

Beyond Clinical Risk: An Experimental Study of Cybersecurity Informed Consent and Patient Choice for Connected Medical Devices

Ronald E. Thompson III
Department of Computer Science
Tufts University
Medford, Massachusetts, USA
rthomp06@cs.tufts.edu

R. Harrison Sweet
Department of Computer Science
Tufts University
Medford, Massachusetts, USA
harrison91sweet@gmail.com

Christian J Dameff
Department of Emergency Medicine
University of California, San Diego
La Jolla, California, USA
cdameff@health.ucsd.edu

Jeffrey L. Tully
Department of Anesthesiology
University of California, San Diego
La Jolla, California, USA
jtully@health.ucsd.edu

Daniel Votipka
Department of Computer Science
Tufts University
Medford, Massachusetts, USA
daniel.votipka@tufts.edu

Abstract

Internet-connected medical devices introduce complex cybersecurity risks that challenge the established practice of informed consent. It remains unclear how patients weigh these abstract, dynamic threats against concrete clinical benefits. We present findings from a large-scale (N=2,666) vignette-based experiment designed to uncover the factors driving patient decision-making. Participants chose whether to adopt a connected pacemaker, weighing its enhanced clinical outcomes against potential vulnerabilities. We systematically varied communication factors, including the source of risk information (e.g., clinician, FDA), risk framing, and the details of a subsequent vulnerability disclosure. Our results reveal patient choice hinges on pre-existing physician trust and risk framing. We did not observe any effect from the information's source. We also find initial choices act as powerful anchors, and that detailed disclosures increase security confidence. Our work provides crucial empirical evidence on this trade-off, offering actionable guidance to better support informed consent for life-critical connected technologies.

CCS Concepts

• **Security and privacy** → **Human and societal aspects of security and privacy**; **Usability in security and privacy**; • **Social and professional topics** → **Medical information policy**; • **Applied computing** → *Life and medical sciences*.

Keywords

Usable Security, Informed Consent, Patient-Centric Security, Medical Devices

ACM Reference Format:

Ronald E. Thompson III, R. Harrison Sweet, Christian J Dameff, Jeffrey L. Tully, and Daniel Votipka. 2026. Beyond Clinical Risk: An Experimental

Study of Cybersecurity Informed Consent and Patient Choice for Connected Medical Devices. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 31 pages. <https://doi.org/10.1145/3772318.3791788>

1 Introduction

The proliferation of internet-connected medical devices represents a paradigm shift in modern healthcare. From insulin pumps to implantable cardiac pacemakers, these technologies offer unprecedented opportunities to improve patient outcomes and provide continuous care [72]. This connectivity, however, introduces a class of risk fundamentally different from traditional clinical complications: cybersecurity vulnerabilities. The security community has extensively documented specific vulnerabilities in medical devices such as pacemakers and infusion pumps [7, 24, 32, 36, 40, 45, 46, 54–56, 64, 66]. A malicious actor could potentially exploit a software flaw to access sensitive patient data, alter a device's function, or cause physical harm, turning a life-saving device into a potential liability and a significant patient safety concern [40, 45, 56].

This new risk landscape poses a profound challenge to one of the foundational pillars of medical ethics: informed consent. The traditional informed consent process is built on a dialogue between a provider and a patient about the known risks, benefits, and alternatives of a proposed treatment [3, 44, 59]. Yet, cybersecurity risks are dynamic, and challenging even for security experts to assess, as exploitation likelihood depends on factors far outside the clinical environment, such as shifting attacker incentives, newly disclosed vulnerabilities, and changes in the broader threat landscape [80]. Healthcare providers, who lack specialized cybersecurity training, currently must communicate these complex, evolving risks to patients during consent discussions. Recognizing these challenges, Tully et al. proposed “cybersecurity informed consent” as a necessary paradigm for connected medical devices, identifying key implementation challenges including clinician knowledge gaps, the involvement of multiple authoritative sources (clinicians, regulators, and manufacturers), and the need for dynamic risk communication [76]. However, their framework is conceptual with little empirical evidence about what factors actually influence patient



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2278-3/26/04
<https://doi.org/10.1145/3772318.3791788>

decisions, how patients weigh clinical benefits against security risk, or whether established risk communication principles operate differently in medical security contexts [10, 37].

A rich body of literature in usable security and risk communication has explored how to present complex technical information to non-experts, often finding that factors like message framing and warning design are critical [2, 61, 62, 77, 82]. However, medical device security risk communication introduces unique aspects due to the highly regulated, high-stakes context of medicine and the existence of regulating bodies, e.g., the US Food and Drug Administration (FDA), and authority figures, i.e., healthcare providers, users are accustomed to communicating. Therefore, it is unknown how patients weigh abstract cybersecurity threats against concrete clinical benefits in this context and how factors such as trust in their healthcare provider and the source of risk information (e.g., provider, FDA, or device manufacturer) influence their choices.

In this paper, we seek to close this gap by empirically examining how patients weigh different clinical and cybersecurity risk information along with information sources when making critical medical decisions. Specifically, we sought to answer the following research questions:

- RQ1:** How do external communication factors, such as the information source and the framing of risk, compared to internal patient attributes, such as their baseline security attitudes and trust in providers, influence a patient’s initial choice to adopt a connected medical device?
- RQ2:** How does a patient’s decision-making change when new risk information (a vulnerability disclosure) is introduced after the initial choice has been made?
- RQ3:** How does the disclosure of a vulnerability, and attributes of that disclosure, impact a patient’s confidence in the device’s security and clinical necessity?

To answer these questions, we conducted a large-scale controlled experiment—using a vignette-based survey of 2,666 US Prolific users. We asked participants to imagine they needed a pacemaker and presented them with information about a connected medical device with better expected patient care outcomes but worse potential security, and a non-connected medical device with worse expected care, but better security. Then, we asked participants to indicate which they would prefer. To understand their decision-making, we varied the information source, the framing of health benefits, the expected likelihood of a hack, and the potential impact of an attack. After making their choice, we informed participants that a vulnerability had been discovered, varying information about the impact and likely exploitability of the vulnerability, and asked them whether they would choose to disable device connectivity (i.e., mitigating the vulnerability, but losing care benefits). This two-phase design allows us to empirically measure the effects of specific communication variables that are central to initially understanding cybersecurity-informed consent, including the role of different authoritative sources and how patients re-evaluate decisions as risks evolve.

Our findings reveal a difference in the factors that influence patient decision-making. We find that internal patient attributes, especially baseline trust in their physician and the framing of a

cyberattack as a threat to physical safety, are far more powerful predictors of choice than external communication factors like who provides the risk information. After an initial choice is made, it creates a powerful anchor, making patients unlikely to change their minds, even when presented with a new vulnerability. Counterintuitively, we find that providing more detailed, verbose disclosures about vulnerabilities does not alarm patients; instead, it significantly increases their confidence in the device’s security, particularly when the notice comes from a regulator like the FDA.

Based on these results, we outline recommendations for healthcare providers, regulators, and other researchers to support improved cybersecurity-informed consent communications.

2 Related Work

Our work is situated at the intersection of medical informed consent, human-computer interaction (HCI) for medical devices, patient-facing security communication, and security advice. We review prior work in these areas to contextualize our contributions.

Medical Informed Consent. The process of informed consent is a cornerstone of ethical medical practice, intended to facilitate shared decision-making between patients and providers [34, 48]. However, research has shown that the implementation often falls short of this ideal. Consent forms are frequently designed more for institutional liability protection than for patient comprehension, often failing to include all basic elements of consent, such as risks, benefits, and alternatives [9]. Recent systematic reviews indicate these challenges are enduring; Sherman et al. found that current measures of informed consent adequacy are often methodologically flawed [67], and Grady argues that the increasing complexity of clinical information continues to widen the gap between the ethical ideal and clinical reality [33]. Furthermore, even when information is provided, patient recall of risks is often poor, a problem that persists even with the use of procedure-specific forms [60]. The communication of evidence, particularly probabilistic risk information, is a known challenge; studies suggest that representing probabilities as natural frequencies and using structured, interactive formats can improve patient understanding [75]. To address these deficits, researchers have examined interventions to improve comprehension, finding that interactive and digital formats—rather than static forms—can significantly enhance patient understanding [30, 44].

Our work builds on this prior literature by considering challenges specific to communicating cybersecurity risks during the informed consent process. Williams and Woodward have framed medical device security as a complex socio-technical ecosystem that complicates patient decision-making [78]. Further, Tully et al. proposed “cybersecurity informed consent” as a new paradigm required to address the unique threat landscape of connected health, especially as many healthcare providers lack security awareness or expertise [76]. Our work operationalizes these theoretical frameworks, providing the empirical data necessary to understand how patients navigate this proposed paradigm. While prior work has focused on evaluating the efficacy of consent through measures such as patient recall, our experiment is designed to investigate the communication of security risks, by systematically varying their framing, impact, and likelihood, influences patients’ choices and

confidence in a simulated consent dialogue. Unlike studies evaluating static consent forms, our two-phase vignette design simulates the evolution of consent after new information (a vulnerability) is discovered, allowing us to examine how patients re-evaluate their initial decisions.

HCI and Medical Devices. The HCI community has a long history of evaluating the safety and usability of medical technologies [21, 42]. The CHI+MED project, for example, extensively mapped how interaction design influences the safety of interactive medical devices [21, 22]. Prior work has used usability heuristics to evaluate patient safety [81] and assessed the usability of specific interfaces, such as pacemaker programmers [18]. However, much of this work focuses on the clinician’s user experience or the prevention of use errors in high-stakes clinical environments [5]. There remains a gap in understanding how patients themselves evaluate the specific security risks of the devices implanted in their own bodies. Our work shifts the HCI lens from the clinician’s interface to the patient’s mental model of risk.

Patient-Facing Security Communication for Medical Devices. There is a growing recognition within both the clinical and security communities of the need to communicate cybersecurity risks to patients with implantable medical devices [4, 70, 80]. This body of work is largely prescriptive, offering guidance to clinicians but lacking empirical data on how patients perceive and act on this information. The study closest to our own is the foundational work by Denning et al., which explored patient views on security for implantable cardiac devices [23]. In their study, Denning et al. conducted semi-structured interviews with 13 patients, using mockups of potential security systems to elicit values and preferences. This qualitative approach was instrumental in identifying key patient concerns, including safety and freedom from unwanted cultural associations. Our research seeks to build directly upon this work by employing a large-scale, quantitative methodology to investigate similar questions. While the qualitative approach of Denning et al. was ideal for initial exploration, our use of a large-scale, between-subjects factorial experiment allows us to systematically isolate and measure the impact of specific communication variables (e.g., framing a risk as a physical safety issue versus a data privacy issue) on patient decisions. Instead of exploring general preferences, our method allows for the quantification and comparison of the relative influence of different factors within a simulated clinical choice.

Post-Breach and Vulnerability Notifications. Our work on vulnerability disclosures draws from research on data breach notifications in the broader security literature. Methodologically, much of this work has relied on content analysis of real-world notices. For instance, Zou et al. analyzed a corpus of breach notifications and found them to be challenging to read, to downplay risks, and to provide unclear guidance [82]. Other work has used surveys and interviews to understand consumer responses after they have been affected by a breach [51]. Our study contributes a different methodological approach to this area. We adapt the context from general consumer data to the high-stakes domain of medical device vulnerabilities, and instead of a post-hoc analysis, we use a

controlled experimental design. This allows us to proactively construct and test the effects of different notification characteristics, such as verbosity, source, and risk likelihood, on patient behavioral intentions and trust.

Security Advice and Persuasion. Finally, our research is informed by prior work on the efficacy of security advice. This work has established that effectiveness depends on factors like users’ mental models [11], the framing and length of the message [31], and the source of the advice [61, 62, 77]. For example, Wash and Cooper used an experiment to show that the ideal presenter of security advice depends on the format, with experts being more persuasive with facts and peers more so with stories [77]. Our study employs a similar experimental logic but translates it to the clinical domain, where the “advice” is part of a medical decision. We directly investigate the role of the advice source by experimentally varying whether security information is attributed to a healthcare provider, the FDA, or the device manufacturer. This methodology allows us to test whether findings from the general security literature translate to the unique context of medical decision-making, where pre-existing roles and trust relationships, such as that between a patient and provider, may interact with the principles of security persuasion.

3 Methodology

Our study was designed to empirically investigate how patients perceive and weigh the novel risks introduced by internet-connected medical devices, following the calls for a more patient-centric approach to medical cybersecurity [10, 76, 78]. Our survey simulated the consent dialogue in a two-phased vignette-based scenario, varying characteristics of the information provided during the consent process to answer each of our research questions. Each scenario consists of an initial device choice (RQ1) and a post-vulnerability notification choice re-evaluation phase (RQ2, RQ3). All our research questions, vignettes, and survey questions were designed in collaboration with two physicians on the research team to ensure that the scenarios, language, and trade-offs accurately reflect a real-world consent dialogue.

3.1 Survey Design

Figure 1 shows all stages of our online survey, and the full survey text is given in Appendix A. We now discuss each part of the survey in turn.

Consent (Fig 1.A). Before answering any survey questions, participants were presented with a consent form detailing the study’s purpose, procedures, potential risks, and data handling policies. They were required to provide informed consent to proceed.

Baseline trust in medical providers (Fig 1.B). After providing consent, participants were asked about their prior interactions with primary care providers and various medical specialists who would provide care relevant to pacemakers (e.g., cardiologists, electrophysiologists). They then completed the Wake Forest Physician Trust Scale (WFPT), which asks participants to rate their agreement with ten statements, to establish a baseline measure of their trust in their own medical providers [35]. This allows us to control for

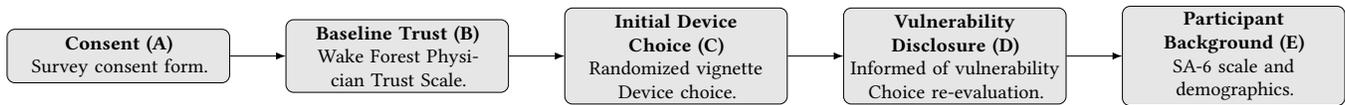


Figure 1: The sequential flow of the survey, from initial consent to final demographic questions.

pre-existing attitudes towards and experience with the medical establishment. Because half of these statements are positive and the other half are negative, we used this as an attention check to see if participants selected the same response for all the statements.

Initial scenario and device choice (RQ1; Fig 1.C). Participants were then introduced to a detailed clinical vignette. We chose a pacemaker as the hypothetical device to build directly upon the existing discourse in the clinical literature, which has often used implantable cardiac devices as a key case study for these emerging challenges [69, 80]. A pacemaker is a widely understood implantable device, has both non-connected and internet-connected variants in real-world use [8, 12, 52, 53], and faces plausible cybersecurity threats that could impact patient safety [24, 36, 46].

The vignette was structured to explicitly include the four essential elements of informed consent: the nature of the procedure (general information), the benefits (clinical outcomes), the risks (cybersecurity and clinical), and the alternatives (non-connected device) [44]. Additionally, we used the scenario generated by the US Food and Drug Administration (FDA) as part of their Patient Engagement Advisory Committee meeting in 2019 as an initial starting point and then further refined based on the experience of the physicians on the research team [50]. The scenario asked participants to imagine they had been diagnosed with a heart block requiring a pacemaker [19] and were being given information to help them decide whether to get an internet-connected or non-connected model. This phase employed a between-subjects factorial design, varying the information presented to assess which features impact participants’ decision-making. Specifically, we vary the information provider (including verbosity), health benefit, initial security likelihood, and impact of an attack (discussed in Section 3.2). An example of the scenario presented to participants is shown in Fig 2.

After reading the scenario, participants were asked to indicate their device preference based on the presented information using a five-point Likert scale from “Definitely the internet-connected device” to “Definitely the non-connected device.” They were also asked to separately rate their confidence in their understanding of the clinical and security information (using a five-point Likert scale for each). Participants were then asked to provide a reason for their confidence in the security and clinical information.

Vulnerability disclosure (RQ2, RQ3; Fig 1.D). Next, we simulated the discovery of a vulnerability in the pacemaker. Participants were asked to assume they had chosen the internet-connected pacemaker so we could capture their reported actions in response to the vulnerability. Then, we told participants that the same party that had initially given them the clinical and security information about the pacemaker had disclosed the vulnerability to them. In this phase, we varied the verbosity of the vulnerability disclosure and risk of exploitation. We describe these variations in detail in

Section 3.2. Based on this disclosure, participants were asked to indicate whether they were “Definitely likely” to “Definitely unlikely” to disable the device’s internet-connectivity or accept the security risk to maintain the healthcare benefits of the connected device (RQ2). We also asked participants to discuss the reason for their decision in an open-ended question. We concluded this section by asking participants to reassess their trust in the security and clinical information they were provided based on the vulnerability disclosure (RQ3). At the end of this section, we included an attention check that asked participants to select the party that provided them with the notification.

Participant background (Fig 1.E). Finally, we asked a series of questions to capture information about participants’ backgrounds that might impact their decision-making. First, participants answered questions about their personal or family experience with medical devices, as this likely influences their awareness of potential medical challenges associated with these devices. Next, they completed the Security Attitudes Six (SA-6) questionnaire to measure their general disposition towards security practices [26]. Then, they provided standard demographic information, including age, race/ethnicity, and gender.

3.2 Conditions

Each participant was randomly assigned to one of our experimental conditions in a between-subjects design. The experiment unfolded in two phases: an *Initial Choice* scenario and a subsequent *Vulnerability Disclosure* scenario. We manipulated several variables within each phase to address our research questions.

Initial Choice Scenario Variables. To answer our first research question (RQ1) about the factors influencing a patient’s initial choice, we varied the information presented about the connected pacemaker along four key dimensions: the information source, the impact of an attack, the device’s health benefits, and the initial likelihood of the attack.

To understand how patients weigh information from different authority figures, we varied the **Information Source**. Participants were told the security information came from one of three sources: their *Healthcare Provider*, a government agency (the *FDA*), or the *Medical Device Manufacturer* (MDM). In each case, participants are told they receive information from their healthcare provider, but we vary whether we tell participants that the information initially came from one of these authorities with more technical expertise. We specifically chose this structure to assess whether patients trust providers directly to provide reliable security information, even though they often lack security expertise, or if it is important that an organization with more technical expertise is cited. Further, we sought to investigate whether patients find this expert information more reliable if it is coming from the company with deep technical

We are going to ask you some questions about how you would make decisions about using medical devices.

We want you to imagine that you have been diagnosed with a heart block, which means that your heart has trouble passing electrical impulses across your heart. This can cause you to have a shortness of breath and cause chest pain as some of the symptoms, but it can also lead to heart failure if not effectively managed. Because of this, your medical provider, i.e. the provider that you routinely see, has said that you require a pacemaker. Your provider is recommending that you use a new pacemaker that is Internet-connected and can be remotely monitored by your care team, which **will have a positive impact on your condition**. The health risks associated with this device are minimal. Additionally, because this device can connect to the Internet, **your medical provider gives you information from a government agency (such as the FDA) that explains there is a low likelihood that the Internet-connected device could be hacked**. If the device was hacked, it could lead to **your private patient data being stolen**. The FDA has approved the clinical use of this device to treat heart blocks.

