Ceremony Meets Usability: An HCI Lens on Authentication

Lorin Schöni
lorin.schoeni@gess.ethz.ch
ETH Zurich
Zurich, Switzerland

Abstract

Passkeys are rapidly moving from niche to mainstream as a phishing-resistant alternative to passwords. Yet, they expose gaps in recovery and cross-device flows, while attackers exploit human fatigue and AI-driven deception to bypass these cryptographic improvements. To address these challenges, we revisit two landmarks in usable security: *Why Johnny Can't Encrypt* (1999) and *Ceremony Design and Analysis* (2007). Our synthesis links *task-level usability* (can users do it?) with *ceremony-level robustness* (is the context safe?), offering a lens for design and evaluation beyond protocol correctness. We conclude by outlining directions for treating authentication flows as familiar HCI patterns, focusing on clarity, consistency, and resilience, so that authentication becomes both usable and hard to misuse.

Keywords

usable security, authentication, passkeys, security ceremonies

1 Introduction

WebAuthn passkeys promise a passwordless, phishing-resistant future, yet their rollout reveals cracks in the experience. Variation in registration, recovery, and cross-device use creates confusion [5, 20], while attackers have shifted tactics to exploit human fatigue rather than encryption. This is evident in the rise of prompt bombing [18, 25], as well as AI-driven phishing and adversarial automation [3, 15]. These attacks target the *ceremony* of authentication, not the protocol.

Usable security research has long recognised that cryptographic correctness is insufficient if users cannot reliably enact the required tasks [10, 14]. Whitten and Tygar's *Why Johnny Can't Encrypt* demonstrated, via an evaluation of PGP 5.0, that apparently "good" interfaces still failed to support novices in performing secure email operations [26]. They argued that security-critical interaction demands a distinct usability standard. Ellison's *Ceremony Design and Analysis* extended this perspective by showing that security protocols are not just technical steps but include people, interfaces, and even physical or social actions [12]. These extra steps, often called "out-of-band" (e.g., confirming a code via SMS or showing an ID), are part of the ceremony and require deliberate design.

Building on these foundations, we revisit *Johnny* and *ceremony* design together to address today's authentication challenges. Our contribution is a synthesis that enables ceremony-aware usability analysis: a view that couples *task-level usability* with *ceremony-level robustness*. We do not prescribe UI guidelines or formal proofs; instead, we offer a conceptual scaffold for researchers and practitioners to derive design patterns, evaluation heuristics, and threat models.

AlpCHI 2026,
Revisiting HCI Research.

2 Two lenses to revisit

Johnny's security usability standard. Johnny established that conventional usability heuristics do not suffice when mistakes have security consequences: the unit of analysis must be the *security task* (e.g., sign-then-encrypt with correct keys), the *mental models* (keys, identities, trust), and the *error surface* where slips lead to a breach [26]. This insight is still relevant for automation. Passkeys shift work from users to authenticators and relying parties, but still leave users with moments of judgement (choice of device, nearby-device approval, account recovery, attestation choices, and consent to risky fallbacks) where misunderstanding is costly [5, 20].

Ceremony design, not protocol fragments. Ellison's ceremony concept insists that human actors, UI channels, and social exchanges are *in band* for analysis [12]. Modern authentication involves actors (user, device, authenticator), channels (prompts, QR codes, recovery flows), and trust assumptions (metadata). Recent work on user-centred ceremony design reinforces this need for structured socio-technical modelling [6, 13].

Combining the lenses: towards ceremony-aware usability. Taken together, these lenses suggest security design must operate on two planes: *task-level comprehensibility* and *ceremony-level robustness* [22]. Johnny highlights learnable paths through security-critical tasks. Ellison shows these tasks sit within ceremonies whose integrity depends on trust and channels. Combined, they offer a conceptual foundation for authentication that anticipates failure modes such as MFA fatigue or recovery exploits [8].

3 Contemporary landscape through the two lenses

We now highlight specific challenges where these lenses reveal why current solutions fall short.

3.1 Passkeys: promise and ceremony gaps

Passkeys provide phishing-resistant authentication by binding public-key credentials to a relying party. While they effectively stop credential stuffing, they introduce new mental model gaps regarding recovery, delegation, and cross-system portability [5, 20]. Recent audits underline this tension: passkeys hold clear security promise, yet implementations reveal uneven patterns across sites and sectors that leave users confused [20, 23].

Seen through *Johnny*, the critical issue is the task model: how do users understand keys that sync across some devices but not others? A *ceremony* lens, in turn, interrogates the context: are the trust checks visible and resilient when a phone is lost or a vendor changes? Updated design guidance seeks to standardise these flows [1, 2], yet insights from industry suggest that without clear ceremony design involving phased rollouts and user involvement, adoption remains fragile [4].

