Security Noir: Critical Designs for Security & Privacy

Didem Demirag
demirag.didem@uqam.ca
UQAM
Montreal, QC, Canada

Leah Zhang-Kennedy
lzhangke@uwaterloo.ca
Stratford School of Interaction Design and Business, University of Waterloo
Waterloo, ON, Canada

Jeremy Clark
j.clark@concordia.ca
Concordia University
Montreal, QC, Canada

Abstract

We introduce *Security Noir*, a critical design approach for engaging with the often hidden and elusive dimensions of security and privacy. We present ten speculative artefacts, expressed through short vignettes that situate each design within everyday contexts of use. These plausible yet provocative scenarios use exaggeration, satire, and juxtaposition to surface socio-technical tensions. We position text-based vignettes as a deliberate design material that enables speculative artefacts to be rapidly produced, shared, and critically engaged with. This work contributes an initial proof-of-concept corpus of critical design artefacts and outlines future work to formalize this approach into a more explicit and transferable design methodology.

CCS Concepts

• **Human-centered computing → Interaction design theory, concepts and paradigms**; • **Security and privacy → Human and societal aspects of security and privacy**;

Keywords

Critical Design, Speculative Design, Design Fiction, Usable Privacy and Security, Threat Modelling, Vignettes, Human-Centered Security, Socio-Technical Systems

ACM Reference Format:
Didem Demirag, Leah Zhang-Kennedy, and Jeremy Clark. 2026. Security Noir: Critical Designs for Security & Privacy. In *Designing Interactive Systems Conference (DIS Companion '26), June 13–17, 2026, Singapore, Singapore*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3802974.3809446

1 Introductory Remarks

Consider the following vignette from a collection of speculative artefacts:

> **Eleventh Finger.** Biometric authentication based on fingerprints is generally user-friendly and fast. Eleventh finger is a 3D printed rubber finger, customized with the user's fingerprint. You can put it on a keychain and use it on cold winter days when you do not want to remove your gloves. Alice lets her friend Bob borrow it when he stays at her house. It can also serve as a backup if the worst happens.

At face value, the design is a convenience-enabling accessory for biometric authentication; yet, the vignette also challenges common assumptions about fingerprints as secure, embodied identifiers. What happens when bodily identity is lent, misplaced, duplicated, or stolen? The scenario does not present a technical solution; instead, it invites reflection.

This mode of engagement is characteristic of critical design, which uses speculative artefacts and narratives to provoke reflection rather than propose deployable products. Early work is commonly attributed to Dunne and Raby [1, 2], and builds on earlier movements in art and design, from Italian radical design to critical practices in HCI [6].

Designing security software and privacy-enhancing technologies requires anticipating threats that have not yet been exploited, a process in computer security known as *threat modelling* [12]. While valuable, this work is often time-consuming for security experts and difficult for practitioners to learn [13]. Prior work has therefore explored bridging speculative design with threat identification, highlighting that threat modelling is inherently speculative, requiring the imagination of possible futures [7]. From this perspective, tools from speculative design and design fiction can support early-stage and exploratory threat identification.

These approaches are also more lightweight than traditional frameworks, making them accessible to practitioners with limited security expertise. This aligns with calls within industry and practice for "fast, cheap, and good" approaches to threat modelling. For example, Shostack + Associates [13] advocate starting with a simple but generative question, *"What can go wrong?"* Speculative methods operationalize this question through narrative and imagination, helping make abstract or overlooked risks more visible.

In this work-in-progress, we investigate critical design as a lightweight, narrative-driven approach to communicating elusive security and privacy risks. We contribute a curated corpus of ten speculative design vignettes that surface a range of security and privacy tensions in everyday life. Each artefact is designed to make a specific socio-technical risk perceptible through exaggeration, satire, or narrative juxtaposition.

