

“Your imaging may be stone-cold normal, but if they look sick, they’re going to get admitted”: An Investigation of Clinicians’ Perceptions of Impact & Likelihood of Security Failures

Ronald E. Thompson III[†], Hamza Khalid[†], Hilary Fisher^{*}, Rhea Votipka[◇], Daniel Votipka[†]
**Brigham and Women’s Hospital; ◇Beth Isreal Lahey Health; †Tufts University*

Abstract

Cyberattacks are a critical patient safety issue, yet security controls often fail to account for the uniqueness of the clinical environment. This paper addresses the gap in understanding clinicians’ security perspectives through a mixed-methods study, with 12 interviews of US clinicians, followed by a 303-participant survey of clinicians across the US, UK, and Canada. Our findings reveal a significant misalignment between perceived threats and deployed controls. Clinicians perceive confidentiality failures (e.g., data breaches) as most likely. They view integrity failures (e.g., manipulated values) as catastrophic but trust their own expertise to ignore anomalous data. Finally, they manage likely and dangerous availability failures with analog workarounds like paper charting, introducing new risks. These results show the need to integrate clinicians into security, highlighting where existing approaches are lacking and providing recommendations for developing more effective, clinician-centered security.

1 Introduction

Cyberattacks against healthcare organizations are not merely IT incidents; they are patient safety crises. High-profile ransomware attacks and network outages have crippled hospital operations, canceled appointments, and delayed critical care, costing single health systems over \$100 million and impacting patient care [35, 39, 47, 50, 66, 68, 70]. Policymakers and healthcare leaders have declared “Cybersecurity *is* patient safety” [31, 51]. While a significant body of work has documented medical devices’ and hospital networks’ technical vulnerabilities [6, 24, 57] and there has been some work into patient perspectives [17], a critical gap remains in understanding another key human element: the clinicians providing care.

In response to the rise in healthcare threats, federal agencies have issued guidance, such as the Cybersecurity Performance Goals [19] and the Health Industry Cybersecurity Practices [18], both from the US Department of Health and Human Services, which focus on technical and procedural

controls designed to engineer out clinician error. While we agree that security burdens on clinicians should be minimized, it is currently impossible to remove them entirely from the process due to their need to access and assess patient data and their central role in patient care. It is essential to understand their perspectives to identify usability challenges they currently face. Further, we expect, due to their significant domain knowledge, rather than sidelining clinicians, a more effective strategy may be to elicit positive security behaviors, treating staff not as a vulnerability but as a potential protection [8]. Research suggests clinicians already provide a critical, albeit often unprepared, last line of defense. For example, high-fidelity simulations have shown clinicians can identify potential errors introduced by security failures, though they often do not recognize these errors as the result of a compromised medical device [16, 71]. However, this prior work has focused on specific populations [21, 27] or explored perceptions of a single technology, such as telehealth [62]. A broader understanding of clinicians’ security practices and perceptions is needed.

To explore these perspectives, we utilize the Confidentiality, Integrity, and Availability (CIA) triad (a foundational model in information security) as a lens to translate technical failures into clinical impacts. In a hospital, these are not abstract concepts but tangible requisites for care: confidentiality ensures patient trust and legal compliance; integrity ensures diagnoses are based on accurate physiological data; and availability ensures life-saving treatments are delivered without delay. By grounding our inquiry in these core tenets, we can identify exactly where clinical needs diverge from traditional enterprise security models. We specifically focus on two central research questions:

- RQ1** What do clinicians see as potential impacts to patient safety and privacy when security failures occur with technologies they rely on?
- RQ2** What practices do clinicians currently follow to protect patient security and privacy? How do these practices impact patient care?

To answer these questions, we first conducted in-depth, semi-structured interviews with 12 practicing inpatient¹ clinicians in the United States, including physicians, nurses, and advanced practice providers. Then, to assess our interview findings' generalizability, we conducted a large-scale survey with 303 clinicians across the US, UK, and Canada.

Through our interviews and survey, we observed that the clinical setting introduces unique opportunities and challenges for security depending on the specific clinical context. For example, we find confidentiality failures (e.g., data breaches) are perceived as the most likely failure. However, this is the only type of failure for which participants described security controls and training being deployed in practice. We found these security controls often failed to account for the high mobility (between units, computers, and hospitals) of clinicians and emergencies. Conversely, integrity failures (e.g., manipulated lab values or imaging) are viewed as potentially catastrophic, due to the potential impact on patient health. While no participants mentioned security controls or security training to mitigate these failures, clinicians dismissed these failures due to their belief they would be mitigated by their clinical knowledge and experience. That is, clinicians take a holistic approach to diagnosis and can disregard manipulated data that does not align with information from other systems, patient history, and physical exams. Finally, availability failures (e.g., system downtime), are seen as both likely and dangerous, a perception grounded in clinicians' frequent, real-world experiences with both security failures (e.g., ransomware) and general system failures or scheduled downtime. Further, for some conditions and treatments, simply delaying care can be catastrophic. To manage these risks, clinicians employ numerous institutionalized process-based workarounds, from reverting to paper charting to using personal devices, each of which introduces their own risks.

Our results provide insights into the security challenges in the clinical setting, highlighting where existing approaches are lacking and identifying misalignments between security failures and security control deployment. Furthermore, we emphasize the importance of involving clinicians in security decisions and the necessity of clear communication with them to determine when security controls are necessary and when existing safety practices can mitigate potential security failures. However, future work is needed to establish the limits of these existing safeguards. Based on these results, we provide recommendations for improvements to healthcare security and outline directions for future work.

2 Related Work

Our work builds on three distinct but related streams of research. We first review prior work on the technical security landscape in healthcare, then discuss studies that engaged

clinicians' perspectives on security, and finally situate our approach within the context of research involving other specialized, high-stakes populations.

Technical security challenges in healthcare. Research has documented healthcare systems' technical vulnerabilities, from network threats such as ransomware [14, 37] to attacks on specific networked medical devices [6, 24, 57]. More recent work frames these vulnerabilities in terms of patient safety, treating cyberattacks as a public health crisis [15, 26, 65]. However, this foundational work focuses on the technical systems, largely overlooking clinicians' impact on the threat model.

Security studies with general healthcare populations. Recent research has begun to examine the human element of healthcare security by focusing on security experts, stakeholders, and the broader hospital workforce. For instance, Thompson et al. investigated the threat modeling practices of medical device security experts [64], while Kustosch et al. examined the stakeholder dynamics of device updates [38]. Similarly, Gutfleisch et al. explored general IoT security attitudes among hospital IT and security staff [29]. Broader population studies, such as Ho et al.'s work on email phishing, considered the entire hospital staff rather than focusing specifically on clinical roles [32]. These studies provide critical context on the ecosystem surrounding clinicians but do not focus on the unique constraints of direct patient care.

Security studies involving clinicians. Some work has studied clinicians directly; however, much of this work has focused on the unique postures of specific clinical specialists. For instance, prior work has assessed audiologists' and speech pathologists' security practices regarding telehealth adoption [21, 62] and security practices in emergency departments [58]. Others have investigated the threat perceptions of cardiologists [27] and reproductive healthcare providers [45].

Perhaps closest to our work, others have used high-fidelity simulations to observe clinicians' behavioral responses to security incidents. This research has shown that while some clinicians can manage the physiological fallout of a device compromise, they often fail to identify the cyberattack as the root cause [16, 71]. Similarly, recent workshop-based studies found clinicians tend to operate with an "all-or-nothing" trust model for their technology [59]. While these simulation-based studies are essential, they capture in-the-moment crisis responses rather than the underlying, day-to-day mental models that drive behavior. Our work complements these studies by explaining the potential clinical impacts of a security failure and providing data on a broader range of threats.

Security studies with specialized and high-stakes populations. Finally, our methodology is informed by security research on other high-stakes, specialized populations, such as journalists [48], political campaigns [12], and industrial control system operators [25]. This literature investigates how a population's unique context, expertise, and threat model shape security behavior. We contribute to this area by ap-

¹Inpatient settings are for patients admitted to the hospital, staying overnight.

plying this specialized lens at scale to clinicians, a critical population whose work has life-or-death consequences.

3 Methodology

We now describe the clinician interview protocol, followed by our survey design based on the interviews' findings, and then our analysis methods. Finally, we outline our work's limitations.

Research team. Our research team consists of an advanced practice provider (APP)² and a nurse, which enabled us to design a protocol accounting for the technologies and terminology familiar to clinicians. Their experience was also essential for recruitment to determine where to recruit and the screening criteria. Both providers are licensed to practice in the US. Another research team member has spent several years working with medical device manufacturers, consulting on product security, and has volunteered in a clinical setting at a large hospital. Another research team member, who is not a clinician but has over a decade of security operations and research experience, spent time shadowing clinicians in an Intensive Care Unit (ICU) at a large hospital to familiarize themselves with technological interactions in this setting and probe for possible security and privacy concerns that the clinicians might not have identified. All the scenarios considered and questions asked in our study were generated through hours of discussions between team members with clinical experience and those with security experience to surface concerns that would be missed based on only one perspective.

3.1 Interview Protocol Design

We conducted semi-structured interviews with clinicians to understand the environment in which they work, how they interact with technology, and the ways technology fails. We outline our interview protocol, recruitment process, and analysis. Interviews took place between January and December 2024. Each interview lasted between 45 to 60 minutes, averaging 55 minutes. For consistency, the lead author conducted all interviews. The full script is in our supplemental materials [1].

Screening survey (Fig 1.0). Participants completed a brief screening survey that asked about demographic information and professional experience. To verify participants' credentials, we asked them to provide a link to a professional website, such as LinkedIn or their hospital webpage, or upload a resume. We included questions about their prescribing privileges and roles. In addition to providing context for their responses, the combination of their responses to these two questions served as an effective screener (in the US, only physicians and APPs have prescribing privileges). Responses that contradicted these facts were screened out as fraudulent.

After validating the participants' provided credentials, we contacted all verified participants for interviews.

Participant background (Fig 1.1). We first asked participants to provide an overview of their healthcare experience, their specific role and unit, and any prior background in computer security or IT. This helped build rapport with the participant and provided context for their responses [4, pg. 94].

General security experiences (Fig 1.2). We then asked about participants' organization's experiences with cyberattacks (if any), their interactions with IT security teams, and whether they received cybersecurity training. This provided context for later responses as we considered whether their prior exposure and training influenced their security and privacy perceptions.

Technology use and security failure exploration (Fig 1.3). We addressed our study's central question, examining participants' security and privacy perceptions regarding technologies used in the clinical setting. We asked participants to list all clinical technologies they interact with daily. We asked participants to consider three security failure scenarios for each technology where the data could be viewed by unauthorized individuals (*confidentiality*), was inaccurate or altered (*integrity*), or became unavailable (*availability*).

For each scenario, we provided a brief example, e.g., "An attacker can modify lab values" for the electronic health record's (EHR) integrity scenario. These examples were intentionally simple to help participants focus on potential patient harms rather than technical implementation details. We then asked them to describe how this failure might occur and its impact on patient safety and privacy in a clinical setting. We used these failure scenarios instead of open-ended questions, as we realized after piloting with the clinicians on our team and from our clinical observations that asking about general security concerns would yield limited responses, due to the lack of security expertise or awareness necessary to consider many potential security issues. However, after providing a scenario and example, participants could offer additional depth and identify other threats based on their clinical operations knowledge. This allowed us to understand how security failures might play out in the clinical setting and focus on their perceptions of security risks. Therefore, we chose to use scenario probes to prime discussion, focusing not on whether medical professionals could recognize failures, but instead on how they perceived possible failures. Each scenario was identified and vetted for feasibility with our clinician team members.

During these discussions, if participants became stuck, we used short probes based on common vulnerabilities, such as unauthorized access to patient data (i.e., a confidentiality attack), or operational realities (e.g., software updates, network outages) to stimulate further discussion. For each failure identified by participants, we asked them to discuss how likely they believed it was that an attacker could successfully exploit the failure and what harm this could cause patients, given other clinical setting controls. In these discussions, when partici-

²The APP role includes Physicians' Assistants and Nurse Practitioners




Figure 1: Diagram of the interview and survey flow.

pants were asked about availability, we also asked them about their experiences with system downtime, i.e., when hospital systems are not functioning, and established alternative procedures (formal or informal) when this occurs. We asked about these events specifically due to their similarity to security events and our clinician team members’ and pilot participants’ assessment of their pervasiveness in practice.

