Journal of Communication, 2026, **00**, 1–13
https://doi.org/10.1093/joc/jqag004
Original Article

International Communication Association OXFORD

Dataveillance inhibits legitimate communication: causal evidence for chilling effects

Noemi Festic[1,*], Moritz Büchi[2], Michael Latzer[1], Céline Odermatt[1]

[1]Department of Communication and Media Research, University of Zurich, Zurich, Switzerland
[2]Institute of Applied Media Studies, ZHAW Zurich University of Applied Sciences, Winterthur, Switzerland
*Corresponding author: Department of Communication and Media Research, University of Zurich, Andreasstrasse 15, 8050 Zurich, Switzerland.
Email: n.festic@ikmz.uzh.ch

Abstract

Chilling effects are an undesirable consequence of the datafication of communication: people may refrain from entirely lawful online activities due to fear of negative repercussions based on their digital traces. Our preregistered, longitudinal online field experiment randomly exposed 827 participants to accounts of dataveillance at 19 prompts across six weeks. Within individuals, this reduced comfort with online opinion sharing, information seeking, and information disclosure. Across individuals, greater total exposure was not robustly associated with lower comfort. Women, older, and less-educated participants reported lower overall comfort, with no evidence of sociodemographic moderation of chilling effects. Follow-up focus groups validated these results. Our study advances chilling effects theory and contributes causal evidence for dataveillance-induced self-censorship in a representative sample. The implications of this subtle relinquishing of autonomy are relevant for individual well-being and democratic participation in societies increasingly shaped by dataveillance.

Keywords chilling effects, dataveillance, self-censorship, surveillance, framed field experiment, online experiment, longitudinal, representative, quantitative, focus groups

Digital communication permeates and benefits most aspects of life. However, an undesirable side effect of the digitalization of everyday life in democratic societies arises when individuals do not feel comfortable or avoid using these technologies out of concern for future repercussions due to dataveillance of their behaviors. These ramifications of dataveillance—the automated, continuous, and unspecific collection, retention, and analysis of digital traces (Büchi et al., 2022)—are addressed by the theorization of chilling effects. This phenomenon has received increasing attention from various disciplines and was initially developed in the context of *political* discourse and free speech (Penney, 2016). Arguably, dataveillance's inhibiting effects on *everyday* digital communication are equally important, potentially hindering people's full utilization of digital infrastructures. Dataveillance by state and corporate actors relies on transnational data infrastructures that transcend political systems. Chilling effects are, thus, relevant beyond specific national contexts and are particularly consequential in democratic societies that depend on free and uninhibited communication (Penney, 2025).

To further theorize and assess the extent to which chilling effects pose a problem, a robust empirical test of the chilling-effects hypotheses is required but currently lacking. Research on chilling effects has predominantly focused on political communication (Stoycheff et al., 2019) or understood them as an undesired side effect of digital advertising (Strycharz & Segijn, 2023), yet chilling effects are theorized to affect more generally mundane online activities (Büchi et al., 2022). The literature also focuses strongly on active self-disclosure, concerning to what extent individuals share personal information when faced with dataveillance. However, chilling effects are theorized to also concern passive forms of online engagement (Stubenvoll & Binder, 2024). Methodologically, existing studies are either not able to fully establish a causal link between the feeling of being watched and subsequent behavioral modifications due to cross-sectional designs, or lack external validity and generalizability in the case of experiments with nonprobability samples for specific contexts. To fill this gap, we empirically test the chilling-effects hypotheses for digital communication for the general online population, i.e., answering the question: *Does stronger exposure to accounts of dataveillance lead to self-inhibited communication behavior?* Because chilling effects depend on the situation (Strycharz & Segijn, 2023), we disaggregate between- from within-person effects. Further, we address differences between societal groups. Traditionally disadvantaged groups use the internet less (Robinson et al., 2015) and reap fewer benefits from their online engagement (Scheerder et al., 2017); and simultaneously, their susceptibility to undesirable consequences such as chilling effects may be higher.

Received: December 5, 2024. **Revised:** March 19, 2026. **Accepted:** March 21, 2026

We employed a mixed-methods design: (1) A longitudinal online field experiment tested the core chilling-effects hypothesis with 827 participants representative of the German-speaking Swiss online population. It included a manipulation of the participants' sense of dataveillance via exposure to accounts of dataveillance and subsequent measurement of the self-inhibition of legitimate digital communication behaviors. (2) Since the validity of the experiment is paramount to the interpretation of the results beyond the study context, follow-up focus groups were conducted with 33 of the participants. Our study provides externally valid insights into the extent and predictors of longer-term chilling effects. The extension to within-person variation, a full range of everyday digital communication, a wide variety of issue contexts, and the integration with digital-inequality research significantly contributes to the theorization of the chilling-effects mechanism.

The chilling effects of digital dataveillance

Theoretical perspective

Traditional definitions of surveillance address individuals or physical spaces as subjects (Clarke, 1988). Given the unprecedented extent of data constantly generated on people's lives in digitized societies, any area of life can now be subjected to surveillance, expanding its overall impact. Therefore, this article uses the term dataveillance (van Dijck, 2014) as the antecedent of chilling effects, which has evolved to an infrastructural practice as it connects data from previously distinct spheres of life (Lyon, 2009). Understanding behavioral effects of dataveillance is integral to surveillance studies and online privacy research. In the latter, studies are typically concerned with privacy management, mostly understood as personal protection or self-disclosure (Baruh et al., 2017). While sharing similarities with self-disclosure literature, chilling effects represent a conceptually separate mechanism addressing the interplay between dataveillance perceptions and self-censorship. This can restrict the exercise of fundamental rights, creating a subtle, yet, cumulative threat to individual autonomy, well-being, and democratic engagement (Véliz, 2020).

Linking dataveillance and self-censored digital communication

In the probably first mention of the chilling-effects phenomenon by this name, White and Zimbardo (1975) proposed two social-psychological theories to explain how surveillance inhibits opinion expression: deindividuation and reactance. According to individuation-deindividuation theory, individuals seek individuation in supportive, but desire anonymity in threatening environments. Surveillance, understood as an "individuating act" (p. 4), makes people feel exposed and uncomfortable, prompting them to deindividuate by censoring their speech. In a digitized society, merely existing can be interpreted as continuously being in such a threatening environment due to pervasive dataveillance, prompting self-inhibition of free digital communication. In line with Bentham's (1995) panopticon analogy, it is secondary whether a person is in fact surveilled; what matters is that they sense they are. Accounting for the affective atmospheres of surveillance, the omnipresent system of dataveillance represents an "almost unnoticed aspect of everyday life […] experienced at the margins of consciousness" (Ellis et al., 2013: 716). The chilling-effects hypotheses capture one possible behavioral response to sensing this dataveillance: self-inhibition of legitimate digital communication.

Büchi et al. (2022) propose a theoretical framework for the individual-level chilling-effects mechanism: Dataveillance practices by states and corporations are made salient to individuals through their everyday online experiences or media reports. This can increase people's sense of dataveillance which can result in the self-inhibition of legitimate digital communication. To empirically test the core chilling-effects hypothesis, it is necessary to especially look at the relationship between exposure to dataveillance, perceptions of dataveillance practices, and the behavioral outcomes. This relationship is at the core of our study, both theoretically and empirically.

