. Some other questions were more common from participants who received an FAQ, like asking for more examples or about guarantees.

An FAQ seems to reduce questions about TEEs. As in the first survey, we received the most questions (117 responses) about TEEs themselves. Although only 34% of people did not receive an FAQ, 45% of the questions about TEEs came from people who did not receive an FAQ. The most common questions were, again, asking for more information about how the TEE (31 responses) or its isolation mechanism (7 responses) work. Unlike in the first survey, we also saw 17 people asking what a TEE is, more generally: "What is a TEE??" (P351S2-XX). Other common questions requested more implementation details (14 responses) or compared TEEs to other technologies (8 responses). We also had 13 requests for a less technical explanation: "Need more details about how they work in general without the use of complicated verbiage." (P52S2-XH). Most of these (69% of the requests) came from people who were forced to wait on the FAQ page.

An FAQ seems to reduce questions about people in the scenario. We also received 74 questions about the scenario. Again, 45% of the questions about the scenario came from people who did not receive an FAQ. The most common questions were about the people involved in the scenario (28 responses) or the data (25 responses). A disproportionate 57% of questions about people come from the participants who did not receive an FAQ, while the questions about data are more evenly distributed between FAQ conditions.

A hidden FAQ seems to reduce questions about hackers. 60 participants had questions about the risks they might encounter. Similar to above, 43% of these questions came from people who did not receive an FAQ. Unlike other questions, though, participants were least likely to ask questions about hackers (28 responses, total) if they got the expandable FAQ: 18% of questions about hackers came from the hidden FAQ condition, while 43% came from the shown FAQ and 39% from the no FAQ condition. Questions about people behaving maliciously (12 responses) were nearly evenly distributed between FAQ conditions. The remaining questions about risks disproportionately came from the people who did not receive an FAQ (54% of the remaining questions asked about risks).

Questions about guarantees or real-world uses seem to be more common with an FAQ. We received 20 questions about guarantees and 29 questions about real-world uses for TEEs. Both questions were more common with an FAQ than without. 45% of questions about guarantees came from participants in the hidden FAQ condition and 45% of questions about real uses of TEEs came from participants who were shown the FAQ. It is possible that some questions came from participants who wanted to write something but couldn't think of anything else to ask: "I can't think of any more questions. Maybe, would be nice to see more real world examples" (P22S2-XS).

A hidden FAQ also seems to reduce other concerns. Similar to our initial survey, 41 participants did not ask a question but used the space to share other thoughts. The most common thoughts were general distrust (23 responses), followed by opinions about the scenario (9 responses). The participants receiving the hidden FAQ condition seemed to be the least likely to use this space to express distrust (these account for 17% of the 23 responses), while the remaining questions were nearly evenly distributed between the shown FAQ and no FAQ conditions (43% and 39%, respectively).

7. Discussion

In this section, we make recommendations for explaining technical concepts to non-experts, navigate the (seemingly) contradictory results between our study and prior work [9], and highlight opportunities for future research. Finally, we revisit the hypothesis from prior work that understanding TEEs would make users more comfortable sharing data with TEE-enhanced technology.

7.1. Explaining technical concepts to non-experts

Avoid technical jargon. Our results in Section 5.1 echo prior work [28], [29], [34] on the importance of avoiding jargon when explaining technical concepts to non-experts. This was also mentioned by participants reading the supplementary technical details we introduced with the FAQ: "Need more details about how they work in general without the use of complicated verbiage." (P52S2-XH).

Be direct and tell users what you want them to know. In Section 5.1, we found that people were more likely to answer questions correctly when the answer was in the explanation or scenario text directly than questions where people had to generalize what they learned and *infer* the answers. For example, participants in the *Prevents* TEE explanation condition were told that the TEE can protect against malicious software on the computer. This group was significantly more likely to answer the question about malicious programs (**Q4**) correctly in three of the four scenarios because the explanation they received gave them the answer to the question. On the other hand, while 87.2% of the participants in Survey 1 knew that the research/development team could access the data (**Q6**), only 57.2% used that knowledge in **Q10** to infer that the same group of people could *steal* their data and use it for personal gain. It is possible that our participants had trouble inferring that even people authorized to access their data might use it for malicious purposes.

Don't tell people what technology they should trust. In Section 6.2, we showed that explaining TEEs does little to address some of the concerns our participants have about technology. In fact, in some cases, information about TEEs made people slightly *more* skeptical. One reason for this might be that TEEs do not address all security and privacy threats, so explaining them, even if they are explained well, does not address all of the concerns people have. These concerns could also explain why some participants were wary of the word "trust" in "Trusted Execution Environment" as we noted in Sections 5.3 and 6.2. Because security technologies are often orthogonal to the concerns people shared with us, it could be counterproductive for users already feeling

skeptical if these solutions are marketed to them as trusted: "I don't trust my information will be secure, especially with the words 'trusted environment' " (P243S1-HNN).

