

SoK: Social Cybersecurity

Yuxi Wu, W. Keith Edwards, Sauvik Das

Georgia Institute of Technology

yuxiwu@gatech.edu, keith@cc.gatech.edu, sauvik@gatech.edu

Abstract—We analyze prior work in social cybersecurity and present a structuring of this literature based on its pertinence to four S&P-relevant social behaviors: (1) negotiating access to shared resources, (2) shared and social authentication, (3) managing self-presentation, and (4) influencing others’ S&P behaviors. We further break down these domains into four scales of social distance—intimate, personal, social, and public—showing that desired access control policies, authentication methods, and privacy and sharing preferences vary across these social scales. We evaluate the current landscape of work through the lens of Ackerman’s social-technical gap in social computing systems, finding that while social behaviors clearly impact S&P preferences and needs, existing S&P systems are designed with little understanding of these behaviors. This mismatch forces users to choose between implementing their ideal S&P policies or reducing social friction. To address this mismatch, we outline a research agenda for social cybersecurity work that better aligns S&P goals with social needs, preferences and behaviors.

I. INTRODUCTION

Many end-user cybersecurity and privacy (S&P) behaviors are inherently social: we share information with other people in our social networks [1], we ask questions of friends and family about best S&P practices when confused [2], [3], and we coordinate with and help others to be safe online [4]. Indeed, work in usable S&P has alluded to the existence of social influences in S&P behaviors as early as the late 1990s [5]. Yet, most tools aimed at helping end-users improve their security and privacy have been designed primarily with individual behaviors in mind. As examples, there has been a vast amount of work focused on improving the usability of individual authentication systems [6] and access control policy interfaces [7] [8], or increasing individual users’ comprehension of S&P warnings [9].

Due to the increasing interconnectedness of people on the Internet, there has been a growing interest in studying end-user S&P beyond the individual actor [10]. These emergent perspectives—variously called social cybersecurity [11] [12], socio-technical cybersecurity [13] [14], community oversight [15] [16], or networked privacy [17]—share a common high-level goal: to understand S&P behaviors and threats in an ecosystem of interconnectedness and influence. Social approaches to S&P vary in the scales of populations they consider and how they orient the experience of the individuals within these populations: from high-level considerations of how influence, (mis)information, and other threats propagate through large social networks, to deeply personalized investigations of the S&P considerations of families, couples, and households. Across these scales, we see not only different

technical approaches to S&P, but also different definitions and theoretical underpinnings across disparate research literature.

In this paper, we synthesize insights from the broad literature on social cybersecurity, highlighting gaps and proposing areas for future exploration. We identify four key behavior domains in the social cybersecurity literature: negotiating access to shared resources, shared and social authentication, managing self-presentation, and influencing others’ S&P behaviors. Within these domains, we categorize specific behaviors into four distinct scales of social organization described in anthropology literature [18]: intimate (e.g., romantic partners), personal (e.g., families, households, and close friends), social (e.g., acquaintances, social media friends, coworkers) and the public (e.g., strangers, advertisers, other institutions).

Much of the literature we explore in this paper covers empirical studies that have shown how human social dynamics push against and complicate the use of technical tools intended to support S&P practices. These studies illustrate the ways in which users must adapt their social practices to fit the affordances of these tools, or, alternately, how users are driven to reappropriate existing technology, using it in new or “unsanctioned” ways, in order to better support their desired social behaviors—sometimes reducing S&P against the threats those technologies are designed to thwart. This distinction between the needs of social groups, and the support provided to them by the technology, has been termed by Ackerman the *social-technical gap*: the tension between what is achievable with existing technology, and what is socially-required by the users of that technology [19].

Many of the challenges of social cybersecurity arise due to this gap: S&P burdens are foisted upon users who must improvise their way through social situations—like sharing digital resources, authenticating social group membership, controlling self-representation online, and helping others with S&P problems—with tools that have been designed with limited understanding of social behavior. But the literature also includes attempts at creating *new* technical tools designed with the social practices of their users in mind, and intended to support or leverage the existing social practices of social groups. Accordingly, we also suggest a few directions for future social cybersecurity work.

II. METHODOLOGY AND SCOPING

The fundamental contribution of our work is a taxonomization and synthesis of existing work on social cybersecurity. Our methodology spanned three phases: gathering relevant

prior work, identifying common themes among these papers, and grouping themes into taxonomically significant domains.

We began our search for existing social cybersecurity work using two index terms on the ACM Digital Library: “human and societal aspects of security and privacy” and “social aspects of security and privacy”. These index terms are typically self-selected by authors to represent their work; we felt that if authors believed that there was a significant social S&P component in their work, they would have self-identified it as such. We supplemented these index terms with known keywords from existing literature, such as “social cybersecurity”, “collaboration”, “community”, “privacy”, and “security”. Quickly realizing that the literature covered disparate social scales and contexts, we expanded searches on those contexts with new keywords as they arose (examples include “couples”, “intimate partner violence”, “family”, “households”, “teenagers”, “social networks”, and “workplace”).

We initially scoped collection to prior work from the last five years of CCS, CHI, CSCW, IEEE S&P, NDSS, PETS, SOUPS, TheWebConf, and USENIX Security. Older work and papers from other venues, e.g., NSPW and UbiComp, that was highly cited within this list was also included. These searches resulted in about 1000 articles, many of which fell under the umbrella of usable security and privacy. We pared this down to approximately 100 by including only works that included some element of *social coordination or cooperation* with end-user S&P behaviors, asking questions like, “does this work advance our knowledge of how social groups jointly navigate S&P decisions, behaviors, threats or tools?”, and “does this information advance our knowledge of how S&P threats, tools, or advice affects groups differently than individuals?”.¹ Recent work by Carley et al. [21], termed “social cyber-security”, was also excluded since we were focused on work that centers end-users, rather than states, as actors.

From each of the papers gathered, we identified the core research question(s), methodology, analysis, and results; research context; targeted stakeholders; and authors’ recommendations for future work. This data was extracted into a spreadsheet: one column per characteristic for each paper row, similar to prior SoKs (e.g., [22]). We then applied a reflexive approach to thematic analysis [23]. One researcher performed the initial coding, updating the codebook as new codes and categories emerged. We included the codebook in the appendix. While we did not formally perform axial coding, two additional researchers participated in iterative discussions to organize codes and generate four taxonomically-significant domains, described in more detail in the next section. These domains are the basis for how we structured and systematized the extant work on social cybersecurity.

III. SYSTEMATIZATION

In our coding process, we found that prior work in social cybersecurity tended to focus on specific user S&P behaviors

¹For example, work on sharing passwords with others qualifies as social cybersecurity, since it involves an element of negotiation and trust in others to protect secrets, but work on general password usage [20] does not).

that were either enabled or complicated by social interaction. We identified four broad behavior domains commonly investigated across the literature:

- **Negotiating access to shared resources** includes sharing media accounts, devices, work files, physical access to workspaces, carpools, cars, and homes. These use cases require trust between sharers, which in turn requires negotiation of mutual S&P practices. Along with a dearth of more socially-nuanced access control systems, this section also includes many password-sharing behaviors, which can function as coarse-grained access control for small groups (more on this in Section IV).
- **Shared and social authentication** includes user reliance on others to help them authenticate, as well as systems that facilitate group authentication.
- **Managing self-presentation** includes disclosing personal information to build trust, self-censorship, audience selection tools on social media, context collapse, curating public profiles, and reckoning with institutional profiling.
- **Influencing others’ S&P behaviors** includes consulting close friends and family for advice, sharing stories about S&P events with others, nudges with social proof to influence S&P behaviors, and cross-cultural considerations for appropriate S&P advice and support.

We further segmented the social S&P these behavioral domains into different interactive social distances, i.e., how frequently and intensely people interact with each other [24]. Work predating the modern computing era suggests that, across social distances, users have different levels of disclosure and privacy preferences [25]. Across intimacy levels, there are distinct types of interpersonal relationships, jointly construed threat models, resources shared, social norms, collaborative capabilities, desired access control policies, and strategies for securing shared digital resources.

To select pertinent levels of social distance, we loosely adapted the four “interpersonal distances of man” proposed by Hall in his formulation of proxemics theory [18], or how people non-verbally communicate their level of intimacy with each other by arranging themselves in a physical environment. Hall describes four different scales of social interaction that are revealed through physical proximity: the intimate space (romantic partners), the personal space (like families and households), the social space (e.g., acquaintances, coworkers), and the public space. This seminal innovation continues to influence work in anthropology [26] and communication theory [27]; it is a canonical reference for work on the relationship between social relations and interpersonal distances, even within computing, where these scales extend beyond physical proximity [28], [29]. For S&P behaviors specifically, as we unpack in the sections to come, the ways romantic partners, families and households, social acquaintances, and the public share resources, credentials, advice, and about themselves with each other are all distinct from one other.

IV. NEGOTIATING ACCESS TO SHARED RESOURCES

In this section, we consider the sharing of digital services (Netflix, banking, work accounts), physical devices, and physical environments. Since many existing solutions for securing shared resources poorly map onto ideal social access control policies, we also discuss situations in which direct account / password sharing is used in place of formal access control as a means of sharing access to a protected digital resource. We also cover work that presents design implications or entirely new technologies focused around supporting socially nuanced access control for jointly owned and shared digital resources.

A. *Intimate relationships*

Intimate partners commonly share a variety of digital devices, accounts and resources, ranging from bank accounts, media subscription services, calendars, and smart devices. However, prior work suggests that existing models for sharing and access control can often complicate these behaviors, resulting in social friction, breaches of desired access control policies, or cumbersome practices to work around technologies designed for single-person use.

1) **Content, device, and resource sharing:** A number of formative studies have examined how social practices impact sharing among intimate partners. Jacobs et al. [30] identified four patterns in the sharing behaviors of couples. One, intentional sharing, usually occurs for practical reasons, e.g., a shared calendar or saving money on subscriptions to media accounts. The second, an explicit lack of sharing, occurs around personal laptops, search histories, and bank accounts, usually motivated as a way for users to maintain personal privacy and identity.² The third, unintentional access, occurs when participants sometimes accidentally see messages on co-located partners' devices. Finally, the fourth, unintentional inhibiting of access, occurs when one partner forgets to share a password with the other.