Generalizing chilling effects in the context of everyday digital communication

So far, chilling effects have mostly been studied for very specific contexts. The relevance of potentially inhibited digital communication goes beyond political participation; it also concerns the broader impact of the cumulation of undue deterrence from everyday acts of digital communication that individuals engage in to fulfill personal or social needs (Büchi et al., 2022).

Following Penney's (2021) understanding of chilling effects as social conformity, chilling effects quietly promote speech and behaviors that are more aligned with the established power structures in society. The spiral-of-silence theory (Noelle-Neumann, 1974) is often referenced to explain such online inhibition, yet the explanatory mechanisms associated with it differ from chilling effects: while studies relying on the spiral-of-silence theory focus on a fear of social isolation or social discomfort as the driver of online inhibition (Sakariassen & Meijer, 2021), the chilling-effects hypothesis adds, and isolates the effect of perceived surveillance. While this understanding initially included fears of government retribution (Stoycheff, 2016), our understanding of chilling effects also includes companies as dataveillance actors. Further, instead of focusing on overtly political or social activities, such as activism or public dissent, our understanding of chilling effects extends to more mundane communication. While such activities can be politicized, it is the context in which they take place and individuals' intentions (Theocharis, 2015) that distinguish them from political or social actions. We, thus, focus on experiences with dataveillance in ordinary, low-stakes contexts that are not understood or intended as primarily political.

Additionally, we adopt a broader definition of self-disclosure, encompassing any action that leaves a digital trace, such as entering search queries, clicking on links, downloading apps, or consuming online content. These activities contribute to the creation of a "datafied self" (Cheney-Lippold, 2017). Accordingly, any behavior that leaves digital traces, as judged by the individual, is viewed as a form of self-disclosure and is potentially subject to chilling effects.

Extending "*who* is affected?" by "*when* are people affected?"

To date, chilling effects have solely been studied at the between-person level: do people with a higher sense of dataveillance inhibit their communication more (than those with a lower sense of dataveillance)? However, chilling effects concern temporary behavioral changes that occur in response to specific circumstances, particularly perceived dataveillance (Murray et al., 2024): When do people inhibit their communication more (than they normally do)? Repeated experiences of chilling effects could potentially accumulate into longer-term behavioral changes (Büchi et al., 2022). However, this is still conceptualized as a dependent series of state-like responses rather than a stable trait. Chilling effects, therefore, concern both inter- and intra-individual change (Thomas et al., 2021). As recently demonstrated in digital well-being research (Beyens et al., 2024), mere population-level, between-person analyses overlook person-specific communication dynamics that can differ in both effect size and direction from population-level averages. Therefore, we disaggregate within- from between-person chilling effects, focusing on the former.

Empirical insights on chilling effects

Empirically studying chilling effects is a relatively recent endeavor, and existing research, while overall scarce, can be grouped into three categories: (1) studies relying on natural (mostly political) stimuli from exogenous events that increase the salience of dataveillance such as the NSA revelations, (2) experiments for specific (mostly political) contexts, or (3) surveys relying on self-reports. The key studies for each category are summarized in Table 1.

A few qualitative accounts of the phenomenon complement these quantitative studies, for instance on surveillance concerns about intelligent personal assistants and potential changes to behavior in the home (Mols et al., 2022) or on imaginaries of dataveillance and personal profiling (Kappeler et al., 2023; Lupton, 2020). Overall, we are lacking causal evidence for the chilling-effects hypothesis for a representative sample and in an internally and externally validated setting, allowing the results' generalizability to everyday digital communication.

Hypotheses

Informed by the chilling-effects framework (Büchi et al., 2022) and prior definitions (Penney, 2021; Stoycheff, 2016), we test the core hypothesis that exposure to surveillance induces a self-censoring mechanism: a decrease in free digital communication across a wide range of everyday topics such as personal health, religion, dating, political debates, or laws. We distinguish between three types of mundane, but relevant and legitimate digital communication behaviors: searching for information, sharing opinions, and disclosing personal information online.

Importantly, in formulating our hypotheses, we distinguish within-person from between-person chilling effects: a within-person chilling effect occurs when increases in an individual's exposure to accounts of dataveillance lead to decreases in that same individual's free digital communication. Empirically, this compares different measurements of the same person, asking whether individuals communicate less freely at times when their exposure is higher than their own typical level. A between-person chilling effect, by contrast, refers to differences between individuals. Such an effect would be present if individuals who are, on average, more exposed to accounts of dataveillance also engage in lower levels of free digital communication than individuals who are less exposed. Here, the reference point is other people in the sample, rather than different observations of the same person. Importantly, these two types of effects do not necessarily occur simultaneously. Within-person effects capture situational adjustments in communication when exposure changes, whereas between-person effects would indicate stable differences in communication behavior across individuals.

Accordingly, we expect these within-person chilling effects:

When their exposure to dataveillance is stronger (than it usually is) …
H1a: … participants feel less comfortable to search for information online.
H1b: … participants feel less comfortable to share their opinion online.
H1c: … participants feel less comfortable to disclose personal information online.

As an addition, the between-subject hypotheses for this article, which address the chilling-effects phenomenon on a global, population level, are the following:

Participants who are more strongly exposed to dataveillance (than other participants) …
H2a: … feel less comfortable to search for information online.
H2b: … feel less comfortable to share their opinion online.
H2c: … feel less comfortable to disclose personal information online.

Mixed-methods design to test the chilling-effects hypotheses

This section details the mixed-methods design, approved by the university's ethics review board. The design for the online field experiment was preregistered[1] (Festic et al., 2023) and all materials including the experimental stimuli are shared on OSF.[2] All data were collected in Switzerland, a democratic country with data protection laws closely aligned with the EU's General Data Protection Regulation (GDPR) (see Steiner, 2023). This study is part of a larger project on chilling effects.

Empirical Step 1: longitudinal online field experiment

The empirical core of this article is a longitudinal online field experiment to test the within- and between-person chilling-effects hypotheses. The study uses a *within-person randomized micro-intervention design* in which each participant is randomized independently at each of the 19 prompts, enabling the estimation

Table 1 Literature overview of key quantitative empirical studies on chilling effects in the context of dataveillance.