There are alternative therapies available to treat this condition, and your provider explains that you have the option of a non-connected pacemaker. However, **there is a slightly increased chance that an adverse event (something bad happening) will not be caught**, as the device cannot be remotely monitored or controlled, you will need to see your provider on a more frequent basis. This pacemaker is very unlikely to be hacked as it is not connected to the Internet. Like the Internet-connected device, this has been approved by the FDA for the treatment of your condition.

(a) Initial information shown to participant. Information Provider: FDA, Health Benefit: Low, Initial Security Likelihood: Low, Impact of Attack: Data

Figure 2: Example vignettes shown to participants. The bolded sections are the parts we varied in our experimental conditions.

Now, assume you have chosen to use the Internet-connected pacemaker. Some time after you have been using the device, you are informed by **the manufacturer of the device** that a cybersecurity vulnerability has been discovered in the device and that **they are unsure if the Internet-connected device** could be hacked. If the device was hacked, it could lead to **your private patient data being stolen**.

As a reminder, the Internet-connected pacemaker **will catch an adverse event (something bad happening) more quickly**.

Assume there is no cost and no surgery is required to disable the Internet connection on the pacemaker, which would remove the benefits of remote monitoring, and **a life-threatening event is much more likely not to be caught until it is too late**. This would involve a one-time short appointment with your medical provider. How likely are you to disable the Internet connection on the pacemaker?

(b) Vulnerability information shown to participant and question about switching devices. Information Source: MDM, Verbosity: Non-Verbose, Health Benefit: High, Vulnerability Security Likelihood: Unsure, Impact of Attack: Data

knowledge of the product they created (i.e., the MDM), but who might have incentives to downplay risks or from a third-party regulator (i.e., the FDA) who also have technical expertise, though not as specific to the device, and do not have the same business incentives.

Next, to understand how the framing of security risks and clinical benefits affects decision-making, we varied two aspects of the scenario's context. First, the **Impact of an Attack** was framed as either a risk to personal data (*Data Privacy Issue*) or a risk to the patient's physical well-being (*Physical Safety Issue*). These risks represent clear, direct impacts of potential attacks. That is, confidentiality violations introduce data privacy issues, and availability and integrity violations would likely both lead to physical safety issues due to the life-critical nature of the device. We chose to present these risks at a high level instead of discussing a specific technical impact to avoid confusing non-technical participants and focus just on the tangible outcome.

In parallel, the **Health Benefit** of the connected device was described as either relatively minor (*Low*, e.g., a positive impact on their condition) or significant (*High*, e.g., the ability to catch a life-threatening event quickly). These health impacts reflect real-world

benefits of a pacemaker [19], and these levels of description represent the spectrum of real-world presentations of health benefits for pacemakers [8, 12, 52, 65].

Finally, we varied the **Initial Security Likelihood** to test the effect of ambiguity. The likelihood of the device being hacked was described as either explicitly *Low* or as *Unsure*. We deliberately excluded a "high" likelihood condition in this initial phase to maintain ecological validity as the FDA would not approve a device that had an unmitigated, known, high-risk vulnerability [17].

Vulnerability Disclosure Scenario Variables. After making their initial choice, participants were presented with a second scenario to address our remaining research questions (RQ2 and RQ3). They were asked to assume they had chosen the connected device and were now being informed of a newly discovered vulnerability. We varied the disclosure of the vulnerability along the same four key dimensions (i.e., the information source, the security impact, the health benefits, and the likelihood of an attack), but added one additional dimension (verbosity of the disclosure) and added an additional level for attack likelihood. Note, the information source, security impact, and health benefits were consistent between the initial information and disclosure for each participant. For example, if a participant was initial told the attack impact would be a

Variable	Value	Text in survey
Notice Type	-	"... you are informed by..."
	Provider	"... your medical provider"
	FDA	"... your medical provider gives you information from a government agency (such as the FDA) that"
Health Benefit	MDM	"... your medical provider gives you information from the manufacturer of the device that"
	-	"... new pacemaker that is Internet-connected...//you have the option of a non-connected pacemaker. However,..."
	Low	"... which will have a positive impact on your condition//... there is a slightly increased risk of negative health outcomes"
Initial Attack Likelihood	High	"... which will catch an adverse event (something bad happening) quickly//... a life-threatening event is much more likely to not be caught until it is too late"
	-	"because this device can connect to the Internet..."
	Low	"... there is a low likelihood the Internet-connected device that could be hacked..."
Post-disclosure Attack Likelihood	Unsure	"... they are unsure if the Internet-connected device that could be hacked..."
	-	"... a cybersecurity vulnerability has been discovered in the device and that..."
	Low	"... there is a low likelihood the Internet-connected device"
	Unsure	"... they are unsure if the Internet-connected device"
Impact of Attack	High	"... there is a high likelihood the Internet-connected device"
	-	"... If the device was hacked, it could lead to..."
	Data privacy	"... your private patient data being stolen"
	Physical safety	"... a safety issue, meaning you could be physically hurt"

Table 1: Text used in the survey (see Appendix A) for each of the conditions. We do not include the verbose text here for brevity, but the verbose version can be found in Appendix A.

data privacy issue, the disclosure also indicated the impact as data privacy. Therefore, we only discuss differences between the initial information and disclosure conditions here.

The first change we made was to vary the **Verbosity** of the disclosure. For participants in the FDA and MDM conditions, the disclosure was presented with either *Standard* verbosity, which included a minimal description of the vulnerability, or as a more detailed *Verbose* letter. The standard descriptions included the same type of information included in the initial description, but indicated that a vulnerability had now been found. Conversely, the verbose disclosures also indicated things that the patient could do to protect themselves, as well as more specific details about what could physically happen to them if the device were compromised. The verbose disclosures were developed to mimic real-world examples of patient communications by the FDA and MDMs, and text was taken directly from specific disclosures [25, 41] to enhance scenario realism. The verbose condition was only assigned to participants who were informed about a potential safety impact, as the FDA (and MDMs) only issues these disclosures in high-risk cases, not in cases involving data privacy concerns [16]. While this creates a partial confounding in our experimental design (i.e., we do not have a *Verbose* condition for data privacy impacts), this was a deliberate

choice to prioritize ecological validity. Presenting a participant with a multi-page, high-urgency regulatory warning for a low-risk data privacy issue would represent an unrealistic scenario that does not reflect current regulatory practices. We therefore restricted the *Verbose* condition to scenarios where such communication is plausible. For participants in the provider condition, only a single, standard level of verbosity was used.

Second, we reported a potentially new **Attack Likelihood** in the disclosure and added the possibility that an attack could be indicated at *High* likelihood. At this phase of the study, participants were told the updated likelihood of attack after the vulnerability was identified was either *Unsure*, *Low*, or *High*. We omitted the *High* condition for the initial information because the FDA would never approve a device with a high attack likelihood, but this restriction no longer applies. That is, this is a new vulnerability that was not known at the time the device went to market, so the FDA has no control over the possible level of impact.

Note that the likelihood of an attack in the disclosure was not assigned fully at random. While participants assigned *Low* for the initial description were randomly assigned any one of the three likelihood levels for the attack likelihood during disclosure, this was not true for participants initially assigned *Unsure*. Those assigned to

Unsure would either be *Unsure* or *High*, as it would be unrealistic for the likelihood to decrease after a vulnerability has been identified and disclosed.

3.3 Recruitment

We used Prolific to recruit participants, specifically those in the United States who speak English and are 18 years or older. To enhance the generalisability of our findings, we utilised Prolific’s census-balancing feature to obtain a sample representative of the US population based on age, sex, and race. Our target sample size was selected to ensure that we had at least 30 participants per condition, we did not conduct any power analysis to determine the sample size as we did not have prior data to estimate the effect sizes for the variables we were testing [28, 39]. We provide details of our population in Section 4.

3.4 Pilot

We piloted the survey with ten people who had limited experience with healthcare. Pilot participants were asked either to “think aloud” while a researcher observed and asked probing questions or completed the survey to ensure question understandability and gauge the average completion time. Based on the pilots, we clarified the phrasing of the questions to remove jargon and specify terms, such as Primary Care Provider (PCP).

3.5 Data Analysis

To answer our research questions, we employed a Bayesian modeling approach. This framework is particularly well-suited for our study, as it aligns with both the principles of modern applied statistics [29] and the goals of researcher-centered design of statistics in HCI [43]. Bayesian inference offers a direct and intuitive approach to quantifying uncertainty. The result of a Bayesian analysis is the full joint posterior distribution of all model parameters, representing our complete knowledge about the parameters conditional on the data and the model [29, Ch. 1]. This enables us to make direct, probabilistic statements about our research questions—for example, quantifying the likelihood that one design is better than another. This aligns with the goals of HCI research by providing richer, more nuanced answers about the magnitude and certainty of effects [43]. Additionally, the framework offers the flexibility to build multi-level models accurately reflecting the data’s structure [29, Ch. 5]. Our confidence measures were taken from each participant twice, meaning these observations are not independent. Our multilevel model estimates the behavior of each participant individually, while simultaneously using the overall patterns from the entire group of participants to inform and moderate these individual estimates.

Our primary dependent variables (given in Table 2), i.e., device choice, likelihood to disable connectivity, and confidence ratings, were measured on five-point Likert scales. To model these outcomes, we used four Bayesian cumulative ordinal regression models [15, 29]. Cumulative models are used for ordered categorical outcomes; in our case, they are used to estimate the probability of a participant’s response falling at or above a certain level on the Likert scale (e.g., the probability of choosing “Probably” or “Definitely” the internet-connected device). Our models for the initial

device choice (*Initial Model*) and the subsequent decision to disable connectivity (*Switch Model*) included all experimental factors and their two-way interactions, along with participant-level covariates for security attitudes (SA-6) and physician trust (WFPT), as shown in Table 2. The confidence models (*Security Confidence* and *Clinical Confidence*) additionally included a participant-level random effect, which explicitly models the baseline confidence level for each individual participant, thereby accounting for the two repeated measurements per person. For all models, we used weakly informative, regularizing priors to prevent overfitting and improve convergence [29, p. 124]. For all population-level effects (regression coefficients), we specified a $\mathcal{N}(0, 1.5)$ prior, which leans the model towards moderate effect sizes and, in the context of our model, represents a shift from one Likert response to another (e.g., a participant switching from ‘Definitely’ to ‘Probably’). We also performed a sensitivity analysis using tighter priors (e.g., $\mathcal{N}(0, 0.5)$), which confirmed that stronger regularization did not alter our substantive conclusions regarding the presence or absence of effects. We estimated each model using four Markov chains with 4,000 iterations each. We assessed model convergence by ensuring the \hat{R} statistic was equal to 1.00 for all parameters, which indicates the model converged across all the Markov chains, meaning there is no meaningful difference in the variance between them. We did not observe any noticeable effects (i.e., estimates around 0 with narrow credible intervals) for race, gender, or whether the participant was Latino; therefore, we excluded these variables from our final models. While we did not observe an apparent effect of age, we included it in our model, given that older participants are more likely to interact with the healthcare system. In contrast, younger participants are more likely to be familiar with technology.

Qualitative data analysis. For our qualitative data, we used an iterative open-coding approach [71]. Two researchers reviewed 50 responses and developed the initial codebook inductively, allowing themes to arise from the data. The codebook was then discussed collaboratively amongst the full research team. The two researchers then independently coded responses in rounds of 50, updating the codebook incrementally. After each round, the two researchers met, assessed inter-rater reliability (IRR) using Krippendorff’s alpha [47], and resolved any disagreements. Krippendorff’s alpha was used to assess IRR, as it provides a conservative metric that accounts for chance agreement [47]. Strong agreement ($\alpha \geq 0.8$) was achieved after four rounds (200 responses). One researcher coded the remaining 2416 responses. The two researchers performed axial coding to identify between-code relationships and identify higher-level themes in responses [20, pg. 123-142]. The final codebook is listed in Appendix B.1

To statistically test for associations between our experimental conditions and the themes that emerged from our qualitative analysis, we employed Pearson’s Chi-square test (χ^2) [57]. After the coding process, the presence or absence of each code for a given participant was treated as a categorical variable. The χ^2 test allowed us to determine whether the frequency of a specific theme was contingent upon the assigned experimental condition. While our primary Bayesian models explain the drivers of participants’ ordinal choices, these tests provide targeted statistical support for

Variable	Description	Initial Model	Switch Model	Security Confidence	Clinical Confidence
<i>Dependent Variables</i>					
Device Choice	5-point ordinal scale (Connected vs. Non-Connected)	X			
Likelihood to Disable	5-point ordinal scale (Unlikely vs. Likely)		X		
Security Confidence	5-point ordinal scale (Not at all confident to Extremely confident)			X	
Clinical Confidence	5-point ordinal scale (Not at all confident to Extremely confident)				X
<i>Experimental Factors</i>					
Device Choice	5-point ordinal scale (Connected vs. Non-Connected)		X		
Notice Type	Categorical: Provider, Government (Verbose, Non-Verbose), or Manufacturer (Verbose, Non-Verbose). Base Case: Provider.	X	X	X	X
Health Benefit	Categorical: High or Low clinical benefit. Base Case: Low.	X	X	X	X
Initial Attack Likelihood	Categorical: Low or Unsure likelihood of hack. Base Case: Low.	X	X	X	X
Post-disclosure Attack Likelihood	Categorical: Low, Unsure, or High likelihood of hack. Base Case: Low.	X	X	X	X
Impact of Attack	Categorical: Data privacy or Physical safety. Base Case: Data privacy.	X	X	X	X
Time	Categorical: Pre-disclosure or Post-disclosure. Base Case: Pre-disclosure.			X	X
<i>Measured Covariates</i>					
Physician Trust	Continuous score from Wake Forest Physician Trust Scale	X	X	X	X
Security Attitude	Continuous score from SA-6 Scale	X	X	X	X
<i>Demographics</i>					
Age	Categorical age brackets	X	X	X	X
Race	Categorical racial identity (<i>Not used in final model</i>)	X	X	X	X
Gender	Categorical gender identity (<i>Not used in final model</i>)	X	X	X	X
Education	Categorical level of education	X	X	X	X

Table 2: Variables Included in Ordinal Regression Models

the patterns in the qualitative data that help explain why those choices were made.

3.6 Ethics

This study was reviewed and approved by the lead author’s organization’s ethics review board. All participants were presented with a detailed consent form before beginning the study and were required to provide consent before participating. Participants were informed that their participation was voluntary and they could withdraw at any time without penalty. Participation through Prolific is provided pseudonymously, with only the participant’s Prolific ID used to identify the response. We did not request further identifying information. Responses were stored in an encrypted, password-protected cloud service accessible only to members of the research team. Participants were compensated \$2.50 for their time.

3.7 Limitations

Our study has several limitations that should be considered when interpreting the results, and which point towards avenues for future research.

Our methodology relies on a scenario-based survey that differs from the real world. While the vignettes were designed in close consultation with clinical experts to reflect reality, hypothetical decisions made in a low-stakes online environment may not perfectly mirror the complex, high-stakes decisions patients face when confronted with a real health diagnosis. The emotional and social pressures of a clinical setting are difficult to simulate; however, introducing this type of experimentation in a real clinical setting introduces significant health risks for the patient. Therefore, we believed this vignette-based approach was ethically important as a first step to assess initial results before testing in a real-world setting. Further, we ensured our sample included participants with

exposure to medical devices so that experience is captured in our data.

Additionally, the traditional informed consent process involves an interactive dialogue between a patient and a healthcare provider. Our survey, by its nature, presents information in a static, one-way format. It cannot capture the nuances of a real-time conversation, where patients can ask clarifying questions and providers can build trust and tailor their communication style. However, this type of dynamic conversation would introduce between-participant variability, limiting our ability to identify the impact of specific changes in information presentation. We attempted to emulate the variance in information by varying message verbosity. We also included all relevant information in the verbose condition to assess an upper-bound approximation on the amount of information a patient might receive.

Additionally, we focus on a single medical device type, a pacemaker. Patient attitudes and risk-benefit calculations may differ from those associated with other types of connected medical devices, such as insulin pumps, neurostimulators, or diagnostic wearables, which have distinct functions and risk profiles. However, we chose pacemakers due to their broad familiarity [79] and the existence of prior examples of vulnerabilities [36]. For these reasons, we believe pacemakers provide a useful general stand-in for medical technologies, providing baseline perceptions of medical device security, especially as our participants are not experts in medical systems.

Although our sample was recruited via Prolific to be census-balanced for age, sex, and race in the United States, Prolific users are often more likely to be more educated and more knowledgeable about security and privacy [73], which may impact generalizability. Additionally, while online participants lack the immediate 'lived authority' of patients in a clinical setting, 51.8% of our sample reported prior exposure to medical devices (including 18.1% as patients or caregivers). To account for these factors, we controlled for education, security attitudes, and prior medical device experience in our regressions.

Our study design involved partial confounding of experimental variables, specifically regarding disclosure verbosity, which was only manipulated in safety-impact conditions. As noted in Section 3.2, this was a design choice to maintain ecological validity. However, it limits our ability to statistically isolate the effect of verbosity independent of risk severity in low-risk scenarios. Additionally, the sequential nature of our vignette (Initial Choice followed by Vulnerability Disclosure) inherently introduces ordering effects. We specifically designed the study this way to measure the anchoring effect of the initial decision, but this means our post-disclosure results are conditional on the initial choice context.

Finally, our survey did not include quantitative manipulation checks to formally verify that participants perceived, for example, 'high' vs. 'low' benefit distinctly. However, our qualitative analysis of open-ended responses indicates that participants actively engaged with and understood the trade-offs presented in the vignettes, suggesting the manipulations were successful in practice. For example, one participant mentioned both the likelihood and the advantages of the connected device, "I think the low likelihood of it being hacked far is outweighed by the benefits of it still being connected and being monitored."

As these limitations apply across all conditions, we consider between-condition comparisons primarily.

4 Participants

We recruited an initial sample of 3,254 participants from the United States via Prolific. We removed participants who did not complete the full survey ($N = 165$), failed one or more attention checks ($N = 414$), or completed the survey in under three minutes, which was more than one standard deviation from the mean and unreasonably quick given the survey's level of detail ($N = 9$). After exclusions, our final sample for analysis consisted of 2,666 participants (at least 30 per condition). The median time to completion was 9.8 minutes (Mean = 11.7, SD = 7.3).

The demographic characteristics of our final sample are detailed in Table 3. Our participants' gender, age, race, and ethnicity were all similar to the 2020 US Census [14]. Our participants were more educated than the US population, i.e., 57.1% had at least a Bachelor's degree. Participants' average SA6 score was 3.8, which is similar to the average score from a prior US Census-representative sample [26]. Participants' average WFPT score (36.3) was also similar to a prior census-representative samples [13, 35].

