USENIX

THE ADVANCED COMPUTING
SYSTEMS ASSOCIATION

Do You See If I See? Investigating Reciprocity in Interpersonal Access-Control Settings (in the U.S.)

Nathan Malkin, *New Jersey Institute of Technology;* Alan F. Luo, Evan J. Zhao, and Michelle L. Mazurek, *University of Maryland*

https://www.usenix.org/conference/soups2025/presentation/malkin

This paper is included in the Proceedings of the Twenty-First Symposium on Usable Privacy and Security.

August 11–12, 2025 • Seattle, WA, USA

ISBN 978-1-939133-51-9

Open access to the Proceedings of the Twenty-First Symposium on Usable Privacy and Security is sponsored by USENIX.

Do You See If I See?
Investigating Reciprocity in Interpersonal Access-Control Settings (in the U.S.)

Nathan Malkin*, Alan F. Luo†, Evan J. Zhao†, Michelle L. Mazurek†
*New Jersey Institute of Technology, †University of Maryland

Abstract

People often share information with each other, motivated by mutual benefit. However, some interfaces force reciprocity by requiring users to reveal the same type of information they want to obtain. For example, in some social networks, a user can view someone's profile only if they allow the other person to access theirs. Read receipts in many messaging apps follow the same pattern. These settings may be detrimental to privacy, since users are forced to reveal information that they may otherwise not wish to share. On the other hand, forced reciprocity may be beneficial, as it keeps interfaces simpler and enforces social norms of fairness. To understand how people perceive these trade-offs and make choices about reciprocal settings, we surveyed 802 participants from the U.S. about interpersonal access-control settings in three domains: read receipts in messaging apps, profile views in social networks, and data visibility settings in smart home devices. We found that forced reciprocity results in privacy losses, but many consider it fair, generally preferring reciprocal access-control settings to interfaces with more options. Our findings suggest that reciprocity is a potent motivator in privacy decision-making and has the potential to be useful as a mechanism in new privacy controls.

1 Introduction

"We did shrooms together, it was beautiful. We felt so open and honest that we both turned on read receipts."

– Overheard in San Francisco [48]

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.
USENIX Symposium on Usable Privacy and Security (SOUPS) 2025.
August 10–12, 2025, Seattle, WA, United States.

Reciprocity—the practice of exchanging things with others for mutual benefit—is a fundamental building block of human society [22]. However, not all reciprocal exchanges are voluntary. For example, when two parties in a car accident exchange contact information, they do so not only for mutual benefit but because they may be legally compelled to [45].

Reciprocity is also fundamental in privacy: people are often more comfortable sharing information with others on a mutual basis [44, 49] ("I'll show you mine if you show me yours"). In the privacy and access-control realm, too, reciprocal information exchange may be motivated not only by goodwill but rather by systemic requirements. For example, in Facebook's early days, seeing someone's profile required adding them as a friend; however, doing so automatically made your own profile, with all its contents, visible to them [47].

This dynamic illustrates a choice faced by system designers. They can make reciprocal sharing required, as Facebook's developers did, or they can allow each party to decide for themselves whether they want to exchange information. Both paths can be seen in different implementations of the read receipts feature in messaging applications. In many apps, including WhatsApp and Signal, users can enable read receipts to find out when their messages have been read [54, 64]. However, if a user has read receipts enabled, their messaging partners can also find out if their messages were read. In contrast, users of Apple's iMessage can choose to send read receipts, but their decision does not affect whether they will receive them, which is determined solely by the sender's settings [7].

Which approach is better for users? There are arguments for both. Requiring reciprocity can send a strong signal about social norms, put people on an equal footing, and simplify users' choices. On the other hand, doing so may entail forcing people to reveal information that they would otherwise not want to share—a potential privacy harm. Even something seemingly minor like a read receipt can communicate to a recipient that their message has been read but is being ignored, which can carry detrimental social consequences [14, 26], and may even be a violation of Contextual Integrity if it runs counter to prevailing norms [44]. But what are those norms?

At present, no guidance is available to system designers about how to best handle potentially reciprocal data sharing among users. Our research aims to shed light on this dilemma. To do so, we set out to understand people's choices and preferences regarding reciprocity, along with the norms for reciprocal interfaces, how they impact decision-making, and how they interact with reciprocity being forced. To that end, we formulated the following research questions:

- RQ1: What are the norms surrounding current real-world reciprocal access-control settings?
- RQ2: Do people make different choices when reciprocity is forced versus when it is not? If so, are these differences consistent between different types of apps?
- RQ3: When configuring potentially reciprocal settings, what reasons drive people's decisions?
- RQ4: Do people prefer optional or forced reciprocity? How do they perceive their fairness and privacy?

To answer these research questions, we conducted a survey study with 802 participants. In order for our findings about reciprocity to better generalize, we investigated participants' choices for existing and hypothetical privacy settings in three independent access-control domains: messaging (read receipts), social media (profile views), and smart homes (reviewing camera footage—a setting where reciprocity is *not* currently used). In each domain, we performed within- and between-subjects comparisons of three variants of settings: (1) forced reciprocity, (2) no reciprocity (users decide independently how to act), and (3) reciprocity is not forced but users can explicitly choose a reciprocal arrangement.

We found that people already frequently change reciprocal settings, and there are substantive disagreements in the population about whether or not reciprocity should be forced (RQ1). Participants shared more information under forced reciprocity, and we observed some variation in the choices in different domains (RQ2). Configuration decisions were driven primarily by perceived value of information and imagined audience (RQ3). Despite the potential privacy loss, most participants perceived forced reciprocity as fair, and many preferred it to settings where they had more flexibility (RQ4). Overall, we conclude that reciprocity plays a large role in users' current access-control choices and can potentially be harnessed in novel privacy controls as well.

2 Related work

Read receipts Read receipts are the most ubiquitous reciprocal privacy setting, having been around for decades. A 2003 study by Tyler and Tang [61] about their presence in email clients already mentioned that they were creating privacy concerns and frustrating users, who did not realize what they were revealing. More recent research has documented the important role read receipts have come to play in interpersonal communications. Hoyle et al. [26] surveyed Facebook Messenger users about their reactions to read receipts. They found that having a message read but not responded to creates various emotions in senders, and people use this channel for social signaling. Similar signaling behaviors, as well positive and negative emotions, were found by Chou et al. [14], who interviewed people about their attitudes and practices surrounding read receipts. Chou et al. [13] also studied ways of mitigating the tension by adding attention management features into instant messaging application, for example communicating a user's expectations about their availability. A slightly earlier attempt to reduce the social pressures created by read receipts came from Cho et al. [12], who developed prototypes of two design concepts: "private status sharing," which reveals a user's status only to those who have sent a message to them, and "sender-controlled notifications," which let senders choose whether to send notifications for their own messages. Perhaps in recognition of the complex and emotional effects of read receipts, most messaging platforms have incorporated a setting that governs sharing of read receipts.

Online status indicators Read receipts are not the only messaging feature to support reciprocity. Sun Microsystems filed a patent in 2000 for a "mechanism for reciprocal awareness of intent to initiate and end interaction among remote users" [58]. Online status indicators (OSIs) serve a similar purpose and became widespread in instant messaging. Cobb et al. [17] surveyed smartphone users about their experiences with OSIs, finding that they are frequently a source of privacy leaks, which users manage by altering their own behavior because of a lack of customizability in the OSI interface. In a separate study, Cobb et al. [16] examined OSI design decisions in 40 mobile apps, observing that they were evenly split as to whether reciprocity was enforced.

Social networks Knowledge about who is viewing you has also been shown to be important in social networking contexts, such as location sharing applications. Tsai et al. [60] found that people consider it very important to know who viewed their location, and Patil et al. [51] discovered a mismatch between users' a priori sharing decisions and ones they wanted to make in the moment. Another type of social media setting provides users with information about who has viewed their profile. Hoyle et al. [25] studied how people engage with this profile views feature on LinkedIn, finding that users have a variety of privacy concerns and behaviors, resulting in self-censorship. These concerns can sometimes be addressed through configuring access-control settings [33].

Smart homes Reciprocity plays an important role in communal usage of smart homes devices [32], for example in the way device owners and bystanders negotiate over the privacy settings of those devices [6]. Researchers have also sought to harness reciprocal relationships to manage privacy in smart homes. Akter et al. [3] experimented with reducing intra-family privacy tensions by developing a prototype app that allowed users to engage in mutual monitoring. Berger

et al. [9] evaluated novel schemes for collaborative interaction with a smart TV, allowing users to embed pre-existing interpersonal dynamics in their access control scheme.

Models and frameworks Besides specific instances of privacy controls, researchers have used more theoretical lenses to consider the relationship between privacy and reciprocity. For example, Chiu et al. [11] applied Social Cognitive Theory and Social Capital Theory to investigate knowledge sharing in virtual communities, identifying reciprocity as one of many key social capital factors that govern knowledge sharing behaviors. Stuart et al. [56] developed a theoretical framework to model people's need for social transparency in online communications. Finally, a number of studies have investigated tensions between wanting to protect and share information using game theory [28, 36, 55].