Communication tools (Fig 1.4). If time allowed, we briefly discussed the tools used for care team communication (e.g., pagers, personal devices, healthcare apps). We explored this area because these tools are integral to coordinating patient care, but often represent a less-controlled technology ecosystem, thus presenting a vector for security and privacy failures.

General perceptions and wrap-up (Fig 1.5). The interview concluded with broader questions to gain a deeper understanding of participants’ general healthcare security and privacy mental models. We specifically asked them to identify what they perceive as the weakest links in hospital security, whether they see themselves as targets for attack, their sense of personal versus organizational responsibility for security, and any perceived conflicts between cybersecurity and patient safety.

Pilot. Before the main set of interviews, we conducted two pilot interviews with clinicians. During these pilots, we solicited feedback on question clarity, flow, and relevance. Following these two initial pilots, we significantly shortened our protocol, as the interviews had taken over 90 minutes. We expected, based on discussions with clinicians, that clinicians would have limited time to participate. Therefore, we trimmed the interview so it could be completed in 60 minutes. To reduce the interview time, we chose to discuss security controls (e.g., authentication) in the context of the scenarios, rather than in a separate, more in-depth discussion. The final version is reflected in the protocol we outlined above. We then piloted the refined interview script with two additional participants. No significant changes were needed after these pilots, confirming the protocol’s suitability. Due to no protocol changes, the latter two pilot interviews are included in our results.

3.2 Survey Protocol

To complement our interviews’ insights, we conducted an online survey targeting a larger and more diverse group of clinicians across multiple countries. We included other patient-

facing roles with fewer responsibilities than the clinicians interviewed, such as nursing assistants and lab technicians, which are also referred to as Allied Health Professionals (AHPs). The full survey is provided in our supplemental materials [1], we discuss the consent process in our ethics section.

Eligibility Screening (Fig 1.A). To ensure our data was collected from qualified clinicians, we implemented a multifaceted screening process. Participants were required to pass (i.e., answer 60% of questions correctly) a knowledge check consisting of 10 Basic Life Support (BLS) questions, which is a necessary certification from the American Heart Association for many clinicians in the United States, adapted from standard guidelines and practice tests [2] to verify baseline clinical knowledge.

Following this, we implemented a rigorous manual validation process acting as a functional attention check. Our clinician authors reviewed open-ended responses to identify nonsensical or potentially AI-generated text that lacked “clinical reality.” We further validated data quality by cross-referencing participants’ self-reported roles against their other responses to identify inconsistent or impossible combinations. For example, we removed participants whose reported role conflicted with legal prescribing privileges (e.g., a US registered nurse claiming to prescribe medication) or whose role did not align with their claimed technology usage (e.g., an AHP using medication dispensing cabinets).

While this stringent validation process may have disqualified some qualified participants, it was necessary to ensure high-quality data from individuals with legitimate and relevant clinical experience. We ensured our screening methods were appropriate for all countries targeted for recruitment (US, UK, and Canada) by consulting clinicians in each.

Medical Experience (Fig 1.B). Next, we collected professional background information. We asked participants to select from a predefined list of all clinical technologies they interact with daily. This list was derived from our interviews. Nine technologies were presented (shown in Table 1). Finally, we asked participants to report their clinical role, i.e., registered nurse (RN), physicians (MD/DO³), APP, AHP, primary specialty (e.g., Cardiology, Emergency Medicine), years of

³We refer to physicians interchangeably with their degree, Doctor of Medicine (MD) or Doctor of Osteopathic Medicine (DO), these degrees are functionally equivalent.

Technology	Description
Records	
EHR	Electronic Health Records (EHR) are comprehensive, longitudinal digital records of a patient’s health information. They serve as the central hub for clinical workflow, containing everything from demographics and notes to orders, medication records, and lab results.
Imaging	
Imaging Devices	Generates diagnostic images of internal structures, e.g., Magnetic Resonance Imaging (MRI), Computed Tomography (CT), X-ray.
PACS	Picture Archiving & Communication System (PACS) are a medical imaging technology providing enterprise-wide storage, retrieval, management, and distribution of images. PACS are networked systems that handle images and integrate with EHRs.
Interventional	
Anaesthetic Equip.	Integrated anaesthesia workstations that deliver inhalational anaesthetic agents and medical gases. They incorporate physiological monitors to continuously track vital functions such as Heart Rate (HR) and Blood Pressure (BP), during surgery.
Digital Surg. Tools	Advanced, computer-assisted systems that enhance surgical precision and capability.
Interventional Dev.	Active treatment or life support devices (e.g., infusion pumps, ventilators, defibrillators).
Medication	
Medication Cabinets	Systems for dispensing medications; integrates with EHR for inventory/safety.
Monitoring	
Non-Continuous	Devices for periodic ‘spot-check’ vitals (e.g. BP); often manual upload to EHR.
Continuous	Real-time streaming of physiological parameters (ECG, ICU monitors); triggers alerts.

Table 1: Technologies used in the survey scenarios.

healthcare experience, and, for physicians, their career stage (resident or attending).

Scenarios (Fig 1.C). The main survey component focused on assessing perceived risks for specific technologies. Each participant was randomly assigned two technologies they indicated they use daily. For each of the two assigned technologies, we first asked about their experience with system downtime, asking how often they had encountered planned or unplanned downtime with each technology in the past year, given that this is a recurring issue with implications for patient safety [40]. Because downtime is a direct and common, real-world example of an availability failure, this question provides context for their evaluation of other security failures.

After we asked about system downtime, participants were presented with three hypothetical failure scenarios corresponding to breaches of *confidentiality*, *integrity*, and *availability* (i.e., six scenarios per participant). The specifics of the CIA scenarios were tailored to each technology, based

on examples commonly provided in our interviews (e.g., the EHR-Integrity example, with altered laboratory values). To mitigate ordering effects, the order of the two technologies and the three CIA scenarios was randomized. This selection was balanced to ensure technologies were assigned evenly across participants.

For each scenario, we asked participants to rate their perception of the scenario’s *likelihood* on a 5-point Likert scale from “Extremely unlikely” to “Extremely likely.” To quantify potential patient safety impact, we asked participants to rate the *worst potential physical harm* on a 5-point scale from “Negligible” to “Catastrophic,” adapted from FDA guidance for medical device cybersecurity [11]. Participants were also given the option to select from a list of *other potential harms* (e.g., financial, emotional). Finally, we asked participants to explain their ratings using a free-text response.

Cybersecurity Controls Challenges (Fig 1.D). Prior work has shown increases in tech-mediated, less direct care positively correlated with clinician burnout [33, 49]. We sought to investigate whether and how security controls contribute to this burnout effect with the aim of quantifying the tension between security controls (e.g., passwords) and their patient care impact. We asked participants to rate the impact on care delivery of the controls interviewees discussed, on a 5-point scale from “Definitely impedes care” to “Definitely improves care.” Our controls were grouped by *login methods* (e.g., Username/Password, Badge access), *communication tools* (e.g., Hospital-issued phones, Secure consumer messaging apps), and *remote hospital access tools* (e.g., via a hospital-issued device or a VPN for remote EHR or scheduling access). For each group, participants were asked to explain their ratings.

Cybersecurity Experience (Fig 1.E). Finally, we gathered information on participants’ general cybersecurity background and attitudes to provide context for their responses. We asked participants about their prior cybersecurity education and whether they had experienced any cyberattacks at work. We measured their security attitudes using the SA6 scale [22].

3.3 Recruitment

Interview Recruitment. To recruit interviewees, we targeted clinicians from inpatient, ICU, and emergency department (ED) units. We chose these settings because clinicians there typically manage patients with more complex or risky medical problems and interact with a broader range of medical technologies compared to outpatient settings. We only recruited participants from the US for this exploratory study component. We did this to limit the complexities from different national regulations, clinical practices, and funding models.

We sought to recruit a mix of physicians, APPs, and registered nurses. Clinicians in each role perform distinct, yet complementary duties. For example, physicians and APPs make diagnoses, decide how patients will be treated, and per-

form complex medical procedures. Conversely, nurses handle routine care tasks and more frequently interact with and monitor patients. These roles cover the range of people interacting with medical technologies and making patient care decisions. By recruiting across roles, we gain a more holistic view of clinical technology use and security perceptions.

We recruited interview participants through the research team’s professional networks, LinkedIn posts, and medical professional development organizations. Additionally, due to the challenges of recruiting from a specialized population, we employed snowball sampling, where interviewees were asked to share information about the study with colleagues. Participants received a \$40 USD gift card for their time.

Survey Recruitment. We broadened recruitment to include participants from the US, UK, and Canada. We added the UK and Canada to capture potential differences from differing healthcare models. We recruited survey participants through Prolific [53]. Recruitment occurred between January 30 and February 22, 2025. Participants who failed screening were directed out of the survey, receiving \$0.50 USD. Qualified participants who completed the survey received \$15.00 USD. Unlike in the interviews, we were not able to ask participants for their CV/credentials without violating Prolific’s Terms of Service; we therefore relied on our rigorous checks to verify our participants’ clinical knowledge as discussed in our Eligibility Screening in Section 3.2.

3.4 Analysis

Interview analysis. Each interview was recorded and transcribed using Zoom’s automated transcription service. We analyzed interview transcripts following an iterative open coding approach [9, 13]. An initial codebook was developed based on the research questions, the script’s structure, and the research team’s collaborative review of two transcripts. Then, two researchers independently coded interviews in rounds of two. After each round, we calculated inter-rater reliability (IRR) for each variable using Krippendorff’s alpha (α) [36]. We assessed agreement per participant response to assess the differences in context between each scenario. This allowed us to account for cases where participants perceived security failures for some scenarios as likely and others as unlikely. We note we did not calculate IRR for the *device* category because these codes were taken directly from the text and no subjective assessment was necessary [46]. This iterative process continued for four rounds until strong agreement (i.e., $\alpha \geq 0.80$) was achieved for each code category [36]. A single coder coded the remaining two interviews. The full codebook is given in our supplemental materials [1].

Survey quantitative data analysis. To investigate the factors influencing clinicians’ perceptions of harm, likelihood, and impact on care, we used mixed-effects ordinal logistic regression models as each dependent variable is ordinal [55]. We

included explanatory variables for the scenarios’ technology and CIA property, as well as participant characteristics. Each model’s included variables are given in Table 2. We did not perform variable selection as all variables were relevant to our research questions. We also developed a secondary model examining the interaction between the technology and CIA property, allowing us to create a more nuanced understanding of the devices’ context. For all models, participant ID was included as a random intercept to account for non-independence of responses from the same individual. Model comparisons using ANOVA were conducted to confirm the appropriateness of including random effects [67]. Because we randomly selected participant scenarios by technology, then performed our analysis by technology *group*, 11 of 303 participants were assigned two scenarios about technology in the same group. We conducted a sensitivity analysis to assess our conclusions by descriptively comparing regression models with and without these 11 participants for harm and likelihood. We compared the models’ direction, magnitude, and statistical significance estimates, as is recommended practice [56, 63]. We confirmed there were no substantive changes in findings.

Survey open-ended response analysis. We employed the same iterative open coding approach used in the interviews to analyze the survey open-ended responses. We developed the initial codebook after reviewing 50 responses. Two researchers then independently reviewed responses in rounds of 50, after each round calculating Krippendorff’s α [30], and updating the codebook as necessary. Strong agreement ($\alpha > 0.80$) was reached for all codes after five rounds. One researcher coded the remaining 1490 responses. The final codebook and α values are in our supplemental materials [1].

3.5 Limitations

While our mixed-methods approach provides a rich, multi-faceted view of clinicians’ security perceptions, here we discuss our work’s limitations and offer avenues for future work.