Most of our participants ($N = 1383$, 51.8%) had some exposure to medical devices. Many reported having a family member or friend who was prescribed a medical device ($N = 899$, 33.7%) or indicated they were the caregiver for someone with a medical device ($N = 323$, 12.1%) or were prescribed a medical device themselves ($N = 161$, 6%). Further, almost all participants had some exposure to healthcare providers ($N = 2633$, 98.8%) and a majority had experience with specialists associated with prescribing medical devices ($N = 1533$, 57.5%).

5 Result

We present our findings in order of the research questions they address. Most of our key findings are taken from our Bayesian ordinal regression models for initial device choice (Table 6), likelihood to disable connectivity after vulnerability disclosure (Table 7), and clinical and security information confidence (Tables 9, 10). For brevity and ease of understanding, we visualize the main effects for device choice in Figure 3, the likelihood to disconnect the device in Figure 4, and the impact of the vulnerability disclosure on participants' confidence in Figure 6. We provide the full model details in Appendix B.

We utilized a Bayesian approach for two primary reasons. First, our dependent variables are ordinal Likert scales. By employing a Bayesian cumulative ordinal model, we explicitly respect the ordered categorical structure of the data, mapping the probability of responses falling into higher or lower categories without assuming the intervals between them are equal. Second, our data has a multi-level structure (repeated measures nested within participants). The Bayesian framework allows us to model this hierarchy intuitively and, crucially, allows us to quantify the uncertainty of our estimates directly through Credible Intervals, rather than relying on binary significance testing. We discuss this further in Section 3.5.

Using a Bayesian cumulative ordinal model is similar to a frequentist ordinal regression, whereby the β coefficients that are

Characteristic	Category	N (%)
Gender	Female	1328 (49.8%)
	Male	1288 (48.3%)
	Non-Binary / Other	37 (1.4%)
	Prefer not to answer	13 (0.5%)
Race	White	1821 (68.3%)
	Black or African American	343 (12.9%)
	Asian or Pacific Islander	182 (6.8%)
	Mixed or Other	296 (11.1%)
	Prefer not to answer	24 (0.9%)
Latino/a	Yes	307 (11.5%)
	No	2359 (88.5%)
Age	18–24	298 (11.2%)
	25–34	505 (18.9%)
	35–44	492 (18.5%)
	45–54	442 (16.6%)
	55–64	599 (22.5%)
	65–74	280 (10.5%)
	75+	44 (1.6%)
Prefer not to answer	6 (0.2%)	

Characteristic	Category	N (%)
Education	High School or Less	341 (12.8%)
	Some College / Associate's	790 (29.6%)
	Bachelor's Degree	988 (37.1%)
	Advanced Degree	534 (20.0%)
	Prefer not to answer	10 (0.5%)

Variable	Mean (Median)	SD	Min, Max	5%, 95%
SA6	3.8 (4)	0.8	1, 5	2.3, 5
WFPT	36.3 (38)	8.6	10, 50	21, 50

Variable	Value	N (%)
Experience with medical devices	No experience	1262 (47.3%)
	Know a friend or family member with device	899 (33.7%)
	Acted as a caregiver	323 (12.1%)
	Participant has device	161 (6.0%)
	No response	21 (0.8%)
Experience with healthcare	Has not seen PCP	33 (1.2%)
	Has seen a PCP	1100 (41.3%)
	Has seen a specialist	1533 (57.5%)

Table 3: The table on the left details participant demographics. The top right table shows distributions for the security attitudes (SA6) and physician trust (WFPT) scales. The bottom-right table details participants' experiences with medical devices and healthcare providers. The total participant count for each table that was used to calculate percentages is $N = 2,666$.

calculated represent the outcome's log-odds. For categorical variables, such as the notice party, the coefficients are the increase (or decrease) in the log-odds compared to the base case. For example, in the initial device choice model, we chose the healthcare provider as the base case, and the β for the FDA was 0.25. This means participants who were given information sourced from the FDA increased the odds of choosing the connected device by a factor of 1.28 (i.e., $e^{0.25} \approx 1.28$), all other variables held constant. Conversely, since the MDM had a $\beta = -0.25$, participants who received information sourced from the MDM were 1.28× less likely, on average, to choose the connected device than those who received information from their healthcare provider. As our models are Bayesian, we do not use p-values. Instead, we quantify uncertainty using a 95% credible interval (CrI). A 95% CrI can be interpreted directly: there is a 95% probability that the true value of β lies within that range, given our data and model. Our practical rule for interpretation is simple: if the CrI contains zero, we cannot be certain of the effect's direction, so we conclude that we did not observe a discernable effect. For example, the CrI for the FDA's effect relative to the provider was [-0.11, 0.61] and the MDM was [-0.61, 0.12], their ranges of plausible values both include a negative and a positive effect, so we report that we do not observe a discernable effect for either. Conversely, the CrI for an "unsure" risk notice was [-0.77, -0.07], the entire interval is below zero, giving us strong evidence that this framing had a negative effect compared to the baseline (i.e., telling participants

the security risk was "low"). We plot the β estimates along with the CrI for our regression factors in Figure 3 and Figure 4 and indicate factors with strong effects by coloring the lines solid blue.

5.1 Initial Device Choice & Drivers of Baseline Confidence (RQ1)

In this section, we discuss the relative influence of how and what security and clinical risks are presented (e.g., security impact, information source) on patients' initial decision whether to use a connected medical device (Fig 1.C), controlling for patients' internal characteristics (e.g., medical experience and trust in their physician). Figure 3 summarizes participants' initial device choices by experimental condition, and Table 11 presents the most common themes in open-ended participant responses that explain their device choices. To provide high-level context, participants showed a strong initial preference for the connected pacemaker, with 66% ($N = 1,746$) reporting that they would select the connected device.

Participants told that an exploit could impact safety or that the likelihood of an exploit was unclear were less likely to use the connected device. We found that the most influential factor in the security information provided to participants during the initial decision was the security impact and the likelihood of a vulnerability. First, participants who were informed that potential hacks could result in medical safety issues were less likely to choose

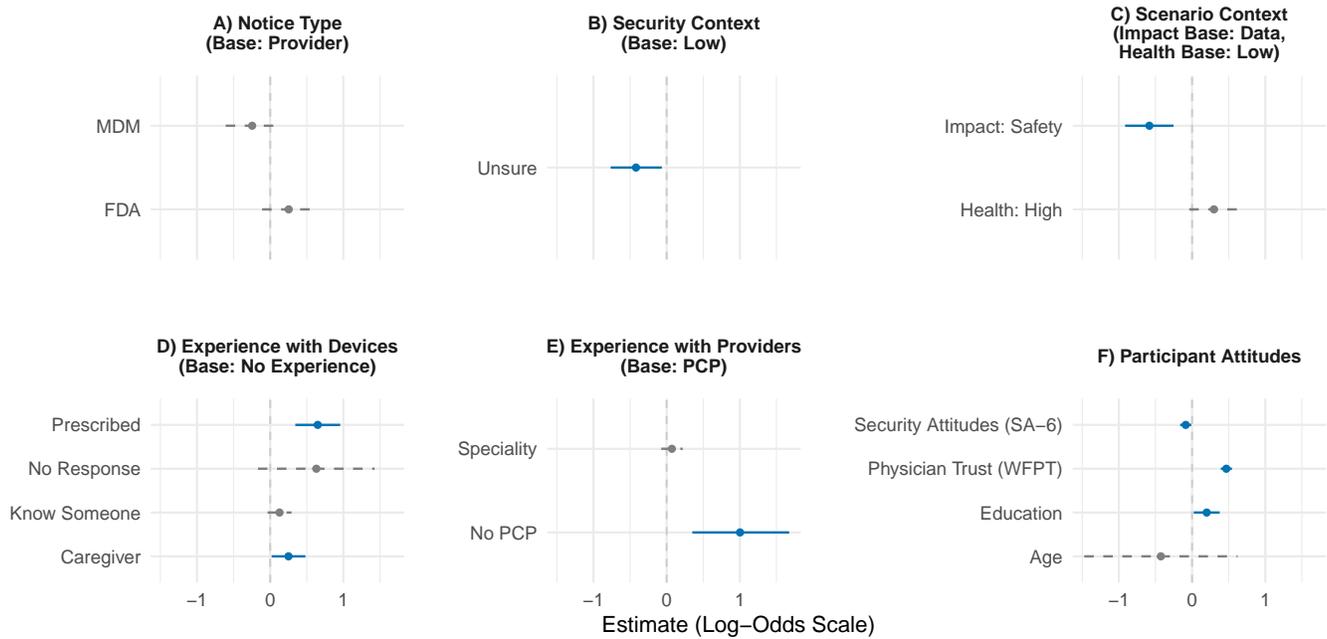


Figure 3: Coefficient plots visualizing the main effects from the Bayesian cumulative ordinal regression model predicting initial device choice. The full regression results can be found in Table 6. Each point represents the median of the posterior distribution for the β coefficient (log-odds), while the horizontal lines span the 95% CrI. Values greater than zero are more likely to choose the connected device compared to the base case, while those less than one are less likely to choose it. Predictors coloured in solid blue have credible intervals that do not overlap with zero, indicating a clear, discernible effect on patient choice. Grey dashed predictors have credible intervals that contain zero, meaning we cannot be certain of their effect’s direction. Each sub-figure represents a feature in the regression, with the base case given under the feature name. The linear features are grouped together in sub-figure F for brevity.

the connected medical device (59.9% said they would at least “Probably” select the connected device) compared to those told hacks could lead to data breaches (74.5%). This difference is clear in our Bayesian regression model ($\beta = -0.58$; 95% CrI: $[-0.91, -0.25]$). Further, participants told that an exploit could cause safety issues were more likely to state that the risks outweighed the medical benefits than those who said the impact was a data breach in their open-ended responses (10.6% vs. 7.4%, $\chi^2 = 7.13$, $p = .008$). For example, one participant said, “I am confident that I know the purposes that pacemakers serve. I trust the sources, but I feel uneasy about the fact that if it gets hacked, it could be a matter of death for [me as the patient].”

Similarly, participants who were told the information provider was “unsure” about the likelihood of an exploit were less likely to choose the connected medical device (59.8% said they would at least “Probably” select the connected device) compared to those told the likelihood of an exploit was “low” (69.3%). This difference is also clear in our Bayesian regression model ($\beta = -0.42$, 95% CrI: $[-0.77, -0.07]$). This seems to suggest participants believe the risk is likely more than low if it is not given.

As we would expect, participants assigned to the high health condition were more likely to choose the connected device given the higher health benefits (69.5% chose the connected device) compared

to the low, or minimal, health benefits (61.3%). While this was not a clear effect, we do see that there is a difference in the regression model, suggesting a relationship, but it is not as strong as the other attributes ($\beta = 0.30$; 95% CrI: $[-0.04, 0.64]$).

Information source has limited impact. In stark contrast to risk framing, the source of the communication—, i.e., whether the information came from a provider, the FDA, or the device manufacturer—had no discernible effect on the initial decision (Figure 3.A). 67.1% who were given the initial notice from the healthcare provider said they would at least “Probably” select the connected device, compared to 65.3% for the FDA and 64.6% for the MDM. We did not observe a discernible difference between the healthcare provider and the FDA ($\beta = 0.25$, 95% CrI: $[-0.11, 0.61]$) or MDM ($\beta = -0.25$, 95% CrI: $[-0.61, 0.11]$) conditions.

Participants wanted addition information about cybersecurity risks in the initial information. When participants were asked what specific information that they wanted in addition to the initial information given, 77.6% who wanted to know how they could minimize the potential for hackers to access the device, 75.4% who wanted to know what could happen in the worst-case scenario if the pacemaker was hacked, 56.6% who wanted to know if

other groups had assessed the security of the device, and 47.6% who wanted to know more about how the assessment was made, only 7.4% explicitly stated that they did not require additional information. Even though most participants requested more information, the majority (66.8% of the 2,471 who requested more information) of those still indicated they would at least “Probably” select the connected device. This reinforces the low impact of the externally provided information on participants’ decision-making. However, it leaves participants in an uncomfortable position where they feel less confident in their decisions, but still choose the connected device.

We observed evidence of this in participants’ open-ended responses, where some explicitly stated that they found the information provided to be clear ($N = 793$, 29.7%). In contrast, other participants explicitly said the information was incomplete ($N = 407$, 15.3%). Those who felt there was not sufficient information expressed a desire for more technical details on the device ($N = 324$, 12.2%). For example, one participant said, “I understand that the Internet-connected pacemaker allows for remote monitoring...[but] I haven’t seen specific technical details on how these devices are protected or how cyberattacks are handled.” Another participant talked about what they could do to stop hackers, “I have never heard of pacemakers being hacked but with growing technology, I am sure it is possible. The idea while a low chance, is still terrifying to me and if I were to choose it, I’d like to know more about it, the risks, and what I could do to stop hackers.”

Physician trust led to a higher likelihood of adopting the connected device and higher confidence. In contrast to the impact of the external factors we assessed (i.e., provided information), we observed that participants’ internal characteristics had a bigger effect on their decision-making. Specifically, a patient’s baseline trust in their provider was a powerful predictor of their willingness to accept a connected device. 76.7% of participants with the highest 25% of WFPT (physician trust) scores indicated they would at least “Probably” select the connected device, whereas of those with the lowest 25% of WFPT scores, only 52.0% chose the connected device. Participants discussed their trust in their provider in their qualitative responses as being a reason to go with the connected device. For example, “I am confident because I trust my care provider and if they explain to me why I need it and the explanation makes sense then I would be confident plus knowing that they would be able to monitor the device remotely is great.”

Interestingly, we also observed that people who reported not having a primary care provider (PCP) were more likely to indicate at least “Probably” selecting the connected device, i.e., 78.8% compared to 63.9% who reported having a PCP. This was a clear difference in our regression ($\beta = 1.00$, 95% CrI: [0.35, 1.67]). At first glance, this suggests contradictory results; however, these experiences are orthogonal, as it is possible to trust physicians without having a PCP. Also, this group comprised a considerably small proportion of participants ($N = 33$, 1.2%), and therefore, we place less emphasis on this effect.

Prior experience with medical devices increases connected device acceptance. Familiarity with medical devices, directly or indirectly through personal relations, was a strong predictor of

choosing the connected device. Specifically, participants who had previously been prescribed a medical device were much more likely to opt for the connected pacemaker (78.3% indicated at least “Probably” selecting the connected device) than those with no medical device experience (62.7%). This difference was clear in our regression ($\beta = 0.65$, 95% CrI: [0.34, 0.96]). Many participants who reported being caregivers for someone with a medical device also indicated at least “Probably” selecting the connected device (67.5%), and this was clearly more likely than those with no experience ($\beta = 0.25$, 95% CrI: [0.02, 0.48]). Finally, we also observed that more participants who knew someone with a medical device indicated at least “Probably” selecting the connected device (66.1%). However, this difference was not quite clear in the regression as the 95% CrI slightly overlapped 0 ($\beta = 0.13$, 95% CrI: [-0.04, 0.29]).

More security-concerned participants were slightly less likely to select the connected device, whereas more educated participants chose the connected device. There was a small negative correlation of the participant’s device choice with their SA6 score ($\beta = -0.09$, 95% CrI: [-0.16, -0.01]), i.e., the more security concerned a participant was, the less likely they were to choose the connected device. We found that participants’ education affected their device choice, such that the more educated a participant was, the more likely they were to choose the internet-connected device ($\beta = 0.20$, 95% CrI: [0.02, 0.38]). Participants showed a level of understanding about trade-offs, and those who were more security-conscious clearly articulated their decision-making, such as one participant saying “I chose the non-connected pacemaker because even though it requires more in-person checkups, I value the reduced risk of cybersecurity threats. The possibility of being hacked, even if small, made me uncomfortable with an Internet-connected device inside my body. Since both options are FDA-approved and effective for treating heart block, I felt safe choosing the one that gave me more control over my data and minimized digital risks.”

No clear effects from participant age. As shown in Figure 3.F, there was no discernible effect from a participant’s age on their initial device choice ($\beta = -0.43$, 95% CrI: [-1.47, 0.62]). In fact, the credible interval was substantial, suggesting any potential effect is highly variable across the population and is overwhelmed by more influential factors like physician trust and security impact.

Summary of Initial Choice Drivers (RQ1).

- **Physician trust trumps source:** Internal attributes, specifically baseline trust in the physician, are far stronger predictors of device adoption than the source of the security notice (FDA vs. MDM vs. Provider).
- **Safety threshold:** Participants exhibit a clear behavioral shift when risks escalate from data privacy to physical safety, identifying a distinct threshold for risk tolerance.
- **Ambiguity aversion:** When the likelihood of a hack is presented as unsure, participants treat the risk similarly to a higher-threat scenario.

5.2 Decision After Vulnerability Disclosure (RQ2)

Our second research question explored participants’ decision whether to disable device connectivity following a vulnerability

disclosure. Figure 4 summarizes participants' likelihood of disconnecting a connected device after receiving a vulnerability disclosure by experimental condition. We found this choice was overwhelmingly anchored to a participant's initial sentiment, a finding that was far more predictive than the specifics of the disclosure notice itself. The proportion of participants likely to disconnect the device was nearly equal to those who were unlikely to do so (44.9% vs. 44.5%). To ensure our findings were not merely driven by participants who initially rejected the connected device, we conducted a sub-group analysis by re-fitting our model on only those participants who had initially chosen the connected option ($N = 1,738$). The results of this analysis were substantively identical to our main model: the credible intervals for all major predictors showed similar magnitudes and directionality. This demonstrates that the factors driving a patient's decision to disconnect are robust and not simply an artifact of the initial anchoring effect. Therefore, for clarity, we present the full model results in our analysis, and the subset model is provided in Table 8.

Initial choice creates a powerful anchor. The model results (Figure 4) show that a participant's initial device choice was the largest predictor of their later decision. The effect of moving one level towards the non-connected device on the initial 5-point scale increased the log-odds of wanting to disconnect by -3.26 (95% CrI: $[-3.02, -3.52]$). Of the 746 participants who initially chose "Definitely" or "Probably" the non-connected device, 82.1% ($N = 612$) later selected "Extremely" or "Somewhat" likely to disconnect. Conversely, of the 1,746 participants who chose "Definitely" or "Probably" the connected device, 61.4% ($N = 1,073$) were "Extremely" or "Somewhat" unlikely to disconnect. This persistent adherence to the initial choice underscores the strong anchoring effect that precedes the vulnerability disclosure.

Safety concerns act as a tipping point towards disconnecting the device. While anchoring was the dominant effect, the security impact continued to exert clear influence, as it did on participants' initial choice (See Section 5.1). Some participants ($N = 315, 31.2%$) who were initially warned that any potential security impact would be a "data breach" indicated they were "Extremely" or "Somewhat" likely to disconnect after the vulnerability was disclosed. However, participants initially told the security impact would be a "safety" were substantially more likely to disconnect ($N = 872, 52.7%$) and this difference was clear in our regression ($\beta 0.96, 95\% \text{ CrI: } [0.52, 1.40]$).