Reciprocity in the social sciences Reciprocity has also seen considerable study in the social sciences. Gouldner's original formulation of reciprocity [22] described it as a sociological phenomenon that contributes to the stability of social systems by enabling cooperation between parties, even if they hold purely self-interested motivations. Fehr and Gächter [20] applied this formulation by examining economic implications of reciprocity and documented numerous cases where it is used to enforce contracts and social norms. Molm et al. [43] showed that the value of reciprocity in a social exchange is primarily governed by its instrumental value as well as the influences it has on trust and solidarity with counterparties.

Our research builds on this related work in several ways. While prior research on read receipts has examined how users feel about them [14, 26], we quantify people's choices and examine how they make decisions for apps without established defaults. We are also able to directly compare people's decisions in different domains, such as messaging and social media [25], thus bridging a research gap. Our work further extends the smart home field, where we contribute a new type of privacy control meant to ease intra-household tensions.

3 Methods

To answer our research questions (see Section 1) and provide guidance for system developers about optional versus forced reciprocity, we developed a survey study. To understand existing norms (RQ1), we asked about respondents' use of existing apps with reciprocal access-control settings. However, our remaining research questions (about people's choices and preferences) would have too many confounds in real-world apps. To systematically investigate factors that affect reciprocity *and* collect ecologically valid data about real-world behaviors, we therefore designed a portion of our survey as an experiment, in which we asked participants to configure hypothetical apps with access-control settings that were, or were not, reciprocal.

We begin with an overview of the study flow (Section 3.1), then explain our methodological choices in greater depth (Section 3.2). Afterwards, we detail our data analysis (Section 3.3), recruitment and ethics (Section 3.4), study limitations (Section 3.5), and participant demographics (Section 3.6).

3.1 Study flow

In our survey, we asked participants to make new choices about hypothetical settings and also asked about choices they had already made in real-world settings (saving the latter for the end of the survey to minimize status quo bias). After obtaining informed consent, our study proceeded as follows.

1. We studied people's preferences for reciprocity in three domains: Messaging, Social Media, and Smart Home. These were used in between-subjects comparisons, with each participant assigned to a single domain. To control for the effect of defaults, we also randomized the setting that was reported as default to the participant (default-enabled vs. default-disabled). (Section 3.2 explains these design choices in greater detail.) Both domain and default were chosen once for each participant and then held constant for that participant for the entire study.

2. We next randomly assigned each participant to one of three reciprocity types: *Forced* (reciprocity is required), *Non-reciprocal* (no reciprocity required), or *Opt-in* (there is an explicit reciprocal option, but it is not required). Section 3.2 also explains these settings in detail.

3. We then instructed participants to role-play that they had just installed a new app with settings they had to configure, for example turning read receipts on or off. We explained each setting as a bullet-point summary and as a table describing what each party would be able to know (see Appendix 6 for complete text).

4. Participants needed to pass two comprehension check questions about the settings to proceed. Those who failed after two attempts were excluded from the study.

5. After the comprehension checks, we asked participants how they would want to configure their new hypothetical app. For example, those assigned to the default-enabled variant of the *Forced* condition in the Messaging domain answered whether they wanted to enable read receipts. These decisions—how participants configured the settings in their apps—represented the primary dependent variable in our experiment. We also asked participants open-ended questions about their reasoning for each of their choices.

6. After this, we randomly selected one of the remaining two reciprocity conditions. We told participants that we wanted them to evaluate a *new* app (in the same domain). We then repeated steps 3–5.

7. After participants had provided input on both apps, we asked which of the two apps they preferred and why.

8. We also asked questions to gauge how much value participants assigned to visibility information—both knowing it about others, and keeping it private about themselves.

9. In the next stage of the survey, we asked about the participant's current usage of reciprocal privacy settings. This varied slightly by domain, because in social media, not all services provide profile views, and in smart homes, we are not aware of existing reciprocal settings, so we asked about data review more generally. To those participants who reported using apps or services with known reciprocal settings, we asked follow-up questions, including whether they chose to enable the settings and the norms they believe to surround them.

10. We then asked respondents how they perceived forced and optional reciprocity from the perspective of fairness and privacy. We saved these questions for the end of the survey, to avoid biasing participants' choices through these considerations or usage of the term "privacy."

11. We concluded with standard demographic questions.

We piloted our survey design with a convenience sample. Participants were encouraged to think aloud as they responded to the survey; using their feedback, we clarified the descriptions of our hypothetical apps and our comprehension checks.

The complete survey instrument used in the Smart Home domain can be found via Appendix 6. The surveys for the remaining domains were substantially similar.

3.2 Experiment design

Next, we explain the rationale for our study's design choices.

Domains To ensure that our results generalized beyond a single context, we wanted to study reciprocal settings in different types of apps. We therefore selected three *domains* for our study, with only one assigned to each participant:

1. **Messaging** (read receipts)
2. **Social Media** (profile views)
3. **Smart Home** (view reports for camera footage)

Read receipts were our first pick because they are extremely widespread, due to being a core feature of messaging apps like WhatsApp and iMessage.

Another relatively common reciprocal privacy feature is profile view reports on social media. TikTok users, for example, can view a list of accounts who have visited their profile. However, if a user enables this feature, then they too will show up in a list of visitors to others' profiles [59]. LinkedIn has a similar feature [25], though its mechanics have become more complex due to the introduction of paid features [35].

Finally, we wanted to see whether any trends from domains with established reciprocal settings would hold in a different domain where such settings are not currently used: smart homes. Users in smart homes frequently encounter privacy tensions or even misuse [15, 42, 65], but privacy controls are often absent, flawed, or underutilized [24, 66], even though devices like cameras collect highly sensitive data [57]. Inspired by proposals for collaborative oversight [3] and optimistic access control [39], we conceptualized a setting, *view reports*, that would allow residents in a smart home to find out whether and when other users in their home reviewed video recordings of them. This setting allows for reciprocity in a way analogous to profile views, because notifications can flow both ways: when viewing someone or being viewed by someone.

This paper uses *visibility information* as a general term for any of read receipts, profile views, and view reports, but the survey itself used only the domain-specific terms.

Reciprocity conditions To understand people's reciprocity choices in a consistent and comparable way, we presented participants with hypothetical apps and asked them how they would configure their settings. The settings offered by these apps represented three different design choices system developers might select when it comes to sending and receiving visibility information:

1. ***Forced*** (reciprocity is required, so there is only one option to enable)
2. ***Non-reciprocal*** (reciprocity is not required, and there are two separate settings for sending and receiving)
3. ***Opt-in*** (reciprocity is not required, and the two separate settings each offer an extra option to opt into reciprocal sharing)

The first possible condition, *Forced*, required reciprocity by providing a single setting for both sending and receiving visibility information. For example, in the Messaging domain, it worked like read receipts in WhatsApp or Signal: turning it on meant that you would see read receipts, if the other party also had them turned on. A detailed explanation of this behavior, as provided to our participants, can be found—for all domains—linked from Appendix 6.

The second type of setting represented *Non-reciprocal* interfaces, as seen, for example, in Apple's iMessage, where read receipts can be received independently of sending them [7]. To capture whether people care about receiving read receipts, we also provided a separate option where people could toggle receiving them, resulting in two controls total: one for sending and the other for receiving (read receipts, profile views, etc.).

We also introduced a third type of setting, *Opt-in*, where participants could express a preference for reciprocity without being forced into it. To actualize this, we built on the *Non-reciprocal* condition; keeping the sending and receiving settings separate, we added a third option—on top of enabling or disabling—to both of them: send (or receive) **conditional** on the other party. Conditional sending meant that the app

would only send (read receipts, profile views, etc.) if the other party was also sending them. Analogously, conditional receiving implied receiving visibility information only if the other user had also chosen to receive.

When participants in our study were assigned to a particular condition, we described it as a hypothetical new app that they had installed, which had a unique name (e.g., in the Messaging domain, these were FastText, QuickMessage, and RapidMessenger). Any questions, including comparative ones, used this name only; the condition identifiers listed above appear only in the paper. We opted to give the apps different names so that we could ask comparison questions that were less confusing than "App A versus App B." While the name stayed consistent for each domain-condition pairing, to minimize confounds, we picked very similar names and made them as neutral as possible.

We wanted each participant to experience more than one condition, so they could tell us which one they preferred. We therefore designed our study to provide both within- and between-subject comparisons of the conditions: each participant experienced two out of the three conditions. (We felt that including all three might fatigue participants.)

Other controlled variables In addition to domain and reciprocity type, our study controlled for other factors that we felt might influence people's choices: ordering and defaults. To control for order effects, we randomized the order of the two different reciprocity types presented to each participant.

We similarly tried to control for the role defaults may play, since they are known to influence choices in privacy settings [2, 29] and human decision-making more generally [31]. To account for this, we designed two possible *defaults* for each setting: **default-enabled** and **default-disabled**. These differed in what was presented as the default behavior of the app in question. For example, in the *Forced* condition of the Messaging domain, the *default-disabled* version presented the choice with the question "Would you like to enable read receipts in [app name]? (The setting is disabled by default.)" The options were "Yes, enable read receipts" and "No, keep read receipts disabled." The *default-enabled* version inverted these choices. Participants were randomly assigned to one of the two defaults for the duration of the study.