Study	Context	Methodological design	Data	Main results
1. Studies relying on natural stimuli from exogenously occurring events				
Penney (2016)	Chilling effects on accessing information online	Interrupted time series design, comparison of Wikipedia traffic to English-language, privacy-sensitive articles before and after the NSA revelations (June 2013)	View counts of 48 privacy-sensitive Wikipedia articles before and after the NSA revelations, $N > 81$ million page views	Significant reduction of traffic for privacy-sensitive articles after NSA revelations, first empirical evidence for chilling effects due to online surveillance revelations
Marthews and Tucker (2017)	Chilling effects on online search behavior	Natural experiment, comparison of Google searches before and after the NSA revelations (June 2013)	Google Trends: searches for privacy-sensitive, embarrassing, and neutral terms in 41 countries, $N = 523{,}340$ observations	Significant reduction in Google searches for personally and government-sensitive terms
Rosso et al. (2020)	Chilling effects as increase in use of privacy-protective search engine DuckDuckGo	Interrupted time series design: comparison of DuckDuckGo traffic before and after the NSA revelations (June 2013)	Weekly direct traffic of DuckDuckGo from Jan 2012 to Dec 2015	Strong and significant effect of the revelations: increase in DuckDuckGo search queries with a breakpoint in June 2013
Maaß et al. (2024)	Chilling effects as proportion of comments deleted, quantity and sentiment of comments	Comparison of posts and comments before and after the NetzDG (aimed to reduce hate speech on social media) law came into effect	$N > 33{,}000$ posts, > 7 million comments on ten public Facebook pages by German parties and media outlets	Slight increase in share of comments deleted per post in two months after law came into effect; no further indication of self-censorship (amount/tonality of comments)
Mak et al. (2024)	Chilling effects on political engagement on Facebook	Comparison of political engagement before and after national security law in Hong Kong came into effect	Two-wave panel survey data, $N = 504$ respondents from Hong Kong	No direct chilling effect on expression, but those opposing government had more fear of the law, which led to stronger privacy management and higher level of political engagement
2. Experiments for specific contexts				
Stoycheff (2023)	Political chilling effects through affective heuristics	Between-subjects online experiment, stimuli: online privacy policy cookie (surveillance) or content moderation (censorship)	$N = 540$ U.S. participants	Both website features activated negative affective responses, but only surveillance led to "problematic chilling effects"
Strycharz and Segijn (2024)	Chilling effects as ethical side effects of digital advertising	Cross-national survey and online experiment (United States and Netherlands), stimulus: different data collection methods	Survey ($N = 334$) and online experiment ($N = 536$)	Chilling effects in response to dataveillance from advertising were context-dependent: method of data collection for digital advertising made a difference and they were more prevalent in the United States
Stubenvoll and Binder (2024)	Chilling effects as a consequence of inducing declarative or	Online between-subjects experiment	$N = 500$ mostly female Austrian students	None of the knowledge treatments had a significant effect on individuals' behaviors, declarative group

(continued)

Table 1 (continued).

Study	Context	Methodological design	Data	Main results
	procedural privacy knowledge			experienced less chilling effects over time; heightening knowledge about data collection does not have side effect of inducing chilling effects
3. Surveys relying on self-reports				
Hampton (2014)	Chilling effects as "online" spiral of silence	Online survey on NSA revelations as a sensitive issue	$N = 1{,}801$ U.S. adults	People were more willing to discuss sensitive topic in person than on social media; and if they thought they shared the minority opinion they were less willing to talk about it in any setting
Latzer et al. (2020)	Self-reported chilling effects on digital communication behaviors due to dataveillance perceptions	Online survey	$N = 1{,}202$ Swiss internet users	Between 20 and 45% of respondents self-reported experiencing chilling effects for different online behaviors with little sociodemographic differences.
Strycharz et al. (2022)	Motivations for (not) changing media use in response to corporate surveillance	Thought-listing exercise and cross-sectional survey	Survey: $N = 156$ in Netherlands, $N = 148$ in United States	Older people perceived less threats and younger people were not willing to change their media use due to their dependency on a mobile device
Oz and Yanik (2024)	Chilling effects of government surveillance perceptions on willingness to express opinion on social media	Online survey	$N = 484$ Turkish X users	Negative relationship between perceived surveillance and willingness to speak out. While perceived majority did not play a role here, online privacy skills positively moderated this effect

of proximal treatment effects (Klasnja et al., 2015). The intensive longitudinal design (Hamaker & Wichers, 2017) allows to account for both the long-term nature and small hypothesized effect sizes of chilling effects (Büchi et al., 2022).

Data and sample

Participants were recruited from an internet-user panel maintained by *YouGov*. The field phase covered six weeks in September and October 2023. The gross sample ($N = 1{,}026$) was representative of German-speaking Swiss internet users aged 16 and over by gender, age, and education. To have enough measurements per person for the within-subject analyses, we only included those who completed at least 12 of the 19 prompts in our analyses. Panel attrition was low, resulting in a net sample of $N = 827$ exceeding our preregistered target of $N = 660$ based on power analysis. The sample was 51% female with an average age of 47 years (SD = 14.9). The distribution of participants' level of education was 5% low, 42% medium, and 53% high.[3] The sample remained representative after exclusion of missing data (<0.1% change for all categories).

Experimental design and procedure

The experiment relied on a one-factor (exposure to dataveillance vs. no exposure) repeated-measures design to manipulate participants' sense of dataveillance. We pretested two stimuli identified in an exploratory interview and a qualitative diary study preceding this experiment: (1) a notification on data usage and (2) a news article on a sensitive topic containing the dataveillance topic. The stimuli were independently pretested with 2,974 Swiss internet users, diverse in gender, age, and education. Results showed that the news articles were the driver of significant differences in a sense of dataveillance, while the notification alone had no additional significant effects. Based on these findings, we opted for a simplified one-factor design in the main experiment reported in this paper, comparing a condition with exposure to dataveillance in both the notification and news article to a control condition with no dataveillance in either element.

Participants received prompts on a mobile or desktop application, depending on their preferences, every three days during the 6-week field phase ($N_{prompts} = 19$). This setting—known as a *framed field experiment* (List, 2011)—involves conducting a structured

experiment in a natural environment, thereby minimizing potential biases associated with a laboratory setting. At each prompt, participants were repeatedly randomly assigned to receive a treatment stimulus with a probability of 70% (both notification and news article with dataveillance) or the control stimulus (both notification and news article without dataveillance), so each person functions as their own control. The unequal probability was chosen to ensure that by the end of the field phase, most participants had been exposed to a substantial amount of dataveillance stimuli (a separate control group at the participant level was not feasible). The prompts had the same structure each day with the notification and news article functioning as sequential components of the experimental manipulation, ensuring exposure to the dataveillance topic (or not):

(1) Participants were first shown a notification and second, a news article on a sensitive topic such as online dating, Netflix watch history, or health apps (see Figure 1 for full list). The notifications and news articles for the treatment and the control groups differed in whether dataveillance was made salient or not.[4] The news articles covered the same sensitive topic for both