7.2. Comparing our findings to prior work

One of the main motivations of our study was the finding from prior work [9] that the presence of cloud-based TEEs can make people more comfortable sharing their data with home IoT, especially if they understand what a TEE is. On the surface, our finding that the TEE explanation has little effect on participants' willingness to use TEE-enhanced technology or their perception of safety (Section 5.2 and Table 3) seem to contradict these results. One explanation for this difference could simply be the methodological differences between the previous study and ours.

The previous study's main goal was to understand whether the presence of a TEE alters existing privacy norms within a smart home environment. In particular, they asked participants to imagine that they own a smart home device (either a smart camera or smart speaker) and how comfortable they feel about having their (or other occupants) data collected under certain conditions (e.g., if the data is shared with law enforcement, if they—the device owner—are notified). In this study, the assumption is that *the participants already own and interact with a smart device* and the goal is to understand how the introduction of a TEE affects participants' perceptions of what makes them comfortable (or uncomfortable) with certain data sharing practices. Our study, on the other hand, does not ask participants to *assume that they will interact* with the TEE-enhanced technology. Instead, we evaluate their comfort by asking them about their *willingness to interact with the technology* and their perception of safety.

7.3. Future research

Explaining limitations, not just features. Our explanations were better at describing the protections provided by TEEs than their limitations. In fact, despite ensuring that our explanations faithfully represent TEE security features, participants believed the TEE would offer some protections that it does not, such as guaranteeing the results of a computation are correct or preventing people with legitimate access from selling their data (**Q9** and **Q10** and in Sections 5.1 and 6.1). A similar phenomenon was observed in prior work [33].

More research is necessary to understand how we can highlight the limitations of security technology. However, as in our study, this research needs to measure comprehension, willingness to use, and beliefs about safety.

Investigating showing vs. hiding the FAQ. In Section 5.3, we explained that many participants in Survey 1 asked for more technical details about TEEs. In Survey 2, we provided participants with those details in an FAQ, and in Section 6.3, we explained that it did lead to fewer questions. However, we also saw that the technical details could be overwhelming to some and that the *Shown* FAQ was more effective for some, *but not all*, of the True/False comprehension questions (Section 6.1). More research is needed to understand why different FAQ models perform differently for some comprehension questions. One hypothesis is that hiding the FAQ allows people to focus their attention on the relevant information, but also makes it more likely that they won't read it at all. Future work could also shed light on how we might balance the trade-offs between providing additional technical details to those who want them and hiding them from those who find them unnecessary.

Revisit prior work with our enhanced explanations. As explained above, the methodological differences between our study and the one from prior work [9] likely explain the seemingly contradictory results about user comfort. Nevertheless, it would be useful to repeat their study using our most effective explanations to see if their results can be reproduced. Repeating the study using their methodology would make it possible to compare results more directly and better understand why even our best explanations seem to have little effect on user comfort.

7.4. How much do users need to know about TEEs?

Our results provide some insights into how we can communicate about technical security concepts more effectively, but suggest that understanding TEEs does not impact decisions about whether or not to use TEE-enhanced technology. While we started off with the hypothesis that understanding TEEs would improve users' trust in technology, our participants' responses drive home the point that, as some of our participants correctly realized, knowledge that a system uses a TEE is insufficient to draw conclusions that a user's data will be adequately protected. We might imagine that the TEE is just one of several components that are being used to protect user data in our scenarios and we could potentially provide a much more detailed explanation of all the protective components to assure users that their data is safe, or to highlight exactly what risks they might face. But this begs the question of whether we should really expect users to understand the inner workings of a security system, or if it should simply be offered to improve transparency around data privacy.

Ultimately, decisions about TEEs are still best left to experts, not end users. Experts' choices about whether and how to use TEEs should revolve around the technology they are developing and the data they require, not whether the TEE would make users more willing to use the technology.

8. Conclusion

In this study, we evaluated strategies for explaining TEEs. Some were more effective at enhancing understanding than others. Our findings highlight the importance of avoiding technical jargon and directly communicating what people should learn. On the other hand, we found that our explanations have limited effects on willingness to use technology or the feeling of safety, likely because TEEs do

not address many of the privacy concerns our participants have. Our results provide insights into how we can communicate more effectively about technical security concepts, but also suggest that explaining security technology might not resolve the concerns users have about data privacy.