First, our interview data was collected exclusively from inpatient clinicians practicing in the United States. We sought to test the generalizability of our interview findings by designing our large-scale survey to span three countries with notably distinct healthcare models: the largely private, market-based system of the US; the public, single-payer National Health Service (NHS) in the UK; and the publicly funded, provincially administered system in Canada. These systems cover a wide range of possible healthcare systems and affect a vast number of patients. However, we are focused on English-speaking countries; our sample may not account for all potential factors.

Next, there is an inherent challenge in studying perceptions of abstract failures. The choice of fixed scenarios based on the CIA triad may limit the scope of threats discussed. However, we found that these scenarios served as effective prompts as participants frequently used them as jumping-off points to identify and elaborate on broader threats and operational

Variable	Definition & Levels	H	L	C	
Outcomes	Harm	Perceived physical harm from a scenario. (Ordinal: <i>Negligible to Catastrophic</i>)	X		
	Likelihood	Perceived likelihood of a scenario. (Ordinal: <i>Ext. Unlikely to Ext. Likely</i>)		X	
	Impact on Care	Perceived impact of a control on care. (Ordinal: <i>Impedes to Improves</i>)			X
Predictors	Technology	Technology in scenario. (Base: <i>Records</i>)	X	X	
	CIA	Security property violated. (Base: <i>Confidentiality</i>)	X	X	
	Control Type	Control. (Base: <i>Username/Password</i>)			X
	Role	Clinician’s role. (Base: <i>Nurse (RN)</i>)	X	X	X
	GEO	Geographical location. (Base: <i>US</i>)	X	X	X
	Unit	Clinical unit. (Base: <i>Inpatient</i>)	X	X	X
	Security Training	Received security training. (Base: <i>False</i>)	X	X	
	Prior Attack Exp.	Prior attack experience. (Base: <i>False</i>)	X	X	
	Tech:CIA	Interaction: technology and CIA.	X	X	

Table 2: Variables Used in the Mixed-Effects Regression Models. The final columns indicate which model used the variable: Harm (H), Likelihood (L), or Controls (C).

challenges beyond the examples provided. Additionally, social desirability bias could lead to under-reporting of insecure behaviors, though the frequent admission of workarounds suggests candor. For integrity failures in particular, we presented hypothetical failure scenarios few clinicians had experienced. This primed them to consider specific failures, and their responses reflect their reasoning about these ideas and are not grounded in their own specific experience. However, as our goal was to understand their operations, mental models, and risk perceptions, this approach was necessary. The nuance of their responses suggests they understood the scenarios and, as we discuss in our results, although not all these failures were directly experienced, they have likely encountered analogous events, such as technology downtime and data manipulation due to system errors. At a minimum, our results provide a valuable framework by applying clinicians’ experiences of real-world healthcare practices to security ideas.

Furthermore, while our work discusses clinicians’ subjective risk perceptions, these should not be viewed as objective security failure frequency. We do not claim that because clinicians perceive a failure as unlikely, it is therefore unlikely in practice. Instead, we argue our findings about clinicians’ perceptions are a crucial reality because they directly shape their behaviour. Future work should seek to correlate these findings with ground-truth data from incident reports. However, the identified descriptions of hospital operations and realities provide valuable directions to support interventions.

Finally, reliably identifying Large Language Model (LLM) generated text remains an open challenge for the community.

	PID	Exp.	Work Setting
Nurses	RN2I	5-10	Academic Medical Center
	RN4I	1-5	Academic Medical Center
	RN8I	1-5	Academic Medical Center
APP	APP3I	5-10	Private Practice, Community Hospital
	APP7E	5-10	Two Academic Medical Centers
Physicians	MD1I	1-5	Academic Medical Center
	MD5E	10-20	Academic Medical Center
	MD6E	1-5	Academic Medical Center, Community Hospital, Rural Hospital
	MD9I	30+	Academic Medical Center, Private Practice, Teleradiology Practice
	MD10I	5-10	Rural, Community, Academic Medical Center
	MD11I	5-10	Academic Medical Center
MD12I	30+	Locums (mainly Rural and Community)	

Table 3: Interviewees’ role/unit, years of relevant experience, and workplace. PID shows the participant’s role (RN - Registered Nurse; APP - Advanced Practice Provider; MD - Physician), and unit (I - Inpatient; E - ED/ICU).

While we performed rigorous manual validation as described in Section 3.2, including verifying the clinical plausibility of qualitative responses and checking for consistency between roles and reported capabilities, we cannot guarantee the absolute exclusion of all AI-generated content.

4 Participants

When reporting results, we use I (at most 12) to denote the number of interviewees and S (at most 303) for the number of survey respondents. Because participants each discussed multiple scenarios, the number of scenarios is higher than the number of participants. We denote the number of per-scenario responses with I_S (at most 237) and S_S (at most 1740).

Interviewees spanned the spectrum of inpatient and ICU/ED settings, and experience levels. Seventeen people volunteered for our interviews. Of these, 12 were qualified inpatient clinicians practicing in the US. The other five were screened out as they all provided obviously fake LinkedIn pages created just before submitting their survey response, with incorrect roles, information that did not match their survey response, or information that conflicted with their survey responses (e.g., nurses with prescribing privileges). The remaining 12 interviewees spanned the range of roles, years of experience, and clinical environments seen in prior large-scale clinician surveys [20, 42]. We interviewed seven physicians, two APPs, as well as three nurses across various specialties. Years of experience ranged from more than 30 years ($I = 2$) to more recent graduates with 1-5 years ($I = 4$), with a median of 10 years. While medical school and nursing school expose clinicians to clinical environments, we only considered professional experience, e.g., starting from residency. Interviewees primarily worked with academic medical cen-

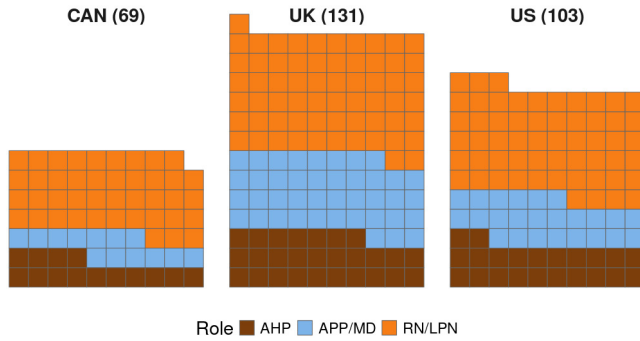


Figure 2: Survey respondents by role across US, UK, and Canada. Each box represents a participant.

ters⁴ ($I = 10$), community hospitals ($I = 4$), private practices ($I = 2$), locum physicians (who rotate among hospital practices in underserved areas) ($I = 1$), and rural healthcare environments ($I = 3$). Four physicians and both APPs work at multiple sites (across different organizations) and provided us insights from the varied practices at each. Participant details are summarized in Table 3. Despite our efforts, almost all participants were recruited from professional networks ($N = 11$), and one participant was recruited via snowball sampling.

Most survey participants were RNs and well-distributed between countries. Our recruitment and screening yielded 303 qualified clinicians from 788 responses. We removed the following participants: 184 failed the BLS screener, 236 did not hold a qualifying role (such as an occupational therapist) or had responses that did not pass validation checks, and 66 did not complete the survey. Participants were geographically distributed, with most from the UK ($S = 131$, 43%), followed by the US ($S = 103$, 34%) and Canada ($S = 69$, 23%). Most survey respondents were nurses (RN/LPN) ($S = 162$, 54%). Other participants were physicians ($S = 57$, 19%), APPs ($S = 21$, 7%) and AHPs ($S = 63$, 21%). Again, participants had a range of clinical experience, from newly practicing clinicians with less than 5 years ($S = 85$, 28%) to veterans with over 15 years ($S = 101$, 33%). This diversity extended to their clinical settings, with participants working across ICU/ED ($S = 41$), inpatient ($S = 151$), and outpatient ($S = 111$) units. The average completion time was 23.1 minutes.

Clinicians’ security education was limited to workplace training. Eight interviewees and 69% ($S = 208$) of survey respondents reported receiving HIPAA or cybersecurity training. As we might expect, no interviewee had any formal computer security education or work experience, and 18% ($S = 57$) of survey respondents reported some cybersecurity experience, which included taking a cybersecurity or IT course or certification, or having related work experience. On average, survey

⁴Academic medical centers serve as both magnet hospitals for the regions in which they are located and represent the nexus of complex patient care, medical training, and clinical research.

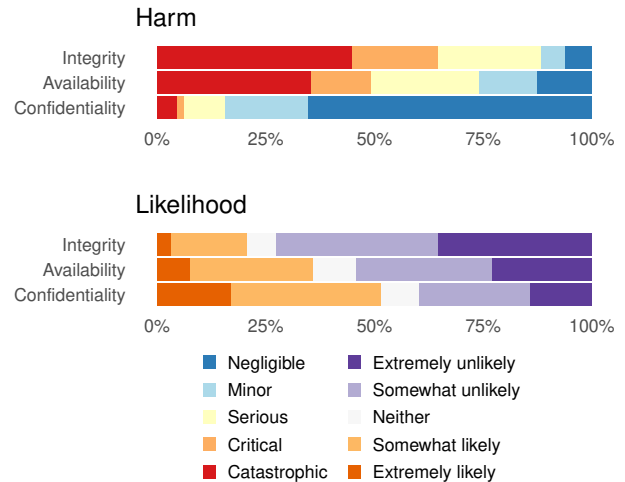


Figure 3: Perceived harm (top graph) and likelihood (bottom) overall. We show the breakdown by technology in Fig 5.

participants had similar SA6 scores (3.8) to prior general US population samples (3.6-4.0) [22]. While interviewees scored lower than average (3.3), we do not believe this impacted our results, given the similarity in other responses to our survey.

Most participants had no prior cyber incident experience. Only two interviewees reported having experienced a cybersecurity incident at work. Similarly, 26% ($S = 79$) of survey respondents reported experiencing an incident. This does not necessarily mean participants had not experienced an incident, but may be unaware. Three additional interviewees speculated they had experienced an incident, but were uncertain, as many of these events are not communicated to clinical staff. RN8I explained, “I actually found out about [the incident] through a patient who works for the IT department of the hospital.”

5 Results

We now discuss our interviews and survey results. We emphasize that these findings reflect clinicians’ subjective *perceptions* of risk rather than objective failure frequencies. Furthermore, to account for the complexity of the clinical environment, we modeled these perceptions using mixed-effects ordinal logistic regression. Consequently, all regression estimates (Odds Ratios) reported below represent the independent association of a factor with the outcome *ceteris paribus*, that is, while statistically holding all other covariates (such as role, unit, and geography) constant [61]. This ensures that reported effects for specific technologies or failure types are not merely artifacts of demographic or operational confounders. For brevity, we will state our results as all else equal. Participants’ perceived harm and likelihood are in Figure 3. Their perception of the impact of login methods is in Figure 4. The models with interactions for perceived harm and likelihood

are in Table 4. The non-interactions model is in Appendix A as the results are similar. The control model is in Table 5.

5.1 Technologies Discussed

EHRs are central to clinical workflows. The EHR was the most central technology discussed, with all 12 interview participants and all 303 survey participants describing it as integral to daily work. These systems, such as Epic, serve as the primary platform for clinical tasks, including reviewing patient histories, placing medication and intervention orders, and documenting care. As MD1I explained, “I use Epic constantly. I use it on my phone. . . I put in orders on my phone. . . In the ICU, like the surgical ICU or the neuro ICU, [vitals monitors] are much more connected to Epic.”

Patient Monitoring Devices. Monitoring devices, e.g., continuous EKG and vitals monitors, were discussed by interviewees ($I = 9$) and used by most survey respondents ($S = 253$, 83.8%). Clinicians use this for data to make care decisions.

Interventional Devices. Interventional devices, such as IV infusion pumps, dialysis machines, and anesthesia equipment, were discussed by seven interview participants and just over half of survey respondents ($S = 157$, 52%). These technologies are used to provide treatments or medications.