3.3 Data analysis

We analyzed participants' open-ended responses using content analysis [53]. For each question, two researchers read through responses and created independent codebooks, which they then combined and refined. They then used the combined codebook to independently code each response, allowing multiple codes per response if applicable. At this stage, the average interrater reliability score, computed using the Kupper and Hafner method [34], was .64. Each response having been double-coded, the two raters then discussed and resolved any differences in their codes, achieving 100% agreement on the final ratings. In the text below, quotes from participants from the Messaging, Social Media, and Smart Home domains are designated as M*X*, S*X*, and H*X*, respectively.

3.3.1 Regression analysis

To supplement our qualitative analysis, we fit two logistic regression models to our participants' configuration decisions in the hypothetical app experiment: whether a participant chose to (1) send and (2) receive visibility information. Each regression model combined all decisions from all participants in all conditions.[1] Because every participant made two decisions (one for each of the two "apps" they considered), each contributed two data points to each regression, which we accounted for by using a model with random effects.

The regression models included the following factors (the reference category is denoted with *):

- domain (Messaging*, Social Media, or Smart Home)
- a binary variable representing whether this was the first or the second* app the participant was hypothetically configuring, to check for order effects
- default (default-enabled* vs default-disabled)
- participant's answer to how much they value obtaining visibility information [numeric score on a 5-point scale]
- participant's answer to how much they value others (not) obtaining visibility information (i.e., privacy) [numeric score on a 5-point scale]
- whether the participant reported enabling visibility information (e.g., read receipts) in a real-world app they were using, if applicable [N/A, yes, no*]
- reciprocity condition (Forced, Non-reciprocal*, Opt-in)
- participant's age, to control for demographic effects

Additionally, based on a priori hypotheses, we included two sets of interaction effects in the models. The first was between age and valuing obtaining information and valuing privacy. We reasoned that, given the social nature of reciprocal settings, different age groups might have different values [63]. The second set tested interactions between the *value* ratings and the domain. Our hypothesis was that people would value information differently across the different domains [19, 52]. More details and complete results are in Appendix 8.

3.4 Recruitment and ethics

Based on a power analysis (details in Appendix 7), we aimed to recruit at least 80 people per condition. We recruited participants between March and September 2023 through the Prolific platform [1], which enforced our requirements of being age 18 or older and from the United States. We offered $4 in compensation for the survey, which took approximately

[1] In the *Forced* condition, where there was a single combined setting, enabling the setting was interpreted as a decision to both send and receive.

15 minutes to complete. Participants provided informed consent before beginning the survey. All study procedures were approved by our Institutional Review Board.

3.5 Limitations

Our work has limitations common to human-centered research. Because participants self-reported their currently configured settings, that data could be subject to failures of recall, experimenter demand bias, or social desirability bias. However, we believe that the neutral word choices and subject matter did not predispose people towards particular choices.

Our experiment relies on hypothetical scenarios; therefore, people's actions might be different in the real world. We attempted to gauge potential behavioral differences by asking respondents, "If you were deciding about enabling [read receipts, profile views, etc.] in real life, how would you make that choice?" Participants mostly stated that they would make choices the same way as they did in our study, though some acknowledged that their choices would depend on prevailing norms and their specific use cases for the given app. The lack of such context represents another limitation of our study, since it is a factor in users' decisions. We chose to limit context to increase generalizability (avoiding overly-specific scenarios that some participants might not relate to), reduce the amount of text participants had to read, and avoid creating new conditions in an already-complex study.

Our study has other ecological limitations, including that real-world users will use interfaces substantially different from our survey-based flow and will not have access to side-by-side comparisons of privacy choices. Thus, like many studies, ours is subject to trade-offs between ecological and internal validity. We prioritized the latter, in order to enable cross-domain comparisons; we see our study as establishing the baseline for people's choices in controlled environments, to be supplemented by future work in more realistic scenarios (e.g., in situ observations of app settings) and beyond the U.S.

3.6 Participants and demographics

1,399 people started our survey and 802 completed it, after accounting for participants who quit midway through (181), failed comprehension checks (392), and manual review to identify low-quality submissions (24).[2] These were allocated to the different domains and reciprocity conditions as shown in Table 1, with at least 80 participants assigned to each condition. Table 2 lists participant demographics.

[2]Comprehension check failures varied by condition. Fewer failed out in *Forced* (3.8%) than *Non-reciprocal* (16.7%) and *Opt-in* (27.4%). This may have had a selection effect on the participants in the respective conditions, which represents a limitation of our study. While some differences between conditions are expected due to forced reciprocity's relative simplicity and familiarity, counteracting this could be possible in future work by pre-testing the comprehension questions and balancing out their difficulty.

Table 1: Number of participants in each study condition

	Messaging	Smart Home	Social
Non-reciprocal + Opt-in	120	80	81
Forced + Opt-in	80	80	80
Forced + Non-reciprocal	120	80	81

Table 2: Demographics of participants across all studies

Gender	Woman	49%
	Man	49%
	Non-binary or unknown	2%
Age	Mean (Median)	36 (33)
	Range	[18, 80]
Race/Ethnicity	White	68%
	Black or African American	12%
	Asian	6%
	Hispanic or Latino	5%
	Other	9%
Education	No college	15%
	Bachelor's or some college	69%
	Post-secondary	16%

4 Results

Below, we detail the results of our study.

4.1 RQ1: What norms surround current reciprocal access-control settings?

To learn about existing norms and expectations, we asked participants about their current usage of forced reciprocal interfaces in the domain assigned to them. We were interested in both their currently configured settings and their perspective on the expectations that surround them.

Equal numbers enable & disable read receipts Aside from SMS (used by 53% of participants in the Messaging domain), the messaging applications most popular with our participants all support read receipts: iMessage (used by 53%), Facebook Messenger (45%), Instagram (33%), Snapchat (25%), WhatsApp (21%), and Telegram (8%).

In their current most-used apps, approximately equal numbers had read receipts enabled (34%) and disabled (37%). Another 8% had them enabled only for some recipients. (The remaining respondents were unsure or used an app without read receipts.) However, current usage of read receipts varied by messenger. For example, 62% of WhatsApp users had read receipts enabled, but only 26% of iMessage users did.

We asked participants whether they actively selected the read receipts setting for their most-used app. A plurality stuck with the default read receipt setting (47%), and more than 90%

of these expressed that they would not change it. In contrast, 38% reported having changed their app's read receipt setting. (The remaining 15% said they were not sure.[3])

No strong consensus about read receipt norms If a participant's most frequently used messaging app supported read receipts (273 out of 320 respondents in the Messaging domain), we asked if they believed there were social norms related to read receipts for that app. (Complete codebooks are in Table C2, see Appendix 6.) Most frequently, participants responded that they were not aware of any norms (58 out of 273, 21.2%) or, more strongly, that there are none (59 out of 273, 21.6%): *"No, I haven't been asked to turn them on or off. I really have not heard anyone mention them"* (M320).

Many did speak about others' preferences. A common sentiment (38 out of 273, 13.9%) was that most others are likely to have their read receipts off. A smaller proportion (22 out of 273, 8.1%) expressed the opposite: *"I feel like most people just have it on and leave it on. I don't know of anyone that has turned it off"* (M176).

Interestingly, despite the findings in literature (and the present study) that read receipts create a pressure to respond [14, 26], only 25 out of 273 participants (9.2%) mentioned normative expectations of a timely response after a message is marked as read: *"Some people get mad if you don't answer them back right away"* (M190).

Profile view norms vary between networks We asked participants in the Social Media domain about which social networks they used. (Instagram was most common, 71%.) For the networks we know to employ forced reciprocity, 37% of participants used LinkedIn and 45% used TikTok.

We asked participants who used either LinkedIn or TikTok about their choices for these networks' profile views settings.[4] Approximately half of LinkedIn users (51% of 89) had profile views turned on; only 24% had them turned off. In contrast, on TikTok, more participants had them turned off (39% of 109) than turned on (25%). (The remaining respondents were unsure about their setting.) Approximately equal numbers said they changed the setting (35%) and left it with the default (37%). However, compared with the Messaging domain, a greater fraction (more than a quarter) expressed the desire to change their setting going forward.

We asked respondents who used LinkedIn or TikTok about the prevailing social expectations for profile views on those networks. (Full results are in Table C2, Appendix 6.) Once again, most commonly, participants did not know (56 out of 150, 37.3%) or thought there are none (27 out of 150, 18.0%).

As with the profile view settings, we observed differences between the platforms. A greater proportion of LinkedIn users (15 out of 71, 21.3%) mentioned that there was an expectation to have profile views enabled when compared to TikTok (1 out of 79, 1.3%). Similarly, a greater proportion of LinkedIn users mentioned that they thought that most other LinkedIn users had profile views enabled (10 out of 71, 14.1%) compared to TikTok users who expressed the same sentiment (6 out of 79, 7.6%). The differences can likely be attributed to users' goals on these platforms: "*Since [LinkedIn is] more for jobs, yes, people want to know who is checking them out there"* (S34).

[3]We consider the relatively high proportion of respondents who were unsure whether they changed a setting to be plausible. WhatsApp was first released in 2009 [46] and iMessage in 2011 [8], so the configuration could have happened years earlier.

[4]Respondents who used both were asked about one of them at random.