2 Related Work and Research Positioning

Much has been written on critical design, and Malpass provides an excellent survey [6]. Critical design has long explored speculative artefacts as a means of provoking reflection rather than proposing deployable solutions. Table 1 presents three examples in vignette form. For example, *Compass Table*, which reflects on the invisible infrastructures surrounding everyday technologies, and *Life Counter*, which explores mortality and perception, both illustrate how critical design can render abstract or imperceptible phenomena visible and open to reflection. In contrast, *Open Informant*, engages with surveillance, suggesting that critical design vignettes may also be a useful form for security research.

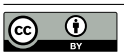

DIS Companion '26, Singapore, Singapore

ACM ISBN 979-8-4007-2632-3/2026/06
https://doi.org/10.1145/3802974.3809446

Table 1: Examples of Critical Design Vignettes

Work	Vignette
Compass Table, A. Dunne and F. Raby	This table reminds you that electronic objects extend beyond their visible limits. The 25 compasses set into its surface twitch and spin when objects like mobile phones or laptop computers are placed on it. The twitching needles can be interpreted as being either sinister or charming, depending on the viewer's state of mind. When we designed the compass table, we wondered if a neat-freak might try to make all the needles line up, ignoring the architectural space of the room in favour of the Earth's magnetic field." Quoted from [2].
Life Counter, I. Matsumoto	"With Life Counter (2001), you choose how many years you would like to or expect to live for and start the counter. Once activated, it counts down the selected time span at four different rates: the number of years, days, hours or seconds to go are shown on different faces. Depending on which face you choose to display, you may feel very relaxed as the years stretch out ahead or begin to panic as you see your life speed away before your eyes. The counter is designed to be visually unassuming and could easily fit into the slightly retro-futuristic style of the moment." Quoted from [2].
Open Informant, J Ardern, Y. Ushigome, and A. Jain	"Open Informant (2013) takes the form of a networked object including a phone app and e-ink badge. The app searches your communications for National Security Agency trigger words and then sends text fragments containing these words to the badge worn by the user for public display. Using the body as an instrument for protest, the badge becomes a means of rendering our own voice visible in an otherwise faceless technological panopticon." Quoted from [6].

While prior work in security and privacy demonstrates the value of critical and speculative design for exposing hidden assumptions, these approaches have primarily been applied within specific domains or participatory contexts. For example, Wong et al.[14] use scenario-based workbooks to elicit user values around smart home surveillance, while Pierce [10] employs speculative scenarios to examine the social implications of domestic technologies. Khovanskaya et al. [5] similarly use critical design to interrogate personal informatics systems, surfacing tensions around visibility and self-tracking. In security-focused work, Merrill [7] introduces "security fictions" to support reasoning about adversarial futures and threat identification. Rossi et al. [11] explore how speculative design and foresight methods can be applied to data protection. These efforts show that speculative and critical design can effectively surface privacy and security concerns through situated narratives, participatory engagement, and domain-specific explorations.

In comparison, our work centers on the construction of a curated corpus of standalone critical design vignettes. Rather than embedding artefacts within workshops or elicitation frameworks, we focus on designing compact, reusable narrative artefacts that independently foreground specific socio-technical tensions.

3 Methodology

This paper reports the first, exploratory phase of a two-phase process. Our goal in this phase was not to present a finalized reproducible method, but to pilot the development and refinement of a small but conceptually diverse set of vignettes. Rather than relying on purely ad hoc brainstorming, we generated concepts by systematically combining four elements:

(1) Security and privacy concerns, misconceptions, or principles
(2) Domains of application
(3) Technological modalities and form factors
(4) Anticipated user tensions and reactions

These combinations were then discussed, interpreted, and iteratively refined into speculative artefacts that aimed to remain plausible while foregrounding critical tension.