References

[1] C. McClain, M. Faverio, M. Anderson, and E. Park, "How Americans view data privacy," *Pew Research Center*, 2023.

[2] Confidential Computing Consortium, "Confidential computing: Hardware-based trusted execution for applications and data," https://confidentialcomputing.io/wp-content/uploads/sites/10/2023/03/CCC_outreach_whitepaper_updated_November_2022.pdf, 2022, [Accessed 10-09-2024].

[3] M. Sabt, M. Achemlal, and A. Bouabdallah, "Trusted execution environment: What it is, and what it is not," in *Proceedings of the 2015 IEEE Trustcom/BigDataSE/Ispa*, 2015.

[4] K. Vaswani, S. Volos, C. Fournet, A. N. Diaz, K. Gordon, B. Vembu, S. Webster, D. Chisnall, S. Kulkarni, G. Cunningham *et al.*, "Confidential computing within an AI accelerator," in *Proceedings of the 2023 USENIX Annual Technical Conference*, 2023.

[5] Y. Chen, F. Luo, T. Li, T. Xiang, Z. Liu, and J. Li, "A training-integrity privacy-preserving federated learning scheme with trusted execution environment," *Information Sciences*, vol. 522, 2020.

[6] F. Mo, Z. Tarkhani, and H. Haddadi, "Machine learning with confidential computing: A systematization of knowledge," *ACM computing surveys*, vol. 56, no. 11, 2024.

[7] G. Ayoade, V. Karande, L. Khan, and K. Hamlen, "Decentralized IoT data management using blockchain and trusted execution environment," in *Proceedings of the 2018 IEEE International Conference on Information Reuse and Integration*, 2018.

[8] Y. Wang, J. Li, S. Zhao, and F. Yu, "Hybridchain: A novel architecture for confidentiality-preserving and performant permissioned blockchain using trusted execution environment," *IEEE access*, vol. 8, 2020.

[9] P. Musale and A. J. Lee, "Trust TEE?: Exploring the impact of trusted execution environments on smart home privacy norms," *Proceedings on Privacy Enhancing Technologies*, vol. 3, 2023.

[10] T. Geppert, S. Deml, D. Sturzenegger, and N. Ebert, "Trusted execution environments: Applications and organizational challenges," *Frontiers in Computer Science*, vol. 4, p. 930741, 2022.

[11] Android Open Source Project, "Android Authentication," https://source.android.com/docs/security/features/authentication, 2024, [Accessed 27-06-2024].

[12] Android, "TrustZone," https://source.android.com/docs/security/features/trusty, 2024, [Accessed 27-06-2024].

[13] Intel, "Intel SGX," https://www.intel.com/content/dam/develop/external/us/en/documents/overview-of-intel-sgx-enclave-637284.pdf, 2024, [Accessed 27-06-2024].

[14] AMD, "Strengthening VM isolation with integrity protection and more," *White Paper*, vol. 53, pp. 1450–1465, 2020.

[15] Intel, "Intel TDX," https://www.intel.com/content/www/us/en/developer/tools/trust-domain-extensions/overview.html, 2024, [Accessed 27-06-2024].

[16] ARM, "ARM CCA," https://www.arm.com/architecture/security-features/arm-confidential-compute-architecture, 2024, [Accessed 27-06-2024].

[17] L. O. Gostin, L. A. Levit, and S. J. Nass, "Beyond the HIPAA privacy rule: enhancing privacy, improving health through research," 2009.

[18] "Health Insurance Portability and Accountability Act of 1996," 45 C.F.R. §§ 160, 162, 164, 1996, united States.

[19] S. Ambiel, "The case for confidential computing," 2024, [Accessed 11-11-2024].

[20] A. Whitten and J. D. Tygar, "Why Johnny can't encrypt: A usability evaluation of PGP 5.0." in *Proceedings of the 8th USENIX Security Symposium*, 1999.

[21] R. Abu-Salma and B. Livshits, "Evaluating the end-user experience of private browsing mode," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020.

[22] H. Habib, J. Colnago, V. Gopalakrishnan, S. Pearman, J. Thomas, A. Acquisti, N. Christin, and L. F. Cranor, "Away from prying eyes: Analyzing usage and understanding of private browsing," in *Proceedings of the 14th Symposium on Usable Privacy and Security*, 2018.

[23] C. Bravo-Lillo, L. F. Cranor, J. Downs, and S. Komanduri, "Bridging the gap in computer security warnings: A mental model approach," *IEEE Security & Privacy*, vol. 9, no. 2, 2010.

