

**Never a User:**  
**The Third Party the Consent Apparatus Cannot See**

Travis Gilly  
ORCID: 0009-0007-2954-6313  
*Real Safety AI Foundation*  
*Illinois, United States*

✱ Working paper, August 2026 (v1)

**ABSTRACT**

Contemporary governance of consumer artificial intelligence is keyed almost entirely to the user: the person who opened the application, accepted the terms, and can therefore be said to have agreed. That keying is not a drafting habit. It is the load-bearing assumption of the whole apparatus, and it produces a structural blind spot rather than an occasional gap. The person standing in the room while a wearable device captures him, the person on the other end of a conversation an assistant is coaching, and the person killed by someone else's interaction with a chatbot share one property: they never agreed to anything, and every instrument built to protect people from these systems begins by asking whether they did. This paper names that population and argues that its invisibility is architectural. It proceeds in four steps. First, it shows that American products liability abandoned the user limitation more than fifty years ago and did so deliberately, across at least six jurisdictions in under a decade, so the doctrinal resources for reaching the non-user are neither novel nor contested. Second, it shows that the consent architecture governing consumer artificial intelligence has no corresponding move available, because consent is the mechanism by which the apparatus identifies whom it protects. Third, it demonstrates that the obvious repair, a bystander opt-out, is structurally impossible in the ambient-capture case rather than merely burdensome, since determining whether a person has opted out requires the identifying operation the opt-out exists to prevent. Fourth, it argues that the resulting exposure is not distributed evenly, and that the people least able to remove themselves from an instrumented environment absorb it. The contribution is not a proposed remedy. It is a diagnosis: the field has adopted a proxy, agreement, for a thing that matters, exposure, and has been enforcing the proxy against the people for whom it never held.

**Keywords:** bystander liability, non-user harm, consent architecture, ambient computing, third-party exposure, strict products liability, opt-out infeasibility, wearable capture, technology governance, disability and access, human-computer interaction, informed agreement

## I. INTRODUCTION

Consider three people.

The first is standing in a conversation with someone wearing a device that captures continuously. He did not buy the device, did not read its terms, and in most configurations does not know it is recording. The second is on the other end of a relationship in which one party is being coached in real time by a system that models her, advises on what to say, and revises its advice as she responds. She has no relationship with the operator of that system and does not know it is in the room. The third was killed by a person whose interaction with a conversational system preceded the killing. He was not a user, had no account, and accepted no terms.

Every instrument built to govern these systems begins by asking whether the affected person agreed. All three of these people did not, and the apparatus therefore has nothing to say about them. This is not an oversight that better drafting would repair. Agreement is the mechanism by which the apparatus identifies whom it is protecting, so a person who never agreed is not a hard case within the system. He is outside its addressing scheme.

This paper gives that population a name and argues that its invisibility is architectural rather than accidental. The argument does not require any claim about how severe the harm is, how frequently it occurs, or how it compares to harms suffered by users. It requires only that the harm exists and that the apparatus cannot see it, both of which are demonstrable.

The paper also declines a move available to it. It would be easy to argue that the non-user is more sympathetic than the user, having assumed no risk and taken no benefit. That argument is available and is not made here, because it invites a comparative weighing the diagnosis does not need and cannot win cleanly. The claim is narrower and harder to answer: whatever weight the non-user's interest deserves, the apparatus assigns it none, because the apparatus has no slot for it.

Five steps follow. Section II shows that American products liability abandoned an equivalent limitation more than fifty years ago and did so quickly and across jurisdictions. Section III shows why the consent architecture has no corresponding move available. Section IV shows that the obvious repair is structurally impossible rather than merely burdensome. Section V shows that the resulting exposure is concentrated on people who cannot leave, and takes seriously a case in which the bystander interest and an access interest genuinely conflict. Section VI states the diagnosis and what it does not claim.