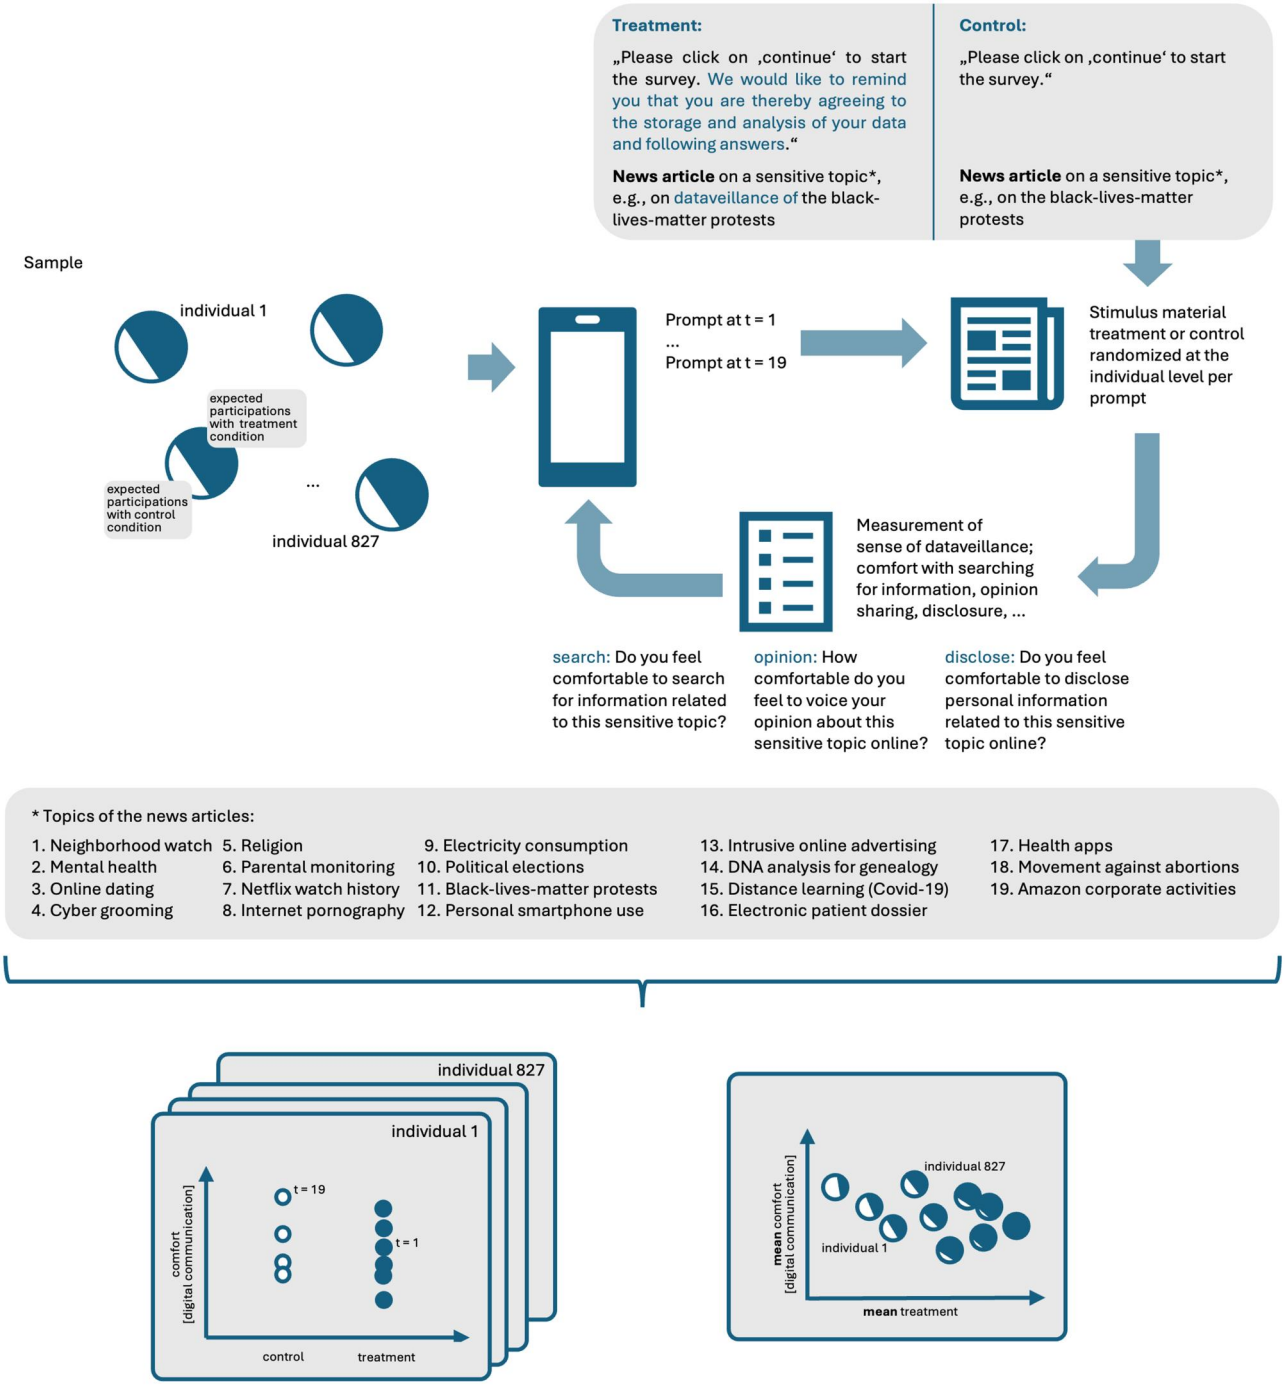

Figure 1 Experimental design: one-factor, repeated measures with randomized allocation to treatment and control conditions at each prompt.

groups. For example, a control article covered a new abortion ban in U.S. states, while the respective treatment article was on women deleting period-tracking apps due to concerns that their data can be used related to this ban. (2) At each prompt, participants then answered questions about their internet use, including four items jointly measuring their sense of dataveillance, which served as a manipulation check. Other items (e.g., on perceptions on internet use amount) were mainly included to conceal the study's purpose. At the end of the field phase, a final survey measured control variables. To conceal the research interest and avoid priming, participants were told the study was about online news consumption and were only debriefed after the field phase.

Measures

Treatment

This variable measured whether the participants were exposed to the treatment (1) or the control (0) condition in the experiment, i.e., whether they were exposed to dataveillance.

Sense of dataveillance

To ensure successful manipulation, we used a validated four-item scale (Segijn et al., 2022) with wording slightly adapted to the context of everyday digital communication (example item: "My every move is watched online"). Participants rated their agreement with the statements on a 5-point Likert scale (1 = do not agree at all, 5 = completely agree).

Comfort with searching for information online

Participants were asked whether they feel comfortable searching for sensitive information related to the news article's topic in their browser and if so, to paste their result and a corresponding link in an open answer field. We used the binary variable distinguishing between participants who said they did not feel comfortable (0) from those who did and proceeded to paste a link into the open answer field (1).

Comfort with opinion sharing online

Participants were asked how comfortable they would feel to share their opinion on the news article's sensitive topic on an online service that they like to use on a 5-point Likert scale (1 = not at all comfortable; 5 = totally comfortable).

Comfort with personal information disclosure online

Participants were asked to provide personal information related to the news article's sensitive topic (e.g., "Were you vaccinated against COVID?") with answer options: 1 = yes, 2 = no, 3 = I don't feel comfortable answering this question online. We recoded this to a binary variable and assigned the participants who chose either "yes" or "no" the value 1 (=feeling comfortable to disclose personal information), and those who chose "I don't feel comfortable answering this question online" value 0. This binary dependent variable was included for topics where it was possible, given ethical and privacy considerations.

Sociodemographic variables

Participants' gender[5] (1 = male, 2 = female), age (in years), and education level were measured. Education was recoded into two categories: up to secondary (primary or secondary school, vocational school, A-levels, or high-school), and tertiary (university, university of applied sciences).

Time

We account for the repeated-measures design by including a wave term in all regression-based analyses. This variable was numeric, ranging from 1 (first prompt) to 19 (last prompt) with intermediate values corresponding to the respective measurement occasion.