[24] B. Ur, F. Noma, J. Bees, S. M. Segreti, R. Shay, L. Bauer, N. Christin, and L. F. Cranor, ""I Added '!' at the end to make it secure"': Observing password creation in the lab," in *Proceedings of the 11th Symposium on Usable Privacy and Security*, 2015.

[25] B. Ur, J. Bees, S. M. Segreti, L. Bauer, N. Christin, and L. F. Cranor, "Do users' perceptions of password security match reality?" in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 2016.

[26] S. Pearman, S. A. Zhang, L. Bauer, N. Christin, and L. F. Cranor, "Why people (don't) use password managers effectively," in *Proceedings of the 15th Symposium on Usable Privacy and Security*, 2019.

[27] E. R. Bouma-Sims, M. Li, Y. Lin, A. Sakura-Lemessy, A. Nisenoff, E. Young, E. Birrell, L. F. Cranor, and H. Habib, "A US-UK usability evaluation of consent management platform cookie consent interface design on desktop and mobile," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023.

[28] T. Wu, R. Zhang, W. Ma, S. Wen, X. Xia, C. Paris, S. Nepal, and Y. Xiang, "What risk? I don't understand. an empirical study on users' understanding of the terms used in security texts," in *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security*, 2020.

[29] S. Zhang, Y. Feng, Y. Yao, L. F. Cranor, and N. Sadeh, "How usable are iOS app privacy labels?" *Proceedings on Privacy Enhancing Technologies*, 2022.

[30] S. Furnell, R. Esmael, W. Yang, N. Li *et al.*, "Enhancing security behaviour by supporting the user," *Computers & Security*, vol. 75, 2018.

[31] European Union, "Regulation 2016/679 general data protection regulation," https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02016R0679-20160504, 2016, accessed: 2024-07-03.

[32] B. Shen, L. Wei, C. Xiang, Y. Wu, M. Shen, Y. Zhou, and X. Jin, "Can systems explain permissions better? Understanding users' misperceptions under smartphone runtime permission model," in *Proceedings of the 30th USENIX Security Symposium*, 2021.

[33] O. Akgul, W. Bai, S. Das, and M. L. Mazurek, "Evaluating In-Workflow messages for improving mental models of End-to-End encryption," in *Proceedings of the 30th USENIX Security Symposium*, 2021.

[34] V. Distler, C. Lallemand, and V. Koenig, "Making encryption feel secure: Investigating how descriptions of encryption impact perceived security," in *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2020, pp. 220–229.

[35] C. Stransky, D. Wermke, J. Schrader, N. Huaman, Y. Acar, A. L. Fehlhaber, M. Wei, B. Ur, and S. Fahl, "On the limited impact of visualizing encryption: Perceptions of E2E messaging security," in *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*, 2021, pp. 437–454.

[36] C. Bravo-Lillo, S. Komanduri, L. F. Cranor, R. W. Reeder, M. Sleeper, J. Downs, and S. Schechter, "Your attention please: Designing security-decision UIs to make genuine risks harder to ignore," in *Proceedings of the 9th Symposium on Usable Privacy and Security*, 2013.

[37] C. Carreira, J. F. Ferreira, A. Mendes, and N. Christin, "Exploring usable security to improve the impact of formal verification: A research agenda," *Electronic Proceedings in Theoretical Computer Science*, vol. 349, Nov. 2021.

[38] A. Xiong, T. Wang, N. Li, and S. Jha, "Towards effective differential privacy communication for users' data sharing decision and comprehension," in *Proceedings of the 2020 IEEE Symposium on Security and Privacy*, 2020.

[39] F. Karegar, A. S. Alaqra, and S. Fischer-Hübner, "Exploring user-suitable metaphors for differentially private data analyses," in *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, 2022, pp. 175–193.

[40] R. Cummings, G. Kaptchuk, and E. M. Redmiles, ""I need a better description": An investigation into user expectations for differential privacy," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021.

[41] S. E. Schechter, R. Dhamija, A. Ozment, and I. Fischer, "The emperor's new security indicators," in *2007 IEEE Symposium on Security and Privacy*. IEEE, 2007, pp. 51–65.

[42] J. Wu, C. Gattrell, D. Howard, J. Tyler, E. Vaziripour, D. Zappala, and K. Seamons, ""Something isn't secure, but I'm not sure how that translates into a problem": Promoting autonomy by designing for understanding in Signal," in *Fifteenth Symposium on Usable Privacy and Security*, 2019, pp. 137–153.