As one participant explained, "If the private data is just the same type of data I can lose now in a data breach, I would be willing to risk that... However, if the hackers can access the actual *workings* of the pacemaker, that would be a different situation." Similarly, another participant said, "I would consider disabling the internet connection because the vulnerability is described as high risk and could lead to a physical safety issue... even though remote monitoring has life saving benefits". The "safety" framing dramatically increased the likelihood participants would describe the risk as "unacceptable" when explaining their reasoning. The proportion of participants using this code more than doubled from 19.4% in the "data" condition to 38.8% in the "safety" condition ($\chi^2 = 107.93, p < .001$). This highlights a critical distinction in

patient risk perception: risks to data are tolerable, while risks to bodily integrity are not.

Participants questioned why an attack would occur or whether it would harm them. Some participants seemed to downplay the likelihood of an attack, questioning an attacker's motivation or explicitly saying that no one would want their data ($N = 312$). One participant said, "I just don't know what benefits a hacker would receive from getting patient information from a pacemaker. If this was my bank information at risk, I would absolutely stop using it but this is patient information in regards to my heart, and to me the benefits outweigh the risks" and another focused on who attackers might be, "While it may be possible for the device to be hacked, the hacker would still need a motive to cause such problems, especially given the consequences that would befall such a hacker if caught. Most hacking is either profit motivated, or state sponsored". Others expressed little care about being a target or having their data stolen ($N = 327$), while some explicitly mentioned that they had previously had data exposed, lessening their concern ($N = 80$). For example, "I chose what I felt was safer. The greater of the good. I'm personally not too worried about my personal information being leaked as it's leaked already and I think my health would be more important." This is something that was observed in a prior interview study [23], and also the FDA has suggested more should be done on highlighting when a vulnerability might affect multiple patients at once, rather than a single device having to be targeted [38].

Verbosity and who provides the notice appeared to have no clear effects. In contrast to the stark effect of the security impact, we did not observe any clear difference in participants' decisions based on the specifics of the disclosure notice, i.e., who it came from and the level of detail. The proportion of participants who were likely (either "Somewhat" or "Extremely") to disconnect was lowest in the verbose conditions (51.8% for FDA Verbose and 49.8% for MDM Verbose) compared to the standard notice conditions (52.7% for the provider, 51.4% for the FDA, and 57.3% for the MDM). However, we did not observe clear differences in our regression when controlling for other factors (shown in Figure 4.A). The credible intervals for the main effects of receiving a standard notice from the FDA ($\beta = 0.20; 95\% \text{ CrI: } [-0.29, 0.70]$) or the MDM ($\beta = 0.27; 95\% \text{ CrI: } [-0.22, 0.77]$) compared to the provider both comfortably included zero, indicating no discernible difference in their direct influence. We also did not observe a noticeable difference for the verbose notices for both the FDA ($\beta = -0.25; 95\% \text{ CrI: } [-2.33, 1.86]$) and MDM ($\beta = -0.06; 95\% \text{ CrI: } [-2.17, 2.00]$).

While the main effects were uncertain, the model revealed a subtle interaction. Participants in the high health-benefit condition, i.e., the ability to catch a life-threatening event quickly, who also received a verbose disclosure from the FDA, were slightly more likely to report wanting to disconnect their device ($\beta = 0.53; 95\% \text{ CrI: } [0.00, 1.04]$). This finding suggests a potential backlash effect: the combination of high clinical stakes and a detailed, formal government warning may have amplified the perceived severity of the vulnerability, overriding the benefit of remote monitoring. Nonetheless, this subtle interaction does not change the primary finding for this research question: communication factors like source and

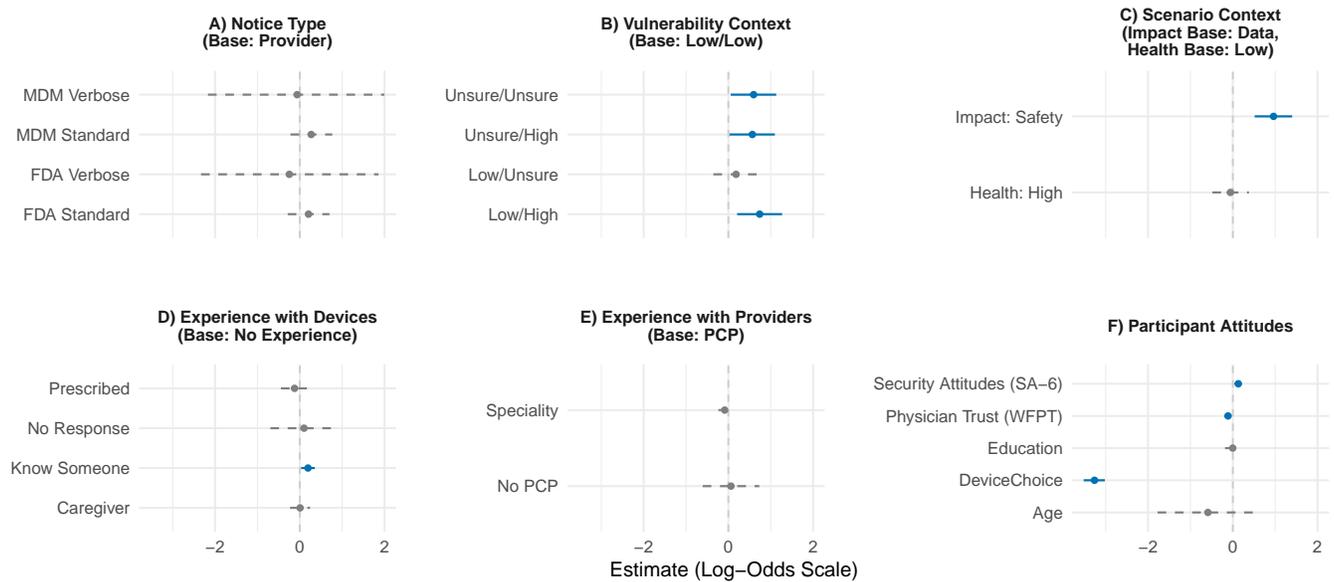


Figure 4: A coefficient plot illustrating the main effects from the Bayesian cumulative ordinal regression model predicting the likelihood of disconnecting the connected device. The full regression results can be found in Table 7. Each point represents the median of the posterior distribution for the β coefficient (log-odds), while the horizontal lines span the 95% CrI. Values greater than zero are more likely to disconnect the device compared to the base case (given below the factor name in each sub-figure), and those less than one are less likely to disconnect the device. Predictors coloured in solid blue have credible intervals that do not overlap with zero, indicating a clear, discernible effect on patient choice. Grey dashed predictors have credible intervals that contain zero, meaning we cannot be certain of their effect's direction.

verbosity were far less influential than the participant's anchored initial choice and the fundamental framing of risk as a threat to physical safety.

Indications that the risk of an exploit of the vulnerability was 'high' or had been 'low' and then was 'unsure' increased the likelihood to of disconnecting. The proportion of participants who were at least "Somewhat" likely to disconnect rose from 33.9% in the stable low-risk condition (i.e., participants were initial told the risk of exploitation was low and were again told the risk was low in the disclosure) to 52.1% when the likelihood of exploitation rose from low (initially) to high (at disclosure) and to 52.8% when the likelihood of exploitation rose from unsure (initially) to high (at disclosure). These differences are clear in our regression as, compared to the baseline where risk remained low throughout, participants who were told the risk was low and is now high (i.e., Low/High condition $\beta = 0.74$; 95% CrI: [0.21, 1.27]) and those who were told the provider was unsure and now believed the risk was high (i.e., Unsure/High condition $\beta = 0.57$; 95% CrI: [0.03, 1.10]) were substantially more likely to opt for disconnecting their device.

Continued uncertainty was also treated as a strong negative signal. A considerable portion of participants (45.2%) who were initially told the information provider was unsure about the risk of exploitation and continued to be unsure in the vulnerability disclosure, reported being at least "Somewhat" likely to disconnect, a

level approaching the conditions where risk was deemed high in disclosures. The model confirms this, showing these participants were clearly more likely to disconnect than those in the stable low-risk condition ($\beta = 0.60$; 95% CrI: [0.06, 1.13]). However, only 38.7% of participants who were first told the risk was low, then informed that there was an unsure risk of exploitation in the disclosure, reported being at least "Somewhat" likely to disconnect. There was no clear difference between this condition and the stable low-risk condition in the regression ($\beta = 0.18$, 95% CrI: [-0.35, 0.71]). This suggests that the initial risk priming has some effect on participants' reactions after disclosure, as they appear to default to a more skeptical view when given ambiguous information, but adopt an optimistic perspective when a specific indication is initially provided.

Patient trust and security posture, but not demographics, continued to influence decisions. The internal patient attributes that were predictive of the initial device choice continued to shape decision-making after the vulnerability disclosure. For every one standard deviation increase in WFPT scores (healthcare provider trust), the log-odds of wanting to disconnect decreased by 0.11 (95% CrI: [-0.19, -0.03]). This suggests that a strong patient-provider relationship can serve as a buffer against alarm when new risks emerge. Conversely, a higher SA6 predicted a greater likelihood of wanting to disconnect ($\beta = 0.13$; 95% CrI: [0.05, 0.21]), reinforcing that participants with stronger security attitudes remained more cautious throughout the scenario. For example, one participant

mentioned wanting to consult with their provider before deciding to switch, “I’m not sure that I would without first discussing the option with my physician.”

In contrast to the initial choice, most other prior experiences and demographic factors had no discernible effect on this subsequent decision. We did not observe a clear difference between participants who reported having a PCP and those reporting no PCP ($\beta = 0.06$; 95% CrI: [-0.61, 0.73]) or having seen medical specialists ($\beta = -0.09$; 95% CrI: [-0.24, 0.06]). Additionally, we did not observe a clear effect on participants’ disconnection decision based on education ($\beta = -0.00$; 95% CrI: [-0.18, 0.18]), or age ($\beta = -0.59$; 95% CrI: [-1.78, 0.62]). The only clear difference we observed for participants’ medical device experience was that participants who knew a friend or family member with a medical device were slightly more likely to indicate that they would disconnect the device after disclosure ($\beta = 0.19$; 95% CrI: [0.03, 0.36]). However, this makes sense in that while they have familiarity with devices, they may not understand the benefits of these devices on quality of life as readily as caregivers or patients who have been previously prescribed a medical device, and so are more willing to disconnect the device.

Summary of Vulnerability Response (RQ2).

- **Anchoring effects:** Initial device choice is the single strongest predictor of post-vulnerability behavior. Post-market disclosures serve more as a nudge for the already hesitant rather than a reset of the decision.
- **Safety as a tipping point:** While data privacy risks are often tolerated, risks to physical safety override the benefits of connectivity for a significant portion of participants.
- **The “Not a Target” mental model:** Qualitative responses reveal a persistent belief among patients that they are not high-value targets for attackers, which dampens the urgency of security warnings.

5.3 Confidence After Vulnerability Disclosure (RQ3)

Our final research question addressed how a vulnerability disclosure impacts patient confidence in both the security and clinical aspects of their device. Figure 5 summarizes participants’ confidence in the information about the security and clinical necessity of the medical devices presented by the experimental condition. Overall, confidence in the security information remained stable before and after disclosure (85.4% and 82.3% reported at least moderate confidence in device security, respectively). Additionally, confidence in the clinical information remained high both before and after the disclosure (91.9% and 92.6%, respectively, were at least moderately confident). This seems an obvious result, as we only provided updates to the security information in the disclosure, so we would not expect to see a change in clinical confidence. However, it is possible there could have been some leakage in confidence effects (i.e., a participant begins to distrust all information because some of it was proven incorrect). Luckily, we did not observe this leakage effect from vulnerability disclosure, as there was no clear difference in clinical confidence between responses before and after disclosure in the model ($\beta = 0.09$; 95% CrI: [-0.30, 0.51]).

Disclosures that do not indicate exploit likelihood decrease confidence in some cases. We identified a clear negative interaction in the clinical confidence model. For participants in the “Unsure/Unsure” vulnerability condition, receiving a verbose notice from the manufacturer was associated with a clear decrease in their confidence in the device’s clinical aspects ($\beta = -1.40$, 95% CrI: [-2.76, -0.06]). This suggests a potential backlash: when a patient is already operating in a state of high uncertainty, a long, detailed letter from the company that made the potentially flawed device may be perceived not as helpful transparency, but as corporate damage control, which in turn erodes trust in the device’s fundamental clinical function. We also observed this in our qualitative data; 61 participants mentioned distrusting the MDM after the disclosure. For example, one participant said “How do I know that the info is accurate? Companies have a vested interest in keeping product flaws quiet.”

We also observed some indications that the uncertainty of exploit likelihood in the disclosure had a marginally negative effect on security confidence. First, participants who were initially told the risk was low and then told it was unsure showed a slight decrease in their security confidence over time compared to the baseline change ($\beta = -0.32$; 95% CrI: [-0.68, 0.04]). We saw a similar effect for participants told the risk was unsure initially and remained unsure after disclosure ($\beta = -0.27$, 95% CrI: [-0.63, 0.09]). However, neither of these effects was observed to be clear, as they slightly overlap with zero.

Internal attributes, not external factors, drive baseline confidence. As with device choice, a patient’s baseline confidence was overwhelmingly predicted by their internal attributes (Figure 5.F) rather than the external communication factors (Figure 5.A-E). Physician trust (WFPT) and security attitudes (SA6) were powerful predictors of higher confidence in both the clinical and security domains. For every one standard deviation increase on the WFPT scale, the log-odds of being more confident increased by 0.91 for clinical (95% CrI: [0.75, 1.08]) and 0.69 for security information (95% CrI: [0.55, 0.84]). The effect for SA6 was similarly large. For every one standard deviation increase on the SA-6 scale, the log-odds of being more confident increased by 0.93 for clinical (95% CrI: [0.76, 1.10]) and 0.78 for security information (95% CrI: [0.63, 0.93]). This demonstrates that patients with high intrinsic trust and security engagement feel more confident in navigating complex health technology information. Conversely, and perhaps counter-intuitively, higher levels of education were associated with lower confidence in both the security ($\beta = -0.67$, 95% CrI: [-1.01, -0.32]) and clinical information ($\beta = -0.57$, 95% CrI: [-0.97, -0.19]), perhaps reflecting a greater awareness of unknown complexities.

Verbose disclosures increased confidence in some cases; the FDA improved security confidence, while MDM improved clinical confidence. While we did not observe any large changes in security or clinical confidence after disclosure, when comparing the different characteristics of disclosure conditions, we did observe some clear effects when a more verbose description of the vulnerability was presented. However, this was not uniform across providers and types of confidence. First, participants shown the verbose disclosures with information provided by the FDA

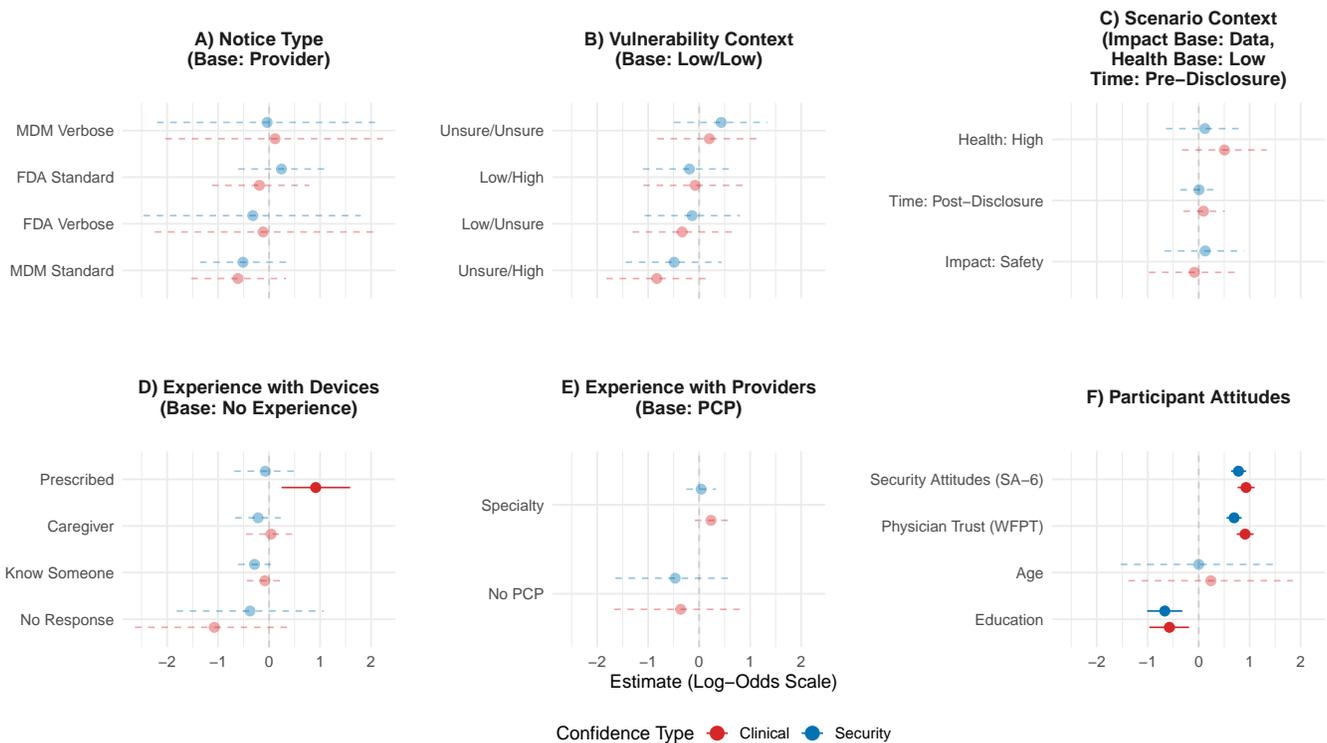


Figure 5: A coefficient plot visualizing and comparing the main effects from two Bayesian cumulative ordinal regression models predicting patient confidence in the information about security (blue) and clinical necessity (red). The full regression results can be found in Tables 9 and 10. Each point represents the median of the posterior distribution for the β coefficient (log-odds), while the horizontal lines span the 95% CrI. Values greater than zero indicate a higher level of confidence compared to the base case (shown below the feature name at the top of each sub-figure). Solid points and lines indicate a clear, discernible effect where the credible interval does not overlap with zero, while faded dashed points and lines represent uncertain effects.

reported clearly higher security confidence after disclosure compared to those shown a non-verbose disclosure and told that their healthcare provider was the information source. The proportion of participants who reported at least moderate confidence in device security rose from 83.6% to 84.9% after disclosure for participants shown the verbose disclosure from the FDA, compared to a change from 87.5% to 84.2% for participants shown standard disclosures from the healthcare provider (with a safety impact). This difference can be seen in the interaction between the time when participant security confidence was recorded (i.e., post-disclosure) and the disclosure verbosity and source ($\beta = 0.58$, 95% CrI: [0.15, 1.00]). We visualize this interaction result in Figure 6.