Table 3: How often data from cameras and other IoT devices is reviewed by owners themselves and others in their homes. (Columns may not add up to 100% due to rounding.)

	Camera		**Other IOT**	
	Self	Others	Self	Others
Daily	32%	20%	5%	4%
Weekly	35%	33%	12%	10%
Monthly	12%	13%	13%	8%
Occasionally	16%	16%	25%	19%
Never	3%	3%	34%	40%
Don't know	1%	14%	11%	19%

Camera but not other IoT owners review data We are not aware of existing smart home devices that employ forced reciprocity; therefore, in the Smart Home domain, we asked IoT owners how often they or others in their household currently review data from their devices. Since cameras are the focus of the hypothetical mechanism in our study, we asked about those separately. Participants' responses are summarized in Table 3. We found that review rates among camera owners were high, with 67% reporting that they viewed footage at least once a week and most believing others in their household did so as well. Review rates were much lower among those who owned other devices, where 34% reported never reviewing device data. This is unsurprising, as this category includes devices like smart TVs and light bulbs, which are popular, but offer no or limited data for users to review. These results support our choice to use view reports for smart cameras as our case study, since, for these products, users are already actively accessing data, necessitating the need for privacy controls.

As in the other domains, we asked participants about any norms surrounding data review. One common expectation is that data should only be reviewed when there is a good reason: "*In general, I think there needs to be some sense of 'cause' to view footage as opposed to watching out of curiosity or general interest"* (H95). For smart cameras, this expectation appeared more frequently (25 out of 91, 27.5%) than for smart home devices generally (12 out of 134, 9.0%), likely due to the more sensitive data collected by the former.

4.2 RQ2: Do people's choices differ under forced reciprocity, and how?

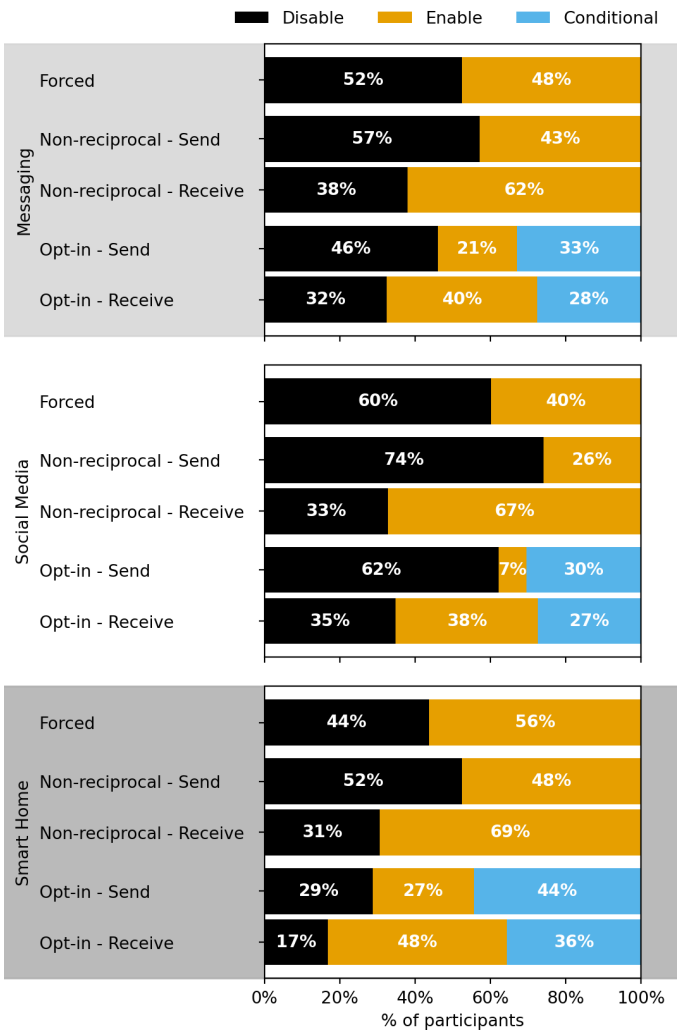

Figure 1: Settings configurations chosen by participants for hypothetical new apps

In Figure 1, we report the settings participants chose in the apps we presented to them. Overall, there was remarkable heterogeneity across all conditions: rather than one choice dominating, participants varied in their preferences.

In the *Forced* condition, where the hypothetical app's users had a binary choice between enabling and disabling, the split was almost even in the Messaging domain: 48% enabled read receipts. In comparison, more respondents enabled view reports in the Smart Home app (56%), but fewer enabled profile views in the Social Media app (40%). These differences between domains were statistically significant ($\chi^2 = 8.77, p = .01$), and we explore participants' reasons for these choices (and the differences) below, in Section 4.3.

People choose reciprocity when it is optional Our results show that, even when reciprocity is not required, and users can send or receive information unconditionally, some still want to do so contingent on the counterparty's behavior. When the Conditional option was available (in the *Opt-in* condition), participants selected it both for sending and receiving. In Messaging and Social Media, around one third chose to send visibility information conditionally (33% and 30%), and the setting was most popular in the Smart Home domain, where 44% enabled it. (These cross-domain differences were statistically significant, $\chi^2 = 7.82, p = .02$.) Conditional receiving was also most popular in the Smart Home domain (36%) relative to the others (28% and 27%), though here the distinction was not found to be significant ($\chi^2 = 3.56, p = .17$). Across all domains, more participants opted in to conditional (reciprocal) sending than enabled unilateral sending. Not only do people choose reciprocity even when it is not forced, but this results in more sharing: our regression model shows that participants in the *Opt-in* condition had the highest likelihood of sending visibility information. (See Appendix 8 for model details and coefficients.)

People are more likely to reveal information when reciprocity is forced We hypothesized that more people would reveal information when reciprocity is forced, because they may prefer to keep their own information private but are forced into giving it up under this condition. Our regression model supports this, showing that users are more likely to send visibility information in the *Forced* condition, as compared with *Non-reciprocal*. As seen in Figure 1, this trend held across all domains: when the choice to send visibility information was separated from receiving it, the number of participants who enabled it went down, from 48% to 43%, from 40% to 26%, and from 56% to 48%, in the Messaging, Social Media, and Smart Home domains, respectively. The differences between the domains themselves are less conclusive: in the regression models, the domain variable did not appear to be significant on its own. However, we observed significant interaction effects between some domains and participants' privacy valuations. While valuing privacy more reduces the odds of disclosure 4.44-fold in the baseline (Messaging) domain, this reduction is significantly less in the Smart Home domain (1.95-fold). For choosing to receive visibility information, the (non-significant) effects of privacy valuation were attenuated, from a 38.4% reduction for Messaging to a 15.6% increase for Smart Home and 1.3% increase for Social Media.

People reveal information even if they are not forced to A naive assumption about information exchange transactions might invoke rationality to predict that people would *never* send visibility information in the *Non-reciprocal* condition, since users no longer have an incentive to do this. On the contrary, from 26% to 48% (depending on the domain, which was significant, $\chi^2 = 17.97, p < .001$), enabled sending, even though it was not necessary for seeing others' information.

Examining respondents' explanations for their choices, we found several social and psychological reasons. Some stated that they would feel discomfort if they did not reciprocate: "*Knowing when others read without letting them know the same just doesn't sit well with me*" (M88). In Smart Home contexts, this extended to a desire to signal prosocial behaviors, such as trust and transparency: "*I'd rather not but if it's my family, I suppose we need the settings open so everyone feels safe and comfortable*" (H9). The absence of view reports could be used as a sign of respecting others' privacy by demonstrating that their data had not been accessed: "*I chose to keep sending enabled just so other members would know that I have not viewed recordings of them*" (H118). Finally, others felt that sending visibility information served as a stand-in for normal communication, as was sometimes the case in Social Media: "*I want them to know that I'm keeping in touch by viewing their profile without having to message them each time*" (S147).

With forced reciprocity, fewer people obtain information they otherwise want We hypothesized that, under a forced-reciprocity policy, some users might opt out (e.g., disabling read receipts), even though they are interested in receiving visibility information, because they may be unwilling to give it up about themselves. Again, our data bears this out, with the numbers for those who enabled the receiving option in the *Non-reciprocal* condition exceeding those in the *Forced* condition in all cases, moving from 48% to 62%, from 40% to 67%, and from 56% to 69%, in each of the domains. The significance of this shift is supported by the the results of our regression model for receiving visibility information, where the reciprocity condition was again significant.

Surprisingly, at least 31% of participants, across all domains, chose to disable receiving visibility information—deliberately cutting off an information source for no apparent gain. One reason for this was to avoid potential anxiety that this information might trigger: "*I don't typically like read receipts as being left on 'read' can cause some anxiety. Better to just not know*" (M28). Others felt that receiving this information was itself a violation of people's privacy: "*I would rather not be caught up in knowing if others are viewing my profile. I feel as though it is like spying on someone else's activity*" (S43). Finally, some disabled receiving because they did not want to exacerbate their notification fatigue: "*I don't need more notifications on my phone, and trust all of those in my household*" (H157).

4.3 RQ3: What reasons drive people's configuration decisions?

We asked participants why they enabled their chosen setting and analyzed their responses using qualitative coding. Complete code frequencies are in Table C1, found in Appendix 6.