Alternately, Park et al. [32] presents a temporal analysis of sharing and security behaviors across the phases of romantic relationships. In the beginning of a relationship, partners are uncertain about sharing, but as trust grows throughout a relationship, partners share more with each other, including passwords and devices. At a relationship's end, however, revocation or disabling of access to shared resources might grow to be non-trivially tedious. Lin et al. [33] add an environmental dimension to this timeline, finding that specifically defined relationship statuses and environments (e.g., relationship formation, cohabitation, engagement/marriage) exhibit sharing behaviors with progressive levels of risk in exposing personal information. For example, before cohabitation, couples share entertainment subscriptions early on. During cohabitation, couples start sharing utilities and shopping accounts, and, through high levels of face-to-face interactions, trust grows (conversely, couples in long-distance relationships are less

²As a counterpoint, more recent work on smart speaker use has found that some users who share devices with their partners feel that they had nothing to hide, and that checking up on one another's information via the devices is par for the course in their relationships [31].

willing to share resources). Finally, the authors found that engagement or marriage was associated with sharing the most personal accounts such as finance and insurance.

On the other hand, when a relationship ends, account sharing does not terminate so cleanly. Obada-Obieh [34] identified abundant psychosocial burdens related to S&P when users end online account sharing with others: uncertainty about whether the sharing had actually stopped, annoyance at having to migrate to a new account (including either possibly losing personal content or not being able to delete all content), and the risk of being hijacked by a sharee. And, to avoid awkward conversation with sharees, users often fail at ending account sharing even when the desire for such sharing had long passed.

2) **Passwords:** There is often a necessary conflation between sharing content, and sharing the authentication information that facilitates access to the content. Even though passwords and PINs generally come with the advice to "never be shared", partners do so anyway. Couples frequently share devices, media, and finances, via password sharing, creating joint accounts, and leaving devices unlocked on purpose [30]. Couples in Australia, for example, frequently share banking PINs and passwords, even though banking systems ask their customers not to do so: married couples often have joint accounts, for which there are only individual passwords [35]. Still, when ending account sharing, users often struggle with password-specific issues: having to remember all the shared users of an account, changing passwords, and remembering which passwords were reused across accounts [34].

In other words, end users appropriate authentication methods like usernames and passwords as a simple, socially frictionless form of access control. Some digital platforms directly build this into their designs: Netflix accounts, among other streaming services, are tied to an individual username and password, but allow for multiple viewers and profiles.

3) **Intimate partner violence:** Not all intimate partners have equal control over S&P policies and practices; indeed, there are unfortunate cases in which one partner abuses S&P configurations to control or stalk the other. Survivors of intimate partner violence (IPV) often share cell phone plans with their abusers, which allows abusers to track devices, activate/deactivate services, and view account activity [36].

The conflation of trust with password sharing can also exacerbate IPV situations. Due to the intimacy and in-person nature of relationships characterized by IPV, abusers can easily access the survivor's devices and accounts. This access is not always seized by brute force: when the relationship is still positive, survivors often willingly share passwords and devices, along with information that could answer account security questions, with their abusers because they trust them [36]. To better organize these social-technical complexities, Matthews et al. [37] describe a three-phase framework for organizing how survivors of IPV make S&P decisions. Initially, survivors must cope with regular physical access by the abuser, who may monitor their devices and accounts or install spyware on their devices; some survivors used alternative devices or accounts in response. When trying to escape, some survivors deactivated

accounts and destroyed their devices to hide themselves from their abusers, making a trade-off between online privacy and access to social support. Once they left, survivors monitored and restricted the activities of their kids and other social networks, blocking contacts if they might threaten their safety.

B. Families and households

Families and households exhibit complex practices around sharing. Past work has explored the implications of both *intended* and *unintended* sharing within the household. Regarding intended sharing, Mazurek et al. [38] found that across diverse household types (couples, roommates, families), users have distinct ideas of what they classified as sensitive, and used a variety of access control mechanisms to combat violations of this sensitivity. However, users often iteratively adapt their ideal access control policies rather than settle on the initial ones they reported, suggesting that *a priori* access control policies may be insufficient within the household context [38].

Generally, device sharing is common in households, even for personal devices like mobile phones. Household device sharing is often characterized by a trust in sharees, i.e., not requiring supervision over the sharee’s usage of a device or account; if a trusted person broke expectations, they would become less trusted and would have less access [39]. Convenience also influenced security behaviors; devices within closer proximity tended to be shared more. Device type and location within the home also influence the degree of sharing within the household: families with more laptops tended toward more individual usage and ownership, while desktop computers, often placed in common areas, were shared across the household [40]. And, as with romantic partners, households frequently share passwords for paid resource accounts to save money. Even users who typically do not share other passwords still share wi-fi passwords with home visitors [41].

Some technical work at this scale has explored how to better support these social practices around intended sharing. For example, Family Accounts, a model for user accounts on shared home computers, proposed making documents and settings shared by default, but allowing individual profiles to personalize settings by making specific folders and documents private to just themselves [42].

The consequences of *unintended* sharing afforded by shared devices are also well-documented. Users of smart speaker devices, often shared within a household, tend to have imperfect understandings of who else in the household can access their data [31]. Even though they are concerned about unauthorized access to personal information via the devices, users’ mitigation strategies often either prevent them from making full use of the speaker’s capabilities or resign themselves to exposing their private information. These problems are exacerbated by the fact that they tend to be purchased and associated with one family member, but shared by the entire home [43]. Unintended access scenarios range diversely from malicious family member access to inadvertent guest access. To address this, Zeng and Roesner [44] designed and deployed a smart home user interface in an in-home user study, finding that

deviations from household social norms might be more to blame than simple inflexibility in access control. For one, an increase in flexibility also means an increase in UI complexity; users also relied more on existing household norms to quell S&P conflicts than directly interacting with the UI controls.

C. Social acquaintances

This scale includes sharing resources within social friend groups, but is primarily focused on workplace sharing. Within social groups, prior work has identified social sharing of resources in a variety of settings, such as digital media accounts, physical items, computing devices, and group messages [45].

Members of social groups individually employ their *personal* S&P practices to combat insider and outsider threats, trusting in other group members to protect shared resources; however, they are also frustrated with the inefficient patchwork nature of such a strategy to represent collective S&P preferences, especially given how diverse the aforementioned resources are [45]. Moju-Igbene et al. [46] try to bridge this inefficiency by engaging groups of users in participatory design workshops to envision social solutions to the access and S&P problems they face as groups. Through these “design jams”, the authors and groups recommend four design dimensions to consider in future work: social transparency, or the ability for the group to keep track of individual behaviors; stakes and responsibility, or the distribution of responsibility and ownership in the S&P of the shared resources; structures of governance, i.e., the collective decision-making process; and motivating pro-group S&P behaviors.

Beyond general social sharing of resources, we specifically identified two domains in the *workplace*: sharing of digital resources like files, mailing lists, and enterprise services; and access control for physical environments.

1) **Workplace digital resource sharing**: Once again, password sharing is utilized as an easy workaround for users to share access to resources needed for their work; the role of passwords in these situations is similar to access sharing behaviors found in intimate relationships and households. For example, colleagues frequently share passwords for shared work resources like mailing lists and digital access to journals [41]. Users share passwords with coworkers in a number of ways: by telling others directly, through email or Slack; writing on a board or post-it note shared in common work areas; using enterprise password managers; or logging into their own accounts for others to use [47]. This sharing not only makes it easier for coworkers to share official work information in a secured environment, but also helps workplaces save money by having fewer accounts for paid services. The practical considerations of sharing access costs blurs the line between common authentication methods and access control.

Some work has focused specifically on novel technologies for workplace sharing, and presented new approaches intended to support more flexible work resource sharing. For example, Volda et al. [48] explored a tool called the Sharing Palette, which attempted to provide the simplicity of email, coupled with greater visibility and control over the sharing process.

2) *Access control for physical environments*: Another group of work at this scale has focused on physical access control policies, and in particular, the ways that static policies may break down with social groups. For example, administrators in an office environment who manage access control policies for a physical space (via swipe cards and physical keys) have several requirements not addressed by currently technology: (1) policies being made by multiple people, (2) policy makers being different from policy implementers, and (3) access control systems being unable to implement desired policies [8]. Bauer et al. [7] compared how well physical keys achieved users' ideal access control policies in an office environment compared to a distributed smartphone system. The smartphone system was able to meet logging, notification, and real-time approval conditions desired by users (which are not supported by physical keys). As such, resource owners have more control and flexibility over access control.

Logas et al. [49] also found that when administrators of makerspaces construct static access control policies for the spaces, they are constrained by four dynamic factors: safety, logistics, experience, and funding. Even though these policies are static, administrators often make exceptions in contextual and social situations, e.g., to keep up with high demand, to maintain good relationships with staff, and to build trust with end-users. Despite these exceptions, end-users were often still frustrated with the static nature of the policies.

D. The public

At the public scale, there is little existing S&P literature focused on negotiating access to shared resources; however, there are a few case studies on sharing resources with strangers via P2P car-, ride-, and accommodation-sharing apps.

Radke et al. [50] investigated collective information sharing in ridesharing contexts, where users are generally strangers to one another, finding that while precise location is necessary for pickups and drop-offs, riders were concerned about being tracked. Riders simultaneously want to provide their personal information to those who would be in the same car as them, but also want to know more about their fellow ridesharers; similarly, riders want both to obscure their personal information in rideshare databases and to maximize transparency about their use. To tread this fine line, Pham et al. [51] built PrivateRide, which anonymizes riders and protects their location data while preserving ride matching accuracy. He et al. [52] go one step further, attempting to protect riders' privacy from their fellow riders as well (e.g., in a Lyft Line or Uber Pool).

There are similar tensions between trust and privacy in other P2P systems. For example, users who rent out their own cars want to ensure they have recorded enough car usage data in the case of wrongdoing, while car renters are more concerned about the amount of information that car owners need to collect from them [53]. There is a negotiation between the two parties: renters are willing to disclose more information under transparently agreed conditions, and owners are willing to respect renters' privacy if they have evaluated them to be trustworthy. Similarly, in Airbnbs, users are less likely to trust

hosts who reveal less about themselves [54]. In the same vein, Lutz et al. [55] also found that Airbnb *hosts* had high levels of trust in the *company itself*, which further boosted their willingness to let strangers into their homes.

E. Takeaways

Sharing practices are inherently socially complex, but this complexity is compounded by the inflexibility of many current tools. For example, the lack of usable access control interfaces methods that understand and support social practices often can lead to behaviors that compromise S&P, e.g., sharing authentication credentials. These observed behaviors serve as an indication of the need for more socially informed approaches to access control.