Imaging and PACS. Both imaging devices and the associated data storage system, i.e., PACS, are common for certain specialities, such as radiology. These include both the devices that generate imaging, such as MRIs or CT scanners, as well as the software used to store and view these images for clinical diagnostics. Seven interviewees discussed these technologies, and slightly less than half of survey respondents use them ($S = 122$, 40.4%). This is to be expected as physicians and APPs are more likely to interact with PACS than other roles, and our survey had a lower proportion of these roles.

Medication Dispensing Cabinets. Medication cabinets are used to provide clinicians on their floor with medications, eliminating the need to visit the pharmacy. These cabinets are usually electronically controlled and have checks to ensure the correct medication is dispensed, as the wrong medication can be fatal. Over a third of survey participants use these daily ($S = 115$, 38.1%), 80.9% of participants who used medication cabinets were nurses. While interviewees discussed these ($I = 9$), our physicians and APPs had limited use, as nurses are primarily responsible for administering medication.

5.2 Confidentiality Failures: A Frequent but Low-Severity Concern (RQ1)

Next, we discuss each security failure type in turn and how clinicians viewed each in the context of the technologies described above. We found clinicians perceived confidentiality failures as the most likely type of security failure (51.4% said “Extremely” or “Somewhat” likely). In our multivariate model,

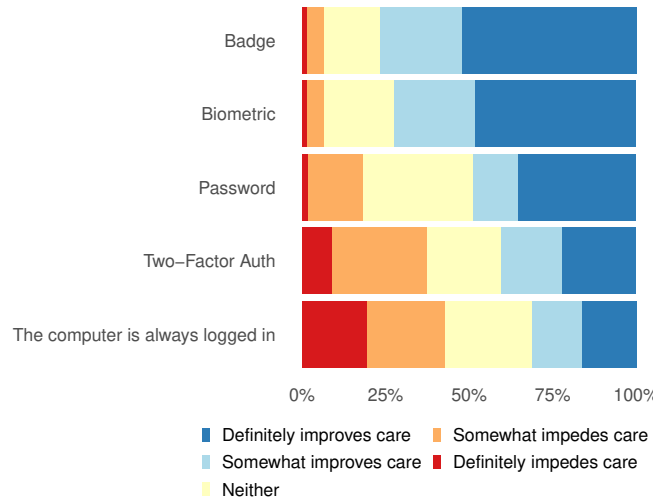


Figure 4: Perceived affect on care for login methods.

holding all else equal, integrity failures ($OR : 0.2, p < .001$) and availability failures ($OR : 0.5, p < .001$) were seen as significantly less likely (Table 7).

EHR confidentiality failures are often due to non-technical records release. While EHR confidentiality failures were perceived to be less likely to occur than the other technologies, except Medication Cabinets (Table 4), respondents still said this type of failure was “Somewhat” or “Extremely” likely (39.6%, $S_S = 42$). While some interviewees said authentication mistakes like not logging out of Epic “could happen pretty easily” (MD11I), they indicated that this was not a major concern ($I = 4$). Instead, participants explained that much of the risk associated with confidentiality failures was due to analog leakage of patient records ($I = 6$). For example, MD11I explained, “everyone prints a patient list every morning. . . like names, date of birth, age. . . sometimes people leave their list somewhere by accident or drop them.” These confidentiality failures can lead to harm through the disclosure of sensitive information. For example, a midwife (APP3I) talked about a recent example of “someone coming in for her first pregnancy visit. She was a teenager, and the father did not know she was pregnant. . . But his phone number was the contact information for this patient and this patient was 20 minutes late, so we called him and. . . said, ‘Oh, your daughter is late for her appointment.’ And he was like, “What appointment?” And so that was kind of, the mother got quite upset with us, but it was his phone number in there.”

Confidentiality perceptions for imaging, monitoring, and interventions. Compared to the EHR, confidentiality failures in specialized medical devices were perceived as significantly less harmful. Both imaging ($OR : 0.4, p = 0.006$) and monitoring systems ($OR : 0.2, p < .001$) were perceived to have lower risk (all else equal) because they produce isolated data points lacking the rich, sensitive context of the EHR. As one

Factor	Value	Harm			Likelihood		
		OR	p-Value	CI	OR	p-Value	CI
Technology	EHR	-	-	-	-	-	-
	Img.	0.4	0.006	0.2, 0.7	1.4	0.229	0.9, 2.2
	Inter.	0.7	0.172	0.4, 1.1	2.6	<0.001	1.7, 4.1
	Cab.	0.7	0.189	0.4, 1.1	0.4	0.004	0.3, 0.7
	Mon.	0.3	<0.001	0.2, 0.4	3.8	<0.001	2.5, 5.7
CIA	C	-	-	-	-	-	-
	I	54.4	<0.001	33.6, 88.1	1.0	0.897	0.7, 1.6
	A	6.4	<0.001	4.1, 9.8	2.1	0.003	1.4, 3.2
Role	RN	-	-	-	-	-	-
	AHP	1.0	0.818	0.7, 1.3	1.2	0.391	0.9, 1.7
	MD/APP	0.8	0.316	0.6, 1.1	0.7	0.1	0.5, 1
Geo	US	-	-	-	-	-	-
	Canada	0.8	0.393	0.6, 1.2	1.0	0.794	0.7, 1.3
	UK	1.1	0.621	0.8, 1.4	0.9	0.641	0.7, 1.2
Unit	Inpatient	-	-	-	-	-	-
	Outpatient	0.8	0.133	0.6, 1.0	0.8	0.195	0.6, 1.1
	ED/ICU	1.5	0.069	1.0, 2.2	1.6	0.042	1.1, 2.4
SA6	I	-	-	-	-	-	-
	+1	1.1	0.205	1.0, 1.3	0.9	0.209	0.7, 1.0
Sec. Train.	False	-	-	-	-	-	-
	True	0.8	0.108	0.6, 1.0	0.9	0.376	0.7, 1.1
Attack Exp.	False	-	-	-	-	-	-
	True	1.3	0.1	1.0, 1.7	1.3	0.138	1.0, 1.8
Technology x CIA	Img : I	0.9	0.825	0.5, 1.8	0.4	0.016	0.2, 0.8
	Inter. : I	1.4	0.376	0.7, 2.7	0.1	<0.001	0.1, 0.2
	Cab. : I	0.8	0.628	0.4, 1.6	0.3	0.006	0.2, 0.7
	Mon : I	1.0	0.916	0.6, 1.9	0.03	<0.001	0.02, 0.1
	Img : A	4.2	<0.001	2.2, 8.1	0.5	0.067	0.3, 0.9
	Iter. : A	8.2	<0.001	4.4, 15.4	0.08	<0.001	0.05, 0.2
	Cab : A	2.4	0.027	1.3, 4.5	0.6	0.117	0.3, 1.0
Mon. : A	8.2	<0.001	4.6, 14.6	0.05	<0.001	0.03, 0.09	

Statistically significant values ($p \leq 0.05$) are **bolded**

- Base case (Odds ratio defined as 1)

Img - Imaging; Inter; Interventions; Cab. - Medication Cabinets; Mon - Monitoring; C - Confidentiality; I - Integrity; A - Availability

Table 4: Results from mixed effects regressions with interactions. An OR > 1 indicates greater harm or higher likelihood relative to the baseline, holding other factors constant.

participant noted, a standalone X-ray is not particularly revealing. Similarly, interventional devices were not a confidentiality concern, with 77.9% of survey respondents anticipating only minor or negligible harm from a breach. This is because clinicians widely believe these devices do not store patient-identifiable data, with MD12I stating there is “No PHI on that system at all.”

Physical design mitigates medication cabinet confidentiality risks. Participants expressed low concern about confidentiality failures for automated medication cabinets; only two interview participants discussed this, and both stated it was less concerning and did not pose any safety risks. Similarly, 78.7% ($S_S = 70$) of survey respondents indicated these failures posed negligible or minor harm. Participants attributed

the low likelihood of harm to the system’s physical placement and design, which naturally limit opportunistic viewing. RN8I explained, “The Pyxis [medication cabinet] is very small, and the screen is very small, and they’re all in little nooks.”

5.3 Managing Confidentiality: Access Controls and Their Challenges (RQ2)

While confidentiality failures were viewed as less concerning, they were also the failure type where most security controls, i.e., access controls, were discussed. On a positive note, all forms of access control were seen as more beneficial for patient care than providing no access control (43% said no access control impedes care, compared to 18% for password, 7% for biometrics, 6% for badges), though two-factor authentication was not much better (37% said it impedes care). Participants’ views and the challenges they faced from access control were shaped by the nature of clinical work.

Mobility issues affecting access controls. One underlying challenge for authentication that clinicians reported was the inherent mobility of clinical work, which makes authentication a frequent, workflow-interrupting event. Clinician mobility occurs at three levels: intra-unit, inter-unit, and inter-organizational. On a given shift, a clinician moves between numerous computers, whether in patient rooms, to workstations in hallways, or on mobile carts. This intra-unit movement requires repeated authentication throughout the day. Due to this mobility, clinicians exhibited a clear preference for methods that streamline workflows. For example, MD12I lauded badge access as the “smoothest way,” explaining, “you logged in once in the morning, and then any other time you sat down at a computer, you just took your badge, tapped it and bam you were right where you left off.” This preference is seen in our regression model, where holding other factors constant, badge access ($OR : 2.8, p < 0.001$) and biometrics ($OR : 2.2, p < 0.001$) were perceived as significantly improving care compared to passwords and two-factor authentication (non-overlapping confidence intervals).

Many clinicians ($I = 6$), such as locum physicians (e.g. MD12I) or travel nurses, work for multiple healthcare organizations (inter-organizational mobility), each with its own distinct credentials and security protocols. Inter-organization mobility introduces significant challenges for clinicians as they have to remember several passwords (or reuse passwords). For example, MD12I, who works across multiple sites, keeps his passwords in a “trusty book” because “there’s no way to keep track of it.” In our model, all else equal, Two-Factor Authentication (2FA) was also perceived as significantly impeding care delivery compared to passwords ($OR : 0.467, p < 0.001$) and other forms of authentication (non-overlapping confidence intervals). MD10I lamented, “I think I have 3 different apps associated with 2 factor authentication. . . is a lot to manage.”

Authentication introduces barriers in emergencies. An-

Factor	Value	OR	p-Value	CI
Control Type	Username/Password	-	-	-
	2FA	0.5	<0.001	0.4, 0.6
	Badge Access	2.8	<0.001	2.1, 3.6
	Biometrics	2.2	<0.001	1.7, 2.9
	No Login Required	0.3	<0.001	0.2, 0.3
Role	RN	-	-	-
	AHP	1.2	0.311	0.9, 1.7
	MD/APP	0.8	0.206	0.6, 1.1
Geo	United States	-	-	-
	Canada	0.9	0.426	0.6, 1.2
	United Kingdom	0.6	<0.001	0.4, 0.7
Unit	Inpatient	-	-	-
	Outpatient	0.9	0.412	0.7, 1.1
	ED/ICU	0.9	0.568	0.6, 1.36
SA6	1	-	-	-
	+1	1.4	<0.001	1.2, 1.7
Security Training	False	-	-	-
	True	0.8	0.127	0.6, 1.0
Prior Attack Experience	False	-	-	-
	True	1.0	0.858	0.8, 1.3

Statistically significant values ($p \leq 0.05$) are **bolded**
 - Base case (Odds ratio defined as 1)

Table 5: Regression results regarding the perceived impact of security controls on patient care. An OR > 1 indicates the control is perceived to improve care relative to the baseline (all else equal), while < 1 indicates a perception of impeding care.

other cause of authentication challenges was the frequent occurrence of emergency situations in the clinical environment. While some interview participants noted computers are left logged in, this is not necessarily a function of laziness, but as a response to emergencies. Many hospital systems will automatically log out users after a period of inactivity to prevent confidentiality failures, but this can get in the way during emergencies as RN8I explained, “And then in an emergency, we need one computer up and then say no one touches it for a few minutes and then it logs everyone out, its can be dangerous. I mean, maybe there should be something like a button that says emergency, keep it open or something.”