Data analysis

To test the hypotheses, we use statistical models for repeated measures and nested data. Mixed-effects (or multilevel) models explicitly model both individual variation and population-level relationships over time (Thomas et al., 2021). This design estimates causal effects by comparing treated and untreated moments within the same individual. Mixed-effects models with participant-level intercepts (and slopes for *comfort with searching for information* and *sharing opinions*) account for the repeated-measures structure and yield unbiased estimates of the proximal treatment effect. To isolate within-person effects, the treatment variable was modeled as a person-mean–centered, time-varying indicator capturing wave-specific deviations from each participant's average treatment exposure. Person-mean centering of time-varying categorical variables is a standard approach for separating within- and between-person components when treatment varies within individuals due to random assignment (Allison, 2009). In this specification, the within-person component represents whether a participant received the treatment at a given wave relative to their own average exposure rate, yielding an unbiased estimate of the causal within-person effect. In these within-person models, the corresponding between-person component of treatment exposure is mechanically determined by the randomization procedure, and therefore, not substantively interpretable. For a secondary between-person perspective, we produce a noncausal estimand by calculating each person's total average exposure and relating it to a person-level dependent variable (person means of the outcomes). Data analysis was conducted in *R* (*lme4* for mixed-effects models).

Empirical Step 2: follow-up focus-group interviews

Considering the online field experiment's methodological limitations, follow-up focus-group interviews were conducted with selected participants of the experiment to validate the results and align interpretations with participants' sense-making as affected media users. The guiding research questions were: *How do the results of the experiment reflect the participants' behavior in their everyday lives? Did the participants engage with the study material as intended?*

Data and sample

Upon designing the study, we intended to organize focus groups based on the experimental results, grouping participants with similar experiences in the experiment: those who do vs. do not experience chilling effects. Based on bivariate correlations between treatment conditions and outcome variables, the group

who did not experience chilling effects was nonexistent, so we changed the grouping condition to form groups with less vs. more affected participants. Based on their personal results in the experiment, the participants for both grouping conditions were individually selected and invited to specific focus groups, which reduces potential self-selection biases. The total sample consisted of 33 participants, balanced by gender, age, and education. We conducted three focus groups with more and two with less affected participants.

Procedure

The focus groups were held in person in April/May 2024 at our university and lasted two hours each. Three research team members were present: #1 as neutral discussion leader, #2 as expert on the online framed field experiment study and the topic of dataveillance, presenting the study-related discussion inputs and responding to questions, #3 for quality control of the discussion and administrative tasks. The moderators fostered an open, accepting environment by establishing ground rules and ensuring anonymity before the discussions. During the sessions, we maintained this environment by posing open and nonsuggestive questions (Krueger & Casey, 2015).

Data analysis

The focus groups were audio recorded and transcribed (automatic transcription in *Whisper*; manual correction and speaker identification). We replaced names with pseudonyms. The transcripts were analyzed in *MAXQDA* using deductive qualitative content analysis to identify themes and patterns related to the research questions. The codes and their interpretation were continuously discussed by the members of the research team.

Results

Quantitative results from the longitudinal online framed field experiment

Manipulation check

The aim of the stimulus material was to expose participants to dataveillance and increase their sense of dataveillance. This multi-item construct was examined using confirmatory factor analysis to assess dimensionality, which supported a single-factor structure with high standardized loadings (.84–.91, all $p <$ .001), $\chi^2(2, N = 14{,}436) = 15.98$. Because the items were measured repeatedly within individuals, the CFA served only to establish one-dimensionality. Measurement reliability was, therefore, assessed using multilevel composite reliability, yielding excellent between-person reliability ($\omega_{between}$ = .99) and acceptable within-person reliability (ω_{within} = .66). To determine whether the stimulus material effectively heightened participants' sense of dataveillance (regression-based factor scores rescaled to the original metric) while accounting for the nested structure of the data, we estimated a linear mixed-effects model with sense of dataveillance as the dependent variable, fixed effects for treatment condition, linear and quadratic time, and a random intercept for participants. The analysis revealed a significant main effect of treatment, $b = 0.04$, $SE = 0.01$, $t(13{,}630) = 6.50$, $p < .001$, 95% CI [0.03, 0.06], indicating that when in a treatment prompt ($M = 3.77$, $SD = 0.87$, $n = 10{,}255$) participants reported a significantly higher sense of dataveillance than when in the control condition ($M = 3.69$, $SD = 0.88$, $n = 4{,}427$). Sense of dataveillance also increased over time ($b = 0.03$, $SE = 0.002$, 95% CI [0.03, 0.04]), at a decreasing rate ($b = -0.0011$, $SE = 0.0001$, 95% CI [−0.0014, −0.0009]). The model-based intraclass correlation coefficient[6] (ICC = .82) indicated that a substantial proportion of the variance in sense of dataveillance was attributable to between-person differences. We further assessed whether the manipulation was successful for each of the 19 topics. Since each participant only responded to the questions about each topic once, independent samples t-tests were conducted here. Cyber grooming and household electricity consumption did *not* show higher means in the stimulus vs. the control condition and were excluded from the analysis. All further results in this article rely on the data of the remaining 17 topics for which the mean sense of dataveillance was higher in the treatment condition.

We now present the results of the linear mixed models for each dependent variable to address the within-person (H1a–H1c) hypotheses and results of linear regressions to address the between-person hypotheses (H2a–H2c).

Results from mixed models: testing the within-person chilling-effects hypotheses

To test the within-person hypotheses, we estimated mixed-effects models that isolated the causal effect of exposure to dataveillance while accounting for time trends (centered wave) and sociodemographic characteristics (age [z-standardized], gender, and education). All models included random intercepts to capture stable between-person differences; for comfort with opinion sharing and comfort with searching for information online, models additionally included random slopes for treatment, allowing the causal effect of dataveillance to vary across individuals.[7]

Within persons, exposure to dataveillance caused a reduction in comfort with sharing opinions, $b = -0.17$, $SE = 0.02$, $t = -8.39$, $p < .001$, 95% CI [−0.21, −0.13]. Comfort increased across measurement occasions, $b = 0.0127$, $SE = 0.0015$, $t = 8.22$, $p < .001$, 95% CI [0.0097, 0.0158]. Women reported lower comfort than men, $b = -0.36$, $SE = 0.06$, $t = -6.05$, $p < .001$, 95% CI [−0.48, −0.25], comfort declined with age, $b = -0.18$, $SE = 0.03$, $t = -5.85$, $p < .001$, 95% CI [−0.23, −0.12], and increased with tertiary education, $b = 0.18$, $SE = 0.06$, $t = 3.02$, $p = .003$, 95% CI [0.06, 0.30]. Baseline comfort differed substantially across individuals (random-intercept SD = 0.82), with participant-level clustering (model-based ICC = .41). Allowing the treatment effect to vary across individuals indicated heterogeneity in chilling effects (random-slope SD = 0.17), implying that individual-specific treatment effects plausibly ranged from approximately −0.50 to +0.17 around the average effect; the intercept–slope correlation was small ($r = -.06$).

A comparable chilling effect was observed for comfort with searching for information online. Exposure to dataveillance reduced the likelihood of reporting comfort with searching, $b = -0.14$, $SE = 0.05$, $z = -2.96$, $p = .003$, 95% CI [−0.24, −0.05] (OR = 0.87). Women were less likely than men to report comfort ($b = -0.73$, $SE = 0.13$, $z = -5.59$, $p < .001$), comfort declined with age ($b = -0.24$, $SE = 0.06$, $z = -3.71$, $p < .001$), and increased with tertiary education ($b = 0.73$, $SE = 0.13$, $z = 5.64$, $p < .001$). Time

trends were not statistically significant ($b = -0.0043$, $SE = 0.0038$, $z = -1.13$, $p = .26$). Participant-level clustering remained substantial (ICC = .47).

