Editors: Yasemin Acar, yasemin.acar@uni-paderborn.de | Jens Grossklags, jens.grossklags@in.tum.de | Julie Haney, julie.haney@nist.gov | Heather Richter Lipford, richter@uncc.edu

Extending the Heilmeier Catechism to Evaluate Security and Privacy Systems

Who is Left Out?

Kevin Butler | University of Florida
Kurt Hugenberg | Indiana University
Eakta Jain | University of Florida
Apu Kapadia | Indiana University
Tadayoshi Kohno | University of Washington
Elissa M. Redmiles | Georgetown University
Franziska Roesner | University of Washington
Mattea Sim | Indiana University
Patrick Traynor | University of Florida
Hanna Barakat | Human Computing Associates

The Heilmeier Catechism consists of a set of questions that researchers and practitioners can consider when formulating research and applied engineering projects. In this article, we suggest explicitly asking who is included and who is left out of consideration.

The Heilmeier Catechism[a] was developed by then DARPA[b] director George H. Heilmeier in the 1970s to help DARPA officials evaluate and refine proposed research programs. The Catechism serves as a systematic framework for researchers and program managers to assess a project's viability and potential impact. The Catechism is still widely used today by both researchers and practitioners to guide foundational engineering research and applied engineering projects. The Heilmeier questions are as follows:

©SHUTTERSTOCK.COM/NICOELNINO

- What are you trying to do? Articulate your objectives using absolutely no jargon.
- How is it done today, and what are the limits of current practice?
- What is new in your approach, and why do you think it will be successful?
- Who cares? If you are successful, what difference will it make?
- What are the risks?
- How much will it cost?
- How long will it take?

[a]catechism: "a set of formal questions put as a test" — Merriam-Webster. https://www.merriam-webster.com/dictionary/catechism
[b]Defense Advanced Research Projects Agency — https://www.darpa.mil/

Digital Object Identifier 10.1109/MSEC.2025.3554848
Date of current version: 13 June 2025

- What are the midterm and final "exams" to check for success?

In this article, we revisit the Heilmeier Catechism in the context of the design of modern computer security and privacy systems. The original Heilmeier Catechism challenges researchers, program managers, product teams, engineers, and others to think in a systematic way about the broader context behind what they are proposing and why what they are proposing matters—similar to traditional computer security threat modeling. But left unspecified in these questions is who is or should be included under the umbrella of those who are affected in the first place. When a security or privacy solution is proposed by a product team or researchers, what assumptions are made about its population of users and therefore about the security risks that are relevant to them? We propose the addition of the question:

- Who is left out? (See also "The Heilmeier Catechism: Extended.")

Why Ask "Who Is Left Out?"

Being "left out" can result in direct harm. In the context of modern computer security, being "left out" has led to prevailing threat models and security-related technology designs that fail to consider the needs and contexts of users beyond the assumed "default user" (i.e., whom designers had in mind during design). Consider the following examples:

Online account security measures, such as account recovery options, password guidelines, and security questions, do not always consider a broader range of use cases, threat models, and cultural contexts. Prior research finds that the use cases or threat models of some user groups make commonly implemented protective mechanisms ineffective.

For example, for victim–survivors of intimate partner violence,[1] notifying an account "owner"—who may be an abusive partner—when recovery details have changed can increase rather than decrease the threat of harm. In many contexts, e.g., cybercafe users in Kenya,[2] users may rely on partially trusted community members to help manage their online accounts, leading to security challenges often not considered in mainstream security research. As another example, early CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart) designs did not sufficiently consider people with visual impairments.[3]

Leaving out stakeholder groups can lead to gaps in a system's threat model, which can in turn result in failure to anticipate harmful dual use of technology.

Leaving out stakeholder groups can lead to gaps in a system's threat model, which can in turn result in failure to anticipate harmful dual use of technology. For example, technologies such as smart home devices can enable dangerous misuse[4] in settings like intimate-partner abuse when an interpersonal threat model is not considered.

These gaps are not unique to computer security, and readers may be familiar with examples from other contexts as well. For example, consider automatic speech recognition tools that fail to account for certain linguistic lexicons or hiring algorithms that unfairly penalize applicants for attributes unrelated to their qualifications for the job (e.g., their names). There is an even longer history of noncomputing technologies that perform more poorly overall because they did not consider all communities in their design and/or testing. Consider, for instance, the first film cameras that were designed to capture only light skin tones or that early automotive crash test dummies were only modeled after men.