Resource sharing between romantic partners presents significant cognitive and psychological hurdles, especially when the relationship ends. Work at the family and household level has mostly covered the fuzziness of users' access control policies, especially the desire for individual personalization, juxtaposed against financial limitations of owning one or only a few devices and accounts. Resource sharing with social acquaintances and co-workers alike is often constrained by patchwork policies and lack of coordination, including the willingness to grant access exceptions to overly strict policies when there is trust in the relationships. The thread of adapting to existing limitations via socialization, rather than through developing new technical solutions that better support these practices, runs through these three scales.

Work on sharing resources at the public scale is sparse; whereas the other scales imply some specific social relationship between actors, the public scale does not. Strangers can be renter and rentee, mutual carpool participants, elected representatives beholden to their constituents, or in any one of an endless number of relationships that require the implicit or explicit specification and negotiation of access control over shared resources. However, this diversity of public relationships and public resources shared means that it is impossible to cover every type of public relationship in one SoK.

Past work has suggested that future designs consider more complete ways of revoking access or better remind users about whom they are sharing with to help users detect account compromise. There is also a gradual shift away from single-user models, in particular disentangling the user-specific preferences and financial information of an account from access to the resources of that account [34], [56]. Even when equipped with interventions that grant greater user flexibility when dealing with S&P conflicts in the home, users resort to existing tools and social interventions. This may suggest that either user requirements have not been properly captured, or that nuances of users desire cannot yet be supported technically.

V. SHARED AND SOCIAL AUTHENTICATION

In this section we consider situations in which authentication credentials are shared for purposes other than giving blanket access to a shared digital resource, as well as work

exploring more socially aware forms of authentication. At multiple scales, password sharing presents the same cognitive burdens: having to remember all the users who know a password, changing passwords and updating users, and remembering which passwords have been reused. Indeed, users commonly reuse passwords across accounts and instead use informal (or nonexistent) methods to manage their password usage [20]. This has led to a push to replace passwords in social settings with a variety of group authentication methods, e.g., those based on shared knowledge, physical knocking patterns, or location-based verification. The work in this section covers these methods, as well as any systems directly related to actually authenticating oneself within an a household, a friend group, a work environment, or a public neighborhood.

A. Intimate relationships

As aforementioned, sharing passwords and sharing access to resources are often conflated; password sharing behaviors are covered in Section IV. However, we did not find specific literature on social authentication by couples, beyond the desire to share the resources granted access by that authentication.

B. Families and households

Singh et al. [35] found that in rural Aboriginal and Torres Strait Islander communities in Australia, which often have poor banking access, ATM cards and PINs alike are often communally shared within a household. Language and technical education constraints mean that elderly members of these communities cannot get cash without sharing their information with others. Users also often have to remember their parents' passwords, either because their parents frequently forget them, or because they helped set up the accounts [41]. We differentiate these instances of password sharing from previous ones covered in Section V because users in these situations face obstacles in the actual authentication process (i.e., language barriers, forgetfulness, less technical knowledge). The primary motivation of their password-sharing in these situations is not a desire to share their resources with others, but a challenge in communicating or proving their identities.

People with visual impairments (PVI) rely on similar sharing of access codes with partners and family members to authenticate themselves for banking services [57]. As a potential technical implementation of this phenomenon, Zhang et al. [58] created WebAlly, a system that help PVI solicit help for task-based visual CAPTCHAs by allowing them to transfer the CAPTCHA to a trusted contact who can help solve it. The authors found that WebAlly brought PVI and their friends closer while preserving the independence of PVI. Thus, assistive transfer systems like WebAlly could be a future design space to help other users who face obstacles in authentication by themselves.

C. Social acquaintances

A lack of formality around password use and reuse extends to the workplace: coworkers share credentials by simply telling other people directly, by sticking Post-It notes on shared

bulletin boards, or logging into their own accounts for other users to use [47]. When password updates are not adequately shared with all relevant employees, employees can get locked out of services required for their jobs. Conversely, when these passwords don't get updated often enough, people who leave the company can still have unauthorized access. This means that login credentials in these environments no longer simply represent the identity of an individual in a workplace, nor do they fully authenticate a user as a member of that workplace.

In place of relying on passwords, a few systems have attempted to facilitate group authentication using methods that are more "socially aware" than traditional passwords. One design that leverages existing social networks is Lineup, a photo-based social authentication system that asks users to identify their Facebook friends in photographs to authenticate their group membership [59]. Yardi et al. found that, while Lineup seemed simple to set up, the boundaries between social groups were often not rigidly defined: for example, a socially-excluded user might still be able to pass low-level security by identifying group members that they recognize.

Some of these social authentication systems require users to delineate social group boundaries or members in advance. For example, Schecter et al. [60] introduced a social authentication system that allows users to designate in advance a list of trusted individuals, or trustees, to help them authenticate themselves in the event of losing access to their online accounts. They found that an overwhelming majority of users who called their trustees could authenticate themselves successfully. Facebook has a real world implementation of this: Trusted Contacts [61], which allows users to designate a few of their Facebook friends to provide account recovery codes.

Existing authentication methods such as PINs and biometrics can be inappropriate for some small social groups where different members might require different levels of access or eventually need access revoked. Toomim et al. [62] proposed a social access control system in which access to shared photos is restricted to viewers who can answer a question designed to test mutually shared knowledge between sharer and intended viewer. Another similar group authentication system, Thumprint [63], uses a shared three-second knock pattern on a surface with a microphone and accelerometer. In both systems, sharers can easily come up with questions or knock patterns that are difficult for strangers to guess. However, Thumprint might be a bit more socially flexible, since individual member expressions of the knocks are identifiable and distinguishable, so access can be revoked or limited.

D. The public

At the public scale, there are, once again, fewer examples of coordinated authentication behaviors. One exception is Nextdoor, a neighborhood-oriented social media system. Nextdoor provides authenticates a user's membership in a specific physical community either by delivering a postcard with a unique code to an address within that neighborhood, or by allowing neighbors to vouch for each others' residence statuses. These methods exploit physical location and users'

own social networks in order to establish access to the system. Masden et al. [64] found that this location verification system made users more confident that others in their local Nextdoor site were real residents of the community, even if users had to delegate trust in Nextdoor to perform this verification.

E. Takeaways

At the intimate, family, and social acquaintance scales, there are necessary confluences between sharing resources and sharing authentication information. These behaviors are primarily covered in Section V: shared passwords and PINs unlock shared media resources, devices, and finances, simultaneously building trust in these relationships through these shared secrets. Beyond a few workarounds that leverage social relationships, however, there have only been a handful of systems designed to more smoothly facilitate group or social authentication without a reliance on passwords. People might resort to these informal password sharing practices because they don't have access to anything better. At the public level, where the informal sharing of passwords becomes untenable, users must place trust in authenticating authorities to verify identities. As such, there is room for future work on grassroots moderation and volunteer/community oversight mechanisms instead of reliance on a top-down system.

VI. MANAGING SELF-PRESENTATION

The third major domain of social S&P behaviors is *managing self-presentation*: the managed sharing that users do of information about themselves to their social circles. Palen and Dourish [65] proposed that privacy was a “dynamic, dialectic process,” building upon privacy regulation theory developed by Irwin Altman, who suggests that privacy is a continuous negotiation of boundaries according to circumstance. They describe three boundaries negotiated in privacy management: disclosure, i.e., between privacy and publicity; identity, i.e., between self and other; and temporality, i.e., between past, present, and future. Users have varying levels of control over these boundaries in their presentation of self.

In this section, when we refer to *sharing*, we specifically mean *how users share information about themselves with others*. Generally, people's willingness to share information about themselves is dependent on both the type of information being shared as well as with whom they share the information [66]. The majority of the work in this section has concerned the social and public spheres—in other words, social scales in which connected individuals may have a more tenuous social connection—and less at the intimate and household scales.

A. Intimate relationships

Within romantic relationships, users disclose and hide personal information for relationship upkeep and to instill trust in partners. Park et al. [32] identified two main themes of sharing information in order to manage self-presentation within such relationships: functional (i.e., for convenience or household maintenance) and emotional (i.e., to establish trust or to improve relationship wellbeing or support). On the flip side,

individuals might hide things in a relationship to conceal what their partner might consider as wrongdoing, maintain personal space, or, more good-naturedly, buy surprise gifts.

But IPV survivors hold a disproportionate burden for managing their online presences to protect themselves from abusers. For example, when leaving an IPV situation, the survivor must block not only the abuser, but also other family and friends who might jeopardize the victim's privacy and safety [36]. Survivors have difficulty navigating both these extensive shared social networks and complex privacy settings on social media platforms. Survivors also create profiles with false information to protect themselves.

B. Families and households

Within families and households, parents often expose information about their children, who have little say in the matter until they are older. When adults post content online, they reveal information about their children that can be linked to other online services to create profiles and inferences about these children; for example, photos of children in adults' Facebook photo albums can be correlated with offline data sources, triggering a “chain reaction of privacy violations” [67]. But adults do it anyway, weakening their children's privacy from strangers and surveillance authorities. Even when parents use surveillance apps to *protect* their children's online privacy, they inadvertently expose data about their children to unknown third parties through the apps themselves [68].

C. Social acquaintances

In social networks, users manage interpersonal relationships as a proxy for controlling the audiences who can see their posts, since context collapse effects mean users' direct control over these audiences is diminished. Wiest et al. [69] found that users' self-reported sense of closeness to their friends (i.e., strength of ties) is the best predictor of how likely they were to share various personal data. And, scenarios that trigger willingness to share often include an exchange of information that a user had in common with their friends—e.g., being within 1 mile of their friend, or socializing with a person they both knew—since users would not have to reveal completely new information to their friends. In the same vein, teens routinely swap their phones with each other to add or update contact information to each others' address books [70].

As for what they *don't* share, users often self-censor to control their presentation of self. For example, Sleeper et al. found that users would potentially share 50% more on Facebook if they could more selectively choose and block the audiences that could see the posts [71]. In a similar vein, Nextdoor users censor themselves when concerned with whether other users in their community would know they were at home: for example, instead of specifying which nights they would need a babysitter, which would reveal when they would not be at home, users solicited general recommendations for babysitters in the area [64]. And, if they don't self-censor their thoughts in time, users frequently go back and delete posts that might be embarrassing or damaging as well [72].

Users also often carefully curate their *digital* availability via online status indicators on social media services. For example, some users avoid opening apps or sign out of apps quickly when they see specific other people online with whom they wanted to avoid communicating. Conversely, users often also open apps and services just to check if specific people were online, and make inferences about these people’s real-world behavior, availability, and emotions. More adversarially, some users describe feeling tracked by others via their indicators, and end up turning off the indicator entirely [73].