Utility of information is primary factor when configuring visibility The most common reason for participants' choices (132 out of all 802 study participants, 16.5%) was finding the underlying visibility information useful: *"it is a very useful tool and I could not run my page without it"* (S217). Respondents did not always clarify how exactly the information would be useful to them, but we observed domain-specific trends. For example, in Smart Home (21 out of 240, 8.8%) and Social Media (13 out of 242, 5.4%) contexts, some of the utility derived from reciprocity was safety-related: *"Safety would be the biggest factor. Any unknown person should not be allowed to go in detail of my profile"* (S111).

These results are supported by the regressions. Those who valued being informed (agreeing with the statement "I strongly value knowing if a recipient has read my message") were more likely to enable receiving. Likewise, sending was correlated with agreeing with the statement "I strongly value my messaging partners knowing if I've read their message."

Context also mattered, for example the target audience and purpose of a social network (31 out of 242, 12.8%), the type of content (images or text) that would be shared on it (24 out of 242, 9.9%), or the placement of the smart camera (41 out of 240, 17.1%): *"what was the camera for, like for my residence or for my work because it makes a difference who it's for and then who it would be viewing"* (H123).

On the other hand, some people find that downsides of visibility information outweigh the benefits, leading them to disable settings. For instance, in the Messaging context, 31 out of 340 (9.7%) expressed concerns about having read receipts on, such as feeling pressure to respond in a timely manner (22 out of 320, 6.9%): *"I would not enable read receipts because [. . .] it adds an almost timed element to messaging that is stressful"* (M188). This pressure was less common but also present in Social Media (6 out of 242, 2.2%): "*Sometimes I just want to view someone's profile without them knowing. If they see it, it opens up a communication opportunity I don't always want to explore*" (S162).

In Smart Home settings, some (12 out of 240, 5.0%) noted that enabling forced reciprocity could actually signal distrust in other members of the household or result in discomfort for them: "*people might get upset knowing that I would be watching them*" (H53).

Counterparty's identity also drives decisions Across all domains, 134 out of 802 (16.7%) respondents cited the role of they party they were interacting with in their decisions: *"If it's a network in which a lot of strangers might view my profile, I would probably not want to know who viewed, but if it's a network that's close with only my friends, then it would be nice to know if they'd viewed my profile and other stuff"* (S6). As illustrated by this quote, the specific relationship (e.g., family members, friends, coworkers, etc.) mattered to respondents (57 out of 802 users overall, 7.1%).

Respondents also cared about which setting the people they

would be interacting with would want them to pick, with 47 out of 802 (5.9%) saying they would make a decision on that basis: *"I would talk with others and try to come up with an agreement regarding viewing/sending reports, set up some basic rules/etiquette, etc."* (H169).

Privacy considerations also influenced participants' choices. In the regression model for sending visibility information, we found that valuing privacy was associated with sending less visibility information, which conforms to intuition. Consistent with this, many participants cited "privacy" as a factor in their decision-making (102 out of 802, 12.7%): *"The main factor that would affect my decision would be my right to privacy"* (M151). This was most common in the Social Media domain (58 out of 242, 24.0%): *"I just don't want people to know how little or often I view their profile"* (S97). While Smart Home data may be more sensitive, people have stronger, more established trust relationships, making privacy less of a concern (14 of 240, 5.8%): *"Who's in the house? If just me and significant other, then it doesn't really matter"* (H91).

People also share visibility information for prosocial reasons, like to provide transparency to those they are communicating with (62 out of 802, 7.7%). This was especially salient in the Smart Home domain, where 32 out of 240 participants mentioned transparency as part of their rationale.

Other contextual factors In our regression models, we also found that users who currently have sending enabled in real-world apps were more likely to do the same in our hypothetical scenarios. For receiving, this choice was again correlated with the user's current real-life setting. Finally, we observed a small positive correlation between a person's age and their desire to receive visibility information.

While the broader literature has found conclusive evidence about the power of defaults [2, 29, 31], our model did not show defaults (default-enabled vs default-disabled) as significant in our experiment. We suspect this may be because we (purposefully) designed our survey questions so that opting in and opting out required equal effort. Therefore, participants *had* to make a choice, whereas in real life, the non-default option requires considerably more effort to enact.

4.4 RQ4: How do people compare forced and optional reciprocity?

In addition to measuring participants' choices, we asked them directly about their preferences and perceptions.

4.4.1 Do people prefer forced or optional reciprocity?

People often, but not always, prefer forced reciprocity After participants saw the two different mechanisms to which they were randomly assigned, we asked them which one they preferred. Their responses are shown in Figure 2.

When comparing the *Forced* and *Non-reciprocal* conditions, participants' preferences varied between domains

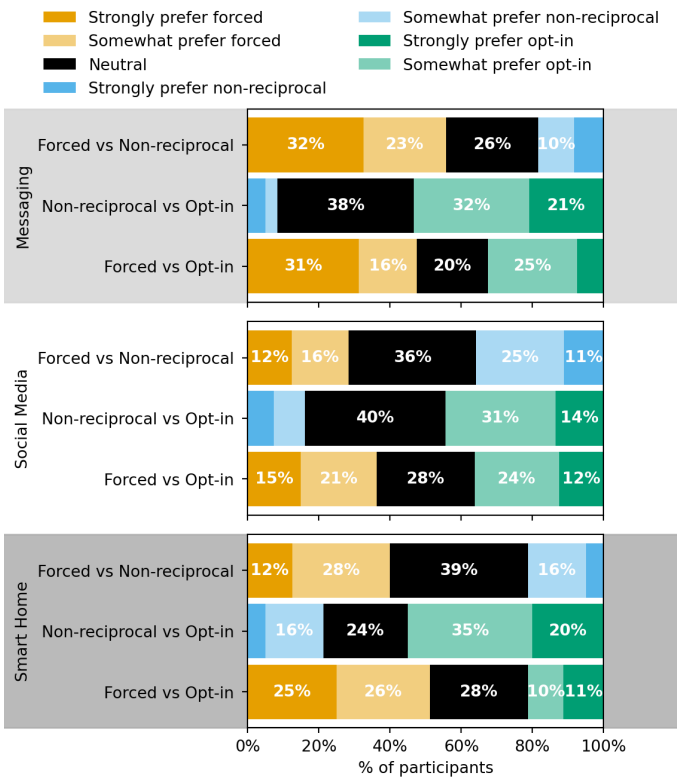

Figure 2: Participants' preferences between types of reciprocity for hypothetical new apps they experienced

(Kruskal-Wallis[5] $H(2) = 17.05, p < .001$). Approximately twice as many preferred *Forced* to *Non-reciprocal* for Smart Home and almost three times as many for Messaging. In contrast, for Social Media, *Forced* was less popular than *Non-reciprocal*. Social media was again the exception when comparing *Forced* and *Opt-in*: more participants in this domain preferred *Opt-in*, while *Forced* was more popular for Messaging and Smart Home. (Here, however, the differences between domains were less pronounced, $H(2) = 4.99, p = .08$.) The one trend that held across domains ($H(2) = 2.82, p = .24$) was that participants overwhelmingly preferred the *Opt-in* condition to the binary *Non-reciprocal* option.

We believe that the popularity of the *Forced* and *Opt-in* conditions, both of which rely on reciprocity, shows clear evidence that people find reciprocity valuable in privacy settings. Additionally, the preference for *Forced* settings is interesting because, as shown above, this mechanism restricts people's choices and produces arguably suboptimal privacy outcomes.

People prioritize simplicity, configurability, reciprocity We asked participants why they preferred their chosen reciprocity condition and analyzed their responses using qualitative coding. (See Table C3, via Appendix 6, for complete

[5] All Kruskal-Wallis tests were corrected for tied ranks.

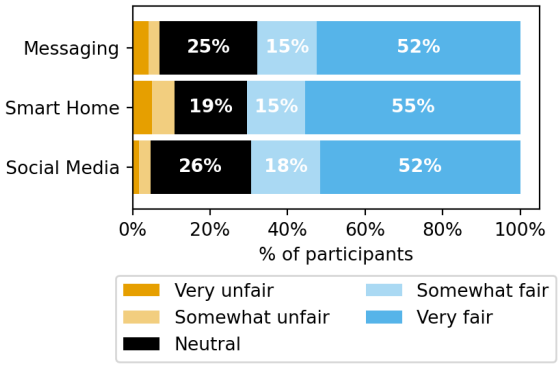

Figure 3: Participants' perceptions of the fairness of forced reciprocity across different domains

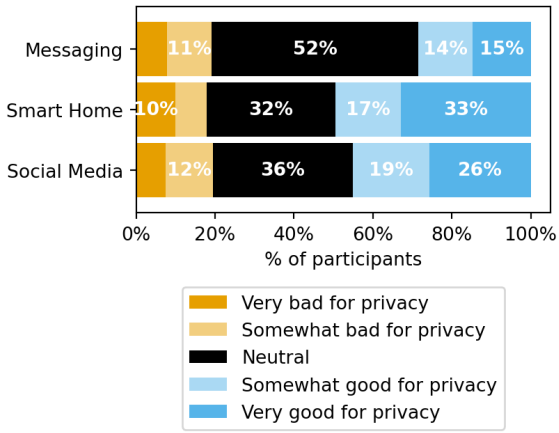

Figure 4: Participants' perceptions of the privacy of forced reciprocity across different domains

codebook.) Across all three domains, participants expressed that they preferred having fewer options to select from—193 out of all 802 participants (24.1%)—because that required less mental effort for themselves and others: *"It seems simpler to only have to enable one setting. I can imagine trying to explain the [*Non-reciprocal *version name] method to my non-technical friends, and they would have A LOT of trouble understanding and implementing it"* (M181). However, the opposite sentiment was also common (174 out of 802, 21.7%), even among those who did not necessarily plan to take advantage of the setting: *"I generally view more features as a benefit, if someone in my home ended up preferring to have conditional sending or receiving it would be nice to have that option, even if it is not very important to me"* (H81).