5.4 Integrity Failures: Catastrophic but Perceived as Less Likely (RQ1)

Integrity failures—the malicious alteration of clinical data or device function—were perceived as more harmful than confidentiality failures in our model, all else equal, (OR : 48.9, $p < .001$, Table 7) and availability failures (non-overlapping confidence intervals) failures. However, they were also consistently rated as the least likely (OR : 0.2, $p < .001$), all else equal, compared to the confidentiality baseline and non-overlapping confidence intervals with availability. Unlike con-

fidentiality and availability, there has been no evidence of a real-world integrity attack. However, like the others, participants did report experiences with similar failures in practice, albeit not due to a malicious actor.

Altered EHR data can lead to fatal outcomes. Integrity failures’ high perceived harm was clear for most technologies clinicians use. The modification of patient records was a widespread concern among interview participants ($I = 9$), who described how altered data could lead to catastrophic outcomes. MD1I provided a stark example, saying “If you could show me the potassium was low every day... I have no way of independently verifying the potassium concentration in someone’s blood... altering people’s electrolytes [by administering potassium] can put them into cardiac arrhythmias, some of which are fatal.” Survey responses underscored this fear; in open-ended responses about EHR integrity failures, participants often described risks that could be categorized as uncontrolled situations ($S = 67$, 65.0%), or said the failure could lead to patient death ($S = 33$, 32.0%). We defined uncontrolled situations as the incident either (a) misled a clinician into providing unnecessary care, such as administering an unneeded drug based on false data, or (b) precipitated an uncontrolled clinical event, such as a patient’s condition worsening without staff being readily aware.

Manipulated images could cause severe misdiagnosis. Imaging system integrity failures were seen as potentially very severe. APP7E, when asked about a chest x-ray manipulated to hide a potential collapsed lung, explained “If someone has a large pneumothorax and the image is edited to make it look like it’s not there or it’s very small... you’re looking for other causes, and you could send someone to the scanner and they could die on the scanner.” The same participant noted the harm is context-dependent, describing an inverse scenario where an image is manipulated to appear worse: “worst case they get a chest tube. That has its own slew of complications of course, but it’s something that isn’t the end of the world.” Similarly, survey responses showed an imaging system’s integrity failure was perceived as just as harmful as one affecting the EHR, all else equal (OR : 0.9, $p = 0.825$), but considered an attack to be less likely than on the EHR (OR : 0.4, $p = 0.016$).

Medication cabinet procedural checks provide a safety net against integrity failures. For medication cabinet integrity failures, e.g., the system providing incorrect medication, survey respondents often mentioned patient death ($S = 23$, 25.3%) as a potential consequence. However, interviewees explained existing secondary checks, which are already standard practice to avoid errors, mitigated this risk ($I = 5$). This included procedural safeguards, such as scanning a medication’s barcode and a patient’s wristband with a scanner that is separate from the medication cabinet. As RN8I explained, if someone grabbed the wrong medication, when they “scan the med, then it’ll be like, ‘Well, this isn’t the

Tylenol,’ and it’ll not match the barcode.” This is supported by the regression results, where the likelihood of medication cabinet integrity failures was seen to be significantly less likely than confidentiality failures ($OR : 0.3, p = 0.006$) all else equal.

Integrity failures leading to clinical errors. Monitoring and interventional device integrity failures were perceived as posing a high risk of physical harm, as corrupted data could lead clinicians to intervene inappropriately or fail to act when necessary. For interventional devices, respondents cited dire outcomes like uncontrolled situations ($S_S = 50$) and patient death ($S_S = 29$). The danger from false monitoring data was most commonly associated with uncontrolled risk to the patient ($S = 110$). MD5E who, when considering a faulty blood pressure reading, noted if a monitor falsely showed a blood pressure of 180, “I would think this patient is really hypertensive. Let me give them a [unnecessary] medication. . . [which] could cause significant harm.”

5.5 Managing Integrity: Clinical Gestalt as a Defense (RQ2)

While access controls are a defense against integrity attacks in general, our participants did not report having the administrative privileges necessary to alter device configurations or data. Instead, editing permissions are reserved for specialized staff, like biomedical engineers, thus our participants did not perceive the authentication controls they use as integrity defenses. Instead, they described relying on their “clinical gestalt,” a holistic and intuitive assessment of the patient. This gestalt is formed by synthesizing multiple, disparate data streams—including direct observation, vital signs, and recent lab results. In the context of security, this gestalt serves as a cognitive defense mechanism, enabling clinicians to perform a rapid plausibility check and identify data that is inconsistent with the patient’s overall state. This reliance on clinical gestalt to manage security failures has also been observed in simulation studies involving clinicians [16, 71].

Clinical gestalt helps identify egregious errors. While concerned about subtle data manipulation, participants expressed confidence in their ability to use their clinical judgement to identify egregious errors ($I = 9$). MD6E explicitly used the term “gestalt” to describe her overall sense of a patient’s condition, explaining this holistic assessment ultimately trumps any single datapoint: “Generally speaking, your labs and imaging are supposed to be tools to help you create differential diagnoses. But ultimately, some of it is based on your gestalt, like if you think the patient looks sick. . . Your imaging may be stone-cold normal, but if they look sick, they’re going to get admitted. That’s what happens.”

Perceived need for medical expertise makes attacks seem implausible. The perception of low likelihood for integrity attacks appears to stem from a shared hypothesis among clin-

	≤ 1mo	2mo	6mo	≤ 1yr	> 1yr
Records ($S=106$)	19.8%	11.3%	21.7%	31.1%	16.0%
Imaging ($S=89$)	11.2%	9.0%	12.4%	29.2%	38.2%
Interventions ($S=112$)	15.2%	5.4%	12.5%	26.8%	40.2%
Medication ($S=89$)	9%	9%	7.9%	27%	47.2%
Monitoring ($S=184$)	14.7%	8.2%	10.3%	19%	47.8%

Table 6: Participants’ downtime rate for each technology.

icians: that a successful attack would require a rare combination of sophisticated technical skill and deep medical expertise. When asked to consider a malicious actor, some interviewees dismissed the scenario as implausible, with integrity failures being “not possible” ($I = 4$) or requiring medical expertise ($I = 2$). As APP7E noted, a successful imaging attack would require “someone who knew what they were looking at for that to be problematic. . . It would have to be a very talented and nefarious person.” This sentiment was echoed by MD9I, who said that with “today’s imaging is harder to do and would be trickier for someone to manipulate.” This suggests a broader belief that clinical data is not easily weaponized without significant contextual knowledge, a barrier that clinicians perceived as being prohibitively high for most malicious actors and suggests a need to prioritize cases where an attacker could produce clinically consistent, but harmful modifications, as these might be missed by clinicians.

5.6 Availability Failures: A Familiar Reality of Downtime (RQ1)

Availability failures—where a system or device is unusable—represented a middle ground in clinicians’ risk perceptions. They were seen as both more harmful, all else equal, than confidentiality failures ($OR : 23.7, p < .001$ in Table 7) and more likely than integrity failures (as shown by the non-overlapping confidence intervals). Unlike more abstract failures, clinicians’ views on availability are firmly grounded in their frequent, direct experiences with non-malicious system downtime, as most have experienced it in the past year across all major technology categories (Table 6).

EHR unavailability causes critical care delays. The primary harm from EHR unavailability, identified consistently across interviews ($I = 8$) and the survey ($S = 93$), was the potential for catastrophic delays in care. As MD10I explained, the inability to access a patient’s record is terrifying in time-sensitive situations: “I don’t know [patients] well enough when the system’s down. . . They call and have an infection. I have to choose an antibiotic. I don’t know their allergies. I don’t know their kidney function. . . I really need that written record to treat them.” Interviewees often situated their unavailability concerns in their own experiences. MD10I reflected on a 24-hour EHR outage and concluded the downtime likely “resulted in some medication errors and probably harmed some

patients.”

Imaging systems’ unavailability is highly disruptive. An availability failure of imaging systems was perceived (all else equal) to be significantly more harmful relative to the baseline EHR ($OR : 4.2, p < .001$), underscoring the critical role of timely diagnostics (Table 4). The primary harm mentioned from this unavailability is a significant diagnosis delay. When viewing software is down, the primary workaround is for a clinician to physically access the scanner’s console to review images. As MD9I explained, “we can go to consoles and review data manually at the scanner. . . to exclude bleeds and other major things by physically walking around [to different imaging systems].” APP7E described resorting to re-scanning a patient when the viewing software was down, and they were unable to access previously captured images, a workaround that unnecessarily introduces additional radiation exposure.

Interventional device downtime can be fatal. Given their direct role in treatment, interventional device availability failures were viewed as very severe—66.3% ($S_S = 75$) reported they could cause critical or catastrophic harm, and this was statistically significantly greater than for EHR systems ($OR : 8.2, p < 0.001$), all else equal. MD10I gave the example that if dialysis centers went down “for more than a day or two, I have patients who would die basically, or have very poor outcomes.”

Medication cabinet availability failures seen as operational headaches, but could lead to fatal outcomes. Medication cabinet availability failures were seen as common, causing frequent operational headaches resulting in significant care delays ($S = 49$). MD12I noted, “The Pyxis going out can be a real pain. . . Because then, hey, I need this drug. You have to go find someone to get it for you.” RN4I described a system-wide failure of the fingerprint scanners used for access, which “really delayed patient care.”

MD6E discussed the widely publicized case of a nurse at Vanderbilt University Medical Center, administering the wrong medication to a patient in 2017, leading to a patient’s death [43, 60]. There were persistent technical issues with the medication cabinets (i.e., an availability failure), which led the unit to formalize overriding the system as a workaround. This workaround created conditions for a catastrophic integrity failure as nursing staff had to select from an ambiguous list of medications, leading one nurse to ultimately administer a fatal dose of a powerful paralytic instead of the intended sedative. This case demonstrates that in clinical settings, failures are rarely confined to a single security principle; instead, usability and availability failures enabled an integrity failure.

Monitoring downtime increases workload and risk. The unavailability of patient monitoring devices was seen as increasing nursing workload and potentially delaying the recognition of a patient’s decline, leading to significant harm—52.7% ($S_S = 97$) reported they could cause critical or catastrophic harm, and this was statistically significantly greater

than for EHR systems ($OR : 8.2, p < 0.001$), all else equal. RN2I described how the inability to connect a patient’s vital sign telemetry to the central monitor, requiring them to go to each patient room to check what might be triggering an alarm, was “so frustrating because things are beeping, and when you go into the patient’s room, they’re fine,” creating a risk of alarm fatigue, a known problem facing clinicians [3, 41].

5.7 Managing Unavailability: A Playbook of Workarounds (RQ2)

Paper charting is the primary EHR downtime procedure. The frequent experience with system downtime, mentioned nearly ten times more frequently in our interviews than integrity failures ($I_S = 63$ vs. $I_S = 7$), has led clinicians to develop several institutionalized workarounds. The most significant is reverting to paper charting during EHR downtime. This practice involves documenting all patient care—from physician orders to medication administration—on physical paper records instead of within the electronic system. Even if clinicians have experienced downtime, they find it highly disruptive. APP7E recalled their most recent downtime shift “was awful. But no one’s familiar with it. So I was giving a verbal order over the phone to one of my nurses, and they’re like, ‘No, no, you have to write this down.’ And so I wound up just writing an order down on a piece of paper.”

Failures necessitate procedural workaround. Failures in hospital infrastructure compel clinicians to adopt workarounds ranging from localized, device-specific solutions to system-wide procedural adaptations. Communication system failures, in particular, frequently necessitated procedural workarounds. APP3I described a multi-day outage of the hospital’s phone system due to a cyberattack, which disabled the emergency paging system and required runners to pass messages. For other technologies, workarounds were more localized, such as MD12I going to a cabinet on a different unit for a “box of emergency drugs” when a medication cabinet was unavailable. Similarly, poor usability of official tools drove clinicians to use alternative communication channels to compensate for system deficiencies ($I = 5$). MD10I discussed this frustration regarding texting images for consults, saying, “You need to take care of somebody, and you’re not willing to wait for technology to not work that well.”