Even though social media sites follow individualized models of privacy control, the nature of privacy online is contextual and networked. In particular, teens are forced to consider alternate models of privacy and challenge conceptions that young people do not care about their privacy online. Teens deftly use audience controls and coded language to limit adult access to intimate digital materials; since they know that they cannot single-handedly control all of their privacy online, they rely on managing their social relationships as a proxy for negotiating the sharing of their information [17]. More broadly, when a user shares information about themselves in this networked model, they can inevitably reveal information about others (e.g., photos of multiple friends); ConsenShare [74], a system that notifies users about information that others share about them and encrypts it, attempts to combat this bystander information leakage.

This goes against adult preferences for “locking everything down”. For example, older adults are largely concerned with unauthorized access to personal information, as opposed to context collapse, unintended sharing, or misuse by large institutions [75]. These concerns, which deal with keeping the details of their daily lives private, differ from those of younger users, who are more concerned with self-presentation. Older adults thus primarily mitigated these concerns by limiting what they posted or staying off social media entirely. However, even some adults have adapted indirect S&P behaviors from teens to form their self-presentation. For example, adults in the workplace share of iTunes playlists not to be able to listen to others’ music, but to explicitly manage the way others form an impression of the musical tastes of the sharer [76].

Lindqvist et al. [77] also found that users of Foursquare use their check-in behavior to signal to others about their personality and self. For example, some users avoid checking in at their own home over and over again because they do not want to present themselves as boring people; others avoid checking into fast food restaurants because they are embarrassed about being seen there. And, like online status indicators, check-ins could also signal availability. However, unlike social-group-level signals, Foursquare users seemed largely unconcerned with public visibility of their profiles.

D. The public

As users conduct more of their lives online, the boundaries between intimate and public social relationships grow increasingly fuzzy. One domain where this is especially pronounced is online dating; a few works have explored how

users manage their privacy and combat context collapse in these environments. For example, Cobb and Kohno [78] explored user conceptions of privacy on online dating platforms, finding that when users saw profiles of people they knew offline, their impressions from the profile lingered in future in-person interactions. Even if men and women tend to present themselves differently on these apps (men to find hookups and relationships, women to self-validate and make friends [79]), online dating users as a whole selectively disclose dating information, both on their profiles and to people they know in person, purposefully trying to hide themselves from people they know [78]. And, similar to ridesharers [50], while users appreciate information on profiles that helps them look up and research other users online, they are also wary of being looked up themselves. Users also regularly take screenshots of profiles and conversation to share with friends, either to shame bad behavior or get opinions about potential matches [78].

These behaviors hint at some set of social expectations for sharing personal information in online dating, but the extent of this sharing is not yet clearly defined. For example, are personal health statuses fair game? Warner et al. [80] explored user reactions to linking HIV-positive status to profiles on Grindr, a geosocial hookup app, finding that while some HIV-positive users hide their status to reduce their potential of being stigmatized, others purposefully disclose their status for the same reasons. Warner et al. also found that when users keep their status private, other users make social assumptions about why they were keeping their status private, revealing social expectations that force them to disclose their statuses.

This series of events is an example of *privacy unraveling*, termed by Peppet [81], who argues that users will eventually face a limited set of options for disclosing personal information, since others might assume that those who withhold or remove information are doing so because they are hiding something unsavory. (Indeed, Minaei et al. [82] have even also begun exploring how to protect users against mass collection of their *deleted* posts). As such, users might be forced into disclosing their information to avoid being socially stigmatized: more broadly, advertising about location features might pressure to overshare their location data with others [83].

This phenomenon lends itself to thought exercises on the future of self-presentation at the public level. TheWebConf, for example, has held workshops devoted to exploring the consequences of technologies in the dystopian science fiction television series *Black Mirror*. One such workshop conjured up future scenarios of users developing anxiety because they cannot match the socially-shared fitness levels of those around them, and fitness-tracker-wearing citizens being ranked by public health agencies for priority access to resources, with those who exercise less often deemed “lazy” and de-prioritized [84]. There is even already evidence for such scenarios. Users of wearable fitness tracker devices tend to lack knowledge of the threats associated with sharing data from their devices [85]. Meanwhile, the wearable fitness tracker company Fitbit has already signed agreements with health insurance companies to provide discounted plans to their users [86].

There have been a few end-user tools that could help circumvent tracking and enable users to more effectively obfuscate their identities. Tor [87], of course, is the well-known onion network that relays users' encrypted traffic through random network nodes to conceal their Internet activity. Since it requires a network of volunteers to run these relay nodes, it is thus cooperative by design: the more volunteers, the more robust the network is against attacks. TrackMeNot [88], a browser extension, aims to help users obfuscate online searches by submitting "ghost queries" to search engine companies to conceal users' actual searches. Similarly, AdNauseum [89] aims to protect users from tracking advertisers by silently clicking on ads that have been blocked and sending noisy data back to advertisers. However, since these tools tend to resist mainstream adoption and thus collective, coordinated use, there have not been demonstrable public effects of hidden browsing information.

E. Takeaways

At each social scale, users continuously weigh the consequences of divulging personal information against building trust and camaraderie with those around them, carefully curating their digital self-presentation while often simultaneously conceding that it is impossible for them to control all digital content about themselves. However, existing audience control technologies do not fully help users navigate the broad range of their social sharing preferences. The upshot is that users are prone to interpersonal privacy violations where unintended audiences encounter sensitive content. Moreover, as divulging personal information online becomes more normative, individuals may feel increasing peer pressure to follow suit or otherwise risk being labeled as suspicious and/or paranoid. We foresee an opportunity for future systems that empower users to challenge these norms by, e.g., affording users greater agency over who can use personal data that they share online.

VII. INFLUENCING OTHERS' S&P BEHAVIORS

A final broad category of study in the literature has focused on how social groups can influence others' S&P behaviors. Users engage in positive S&P behaviors more frequently when spurred on by social triggers than when forced to do so, and are more likely to share their own S&P experiences with others when the experiences were socially-triggered [90]. This behavior is reflected in work across different social scales: reliance on partners for S&P knowledge, children consulting parents for help with S&P online, older adults preferring S&P aid from in-person interactions over online searches, along with a slew of S&P behavioral nudges via social proof.

A. Intimate relationships

Within romantic relationships, users tend to follow advice and influence behaviors similar to those within households: the user with the more technical knowledge tends to set up device and network settings within the home [91]. In IPV situations, abusers tend to fill the more "tech-savvy" role, from purchasing the devices to setting them up with

restrictive surveillance controls [36]. As such, some survivors avoid technology altogether after leaving these relationships, even though they need it to register for social services or look for employment. With no one else to turn to, survivors rely on professionals in the IPV ecosystem, like social case managers and attorneys, to help them with their S&P [36]. These professionals, in turn, feel limited in the S&P advice they can offer since they themselves are not S&P experts; instead, they are forced to learn on the job on behalf of their survivor clients.

B. Families and households

Much research at this scale has explored how parents influence (and control) children's S&P behavior in the home. For example, young children often recognize S&P behaviors (e.g., identifying information as sensitive, understanding what was appropriate to share) and develop strategies to manage their concerns, but ask their parents for help anyway [92]. Even as they understand the importance of authentication, e.g., putting passwords on their devices to deter siblings from viewing their information, they still need their parents to teach them about S&P events online.

Conversely, parents prefer deferring action on their young kids' online privacy to the future, i.e., when they were older and more engaged socially online, instead of building foundational S&P behaviors [92]. When parents do directly intervene on their *teenage* children's privacy, e.g., through parental controls or setting up the teen's social media privacy settings, there is a potential suppressive effect: while the teens' exposure to online risks is reduced, so are their opportunities to engage with others online and learn how to cope with these risks [93]. Generally, work on the parent-child S&P relationship recommends that parents engage in more active parental mediation, e.g., discussing their posts or commenting on Facebook, rather than "locking everything down", in order to empower children to learn to engage with others online.

Murthy et al. [94] also observed self-appointed technology managers within Indian households who help establish S&P guidelines within families. These managers assume that older members of the household are less technically literate and more vulnerable to S&P threats; accordingly, they often unilaterally control and make changes to S&P settings and device settings on behalf of older users, suggesting paternalism and removing digital agency from older users. Older adult users can look outside of the household, though: they surprisingly also trust local after-sales support staff, since they almost exclusively only seek S&P information when they needed to fix a problem [95]. The in-person aspect of such advice is most key: older adults tend to avoid searching on the Internet for help and are skeptical of the trustworthiness of information online.

C. Social acquaintances

Digioia and Dourish [96] first introduced social navigation as a way to help visualize user activities within a system and better incorporate the user into making security decisions.

They specifically refer to this as a way for users to conceptualize how their own security behaviors compare to and are seen by others. Richter Lipford and Zurko then proposed a community oversight paradigm for security-related behaviors that would take into account social context and processes [16]. Chouhan et al. [15] further formalized this into Community Oversight for Privacy and Security (CO-oPS), a framework for guiding users to interact with the people they trust in their social communities when making S&P decisions. We find that such behaviors are triggered via two main methods of influence: (1) peer-level influence through conversations about and observations of positive S&P behaviors, and (2) top-down nudges to make positive S&P choices.

1) **Peer influence:** Prior work has found that social influence affects end-user S&P behaviors, be it through conversation, observation, or peer connection more broadly.

Conversations about S&P are driven by a sense of personal responsibility for others around us, e.g., to protect or warn them about threats, or to gather more info about something we are experiencing ourselves. Negative personal experiences are frequent catalysts for conversations about security, as well as news articles and personal observations of security-related behaviors, since users often try to notify or warn those around them about what they’ve experienced or learned [2]. Rader et al. [3] found that undergraduates at Michigan State University often learned lessons from stories about security events they heard from family and friends, which impacted their subsequent S&P decisions. These students then retold these stories to others to inform even more people [3]. Similarly, Facebook users who seek support in their friends and family upon receiving suspicious login alerts from Facebook find a sense of camaraderie in these conversations (and those who do not, feel embarrassment) [97]. Yet, those who seek information about others’ S&P experiences also feel less inclined to share their own experiences and help others, preferring passive participation in conversation [15]. Watson et al. [45] also observed that group interactions about S&P were rare and tended to focus on abstract concepts found in the news instead of personal experiences. These groups might see their S&P behaviors as lacking in conversational value, even if they end up enjoying talking about them.