3.2 MFA fatigue as ceremony failure

"Prompt bombing" exploits the approval *ceremony* rather than any cryptographic weakness: repeated, interruptive prompts shift the user's goal from validation to stopping the noise [18, 25]. Effective mitigations (e.g., number matching or contextual prompts) are therefore ceremony-level redesigns. They must bind context to action while respecting the user's compliance budget, ensuring that users validate the login source rather than blindly clicking "Approve" just to clear the screen [8, 18, 25].

3.3 Indicators in an era of AI-assisted deception

The Emperor's New Security Indicators showed that users often overlook passive cues (like HTTPS locks) and that behavioural stakes matter [24]. With AI-assisted phishing generating polished, context-aware threats, the salience and timing of indicators become even more critical. A ceremony-first stance argues that we cannot rely on users noticing subtle changes in the browser's UI, such as the address bar. Instead, decisive cues (such as the requesting device, location, or transaction scope) must be bound directly into the approval step, aligning the security check with the user's immediate goal [16, 24].

4 Actionable Directions

Rather than prescribing a rigid UI guideline, we outline three actionable directions that connect our synthesis to current challenges and HCI implications:

1. Map ceremonies to interaction patterns. Authentication flows should be treated as structured interaction sequences, similar to onboarding or consent dialogs in HCI. Mapping ceremonies to these patterns can reveal where trust cues and decision points occur, enabling designers to apply familiar usability heuristics (such as visibility or feedback) under adversarial conditions.

2. Prioritise security-critical choice points. Passkey adoption and MFA fatigue show that user errors cluster at specific junctures: recovery, cross-device migration, and approval prompts. Identifying these choice points and designing for clarity and context-binding (e.g., number matching, verifiable RP identity) can reduce risk without overwhelming users.

3. Harmonise flows and rehearse recovery. Contemporary audits reveal inconsistent passkey UX across sites. Harmonising flows through shared design patterns and making recovery legible and rehearsable, transforming it from a panic-driven exception into a familiar routine, supports both usability and ceremony robustness. This echoes HCI's emphasis on learnability and pattern libraries, even in adversarial contexts.

5 Implications for HCI

While our focus has been on usable security, the synthesis of task-level usability and ceremony-level robustness also resonates with broader HCI concerns:

Ceremony as an HCI interaction unit. Treating authentication as a *ceremony* aligns with long-standing HCI theories that broaden the unit of analysis beyond single screens and single users. In *distributed cognition*, cognition is organised across people, artefacts, and environments [17]. Similarly, *activity theory* frames purposeful, mediated activity through tools, rules, and roles [19]. Reading authentication through these lenses helps HCI model multi-device, cross-channel flows and the orchestration of cues, roles, and artefacts to make ceremony structure a primary interaction concern, rather than isolated UI states.

Designing and evaluating under adversarial conditions. HCI traditionally optimises for effectiveness and satisfaction under benign conditions; ceremonies foreground interaction under deception, interruption, and uncertainty. Classic design principles remain vital for making "the right action at the right time" legible in security-critical moments [21]. Human-AI Interaction (HAII) guidelines add concrete patterns for expectation-setting, error handling, and graceful failure [7]. Together they suggest evaluation heuristics beyond time-on-task: robustness under misleading cues, recovery without unsafe shortcuts, and user comprehension of risk at decision points.

HAII synergy: trust calibration at the point of action. AI-assisted deception raises the stakes for calibrated trust. HAII research shows explanations and confidence signals must be *situated* to support decisions within a clear context, rather than generic model expositions [9, 11]. Ceremony-aware designs move decisive cues into the approval step (binding verifiable context to the user's action) and borrow from human-centred explainable AI research to present uncertainty and rationale in ways that prevent over- or under-reliance. This bridges HAII with socio-technical structuring of critical interactions and offers HCI testable patterns for explanation placement, content, and timing.

6 Scope and Conclusion

Our stance is intentionally conceptual. We do not prescribe a single UI nor a fixed set of metrics. Instead, we surface *where* design effort belongs (ceremony structure, context-binding, recovery) and *how* to reason about human-in-the-loop authentication without collapsing into either protocol formalism or purely aesthetic UI concerns [12, 26].

Our synthesis of Johnny's usability lens and Ellison's ceremony perspective highlights that authentication is not only a protocol problem but an interaction design challenge. Passkeys, MFA fatigue, and AI-driven phishing expose gaps where human decisions matter. Treating authentication as a structured interaction, similar to onboarding or consent flows, helps designers apply HCI principles such as clarity, feedback, and consistency under adversarial conditions. A ceremony-aware agenda that prioritises context-binding, principled friction, and rehearsable recovery can make phishing-resistant, passkey-ready authentication both usable and resilient.