Informal knowledge is critical for resilience. Another key theme was that many of these resilience strategies are not formally taught, but passed down informally between colleagues. This was especially true for technologies that frequently fail. RN8I noted, “it’s like a thing I have to teach to the [new] nurses. I’m like, ‘Oh, this frequently doesn’t work. Just do it this way.’” This reliance on informal knowledge demonstrates clinicians’ adaptability but also highlights a systemic fragility where care continuity depends on passed-down wisdom rather than robust, formal procedures.

5.8 Cross-Cutting Themes in Risk Perception

Beyond the specific CIA failures, we found three underlying themes shaping clinicians' perceptions and responses to cybersecurity risks.

Prioritizing patient care drives security trade-offs. When faced with a conflict between security protocols and the immediate needs of patients, clinicians will prioritize patient care. This often leads to conscious security trade-offs. MD10I explained because the official telehealth platform “really doesn't work well,” he and his colleagues will resort to unsanctioned tools: “we end up just facetimeing. I probably shouldn't really say that on the record. But we end up just using our personal iPhones.” He justified this by explaining his priorities: “my priority is taking care of patients. And sometimes I feel like that desire... outweighs... the importance I place on security.” Similarly, many clinicians mentioned a need to access systems remotely to provide prompt care ($I = 9$). These clinicians' hospitals require them to use VPNs to access patient records, which protect against man-in-the-middle attacks. Participants described how this remote access capability introduces new potential threats from opportunistic shoulder surfers. APP7E recounted, “I've been at conferences before and I've seen people... open up their personal laptops and are charting or following up on a patient thing... everyone can see you.”

Lack of communication between IT and clinicians. As we mentioned in Section 4, interviewees were uncertain about prior security incidents occurring ($I = 5$), and there appears to be a lack of details about the causes of incidents. APP3I described a situation where “I think [the hospital] recently had their phones hacked... so because the phone lines were down, we couldn't page even for emergencies.” APP3I explained how there was no follow-up about what occurred. This lack of communication extended to more routine maintenance. RN8I talked about how they try to keep at least one infusion pump per bed on the floor, but the pumps are taken by IT/biomedical engineering for patching without consulting the nursing staff, finding “ourselves hiding pumps in cabinets because we know these people come around, [and take] the pumps.”

Training's focus on compliance shapes risk perception. The heightened focus on confidentiality failures' likelihood appears to be strongly shaped by hospital-mandated security training. Almost all interviewees ($I = 11$) described their annual training as focused on HIPAA compliance, preventing phishing, and password hygiene. As RN8I summarized her training's core message, it is “mostly privacy... they don't want them getting access to our EHR... they're like, 'Don't put out your password because then they'll get access to all your stuff.'” This constant reinforcement of privacy rules and personal culpability likely frames clinicians' risk perception, making them highly attuned to common privacy violations while focusing less on less frequent but more dangerous integrity and availability failures.

6 Discussion

In this section, we discuss future work necessary to understand the strengths and weaknesses of clinicians' existing natural defense mechanisms—their clinical gestalt and manual workarounds. We also provide recommendations to build security controls steeped in the realities of clinical work that augment these defenses without introducing unnecessary burdens on patient care—as is the case for current security controls.

Designing for clinical reality. The tension between security and patient care is most acute in the context of clinician mobility and emergencies. The current paradigm often forces a false choice: follow security protocol or deliver immediate care. Our findings regarding the friction caused by static authentication mechanisms align with Stobert et al.'s observations in emergency departments, where the high-tempo environment necessitates fluid interaction with technology [58]. Similarly, the prevalence of unsanctioned workarounds we observed echoes Wani et al.'s findings on BYOD usage, where clinicians bypass hospital-issued devices in favor of personal tools that offer better usability and efficiency [69]. Unsurprisingly, clinicians choose care. This is a design failure, not a user failure; one that is particularly critical as the US is considering requiring 2FA in a planned HIPAA update [23]. While there are known attacks against badge-based access schemes [34, 72], clinicians significantly preferred them over authentication mechanisms, such as 2FA and passwords, as they were the only authentication mechanism that fit clinicians' highly mobile work (Section 5.3). Further work is necessary to develop 2FA schemes that support mobility. Alternatively, we believe it is worth considering whether the threat models utilized in prior badge-based schemes still hold in the more controlled clinical setting, or if some addition of sweeps for skimming [7, 52, 54] and regular badge code recycling might be sufficient to mitigate these attacks.

Additionally, clinical environments require security that is fluid and emergency-aware (Section 5.3). One potential solution is to create the digital equivalent of the “crash cart” clinicians use in medical emergencies, which provides quick access to medications that would otherwise require additional authorization. Systems could be designed with an explicit, auditable “break-glass” configuration. In an emergency, non-essential security controls could be temporarily relaxed, e.g., extending a session timeout on a specific terminal, with all actions subjected to heightened logging and post-event review.

Current healthcare security standards, such as the Health Industry Cybersecurity Practices (HICP) [18] and NIST SP 800-66 [44], heavily emphasize technical controls like 2FA and strict access management. While effective in enterprise environments, our study highlights that these standards often lack the flexibility required for clinical workflows. For instance, while HICP recommends 2FA to mitigate credential theft, our participants viewed it as a significant barrier to care during emergencies. This suggests that standards bodies

should evolve to include “break-glass” provisions or context-aware authentication that accounts for the physical security and urgent nature of the inpatient setting.

Understanding clinical gestalt as a security defense. A central theme from our interviews is that clinicians rely on their “clinical gestalt,” a holistic and intuitive assessment of the patient, as their primary defense against data integrity failures (Section 5.4). Their focus is rightly on immediate issues to patient safety, and they are skilled at identifying data clearly inconsistent with a patient’s condition [16, 71]. This gestalt can be a security strength, but more work is necessary to understand its bounds. While security experts often model threats based on technical feasibility and worst-case scenarios [64], our results suggest that clinicians filter these threats through a probability lens based on their clinical knowledge. A key misalignment exists here: experts design for the sophisticated attacker who can spoof data perfectly, while clinicians rely on a defense that assumes attacks will lack medical nuance.

Determining whether a data anomaly is malicious or merely a benign error, which our participants noted is common, is non-trivial. Simply flagging all data inconsistencies implies a high risk of exacerbating alert fatigue, already a significant patient safety hazard [3, 41]. Therefore, we recommend leveraging clinical gestalt to filter obvious errors, thereby reducing the cognitive overhead of real-time integrity checks. To operationalize this, research must interrogate the bounds of this defense. Existing clinical protocols, such as the “Pneumothorax” checklist in the Stanford Emergency Manual [28], already encode methods for verifying data consistency (e.g., validating machine readings against physical breath sounds). These protocols offer a roadmap for designing security tools that support, rather than override, clinician intuition.

Bridging the communication gap. A recurring theme is the breakdown in communication between IT/security teams and clinicians. Participants noted system updates are often poorly communicated (Section 5.8), and information about prior security incidents is rarely shared (Sections 4, 5.8). This lack of transparency leaves clinicians unprepared and fosters a reliance on informal, ad-hoc workarounds. To build a truly resilient system, we must foster a culture of open communication between clinicians and security staff. This reflects broader stakeholder challenges identified by Kustosch et al., who found that the “patching” process for medical devices is often opaque to the clinical staff relying on them, creating friction when availability is impacted [38]. Furthermore, as Gutfleisch et al. note, there is often a separation of concerns where clinical staff view security as solely an IT responsibility, disengaging from the shared ownership required for a resilient safety culture [29].

One option is to adopt a common medical community tradition: the Morbidity and Mortality (M&M) conference [5, 10]. M&Ms are blameless post-mortem incident discussions to identify root causes and determine process changes. Clini-

cians conduct these post-mortems anytime there is unexpected patient harm or a severe risk to patient safety. Cybersecurity should be incorporated into these discussions by having IT/Security teams join or prompt M&Ms when appropriate to discuss technology and security failures, with an emphasis on specific impacts to patient care. This would allow clinicians to understand how security failures manifest in their environment, contribute their expertise to developing robust resilience strategies, and inform clinical simulations [16, 71].

Acknowledgements

We thank Angela Yan for all her help with this paper. We want to thank Joanne Wozniak at Beth Israel Lahey and the teams at Emory Healthcare, Children’s Hospital of Atlanta, and Vitruvian Health for helping with recruitment. We thank the interviewees and survey respondents for their participation. We thank the anonymous reviewers who provided helpful comments on this paper’s drafts. This project was supported by a MedCrypt gift, ARPA-H contract (140D042590046), and US Air Force contract (FA864924P0557).

Ethical Considerations

This study was reviewed and approved as exempt research by the primary author’s Institutional Review Board (IRB). The entire study was conducted in a manner consistent with the standards defined in the Menlo Report.

Our analysis identified several stakeholders impacted by this research and its publication. These include the clinician participants—the 12 interviewees and 303 survey respondents who provided data—the research team, and indirectly, patients, whose safety this work aims to improve. Other stakeholders include healthcare organizations, the broader security and healthcare communities that will utilize these findings, and society at large, which benefits from a more resilient healthcare system.

We obtained informed consent from all participants before data collection began. For the online survey, participants were required to review and agree to the consent form before proceeding. For participants in the UK, this process included all necessary GDPR disclosures. For the interviews, after a participant had reviewed the consent form, we reaffirmed key information at the start of the session. This included reminding them that we would remove from the transcript any mention of their name, their institution, or any other potentially identifying information. We asked participants to reaffirm their consent before we began recording. We also made clear that participation was voluntary, that they could skip any question or withdraw at any time without penalty, and that they would still receive full compensation. All collected data was anonymized to protect participant privacy and confidentiality.

We provided compensation for participants' time and expertise. Interview participants were given a \$40 Amazon gift card, while survey participants were compensated with \$15 (\$45/hour) on Prolific if they completed the survey. Participants who did not complete the survey or were removed for failing the BLS check were compensated \$0.50.

We also note that we collected copies of the interview participants' CVs/resumes or public profiles (such as LinkedIn or a hospital profile page) to validate their credentials. We were not able to collect similar information with participants from Prolific, as this would violate their Terms of Service and deanonymize our participants. We instead elected to screen participants with both knowledge checks (BLS) and verifying that their responses were in line with their stated roles and made clinical sense in consultation with our clinical team members and external clinicians we spoke to. We broadened our survey recruitment to include the US, UK, and Canada to ensure our findings were not limited to a single healthcare system. Our research serves a clear public interest by aiming to strengthen the security of critical healthcare infrastructure, a direct benefit to patients and society at large.

The only potential harm we envision from publishing this work is that a malicious actor could exploit our findings on clinician mental models and common workarounds to design more subtle attacks. However, we believe this risk is low compared to the significant benefit of improving healthcare security. These operational realities must be understood by security designers and policymakers to be fixed. Our work offers actionable recommendations for developing more effective, clinician-centered security that can directly improve patient safety. Therefore, we believe that the benefits of this work substantially outweigh the potential limited harms to the clinical community and the patients they serve.

Open Science

In the spirit of transparency and reproducibility, we provide the following artifacts in our supplemental materials [1]: (1) interview script and screening survey, (2) survey questions, including screening questions, (3) the codebooks associated with both our interviews and survey, and (4) anonymized quantitative data from our survey with analysis scripts used to generate the results in our paper. Our appendix contains additional information related to the survey, including a visual representation of the responses to the perceived harm and likelihood across the scenarios for the technology groups (Fig 5, this can be recreated with the data provided in the supplemental materials [1]), and we include a similar visual representation for the various controls (Fig 6).

To protect participants' privacy and confidentiality, we do not release raw interview recordings or full transcripts. This decision is in accordance with our IRB-approved protocol and ethical research practices for studies involving sensitive discussions with a specialized population. Per our protocol, all

raw audio recordings were permanently deleted immediately after transcript validation for accuracy. We instead present our qualitative findings through thematic analysis supported by illustrative anonymized quotes.