Adoption rates of security features on social media sites can also be affected by the security practices of potential adopters’ friends. For example, Das et al. [98] found that social influence affected security feature adoption when users observed Facebook friends from multiple different social circles also adopting the features; and, while a person with more feature-adopting friends is likely to adopt features themselves, a person with just a few such friends might be negatively influenced. Similarly, De Luca et al. [99] found that users’ adoption of security messaging tools like Threema was significantly dependent on whether their friends were also using these tools, rather than their inherent improved security. Building on the power of peer influence, Bonneau et al. [100] envision Privacy Suites, a system that allows users to import “suites” of privacy settings that have been pre-selected by their friends

or trusted experts—modifying them if they wish—and that supports public reviews of the suites to establish trust in them.

2) **Social nudges:** Other work has explored how social proof can be employed in top-down nudges to influence S&P behavior. For example, Das et al. [101] showed Facebook users announcements about extra features to help secure their accounts, including prompts about the number of their friends who used these features. Announcements that included these social proof prompts were more effective at getting them to explore these features than those that did not; more people also adopted these features. This recalls similar prior findings [102] that password meters based on social pressure resulted in users generating slightly stronger passwords.

Emami-Naeini et al. [103] also investigated the role of social cues in user decisions about data collection by IoT devices, showing users messages about the actions that either “friends” or “experts” took in various data collection scenarios. They found that participants were influenced by when friends denied data collection and when experts allowed data collection (and not so vice versa). They also observed that this influence could change with repeated exposure to social proof; for example, when friends repeatedly allowed risky data collection, participants were less influenced by them. In the same vein, when adolescent users on the Japanese social networking site Himabu were presented with negative framing of choices, e.g., “90% of users would not share a photo without permission,” they were more likely to avoid potentially risky choices than vice versa (conversely, when users saw affirmative framings, they tended to make riskier choices) [104].

D. The public

Work at this scale includes both empirical work on social influence at demographic levels, as well as more systems that use social influence to guide better S&P behaviors.

In the empirical category, a number of papers have explored demographic variations in how users approach S&P behaviors and the role that influence plays. Ur and Wang [105] argued that considerations and support for user privacy has not been equitably distributed internationally, and proposed a framework for evaluating how well a social network site’s privacy settings supports cross-cultural user bases. The framework includes questions about local cultural expectations of privacy, governmental restrictions and requirements on data collection, and local language availability. Prior work [97] finding that users from more collectivistic countries (e.g., Brazil, Vietnam, India) seek out information from others at higher rates than those in the US corroborates the need for such a framework.

More broadly, male and female users might think about privacy in different ways: whereas females are more likely to mention other people and bring up issues of safety and respect, males tend to refer to privacy as having freedom or being anonymous [106]. There are also potential age difference effects: older adults describe privacy in terms of space (e.g., home invasion) rather than information, and tend to view private information as concrete objects like documents or specific secrets much more than young adults do [106].

Other work has explored the nuances of end-user adoption of S&P advice. Herley [107] argues that even though users are often portrayed as lazy and unmotivated about their personal security, their rejection of security advice is economically rational. He suggests that most security advice offers a poor trade-off between the costs of implementation versus the actual user benefits, so users choose not to take it, especially if experts themselves do not necessarily know the full extent of security risks and harms. More work on the actual harm that users endure, as well as user education technologies that target just those who are at risk, might be more effective, especially given that users tend to reject advice that is too marketing-oriented, especially if they did not feel like they were personally at risk [4]. Specifically, users of lower socio-economic status, who seek advice from different sources than those with more resources and technical education, might be more vulnerable and need more attention [108].

On the technical front, Goecks et al. [109] expand the scope of social-influence S&P nudges to the public, presenting two systems that leverage social navigation for public-scale S&P interventions. The first, Acumen, shows users the prior decisions taken by others in deciding whether to accept a website’s cookie settings; Bonfire, the second, shows prior decisions in the context of firewall settings. A particular focus of these systems is avoiding “information cascades,” a sort of herd mentality where users blindly follow the decisions of others, resulting in incorrect choices.

E. Takeaways

At the intimate relationship scale, we saw that user behaviors mirror that of the household scale, down to the imbalance in technical education between abuser and survivor in IPV situations. In particular, survivors seek S&P advice from IPV professionals who may or not may not be qualified to give it, and who must often instead negotiate on the survivor’s behalf with technology companies to get abusive content removed. There is thus a future design space for systems that allow IPV professionals to more efficiently advocate for their survivor clients with technology companies, rather than the companies outright implementing new S&P controls that neither the professional nor the survivor client might know how to use.

This imbalance and advocacy continues at the household level, where parents are responsible for giving S&P advice to both their younger children and to their own older parents. In both cases, parents tend to prefer “lockdown” approaches to protect what they perceive as less technically-literate, more vulnerable members of their family, taking away S&P agency from both the younger children and older adults.

At the social acquaintance level, much work has been done on designing ways to inform users about S&P behaviors via socially-flavored nudges and observational alerts, as well as making sense of how users have conversations with those around them about S&P. While this work hints at the latent power of social responsibility that users feel for their communities’ S&P, they tend to recommend conversations and social nudges and alerts as ways to educate users only about existing

ways to improve S&P. Observing friends is the most frequent trigger of change in S&P behaviors, but many security settings are inherently private and not observable. Potential adopters might feel that existing users are simply paranoid [98], [110]. Instead, we suggest that future work could provide outlets for these users to collectively demand or construct better methods.

The threads of institutional entrenchment and near-paternalism of smaller social scales continue at the public level. S&P researchers have often tended to believe that users as a whole irrationally don’t adopt S&P behaviors, without considering the social and cultural contexts of heeding such advice, as well as the risk assessments that users make about their S&P. Past work has investigated how to prevent a cascade of users from making S&P decisions deemed “wrong”, but does not entertain the possibility of a cascade of users who want to make choices outside of the existing set of options.

VIII. DISCUSSION

Throughout our systematization of prior literature in social cybersecurity, we used Ackerman’s social-technical gap, or the difference between what users desire socially and what is supported technically [19], as a lens to outline potential areas for future work. The social-technical gap is extensively cited in the broader HCI literature [111], and has been specifically referenced in the usable S&P [112]–[114] and online communities [115], [116] literature to illustrate how technical systems may fail to account for human social behavior.

Based on Ackerman’s [19] assertions that (1) the gap is enduring, so we should always strive to do *something* about it, and (2) we should not force users to adapt to technology but adapt technology to its users, we developed three questions to assess the social-technical gap for each social S&P use case identified in prior literature:

- **Are there existing systems that help facilitate this social S&P use case?** For example, the Thumprint system [63] supports group authentication by using secret knocking patterns to authenticate members of a social group, so the box for that use case’s row is checked in Figure 1. As another example, while audience selection tools exist on social media platforms, teens choose to use coded language instead to read their intended audience [17], so that row gets a starred check. In contrast, our review indicated that there are no technical systems to facilitate giving out S&P advice that directly targets vulnerable populations [107], [108], so it does not get a check mark.
- **Can users fit the affordances of existing S&P systems without altering their ideal social behaviors?** For example, in supporting group authentication, Thumprint knocking patterns are simple for users to collectively devise and share. Conversely, for social media audience selection, existing tools must be pre-defined, but teens have dynamic audiences in mind when posting.
- **Can users use these existing systems, as intended, to meet both their ideal social behaviors and S&P goals?** For example, Thumprint uses a single shared group secret, the expression of which is individually

Behavior Domain	Non-exhaustive List of Specific Social Use Cases	Are there existing systems that help facilitate this social S&P use case?	Can users fit the affordances of existing systems without altering their ideal social behaviors?	Can users use these existing systems, as intended, to meet both their ideal social behavior and S&P goals?
Shared resources	Anonymized rider and driver data in ridesharing [51,52]	✓	✓	✓
	Car-, ride- and home-sharing [50-55]	✓		✓
	Makerspace physical access controls [49]	✓		
	Static work environment physical access control policies [7,8]	✓		
	Posting passwords on a Post-It note in a shared workspace [47]	✓*		
	Sharing passwords with coworkers via Slack, email [41,47]	✓*		
	Managing access to files shared with coworkers via email [48]	✓		✓
	Updating passwords upon ending account sharing [34]	✓*		
	Having awkward conversations to end account sharing [34]	✓*		
	Sharing smart speakers and leaking personal information [31,43,44]	✓		
	Sharing tablets and computing devices with different profiles [30,39,40,42]	✓		
	Leaving devices unlocked to signal trust [30]	✓*		
	Sharing access to media, wi-fi, devices via passwords [30,34,35,38,41]	✓		
	Sharing calendars and task lists [30,33]	✓		✓
	Abusers surveilling partner devices and online activity [36,37]	✓*		
	Abusers accessing IPV survivors' accounts [36,37]	✓*		
IPV survivors destroying devices after leaving relationships [37]	✓*			
Joint partner bank accounts with one password [35]	✓*			
Shared authentication	Nextdoor location-based authentication [64]	✓	✓	✓
	Group authentication via secret knocking patterns [63]	✓	✓	✓
	Group authentication via shared knowledge [62]	✓	✓	✓
	Last-resort authentication via list of trusted individuals [60,61]	✓	✓	✓
	Group authentication via identifying friends in photos [59]	✓	✓	✓
	Updating shared passwords/credentials with coworkers [47]	✓*		
	Removing credentials when employees leave a company [47]	✓*		
Users sharing banking PINs w/ family due to accessibility challenges [35,41,57]	✓*			
Managing self-representation	Volunteering to run relays for Tor to help others obscure traffic [87]	✓	✓	✓
	Sharing health status in public due to social pressure [80,84-86]	✓		
	Online dating users hiding profiles from people they know [78]	✓		
	Posting public location check-ins on social media to signal personality [77]	✓		✓
	Self-censoring on Nextdoor to hide when home is empty [64]	✓*		✓
	Self-censoring on social media to avoid seeming silly or offending others [71,72]	✓*		✓
	Using coded language on social media instead of audience selection tools [17]	✓		
	Sharing digital music playlists to manage coworkers' impressions [76]	✓*		
	Notifications when others share information about yourself [74]	✓		
	Parents leaking privacy of children online [67,68]	✓		
	IPV survivors creating profiles with false information [36]	✓*		
Sharing private information with partner to establish trust [32]			n/a	
Hiding from partner due to embarrassment or wrongdoing [32]			n/a	
Advice and influence	Consideration of cross-cultural S&P practices and communication [97,105,106]			n/a
	Public advice that targets those at risk of S&P attacks [4,107,108]			n/a
	Top-down social proof nudges [101-104]	✓		✓
	Observing and emulating others' S&P behaviors [98-100]	✓		✓
	Sharing stories about S&P experiences and finding camaraderie [2,3,15,45,97]			n/a
	Older users consulting adult children for S&P advice [95]			n/a
	Older users losing S&P agency within households [94]			n/a
	Delegating home network and S&P setup to one household member [91]	✓*		✓
	Younger children consulting parents for S&P advice [92]			n/a
Parents locking down children's content access/privacy settings [68,92,93]	✓		✓	
IPV survivors seeking S&P advice from outside the relationship [36]			n/a	

Social Scale Key: Intimate relationships Families and households Social acquaintances The public

Fig. 1: ✓* indicates that yes, there is a relevant technical system, but the system does not directly facilitate this use case.

distinguishable, allowing outsiders to be easily identified without requiring insiders to keep individual secrets from each other. On the other hand, for social media audience selection, teens have no usable way of dynamically specifying a desired audience, outside of manually selecting individual audience members from a picklist of hundreds or thousands of individuals.