The Heilmeier Catechism: Extended

- What are you trying to do? Articulate your objectives using absolutely no jargon.
- How is it done today, and what are the limits of current practice?
- What is new in your approach, and why do you think it will be successful?
- Who cares? If you are successful, what difference will it make?
- What are the risks?
- How much will it cost?
- How long will it take?
- What are the midterm and final "exams" to check for success?
- Who is left out?

Acknowledging Biases

The question "who is left out?" compels us to confront both explicit and implicit omission. Research has shown that system designers and practitioners often hold both explicit and implicit biases. Explicit bias refers to conscious and deliberate attitudes or beliefs that individuals are aware of and can express openly. Implicit bias, on the other hand,

involves attitudes or stereotypes that can unconsciously and unintentionally affect one's understanding, actions, and decisions.[5] Often, the implicit and explicit biases that emerge reflect who is building the systems.[6] For example, affinity bias (an implicit bias) describes the phenomenon of individuals favoring or implicitly prioritizing people who are similar to themselves.

To mitigate the potential influence of biases, we must begin by acknowledging them. We can see these biases unaccounted for in the logic of the questions of Heilmeier's Catechism. For example, "what are the risks?" is often interpreted as "what are the risks to the (assumed) primary stakeholders?" or "what are the risks that this might not succeed?" However, no question prompts researchers to think about who might be left out (or excluded) of consideration when a system's benefits or impacts are evaluated. In short, the way a researcher interprets and answers Heilmeier's "who cares?" may be different from "who should care?", "who might care?" or "who might be impacted?"

Addressing the question "Who is left out?" is rarely straightforward. A generative approach to this challenge might involve reflecting on its inverse: "Who is included?" That is, identifying the subjectivity of those conducting the research or designing the technology and identifying the stakeholders on which they focused during the design process (see also "Sample Approaches for Identifying and Considering Perspectives "Left Out"""). Reflecting on one's subjectivity provides a valuable lens to address how our subjectivities (and implicit biases) become embedded in systems, allowing space to step back and consider the question of "Who is left out?"

Addressing Who Is Left Out

How might one go about identifying which stakeholder group is "left

Sample Approaches for Identifying and Considering Perspectives "Left Out"

Participatory action research (PAR):[S1] PAR encourages the active involvement of community members directly affected by the research topic. This approach aims to foster cocreation and empower participants to contribute to the research design, data collection, analysis, and decision-making. PAR methods emphasize active stakeholder involvement throughout the research process. Implementations of participatory research in security include collaborating with community members as peers on research teams[S2] and participatory threat modeling.[S3]

Design Thinking: Design thinking emphasizes empathy, ideation, and cocreation to develop user-centered technologies. The primary focus is on understanding user needs and preferences to create innovative solutions. This approach prompts the expansion of the question "Who is left out?" to consider a range of design (and research) questions, outlined by Sasha Costanza-Chock. These include questions about beneficiaries (who do we design for or with?), values (what values do we encode or reproduce in the objects and systems we design?), sites (where do we design; what design sites are considered or ignored?), and ownership (who owns and profits from outcomes?).

Value sensitive design (VSD): VSD is an approach to designing technology that addresses the ethical considerations of direct and indirect stakeholders throughout the design process. A key component of VSD is stakeholder identification. The process of identifying and incorporating a broad spectrum of stakeholders gives rise to a range of conceptual, empirical, and technical methodologies. Similar to PAR, VSD often involves an iterative process where technology is cocreated and designs refined based on feedback. In the context of security and privacy research, specifically, Bellini et al. offer pragmatic strategies for digital-safety research involving at-risk users,[10] and Sim et al. offer evidence that a simple intervention encouraging students to think about a broader range of stakeholders during threat modeling exercises can have a measurable impact.[S4]

References

S1. P. Reason and H. Bradbury, Eds., *The SAGE Handbook of Action Research: Participative Inquiry and Practice*, 2nd ed. London, U.K.: Sage, 2007.

S2. R. Bhalerao, V. Hamilton, A. McDonald, E. M. Redmiles, and A. Strohmayer, "Ethical practices for security research with at-risk populations," in *Proc. IEEE Eur. Symp. Secur. Privacy Workshops*, 2022, pp. 546–553, doi: 10.1109/EuroSPW55150.2022.00065.