As an example, consider the *Eleventh Finger* vignette. It emerged from a combination of the four elements above. First, we aimed to address a dimension of security beyond privacy. To broaden the range of concerns under consideration, we drew on the STRIDE model [12] and selected spoofing as the focal issue. Second, authentication processes, including passwords, biometrics, and passkeys, provided an appropriate domain. Third, technologies such as smart locks, smart homes, and 3D printing informed the modality of the design. Fourth, the vignette was developed to foreground a user tension between convenience and loss of control.

In practice, each element was represented on an ideation card created by the authors as sticky notes in an online whiteboarding platform. Designs such as *Eleventh Finger* emerged by drawing and recombining cards across categories, then refining the resulting combinations into plausible but provocative vignettes. Not every combination produced a viable design. We typically drew one card from each category and used open-ended discussion to assess whether the combination could support a coherent vignette. Some were too awkward, too literal, or did not surface a meaningful security or privacy tension, and were set aside. Others required reinterpretation, recombination, or simplification before yielding a plausible near-future artefact. We retained concepts that achieved three goals simultaneously: conceptual clarity, plausibility, and distinct critical tension. While this process was effective for exploratory ideation, this process relied on open-ended interpretation and collaborative judgment, and is therefore better understood as a structured pilot than as a finalized reproducible method. The ten vignettes presented here are those that survived multiple rounds of discussion and revision. A central goal of the second phase of this project is to build on this exploratory process and formalize it into a more explicit and transferable design methodology.

4 Initial Set of Designs

Next, we present the initial set of vignettes developed through our exploratory process, illustrating how speculative artefacts can surface diverse security and privacy tensions in everyday contexts.

No(i)sy elevator. Elevator music has never been regarded as elevating your mood. This is why employees will be happier riding the noisy elevator. When it recognizes a user, it spins up a selection from their most-played songs on major music streaming platforms. Alice finds hearing her favourite Miles Davis song grounds her at the start of each workday, while Bob and Carol, coincidentally riding the elevator together, discover their shared love for 90s Britpop. David is a bit more concerned about the profanity-ridden banger that the nosy elevator plays for him and his boss.

Wifi Projector. Minimalist wifi routers have a single light indicator to show internet connectivity, such as a green light as opposed to red light. This understates the interesting and intriguing data flowing through the device as users browse the internet. By contrast, the wifi projector is a talkative router that projects visualizations onto the ceiling above where it is placed. Every website that pulls in scripts and cookies from other domains is displayed as a constellation of stars of various colours and sizes, which slowly fade away. While Alice's visit to dictionary.com is a beautiful galaxy, she is also perturbed by the number of ads and tracking companies surveilling her.

Crystal Avatars. Crystal avatars are created instantly but only get better over time. At first, they are not recognizable and cannot be differentiated from other users, but as users browse the site, small details crystallize. Alice is happy to save time registering for the site and have a personal avatar, which automatically adds a cat accessory after she spent a lunch hour reading about hairball treatments. However, she grows uneasy about the "Woman, Life, Freedom" badge, a cause she believes in but only vocalizes discriminately.

Chatty. The latest software update for Alice's voice-activated home assistant Chatty adds machine learning to better infer Alice's commands, such as adding items to her shopping list and suggesting music she might like. It does not always know whether Alice is talking to it or not, so over time it picks up on pieces of Alice's life and has trouble unlearning them. Alice is startled but pleased when Chatty chimes in with the name of 'that actor who was in that thing' she was telling her friend about. She is bemused to find a specialized vinyl cleaning solution on her shopping list after she said Bob's records smell fishy.

Dynamic Laughtrack. Digital television content, such as sitcoms, encode laughtracks as a series of cues rather than an actual recording. Smart TVs listen and classify the viewer's level of laughter on a 10-point scale (with fine-grained training over time). The TV adds laughter to the content in a sensitive and considerate manner, where the laughtrack is only marginally higher on the scale than the viewer's own laughter. This gently nudges the viewer to greater enjoyment of the program without bombarding her. It also mixes in actual past recordings of the viewer's own laughter to capitalize on emotional mirroring.