We answer these questions broadly in the subsections to follow. Figure 1 summarizes our systematization of prior work across behavioral domains and social scales, addressing the three questions listed above.

A. Are there existing systems that help facilitate this social S&P use case?

A great majority of behaviors and use cases observed in social cybersecurity work involve *some* technical system, e.g., passwords, physical devices, shared knowledge authentication, audience selection tools on social media. However, we note a distinction between novel systems designed to directly facilitate social behaviors, and extant systems that were worked around or modified to fit social needs. Most examples of the former were developed for use-cases that required little direct social interaction, e.g., social authentication systems [59], [60], [62], [63], the Tor relay network [87] and Nextdoor’s location-based authentication [64]. There were also two examples of domain-specific social S&P systems for carsharing [50] and ridesharing [50]. But, overall, there were few examples of novel S&P systems that required direct social interaction, specifically in the advice and influence behavior domain.

B. Can users fit the affordances of existing S&P systems without altering their ideal social behaviors?

Many S&P systems are designed to be ignorant of social context, and force users to choose between security and social acceptability. For example, when a user shares a Netflix account, they are burdened with making their private account details visible, having their personalized content recommendations being disrupted, exposing their password to unknown actors, and risking their account being commandeered by the secondary user; in turn, they may change how often they use the account or stop using it entirely [34]. As another example, teenagers use coded language instead of existing audience selection tools on social media to hide from adults in plain sight; they do this with the expectation that they cannot control all of their privacy anyway [17]. And, even in domains without direct technical systems, like influencing others’ S&P behaviors, users are expected to abide by institutionally-set advice for *existing* technologies. Such cases are ripe for exploration on how to better technically integrate social ties and support user agency.

C. Can users use these existing systems, as intended, to meet both their ideal social behaviors and S&P goals?

There are nuances in social behaviors that existing technical systems are ill-equipped to navigate. For example, users often share passwords or elide authentication altogether to facilitate

sharing digital resources. Indeed, passwords and personal devices, though not designed to be shared, are regularly posted by users in common work areas and left unlocked for others on purpose, respectively. These use cases reveal limitations in existing S&P systems: by failing to account for human social behaviors, these systems no longer serve their intended purpose. For example, users use single-user authentication methods (e.g., passwords, keycards) as proxies for group access control. Existing examples of production-ready social cybersecurity systems are relatively few: those that employ collectives for obfuscation but assume no direct user interaction—e.g., Tor [87], TrackMeNot [88], or AdNauseum [89]—and those that allow for social fallback authentication for individual accounts (e.g., Facebook Trusted Contacts [61]).

D. Future work

We foresee two key directions for future work on addressing the social-technical gap in S&P systems. First, we have analyzed a rich literature describing the social inadequacies of existing S&P systems, but we found little systems work addressing these inadequacies. Following recommendations from prior work [12], we see a need for the design of social S&P systems that foster greater: observability—systems that make it easy for users to see how others are protecting their S&P; cooperation—systems that allow groups to act in mutual benefit of everyone’s S&P; and, stewardship—systems that allow individuals to act in benefit of others’ S&P. We also found little extant work—theory or systems—on the “public” social scale in the context of S&P. Collective action systems may be explored here as a mechanism for users to directly advocate for more human-centered S&P protections against, e.g., web tracking and surveillance technologies [117].

IX. CONCLUSION

Throughout this structuring process we identified four key behavior domains of social cybersecurity work, and broke them each down by social scale. We found extensive descriptive evidence that, today, end-users must often either adapt their ideal social behaviors to realize their S&P settings, or adapt their ideal S&P behaviors to reduce social friction. However, as illustrated in Figure 1, there has been comparatively little *prescriptive* systems work on addressing this social-technical gap. In short, ignoring human social behaviors in designing S&P systems leads to maladaptive user behaviors that either reduce security, cause social friction, or both. In contrast, by designing for and leveraging human social behaviors in S&P systems, there is an opportunity to both increase the efficacy *and* the widespread adoption of those systems.

ACKNOWLEDGMENTS

This work was generously supported, in part, by NSF Grant #1755625. We also thank our anonymous reviewers for being generous with their time; their comments helped us improve the paper.

REFERENCES

- [1] S. Das, J. Lo, L. Dabbish, and J. I. Hong, "Breaking! a typology of security and privacy news and how it's shared," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–12.
- [2] S. Das, T. H.-J. Kim, L. A. Dabbish, and J. I. Hong, "The effect of social influence on security sensitivity," in *10th Symposium On Usable Privacy and Security ({SOUPS} 2014)*, 2014, pp. 143–157.
- [3] E. Rader, R. Wash, and B. Brooks, "Stories as informal lessons about security," in *Proceedings of the Eighth Symposium on Usable Privacy and Security*, 2012, pp. 1–17.
- [4] E. M. Redmiles, A. R. Malone, and M. L. Mazurek, "I think they're trying to tell me something: Advice sources and selection for digital security," in *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2016, pp. 272–288.
- [5] A. Adams and M. A. Sasse, "Privacy issues in ubiquitous multimedia environments: Wake sleeping dogs, or let them lie?" 1999.
- [6] J. Bonneau, C. Herley, P. C. Van Oorschot, and F. Stajano, "The quest to replace passwords: A framework for comparative evaluation of web authentication schemes," in *2012 IEEE Symposium on Security and Privacy*. IEEE, 2012, pp. 553–567.
- [7] L. Bauer, L. F. Cranor, R. W. Reeder, M. K. Reiter, and K. Vanica, "A user study of policy creation in a flexible access-control system," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2008, pp. 543–552.
- [8] —, "Real life challenges in access-control management," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2009, pp. 899–908.
- [9] C. Bravo-Lillo, L. Cranor, S. Komanduri, S. Schechter, and M. Sleeper, "Harder to ignore? revisiting pop-up fatigue and approaches to prevent it," in *10th Symposium On Usable Privacy and Security ({SOUPS} 2014)*, 2014, pp. 105–111.
- [10] J. Hong, S. Das, T. H.-J. Kim, and L. Dabbish, "Social cybersecurity: Applying social psychology to cybersecurity," *Human Computer Interaction Institute, Carnegie Mellon University*, 2015.
- [11] S. Das, "Social cybersecurity: Reshaping security through an empirical understanding of human social behavior. dissertations (may 2017)," 2017.
- [12] —, "Social cybersecurity: Understanding and leveraging social influence to increase security sensitivity," *it-Information Technology*, vol. 58, no. 5, pp. 237–245, 2016.
- [13] F. Dalpiaz, E. Paja, and P. Giorgini, "Security requirements engineering via commitments," in *2011 1st Workshop on Socio-Technical Aspects in Security and Trust (STAST)*. IEEE, 2011, pp. 1–8.
- [14] A. Ferreira, J.-L. Huynen, V. Koenig, and G. Lenzini, "A conceptual framework to study socio-technical security," in *International Conference on Human Aspects of Information Security, Privacy, and Trust*. Springer, 2014, pp. 318–329.
- [15] C. Chouhan, C. M. LaPerriere, Z. Aljallad, J. Kropczynski, H. Lipford, and P. J. Wisniewski, "Co-designing for community oversight: Helping people make privacy and security decisions together," *Proceedings of the ACM on Human-Computer Interaction*, vol. 3, no. CSCW, pp. 1–31, 2019.
- [16] H. R. Lipford and M. E. Zurko, "Someone to watch over me," in *Proceedings of the 2012 New Security Paradigms Workshop*, 2012, pp. 67–76.
- [17] A. E. Marwick and D. Boyd, "Networked privacy: How teenagers negotiate context in social media," *New media & society*, vol. 16, no. 7, pp. 1051–1067, 2014.
- [18] E. T. Hall, R. L. Birdwhistell, B. Bock, P. Bohannon, A. R. Diebold Jr, M. Durbin, M. S. Edmonson, J. Fischer, D. Hymes, S. T. Kimball *et al.*, "Proxemics [and comments and replies]," *Current anthropology*, vol. 9, no. 2/3, pp. 83–108, 1968.
- [19] M. S. Ackerman, "The intellectual challenge of cscw: the gap between social requirements and technical feasibility," *Human-Computer Interaction*, vol. 15, no. 2-3, pp. 179–203, 2000.
- [20] E. Hayashi and J. Hong, "A diary study of password usage in daily life," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2011, pp. 2627–2630.
- [21] K. M. Carley, G. Cervone, N. Agarwal, and H. Liu, "Social cyber-security," in *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*. Springer, 2018, pp. 389–394.
- [22] T. Schnitzler, M. S. Mirza, M. Dürmuth, and C. Pöpper, "Sok: Managing longitudinal privacy of publicly shared personal online data." *Proc. Priv. Enhancing Technol.*, vol. 2021, no. 1, pp. 229–249, 2021.
- [23] V. Braun and V. Clarke, "Reflecting on reflexive thematic analysis," *Qualitative Research in Sport, Exercise and Health*, vol. 11, no. 4, pp. 589–597, 2019.
- [24] C. Kadushin, "Social distance between client and professional," *American Journal of Sociology*, vol. 67, no. 5, pp. 517–531, 1962.
- [25] R. F. Murphy, "Social distance and the veil," *American Anthropologist*, vol. 66, no. 6, pp. 1257–1274, 1964.
- [26] D. L. Lawrence and S. M. Low, "The built environment and spatial form," *Annual review of anthropology*, vol. 19, no. 1, pp. 453–505, 1990.
- [27] S. Niemeier, C. P. Campbell, and R. Dirven, *The cultural context in business communication*. John Benjamins Publishing, 1998.
- [28] G. M. Olson and J. S. Olson, "Distance matters," *Human-computer interaction*, vol. 15, no. 2-3, pp. 139–178, 2000.
- [29] J. M. Wilson, M. Boyer O'Leary, A. Metiu, and Q. R. Jett, "Perceived proximity in virtual work: Explaining the paradox of far-but-close," *Organization studies*, vol. 29, no. 7, pp. 979–1002, 2008.
- [30] M. Jacobs, H. Cramer, and L. Barkhuus, "Caring about sharing: Couples' practices in single user device access," in *Proceedings of the 19th International Conference on Supporting Group Work*, 2016, pp. 235–243.
- [31] Y. Huang, B. Obada-Obieh, and K. Beznosov, "Amazon vs. my brother: How users of shared smart speakers perceive and cope with privacy risks," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–13.
- [32] C. Y. Park, C. Faklaris, S. Zhao, A. Sciuto, L. Dabbish, and J. Hong, "Share and share alike? an exploration of secure behaviors in romantic relationships," in *Fourteenth Symposium on Usable Privacy and Security ({SOUPS} 2018)*, 2018, pp. 83–102.
- [33] J. Lin, J. I. Hong, and L. Dabbish, "'it's our mutual responsibility to share" the evolution of account sharing in romantic couples," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW1, pp. 1–27, 2021.
- [34] B. Obada-Obieh, Y. Huang, and K. Beznosov, "The burden of ending online account sharing," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–13.
- [35] S. Singh, A. Cabraal, C. Demosthenous, G. Astbrink, and M. Furlong, "Password sharing: implications for security design based on social practice," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2007, pp. 895–904.
- [36] D. Freed, J. Palmer, D. E. Minchala, K. Levy, T. Ristenpart, and N. Dell, "Digital technologies and intimate partner violence: A qualitative analysis with multiple stakeholders," *Proceedings of the ACM on Human-Computer Interaction*, vol. 1, no. CSCW, pp. 1–22, 2017.
- [37] T. Matthews, K. O'Leary, A. Turner, M. Sleeper, J. P. Woelfer, M. Shelton, C. Manthorne, E. F. Churchill, and S. Consolvo, "Stories from survivors: Privacy & security practices when coping with intimate partner abuse," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, pp. 2189–2201.
- [38] M. L. Mazurek, J. Arsenaault, J. Bresee, N. Gupta, I. Ion, C. Johns, D. Lee, Y. Liang, J. Olsen, B. Salmon *et al.*, "Access control for home data sharing: Attitudes, needs and practices," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2010, pp. 645–654.
- [39] T. Matthews, K. Liao, A. Turner, M. Berkovich, R. Reeder, and S. Consolvo, "'she'll just grab any device that's closer" a study of everyday device & account sharing in households," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 2016, pp. 5921–5932.
- [40] A. B. Brush and K. M. Inkpen, "Yours, mine and ours? sharing and use of technology in domestic environments," in *International Conference on Ubiquitous Computing*. Springer, 2007, pp. 109–126.
- [41] J. Kaye, "Self-reported password sharing strategies," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2011, pp. 2619–2622.
- [42] S. Egelman, A. B. Brush, and K. M. Inkpen, "Family accounts: A new paradigm for user accounts within the home environment," in *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, 2008, pp. 669–678.