Similarly, we found a second clear effect for clinical confidence with verbose disclosures from the MDM. The proportion of participants who reported at least moderate confidence in the device's clinical necessity rose from 90.3% to 90.9% after disclosure for participants shown the verbose disclosure from the MDM, compared to a change from 92.6% to 92.8%. This difference can be seen in the interaction between the time when participants' clinical confidence was recorded (i.e., post-disclosure) and the disclosure verbosity and source ($\beta = 0.58$, 95% CrI: [0.15, 1.00]). Additionally, while the same

interaction effect for the verbose MDM disclosure did not indicate a clear difference in the security confidence regression ($\beta = 0.40$, 95% CrI: [-0.03, 0.82]), it only slightly crossed 0. Together, these results regarding disclosure verbosity suggest that detailed communication, particularly from the entity that made the device, can bolster both security and clinical trust. This finding demonstrates that transparency about security risks does not need to undermine a patient's trust in the clinical value of their care.

Participants' open-ended responses also supported this result. One participant, who was assigned the verbose disclosure from the FDA, wrote, "I feel very confident in the information about the pacemaker because the email that was sent to the patients was very detailed and it honestly reflected how significant a risk it actually is having the device connected to the Internet. The email also provided further information on how patients could reduce their risk by following some cybersecurity precautions." Comparatively, a participant who did not receive a verbose disclosure said they wanted more information that included the specifics that are being done to mitigate it, "I remain only moderately confident; while I understand a vulnerability exists with low risk of data theft, I still lack detailed information regarding the nature of the vulnerability, the

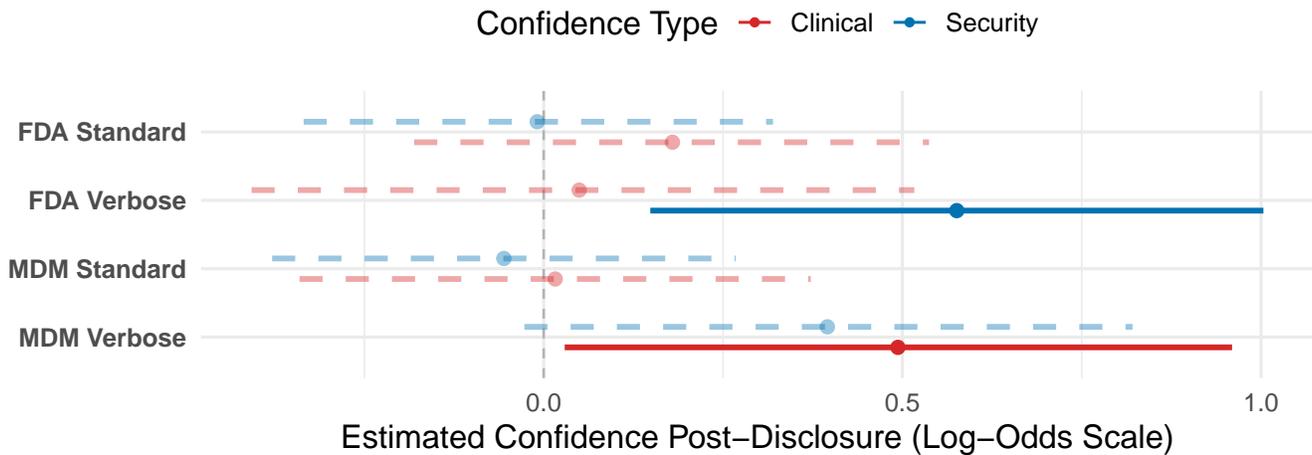


Figure 6: A coefficient plot visualizing the confidence following the vulnerability disclosure, comparing the effect of different notice types on Security confidence (blue) and Clinical confidence (red). The plot displays the posterior distributions for the key interaction terms from our models (full results in Tables 9 and 10). Each point represents the median of the posterior distribution for the β coefficient, while the horizontal lines span the 95% CrI. Values credibly greater than zero indicate that a specific notice type produced a larger increase in confidence compared to the baseline (a standard notice from the provider).

specific measures being taken to address it.” Further, participants shown the verbose disclosures were statistically significantly less likely than those shown non-verbose disclosures to describe the disclosure as “incomplete” when asked to explain their indicated level of confidence (12.9% vs. 18.8%, respectively; $\chi^2 = 11.84$, $p < .001$).

Direct personal experience bolsters clinical, but not security, confidence. As shown in Figure 5.D, participants who reported having previously been prescribed a medical device were clearly more confident in their understanding of the clinical information ($\beta = 0.91$, 95% CrI: [0.25, 1.59]). The majority of these participants reported at least moderate confidence in the clinical necessity information provided before (94.4%) and after (95.1%) disclosure compared to participants with no medical device experience (92.0% before and 92.2% after). However, this same experience had no discernible effect on their confidence in the security information, and we did not observe any other clear effects between participants with less direct experience with medical devices or based on their experience with medical providers (Figure 5.E).

Summary of Post-Disclosure Confidence (RQ3).

- **Transparency builds confidence:** Detailed, verbose disclosures from regulators (FDA) increased patient security confidence.
- **No clinical spillover:** Vulnerability disclosures did not erode patients’ confidence in the clinical necessity of the device, suggesting patients can distinguish between security hygiene and clinical validity.
- **Internal drivers:** As with device choice, a patient’s baseline trust in their physician and security attitudes are stronger predictors of their confidence levels than the external communication materials.

6 Discussion

Our study reveals a clear delineation of factors that influence a patient’s decision-making process for connected medical devices. The findings demonstrate that a patient’s internal attributes, particularly their trust in their physician and their framing of risk in terms of physical safety, are far more influential than the external characteristics of how risk information is communicated, such as its source or level of detail. Furthermore, we find that the initial decision to adopt a device creates a powerful psychological anchor that is difficult to shift, even when new risk information is introduced. Counter-intuitively, our results also suggest that greater transparency in the form of detailed vulnerability disclosures can enhance, rather than erode, patient confidence.

These findings must be contextualized within the broader landscape of HCI and medical security research. While Tully et al. proposed the theoretical necessity of “cybersecurity informed consent” [76], our work provides the first large-scale empirical operationalization of this concept. We demonstrate that the theoretical ideal of a patient weighing all source information equally does not match the empirical reality: patients rely heavily on trusted heuristics (the physician) rather than processing new information sources (the FDA/MDM) independently. Consequently, the complexity of the regulatory ecosystem is, for the patient, effectively treated as a single channel of trust: their provider. Our results should be interpreted as exploratory evidence mapping these perceptions, offering parameters for future experimental designs rather than definitive behavioral predictions.

In this section, we discuss the biggest takeaways from our work, followed by the practical implications of these findings for three key groups: clinicians who conduct the consent process, regulators and manufacturers who oversee device safety and communication, and

researchers who seek to further improve the security and usability of medical systems.

6.1 Rethinking the Knowledge Deficit Model in Security Consent

Our findings suggest a fundamental misalignment between the current regulatory approach to patient communication and the empirical reality of patient decision-making. The prevailing strategy, the “knowledge deficit model” in science communication theory, assumes that risk-averse behavior stems from a lack of technical knowledge, and that providing patients with accurate, authoritative data (e.g., from the FDA or MDM) will enable them to perform a rational risk calculation [27, 37, 68]. However, our data reveal that this model misses important nuances in the context of cybersecurity-informed consent, including the reliance on trust in the provider, and that the type of impact leads to different decision-making.

The Failure of Direct-to-Patient Authority. Participants’ decisions were overwhelmingly predicted by their baseline trust in their physician [35]. We found that the source of the information (FDA, MDM, or Provider) had no discernible effect on the initial decision. This implies that patients treat the complex regulatory ecosystem as a “black box”; they do not independently evaluate the credibility of a federal regulator versus a manufacturer. Instead, they rely on the clinician as a trusted proxy to filter this complexity. The current FDA framework relies on direct-to-patient notices to inform risk [37], but our work suggests patients interpret risk through the lens of their physician. Therefore, the current framework likely does not utilize an essential component for patient decision-making and confidence, the provider.

Unfortunately, incorporating providers into this process presents a challenge as prior work has shown providers lack the security expertise to evaluate that risk [80]. Our work suggests it is important to help educate providers so they can provide this type of support. However, we are not suggesting the FDA, HHS, or other healthcare organizations seek to make all providers cybersecurity experts, as this is not a viable solution. Instead, it is important that the FDA focus on providing vulnerability information tailored to educating providers, not patients, about potential medical risks and allow providers to translate those medical risks to their patients’ specific context, as they already do with more traditional health risks.

Operationalising the Safety Threshold. Furthermore, our results identify a distinct “safety threshold” in patient heuristics, extending prior qualitative work on patient values [23]. The shift from data privacy risk to physical safety risk increased the likelihood of disconnecting ($\beta = 0.96$), suggesting that patients operate on a tiered model of risk tolerance rather than making linear trade-offs. In this hierarchy, data privacy risks are often viewed as tolerable low-impact events that align with the “I’m not a target” mental models found in prior medical device work [23], and broader security literature [11, 61]. Conversely, physical safety risks constitute an intolerable situation that triggers an immediate defensive response. This sharp behavioral divergence confirms that patients view device cybersecurity not just as an information problem, but as a critical patient safety “never event” [45].

Current disclosures often conflate these, presenting “cybersecurity risk” as a monolithic concept [37]. To support informed consent, communication frameworks must explicitly categorize vulnerabilities into these two distinct buckets, rather than relying on technical severity scores, which do not map cleanly to this patient-centric mental model.

6.2 Recommendations for Clinicians

The patient-provider relationship is the bedrock upon which consent is built. Our findings point to several actionable strategies for clinicians navigating cybersecurity discussions.

Prioritise the trust heuristic over communication details. The most powerful predictor of a patient’s willingness to adopt a connected device was their baseline trust in their provider. This indicates that when faced with unfamiliar cybersecurity risks, patients fall back on a known, trusted heuristic: the guidance of their clinician. The informed consent process for these devices is therefore not just a technical explanation but a relational act. Efforts to maintain and build this trust are paramount, as it serves as the primary lens through which patients will evaluate complex risk-benefit trade-offs.

Frame risk around the safety threshold. The distinction between a risk to data and a risk to physical safety was the most potent communication factor in our study. Clinicians should be aware that patients possess a clear mental threshold; risks of data breaches are often viewed as a manageable, almost familiar, inconvenience, whereas the potential for physical harm is a clear tipping point. Consent conversations should explicitly address this distinction, clarifying what a “hack” could mean in concrete terms for the patient’s bodily integrity versus their personal information.

Leverage the initial conversation. Our results show that a patient’s initial choice is remarkably “sticky” due to a powerful anchoring effect. This places immense importance on the initial consent conversation as the highest-leverage moment in the decision-making process. Post-market disclosures serve more as a nudge for the already hesitant rather than a full reset of the decision. Therefore, ensuring the first conversation is clear and addresses the safety threshold is the most effective way to achieve meaningful informed consent.

6.3 Recommendations for Regulators and Manufacturers

Our findings both reinforce and add critical nuance to existing federal guidance, providing empirical evidence to help shape future communication strategies that better align with how patients make decisions [27].

Embrace verbosity and transparency. The FDA’s guidance correctly encourages “clear and plain language” but worries that technical jargon can confuse patients [37]. However, our results suggest that detail, when well-presented, is beneficial. A key, counterintuitive finding from our study is that detailed, verbose vulnerability disclosures increased patient confidence, particularly when issued

by the FDA. This provides empirical support for the idea that transparency builds trust. The common fear that providing more detail will only confuse or frighten patients is not supported in our data. Therefore, regulators push manufacturers beyond minimalistic disclosures and toward providing richer information about risks and mitigations, as we expect it will empower, not alarm, patients.

Highlight that security disclosures do not erode clinical trust. The FDA guidance recommends a “balanced discussion between risk and benefits,” especially for life-saving devices [37]. Our study suggests this balance is achievable. We found no negative spillover effect; a vulnerability disclosure had no discernible impact on a patient’s confidence in the clinical necessity of their device. This evidence should embolden the entire ecosystem to adopt more proactive and open communication.

6.4 Directions for Future Research

This study opens several avenues for future work at the intersection of HCI, usable security, and clinical practice.

EHR-Integrated Device Dashboards. Future work should explore the integration of device security information directly into patient-facing Electronic Health Record (EHR) portals (e.g., MyChart). This leverages the trust patients have in their providers, and because the source of the information has little effect, presenting it as coming from the provider does not directly affect a patient’s decision-making. Just as patients currently view pharmaceutical prescriptions and lab results within these portals, a dedicated Device tab could provide a centralized inventory of a patient’s implanted hardware and real-time vulnerability status. However, designing this interface presents a critical challenge regarding the timing of information delivery. Much like the tension observed with the 21st Century Cures Act—where patients may access sensitive diagnostic results (e.g., oncology reports) before a clinician can provide context—pushing raw vulnerability data to a patient without clinical mediation risks inducing unnecessary panic. Future studies must therefore investigate design patterns that balance transparency while ensuring that technical alerts are synchronized with provider availability to prevent the “stumbling upon” of alarming information in isolation.

Centralized Vulnerability Reporting Channels. Additionally, we see an opportunity to investigate the efficacy of using the patient portal as a centralized middleware for reporting issues to MDMs and regulators. Currently, the reporting ecosystem is fragmented, requiring patients to navigate unfamiliar external portals to contact an MDM or the FDA. By embedding these reporting mechanisms within the hospital’s existing digital infrastructure, the system can leverage the “trust transfer” effect observed in our results, in which the high baseline trust in the provider extends to the tools hosted within the ecosystem. While this architecture places a greater implementation burden on hospitals, it significantly reduces the cognitive load on the patient. Research is needed to quantify whether this centralization enhances the volume and quality of patient-initiated adverse event reporting, effectively utilizing the hospital’s digital environment to bridge the gap between patients, manufacturers, and regulators.

Addressing “I’m Not a Target” and multi-patient harm. A recurring theme in our qualitative data, consistent with prior work [23], was patients’ belief that an attacker would not personally target them. This belief is fundamentally misaligned with the modern threat landscape, where attacks are often designed for scale, creating the potential for multi-patient harm. This shift from targeted single-patient exploits to systemic risk is now a primary concern for regulators, with recent FDA pre-market guidance explicitly requiring manufacturers to address threats that could impact multiple devices at once [38]. There is now a critical gap between the patient’s mental model (“Why would anyone target me?”) and the regulator’s core concern (systemic, multi-patient risk). Future research is needed to develop communication strategies to help patients understand this concept of scalable risk, thereby aligning their perceptions with the reality of the modern threat landscape. Future work can bridge this gap by adopting a cyber public health framework, a concept explored by organizations like CyberGreen to help differentiate between singular incidents and systemic risks [1, 6]. Communicating these vulnerabilities as being analogous to one’s susceptibility to a widespread illness like the flu, rather than a targeted attack, may be a crucial step in aligning patient risk perception with the reality of these scalable threats.

Patient-engaged Threat Modeling. Finally, our results highlight a disconnect between the technical view of risk and the patient view of safety thresholds. The issue of integrating safety into security risk assessments has been identified in prior work looking at medical device security [74]. Future research should engage with patient-led advocacy groups, such as the Light Collective (a non-profit dedicated to advancing the collective rights and voices of patient communities in health technology), to co-design vulnerability communications. Our data shows that patients can digest verbose, transparent information without panic; engaging them directly in the threat modeling process could further align technical disclosures with patient values. Crucially, this engagement is necessary to define the ethical boundaries of “context-aware” or “personalized” interventions. While dynamic tailoring is an effective strategy in general-purpose usable security [49, 58, 63], its application in clinical informed consent risks violating the ethical principle of equal access to information. An algorithmic approach that dynamically filters or simplifies risk disclosures for specific demographics could inadvertently institutionalize information asymmetries, ensuring that only certain patients receive the full fidelity of clinical risk data. Future frameworks must therefore balance usability with equity, ensuring that efforts to reduce cognitive load do not exacerbate existing health disparities.

Conducting longitudinal and in-situ studies. Our study relied on a vignette-based survey, which provides high internal validity but may not fully capture the complexities of real-world decision-making. Longitudinal studies that follow actual patients from the point of consent through their experience of living with a connected device are needed. Such research could explore how patient attitudes, concerns, and decisions evolve over months or years, and whether the powerful anchoring effect we observed persists over time. Furthermore, future qualitative work should seek out

and interview patients who have been directly impacted by real-world cybersecurity events, such as the recall of Medtronic's MiniMed insulin pumps [54], to understand their lived experiences and decision-making processes. Our work enables these future studies by providing specific, previously unknown insights into patient decision-making. We now know to focus on measuring a patient's baseline trust in their physician, as our results show it outweighs the influence of an information source's formal authority (e.g., the FDA), and to track the powerful anchoring effect of their initial choice. Critically, our discovery that detailed vulnerability disclosures can increase confidence rather than cause panic provides a validated, ethical path forward for discussing these sensitive topics with patients in a real-world setting.

Acknowledgments

We thank the survey respondents for their participation. We thank the anonymous reviewers who provided helpful comments on this paper's drafts. This project was supported by an ARPA-H contract (140D042590046) and NSF grant CNS-2440353.