The opt-in setting also provoked contradictory reactions. Some, especially in the Social Media domain (47 out of 242, 19.4%), felt that it was complex and redundant: *"The conditional viewing setting is sort of a redundant setting and I can't see why anyone would click it on with the other setting being available"* (S16). But others believed the complexity was justified: *"Though it is a bit more complicated I feel it requires more of a 'mutual agreement' which I prefer"* (M55).

4.4.2 How do people perceive forced reciprocity?

Most feel that forced reciprocity is fair We asked our participants whether they thought it was fair that, in the forced reciprocity condition, sending and receiving visibility information were mutually linked. Consistently across all domains (Kruskal-Wallis $H(2) = .18, p = .91$), a majority considered this trade-off to be "very fair" (Figure 3).

Many feel that forced reciprocity helps protect privacy We also asked participants whether requiring reciprocity was good or bad for privacy. Overall, more participants thought forced reciprocity was better rather than worse for privacy (Figure 4), though these results were less asymmetric than the fairness ratings, with a sizable minority feeling that forced reciprocity was worse rather than better for privacy. There was also a difference between domains ($H(2) = 19.04, p < .001$), as fewer participants in the Messaging domain thought that forced reciprocity was better for privacy.

5 Discussion

Our study is among the first to isolate and experimentally examine reciprocity in privacy settings, which is important to understand as one widespread but understudied frontier in how people negotiate interpersonal privacy. Unlike prior work on read receipts [14, 26], ours is the first to collect data on people's preferred settings and real-world configurations, and compare these across domains. Our findings offer several practical implications for system designers.

Norms are app-specific; having choices helps Our research aims to provide guidance to system designers about how they should handle reciprocity. The most straightforward approach would be to defer to overarching social norms. Unfortunately, no consensus is evident in people's behaviors or expectations. Instead, norms appear to differ between platforms and subcommunities. For example, many more participants had read receipts turned on in WhatsApp than in iMessage. These differences are consistent with the variance in reciprocity defaults among various online status indicators [16]. New services should therefore carefully consider the expectations and precedent set by their initial choices.

The heterogeneity of people's preferences itself signals an important design implication: users benefit from choice. Some people want read receipts or profile views on, and others feel just as strongly that they should be off. Therefore services that do not provide choices risk alienating a large fraction of their users who prefer the other option. This is a lesson for apps that currently do not support disabling read receipts, like

Instagram, Facebook Messenger, and Telegram [50].

Forced reciprocity hurts privacy (but might also help) Our results demonstrate that forced reciprocity introduces a privacy loss, as users configuring access-control settings end up sharing information that they would have otherwise chosen not to. Participants also explicitly articulated privacy concerns, which date back decades [61]. In spite of this, a number of participants described forced reciprocity as being good for privacy. We hypothesize that this is because reciprocal features, especially in smart homes and social media, may deter snooping, which can be seen as a win for privacy.

People prefer reciprocity, whether or not it is forced In general, participants had positive perceptions of forced reciprocity. They emphasized its fairness and, when given a choice, typically preferred forced reciprocity to *Non-reciprocal* settings, often citing relative simplicity as more important than greater flexibility or the potential for increased privacy. Thus, forced reciprocity may be preferable to users, despite its drawbacks. This takeaway may be particularly relevant to social media, where visibility [25, 35] and other settings [37, 38] are known for being particularly confusing.

These are not the only design choices, however. Reciprocity seems even more important than simplicity: participants overwhelmingly preferred the app that offered the more complex *Opt-in* setting to the version that only had the relatively simpler binary *Non-reciprocal* option. Indeed, many participants enabled the optional reciprocity setting when it was available to them, indicating that they wanted to match their sending or receiving to the actions of their counterparty.

Non-binding reciprocity may also lead to more information being shared. In settings with *Opt-in* reciprocity, the total number of participants who would send (or receive) visibility information was typically higher than in other conditions. This suggests that system designers might maximize sending or receiving by offering non-binding reciprocal choices.

Design for virtue, not selfishness Our quantitative and qualitative results point to individuals' perceived valuations of visibility information as a critical factor when configuring potentially reciprocal settings. Yet it is not the case that obtaining this information is the sole driver of people's actions. Across domains, participants said they would send visibility information even if it would not get them new information in return. Their motivations included altruism and respect for others, but sometimes also self-interest—seeing benefits in others knowing that their message had been read. A takeaway for system designers is that they should not expect users to behave in a purely transactional manner—or think that others will act that way. Even if users are not forced to share information, they are quite likely to do so voluntarily.

We also saw a significant fraction of participants decline to receive visibility information, even though they had nothing to lose by getting it. They motivated their choice by explaining that they did not find value in this information and, in fact, found it stressful or otherwise unwelcome—consistent with the complicated feelings other research has found to be associated with read receipts [14, 26]. A takeaway is that not everyone wants to receive read receipts or other visibility information; for many, having more information is not necessarily better. One way of summarizing our observations is as empirical evidence in support of the philosophical framework of design for human flourishing, by way of virtue ethics, in which reciprocity is one of the virtues [62].

Reciprocity holds promise in new domains and settings Messaging read receipts and social media profile views represent settings that have a long history and hundreds of millions of active users. To understand whether forced reciprocity could potentially be used in new settings, we introduced a third domain to our study: view reports in smart home cameras. Despite the concept's novelty, the mechanism was well received, and more participants wanted to enable it than equivalent features in other domains. Smart homes may be a particularly promising area for new access control models because demand for greater privacy is high [10, 21, 40, 65] but uptake of existing settings is low [24, 66]. In particular, participants saw the view reports we proposed as a safety feature, suggesting that these could be the basis for differentiated access control features for household members and visitors.

Overall, how people perceived and engaged with the smart home reciprocal setting was in line with other, more established domains. These results suggest that reciprocity, whether forced or optional, can be a promising tool for new privacy settings, potentially in further domains. One advantage of forced reciprocity as a privacy mechanism is that it reduces the effort required from users to configure privacy settings. Rather than configuring access control on a per-user basis, forced reciprocity instead makes it possible to set a more general policy—reciprocity—and customize behavior as needed.

Future work should explore contexts behind reciprocity Our findings may be seen through the lens of the theory of Contextual Integrity [44]. In CI, norms about reciprocity represent the "transmission principle" that makes particular information flows acceptable in certain contexts. Because norms differ between contexts, future work should further examine how expectations about reciprocity vary in different contexts.

Another context future work could explore is group settings. In our study, reciprocal settings served to manage one-on-one relationships or those in relatively small, static groups (households in the smart home domain). People's choices may be different in one-to-many or many-to-many interactions.

More generally, notions of reciprocity are deeply interwoven into the fabric of human relationships. By continuing to study reciprocity, we can learn how to harness it to yield better privacy outcomes.

Acknowledgments

We are grateful to Phoebe Moh, Julio Poveda, and Wentao Guo for their help with piloting; and all participants, without whom this research would not have been possible. This paper results from the SPLICE research program, supported by a collaborative award from the National Science Foundation (NSF) SaTC Frontiers program under award number 1955805.

References

[1] Prolific. https://www.prolific.com/.

[2] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online. *ACM Computing Surveys*, 50(3):1–41, May 2018.

[3] Mamtaj Akter, Amy J. Godfrey, Jess Kropczynski, Heather R. Lipford, and Pamela J. Wisniewski. From Parental Control to Joint Family Oversight: Can Parents and Teens Manage Mobile Online Safety and Privacy as Equals? *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW1):1–28, March 2022.

[4] John Aldrich and Forrest Nelson. *Linear Probability, Logit, and Probit Models*. SAGE Publications, Inc., 2455 Teller Road, Thousand Oaks, California 91320 United States of America, 1984.

[5] Lujayn Alhddad. *Read Receipts Feature in Mobile Platform: An Investigation Study Based on Social Tie between the Sender and Receiver*. PhD thesis, Rochester Institute of Technology, December 2015. https://scholarworks.rit.edu/theses/8946/.

[6] Ahmed Alshehri, Eugin Pahk, Joseph Spielman, Jacob T Parker, Benjamin Gilbert, and Chuan Yue. Exploring the Negotiation Behaviors of Owners and Bystanders over Data Practices of Smart Home Devices. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–27, Hamburg Germany, April 2023. ACM.

[7] Apple. Send read receipts in Messages on Mac. https://support.apple.com/guide/messages/send-read-receipts-ichtee6857f6/mac.