- [43] W. Jang, A. Chhabra, and A. Prasad, "Enabling multi-user controls in smart home devices," in *Proceedings of the 2017 Workshop on Internet of Things Security and Privacy*, 2017, pp. 49–54.
- [44] E. Zeng and F. 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)*, 2019, pp. 159–176.
- [45] H. Watson, E. Moju-Igbene, A. Kumari, and S. Das, "'we hold each other accountable': Unpacking how social groups approach cybersecurity and privacy together," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–12.
- [46] E. Moju-Igbene, H. Abdi, A. Lu, and S. Das, "'how do you not lose friends?': Synthesizing a design space of social controls for securing shared digital resources via participatory design jams," *To Appear In Proceedings of the 31st USENIX Security Symposium (SEC)*, 2022.
- [47] Y. Song, C. Faklaris, Z. Cai, J. I. Hong, and L. Dabbish, "Normal and easy: Account sharing practices in the workplace," *Proceedings of the ACM on Human-Computer Interaction*, vol. 3, no. CSCW, pp. 1–25, 2019.
- [48] S. Volda, W. K. Edwards, M. W. Newman, R. E. Grinter, and N. Ducheneaut, "Share and share alike: exploring the user interface affordances of file sharing," in *Proceedings of the SIGCHI conference on Human Factors in computing systems*, 2006, pp. 221–230.
- [49] J. Logas, R. Zhong, S. Almeida, and S. Das, "Tensions between access and control in makerspaces," *Proceedings of the ACM on Human-Computer Interaction*, vol. 4, no. CSCW3, pp. 1–33, 2021.
- [50] K. Radke, M. Brereton, S. Mirisae, S. Ghelawat, C. Boyd, and J. G. Nieto, "Tensions in developing a secure collective information practice—the case of agile ridesharing," in *IFIP Conference on Human-Computer Interaction*. Springer, 2011, pp. 524–532.
- [51] A. Pham, I. Dacosta, B. Jacot-Guillarmod, K. Huguenin, T. Hajar, F. Tramèr, V. Gligor, and J.-P. Hubaux, "Privateride: A privacy-enhanced ride-hailing service," *Proceedings on Privacy Enhancing Technologies*, vol. 2017, no. 2, pp. 38–56, 2017.
- [52] Y. He, J. Ni, X. Wang, B. Niu, F. Li, and X. Shen, "Privacy-preserving partner selection for ride-sharing services," *IEEE Transactions on Vehicular Technology*, vol. 67, no. 7, pp. 5994–6005, 2018.
- [53] P. Bossauer, T. Neifer, G. Stevens, and C. Pakusch, "Trust versus privacy: Using connected car data in peer-to-peer carsharing," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–13.
- [54] X. Ma, J. T. Hancock, K. Lim Mingjie, and M. Naaman, "Self-disclosure and perceived trustworthiness of airbnb host profiles," in *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*, 2017, pp. 2397–2409.
- [55] C. Lutz, C. P. Hoffmann, E. Bucher, and C. Fieseler, "The role of privacy concerns in the sharing economy," *Information, Communication & Society*, vol. 21, no. 10, pp. 1472–1492, 2018.
- [56] J. Barkley, K. Beznosov, and J. Uppal, "Supporting relationships in access control using role based access control," in *Proceedings of the fourth ACM workshop on Role-based access control*, 1999, pp. 55–65.
- [57] J. Hayes, S. Kaushik, C. E. Price, and Y. Wang, "Cooperative privacy and security: Learning from people with visual impairments and their allies," in *Fifteenth Symposium on Usable Privacy and Security ({SOUPS} 2019)*, 2019.
- [58] Z. Zhang, Z. Zhang, H. Yuan, N. M. Barbosa, S. Das, and Y. Wang, "Webally: Making visual task-based captchas transferable for people with visual impairments," in *Seventeenth Symposium on Usable Privacy and Security ({SOUPS} 2021)*, 2021, pp. 281–298.
- [59] S. Yardi, N. Feamster, and A. Bruckman, "Photo-based authentication using social networks," in *Proceedings of the first workshop on Online social networks*, 2008, pp. 55–60.
- [60] S. Schechter, S. Egelman, and R. W. Reeder, "It's not what you know, but who you know: a social approach to last-resort authentication," in *Proceedings of the sigchi conference on human factors in computing systems*, 2009, pp. 1983–1992.
- [61] "How can i choose friends to help me log in if i ever get locked out of my facebook account?" 2021. [Online]. Available: <https://www.facebook.com/help/119897751441086>
- [62] M. Toomim, J. Fogarty, J. Landay, N. Morris, X. Zhang, and T. Kohno, "Access control by testing for shared knowledge," Feb. 26 2013, uS Patent 8,387,122.
- [63] S. Das, G. Laput, C. Harrison, and J. I. Hong, "Thumprint: Socially-inclusive local group authentication through shared secret knocks," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, pp. 3764–3774.
- [64] C. A. Masden, C. Grevet, R. E. Grinter, E. Gilbert, and W. K. Edwards, "Tensions in scaling-up community social media: a multi-neighborhood study of nextdoor," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2014, pp. 3239–3248.
- [65] L. Palen and P. Dourish, "Unpacking" privacy" for a networked world," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2003, pp. 129–136.
- [66] J. S. Olson, J. Grudin, and E. Horvitz, "A study of preferences for sharing and privacy," in *CHI'05 extended abstracts on Human factors in computing systems*, 2005, pp. 1985–1988.
- [67] T. Minkus, K. Liu, and K. W. Ross, "Children seen but not heard: When parents compromise children's online privacy," in *Proceedings of the 24th international conference on World Wide Web*, 2015, pp. 776–786.
- [68] A. Feal, P. Calciati, N. Vallina-Rodriguez, C. Troncoso, and A. Gorla, "Angel or devil? a privacy study of mobile parental control apps," *Proceedings of Privacy Enhancing Technologies (PoPETS)*, vol. 2020, 2020.
- [69] J. Wiese, P. G. Kelley, L. F. Cranor, L. Dabbish, J. I. Hong, and J. Zimmerman, "Are you close with me? are you nearby? investigating social groups, closeness, and willingness to share," in *Proceedings of the 13th international conference on Ubiquitous computing*, 2011, pp. 197–206.
- [70] R. Grinter and M. Eldridge, "Wan2tlk? everyday text messaging," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2003, pp. 441–448.
- [71] M. Sleeper, R. Balebako, S. Das, A. L. McConahy, J. Wiese, and L. F. Cranor, "The post that wasn't: exploring self-censorship on facebook," in *Proceedings of the 2013 conference on Computer supported cooperative work*, 2013, pp. 793–802.
- [72] M. Mondal, J. Messias, S. Ghosh, K. P. Gummadi, and A. Kate, "Forgetting in social media: Understanding and controlling longitudinal exposure of socially shared data," in *Twelfth Symposium on Usable Privacy and Security ({SOUPS} 2016)*, 2016, pp. 287–299.
- [73] C. Cobb, L. Simko, T. Kohno, and A. Hiniker, "User experiences with online status indicators," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–12.
- [74] A.-M. Olteanu, K. Huguenin, I. Dacosta, and J.-P. Hubaux, "Consensual and privacy-preserving sharing of multi-subject and interdependent data," in *Proceedings of the 25th Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2018, pp. 1–16.
- [75] A. Quan-Haase and I. Elueze, "Revisiting the privacy paradox: Concerns and protection strategies in the social media experiences of older adults," in *Proceedings of the 9th international conference on social media and society*, 2018, pp. 150–159.
- [76] A. Volda, R. E. Grinter, N. Ducheneaut, W. K. Edwards, and M. W. Newman, "Listening in: practices surrounding itunes music sharing," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2005, pp. 191–200.
- [77] J. Lindqvist, J. Cranshaw, J. Wiese, J. Hong, and J. Zimmerman, "I'm the mayor of my house: examining why people use foursquare—a social-driven location sharing application," in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2011, pp. 2409–2418.
- [78] C. Cobb and T. Kohno, "How public is my private life? privacy in online dating," in *Proceedings of the 26th International Conference on World Wide Web*, 2017, pp. 1231–1240.
- [79] G. Ranzini and C. Lutz, "Love at first swipe? explaining tinder self-presentation and motives," *Mobile Media & Communication*, vol. 5, no. 1, pp. 80–101, 2017.
- [80] M. Warner, A. Gutmann, M. A. Sasse, and A. Blandford, "Privacy unraveling around explicit hiv status disclosure fields in the online geosocial hookup app grindr," *Proceedings of the ACM on Human-Computer Interaction*, vol. 2, no. CSCW, pp. 1–22, 2018.
- [81] S. R. Peppet, "Unraveling privacy: The personal prospectus and the threat of a full-disclosure future," *Nw. UL Rev.*, vol. 105, p. 1153, 2011.
- [82] M. Minaei, S. C. Mouli, M. Mondal, B. Ribeiro, and A. Kate, "Deceptive deletions for protecting withdrawn posts on social media platforms," in *NDSS*, 2021.
- [83] A. M. Olteanu, M. Humbert, K. Huguenin, and J.-P. Hubaux, "The (co)-location sharing game," -, Tech. Rep., 2019.