[43] B. Ur, F. Alfieri, M. Aung, L. Bauer, N. Christin, J. Colnago, L. F. Cranor, H. Dixon, P. Emami Naeini, H. Habib *et al.*, "Design and evaluation of a data-driven password meter," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017.

[44] P. Emami-Naeini, Y. Agarwal, L. F. Cranor, and H. Hibshi, "Ask the experts: What should be on an IoT privacy and security label?" in *Proceedings of the 2020 IEEE Symposium on Security and Privacy*. IEEE, 2020.

[45] P. G. Kelley, J. Bresee, L. F. Cranor, and R. W. Reeder, "A "nutrition label" for privacy," in *Proceedings of the 5th Symposium on Usable Privacy and Security*, 2009.

[46] L. Schaewitz, D. Lakotta, M. A. Sasse, and N. Rummel, "Peeking into the black box: Towards understanding user understanding of E2EE," in *Proceedings Of The 2021 European Symposium On Usable Security*, 2021, pp. 129–140.

[47] "Supplementary Material," https://anonymous.4open.science/r/TEE-2024-D1A4/README.md, 2024.

[48] Pew Research Center, "Mobile fact sheet," https://www.pewresearch.org/internet/fact-sheet/mobile/, accessed: 2024-05-17.

[49] O. E. Karpov, E. N. Pitsik, S. A. Kurkin, V. A. Maksimenko, A. V. Gusev, N. N. Shusharina, and A. E. Hramov, "Analysis of publication activity and research trends in the field of AI medical applications," *International Journal of Environmental Research and Public Health*, vol. 20, no. 7, p. 5335, 2023.

[50] Google, "Google Home Gemini," https://support.google.com/gemini/answer/15335456, accessed: 2024-05-17.

[51] Amazon, "LLM-powered Alexa experiences," https://developer.amazon.com/en-US/alexa/alexa-ai, accessed: 2024-05-17.

[52] Arm Developer Hub, "Arm TrustZone," https://developer.arm.com/documentation/102418/0101/What-is-TrustZone-?lang=en, 2024, [Accessed 27-06-2024].

[53] Android Open Source Project, "Android Gatekeeper," https://source.android.com/docs/security/features/authentication/gatekeeper, 2024, [Accessed 17-06-2024].

[54] Confidential Computing Consortium, "Common terminology for confidential computing," https://confidentialcomputing.io/wp-content/uploads/sites/10/2023/03/Common-Terminology-for-Confidential-Computing.pdf, 2024, [Accessed 17-06-2024].

[55] Prolific, "How do I balance my sample within demographics?" https://researcher-help.prolific.com/hc/en-gb/articles/360009221213-How-do-I-balance-my-sample-within-demographics, accessed: 2024-05-17.

[56] M. van Smeden, K. G. Moons, J. A. de Groot, G. S. Collins, D. G. Altman, M. J. Eijkemans, and J. B. Reitsma, "Sample size for binary logistic prediction models: Beyond events per variable criteria," *Statistical Methods in Medical Research*, vol. 28, no. 8, 2019.

[57] B. D. Douglas, P. J. Ewell, and M. Brauer, "Data quality in online human-subjects research: Comparisons between MTurk, Prolific, CloudResearch, Qualtrics, and SONA," *Plos one*, vol. 18, no. 3, 2023.

[58] J. Tang, E. Birrell, and A. Lerner, "Replication: How well do my results generalize now? The external validity of online privacy and security surveys," in *Proceedings of the 18th Symposium on Usable Privacy and Security*, 2022.

Appendix A. Candidate TEE Explanations

Candidate TEE explanations are composed of 2-3 sentences, where each sentence has a different theme. We evaluate every combination of the 2-3 sentences in our surveys. Each theme is shown below in *italics*, followed by the corresponding sentence from our evaluation.

Sentence 1: Introducing TEEs
Hardware: A Trusted Execution Environment (TEE) is a technique for running programs and interacting with data securely using a protected area of the physical computer.
Trust: A Trusted Execution Environment (TEE) is a technique for running programs and interacting with data securely, even if the rest of the computer is not trustworthy.
Unsubstantial: A Trusted Execution Environment (TEE) is a technique for running programs and interacting with data securely.

Sentence 2: Isolation, confidentiality, and integrity
Technical: A program running in a TEE is isolated from the rest of the computer to protect the confidentiality and integrity of the program and data.
Non-Technical: A program running in a TEE is isolated from the rest of the computer to allow only authorized people to view or change the program and data.

Sentence 3: (Optional) threat prevented by TEE
Prevents: The TEE protects the program and data even when other software on the computer is behaving maliciously.
No Prevents: (No third sentence)

Appendix B. Additional Results and Statistics

The complete set of comprehension questions from Survey 1 and 2 are shown in Table 4. Tables 5- 8 include details about our statistics.