References

[1] 2024. Design Guidelines for Passkeys - Windows Apps. https://learn.microsoft.com/en-us/windows/apps/develop/security/design
[2] 2024. FIDO Alliance Design Guidelines for Passkeys. https://fidoalliance.org/new-design-guidelines-optimizing-user-sign-in-experience-with-passkeys/
[3] 2025. Cybersecurity Implications of AI: Pulse Report. https://ismg.io/resource/cybersecurity-implications-of-ai-pulse-report-2025/
[4] 2025. Passkeys for Enterprise Report from FIDO. https://www.biometricupdate.com/202503/passkeys-for-enterprise-report-from-fido-says-adoption-is-growing
[5] 2025. WebAuthn and Passkeys Primer. https://www.webauthn.me/passkeys

[6] A. Alatawi et al. 2023. SoK: An Analysis of End-to-End Encryption and Authentication Ceremonies. In *ACM CCS*. https://doi.org/10.1145/3558482.3581773

[7] Saleema Amershi and et al. 2019. Guidelines for Human-AI Interaction. In *Proceedings of CHI*. 1–13. doi:10.1145/3290605.3300233

[8] A. Beautement, M. A. Sasse, and M. Wonham. 2008. The Compliance Budget: Managing Security Behaviour in Organisations. In *NSPW*. https://doi.org/10.1145/1595676.1595684

[9] Andrea Bussone and et al. 2015. The Role of Explanations on Trust and Reliance in Clinical Decision Support Systems. *International Journal of Human-Computer Studies* 72, 12 (2015), 941–953. doi:10.1016/j.ijhcs.2014.09.002

[10] F. Di Nocera et al. 2023. Usable Security: A Systematic Literature Review. *Information* 14, 12 (2023), 641. https://doi.org/10.3390/info14120641

[11] Upol Ehsan and Mark O. Riedl. 2020. Human-Centered Explainable AI: Towards a Research Agenda. In *Proceedings of CHI Workshop on Human-Centered XAI*. https://doi.org/10.1007/978-3-030-60117-1_33

[12] Carl Ellison. 2007. Ceremony Design and Analysis. IACR ePrint. https://eprint.iacr.org/2007/399

[13] M. Fassl et al. 2021. Exploring User-Centered Security Design for Usable Authentication Ceremonies. In *CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3411764.3445164

[14] Steven Furnell. 2024. Usable Cybersecurity: a Contradiction in Terms? *Interacting with Computers* 36, 1 (2024), 3–20. https://doi.org/10.1093/iwc/iwad035

[15] M. Hashmi et al. 2024. Securing Tomorrow: A Comprehensive Survey on AI and Information Security. *Journal of Ambient Intelligence and Humanized Computing* (2024). https://doi.org/10.1007/s43681-024-00529-z

[16] F. Heiding, B. Schneier, A. Vishwanath, and J. Bernstein. 2023. Devising and Detecting Phishing: Large Language Models vs. Smaller Human Models. Black Hat USA whitepaper. https://doi.org/10.48550/arXiv.2308.12287

[17] James Hollan, Edwin Hutchins, and David Kirsh. 2000. Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research. *ACM Transactions on Computer-Human Interaction* 7, 2 (2000), 174–196. doi:10.1145/353485.353487

[18] D. James. 2023. Beating MFA Fatigue: Why Hackers Have Resorted to Prompt Bombing. CPO Magazine. https://www.cpomagazine.com/cyber-security/beating-mfa-fatigue-why-hackers-have-resorted-to-prompt-bombing/

[19] Victor Kaptelinin and Bonnie Nardi. 2012. *Activity Theory in HCI: Fundamentals and Reflections*. Morgan & Claypool. https://dl.acm.org/doi/10.5555/2361921

[20] A. Matzen et al. 2025. Challenges and Potential Improvements for Passkey Adoption—A Literature Review with a User-Centric Perspective. *Applied Sciences* 15, 8 (2025), 4414. https://doi.org/10.3390/app15084414

[21] Donald A. Norman. 2013. *The Design of Everyday Things (Revised and Expanded Edition)*. Basic Books.

[22] A. Ortloff and G. Martius. 2025. Meta-Science in Usable Security and Privacy and HCI. https://doi.org/10.1145/3706598.3714022

[23] B. Ramat, D. Kartchner, and K. Seamons. 2025. A Systematic Analysis of the Passkey User Experience (poster abstract). USENIX SOUPS. https://www.usenix.org/system/files/soups2025_poster46_abstract-ramat_v2.pdf

[24] Stuart E. Schechter, Rachna Dhamija, Andy Ozment, and Ian Fischer. 2007. The Emperor's New Security Indicators. In *IEEE Symposium on Security and Privacy*. https://doi.org/10.1109/SP.2007.35

[25] University of Chicago Information Security. 2024. MFA Fatigue Attacks: What to Watch For and What To Do. https://security.uchicago.edu/2024/10/01/mfa-fatigue-attacks/

[26] Alma Whitten and J. D. Tygar. 1999. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. In *USENIX Security Symposium*. https://www.usenix.org/conference/8th-usenix-security-symposium/why-johnny-cant-encrypt-usability-evaluation-pgp-50