- [84] L. Anticoli and M. Basaldella, "Shut up and run: the never-ending quest for social fitness," in *Companion Proceedings of the The Web Conference 2018*, 2018, pp. 1553–1556.
- [85] S. Gabriele and S. Chiasson, "Understanding fitness tracker users' security and privacy knowledge, attitudes and behaviours," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–12.
- [86] "Wearables help play a key role in the future of member engagement," Mar 2021. [Online]. Available: <https://www.uhc.com/employer/news/mid-sized-business/wearables-help-play-a-key-role-in-the-future-of-member-engagement>
- [87] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," Naval Research Lab Washington DC, Tech. Rep., 2004.
- [88] H. Nissenbaum and H. Daniel, "Trackmenot: Resisting surveillance in web search," 2009.
- [89] D. C. Howe and H. Nissenbaum, "Engineering privacy and protest: A case study of adnauseam."
- [90] S. Das, L. A. Dabbish, and J. I. Hong, "A typology of perceived triggers for end-user security and privacy behaviors," in *Fifteenth Symposium on Usable Privacy and Security ({SOUPS} 2019)*, 2019.
- [91] R. E. Grinter, W. K. Edwards, M. W. Newman, and N. Ducheneaut, "The work to make a home network work," in *ECSCW 2005*. Springer, 2005, pp. 469–488.
- [92] P. Kumar, S. M. Naik, U. R. Devkar, M. Chetty, T. L. Clegg, and J. Vitak, "'no telling passcodes out because they're private' understanding children's mental models of privacy and security online," *Proceedings of the ACM on Human-Computer Interaction*, vol. 1, no. CSCW, pp. 1–21, 2017.
- [93] P. Wisniewski, H. Jia, H. Xu, M. B. Rosson, and J. M. Carroll, "'preventative' vs. 'reactive' how parental mediation influences teens' social media privacy behaviors," in *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*, 2015, pp. 302–316.
- [94] S. Murthy, K. S. Bhat, S. Das, and N. Kumar, "Individually vulnerable, collectively safe: The security and privacy practices of households with older adults," *Proceedings of the ACM on Human-Computer Interaction*, no. CSCW, 2020.
- [95] J. Nicholson, L. Coventry, and P. Briggs, "'if it's important it will be a headline' cybersecurity information seeking in older adults," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–11.
- [96] P. DiGioia and P. Dourish, "Social navigation as a model for usable security," in *Proceedings of the 2005 symposium on Usable privacy and security*, 2005, pp. 101–108.
- [97] E. M. Redmiles, "'should i worry?' a cross-cultural examination of account security incident response," in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 920–934.
- [98] S. Das, A. D. Kramer, L. A. Dabbish, and J. I. Hong, "The role of social influence in security feature adoption," in *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*, 2015, pp. 1416–1426.
- [99] A. De Luca, S. Das, M. Ortlieb, I. Ion, and B. Laurie, "Expert and non-expert attitudes towards (secure) instant messaging," in *Twelfth Symposium on Usable Privacy and Security ({SOUPS} 2016)*, 2016, pp. 147–157.
- [100] J. Bonneau, J. Anderson, and L. Church, "Privacy suites: shared privacy for social networks," in *SOUPS*, vol. 9, 2009, pp. 1–2.
- [101] S. Das, A. D. Kramer, L. A. Dabbish, and J. I. Hong, "Increasing security sensitivity with social proof: A large-scale experimental confirmation," in *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, 2014, pp. 739–749.
- [102] S. Egelman, A. Sotirakopoulos, I. Muslukhov, K. Beznosov, and C. Herley, "Does my password go up to eleven? the impact of password meters on password selection," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2013, pp. 2379–2388.
- [103] P. Emami Naeini, M. Degeling, L. Bauer, R. Chow, L. F. Cranor, M. R. Haghighat, and H. Patterson, "The influence of friends and experts on privacy decision making in iot scenarios," *Proceedings of the ACM on Human-Computer Interaction*, vol. 2, no. CSCW, pp. 1–26, 2018.
- [104] H. Masaki, K. Shibata, S. Hoshino, T. Ishihama, N. Saito, and K. Yatani, "Exploring nudge designs to help adolescent sns users avoid privacy and safety threats," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–11.
- [105] B. Ur and Y. Wang, "A cross-cultural framework for protecting user privacy in online social media," in *Proceedings of the 22nd International Conference on World Wide Web*, 2013, pp. 755–762.
- [106] M. Kwasny, K. Caine, W. A. Rogers, and A. D. Fisk, "Privacy and technology: folk definitions and perspectives," in *CHI'08 Extended Abstracts on Human Factors in Computing Systems*, 2008, pp. 3291–3296.
- [107] C. Herley, "So long, and no thanks for the externalities: the rational rejection of security advice by users," in *Proceedings of the 2009 workshop on New security paradigms workshop*, 2009, pp. 133–144.
- [108] E. M. Redmiles, S. Kross, and M. L. Mazurek, "How i learned to be secure: a census-representative survey of security advice sources and behavior," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 666–677.
- [109] J. Goecks, W. K. Edwards, and E. D. Mynatt, "Challenges in supporting end-user privacy and security management with social navigation," in *Proceedings of the 5th Symposium on Usable Privacy and Security*, 2009, pp. 1–12.
- [110] S. Gaw, E. W. Felten, and P. Fernandez-Kelly, "Secrecy, flagging, and paranoia: adoption criteria in encrypted email," in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2006, pp. 591–600.
- [111] M. Ratto, "Critical making: Conceptual and material studies in technology and social life," *The information society*, vol. 27, no. 4, pp. 252–260, 2011.
- [112] P. Dourish, R. E. Grinter, J. D. De La Flor, and M. Joseph, "Security in the wild: user strategies for managing security as an everyday, practical problem," *Personal and Ubiquitous Computing*, vol. 8, no. 6, pp. 391–401, 2004.
- [113] J. I. Hong and J. A. Landay, "An architecture for privacy-sensitive ubiquitous computing," in *Proceedings of the 2nd international conference on Mobile systems, applications, and services*, 2004, pp. 177–189.
- [114] S. Lederer, J. I. Hong, A. K. Dey, and J. A. Landay, "Personal privacy through understanding and action: five pitfalls for designers," *Personal and ubiquitous computing*, vol. 8, no. 6, pp. 440–454, 2004.
- [115] C. Lampe, N. B. Ellison, and C. Steinfield, "Changes in use and perception of facebook," in *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, 2008, pp. 721–730.
- [116] C. S. De Souza and J. Preece, "A framework for analyzing and understanding online communities," *Interacting with computers*, vol. 16, no. 3, pp. 579–610, 2004.
- [117] S. Das, W. K. Edwards, D. Kennedy-Mayo, P. Swire, and Y. Wu, "Privacy for the people? exploring collective action as a mechanism to shift power to consumers in end-user privacy," *IEEE Security & Privacy*, vol. 19, no. 5, pp. 66–70, 2021.

APPENDIX A
CODEBOOK

Fig. 2: Codebook used in analysis of social cybersecurity work.

LABEL	DESCRIPTION	EXAMPLE WORK
advice	seeking or giving S&P advice	[3]
dating	related to online dating	[78]
demos	differences in behavior based on demographics	[97]
devices	sharing devices	[30]
emp-obs	an empirical or observational paper	[35]
end-share	ending account sharing, like at the end of a relationship	[34]
entertain	sharing netflix/spotify/news accounts	[30]
ext-percep	how others perceive your S&P behaviors	[45]
gen-diffs	young vs. old users	[17]
fitbit	fitbits, exercise data, health data	[85]
intended	using a technology as intended	[1]
intimate	romantic relationships	[32]
IPV	intimate partner violence	[37]
location	sharing location data	[83]
nudges	a top-down nudge	[101]
obf-self	hiding info about self through typical user tools (audience selection, disclosing less) from other users	[75]
obf-tech	hiding info about self from other users and institutions through less typical tools	[88]
parents	parental controls	[93]
passwords	contains password usage	[35]
personal	family and close friends	[94]
public	the public space	[81]
pw-sharing	sharing passwords	[41]
self-cens	refraining from sharing information about oneself	[71]
self-rep	related to managing self-representation	[17]
share-econ	ride/car/home-sharing	[54]
social	friends, acquaintances	[69]
social-auth	social authentication systems	[60]
spec-theory	a theory or speculative paper	[16]
system	a systems paper	[58]
teens	teenagers using the internet	[70]
trust	S&P behaviors influenced by trust in others	[33]
unintended	using a technology outside of its intended use	[76]
workplace	occurring in the workplace	[47]