For disclosing personal information online, exposure to dataveillance also reduced the likelihood of reporting comfort, $b = -0.26$, $SE = 0.10$, $z = -2.68$, $p = .007$, 95% CI [−0.44, −0.07] (OR = 0.77). Comfort with disclosure decreased over time ($b = -0.0231$, $SE = 0.0075$, $z = -3.07$, $p = .002$), whereas age, gender, and education were not statistically significant predictors. Baseline differences between participants were sizeable (ICC = .39).

Figure 2 depicts these within-person treatment effects of dataveillance on the three outcomes. The left panel shows estimated treatment coefficients (log-odds for the binary outcomes search and disclose). The right panel displays predicted outcomes for the control and treatment conditions, shown on the original scale for opinion and on a probability scale for the binary outcomes.

Taken together, the within-person analyses provide causal evidence that exposure to dataveillance produces chilling effects for comfort with searching for information online (H1a supported), sharing opinions online (H1b supported), and disclosing personal information online (H1c supported). Across outcomes, a substantial proportion of variance was attributable to stable between-person differences, underscoring the importance of accounting for individual heterogeneity when estimating effects.

Results from linear regressions: testing between-person associations

We estimated between-person associations using ordinary least squares (OLS) regression with person-level mean outcomes. Exposure was operationalized as the proportion of measurement occasions assigned to the dataveillance condition; its distribution was approximately symmetric, with a mode around 13 treated occasions, consistent with the 70% treatment probability across the 19 waves. All models included gender, age, and education as covariates.

Greater exposure to dataveillance was not significantly associated with comfort with sharing opinions, $b = -0.42$, $SE = 0.26$, $t(818) = -1.59$, $p = .11$, $R^2 = .09$. Women reported lower comfort than men ($b = -0.36$, $p < .001$), comfort decreased with age ($b = -0.01$, $p < .001$), and increased with higher education ($b = 0.18$, $p = .003$). Greater dataveillance exposure was associated with lower comfort with searching for information, $b = -0.19$, $SE = 0.09$, $t(819) = -2.16$, $p = .031$, $R^2 = .09$. Women reported lower comfort than men ($b = -0.11$, $p < .001$), comfort declined with age ($b = -0.003$, $p < .001$), and increased with education ($b = 0.11$, $p < .001$). For comfort with disclosing personal information, exposure was not significantly associated with comfort, $b = -0.03$, $SE = 0.03$, $t(819) = -0.87$, $p = .39$, with minimal explained variance ($R^2 = .01$). Being female ($b = -0.01$, $p = .051$), age ($b = .0001$, $p = .61$), and education ($b = .0004$, $p = .96$) were not associated with disclosure comfort.

Regarding the between-person hypotheses, greater cumulative exposure to dataveillance was associated with lower comfort with searching for information (H2a supported), but was not significantly associated with comfort with sharing opinions (H2b not supported) or with comfort with disclosing personal information online (H2c not supported).

Differential susceptibility to chilling effects: exploratory analyses

We further examined whether susceptibility to chilling effects varied by sociodemographic characteristics. At the between-person level, models including interactions between cumulative exposure to dataveillance and age, gender, and education yielded no evidence of moderation for opinion sharing, information searching, or personal information disclosure (all interaction terms not statistically significant). At the within-person level, moderation analyses similarly showed no significant interactions for opinion sharing. For the binary outcomes (search and disclosure), models including treatment × sociodemographic interactions did not converge, precluding firm conclusions about within-person moderation for these outcomes. Overall, the results provide no consistent evidence that chilling

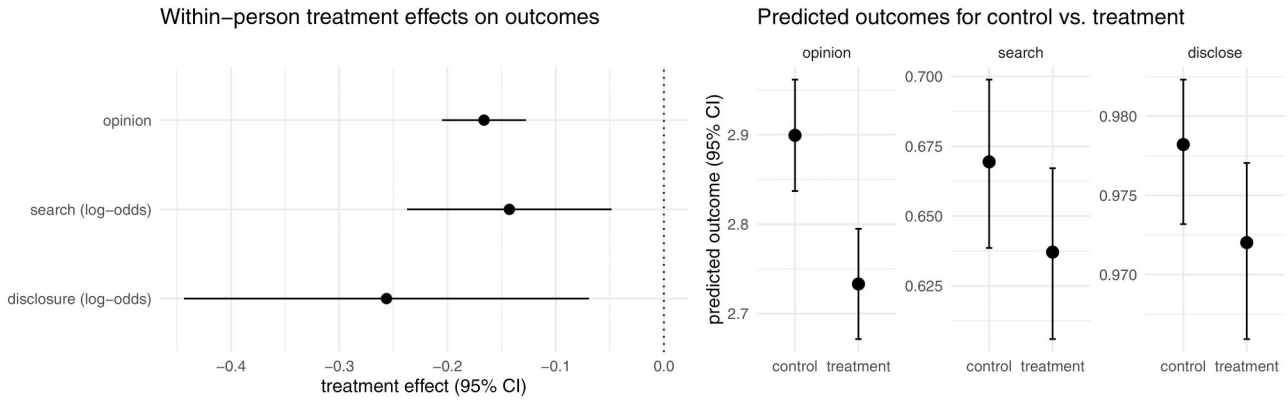

Figure 2 Within-person treatment effects and predicted outcomes. *Note*. The continuous outcome (*opinion*) was modeled using a linear mixed-effects model and is plotted on the original scale. The binary outcomes (*search*, *disclose*) were modeled using logistic mixed-effects models. In the left panel, points represent the estimated within-person treatment coefficients from the respective models (i.e., log-odds coefficients for *search* and *disclose*). Confidence intervals exclude zero. In the right panel, predicted outcomes for the control and treatment conditions are shown on the probability scale for the binary outcomes, obtained via the inverse logit transformation. Apparent overlap of confidence intervals in the right panel reflects uncertainty in the predicted means, not the effect; the treatment effects are estimated as within-model contrasts.

effects systematically differ by age, gender, or education. Further, none of the specific topics were significant predictors in any of the models.

Validity check from the qualitative focus-group interviews

Considerations on the online field experiment's *internal validity* based on the focus group results revealed no concerns. Potential confounding variables were well-controlled, and no significant threats to the cause-and-effect relationship between the experimental manipulations and the outcomes were identified. The cover story (i.e., study on news consumption) was successful, as participants did not discern the study's true objective. Although some speculated about the study goal, their guesses revolved around their ability to locate information online or the time required to find information. Further, the experimental manipulations were effectively concealed, as participants largely did not notice any differences between the notifications and news articles in the control and treatment conditions. Those who reported noticing differences did not consider them significant or systematic, and they reported continuing their participation in the study without altering their behavior. Nora (f, 24) reflected on her observation of the notification in the treatment condition one day: "I only noticed once that the text at the bottom [of the notification] was different. I just noticed it once. [And thought:] Was it already there yesterday? (…) And I didn't really think about it either." When asked to reflect on their study participation, it became clear that participants interacted with the study material as intended. They completed tasks responsibly and adhered to instructions. Aside from minor issues (fatigue due to repeated experimental measures, occasional technical difficulties when moving from app to browser), no significant issues that would impact our interpretation of the experiment results were mentioned.