References

- [1] Adam Shostack. 2023. Technical Report 23-01: A Cyber Belief Model.
- [2] Devdatta Akhawe and Adrienne Porter Felt. 2013. Alice in Warningland: A {Large-Scale} Field Study of Browser Security Warning Effectiveness. 257–272. <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/presentation/akhawe>
- [3] American Medical Association. [n. d.]. Informed Consent - Code of Ethics. <https://code-medical-ethics.ama-assn.org/ethics-opinions/informed-consent>.
- [4] Adrian Baranchuk, Marwan M. Refaat, Kristen K. Patton, Mina K. Chung, Kousik Krishnan, Valentina Kutiyfa, Gaurav Upadhyay, John D. Fisher, Dhanunjaya R. Lakkireddy, and null null. 2018. Cybersecurity for Cardiac Implantable Electronic Devices. *JACC* 71, 11 (2018), 1284–1288. doi:10.1016/j.jacc.2018.01.023 arXiv:<https://www.jacc.org/doi/pdf/10.1016/j.jacc.2018.01.023>
- [5] Roman Bednarik, Ann Blandford, Feng Feng, Antti Huotari, Matti Iso-Mustajärvi, Ahreum Lee, Federico Nicolosi, Jeremy Opie, Soojeong Yoo, and Bin Zheng. 2022. Integration of human factors in surgery: Interdisciplinary collaboration in design, development, and evaluation of surgical technologies. In *CHI conference on human factors in computing systems extended abstracts*. 1–7.
- [6] Bill Reid and Taylor Lehmann. 2024. Cyber Public Health: A New Approach to Cybersecurity. <https://cloud.google.com/blog/products/identity-security/cyber-public-health-a-new-approach-to-cybersecurity>.
- [7] Tamara Bonaci, Jeffrey Herron, Tariq Yusuf, Junjie Yan, Tadayoshi Kohno, and Howard Jay Chizeck. 2015. To Make a Robot Secure: An Experimental Analysis of Cyber Security Threats Against Teleoperated Surgical Robots. arXiv:1504.04339 [cs] (May 2015). <http://arxiv.org/abs/1504.04339> arXiv: 1504.04339.
- [8] Boston Scientific. [n. d.]. How Remote Monitoring System Works. <https://www.bostonscientific.com/en-US/patients-caregivers/device-support/remote-monitoring-system/how-remote-monitoring-system-works.html>.
- [9] Melissa M. Bottrell, Hillel Alpert, Ruth L. Fischbach, and Linda L. Emanuel. 2000. Hospital Informed Consent for Procedure Forms: Facilitating Quality Patient-Physician Interaction. *Archives of Surgery* 135, 1 (01 2000), 26–33. doi:10.1001/archsurg.135.1.26
- [10] Aaron F. Brantly and Nataliya D. Brantly. 2020. Patient-centric cybersecurity. *Journal of Cyber Policy* 5, 3 (Sept. 2020), 372–391. doi:10.1080/23738871.2020.1856902
- [11] Cristian Bravo-Lillo, Lorrie Faith Cranor, Julie Downs, and Saranga Komanduri. 2011. Bridging the Gap in Computer Security Warnings: A Mental Model Approach. *IEEE Security & Privacy* 9, 02 (March 2011), 18–26. doi:10.1109/MSP.2010.198
- [12] Brigham and Women's Hospital. [n. d.]. Cardiac Remote Monitoring. <https://www.brighamandwomens.org/heart-and-vascular-center/procedures/remot-monitoring>.
- [13] Carolyn Brown, Amy Morlock, Karin Blakolmer, Elham Heidari, and Robert Morlock. 2022. COVID-19 vaccination and race—A nationwide survey of vaccination status, intentions, and trust in the US general population. *Journal of Managed Care & Specialty Pharmacy* 28, 12 (2022), 1429–1438.
- [14] US Census Bureau. 2023. American Community Survey. https://www.himss.org/sites/hde/files/handout-35FINAL_111.pdf.
- [15] Paul-Christian Bürkner. 2017. Advanced Bayesian Multilevel Modeling with the R Package brms. arXiv:1705.11123 [stat.CO] <https://arxiv.org/abs/1705.11123>
- [16] Center for Devices and Radiological Health, Food & Drug Administration. 2016. Postmarket Management of Cybersecurity in Medical Devices. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/postmarket-management-cybersecurity-medical-devices>
- [17] Center for Devices and Radiological Health, Food & Drug Administration. 2023. Cybersecurity in Medical Devices: Refuse to Accept Policy for Cyber Devices Under Section 524B of the FD&C Act.
- [18] Christine C Chiu, Kim J Vicente, ILAN BUFFO-SEQUEIRA, Robert M Hamilton, and BRIAN W McCRINDLE. 2004. Usability assessment of pacemaker programmers. *Pacing and clinical electrophysiology* 27, 10 (2004), 1388–1398.
- [19] Cleveland Clinic. 2024. Do You Know the Symptoms of Heart Block? <https://my.clevelandclinic.org/health/diseases/17056-heart-block>.
- [20] Juliet Corbin and Anselm Strauss. 2014. *Basics of qualitative research: Techniques and procedures for developing grounded theory*. Sage publications.
- [21] Paul Curzon, Ann Blandford, Harold Thimbleby, and Anna Cox. 2015. Safer interactive medical device design: insights from the CHI+ MED project. In *Proceedings of the 5th EAI International Conference on Wireless Mobile Communication and Healthcare*. 34–37.
- [22] Paul Curzon, Paolo Masci, Patrick Oladimeji, Rimvydas Rukšenas, Harold Thimbleby, and Enrico D'Urso. 2014. Human-computer interaction and the formal certification and assurance of medical devices: the CHI+ MED project. In *2nd Workshop on verification and assurance (Verisure2014), in association with computer-aided verification (CAV), part of the Vienna summer of logic*.
- [23] Tamara Denning, Alan Borning, Batya Friedman, Brian T. Gill, Tadayoshi Kohno, and William H. Maisel. 2010. Patients, pacemakers, and implantable defibrillators: human values and security for wireless implantable medical devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 917–926. <https://doi.org/10.1145/1753326.1753462>
- [24] Tamara Denning, Kevin Fu, and Tadayoshi Kohno. 2008. Absence Makes the Heart Grow Fonder: New Directions for Implantable Medical Device Security.. In *HotSec*.
- [25] Division of Industry and Consumer Education. 2019. Certain Medtronic MiniMed Insulin Pumps Have Potential Cybersecurity Risks: FDA Safety Communication | FDA. <https://web.archive.org/web/20211020180523/https://www.fda.gov/medical-devices/safety-communications/certain-medtronic-minimed-insulin-pumps-have-potential-cybersecurity-risks-fda-safety-communication>.
- [26] Cori Faklaris, Laura A Dabbish, and Jason I Hong. 2019. A {self-report} measure of {end-user} security attitudes ({{{SA-6}}}). In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 61–77.
- [27] FDA. 2019. Priority Area 8: Strengthen Social and Behavioral Science to Promote Informed Decision-Making About FDA-Regulated Products. <https://www.fda.gov/science-research/advancing-regulatory-science/priority-area-8-strengthen-social-and-behavioral-science-promote-informed-decision-making-about-fda> Publisher: FDA.
- [28] Andrew Gelman and John Carlin. 2014. Beyond power calculations: Assessing type S (sign) and type M (magnitude) errors. *Perspectives on psychological science* 9, 6 (2014), 641–651.
- [29] Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. 1995. *Bayesian data analysis*. Chapman and Hall/CRC.
- [30] Johanna Glaser, Sarah Nouri, Alicia Fernandez, Rebecca L Sudore, Dean Schillinger, Michele Klein-Fedyshin, and Yael Schenker. 2020. Interventions to improve patient comprehension in informed consent for medical and surgical procedures: an updated systematic review. *Medical Decision Making* 40, 2 (2020), 119–143.
- [31] Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib, Norman Sadeh, Lorrie Faith Cranor, and Yuvraj Agarwal. 2016. How short is too short? implications of length and framing on the effectiveness of privacy notices. In *Proceedings of the Twelfth USENIX Conference on Usable Privacy and Security (Denver, CO, USA) (SOUPS '16)*. USENIX Association, USA, 321–340.
- [32] Dan Goodin. 2019. Critical flaw lets hackers control lifesaving devices implanted inside patients. <https://arstechnica.com/information-technology/2019/03/critical-flaw-lets-hackers-control-lifesaving-devices-implanted-inside-patients/>
- [33] Christine Grady. 2015. Enduring and emerging challenges of informed consent. *New England Journal of Medicine* 372, 9 (2015), 855–862.
- [34] Daniel E. Hall, Allan V. Prochazka, and Aaron S. Fink. 2012. Informed consent for clinical treatment. *CMAJ* 184, 5 (2012), 533–540. doi:10.1503/cmaj.112120 arXiv:<https://www.cmaj.ca/content/184/5/533.full.pdf>
- [35] Mark A. Hall, Beiyao Zheng, Elizabeth Dugan, Fabian Camacho, Kristin E. Kidd, Aneil Mishra, and Rajesh Balkrishnan. 2002. Measuring Patients' Trust in their Primary Care Providers. *Medical Care Research and Review* 59, 3 (2002), 293–318. doi:10.1177/1077558702059003004 arXiv:<https://doi.org/10.1177/1077558702059003004> PMID: 12205830.
- [36] Daniel Halperin, Thomas S Heydt-Benjamin, Benjamin Ransford, Shane S Clark, Benessa Defend, Will Morgan, Kevin Fu, Tadayoshi Kohno, and William H Maisel. 2008. Pacemakers and implantable cardiac defibrillators: Software radio attacks

- and zero-power defenses. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*. IEEE, 129–142.
- [37] Center for Devices and Radiological Health. 2021. *Best Practices for Communicating Cybersecurity Vulnerabilities to Patients*. Technical Report 152608. FDA. <https://www.fda.gov/about-fda/cdrh-patient-science-and-engagement-program/best-practices-communicating-cybersecurity-vulnerabilities-patients>. Publisher: FDA.
- [38] Center for Devices and Radiological Health. Thu, 06/26/2025 - 13:58. Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-medical-devices-quality-system-considerations-and-content-premarket-submissions>.
- [39] John M Hoenig and Dennis M Heisey. 2001. The abuse of power: the pervasive fallacy of power calculations for data analysis. *The American Statistician* 55, 1 (2001), 19–24.
- [40] Mohammad S Jalali and Jessica P Kaiser. 2018. Cybersecurity in hospitals: a systematic, organizational perspective. *Journal of medical Internet research* 20, 5 (2018), e10059.
- [41] James Dabbs. 2019. MiniMed™ 508 Insulin Pump and MiniMed™ Paradigm™ Series Insulin Pumps | Medtronic Diabetes. <https://web.archive.org/web/20211020215638/https://www.medtronicdiabetes.com/customer-support/product-and-service-updates/notice11-letter>.
- [42] Todd R Johnson, Harold Thimbleby, Peter Killoran, and Franck Diaz-Garelli. 2024. Human-Computer Interaction in Medical Devices. In *Human Computer Interaction in Healthcare: The Role of Cognition*. Springer, 319–343.
- [43] Matthew Kay, Gregory L Nelson, and Eric B Hekler. 2016. Researcher-centered design of statistics: Why Bayesian statistics better fit the culture and incentives of HCI. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 4521–4532.
- [44] Aoife Kiernan, Brian Fahey, Shaista S Guraya, Fiona Boland, Daragh Moneley, Frank Doyle, and Denis W Harkin. 2023. Digital technology in informed consent for surgery: systematic review. *BJS Open* 7, 1 (01 2023), zrac159. doi:10.1093/bjsopen/zrac159 arXiv:https://academic.oup.com/bjsopen/article-pdf/7/1/zrac159/48846663/zrac159.pdf
- [45] Daniel B Kramer, Jennifer R Amos, Julian M Goldman, and Kevin Fu. 2025. Threats to Patient Safety From Cybersecurity Flaws—A New Never Event. *JAMA* (2025).
- [46] Daniel B Kramer and Kevin Fu. 2017. Cybersecurity concerns and medical devices: lessons from a pacemaker advisory. *Jama* 318, 21 (2017), 2077–2078.
- [47] Klaus Krippendorff. 2018. *Content analysis: An introduction to its methodology*. Sage publications.
- [48] Charles W Lidz, Paul S Appelbaum, and Alan Meisel. 1988. Two models of implementing informed consent. *Archives of Internal Medicine* 148, 6 (1988), 1385–1389.
- [49] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuheimi, Shikun Aerin Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. 2016. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Twelfth symposium on usable privacy and security (SOUPS 2016)*. 27–41.
- [50] Lowell J. Schiller. 2019. September 10, 2019: Patient Engagement Advisory Committee Meeting Announcement.
- [51] Peter Mayer, Yixin Zou, Florian Schaub, and Adam J. Aviv. 2021. “Now I’m a bit angry:” Individuals’ Awareness, Perception, and Responses to Data Breaches that Affected Them. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 393–410. <https://www.usenix.org/conference/usenixsecurity21/presentation/mayer>
- [52] Medtronic. 2025. Remote Monitoring for Implanted Heart Devices. <https://www.medtronic.com/en-us/l/patients/treatments-therapies/remotemonitoring.html>.
- [53] Medtronic. 2025. Types of Pacemakers. <https://www.medtronic.com/en-us/l/patients/treatments-therapies/pacemakers/options-types.html>.
- [54] Medtronic Diabetes. 2021. URGENT MEDICAL DEVICE RECALL MiniMed™ Remote Controller (MMT-500 or MMT-503). <https://www.medtronicdiabetes.com/customer-support/product-and-service-updates/notice16-letter>
- [55] Peter Ney, Karl Koscher, Lee Organick, Luis Ceze, and Tadayoshi Kohno. [n. d.]. Computer Security, Privacy, and DNA Sequencing: Compromising Computers with Synthesized DNA, Privacy Leaks, and More. ([n. d.]), 15.
- [56] Niki O’Brien, Saira Ghafur, and Mike Durkin. 2021. Cybersecurity in health is an urgent patient safety concern: we can learn from existing patient safety improvement strategies to address it. *Journal of Patient Safety and Risk Management* 26, 1 (2021), 5–10.
- [57] Karl Pearson. 1900. X. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 50, 302 (1900), 157–175.
- [58] Eyal Peer, Serge Egelman, Marian Harbach, Nathan Malkin, Arunesh Mathur, and Alisa Frik. 2020. Nudge me right: Personalizing online security nudges to people’s decision-making styles. *Computers in Human Behavior* 109 (2020), 106347.
- [59] PennState Health. 23/06/22. Consent for Permanent Pacemaker Insertion. <https://psh.myprintdesk.net/DSF/~/PreviewPdf.aspx?+LoStWa3IPQAGwWagB2J9YTXESQ8+oNn6Clmx1XPjUDAvEnMCc5cXyLADeGgR0l9Q9Pi6l8mYZ32izlouFwR1/ADHqciDtsILJv5aeVy0ecIkX3Et5jNfyLkJbj5hi5Ngv23kvBsztQ5ja3sl2EZLqJZhuQK+LmVQMa8Ofz+WoNjhj6futtZQ==>.
- [60] Eoghan Pomeroy, Shahril Shaarani, Robert Kenyon, and James Cashman. 2021. Patient Recall of Informed Consent at 4 Weeks After Total Hip Replacement With Standardized Versus Procedure-Specific Consent Forms. *Journal of patient safety* 17, 6 (September 2021), e575–e581. doi:10.1097/pts.0000000000000412
- [61] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. 2016. How I Learned to be Secure: a Census-Representative Survey of Security Advice Sources and Behavior. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (Vienna, Austria) (CCS ’16)*. Association for Computing Machinery, New York, NY, USA, 666–677. doi:10.1145/2976749.2978307
- [62] Elissa M. Redmiles, Amelia R. Malone, and Michelle L. Mazurek. 2016. I Think They’re Trying to Tell Me Something: Advice Sources and Selection for Digital Security. In *2016 IEEE Symposium on Security and Privacy (SP)*. 272–288. doi:10.1109/SP.2016.24
- [63] Elissa M Redmiles, Noel Warford, Amritha Jayanti, Aravind Koneru, Sean Kross, Miraida Morales, Rock Stevens, and Michelle L Mazurek. 2020. A comprehensive quality evaluation of security and privacy advice on the web. In *29th USENIX Security Symposium (USENIX Security 20)*. 89–108.
- [64] Michael Rushanan, Aviel D Rubin, Denis Foo Kune, and Colleen M Swanson. 2014. Sok: Security and privacy in implantable medical devices and body area networks. In *2014 IEEE symposium on security and privacy*. IEEE, 524–539.
- [65] Leslie A Saxon, David L Hayes, F Roosevelt Gilliam, Paul A Heidenreich, John Day, Milan Seth, Timothy E Meyer, Paul W Jones, and John P Boehmer. 2010. Long-term outcome after ICD and CRT implantation and influence of remote device follow-up: the ALTITUDE survival study. *Circulation* 122, 23 (2010), 2359–2367.
- [66] Ben Seri and Barak Hadad. 2021. PwnedPiper. <https://www.armis.com/research/pwnedpiper/>
- [67] Kerry A Sherman, Christopher Jon Kilby, Melissa Pehlivan, and Brittany Smith. 2021. Adequacy of measures of informed consent in medical practice: A systematic review. *Plos one* 16, 5 (2021), e0251485.
- [68] Molly J Simis, Haley Madden, Michael A Cacciatore, and Sara K Yeo. 2016. The lure of rationality: Why does the deficit model persist in science communication? *Public Understanding of Science* 25, 4 (2016), 400–414.
- [69] David J Slotwiner, Thomas F Deering, Kevin Fu, Andrea M Russo, Mary N Walsh, and George F Van Hare. 2018. Cybersecurity vulnerabilities of cardiac implantable electronic devices: communication strategies for clinicians—Proceedings of the Heart Rhythm Society’s Leadership Summit. *Heart Rhythm* 15, 7 (2018), e61–e67.
- [70] David J. Slotwiner, Thomas F. Deering, Kevin Fu, Andrea M. Russo, Mary N. Walsh, and George F. Van Hare. 2018. Cybersecurity vulnerabilities of cardiac implantable electronic devices: Communication strategies for clinicians—Proceedings of the Heart Rhythm Society’s Leadership Summit. *Heart Rhythm* 15, 7 (2018), e61–e67. doi:10.1016/j.hrthm.2018.05.001
- [71] Anselm Strauss and Juliet Corbin. 1990. *Basics of qualitative research*. Sage publications.
- [72] Wencheng Sun, Zhiping Cai, Yangyang Li, Fang Liu, Shengqun Fang, and Guoyan Wang. 2018. Security and privacy in the medical internet of things: a review. *Security and Communication Networks* 2018, 1 (2018), 5978636.
- [73] Jenny Tang, Eleanor Birrell, and Ada Lerner. 2022. Replication: How well do my results generalize now? The external validity of online privacy and security surveys. In *Eighteenth symposium on usable privacy and security (SOUPS 2022)*. 367–385.
- [74] Ronald E. Thompson, Madline McLaughlin, Carson Powers, and Daniel Votipka. 2024. “There are rabbit holes I want to go down that I’m not allowed to go down”: An Investigation of Security Expert Threat Modeling Practices for Medical Devices. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association, Philadelphia, PA, 4909–4926. <https://www.usenix.org/conference/usenixsecurity24/presentation/thompson>
- [75] Lyndal J Trevena, , Alexandra Barratt, Phyllis Butow, and Patrina Caldwell. 2006. A systematic review on communicating with patients about evidence. *Journal of Evaluation in Clinical Practice* 12, 1 (2006), 13–23. doi:10.1111/j.1365-2753.2005.00596.x arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1365-2753.2005.00596.x
- [76] Jeffrey Tully, Andrea Coravos, Megan Doerr, and Christian Dameff. 2020. Connected medical technology and cybersecurity informed consent: a new paradigm. *Journal of medical internet research* 22, 3 (2020), e17612.
- [77] Rick Wash and Molly M. Cooper. 2018. Who Provides Phishing Training? Facts, Stories, and People Like Me. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal, QC, Canada) (CHI ’18)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3173574.3174066
- [78] Patricia AH Williams and Andrew J Woodward. 2015. Cybersecurity vulnerabilities in medical devices: a complex environment and multifaceted problem. *Medical Devices: Evidence and Research* (2015), 305–316.