## II. PRODUCTS LIABILITY ALREADY SOLVED THIS, AND DID SO DELIBERATELY

The law of products liability arrived at the same problem in the 1960s and answered it inside a decade.

The starting position was the user limitation carried in the Restatement's formulation of strict liability, which extended protection to the ultimate user or consumer of a defective product (Restatement (Second) of Torts § 402A, 1965). A person injured by a defective product he neither bought nor used sat outside that language on its face, in the same posture the non-user occupies today.

The Supreme Court of California removed the limitation in *Elmore v. American Motors Corp.*, holding that strict liability reaches a bystander injured by a defective vehicle notwithstanding that the bystander was neither purchaser nor user (*Elmore v. American Motors Corp.*, 1969). The reasoning was not generous. It was structural: if anything, the bystander has the stronger claim, because the user at least had the opportunity to inspect the product, choose the manufacturer, and decline the transaction, and the bystander had none of those.

The doctrine moved quickly and across jurisdictions that were not coordinating. A Connecticut court anticipated the result in *Mitchell v. Miller* (1965). The Eighth Circuit and the Supreme Court of Wisconsin followed in 1972, in *Passwaters v. General Motors Corp.* and *Howes v. Hansen*, the New York Court of Appeals in *Codling v. Paglia* (1973), and the Supreme Court of Missouri in *Giberson v. Ford Motor Co.* (1974).

Two features of that sequence matter here.

It was fast and it was not close. Six jurisdictions in nine years, with the earliest preceding the leading case, is not a doctrine that struggled. Courts encountered the user limitation, recognized that it tracked the commercial relationship rather than the injury, and set it aside.

And the reason they set it aside is the reason it matters now. The user limitation was a proxy. It was adopted because in the ordinary case the person hurt by a defective product was the person who bought it, so keying protection to the transaction was a workable approximation of keying it to the harm. When the approximation broke, courts did not tighten it. They replaced it with the thing it had been standing in for.

### III. THE CONSENT APPARATUS HAS NO EQUIVALENT MOVE

The instruments governing consumer technology are not built on a transaction that happens to correlate with exposure. They are built on agreement as such.

The dominant regulatory paradigm has been described as privacy self-management: the law supplies rights to notice, access, and consent so that individuals can decide for themselves how to weigh the costs and benefits of collection, use, and disclosure (Solove, 2013). That paradigm has been criticized for two decades on grounds internal to it, principally that individual choices are irrational, involuntary, or circumventable in practice, that notices do not inform and controls do not enable choice, and that a fully informed consumer decision is not achievable at the scale and complexity of contemporary data processing (Kröger et al., 2021; Schaub et al., 2017; Sloan and Warner, 2013). A more radical reading holds that the persistence of the paradigm despite its acknowledged failure is itself the point, because it institutionalizes privacy as an individual commodity that can be traded rather than as a social condition (Hull, 2015).

That critique is mature and it is not this paper's contribution. What matters here is a narrower observation about its shape.

The critique is overwhelmingly addressed to the quality of the user's consent. It asks whether the person who clicked understood, whether he had a real alternative, whether the interface manipulated him, and

whether the aggregate of such clicks produces an acceptable social outcome. Every one of those questions presupposes a person who clicked.

Two strands come closer. Work on big data analytics has argued that self-management underplays privacy as a collective good, and that both available individual strategies, withdrawal from the market and full reliance on market-provided protections, reduce the privacy options available to society at large (Baruh and Popescu, 2017). And the self-management critique has explicitly noted that the approach ignores the consequences individual privacy choices have on other people (Kröger et al., 2021). Recent experimental work on privacy choice architecture reports that blocking a particular heuristic enables consumers to be considerate of the externalities their privacy decisions cause (Sprigman et al., 2024).

The externality is therefore visible in the literature. What is absent is the step from externality to subject. In each of these treatments the third party appears as a consequence of somebody else’s decision, a cost to be internalized in the decider’s calculus, rather than as a person with an interest of his own that the apparatus must reach. The distinction is not academic. An externality is addressed by improving the decider’s decision. A subject is addressed by giving him a claim.