Receiver plant. The receiver plant in the front yard of Alice records the information of the people that pass by. It is a greedy plant that needs to absorb enough bluetooth data to grow. Bob is singlehandedly responsible for its lustrous canopy from walking by Alice's house every day, mostly while checking his work emails on his phone. Perhaps the plant overstepped when it displayed, "you are 15 minutes late today, are you sure you can make it to work on time?"

GOTTCHA. GOTTCHA is a human-detection system to protect online accounts being accessed by bots. It invokes the device camera to analyze if a live human is using the device. To protect against video replays or machine-generated video, it unexpectedly prompts the user with a randomly selected image and carefully measures their reaction to it. Micro-expressions are involuntary facial displays of emotion that are too fast to mimic artificially. GOTTCHA's image bank can provoke disgust, anger, fear, sadness, happiness, surprise or contempt. Over time, Alice stops visiting websites that use GOTTCHA because of its capricious tendency to mix in disturbing images.

ToS Fishing. ToS Fishing is a digital game where players earn points by hunting down excessive legalese on websites, including terms of service (ToS), privacy policies, and cookie policies. To play, users simply copy and paste the URL of the legalese. If the ToS has been seen before, the users are awarded points based on its word count. If it has not been seen before, it is reviewed by a human for validity and word count—and in this case, the user gets a finder's bonus. To mitigate people from launching custom websites, the game only accepts sites from the Alexis top million. Users display their aggregate score on a leaderboard with a profile showing their top catches (word count only, not the website), and can earn various badges for playing consistently.

Eleventh Finger. Biometric authentication based on fingerprints is generally user-friendly and fast. Eleventh finger is a 3D printed rubber finger, customized with the user's fingerprint. You can put it on a keychain and use it on cold winter days when you do not want to remove your gloves. Alice lets her friend Bob borrow it when he stays at her house. It can also serve as a backup if the worst happens.

Calm Watch. Calm Watch is a smartwatch that tracks anxiety levels based on several biomarkers. When the anxiety level of the wearer crosses a threshold, a haptic vibration pulse prompts the user to look at the watch. From there, they can select a variety of calming techniques, including breathing exercises and guided meditation. Bob is nervous to talk to Alice and becomes flustered when his watch starts vibrating mid-conversation, just loud enough for Alice to hear it and notice his fidgeting with the watch.

4.1 Discussion

Vignettes. Critical designs are often realized as physical or visual artefacts that are built, exhibited, or otherwise materialized to support engagement and interpretation. For example, Dunne and Raby's *Compass Table* in Table 1 (and the seven other objects in their *Placebo* project [2]) were fabricated and placed into participants' homes for everyday use, followed by exit surveys capturing how individuals interpreted and experienced the objects. At the same time, critical designs are frequently accompanied by narrative descriptions that situate the artefact within a context of use, making their intended critique legible.

Here, we adopt text-based vignettes not as a substitute for design, but as a deliberate design material. Each vignette describes both the form and situated use of a speculative artefact, enabling readers to imagine its integration into everyday life and engage with it critically. When successful, these descriptions support what Malpass describes as the 'rhetorical use' of critical design [6], where the artefact functions to provoke reflection rather than enable practical deployment. The designs are intentionally transgressive, crafted to evoke emotional and interpretive responses through humour, satire, ridicule, poetry, playfulness, lewdness, appropriation, irreverence, or deliberate overcompensation.

Our use of text aligns with traditions in design fiction and scenario-based design, where narrative is treated as a legitimate medium for constructing and encountering speculative artefacts. Design fiction, in particular, often employs "diegetic artefacts"; objects situated within a fictional world that make speculative technologies tangible and support reasoning about their implications, whether or not they are fully realized as physical artefacts. Similarly, value scenarios [9] use textual narratives to explore the societal and ethical implications of design decisions. In this sense, our vignettes can be understood as lightweight, narrative instantiations of critical designs that foreground use, context, and consequence.