References

- [1] Supplemental materials. https://osf.io/ahjp3/overview?view_only=8442f58f00ca4dfdbf2ed2dfa90baba0.
- [2] American Heart Association. Highlights of the 2020 AHA guidelines for CPR and ECC, 2020.
- [3] Jessica S. Ancker, Alison Edwards, Sarah Nosal, Diane Hauser, Elizabeth Mauer, Rainu Kaushal, et al. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. *BMC Medical Informatics and Decision Making*, 17(1):36, 2017.
- [4] J.D. Aurini, M. Heath, and S. Howells. *The How To of Qualitative Research*. SAGE Publications, 2021.
- [5] Brendin R. Beaulieu-Jones, Spencer Wilson, Daniel S. Howard, Gordana Rasic, Ben Rembetski, Erica A. Brotschi, and Luise I. Pernar. Defining a high-quality and effective morbidity and mortality conference: A systematic review. *JAMA Surgery*, 158(12):1336–1343, 2023.
- [6] Tamara Bonaci, Jeffrey Herron, Tariq Yusuf, Junjie Yan, Tadayoshi Kohno, and Howard Jay Chizeck. To make a robot secure: An experimental analysis of cyber security threats against teleoperated surgical robots. *arXiv preprint arXiv:1504.04339*, 2015.
- [7] Steve Bono, Matthew Green, Adam Stubblefield, Ari Juels, Aviel D. Rubin, and Michael Szydlo. Security analysis of a cryptographically-enabled RFID device. In *USENIX Security Symposium*, volume 31, pages 1–16, 2005.
- [8] Dawn Branley-Bell, Lynne Coventry, Elizabeth Sillence, Sabina Magalini, Pasquale Mari, Aimilia Magkanaraki, and Kalliopi Anastasopoulou. Your hospital needs you: Eliciting positive cybersecurity behaviours from healthcare staff. *Annals of Disaster Risk Sciences*, 3(1), 2020.
- [9] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, 2006.
- [10] J.S. Carroll and Amy C. Edmondson. Leading organisational learning in health care. *BMJ Quality & Safety*, 11(1):51–56, 2002.
- [11] Center for Devices and Radiological Health, FDA. Postmarket management of cybersecurity in medical devices, 2016.
- [12] Sunny Consolvo, Patrick Gage Kelley, Tara Matthews, Kurt Thomas, Lee Dunn, and Elie Bursztein. “why wouldn’t someone think of democracy as a target?”: Security practices & challenges of people involved with US political campaigns. In *30th USENIX Security Symposium*, pages 1181–1198, 2021.
- [13] Juliet Corbin and Anselm Strauss. *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. Sage Publications, 4th edition, 2014.

- [14] Lynne Coventry and Dawn Branley. Cybersecurity in health-care: A narrative review of trends, threats and ways forward. *Maturitas*, 113:48–52, 2018.
- [15] Christian Dameff, Jennifer Farah, James Killeen, and Theodore Chan. Cyber disaster medicine: A new frontier for emergency medicine. *Annals of Emergency Medicine*, 75(5):642–647, 2020.
- [16] Christian J. Dameff, Jordan A. Selzer, Jonathan Fisher, James P. Killeen, and Jeffrey L. Tully. Clinical cybersecurity training through novel high-fidelity simulations. *The Journal of Emergency Medicine*, 56(2):233–238, 2019.
- [17] Tamara Denning, Alan Borning, Batya Friedman, Brian T. Gill, Tadayoshi Kohno, and William H. Maisel. 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*, pages 917–926, 2010.
- [18] Department of Health and Human Services. Health industry cybersecurity practices: Managing threats and protecting patients, 2023.
- [19] Department of Health and Human Services. HPH cybersecurity performance goals, January 2025.
- [20] Diane Irvine Doran, Souraya Sidani, Margaret Keatings, and Doris Doidge. An empirical test of the nursing role effectiveness model. *Journal of Advanced Nursing*, 38(1):29–39, 2002.
- [21] Josiah Dykstra, Rohan Mathur, and Alicia Spoor. Cybersecurity in medical private practice: Results of a survey in audiology. In *2020 IEEE 6th International Conference on Collaboration and Internet Computing (CIC)*, pages 169–176. IEEE, 2020.
- [22] Cori Faklaris, Laura A. Dabbish, and Jason I. Hong. A self-report measure of end-user security attitudes (SA-6). In *15th Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 61–77, 2019.
- [23] Federal Register. HIPAA security rule to strengthen the cybersecurity of electronic protected health information, January 2025.
- [24] Kevin Fu and James Blum. Controlling for cybersecurity risks of medical device software. *Communications of the ACM*, 56(10):35–37, 2013.
- [25] Andrea Gallardo, Robert Erbes, Katya Le Blanc, Lujo Bauer, and Lorrie Faith Cranor. Interdisciplinary approaches to cyber vulnerability impact assessment for energy critical infrastructure. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–24, 2024.
- [26] Saira Ghafur, Emilia Grass, Nick R. Jennings, and Ara Darzi. The challenges of cybersecurity in health care: The UK national health service as a case study. *The Lancet Digital Health*, 1(1):e10–e12, 2019.
- [27] Daniele Giansanti and Lisa Monoscalco. The cyber-risk in cardiology: Towards an investigation on the self-perception among the cardiologists. *mHealth*, 7:28, 2021.
- [28] Sara N. Goldhaber-Fiebert, Naola S. Austin, Ellile Sultan, Barbara K. Burian, Amanda Burden, Steven K. Howard, David M. Gaba, and T. Kyle Harrison. *Emergency Manual: Cognitive Aids for Perioperative Crises*. Stanford Anesthesia Cognitive Aid Program, 2021. Version 4.
- [29] Marco Gutfleisch, Markus Schöps, Jonas Hielscher, Mary Cheney, Sibel Sayin, Nathalie Schuhmacher, Ali Mohamad, and M. Angela Sasse. Caring about IoT-security—an interview study in the healthcare sector. In *Proceedings of the 2022 European Symposium on Usable Security*, pages 202–215, 2022.
- [30] Andrew F. Hayes and Klaus Krippendorff. Answering the call for a standard reliability measure for coding data. *Communication Methods and Measures*, 1(1):77–89, 2007.
- [31] Healthcare Sector Coordinating Council. Health industry cybersecurity practices: Managing threats and protecting patients, 2019.
- [32] Grant Ho, Ariana Mirian, Elisa Luo, Khang Tong, Euyhyun Lee, Lin Liu, Christopher A. Longhurst, Christian Dameff, Stefan Savage, and Geoffrey M. Voelker. Understanding the efficacy of phishing training in practice. In *2025 IEEE Symposium on Security and Privacy (SP)*, page 76. IEEE Computer Society, 2024.
- [33] Mi Ok Kim, Enrico Coiera, and Farah Magrabi. Problems with health information technology and their effects on care delivery and patient outcomes: A systematic review. *Journal of the American Medical Informatics Association*, 24(2):246–250, 2017.
- [34] Thomas Korak and Thomas Plos. Applying remote side-channel analysis attacks on a security-enabled NFC tag. In *Cryptographers’ Track at the RSA Conference*, pages 207–222. Springer, 2013.
- [35] Daniel B. Kramer, Jennifer R. Amos, Julian M. Goldman, and Kevin Fu. Threats to patient safety from cybersecurity flaws—a new never event. *JAMA*, July 2025.
- [36] Klaus Krippendorff. *Content Analysis: An Introduction to Its Methodology*. Sage Publications, 2nd edition, 2004.
- [37] Clemens Scott Kruse, Benjamin Frederick, Taylor Jacobson, and D. Kyle Monticone. Cybersecurity in healthcare: A systematic review of modern threats and trends. *Technology and Health Care*, 25(1):1–10, 2017.
- [38] Lorenz Kustosich, Carlos Gañán, Michel van Eeten, and Simon Parkin. Patching up: Stakeholder experiences of security updates for connected medical devices. In *34th USENIX Security Symposium*, pages 2265–2281, 2025.
- [39] Heather Landi. Scripps health cyberattack cost the system \$113m in lost revenue, remediation costs. *Fierce Healthcare*, 2021.
- [40] Ethan Larsen, Allan Fong, Christian Wernz, and Raj M. Ratwani. Implications of electronic health record downtime: An analysis of patient safety event reports. *Journal of the American Medical Informatics Association*, 25(2):187–191, 2018.
- [41] Eva K. Lee, Tsung-Lin Wu, Tal Senior, and James Jose. Medical alert management: A real-time adaptive decision support tool to reduce alert fatigue. In *AMIA Annual Symposium Proceedings*, volume 2014, page 845, 2014.

- [42] Peter K. Lindenauer, Steven Z. Pantilat, Patricia P. Katz, and Robert M. Wachter. Hospitalists and the practice of inpatient medicine: Results of a survey of the national association of inpatient physicians. *Annals of Internal Medicine*, 130(4_Part_2):343–349, 1999.
- [43] Connor Lusk, Elise DeForest, Gabriel Segarra, David M. Neyens, James H. Abernathy, and Ken Catchpole. Reconsidering the application of systems thinking in healthcare: The RaDonda Vaught case. *British Journal of Anaesthesia*, 129(3):e61–e62, 2022.
- [44] Jeffrey Marron. Implementing the health insurance portability and accountability act (HIPAA) security rule: A cybersecurity resource guide. Technical Report NIST SP 800-66r2, National Institute of Standards and Technology, February 2024.
- [45] Nora McDonald, Alan Luo, Phoebe Moh, Michelle L. Mazurek, and Nazanin Andalibi. Threat modeling healthcare privacy in the united states. *ACM Transactions on Computer-Human Interaction*, 32(1):1–37, 2025.
- [46] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–23, 2019.
- [47] Claire C. McGlave, Hannah Neprash, and Sayeh Nikpay. Hacked to pieces? the effects of ransomware attacks on hospitals and patients. *The Effects of Ransomware Attacks on Hospitals and Patients*, October 2023.
- [48] Susan E. McGregor, Polina Charters, Tobin Holliday, and Franziska Roesner. Investigating the computer security practices and needs of journalists. In *24th USENIX Security Symposium*, pages 399–414, 2015.
- [49] Daniel R. Murphy, Tyler Satterly, Traber D. Giardina, Dean F. Sittig, and Hardeep Singh. Practicing clinicians’ recommendations to reduce burden from the electronic health record inbox: A mixed-methods study. *Journal of General Internal Medicine*, 34:1825–1832, 2019.
- [50] Hannah T. Neprash, Claire C. McGlave, Dori A. Cross, Beth A. Virnig, Michael A. Puskarich, Jared D. Huling, Alan Z. Rozenstein, and Sayeh S. Nikpay. Trends in ransomware attacks on US hospitals, clinics, and other health care delivery organizations, 2016-2021. In *JAMA Health Forum*, volume 3, page e224873. American Medical Association, 2022.
- [51] Office of U.S. Senator Mark R. Warner. Cybersecurity is patient safety: Policy options in the health care sector, 2022.
- [52] Jae Hoon Park, HoonJae Lee, and ManKi Ahn. Side-channel attacks against ARIA on active RFID device. In *2007 International Conference on Convergence Information Technology (ICCIT 2007)*, pages 2163–2168, 2007.
- [53] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. Beyond the turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70:153–163, 2017.
- [54] Pedro Peris-Lopez, Julio Cesar Hernandez-Castro, Juan M. Estevez-Tapiador, and Arturo Ribagorda. Attacking RFID systems. *Security in RFID and Sensor Networks*, 29, 2016.
- [55] R. H. B. Christensen. *ordinal - Regression Models for Ordinal Data*, 2023. R package version 2023.12-4.
- [56] Kenneth J. Rothman, Sander Greenland, Timothy L. Lash, Charles Poole, and Zirui Geng. *Modern Epidemiology*. Wolters Kluwer, 4th edition, 2021.
- [57] Michael Rushanan, Aviel D. Rubin, Denis Foo Kune, and Colleen M. Swanson. SoK: Security and privacy in implantable medical devices and body area networks. In *2014 IEEE Symposium on Security and Privacy*, pages 524–539. IEEE, 2014.
- [58] Elizabeth Stobert, David Barrera, Valérie Homier, and Daniel Kollek. Understanding cybersecurity practices in emergency departments. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–8, 2020.
- [59] Isabel Straw, Irina Brass, Andrew Mkwashi, Inika Charles, Amelie Soares, and Caroline Steer. Insights from a clinically orientated workshop on health care cybersecurity and medical technology: Observational study and thematic analysis. *Journal of Medical Internet Research*, 26:e50505, 2024.
- [60] Susan Swart. Empowering future nurses: RaDonda Vaught’s journey and insights at SNPAD, January 2025.
- [61] Jenny Tang, Lujo Bauer, and Nicolas Christin. Misuse, misreporting, misinterpretation of statistical methods in usable privacy and security papers. In *21st Symposium on Usable Privacy and Security (SOUPS 2025)*, pages 475–493, 2025.
- [62] Faiza Tazi, Josiah Dykstra, Prashanth Rajivan, and Sanchari Das. “we have no security concerns”: Understanding the privacy-security nexus in telehealth for audiologists and speech-language pathologists. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–20, 2024.
- [63] Lehana Thabane, Lawrence Mbuagbaw, Shiyuan Zhang, Zainab Samaan, Maura Marucci, Chenglin Ye, Masuzyo Thabane, et al. A tutorial on sensitivity analyses in clinical trials: The what, why, when and how. *BMC Medical Research Methodology*, 13(1):1–12, 2013.
- [64] Ronald E. Thompson, Madline McLaughlin, Carson Powers, and Daniel Votipka. “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*, pages 4909–4926, 2024.
- [65] Jeff Tully, Jordan Selzer, James P. Phillips, Patrick O’Connor, and Christian Dameff. Healthcare challenges in the era of cybersecurity. *Health Security*, 18(3):228–231, 2020.
- [66] Jeffrey L. Tully, Sumanth Rao, Isabel Straw, Rodney A. Gabriel, Chris Longhurst, Stefan Savage, Geoffrey M. Voelker, and Christian J. Dameff. Patient care technology disruptions associated with the crowdstrike outage. *JAMA Network Open*, 8(7):e2530226, July 2025.
- [67] Gerhard Tutz and Wolfgang Hennevoogl. Random effects in ordinal regression models. *Computational Statistics & Data Analysis*, 22(5):537–557, 1996.
- [68] Liselotte S. van Boven, Renske W.J. Kusters, Derrick Tin, Frits H.M. van Osch, Harald De Cauwer, Lindsay Ketelings, Madhura Rao, Christian Dameff, and Dennis G. Barten. Hacking acute care: A qualitative study on the health care impacts