[8] Apple. See new features included in iOS 5. https://web.archive.org/web/20120308031950/https://www.apple.com/ios/features.html, 2011.

[9] Melanie Berger, Rutger Verstegen, Bahareh Barati, Harm Van Essen, and Regina Bernhaupt. Collaborative TV Control: Towards Co-experience and Social Connectedness. In José Abdelnour Nocera, Marta Kristín Lárusdóttir, Helen Petrie, Antonio Piccinno, and Marco Winckler, editors, *Human-Computer Interaction – INTERACT 2023*, volume 14144, pages 369–392. Springer Nature Switzerland, Cham, 2023.

[10] Julia Bernd, Ruba Abu-Salma, and Alisa Frik. Bystanders' privacy: The perspectives of nannies on smart home surveillance. In *10th USENIX Workshop on Free and Open Communications on the Internet (FOCI 20)*. USENIX Association, August 2020. https://www.usenix.org/conference/foci20/presentation/bernd.

[11] Chao-Min Chiu, Meng-Hsiang Hsu, and Eric T.G. Wang. Understanding knowledge sharing in virtual communities: An integration of social capital and social cognitive theories. *Decision Support Systems*, 42(3):1872–1888, December 2006.

[12] Hyunsung Cho, Jinyoung Oh, Juho Kim, and Sung-Ju Lee. I Share, You Care: Private Status Sharing and Sender-Controlled Notifications in Mobile Instant Messaging. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW1):1–25, May 2020.

[13] Yu-Ling Chou, Yu-Ling Chien, Yu-Hsin Lin, Kung-Pai Lin, Faye Shih, and Yung-Ju Chang. Because I'm Restricted, 2 – 4 PM Unable to See Messages: Exploring Users' Perceptions and Likely Practices around Exposing Attention Management Use on IM Online Status. In *CHI Conference on Human Factors in Computing Systems*, pages 1–18, New Orleans LA USA, April 2022. ACM.

[14] Yu-Ling Chou, Yi-Hsiu Lin, Tzu-Yi Lin, Hsin Ying You, and Yung-Ju Chang. Why Did You/I Read but Not Reply? IM Users' Unresponded-to Read-receipt Practices and Explanations of Them. In *CHI Conference on Human Factors in Computing Systems*, pages 1–15, New Orleans LA USA, April 2022. ACM.

[15] Camille Cobb, Sruti Bhagavatula, Kalil Anderson Garrett, Alison Hoffman, Varun Rao, and Lujo Bauer. "I would have to evaluate their objections": Privacy tensions between smart home device owners and incidental users. *Proceedings on Privacy Enhancing Technologies*, 2021(4):54–75, 2021.

[16] Camille Cobb, Lucy Simko, Tadayoshi Kohno, and Alexis Hiniker. A Privacy-Focused Systematic Analysis of Online Status Indicators. *Proceedings on Privacy Enhancing Technologies*, 2020(3):384–403, 2020.

[17] Camille Cobb, Lucy Simko, Tadayoshi Kohno, and Alexis Hiniker. User Experiences with Online Status Indicators. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–12, Honolulu HI USA, April 2020. ACM.

[18] Jacob Cohen. Statistical Power Analysis. *Current Directions in Psychological Science*, 1(3):98–101, June 1992.

[19] Avinash Collis, Alex Moehring, Ananya Sen, and Alessandro Acquisti. Information Frictions and Heterogeneity in Valuations of Personal Data, November 2021.

[20] Ernst Fehr and Simon Gächter. Fairness and Retaliation: The Economics of Reciprocity. *Journal of Economic Perspectives*, 14(3):159–182, August 2000.

[21] Christine Geeng and Franziska Roesner. Who's In Control?: Interactions In Multi-User Smart Homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, pages 268:1–268:13, Glasgow, Scotland UK, 2019. ACM.

[22] Alvin W. Gouldner. The Norm of Reciprocity: A Preliminary Statement. *American Sociological Review*, 25(2):161, April 1960.

[23] Timothy M. Hagle and Glenn E. Mitchell. Goodness-of-Fit Measures for Probit and Logit. *American Journal of Political Science*, 36(3):762, August 1992.

[24] Weijia He, Maximilian Golla, Roshni Padhi, Jordan Ofek, Markus Dürmuth, Earlence Fernandes, and Blase Ur. Rethinking access control and authentication for the home internet of things (IoT). In *27th USENIX Security Symposium (USENIX Security 18)*, pages 255–272, Baltimore, MD, August 2018. USENIX Association. https://www.usenix.org/conference/usenixsecurity18/presentation/he.

[25] Roberto Hoyle, Srijita Das, Apu Kapadia, Adam J. Lee, and Kami Vaniea. Viewing the Viewers: Publishers' Desires and Viewers' Privacy Concerns in Social Networks. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pages 555–566, Portland Oregon USA, February 2017. ACM.

[26] Roberto Hoyle, Srijita Das, Apu Kapadia, Adam J. Lee, and Kami Vaniea. Was my message read?: Privacy and Signaling on Facebook Messenger. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 3838–3842, Denver Colorado USA, May 2017. ACM.

[27] F. Y. Hsieh, Daniel A. Bloch, and Michael D. Larsen. A simple method of sample size calculation for linear and logistic regression. *Statistics in Medicine*, 17(14):1623–1634, July 1998.

[28] Hongxin Hu, Gail-Joon Ahn, Ziming Zhao, and Dejun Yang. Game theoretic analysis of multiparty access control in online social networks. In *Proceedings of the 19th ACM Symposium on Access Control Models and Technologies*, pages 93–102, London Ontario Canada, June 2014. ACM.

[29] Athina Ioannou, Iis Tussyadiah, Graham Miller, Shujun Li, and Mario Weick. Privacy nudges for disclosure of personal information: A systematic literature review and meta-analysis. *PLOS ONE*, 16(8):e0256822, August 2021.

[30] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An Introduction to Statistical Learning*, volume 103 of *Springer Texts in Statistics*. Springer New York, New York, NY, 2013.

[31] Eric J. Johnson and Daniel Goldstein. Do Defaults Save Lives? *Science*, 302(5649):1338–1339, November 2003.

[32] Martin J Kraemer, Ivan Flechais, and Helena Webb. Exploring Communal Technology Use in the Home. In *Proceedings of the Halfway to the Future Symposium 2019*, pages 1–8, Nottingham United Kingdom, November 2019. ACM.

[33] Isadora Krsek, Kimi Wenzel, Sauvik Das, Jason I. Hong, and Laura Dabbish. To Self-Persuade or be Persuaded: Examining Interventions for Users' Privacy Setting Selection. In *CHI Conference on Human Factors in Computing Systems*, pages 1–17, New Orleans LA USA, April 2022. ACM.

[34] Lawrence L. Kupper and Kerry B. Hafner. On Assessing Interrater Agreement for Multiple Attribute Responses. *Biometrics*, 45(3):957, September 1989.

[35] LinkedIn. Who's viewed your profile visibility settings. https://www.linkedin.com/help/linkedin/answer/a568195.

[36] Hua Liu, Bhaskar Krishnamachari, and Murali Annavaram. Game theoretic approach to location sharing with privacy in a community-based mobile safety application. In *Proceedings of the 11th International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, pages 229–238, Vancouver British Columbia Canada, October 2008. ACM.

[37] Byron Lowens, Sean Scarnecchia, Jane Im, Tanisha Afnan, Annie Chen, Yixin Zou, and Florian Schaub.

Misalignments and Demographic Differences in Expected and Actual Privacy Settings on Facebook. *Proceedings on Privacy Enhancing Technologies*, 2025. https://petsymposium.org/popets/2025/popets-2025-0025.php.

[38] Michelle Madejski, Maritza Johnson, and Steven M. Bellovin. A study of privacy settings errors in an online social network. In *2012 IEEE International Conference on Pervasive Computing and Communications Workshops*, pages 340–345, March 2012.

[39] Nathan Malkin, Alan F. Luo, Julio Poveda, and Michelle L. Mazurek. Optimistic Access Control for the Smart Home. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 3043–3060, May 2023.

[40] Karola Marky, Alexandra Voit, Alina Stöver, Kai Kunze, Svenja Schröder, and Max Mühlhäuser. "I don't know how to protect myself": Understanding privacy perceptions resulting from the presence of bystanders in smart environments. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*, NordiCHI '20, New York, NY, USA, 2020. Association for Computing Machinery.

[41] Scott Menard. *Applied Logistic Regression Analysis*. SAGE Publications, Inc., 2455 Teller Road, Thousand Oaks California 91320 United States of America, 2002.

[42] Phoebe Moh, Pubali Datta, Noel Warford, Adam Bates, Nathan Malkin, and Michelle L. Mazurek. Characterizing Everyday Misuse of Smart Home Devices. In *2023 IEEE Symposium on Security and Privacy (SP)*, San Francisco, CA, USA, May 2023. IEEE. https://doi.ieeecomputersociety.org/10.1109/SP46215.2023.00089.

[43] Linda D. Molm, David R. Schaefer, and Jessica L. Collett. The Value of Reciprocity. *Social Psychology Quarterly*, 70(2):199–217, June 2007.

[44] Helen Nissenbaum. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford Law Books, Stanford, Calif, 2009.