Vignettes are also widely used in research contexts, defined as "short stories about hypothetical characters in specified circumstances, to whose situation the interviewee is invited to respond" [3]. In HCI and usable security research, such narrative constructions are used to surface assumptions, value tensions, and reflections on possible futures when direct observation or deployment is infeasible, unsafe, or premature (e.g., [4, 8]). By leveraging text as a design material, our approach prioritizes accessibility, scalability, and interpretive flexibility, enabling speculative artefacts to be rapidly generated, shared, and adapted while still supporting meaningful critical engagement.

Usage Contexts. We propose vignettes as a methodological vehicle for communicating speculative artefacts rather than as experimental instruments intended to elicit measurable responses. The vignettes function as shareable analytic artefacts that others may appropriate for reflection, inquiry, and design exploration. This approach therefore aligns with established HCI practices in which scenarios are used to probe socio-technical implications.

Beyond their role as speculative artefacts, we foresee several practical use cases of critical design vignettes for security and privacy. First, they may serve as reflective prompts in educational, participatory, or professional contexts, where speculative scenarios facilitate discussion of assumptions, values, and mental models surrounding security and privacy practices. Second, they may function as research instruments in future speculative design or vignette-based studies, where they can be adapted as stimuli to explore perceptions of risk, trust, or technological futures. Finally, we envision their use as generative tools in design processes to inspire ideation, highlight overlooked attack vectors, or prompt reconsideration of design assumptions before technologies are implemented.

We therefore position *Security Noir* (an homage to Dunne and Raby's term 'Design Noir' [2]) as an approach for engaging with the often uncomfortable or hidden aspects of security and privacy technologies and their impacts on human experience. Methodologically, this approach takes the form of a lightweight, imaginable, and scalable medium that allows speculative artefacts to be encountered without technological implementation, while remaining grounded in established HCI traditions of using vignettes in survey research, design fiction, and scenario-based studies. These critical designs intend to activate readers' existing mental models and experiences, supporting reflective engagement with security and privacy risks that are otherwise difficult to visualize or experience directly.

4.2 Future Work

As a work-in-progress, this paper reports on the initial exploratory phase of our approach. Future work will focus on formalizing the design process into a reproducible method, expanding the space of concerns and domains considered, and evaluating how such vignettes function in practice (e.g., in workshops, educational settings, or participatory design contexts). Through this progression, we aim to establish critical design vignettes as a viable methodological tool within usable security and HCI research.

While this work demonstrates our approach based on concerns, domains, modalities, and user tensions, future work will expand and refine these dimensions through a more comprehensive synthesis of usable security and privacy literature. This includes developing more exhaustive and theoretically grounded sets of security and privacy concerns, misconceptions, and principles, as well as broadening the range of stakeholders and contexts.

We plan to more systematically document and analyze the design process itself. In this work, ideation and refinement were conducted through iterative discussion and selection; future work will more explicitly capture this process through structured design tools (e.g., ideation cards or generative frameworks) and reflective documentation. This will enable us to better articulate how concepts are generated, filtered, and developed, addressing current limitations in methodological transparency and supporting adoption by other researchers. For example, we have begun to develop structured ideation cards representing concepts, misconceptions, domains, and interaction modalities. These cards will be combined in small sets to scaffold exploration and examine how constrained prompts shape design diversity, thematic coverage, and narrative plausibility. We will also explore generative AI as an ideation partner (the designs in this paper were developed without AI assistance). To better understand these processes, we plan to document the design work through autoethnography to identify recurring practices and challenges, including topic selection and balancing plausibility with provocation.

We aim to evaluate how these vignettes function in practice in user studies. While this paper positions them as communicative probes, our future work will examine how different audiences interpret and engage with the designs, such as software develpers and other non-specialists. In particular, we are interested in how vignettes shape users' threat identification of security and privacy risks, what kinds of discussions they provoke and whether they reveal gaps, misconceptions, or alternative perspectives that are not easily captured through traditional methods.