of ransomware attacks against hospitals. *Annals of Emergency Medicine*, 83(1):46–56, 2024.

- [69] T.A. Wani, Antonette Mendoza, and Kathleen Gray. BYOD use and perception among hospital clinicians—a qualitative study. *Ethics, Medicine and Public Health*, 33:101031, 2025.
- [70] Jess Warren. NHS ransomware attack contributed to patient’s death. *BBC*, June 2025.
- [71] Markus Willing, Christian Dresen, Eva Gerlitz, Maximilian Haering, Matthew Smith, Carmen Binnewies, Tim Guess, Uwe Haverkamp, and Sebastian Schinzel. Behavioral responses to a cyber attack in a hospital environment. *Scientific Reports*, 11(1):19352, 2021.
- [72] Qing Yang and Lin Huang. RFID/NFC security. In *Inside Radio: An Attack and Defense Guide*, pages 71–121. Springer, 2018.

A Survey Data

We provide a visual representation of the perceived harm and likelihood for each scenario (Fig 5) as well as controls (Fig 6). We also provide the results from our non-interaction models for both perceived harm and likelihood (Table 7). The complete data for the scenarios and controls can be found in our supplemental materials [1].

Factor	Value	Harm			Likelihood		
		OR	p-Value	CI	OR	p-Value	CI
Technology	EHR	-	-	-	-	-	-
	Imaging	0.7	0.033	0.5, 0.9	0.9	0.394	0.6, 1.2
	Interventions	1.6	0.016	1.2, 2.1	0.6	0.001	0.4, 0.7
	Med. Cabinets	0.9	0.424	0.6, 1.2	0.3	<0.001	0.2, 0.4
	Monitoring	0.6	<0.001	0.4, 0.7	0.5	<0.001	0.4, 0.6
CIA	Confidentiality	-	-	-	-	-	-
	Integrity	48.9	<0.001	38.0, 63.0	0.2	<0.001	0.2, 0.2
	Availability	23.7	<0.001	18.8, 29.9	0.5	<0.001	0.4, 0.6
Role	RN	-	-	-	-	-	-
	AHP	1.0	0.880	0.7, 1.3	1.2	0.362	0.9, 1.7
	MD/APP	0.8	0.317	0.6, 1.1	0.7	0.085	0.6, 1.0
Geo	United States	-	-	-	-	-	-
	Canada	0.9	0.426	0.6, 1.2	0.9	0.746	0.7, 1.3
	United Kingdom	1.1	0.649	0.8, 1.4	0.9	0.594	0.7, 1.2
Unit	Inpatient	-	-	-	-	-	-
	Outpatient	0.8	0.120	0.6, 1.0	0.8	0.239	0.6, 1.1
	ED/ICU	1.5	0.064	1.0, 2.2	1.6	0.038	1.1, 2.3
SA6	1	-	-	-	-	-	-
	+1	1.1	0.200	1.0, 1.3	0.9	0.215	0.8, 1.0
Security Training	False	-	-	-	-	-	-
	True	0.8	0.106	0.6, 1.0	0.9	0.374	0.7, 1.1
Prior Attack Exp.	False	-	-	-	-	-	-
	True	1.3	0.117	1.0, 1.7	1.3	0.141	1.0, 1.7

Statistically significant values ($p \leq 0.05$) are **bolded**

- Base case (Odds ratio defined as 1)

Table 7: Results from mixed effects regressions for perceived harm and the likelihood of a scenario occurring. OR above one implies more potential harm/likely than the base case.

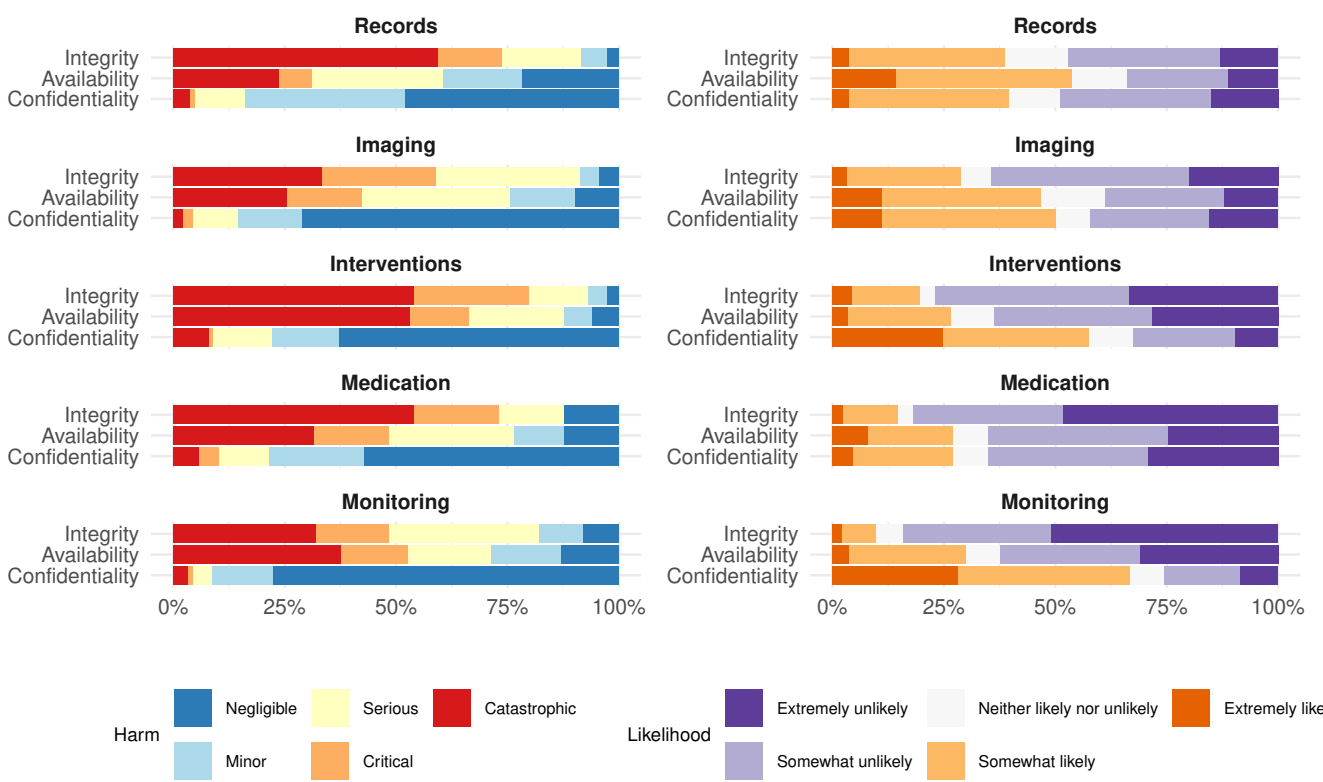


Figure 5: Perceived harm and likelihood for each scenario across the technology groups.

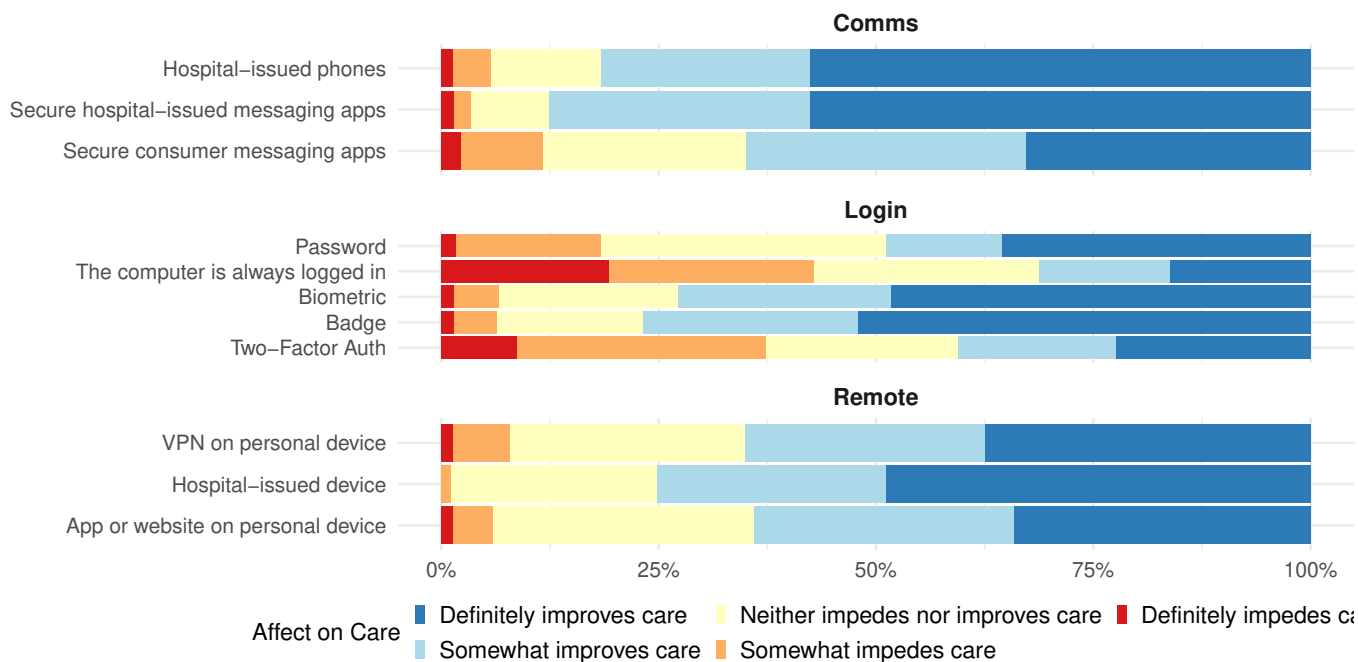


Figure 6: Perceived impact on care from security controls, communications devices, and remote access tools.