[45] Ohio. Ohio Revised Code. https://codes.ohio.gov/ohio-revised-code/section-4549.02, September 2016.

[46] Parmy Olson. The Rags-To-Riches Tale Of How Jan Koum Built WhatsApp Into Facebook's New $19 Billion Baby. *Forbes*, February 2014. https://www.forbes.com/sites/parmyolson/2014/02/19/exclusive-inside-story-how-jan-koum-built-whatsapp-into-facebooks-new-19-billion-baby/.

[47] Kurt Opsahl. Facebook's Eroding Privacy Policy: A Timeline. https://www.eff.org/deeplinks/2010/04/facebook-timeline, April 2010.

[48] Overheard in San Francisco. https://www.instagram.com/p/ClrVJgPv4Qn/, December 2022.

[49] Peiyu Pai and Hsien-Tung Tsai. Reciprocity norms and information-sharing behavior in online consumption communities: An empirical investigation of antecedents and moderators. *Information & Management*, 53(1):38–52, January 2016.

[50] Khamosh Pathak. How to Turn Off 'Read Receipts' on Most Popular Messaging Apps. https://lifehacker.com/how-to-turn-off-read-receipts-on-most-popular-messagi-1848438633, January 2022.

[51] Sameer Patil, Roman Schlegel, Apu Kapadia, and Adam J. Lee. Reflection or action?: How feedback and control affect location sharing decisions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 101–110, Toronto, Ontario, Canada, April 2014. ACM.

[52] Jeffrey T. Prince and Scott Wallsten. How much is privacy worth around the world and across platforms? *Journal of Economics & Management Strategy*, 31(4):841–861, 2022.

[53] Johnny Saldaña. *The Coding Manual for Qualitative Researchers*. SAGE Publications, 2021.

[54] Signal. Read Receipts. https://support.signal.org/hc/en-us/articles/360007059812-Read-Receipts.

[55] Anna C. Squicciarini, Mohamed Shehab, and Joshua Wede. Privacy policies for shared content in social network sites. *The VLDB Journal*, 19(6):777–796, December 2010.

[56] H. Colleen Stuart, Laura Dabbish, Sara Kiesler, Peter Kinnaird, and Ruogu Kang. Social transparency in networked information exchange: A theoretical framework. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work - CSCW '12*, page 451, Seattle, Washington, USA, 2012. ACM Press.

[57] Neilly H. Tan, Richmond Y. Wong, Audrey Desjardins, Sean A. Munson, and James Pierce. Monitoring Pets, Deterring Intruders, and Casually Spying on Neighbors: Everyday Uses of Smart Home Cameras. In *CHI Conference on Human Factors in Computing Systems*, pages 1–25, New Orleans, LA, USA, April 2022. ACM.

[58] John C. Tang, Nicole Y. Mordecai, James M. A. Begole, Janak R. Bhalodia, and Max G. Van Kleek. Mechanism for reciprocal awareness of intent to initiate and end interaction among remote users, May 2004. https://patents.google.com/patent/US6731308B1/en.

[59] TikTok. Profile view history. https://support.tiktok.com/en/account-and-privacy/account-privacy-settings/tiktok-profile-visit-history.

[60] Janice Y. Tsai, Patrick Kelley, Paul Drielsma, Lorrie Faith Cranor, Jason Hong, and Norman Sadeh. Who's viewed you?: The impact of feedback in a mobile location-sharing application. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2003–2012, Boston, MA, USA, April 2009. ACM.

[61] Joshua R. Tyler and John C. Tang. When can I expect an email response? A study of rhythms in email usage. In Kari Kuutti, Eija Helena Karsten, Geraldine Fitzpatrick, Paul Dourish, and Kjeld Schmidt, editors, *ECSCW 2003*, pages 239–258, Dordrecht, 2003. Springer Netherlands.

[62] Shannon Vallor. Social networking technology and the virtues. *Ethics and Information Technology*, 12(2):157–170, June 2010.

[63] Miranda Wei, Jaron Mink, Yael Eiger, Tadayoshi Kohno, Elissa M. Redmiles, and Franziska Roesner. SoK (or SoLK?): On the Quantitative Study of Sociodemographic Factors and Computer Security Behaviors. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 7011–7030, 2024. https://www.usenix.org/conference/usenixsecurity24/presentation/wei-miranda-solk.

[64] WhatsApp. How to change your privacy settings. https://faq.whatsapp.com/3307102709559968.

[65] Yaxing Yao, Justin Reed Basdeo, Oriana Rosata Mcdonough, and Yang Wang. Privacy perceptions and designs of bystanders in smart homes. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW), November 2019.

[66] Eric Zeng and Franziska Roesner. Understanding and improving security and privacy in multi-user smart homes: A design exploration and in-home user study. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 159–176, Santa Clara, CA, August 2019. USENIX Association. https://www.usenix.org/conference/usenixsecurity19/presentation/zeng.

Appendices

6 Supplementary Materials

A sample survey instrument, complete explanations for all conditions, and codebooks from qualitative coding can be found at: https://osf.io/e574g/

7 Power analysis

To determine the number of participants in our study, we conducted a power analysis [18, 27]. We performed an *a priori* one-tailed z-test for a logistic regression with standard values of .05 for confidence and .8 for power, and estimated the R^2 *other X* (amount of variability in the main predictor that is accounted for by covariates) to be low association, using a standard value of .04. We estimated the probability that a user is willing to share information if forced reciprocity has no effect on sharing behaviors (H_0) using .45 as a lower bound and .65 as an upper bound, based on the results of a survey performed by Alhddad [5] measuring read receipt behaviors. We then estimated the probability that a user is more willing to share information if forced reciprocity does have an effect on sharing behaviors using .55 as a lower bound, and .75 as an upper bound—this is because we estimated our effect size would be small-to-medium. We found that the number of people per condition we would need to be 28–174, depending on the effect size. Based on an estimate of medium effect sizes, prior experience, and budgetary considerations, we targeted at least 80 people per condition.

8 Regression analysis

To quantitatively analyze how participants configured the hypothetical apps in our experiment, we created two logistic regression models, one for sending and one for receiving visibility information, as described in §3.3.1.

To check for multicollinearity, we verified that all Variance Inflation Factors (VIF) were below 5, the level considered concerning [30, 41]. Because the VIF for the "value" variables were the highest we observed (3.39, 3.67), we further tested models that omitted one of these variables and found that they performed worse on BIC, AIC and Aldrich-Nelson pseudo-R^2, while yielding nearly identical results. As a result, we included all variables mentioned in §3.3.1 in the final models.

In addition to the factors and interaction effects mentioned there, we also tested several additional interaction effects: between reciprocity condition and domain (to explore the effects of context), between reciprocity condition and defaults (since default effects could be stronger, for example, in a novel domain like smart homes), and between reciprocity condition and choice order (to test for potential asymmetric order effects). However, all of these models had a higher

BIC, suggesting a worse fit with the data, so we omit these interaction effects from the final model.

We computed Aldrich-Nelson pseudo-R^2 [4] due to its superior characteristics [23]. The pseudo-R^2 values for the sending and receiving models were .48 and .51, respectively, suggesting that our factors explain a relatively high amount of variance in the dependent variables.

The regression results for sending and receiving visibility information appear below. Asterisks indicate significance at the .05, .01, and .001 level, respectively.

	Sending				**Receiving**			
	Odds Ratio	Std. Error	z value	p-value	Odds Ratio	Std. Error	z value	p-value
Domain – Smart Home	0.704	0.463	-0.757	0.449	0.625	0.398	-1.181	0.238
Domain – Social Media	0.434	0.434	-1.925	0.054	0.648	0.369	-1.176	0.239
Choice order	0.770	0.164	-1.589	0.112	1.212	0.160	1.197	0.231
Choice framing	0.722	0.246	-1.329	0.184	0.723	0.215	-1.506	0.132
Value being informed	1.511	0.381	1.084	0.278	4.189	0.335	4.278	<.001***
Value privacy	0.225	0.398	-3.748	<.001***	0.616	0.330	-1.466	0.143
Reciprocal setting in real-world app – N/A	2.698	0.438	2.267	0.023*	2.339	0.337	2.526	0.012*
Reciprocal setting in real-world app – Enabled	6.255	0.449	4.081	<.001***	4.655	0.367	4.191	<.001***
Reciprocity condition – Opt-in	4.257	0.219	6.610	<.001***	1.806	0.217	2.721	0.007**
Reciprocity condition – Forced	2.817	0.216	4.800	<.001***	0.299	0.200	-6.029	<.001***
Age	1.022	0.012	1.923	0.055	1.021	0.010	2.148	0.032*
Value info + Age	0.996	0.008	-0.445	0.656	0.989	0.008	-1.374	0.169
Value privacy + Age	0.994	0.009	-0.735	0.462	0.990	0.008	-1.323	0.186
Value info + Smart Home	1.270	0.287	0.833	0.405	1.253	0.258	0.875	0.382
Value info + Social Media	1.607	0.320	1.480	0.139	1.644	0.269	1.848	0.065
Value privacy + Smart Home	2.281	0.292	2.824	0.005**	1.877	0.249	2.535	0.011*
Value privacy + Social Media	1.640	0.310	1.598	0.110	1.645	0.253	1.970	0.049*