- [79] Mark A Wood and Kenneth A Ellenbogen. 2002. Cardiac pacemakers from the patient's perspective. *Circulation* 105, 18 (2002), 2136–2138.
- [80] Emily P. Zeitler and Daniel B. Kramer. 2021. What Should Cardiac Patients Know About Device Cybersecurity Prior to Implantation? *AMA Journal of Ethics* 23, 9 (Sept. 2021), 705–711. doi:10.1001/amajethics.2021.705 Publisher: American Medical Association.
- [81] Jiajie Zhang, Todd R Johnson, Vimla L Patel, Danielle L Paige, and Tate Kubose. 2003. Using usability heuristics to evaluate patient safety of medical devices. *Journal of biomedical informatics* 36, 1-2 (2003), 23–30.
- [82] Yixin Zou, Shawn Danino, Kaiwen Sun, and Florian Schaub. 2019. You 'Might' Be Affected: An Empirical Analysis of Readability and Usability Issues in Data Breach Notifications. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3290605.3300424

A Survey

This appendix contains the full text of the survey presented to participants. Before the initial questions, participants were given a copy of our consent form and asked to affirm their eligibility and willingness to participate.

Initial Questions

QProviders. Have you ever interacted with the following types of medical providers? Please select all that apply.

- Primary Care Provider
- Cardiologist
- Electrophysiologist
- Endocrinologist
- Neurologist
- None of the above

QTrust (Wake Forest Physician Trust Scale). We want to understand your relationship with YOUR medical provider(s), i.e., your Primary Care Provider or other providers you routinely see. Please rate your level of agreement with the following statements. (5-point Likert scale from *Strongly Disagree* to *Strongly Agree*)

- (1) Your provider will do whatever it takes to get you all the care you need.
- (2) Sometimes your provider cares more about what is convenient for him/her than about your medical needs.
- (3) Your provider's medical skills are not as good as they should be.
- (4) Your provider is extremely thorough and careful.
- (5) You completely trust your provider's decisions about which medical treatments are best for you.
- (6) Your provider is totally honest in telling you about all of the different treatment options available for your condition.
- (7) Your provider only thinks about what is best for you.
- (8) Sometimes your provider does not pay full attention to what you are trying to tell him/ her.
- (9) You have no worries about putting your life in your provider's hands.
- (10) All in all, you have complete trust in your provider.

Scenario: Pacemaker Installation

[The following text was presented to participants. The italicized portions in brackets represent the text that was varied according to the participant's randomly assigned experimental condition.]

We are going to ask you some questions about how you would make decisions about using medical devices.

We want you to imagine that you have been diagnosed with a heart block, which means that your heart has trouble passing electrical impulses across your heart. This can cause you to have a shortness of breath and cause chest pain as some of the symptoms, but it can also lead to heart failure if not effectively managed. Because of this, your medical provider, i.e., the provider that you routinely see, has said that you require a pacemaker.

Your provider is recommending that you use a new pacemaker that is Internet-connected and can be remotely monitored by your

care team, which *<Health Information>*. The health risks associated with this device are minimal. Additionally, because this device can connect to the Internet, *<Party Providing Information>* explains *<Security Information>* that could be hacked. If the device was hacked, it could lead to *<Impact of Attack>*. The FDA has approved the clinical use of this device to treat heart blocks.

There are alternative therapies available to treat this condition, and your provider explains that you have the option of a non-connected pacemaker. However, *<Health Information Non-Connected>*, as the device cannot be remotely monitored or controlled, you will need to see your provider on a more frequent basis. This pacemaker is very unlikely to be hacked as it is not connected to the Internet. Like the Internet-connected device, this has been approved by the FDA for the treatment of your condition.

QInstall-1. Which of the two devices would you choose to use?

- Definitely the Internet and smartphone connected device
- Probably the Internet and smartphone connected device
- Either the Internet and smartphone connected device or the non-connected device
- Probably the non-connected device
- Definitely the non-connected device

QInstall-2. How confident are you in your understanding the clinical need for the pacemaker? (5-point Likert scale from *Not confident at all* to *Extremely confident*)

QInstall-3. How confident are you in your understanding of the security of the pacemaker? (5-point Likert scale from *Not confident at all* to *Extremely confident*)

QInstall-4. Please briefly explain why you feel confident or not confident in the information about the pacemaker. (Free text response)

QInstall-5. Which of the following would you want to know more about the security of the pacemaker? Please select all that apply.

- Things that you can do to minimize the potential for hackers to access your pacemaker
- How the *<Party Providing Information>* made their assessment
- What other groups have evaluated the security of the pacemaker
- More details on what would happen in the worst-case if the pacemaker was hacked
- Other: [Free Text]
- I do not need more information about the security of the pacemaker

Scenario: Vulnerability Disclosure

[The following text was presented to participants. The italicized portions in brackets represent the text that was varied according to the participant's randomly assigned experimental condition.]

Now, assume you have chosen to use the Internet-connected pacemaker. Some time after you have been using the device, you are informed by *<Party Providing Information>* that a cybersecurity vulnerability has been discovered in the device and that *<Vulnerability*

Information> could be hacked. If the device was hacked, it could lead to *<Impact of Attack>*.

As a reminder, the Internet-connected pacemaker *<Health Information>*.

[For participants assigned the verbose condition, the text above was modified to the following.]

Now, assume you have chosen to use the Internet-connected pacemaker. Some time after you have been using the device, your provider gives a letter from the *<Party Providing Information>* that a cybersecurity vulnerability has been discovered in the device and that *<Vulnerability Information>* could be hacked, which could lead to *<Impact of Attack>*. As a reminder, the Internet-connected pacemaker *<Health Information>*. Here is the letter you received:

Dear Patient,

[For FDA Notice] The FDA is warning patients and health care providers that certain pacemakers have a potential cybersecurity vulnerability related to their wireless connectivity. You are receiving this letter because our records indicate you have an affected pacemaker.

[For MDM Notice] You are receiving this letter because our records indicate you have a pacemaker that we manufactured. Because your safety is our top priority, we are making you aware of a potential cybersecurity risk.

Potential Cybersecurity Risk

These pacemakers are designed to communicate using a wireless radio frequency (RF) with other devices, such as your home monitoring system and clinic programmers. This remote monitoring allows your doctor to receive important health data without requiring an in-person visit.

Security researchers have identified a potential cybersecurity vulnerability related to the wireless feature on these pacemakers. An unauthorized person could potentially connect wirelessly to a nearby pacemaker to change its settings and control pacing therapy. This could lead to a dangerously slow heart rate (bradycardia) if pacing is stopped, or other dangerous heart rhythms if pacing is delivered inappropriately.

[For FDA Notice] IMPORTANT NOTE: The FDA has worked with the medical device manufacturer and determined that *<Vulnerability Information>* could be hacked.

[For MDM Notice] IMPORTANT NOTE: We have determined that *<Vulnerability Information>* could be hacked.

Cybersecurity Precautions for All Patients

- Keep your home monitor and any other pacemaker-related equipment in your personal control at all times.
- Do not share your pacemaker's serial number.
- Be attentive to any device alerts, vibrations, or notifications from your home monitor.
- Report any new or unexpected symptoms (e.g., dizziness, fainting, palpitations) to your doctor immediately.
- Only use software and home monitoring equipment that your care team has provided.

- Disconnect your home monitor from your computer or internet connection when not in use.

[All participants saw the text below]

QVuln-1. Assume there is no cost and no surgery is required to disable the Internet connection on the pacemaker, which would remove the benefits of remote monitoring and <Health Information Non-Connected>. This would involve a one-time short appointment with your medical provider. How likely are you to disable the Internet connection on the pacemaker? (5-point Likert scale from *Extremely unlikely to Extremely likely*)

QVuln-2. Why would or would you not disable the Internet connection after hearing about this vulnerability? (*Free text response*)

QVuln-3. How confident are you in your understanding the clinical need for the pacemaker? (5-point Likert scale from *Not confident at all to Extremely confident*)

QVuln-4. How confident are you in your understanding of the security of the pacemaker? (5-point Likert scale from *Not confident at all to Extremely confident*)

QVuln-5. Please briefly explain why you feel confident or not confident in the information about the pacemaker? (*Free text response*)

QAttention. In the description above, which group told you about the vulnerability for this device?

- Your medical provider
- The manufacturer of the device
- A government agency (such as the FDA)
- An internet company

QSources. For each of the following possible information sources, please indicate how likely you would be to trust cybersecurity information about medical devices they shared. (5-point Likert scale from *Extremely unlikely to Extremely likely*)

- Healthcare Providers (e.g. Physicians or Nurse Practitioners)
- Friends/family
- Social media users I follow
- Healthcare specific websites (e.g. WebMD or Healthline)
- AI chat bots (e.g. ChatGPT or Claude)
- The government (such as the FDA)
- Other: [Free Text]

Demographics

In the next section, we are going to ask you questions about your prior exposure to medical devices and your security perceptions.

QMed-Experience. Have you, someone you care for, or a close friend or family member ever been prescribed a medical device?

- Yes, I have been prescribed a medical device
- Yes, someone I care for has been prescribed a medical device
- Yes, a close friend or family member has been prescribed a medical device
- No

SA-6. For the following statements we want you to think about your general security practices. Please indicate your level of agreement or disagreement. (5-point Likert scale from *Strongly Disagree to Strongly Agree*)

- (1) I seek out opportunities to learn about security measures that are relevant to me.
- (2) I am extremely motivated to take all the steps needed to keep my online data and accounts safe.
- (3) Generally, I diligently follow a routine about security practices.
- (4) I often am interested in articles about security threats.
- (5) I always pay attention to experts' advice about the steps I need to take to keep my online data and accounts safe.
- (6) I am extremely knowledgeable about all the steps needed to keep my online data and accounts safe.

QAge. How old are you? (*Response options: 18-24, 25-34, ..., 85 or older, Prefer not to answer*)

QRace. How would you describe your race? (*Select all that apply*)

QLatino. Do you identify as Hispanic or Latino?

QGender. What term best describes you? (*Response options: Female, Male, Non-Binary, Other, Prefer Not To Answer*)

QEducation. What is the highest grade or year of school you have completed? (*Response options: Less than high school/no GED, ..., Advanced degree, Prefer not to answer*)

B Additional Data Related to Results

B.1 Codebooks

Table 4: Codebook for why participant felt confident (pre- and post-disclosure).

Concept	Name	Definition
<i>Occupation/Prior Knowledge</i>		
	Familiar with Technology	Has some familiarity with technology either through a job or personal experience.
	Familiar with Security	Has some familiarity with security or cybersecurity.
	Prior medical experience	Either personal, familial, or friend relationship with medical experience.
	Knowledge matches previous experience	Believes that the explanation is sufficient and matches previous experiences.
	Knowledge will improve/Can learn	Believes their knowledge will improve or that they can learn more.
	No Knowledge about Tech	No knowledge in tech-related fields.
<i>Trust/Distrust in Authority</i>		
	Trusts Doctor	Believes that the doctor would know best.
	Trusts MDM	Believes that the medical device manufacturer has the correct information.
	Trusts FDA/Government	Believes that the government would properly approve the device.
	Trust in Device/Security	Believes that the device itself is secure and that the company will monitor for vulnerabilities.
	Distrust in Doctor	Does not believe the information given by the doctor.
	Distrust in MDM	Does not believe the information given by the device manufacturer.
	Distrust in Federal Administration	Has little faith in the current federal administration.
	Wants a second opinion	Would want another entity to approve or recommend the device.
	Doctor lacks security knowledge	Believes the doctor would not know about the technical specifications of the device.
<i>Trust/Distrust in Technology</i>		
	Trust/Embrace of Technology	Has a positive outlook on technology and connected devices.
	Scared/Worried about device	Psychologically fearful or surprised by the potential of being hacked.
	Distrust in Technology	Has a distrust or wariness of technology generally.
<i>Information & Knowledge Gaps</i>		
	Information is clear	Believes that the information given is sufficient to make a decision.
	Information is incomplete	Would want more information in order to make a decision.
	Information is confusing	Doesn't believe the information is clear or makes sense.
	Technical aspects of device	Wants information on device hardware, software, or security features.
	Medical aspects of device	Wants information on medical or functional complications from the implant.
	Cost/Insurance aspects of device	Wants information on monetary or financial costs of the device.
	Misconception that device was hacked	Believes that the device has already been hacked and made the decision on that assumption.
<i>General Risk Perception (Not Worried)</i>		
	Not Concerned	Is not concerned about attackers or believes the chances are very low.
	Risks are manageable	Believes that they or the device can manage the security risks.
	Data leak won't hurt them	Belief that a data breach won't hurt them.
	Low Likelihood of being hacked	Disbelief in the value of their own information.
	Unclear attacker motivation	Confusion about why someone would want to attack them.
<i>Cost-Benefit Analysis</i>		
	Benefits Outweigh Costs	Believes the benefits of the device outweigh the costs of a potential vulnerability.
	Costs Outweigh Benefits	Believes the costs of hackers accessing the device outweighs the benefits.
	Condition/Benefits are too low	Doesn't believe that their condition needs an internet-connected pacemaker.
	Internet is not available	The internet is not available where they live or travel.

Table 5: Codebook for why participants did/did not disconnect the device.

Concept	Name	Definition
<i>Privacy-Related Concerns</i>		
	Values privacy	Belief that their data is important to keep private.
	Low value on privacy	Belief that their data is not that important to keep private.
	Privacy has already been breached	Believes their data is already exposed so it doesn't matter.
<i>Risk Tolerance</i>		
	Benefits outweigh fear	Believes that their health is more important than their data being vulnerable.
	General low trust	Has low trust in the government, doctor, and MDM about the device.
	Unacceptable risk	The fear of the device or lack of benefits means that it is not a risk worth taking.
	Depends on type of data	Answer changes depending on the type of data that would be vulnerable.
	Hacking is inevitable	Believes that hacking is inevitable no matter what happens.
<i>Specific Fears</i>		
	General Fear	Has a general fear over an unsecured device.
	Health Data	Afraid that their health data (HIPAA or unspecified) will be exposed.

Continued on next page

Table 5: (Continued) Codebook for why participants did/did not disconnect the device.

Concept	Name	Definition
	Financial Data	Afraid that their financial data will be exposed.
	Personal Data	Afraid that their personal identifying or contact data will be exposed.
	Physical Harm	Afraid that physical harm will come from the device.
	Access to Internals	Believes the attacker would have the ability to change the device's function.
	Not fixable	Expresses doubt that the vulnerability can or will be fixed properly.
	Ambiguity Aversion	The decision is driven by a desire to avoid uncertainty, preferring a known outcome.
<i>Reasons for Not Switching (Not Concerned)</i>		
	Generally Not Concerned	Not concerned about the risk of the vulnerability.
	OK with financial data exposure	Believes that their financial data being exposed is an acceptable risk.
	OK with personal data exposure	Believes that their personal data being exposed is an acceptable risk.
	OK with health data exposure	Believes that their health data being exposed is an acceptable risk.
	Can be fixed with update	Believes the vulnerability could be patched with a software update.

B.2 Device Choice

Table 6: Posterior summaries for the cumulative logit model of Device Choice. The model includes main effects for all conditions and covariates, plus all two-way interactions between the primary experimental conditions. Estimates are on the log-odds scale.

Predictor	Estimate (β)	SE	95% Credible Interval
<i>Main Effects</i>			
Notice: FDA	0.25	0.19	[-0.11, 0.61]
Notice: MDM	-0.25	0.18	[-0.61, 0.12]
Health: High	0.30	0.17	[-0.04, 0.64]
Attack Likelihood: Unsure	-0.42	0.18	[-0.77, -0.07]
Impact: Safety	-0.58	0.17	[-0.91, -0.25]
SA-6 (scaled)	-0.09	0.04	[-0.16, -0.01]
WFPT (scaled)	0.47	0.04	[0.39, 0.54]
Device Exp: Caregiver	0.25	0.12	[0.02, 0.48]
Device Exp: Knows Someone	0.13	0.08	[-0.04, 0.29]
Device Exp: No Response	0.63	0.41	[-0.17, 1.43]
Device Exp: Prescribed	0.65	0.16	[0.34, 0.96]
Provider Exp: No PCP	1.00	0.34	[0.35, 1.67]
Provider Exp: Specialty	0.07	0.08	[-0.07, 0.22]
Age	-0.43	0.54	[-1.47, 0.62]
Education	0.20	0.09	[0.02, 0.38]
<i>Two-Way Interactions</i>			
Notice: FDA \times Health: High	-0.19	0.18	[-0.54, 0.17]
Notice: MDM \times Health: High	0.16	0.18	[-0.20, 0.51]
Notice: FDA \times Attack Likelihood: Unsure	0.11	0.19	[-0.26, 0.47]
Notice: MDM \times Attack Likelihood: Unsure	0.05	0.19	[-0.32, 0.42]
Notice: FDA \times Impact: Safety	-0.27	0.19	[-0.63, 0.09]
Notice: MDM \times Impact: Safety	0.07	0.19	[-0.30, 0.43]
Health: High \times Attack Likelihood: Unsure	0.05	0.15	[-0.23, 0.34]
Health: High \times Impact: Safety	0.13	0.15	[-0.17, 0.42]
Attack Likelihood: Unsure \times Impact: Safety	-0.09	0.15	[-0.39, 0.21]

B.3 Switch

Table 7: Posterior summaries for the cumulative logit model of Device Switch. The model includes main effects, all two-way interactions between the primary experimental conditions, and several covariates including prior device choice. Estimates are on the log-odds scale.