Based on the focus-group results, we can further conclude that the *external validity* of the experimental results is high. Participants mostly stated that they behaved as they would in their typical everyday life. The participants were shown their own, person-level results of the experiment and most of them were not surprised because they aligned with their self-perception of their communication behaviors. Gabriel (m, 52) emphasized this similarity. Other participants pointed out discrepancies between their personal results and perceived everyday-life behavior. They perceived their self-inhibition in everyday life as higher than in the study, and attributed this to the specific study context because their trust in the university conducting the research was high and they have given consent to their participation. These results indicate that the experimental results may be generalized beyond the specific context of the experiment, and that we likely *underestimate* the extent of chilling effects since the participants, if at all, only mentioned differences in one direction: they believe they inhibit their digital communication behavior more strongly in their everyday lives as compared to the study environment that they perceived to be safer.

Discussion

This article advances the theorization of the chilling effects of dataveillance—a largely unintended side effect of digitized everyday lives in democratic societies. With roots in psychology (White & Zimbardo, 1975) and legal studies (Solove, 2006), the chilling-effects hypothesis has found entry into communication science where it has mostly been studied with a focus on political communication (Stoycheff, 2016) or digital advertising (Strycharz et al., 2022). We synthesize these interdisciplinary insights and build upon a systematic theoretical model of causal mechanisms, employing a generalized perspective. Specifically, we empirically establish chilling effects as a threat to everyday digital communication behaviors across various contexts (Büchi et al., 2022). We provide the first causal evidence of chilling effects in ordinary, low-stakes digital communication settings, derived from a longitudinal and representative online framed field experiment and validated through a follow-up focus-group study. For our focus on intra-individual change and the hypothesized small effect sizes, we extended the usual focus on *who* is affected (between-person level) to include *when* people are affected (within-person level).

The quantitative results generally support the existence of chilling effects. To address the initial research question ("Does stronger exposure to accounts of dataveillance lead to self-inhibited communication?"), the answer is largely yes—the within-person analyses provide causal evidence for chilling effects, as they leverage experimental variation in exposure to the dataveillance topic. When individuals were exposed to accounts of dataveillance, they became less comfortable sharing opinions, searching for information, and disclosing personal information online. Thus, chilling effects extend beyond overt opinion expression and include more passive forms of digital engagement. The between-person analyses complement this causal evidence by showing that greater cumulative exposure to dataveillance was associated with lower overall comfort with information seeking. Together, these results suggest that chilling effects operate mainly as immediate, situational responses to dataveillance, but at least for information seeking, also as longer-term differences in communication comfort that accumulate with repeated exposure. Random intercepts were significant across models, highlighting meaningful variation in baseline communication comfort.

The divergence between within- and between-person findings provides additional insight into the dynamics of chilling effects. The within-person results indicate that individuals adjust their communication behavior when their exposure to accounts of dataveillance increases relative to their own typical level, suggesting that chilling effects operate as situational responses to perceived surveillance. In contrast, the weaker and more selective between-person patterns indicate that these reactions do not necessarily translate into broad or stable differences in communication behavior across individuals. These findings highlight the value of within-person analyses for studying chilling effects, as such approaches capture dynamic adjustments in communication behavior that may remain invisible in designs focusing solely on differences between individuals.

To explore digital inequalities in the context of chilling effects, we examined whether exposure to dataveillance and comfort with digital communication varied by age, education, and gender. Contrary to expectations of differential susceptibility, we found no consistent evidence that chilling effects systematically differed across sociodemographic groups. Instead, sociodemographic characteristics were primarily associated with overall

levels of comfort: women, older participants, and less-educated individuals reported lower baseline comfort with several forms of digital communication, independent of experimental exposure. Situated within scholarship on digital inequalities (Robinson et al., 2015) and differential susceptibility to media effects (Valkenburg & Peter, 2013), these findings suggest that disparities in digital participation may stem less from heightened responsiveness to dataveillance cues than from preexisting differences in comfort and engagement. Thus, while dataveillance exerts chilling effects across groups, it does not appear to disproportionately amplify inequalities by differentially affecting specific sociodemographic populations. At the same time, reduced engagement in certain online behaviors—particularly those involving the disclosure of personal information—may reflect a considered response to perceived risks rather than disengagement *per se*, underscoring the importance of distinguishing between protective restraint and chilling effects in future research.

Additionally, future research should address a possible societal polarization process, which, over time, could lead to a decrease in diversity and representation in the digital world. This shift could mean that online spaces may increasingly reflect the perspectives of those who are more vocal, marginalizing quieter voices and narrowing the spectrum of views present in digital discourse (see Noelle-Neumann, 1974).

The chilling effects we found in our analyses were small, consistent with the theory describing a subtle and long-term process (Büchi et al., 2022). The practical relevance of these effects warrants further attention, yet, was evidenced by qualitative accounts in the focus-group interviews. The focus-group results also suggest that the validity of our experimental results is likely high, and we find indications for underestimations of the scope of chilling effects: all participants reported adjusting their free digital communication. Another reason why self-censorship of digital communication is likely higher in the real world was that participants reported trusting the university conducting the study and felt comfortable to communicate. Dataveillance seems to be engrained into people's considerations on what (not) to do online. Additionally, merely by living in digitized societies, participants of our experiment were already exposed to constant dataveillance—a kind of baseline treatment or "threatening" environment. Our experiment solely measured the additional treatment effect from exposure to the stimulus, with the real difference between people's intended levels of free digital communication behaviors and their actual behaviors accounting for their sense of dataveillance likely being more pronounced.

From a policy perspective, especially the focus-group results align with chilling effects as a constant, yet, often subconscious (see Ellis et al., 2013) companion in individuals' digital communication deliberations. Since highly educated individuals seem more comfortable with free digital communication, increasing dataveillance knowledge as a form of governance by information without inducing chilling effects as a side effect (Stubenvoll & Binder, 2024) could be a promising way to reduce self-censored behaviors in digitized societies.

Our study offers some directions for future research. According to the focus-group results and in line with the theory of planned behavior (Ajzen, 1991), we find that chilling effects sometimes manifest subconsciously, as second thoughts, or slight adaptations that are difficult to measure. These nuances of self-inhibited behaviors require further investigation. The effects of mediating variables such as privacy management, personality characteristics or different dataveillance sources on the chilling effects of dataveillance require further attention.

The results of the online field experiment and focus groups revealed rare instances of a "warming effect": an *increase* in comfort with engaging in the behaviors as a result to being exposed to dataveillance. Considering that responses to the threatening dataveillance environment are cognitive and affective (Ratcliff, 2021), the theory of psychological reactance (Brehm, 1966) offers a possible explanation: individuals with certain behavioral freedoms might be motivated to regain them upon threats to these freedoms to restore their agency. Affective or cognitive reactance could mediate the relationship between a sense of dataveillance and self-inhibition.