Q#	T/F	Question Text
Q1	F	A member of the general public can access your data
Q2	F	___ can access your data *Medical*: A hospital employee unrelated to the research team *Smart home*: Someone working at the company on an unrelated team
Q3	F	If there were a bug in other software on the computer, outside of the TEE storing your data, then a hacker could use the bug to access your data
Q4	F	If a disgruntled [___] installed a malicious program on the computer storing your data, then they could access your data *Medical*: hospital employee unrelated to the research team *Smart home*: employee on an unrelated team
Q5	F	Other [___] working on different projects on the same computer can access your data *Medical*: researchers *Smart home*: developers
Q6	T	A member of the [___] can access your data *Medical*: research team *Smart home*: development team
Q7	T	If [___] makes a mistake collecting your data, then your data could be incorrect *Medical*: a member of the research team *Smart home*: the light bulb / the voice assistant
Q8	T	A [___] could later use your data to [___] *Medical without AI*: A member of the research team / choose the location for a new fire station *Medical with AI*: A member of the research team / train another AI diagnosis tool for a different medical condition *Smart home*: Someone on the development team / develop a smart vacuum
Q9	F	The TEE ensures [___] *Medical without AI*: the hospital being constructed will be closer to the patients who most need it *Medical with AI*: the diagnosis made by the AI tool will always be correct *Smart home without AI*: the new light bulbs will have features relevant to you *Smart home with AI*: your voice will always be recognized by the improved AI
Q10	T	[___] could steal your data and sell it on the dark web *Medical*: A member of the research team *Smart home*: Someone on the development team
Q11	T	When you unlock your Android phone with a PIN, the PIN is verified in a TEE
Q12	F	We cannot be sure that a TEE is configured correctly

TABLE 4: Questions for evaluating TEE concept comprehension. The expected answer (**T**rue or **F**alse) is shown in the second column. **Q11** and **Q12** only appear in the follow-up survey. We note the places where the questions differ between scenarios.

	Medical				Smart Home			
	With AI		Without AI		With AI		Without AI	
Aspects	W	p-value	W	p-value	W	p-value	W	p-value
Use of TEE	1421	0.002597^{**}	1638.5	0.0002417^{***}	1483.5	0.0006844^{***}	1547.5	0.0005205^{***}
Purpose	2168	0.636	2689	0.8193	2231	0.3749	2258.5	0.3077
Data Collect	2258	0.8642	2815.5	0.8709	2363	0.6407	2384	0.5405
Hospital/Company	2090	0.4634	3285	0.09689	2637.5	0.6806	2263	0.3194
People on the team	1965	0.2447	3412.5	0.04132^{*}	2580	0.8183	2309.5	0.3909

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

TABLE 5: Table with results of Mann-Whitney U tests for the aspects contributing to the belief that the data would be safe/unsafe, with information about the TEE (TEE explanation or FAQ) vs no TEE information provided. Each scenario was treated separately. Statistical significance is noted with asterisks.

	Medical				Smart Home			
	With AI		Without AI		With AI		Without AI	
	W	p-value	W	p-value	W	p-value	W	p-value
Average Score	8226	0.5146	11723	$1.505e-06^{***}$	10020	0.004367^{**}	13756	$4.378e-09^{***}$

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

TABLE 6: Table with results of Wilcoxon signed-rank tests comparing the average score for questions Q1-Q5, to the average score for questions Q6-Q10 in the follow-up study. Each scenario was treated separately. Statistical significance is noted with asterisks.