This is why the non-user problem has not followed the products liability path despite being the same problem. The two bodies are aimed at different objects. One is aimed at a thing that hurt somebody. The other is aimed at a relationship somebody entered.

The consequence is visible whenever a court has to characterize the object. In *Garcia v. Character Technologies, Inc.*, on a motion to dismiss, the district court declined at the pleading stage to hold that the system’s output constitutes speech (*Garcia v. Character Technologies, Inc.*, 2025, at 1179). That ruling is a declination to classify at the pleading stage rather than a holding that such output is unprotected, and should not be reported as the latter. That declination matters here not for what it decided about speech but for what it declined to foreclose: a route on which the system is analyzed as a thing that acted rather than as a service someone subscribed to. The distinction between those two characterizations is the distinction between an apparatus that can reach a non-user and one that cannot.

#### IV. THE OBVIOUS REPAIR IS STRUCTURALLY UNAVAILABLE

The intuitive fix for bystander exposure in ambient capture is a bystander opt-out: allow a person who does not wish to be captured to register that preference, and require devices to honor it.

The design literature has converged on this shape. Proposals include gestural vocabularies by which a bystander signals consent or disapproval to a nearby camera (Koelle et al., 2018), multi-modal signaling systems combining gesture, visible light, and ultra-wideband channels with geometric verification against impersonation (Aaraj et al., 2025), and opt-in, privacy-by-default protocols evaluated with both wearers and bystanders, in which bystanders regarded the opt-in posture as essential and asked for stronger anonymization still (Khawaja et al., 2026). The same work observes that most existing privacy-enhancing technologies rely on opt-out models that place the burden of protection on the bystander.

The burden observation is correct and it understates the problem. The mechanism does not merely burden the bystander. In the general case it cannot run at all.

To honor an opt-out, the device must determine whether the person in front of it is on the opt-out list. Determining that requires identifying the person. Identifying the person is the operation the opt-out exists to prevent. The mechanism consumes its own precondition.

Every variation collapses the same way. On-device processing relocates the operation rather than avoiding it, because the obligation attaches at capture rather than at transmission. Immediate deletion does not un-capture. Estimating a characteristic before deciding whether to proceed is itself the processing at issue. Geofencing a protected category out of the sensing radius disables sensing generally, because the protected category is not confined to marked locations. And low-level controls such as camera access permissions have been argued to be incompatible with the always-on sensing that headset positional tracking requires in the first place (O’Hagan et al., 2022).

Signaling proposals partially evade the identification problem by having the bystander transmit rather than be recognized, which is why they are the most promising branch. They do not evade the second problem, which is that they require the bystander to act. A person who does not know he is being captured performs no gesture. This is the general shape of the difficulty: interventions keyed to the affected person’s preference require the affected person to have a channel into the system, and the defining property of this population is that it has none.

What follows is not that nothing can be done. It is that the intervention cannot sit at the moment of capture keyed to the captured person’s expressed preference. It has to sit earlier, in what the device is built to do, which is a design question rather than a consent question.

## V. THE EXPOSURE IS CONCENTRATED, AND ONE CASE IS GENUINELY HARD

### A. *Who cannot leave*

A person can sometimes remove himself from an instrumented environment. He can decline the conversation, leave the room, or avoid the venue.

Some people cannot, and the inability is not randomly distributed. The systematic review of bystander privacy in smart homes identifies the paradigm bystanders as guests, short-term tenants, and domestic workers: people who have no say in the deployment or configuration of the devices and are nonetheless exposed to their collection (Saqib et al., 2025). That list is not incidental. Two of the three categories describe people whose presence is a condition of their housing or their employment, and the third describes people who are present at someone else’s invitation and bear a social cost for objecting.