S3. J. Slupska, S. D. Dawson Duckworth, L. Ma, and G. Neff, "Participatory threat modelling: Exploring paths to reconfigure cybersecurity," in *Proc. Extended Abstr. CHI Conf. Human Factors Comput. Syst.*, 2021, pp. 1–6.

S4. M. Sim, K. Hugenberg, T. Kohno, and F. Roesner, "A scalable inclusive security intervention to center marginalized & vulnerable populations in security & privacy design," in *Proc. New Secur. Paradigms Workshop*, 2023, pp. 102–115.

out"? And, once a stakeholder or group has been identified as "left out", how should these perspectives best be included? In the past several years, the computer security and privacy research community has increasingly foregrounded studies of marginalized and vulnerable populations.[7] Calls to consider users beyond the "default" are also (and were already) commonplace in other research communities (e.g., the human-computer interaction community). We provide a sidebar with a small sample of approaches and resources drawn from these communities that can be used to broaden the perspectives considered in security and privacy research and design: for example, participatory research methods that ensure engagement from affected stakeholders, and simple interventions for broadening designers' perspectives beyond a default persona. Such approaches can allow for more impactful innovation and creative ideation around privacy and security solutions, benefiting more stakeholders.

That said, these approaches are not free from their own socio-technical biases. For instance, while VSD offers pragmatic strategies for at-risk users, at least early instantiations assumed a set of "universal values" that subscribed to Western norms and could overlook broader cultural nuances.[8] Prior research also sheds light on how well-intended efforts to include voices typically "left out" can unintentionally place a burden on marginalized voices and lead to more harm[9]—for example, conducting usability testing to "check a box" without a genuine effort to address a population's specific needs, or asking to collaborate with someone because the research needs translation services and they speak the desired language. Thus, there is a growing push to move beyond "incorporating" the perspectives of those left out toward materially empowering left-out communities vis-a-vis education, resource distribution, coownership, and so on (as seen in PAR).

Reflecting on "Who is left out?" is a critical first step toward developing even more robust security and privacy systems that work well for more people.

Finally, it is important to keep in mind that, once "Who is left out?" is considered, ethical concerns or other considerations may lead to the ultimate conclusion to not explore a particular research direction at all.

Reflecting on "Who is left out?" is a critical first step toward developing even more robust security and privacy systems that work well for more people. With this piece, we encourage anyone using the Heilmeier Catechism — and, more broadly, anyone proposing or assessing a product, project, or solution, whether or not using the Heilmeier Catechism — to explicitly consider "Who is left out?" This perspective is especially important in the context of the design of security and privacy mechanisms where design decisions can have far-reaching implications for people's security, privacy, and safety. We join our colleagues in computer security and beyond[10,11] calling for effective consideration of those perspectives to ensure safer computing for all. ■

References

1. A. Daffalla, M. Bohuk, N. Dell, R. Bellini, and T. Ristenpart, "Account security interfaces: Important, unintuitive, and untrustworthy," in *Proc. USENIX Security Symp.*, 2023, pp. 3601–3618.
2. C. W. Munyendo, Y. Acar, and A. J. Aviv, ""In eighty percent of the cases, I Select the password for them": Security and privacy challenges, advice, and opportunities at cybercafes in Kenya," in *Proc. IEEE Symp. Secur. Privacy*, 2023, pp. 570–587, doi: 10.1109/SP46215.2023.10179410.
3. World Wide Web Consortium (W3C), "Inaccessibility of CAPTCHA: Alternatives to visual Turing tests on the web," Dec. 2021. [Online]. Available: https://www.w3.org/TR/turingtest/
4. S. Stephenson, M. Almansoori, P. Emami-Naeini, D. Y. Huang, and R. Chatterjee, "Abuse vectors: A framework for conceptualizing IoT-enabled interpersonal abuse," in *Proc. USENIX Secur. Symp.*, 2023, pp. 69–86.
5. Y. Wang and D. Redmiles, "Implicit gender biases in professional software development: An empirical study," in *Proc. IEEE/ACM 41st Int. Conf. Softw. Eng. Softw. Eng. Soc. (ICSE-SEIS)*, 2019, pp. 1–10, doi: 10.1109/ICSE-SEIS.2019.00009.
6. S. Wachter-Boettcher, *Technically Wrong: Sexist Apps, Biased Algorithms, and Other Threats of Toxic Tech*. New York, NY, USA: Norton, 2017.
7. N. Warford et al., "SoK: A framework for unifying at-risk user research," in *Proc. IEEE Symp. Secur. Privacy*, 2022, pp. 2344–2360, doi: 10.1109/SP46214.2022.9833643.
8. A. Borning and M. Muller, "Next steps for value sensitive design," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2012, pp. 1125–1134, doi: 10.1145/2207676.2208560.
9. C. A. Liang, S. A. Munson, and J. A. Kientz, "Embracing four tensions