Variable	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10
Medical Scenario Without AI										
Explanation sentence 1 [Baseline = Unsubstantial]										
Hardware	0.42	−0.96	0.28	−0.01	−0.85	0.03	−0.76	−0.09	−0.90**	0.25
Trust	0.38	−0.62	0.68	−0.01	−0.58	0.41	−0.35	0.41	−0.24	−0.07
Explanation sentence 2 [Baseline = Technical]										
Non-technical	0.16	2.14**	−0.07	0.17	0.33	0.49	−0.12	0.14	0.08	0.20
Explanation sentence 3 [Baseline = No Prevents]										
Prevents	−0.38	0.10	0.64*	1.30***	0.04	−0.11	−0.30	−0.01	0.13	−0.64*
Medical experience	−1.43	−0.08	0.09	0.75	0.29	−0.18	0.02	0.02	0.27	0.34
CS experience	−0.31	0.49	−0.03	0.19	0.35	0.16	0.08	0.53	0.37	−0.07
Medical Scenario With AI										
Variable	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**	**Q6**	**Q7**	**Q8**	**Q9**	**Q10**
Explanation sentence 1 [Baseline = Unsubstantial]										
Hardware	0.03	1.60	0.07	0.11	0.11	1.04	0.62	0.41	−0.13	−0.31
Trust	−1.77	−0.78	−0.15	−0.18	0.15	0.36	0.49	0.18	−0.18	0.04
Explanation sentence 2 [Baseline = Technical]										
Non-technical	−0.24	0.47	0.10	−0.47	0.43	−0.12	−0.47	−0.49	−0.55	0.71*
Explanation sentence 3 [Baseline = No Prevents]										
Prevents	−0.56	−0.93	0.53	1.04**	0.24	−0.04	0.56	0.51	−0.61	0.01
Medical experience	0.32	0.79	0.21	0.01	−0.23	0.27	−0.05	1.01	1.13	1.28***
CS experience	−1.45*	−0.82	−0.14	−0.78*	−0.95*	−0.07	0.24	0.40	−0.23	0.26
Smart Home Scenario Without AI										
Variable	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**	**Q6**	**Q7**	**Q8**	**Q9**	**Q10**
Explanation sentence 1 [Baseline = Unsubstantial]										
Hardware	−0.04	−0.65	−0.28	−0.35	0.02	0.19	−0.10	−0.11	−0.08	0.32
Trust	−0.48	−0.46	0.33	0.19	−0.29	−0.07	−0.34	−0.01	−0.13	0.19
Explanation sentence 2 [Baseline = Technical]										
Non-techncial	−0.18	0.94	0.23	−0.08	0.04	0.28	−0.14	0.11	−0.25	0.24
Explanation sentence 3 [Baseline = No Prevents]										
Prevents	−0.99	−0.01	0.07	0.51	0.16	0.03	−0.53	0.03	0.11	0.01
Smart home experience	−15.97	−1.46	−1.26*	−1.02	−0.46	−0.57	−0.51	−0.32	−0.79*	−0.03
CS experience	−1.76*	0.41	−0.36	−0.52	−0.16	−1.00*	−0.05	−0.68*	0.66	−0.43
Smart Home Scenario With AI										
Variable	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**	**Q6**	**Q7**	**Q8**	**Q9**	**Q10**
Explanation sentence 1 [Baseline = Unsubstantial]										
Hardware	−0.81	−0.10	0.31	0.39	0.29	−0.35	−0.41	−0.11	−0.49	−0.42
Trust	−1.72	−1.06	0.35	−0.21	−0.22	−0.22	0.16	0.40	−0.70*	0.02
Explanation sentence 2 [Baseline = Technical]										
Non-technical	0.13	0.71	0.14	−0.36	−0.05	0.02	−0.61	0.21	0.11	0.35
Explanation sentence 3 [Baseline = No Prevents]										
Prevents	−0.60	0.73	−0.39	0.99**	0.80	0.15	0.26	0.11	−0.25	−0.40
Smart home experience	−0.60	0.73	−0.39	0.23	0.18	−0.65	0.06	−0.36	0.01	−0.49
CS experience	−0.79	−0.46	0.16	0.23	−0.39	0.07	0.33	0.24	0.42	0.29

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

TABLE 7: Regression table for True/False comprehension questions in **Survey 1** where each question in each scenario is a different model (40 models total). The numbers are the log-odds coefficients for each predictor, with the baseline used in each model noted in *italics*. Statistical significance is noted with asterisks and shaded cells: blue for positive and orange for negative coefficients.

Medical Scenario Without AI												
Variable	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**	**Q6**	**Q7**	**Q8**	**Q9**	**Q10**	**Q11**	**Q12**
Explanation [Baseline = No Explanation]												
Unsubstantial	−0.42	−1.17	0.76*	0.70	0.95*	−1.33	0.35	−0.02	0.23	−0.40	−0.50	0.21
Hardware	0.66	−0.27	0.93*	1.44**	0.65	−0.71	0.29	−0.13	0.44	−0.33	−0.55	0.34
Trust	−0.34	−1.67*	1.36**	0.99*	0.50	−0.50	−0.11	−0.20	0.14	−0.36	−0.40	0.16
FAQ [Baseline = No FAQ]												
Hidden	18.07	0.29	1.29**	0.33	0.72	0.24	−0.75	0.12	−0.07	−0.83*	2.96***	1.83***
Shown	0.53	0.84	0.37	0.68	0.71	−0.29	−1.00	−0.30	0.18	−1.01**	3.17***	1.86***
Medical Exp	−0.98	0.97	−0.12	−0.55	−0.46	−0.76	0.04	0.42	−0.10	0.12	0.31	−0.11
CS Exp	−0.84	−0.09	0.10	0.34	0.09	−0.38	−0.11	0.20	0.13	−0.02	0.04	−0.34