Finally, we will explore how critical design vignettes can be integrated into existing security practices, including threat modelling and design ideation workflows. Given the inherently speculative nature of threat identification, we see potential for these artefacts to complement conventional approaches by supporting early-stage exploration of *"what can go wrong"*, as suggested by Shostack + Associates [13], in ways that are accessible to non-specialists. Future work will investigate how such methods can be incorporated into interdisciplinary teams and whether they can meaningfully inform design and decision-making processes. Through these directions, we aim to move from an initial corpus of designs toward a more robust methodological framework for using speculative artefacts as tools for reflection, communication, and inquiry within usable security and HCI research.

Acknowledgments

J. Clark acknowledges support for this research project from the National Sciences and Engineering Research Council of Canada (NSERC) through the NSERC, Raymond Chabot Grant Thornton, and Catallaxy Industrial Research Chair in Blockchain Technologies IRCPJ/545498-2018 (https://www.nserc-crsng.gc.ca/Chairholders-TitulairesDeChaire/Chairholder-Titulaire_eng.asp?pid=1045) and a Discovery Grant RGPIN/04019-2021

References

[1] Anthony Dunne. 2005. *Hertzian Tales.* MIT.
[2] Anthony Dunne and Fiona Raby. 2000. *Design Noir: The Secret Life of Electronic Objects.* Birkhauser.
[3] Janet Finch. 1987. The vignette technique in survey research. *Sociology* 21, 1 (1987), 105–114.
[4] Hilda Hadan, Derrick M Wang, Lennart E Nacke, and Leah Zhang-Kennedy. 2024. Privacy in immersive extended reality: Exploring user perceptions, concerns, and coping strategies. In *ACM CHI Conference on Human Factors in Computing Systems.*
[5] Vera Khovanskaya, Eric PS Baumer, Dan Cosley, Stephen Voida, and Geri Gay. 2013. "Everybody knows what you're doing" a critical design approach to personal informatics. In *ACM CHI Conference on Human Factors in Computing Systems.* 3403–3412.
[6] M Malpass. 2017. *Critical Design in Context: History, Theory, and Practices.* Bloomsbury.
[7] Nick Merrill. 2020. Security fictions: Bridging speculative design and computer security. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference.* 1727–1735.
[8] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. 2017. Privacy expectations and preferences in an {IoT} world. In *Symposium on Usable Privacy and Security (SOUPS).*
[9] Lisa P Nathan, Batya Friedman, Predrag Klasnja, Shaun K Kane, and Jessica K Miller. 2008. Envisioning systemic effects on persons and society throughout interactive system design. In *Proceedings of the 7th ACM conference on Designing interactive systems.* 1–10.
[10] James Pierce. 2019. Lamps, Curtains, Robots: 3 scenarios for the future of the smart home. In *Proceedings of the 2019 Conference on Creativity and Cognition.* 423–424.
[11] Arianna Rossi, Regis Chatellier, Stefano Leucci, Rossana Ducato, and Estelle Hary. 2022. What if data protection embraced foresight and speculative design? *DRS2022: Bilbao* (2022).
[12] Adam Shostack. 2014. *Threat modeling: Designing for security.* John wiley & sons.
[13] Adam Shostack. 2021. *Fast, Cheap, and Good: An Unusual Trade-off Available in Threat Modeling.* Technical Report. Shostack + Associates. https://shostack.org/files/papers/Fast-Cheap-and-Good.pdf Accessed: 2026-04-10.
[14] Richmond Y Wong, Jason Caleb Valdez, Ashten Alexander, Ariel Chiang, Olivia Quesada, and James Pierce. 2023. Broadening privacy and surveillance: Eliciting interconnected values with a scenarios workbook on smart home cameras. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference.* 1093–1113.