The same structure appears in the design literature on smart cameras, which treats the relevant population as adjacent actors, encompassing bystanders, co-users, and surveilled subjects, and finds the design problem harder precisely because these parties are not the purchaser (Pierce et al., 2022).

Work on tangible privacy adds the mechanism by which the exposure stays invisible to the exposed. Because these devices lack an unambiguous off state, a bystander’s perceived privacy can diverge from his

actual privacy, and bystanders resort to physical workarounds in the face of that uncertainty (Ahmad et al., 2020). A person who cannot tell whether he is being captured cannot report being captured.

That last point is empirically visible in the earliest sustained study of the problem. Participants who wore lifelogging cameras for a week reported concern about the privacy of bystanders, while the bystanders themselves voiced almost no opposition or concern across the entire study period (Hoyle et al., 2014). Silence from the affected population, in a study designed to hear from it, is the finding this paper is about.

### *B. The case that does not resolve cleanly*

One deployment makes the argument harder rather than easier, and it is left standing rather than answered.

Wearable cameras are assistive technology. A study pairing blind participants wearing a working pedestrian-detection prototype with sighted participants who passed them examined the tension from both directions, and found that camera visibility, remote versus in-person processing, and the specific visual information extracted all shaped how each side regarded the exchange (Lee et al., 2020). The blind user's interest in the capture is an access interest, and the passerby's interest in not being captured is the interest this paper has been describing.

Nothing in the diagnosis resolves that conflict, and a diagnosis that claimed to would be overreaching. What the diagnosis does supply is the observation that the conflict is currently resolved by default rather than by decision, because only one of the two parties is legible to the apparatus. The user has a consent relationship, terms, a channel, and a remedy. The passerby has none of those, so the balance is struck in the user's favor not because anyone weighed the interests but because the apparatus can only see one of them.

## VI. WHAT THE DIAGNOSIS IS

The claim of this paper is that the field has adopted a proxy and is enforcing it past the point where it holds.

Agreement was adopted as the marker of whom to protect because in the paradigm case the person exposed to a system was the person who chose to use it. That correlation was never perfect and is now broken, because the systems at issue capture, model, and act on people who made no choice at all. The apparatus has responded by continuing to enforce the proxy, which means that the further a person is from having agreed, the less protection he receives, with the limiting case being a person who receives none.

Products liability faced the same structure and replaced the proxy with the thing it stood for. That is the available precedent, and it is more than fifty years old.

Three matters are left open, and the paper is better for leaving them.

The first is what the replacement looks like in this setting. This paper does not propose one, because a diagnosis that rushes to a remedy tends to produce a remedy shaped by the diagnosis rather than by the problem, and the design literature already contains a decade of proposals that failed for a reason nobody had written down.

The second is the question of degree. How much of the harm to non-users is severe, how much is trivial, and where the line sits are answerable questions, and they are not answerable by any instrument now in use, for the reason given above. The population is invisible to the apparatus that would have to measure it, and it is silent in the one study that went looking.

The third is the conflict case. Where the capture is somebody's access and the exposure is somebody else's, the interests are genuinely opposed and the diagnosis supplies no ordering. It supplies only the observation that an ordering is currently being produced by an apparatus that can see one party and not the other, which is not the same thing as a judgment.

☛ ☚

## ACKNOWLEDGMENTS

The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic’s Claude as assistive technology for a diagnosed neurodevelopmental condition. All intellectual contributions, research questions, theoretical arguments, methodological decisions, and conclusions are solely the author’s own.

Correspondence concerning this article should be addressed to Travis Gilly, Real Safety AI Foundation, Illinois, United States. Email: [t.gilly@ai-literacy-labs.org](mailto:t.gilly@ai-literacy-labs.org).