in human-computer interaction research with marginalized people," *ACM Trans. Comput.-Hum. Interact.*, vol. 28, no. 2, pp. 1–47, 2021, doi: 10.1145/3443686.
10. R. Bellini et al., "SoK: Safer digital-safety research involving at-risk users," in *Proc. IEEE Symp. Secur. Privacy*, 2024, pp. 635–654, doi: 10.1109/SP54263.2024.00071.
11. Y. Wang, "The third wave? Inclusive privacy and security," in *Proc. New Secur. Paradigms Workshop*, 2017, pp. 122–130.

Kevin Butler is a professor of computer and information science and engineering and director of the Florida Institute for Cybersecurity Research at the University of Florida, Gainesville, FL 32611 USA. His research interests include the security and trustworthiness of systems and the security of users accessing them. Butler received a Ph.D. in computer science and engineering from The Pennsylvania State University. He is a Senior Member of IEEE. Contact him at butler@ufl.edu.

Kurt Hugenberg is the James H. Rudy Professor of Psychology in the Department of Psychological and Brain Sciences at Indiana University, Bloomington, IN 47405 USA. His research interests include intergroup relations and social cognition, with a particular focus on how stereotypes can influence or distort how we perceive other people and groups. Hugenberg received a Ph.D. in social psychology from Northwestern University. Contact him at khugenb@iu.edu.

Eakta Jain is an associate professor of computer and information science and engineering at the University of Florida, Gainesville, FL 32611 USA. Her research interests include safety, privacy, and security of data gathered for user modeling, particularly eye tracking data. Jain received a Ph.D. in robotics from Carnegie Mellon University. She is a Member of IEEE and an ACM Senior Member. Contact her at ejain@ufl.edu.

Apu Kapadia is a professor of computer science at the Luddy School of Informatics, Computing, and Engineering, Indiana University Bloomington, Bloomington, IN 47405 USA. His research interests include computer security and privacy, with an emphasis on usable security and human-centered computing. Kapadia received a Ph.D. in computer science from the University of Illinois at Urbana-Champaign. He is a Member of IEEE. Contact him at kapadia@iu.edu.

Tadayoshi Kohno is a professor in the Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, WA 98195 USA. His research interests include helping protect the security, privacy, and safety of users of current- and future-generation technologies. Kohno received a Ph.D. in computer science from the University of California, San Diego. He is a Fellow of IEEE. Contact him at yoshi@cs.washington.edu.

Elissa M. Redmiles is the Clare Luce Booth Assistant Professor of Computer Science at Georgetown University in Washington, DC 20007 USA. Her research interests include security, privacy and ethics. She received a Ph.D. in computer science from the University of Maryland. Contact her at elissaredmiles.com.

Franziska Roesner is a professor in the Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, WA 98195 USA. Her research interests include identifying and addressing security, privacy, and safety challenges faced by end users of existing and emerging technologies. She received a Ph.D. in computer science and engineering from the University of Washington. She is a Member of IEEE. Contact her at franzi@cs.washington.edu.

Mattea Sim is a postdoctoral researcher in the Department of Psychological and Brain Sciences at Indiana University, Bloomington, IN 47405 USA. Her research interests include computing practices at the interface of social psychology and computer security and privacy. She received a Ph.D. in psychology from Indiana University. Contact her at matsim@iu.edu.

Patrick Traynor is the John and Mary Lou Dasburg Preeminent Chair in Engineering in the Department of Computer and Information Science and Engineering at the University of Florida, Gainesville, FL 32611 USA. His research interests include systems and network security, particularly in relation to cellular and mobile systems. Traynor received a Ph.D. in computer science and engineering from The Pennsylvania State University. He is a Senior Member of IEEE and ACM. Contact him at traynor@ufl.edu.

Hanna Barakat is a researcher at Human Computing Associates, Washington, DC 20012 USA. Her research interests include participatory technology design and critical approaches to trust and safety. She received a B.A. in international development studies from Brown University. Contact her at hanna_barakat@alumni.brown.edu.