Medical Scenario With AI												
Variable	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**	**Q6**	**Q7**	**Q8**	**Q9**	**Q10**	**Q11**	**Q12**
Explanation [Baseline = No Explanation]												
Unsubstantial	1.27	0.39	0.56	1.21**	1.28*	0.92	1.05	0.27	1.11*	0.45	−0.59	−0.54
Hardware	1.61	0.88	0.79*	0.88*	0.80	0.93	0.55	0.44	0.36	−0.01	−0.08	−0.24
Trust	0.24	0.19	0.52	0.99*	−0.01	1.66*	0.63	0.33	0.71	0.35	0.13	−0.64
FAQ [Baseline = No FAQ]												
Hidden	−0.01	0.77	−0.14	0.20	0.25	−1.42*	0.43	−1.36**	−0.23	−0.48	2.30***	1.20***
Shown	0.47	1.31*	0.13	0.25	0.06	−0.54	0.76	−0.82	−0.08	−0.15	3.27***	1.56***
Medical Exp	−0.69	−0.49	−0.35	−0.78*	−0.25	−0.64	0.25	0.56	−0.59	0.45	−0.04	−0.41
CS Exp	0.84	−0.13	−0.45	−0.93**	0.12	−0.88	0.03	0.63	0.11	0.03	0.54	−0.18

Smart Home Scenario Without AI												
Variable	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**	**Q6**	**Q7**	**Q8**	**Q9**	**Q10**	**Q11**	**Q12**
Explanation [Baseline = No Explanation]												
Unsubstantial	−1.45	−0.01	0.78	1.07*	−0.52	0.07	0.19	0.16	−0.05	0.35	−0.63	−0.52
Hardware	−1.21	0.52	0.99*	1.05*	0.20	−0.38	−0.02	−0.41	−0.11	0.16	−0.55	−0.21
Trust	−0.61	0.36	0.61	0.70	−0.10	0.78	0.49	0.33	−0.14	0.27	−0.08	−0.30
FAQ [Baseline = No FAQ]												
Hidden	0.77	0.14	0.10	0.31	0.48	0.13	−0.05	0.08	−0.53	−0.62*	2.99***	1.14***
Shown	1.61	1.13	0.13	0.19	0.93*	0.69	−0.78	−0.39	−0.21	−0.22	2.93***	1.43***
Smart Home Experience	0.43	0.02	0.68	0.75	0.73	0.05	0.02	0.05	−0.30	−0.31	0.35	0.16
CS Exp	−1.04	0.21	−0.54	−0.61	−0.39	−0.27	−0.18	0.74	0.86**	0.14	0.51	−0.42

Smart Home Scenario With AI												
Variable	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**	**Q6**	**Q7**	**Q8**	**Q9**	**Q10**	**Q11**	**Q12**
Explanation [Baseline = No Explanation]												
Unsubstantial	0.89	0.88	1.24**	1.02*	1.64***	−0.95	−0.01	−0.15	−0.43	−0.01	−0.46	−0.06
Hardware	0.64	0.72	1.10**	1.12**	1.17**	−0.07	0.13	−0.12	0.19	−0.70	−0.34	0.06
Trust	−0.67	−0.05	1.14**	0.83*	0.77	−0.64	0.49	−0.25	0.17	−0.34	−0.39	−0.15
FAQ [Baseline = No FAQ]												
Hidden	0.05	−0.19	0.07	0.16	−0.29	−0.33	−0.54	−0.26	−0.02	−0.63	2.53***	1.25***
Shown	−0.18	−0.71	0.86*	0.23	0.14	−0.61	−0.40	−0.81*	−0.37	−0.51	3.77***	1.30***
Smart Home Experience	1.18	−1.32	0.24	−0.53	−1.27*	−0.97	−0.99	−0.69	−0.76*	−0.42	0.78	−0.16
CS Exp	−1.42	−0.64	−0.34	−0.60	0.04	0.24	−0.24	−0.19	0.52	0.29	0.43	−0.41

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

TABLE 8: Regression table for True/False comprehension questions in **survey 2**. There is one logistic regression model for each question in each scenario (48 models total). The numbers in this table are the log-odds coefficients for each predictor, with the baseline explanations used in each model noted in *italics*. Statistical significance is noted with asterisks and shaded cells: blue for positive coefficients and orange for negative coefficients.