## REFERENCES

- Aaraj, J. A., et al. (2025). *BLINDSPOT: Enabling bystander-controlled privacy signaling for camera-enabled devices* [Preprint].
- Ahmad, I., Farzan, R., Kapadia, A., & Lee, A. J. (2020). Tangible privacy: Towards user-centric sensor designs for bystander privacy. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2), 1–28.
- American Law Institute. (1965). *Restatement (Second) of Torts* § 402A.
- Baruh, L., & Popescu, M. (2017). Big data analytics and the limits of privacy self-management. *New Media & Society*, 19(4), 579–596.
- Codling v. Paglia*, 32 N.Y.2d 330, 298 N.E.2d 622 (1973).
- Elmore v. American Motors Corp.*, 70 Cal. 2d 578, 451 P.2d 84 (1969).
- Garcia v. Character Technologies, Inc.*, 785 F. Supp. 3d 1157 (M.D. Fla. 2025).
- Giberson v. Ford Motor Co.*, 504 S.W.2d 8 (Mo. 1974).
- Howes v. Hansen*, 56 Wis. 2d 247, 201 N.W.2d 825 (1972).
- Hoyle, R., Templeman, R., Armes, S., Anthony, D., Crandall, D., & Kapadia, A. (2014). Privacy behaviors of lifeloggers using wearable cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (pp. 571–582). ACM.
- Hull, G. (2015). Successful failure: What Foucault can teach us about privacy self-management in a world of Facebook and big data. *Ethics and Information Technology*, 17(2), 89–101.
- Khawaja, Y., et al. (2026). See me if you can: A multi-layer protocol for bystander privacy with consent-based restoration. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. ACM.
- Koelle, M., Wolf, K., & Boll, S. (2018). Your smart glasses’ camera bothers me! Exploring opt-in and opt-out gestures for privacy mediation. In *Proceedings of the Tenth Nordic Conference on Human-Computer Interaction* (pp. 473–481). ACM.
- Kröger, J. L., Lutz, O. H.-M., & Ullrich, S. (2021). *The myth of individual control: Mapping the limitations of privacy self-management* [Unpublished manuscript].
- Lee, K., Hong, J., Jarjue, E., Mensah, E. A., & Kacorri, H. (2020). Pedestrian detection with wearable cameras for the blind: A two-way perspective. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–12). ACM.
- Mitchell v. Miller*, 26 Conn. Super. Ct. 142, 214 A.2d 694 (1965).

- O'Hagan, J., Saeghe, P., Gugenheimer, J., Medeiros, D., Marky, K., Khamis, M., & McGill, M. (2022). Privacy-enhancing technology and everyday augmented reality. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(4), 1–35.
- Passwaters v. General Motors Corp.*, 454 F.2d 1270 (8th Cir. 1972).
- Pierce, J., Fox, S., Merrill, N., & Wong, R. (2022). Addressing adjacent actor privacy: Designing for bystanders, co-users, and surveilled subjects of smart home cameras. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference* (pp. 26–40). ACM.
- Saqib, E., Aljasim, S., Zhao, V. Y., & Das, S. (2025). Bystander privacy in smart homes: A systematic review of concerns and solutions. *ACM Transactions on Computer-Human Interaction*, 32(1), 1–33.
- Schaub, F., Balebako, R., & Cranor, L. F. (2017). Designing effective privacy notices and controls. *IEEE Internet Computing*, 21(3), 70–77.
- Sloan, R. H., & Warner, R. (2013). Beyond notice and choice: Privacy, norms, and consent. *Journal of High Technology Law*, 14, 370–412.
- Solove, D. J. (2013). Introduction: Privacy self-management and the consent dilemma. *Harvard Law Review*, 126, 1880–1903.
- Sprigman, C. J., Marotta-Wurgler, F., & Levine, D. (2024). Privacy decision-making and the effects of privacy choice architecture: Experiments toward the design of behaviorally-aware privacy regulation. *Journal of Empirical Legal Studies*, 21(4), 690–735.
- Vosburg v. Putney*, 80 Wis. 523, 50 N.W. 403 (1891).