Predictor	Estimate (β)	SE	95% Credible Interval
<i>Main Effects</i>			
Notice: FDA Standard	0.20	0.25	[-0.29, 0.70]
Notice: FDA Verbose	-0.25	1.07	[-2.33, 1.86]
Notice: MDM Standard	0.27	0.25	[-0.22, 0.77]
Notice: MDM Verbose	-0.06	1.07	[-2.17, 2.00]
Health: High	-0.06	0.22	[-0.48, 0.38]
Attack Likelihood: Low/Unsure	0.18	0.27	[-0.35, 0.71]
Attack Likelihood: Low/High	0.74	0.27	[0.21, 1.27]
Attack Likelihood: Unsure/Unsure	0.60	0.27	[0.06, 1.13]
Attack Likelihood: Unsure/High	0.57	0.27	[0.03, 1.10]
Impact: Safety	0.96	0.23	[0.52, 1.40]
Device Choice	-3.26	0.13	[-3.52, -3.02]

Continued on next page

Table 7 – continued from previous page

Predictor	Estimate (β)	SE	95% Credible Interval
SA-6 (scaled)	0.13	0.04	[0.05, 0.21]
WFPT (scaled)	-0.11	0.04	[-0.19, -0.03]
Device Exp: Caregiver	0.01	0.12	[-0.23, 0.24]
Device Exp: Knows Someone	0.19	0.08	[0.03, 0.36]
Device Exp: No Response	0.10	0.40	[-0.69, 0.89]
Device Exp: Prescribed	-0.12	0.16	[-0.45, 0.20]
Provider Exp: No PCP	0.06	0.34	[-0.61, 0.73]
Provider Exp: Specialty	-0.09	0.08	[-0.24, 0.06]
Age	-0.59	0.60	[-1.78, 0.62]
Education	-0.00	0.09	[-0.18, 0.18]
<i>Two-Way Interactions</i>			
Notice: FDA Standard \times Health: High	-0.23	0.20	[-0.63, 0.17]
Notice: FDA Verbose \times Health: High	0.53	0.26	[0.00, 1.04]
Notice: MDM Standard \times Health: High	-0.31	0.20	[-0.71, 0.08]
Notice: MDM Verbose \times Health: High	-0.41	0.26	[-0.92, 0.12]
Notice: FDA Standard \times Attack Likelihood: Low/Unsure	-0.03	0.30	[-0.62, 0.55]
Notice: FDA Verbose \times Attack Likelihood: Low/Unsure	0.03	0.39	[-0.74, 0.78]
Notice: MDM Standard \times Attack Likelihood: Low/Unsure	0.31	0.30	[-0.28, 0.91]
Notice: MDM Verbose \times Attack Likelihood: Low/Unsure	0.01	0.38	[-0.74, 0.75]
Notice: FDA Standard \times Attack Likelihood: Low/High	0.33	0.31	[-0.28, 0.94]
Notice: FDA Verbose \times Attack Likelihood: Low/High	-0.06	0.38	[-0.81, 0.69]
Notice: MDM Standard \times Attack Likelihood: Low/High	-0.00	0.30	[-0.59, 0.60]
Notice: MDM Verbose \times Attack Likelihood: Low/High	0.38	0.40	[-0.41, 1.17]
Notice: FDA Standard \times Attack Likelihood: Unsure/Unsure	-0.09	0.30	[-0.67, 0.51]
Notice: FDA Verbose \times Attack Likelihood: Unsure/Unsure	0.55	0.39	[-0.20, 1.31]
Notice: MDM Standard \times Attack Likelihood: Unsure/Unsure	-0.04	0.30	[-0.63, 0.57]
Notice: MDM Verbose \times Attack Likelihood: Unsure/Unsure	0.20	0.38	[-0.54, 0.95]
Notice: FDA Standard \times Attack Likelihood: Unsure/High	-0.28	0.30	[-0.88, 0.32]
Notice: FDA Verbose \times Attack Likelihood: Unsure/High	0.17	0.38	[-0.58, 0.92]
Notice: MDM Standard \times Attack Likelihood: Unsure/High	0.14	0.31	[-0.46, 0.75]
Notice: MDM Verbose \times Attack Likelihood: Unsure/High	0.20	0.40	[-0.57, 0.99]
Notice: FDA Standard \times Impact: Safety	-0.18	0.20	[-0.58, 0.21]
Notice: FDA Verbose \times Impact: Safety	-0.27	1.07	[-2.35, 1.83]
Notice: MDM Standard \times Impact: Safety	-0.17	0.20	[-0.57, 0.22]
Notice: MDM Verbose \times Impact: Safety	-0.07	1.08	[-2.15, 2.05]
Health: High \times Attack Likelihood: Low/Unsure	-0.06	0.23	[-0.51, 0.38]
Health: High \times Attack Likelihood: Low/High	-0.37	0.23	[-0.82, 0.07]
Health: High \times Attack Likelihood: Unsure/Unsure	-0.23	0.22	[-0.67, 0.21]
Health: High \times Attack Likelihood: Unsure/High	0.27	0.23	[-0.17, 0.71]
Health: High \times Impact: Safety	-0.11	0.17	[-0.43, 0.22]
Attack Likelihood: Low/Unsure \times Impact: Safety	-0.07	0.25	[-0.57, 0.42]
Attack Likelihood: Low/High \times Impact: Safety	0.39	0.25	[-0.10, 0.90]
Attack Likelihood: Unsure/Unsure \times Impact: Safety	-0.32	0.26	[-0.82, 0.18]
Attack Likelihood: Unsure/High \times Impact: Safety	0.09	0.26	[-0.41, 0.58]

Table 8: Posterior summaries for the cumulative logit model of Device Switch, for participants who initially chose the connected device. The model includes main effects, all two-way interactions between the primary experimental conditions, and several covariates. Estimates are on the log-odds scale.

Predictor	Estimate (β)	SE	95% Credible Interval
<i>Main Effects</i>			
Notice: FDA Standard	0.10	0.29	[-0.47, 0.68]
Notice: FDA Verbose	-0.24	1.07	[-2.35, 1.86]
Notice: MDM Standard	0.16	0.29	[-0.40, 0.73]
Notice: MDM Verbose	-0.31	1.09	[-2.44, 1.87]
Health: High	-0.22	0.26	[-0.72, 0.29]
Attack Likelihood: Low/Unsure	0.09	0.31	[-0.51, 0.69]
Attack Likelihood: Low/High	0.95	0.31	[0.35, 1.56]
Attack Likelihood: Unsure/Unsure	0.32	0.33	[-0.32, 0.97]
Attack Likelihood: Unsure/High	0.62	0.32	[-0.00, 1.26]
Impact: Safety	1.05	0.27	[0.50, 1.57]
Device Choice	-0.81	0.07	[-0.95, -0.67]
SA-6 (scaled)	0.11	0.05	[0.01, 0.20]
WFPT (scaled)	-0.13	0.05	[-0.23, -0.03]
Device Exp: Caregiver	0.02	0.15	[-0.27, 0.30]
Device Exp: Knows Someone	0.26	0.10	[0.06, 0.47]
Device Exp: No Response	0.27	0.46	[-0.61, 1.17]
Device Exp: Prescribed	0.01	0.19	[-0.35, 0.38]
Provider Exp: No PCP	0.45	0.39	[-0.31, 1.21]
Provider Exp: Specialty	-0.07	0.10	[-0.26, 0.12]
Age	-0.64	0.61	[-1.83, 0.55]

Continued on next page

Table 8 – continued from previous page

Predictor	Estimate (β)	SE	95% Credible Interval
Education	0.12	0.11	[-0.10, 0.34]
<i>Two-Way Interactions</i>			
Notice: FDA Standard \times Health: High	-0.28	0.24	[-0.76, 0.19]
Notice: FDA Verbose \times Health: High	0.10	0.33	[-0.55, 0.74]
Notice: MDM Standard \times Health: High	-0.36	0.24	[-0.84, 0.12]
Notice: MDM Verbose \times Health: High	-0.26	0.32	[-0.90, 0.35]
Notice: FDA Standard \times Attack Likelihood: Low/Unsure	0.10	0.35	[-0.60, 0.79]
Notice: FDA Verbose \times Attack Likelihood: Low/Unsure	-0.23	0.46	[-1.15, 0.67]
Notice: MDM Standard \times Attack Likelihood: Low/Unsure	0.43	0.36	[-0.27, 1.13]
Notice: MDM Verbose \times Attack Likelihood: Low/Unsure	0.11	0.45	[-0.77, 1.00]
Notice: FDA Standard \times Attack Likelihood: Low/High	0.44	0.35	[-0.25, 1.14]
Notice: FDA Verbose \times Attack Likelihood: Low/High	-0.30	0.47	[-1.21, 0.62]
Notice: MDM Standard \times Attack Likelihood: Low/High	0.06	0.35	[-0.62, 0.73]
Notice: MDM Verbose \times Attack Likelihood: Low/High	0.55	0.47	[-0.36, 1.49]
Notice: FDA Standard \times Attack Likelihood: Unsure/Unsure	-0.00	0.37	[-0.72, 0.72]
Notice: FDA Verbose \times Attack Likelihood: Unsure/Unsure	0.53	0.49	[-0.41, 1.49]
Notice: MDM Standard \times Attack Likelihood: Unsure/Unsure	0.28	0.36	[-0.43, 0.99]
Notice: MDM Verbose \times Attack Likelihood: Unsure/Unsure	0.16	0.47	[-0.75, 1.10]
Notice: FDA Standard \times Attack Likelihood: Unsure/High	-0.28	0.37	[-0.99, 0.45]
Notice: FDA Verbose \times Attack Likelihood: Unsure/High	0.35	0.48	[-0.60, 1.29]
Notice: MDM Standard \times Attack Likelihood: Unsure/High	0.41	0.37	[-0.31, 1.13]
Notice: MDM Verbose \times Attack Likelihood: Unsure/High	0.36	0.47	[-0.57, 1.29]
Notice: FDA Standard \times Impact: Safety	-0.33	0.25	[-0.82, 0.15]
Notice: FDA Verbose \times Impact: Safety	-0.22	1.07	[-2.31, 1.89]
Notice: MDM Standard \times Impact: Safety	-0.41	0.25	[-0.89, 0.08]
Notice: MDM Verbose \times Impact: Safety	-0.30	1.09	[-2.47, 1.81]
Health: High \times Attack Likelihood: Low/Unsure	0.05	0.26	[-0.46, 0.56]
Health: High \times Attack Likelihood: Low/High	-0.44	0.27	[-0.97, 0.07]
Health: High \times Attack Likelihood: Unsure/Unsure	0.01	0.28	[-0.54, 0.56]
Health: High \times Attack Likelihood: Unsure/High	0.40	0.28	[-0.14, 0.94]
Health: High \times Impact: Safety	-0.02	0.21	[-0.42, 0.38]
Attack Likelihood: Low/Unsure \times Impact: Safety	0.12	0.29	[-0.45, 0.70]
Attack Likelihood: Low/High \times Impact: Safety	0.42	0.29	[-0.15, 0.99]
Attack Likelihood: Unsure/Unsure \times Impact: Safety	-0.21	0.31	[-0.81, 0.40]
Attack Likelihood: Unsure/High \times Impact: Safety	0.21	0.31	[-0.40, 0.83]

B.4 Confidence Models

Table 9: Posterior summaries for the cumulative logit model of Security Confidence. The model includes main effects, all two-way interactions, and a random intercept for each participant. Estimates are on the log-odds scale.

Predictor	Estimate (β)	SE	95% Credible Interval
<i>Main Effects</i>			
Time: Vulnerability	0.01	0.19	[-0.36, 0.37]
Notice: FDA Standard	0.24	0.43	[-0.61, 1.10]
Notice: FDA Verbose	-0.32	1.09	[-2.47, 1.80]
Notice: MDM Standard	-0.51	0.43	[-1.35, 0.34]
Notice: MDM Verbose	-0.04	1.09	[-2.20, 2.08]
Health: High	0.12	0.39	[-0.64, 0.88]
Attack Likelihood: Low/Unsure	-0.14	0.47	[-1.07, 0.80]
Attack Likelihood: Low/High	-0.19	0.47	[-1.10, 0.72]
Attack Likelihood: Unsure/Unsure	0.43	0.47	[-0.49, 1.34]
Attack Likelihood: Unsure/High	-0.49	0.48	[-1.44, 0.44]
Impact: Safety	0.13	0.40	[-0.67, 0.90]
WFPT (scaled)	0.69	0.07	[0.55, 0.84]
SA-6 (scaled)	0.78	0.08	[0.63, 0.93]
Device Exp: Caregiver	-0.22	0.23	[-0.67, 0.23]
Device Exp: Knows Someone	-0.29	0.16	[-0.61, 0.04]
Device Exp: No Response	-0.37	0.74	[-1.82, 1.07]
Device Exp: Prescribed	-0.08	0.31	[-0.68, 0.54]
Provider Exp: No PCP	-0.47	0.60	[-1.64, 0.70]
Provider Exp: Specialty	0.04	0.15	[-0.25, 0.33]
Age	0.00	0.78	[-1.53, 1.51]
Education	-0.67	0.18	[-1.01, -0.32]
<i>Two-Way Interactions</i>			
Time \times Notice: FDA Standard	-0.01	0.17	[-0.33, 0.32]
Time \times Notice: FDA Verbose	0.58	0.21	[0.15, 1.00]
Time \times Notice: MDM Standard	-0.06	0.16	[-0.38, 0.27]
Time \times Notice: MDM Verbose	0.40	0.22	[-0.03, 0.82]
Time \times Health: High	-0.02	0.12	[-0.26, 0.21]
Time \times Attack Likelihood: Low/Unsure	-0.32	0.19	[-0.68, 0.04]

Continued on next page

Table 9 – continued from previous page

Predictor	Estimate (β)	SE	95% Credible Interval
Time \times Attack Likelihood: Low/High	-0.10	0.18	[-0.46, 0.25]
Time \times Attack Likelihood: Unsure/Unsure	-0.27	0.19	[-0.63, 0.09]
Time \times Attack Likelihood: Unsure/High	0.17	0.18	[-0.20, 0.53]
Time \times Impact: Safety	-0.15	0.13	[-0.40, 0.12]
Notice: FDA Standard \times Health: High	-0.12	0.37	[-0.85, 0.61]
Notice: FDA Verbose \times Health: High	0.27	0.47	[-0.66, 1.18]
Notice: MDM Standard \times Health: High	-0.09	0.37	[-0.81, 0.63]
Notice: MDM Verbose \times Health: High	0.22	0.46	[-0.70, 1.13]
Notice: FDA Standard \times Attack Likelihood: Low/Unsure	0.36	0.52	[-0.65, 1.37]
Notice: FDA Verbose \times Attack Likelihood: Low/Unsure	-0.40	0.66	[-1.68, 0.91]
Notice: MDM Standard \times Attack Likelihood: Low/Unsure	0.66	0.51	[-0.33, 1.67]
Notice: MDM Verbose \times Attack Likelihood: Low/Unsure	0.46	0.64	[-0.76, 1.71]
Notice: FDA Standard \times Attack Likelihood: Low/High	0.48	0.51	[-0.53, 1.50]
Notice: FDA Verbose \times Attack Likelihood: Low/High	0.98	0.63	[-0.26, 2.21]
Notice: MDM Standard \times Attack Likelihood: Low/High	0.65	0.51	[-0.35, 1.65]
Notice: MDM Verbose \times Attack Likelihood: Low/High	-0.94	0.65	[-2.23, 0.34]
Notice: FDA Standard \times Attack Likelihood: Unsure/Unsure	-0.76	0.51	[-1.77, 0.25]
Notice: FDA Verbose \times Attack Likelihood: Unsure/Unsure	-0.08	0.65	[-1.35, 1.18]
Notice: MDM Standard \times Attack Likelihood: Unsure/Unsure	0.15	0.51	[-0.87, 1.14]
Notice: MDM Verbose \times Attack Likelihood: Unsure/Unsure	-0.61	0.65	[-1.89, 0.66]
Notice: FDA Standard \times Attack Likelihood: Unsure/High	0.15	0.51	[-0.86, 1.15]
Notice: FDA Verbose \times Attack Likelihood: Unsure/High	-0.54	0.64	[-1.79, 0.70]
Notice: MDM Standard \times Attack Likelihood: Unsure/High	0.15	0.52	[-0.86, 1.17]
Notice: MDM Verbose \times Attack Likelihood: Unsure/High	-0.53	0.66	[-1.82, 0.77]
Notice: FDA Standard \times Impact: Safety	-0.14	0.38	[-0.88, 0.61]
Notice: FDA Verbose \times Impact: Safety	-0.31	1.09	[-2.43, 1.84]
Notice: MDM Standard \times Impact: Safety	-0.03	0.37	[-0.74, 0.70]
Notice: MDM Verbose \times Impact: Safety	-0.02	1.09	[-2.14, 2.12]
Health: High \times Attack Likelihood: Low/Unsure	-0.04	0.41	[-0.85, 0.76]
Health: High \times Attack Likelihood: Low/High	-0.07	0.40	[-0.83, 0.72]
Health: High \times Attack Likelihood: Unsure/Unsure	-0.13	0.41	[-0.95, 0.67]
Health: High \times Attack Likelihood: Unsure/High	0.26	0.41	[-0.54, 1.08]
Health: High \times Impact: Safety	-0.05	0.31	[-0.65, 0.56]
Attack Likelihood: Low/Unsure \times Impact: Safety	-0.25	0.44	[-1.11, 0.61]
Attack Likelihood: Low/High \times Impact: Safety	-0.38	0.45	[-1.27, 0.51]
Attack Likelihood: Unsure/Unsure \times Impact: Safety	-0.36	0.45	[-1.25, 0.51]
Attack Likelihood: Unsure/High \times Impact: Safety	0.33	0.44	[-0.52, 1.20]

Table 10: Posterior summaries for the cumulative logit model of Clinical Confidence. The model includes main effects, all two-way interactions, and a random intercept for each participant. Estimates are on the log-odds scale.

Predictor	Estimate (β)	SE	95% Credible Interval
<i>Main Effects</i>			
Time: Vulnerability	0.09	0.21	[-0.30, 0.51]
Notice: FDA Standard	-0.19	0.49	[-1.12, 0.79]
Notice: FDA Verbose	-0.12	1.10	[-2.24, 2.07]
Notice: MDM Standard	-0.61	0.48	[-1.53, 0.34]
Notice: MDM Verbose	0.12	1.09	[-2.04, 2.23]
Health: High	0.51	0.43	[-0.33, 1.34]
Attack Likelihood: Low/Unsure	-0.33	0.51	[-1.31, 0.68]
Attack Likelihood: Low/High	-0.08	0.52	[-1.10, 0.97]
Attack Likelihood: Unsure/Unsure	0.20	0.51	[-0.82, 1.19]
Attack Likelihood: Unsure/High	-0.83	0.51	[-1.82, 0.21]
Impact: Safety	-0.08	0.45	[-0.98, 0.80]
WFPT (scaled)	0.91	0.09	[0.75, 1.08]
SA-6 (scaled)	0.93	0.09	[0.76, 1.10]
Device Exp: Caregiver	0.04	0.26	[-0.45, 0.55]
Device Exp: Knows Someone	-0.08	0.18	[-0.44, 0.27]
Device Exp: No Response	-1.08	0.80	[-2.63, 0.46]
Device Exp: Prescribed	0.91	0.34	[0.25, 1.59]
Provider Exp: No PCP	-0.36	0.66	[-1.67, 0.93]
Provider Exp: Specialty	0.23	0.17	[-0.09, 0.56]
Age	0.24	0.82	[-1.38, 1.85]
Education	-0.57	0.20	[-0.97, -0.19]
<i>Two-Way Interactions</i>			
Time \times Notice: FDA Standard	0.18	0.18	[-0.18, 0.54]
Time \times Notice: FDA Verbose	0.05	0.23	[-0.41, 0.52]
Time \times Notice: MDM Standard	0.02	0.18	[-0.34, 0.37]
Time \times Notice: MDM Verbose	0.49	0.24	[0.03, 0.96]
Time \times Health: High	0.09	0.13	[-0.16, 0.35]
Time \times Attack Likelihood: Low/Unsure	-0.19	0.20	[-0.58, 0.20]

Continued on next page

Name	Count	Percentage
Benefits outweigh fear	998	37.40
Unacceptable risk	838	31.40
Ambiguity Aversion	282	10.60
Physical Harm	248	9.30
Can be fixed with update	178	6.68
Ok with health data	177	6.64
General Fear	163	6.11
Generally Not Concerned	122	4.58
Health Data	103	3.86
Already wouldn't choose connected device	100	3.75

Table 13: Top 10 codes for why switching/not switching