The following study design considerations warrant attention in interpreting the findings: Regarding external validity, the news article stimuli were based on real news articles that we shortened, slightly rephrased, or to which we (in the treatment case) added the dataveillance theme. This ensured high external validity of the stimuli. Additionally, we validated our results with follow-up focus-group interviews where we asked whether the news articles and dataveillance notifications felt natural and were similar to what participants encountered during typical everyday internet use. However, the manipulation did not work for two topics (grooming and household energy consumption), so data from these topics were excluded from further analysis. Also, the experimental results maintain a certain artificiality that needs to be considered when interpreting the results. Further, our data was collected in a democratic country and cannot be generalized to other political contexts, where inhibition of free communication is not seen as a largely unintended consequence of dataveillance, and chilling effects might unfold differently.

This article shows that "living digitally" carries an implied threat. In a democratic society, individuals should feel free to share opinions, search for information, and disclose personal information, even if related to sensitive topics. In a time where most of these behaviors are mediated through digital media, the constant exposure to dataveillance and diffuse fears of potential repercussions deter individuals from these entirely legitimate behaviors. On an individual level, experiencing this impairment could decrease autonomy and well-being. On a societal level, chilling effects are problematic for democratic processes in terms of participation, opinion formation, and access to information. On both levels, challenges arise for the governance of dataveillance practices and polarization.

Notes

1. The preregistration can be found here: https://doi.org/10.17605/OSF.IO/JR2Y7
2. The OSF page can be found here: https://osf.io/4ma8k
3. Due to the small low-education subsample (n = 38), low and medium education were combined in the multivariate analyses, yielding two categories: "up to secondary education" and "tertiary education."
4. See OSF (https://osf.io/4ma8k) for stimulus materials.

5. Since YouGov collected basic sociodemographic variables of the participants prior to our field phase, it was not possible to include a third category for gender.
6. Intraclass correlation coefficients (ICCs) were calculated from variance components of the mixed-effects models as the ratio of between-participant variance to total variance. The intercept-only model yields the unconditional ICC, whereas models including predictors yield conditional ICCs. For linear models, residual variance was estimated as σ^2; for logistic models, ICCs were computed on the latent scale using the standard logistic residual variance ($\pi^2/3$).
7. Each model was first estimated with random intercepts, and then, with both random intercepts and slopes. Allowing the treatment effect to vary across individuals improved model fit for opinion sharing and information searching, as indicated by successful convergence and nonzero random-slope variance estimates. In contrast, models including random slopes for personal information disclosure did not converge and produced boundary solutions, so a random-intercept–only model was retained (see https://osf.io/4ma8k for respective outputs).

Data availability

All shareable components of the empirical process (preregistration, analysis code, outputs, and materials) are available on OSF (https://osf.io/4ma8k).

Funding

This project received funding from the Swiss National Science Foundation (grant number: 201176).

Conflicts of interest

None declared.

Acknowledgments

We thank Kiran Kappeler for her support in preparing the field experiment, Sarah Daoust-Braun and Sophia Tan for their assistance with the focus group interviews, and Reto Bürgin for his guidance on data analysis. We also thank the anonymous reviewers for their constructive comments.

References

Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, *50*, 179–211. https://doi.org/10.1016/0749-5978(91)90020-T

Allison, P. D. (2009). *Fixed effects regression models*. SAGE.

Baruh, L., Secinti, E., & Cemalcilar, Z. (2017). Online privacy concerns and privacy management. *Journal of Communication*, *67*, 26–53. https://doi.org/10.1111/jcom.12276

Bentham, J. (1995). *The panopticon writings, 1748-1832*. Verso.

Beyens, I., Pouwels, J. L., van Driel, I. I., Keijsers, L., & Valkenburg, P. M. (2024). Social media use and adolescents' well-being: Developing a typology of person-specific effect patterns. *Communication Research*, *51*, 691–716. https://doi.org/10.1177/00936502211038196

Brehm, J. W. (1966). *A theory of psychological reactance*. Academic Press.

Büchi, M., Festic, N., & Latzer, M. (2022). The chilling effects of digital dataveillance. *Big Data & Society*, *9*. https://doi.org/10.1177/20539517211065368

Cheney-Lippold, J. (2017). *We are data: Algorithms and the making of our digital selves*. NYU Press.

Clarke, R. (1988). Information technology and dataveillance. *Communications of the ACM*, *31*, 498–512. https://doi.org/10.1145/42411.42413

Ellis, D., Tucker, I., & Harper, D. (2013). The affective atmospheres of surveillance. *Theory & Psychology*, *23*, 716–731. https://doi.org/10.1177/0959354313496604

Festic, N., Latzer, M., Büchi, M., Kappeler, K., & Odermatt, C. (2023, November 24). The Chilling Effects of Dataveillance: A Longitudinal Field Experiment. https://doi.org/10.17605/OSF.IO/JR2Y7

Hamaker, E. L., & Wichers, M. (2017). No time like the present: Discovering the hidden dynamics in intensive longitudinal data. *Current Directions in Psychological Science*, *26*, 10–15. https://doi.org/10.1177/0963721416666518

Hampton, K, Rainie, L, Lu, W, Dwyer, M, Shin, I., & Purcell, K. (2014, August 26). *Social media and the 'spiral of silence.'* Pew Research Center. https://www.pewresearch.org/internet/2014/08/26/social-media-and-the-spiral-of-silence/

Kappeler, K., Festic, N., & Latzer, M. (2023). Dataveillance imaginaries and their role in chilling effects online. *International Journal of Human-Computer Studies*, *179*, 103120. https://doi.org/10.1016/j.ijhcs.2023.103120

Klasnja, P., Hekler, E. B., Shiffman, S., Boruvka, A., Almirall, D., Tewari, A., & Murphy, S. A. (2015). Microrandomized trials: An experimental design for developing just-in-time adaptive interventions. *Health Psychology*, *34S*, 1220–1228. https://doi.org/10.1037/hea0000305

Krueger, R. A., & Casey, M. A. (2015). *Focus group: A practical guide for applied research* (5th ed.). SAGE.

Latzer, M., Festic, N., & Kappeler, K. (2020). *Coping practices related to algorithmic selection in Switzerland.* University of Zurich. https://mediachange.ch/media//pdf/publications/Report_4_Practices.pdf

List, J. A. (2011). Why economists should conduct field experiments and 14 tips for pulling one off. *Journal of Economic Perspectives*, *25*, 3–16. www.doi.org/10.1257/jep.25.3.3

Lupton, D. (2020). Thinking with care about personal data profiling: A more-than-human approach. *International Journal of Communication*, *14, 3165–3183.*

Lyon, D. (2009). Surveillance, power, and everyday life. In C. Avgerou, R. Mansell, D. Quah, & R. Silverstone (Eds.), *The Oxford handbook of information and communication technologies*. Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199548798.003.0019

Maaß, S., Wortelker, J., & Rott, A. (2024). Evaluating the regulation of social media. *Telecommunications Policy*, *48*, 102719. https://doi.org/10.1016/j.telpol.2024.102719

Mak, M. K., Koo, A. Z.-X., & Rojas, H. (2024). Social media engagement against fear of restrictions and surveillance. *New Media & Society*, *26*, 1984–2005. https://doi.org/10.1177/14614448221077240

Marthews, A., & Tucker, C. E. (2017). *Government surveillance and internet search behavior*. https://ssrn.com/abstract=2412564 or http://dx.doi.org/10.2139/ssrn.2412564

Mols, A., Wang, Y., & Pridmore, J. (2022). Household intelligent personal assistants in the Netherlands. *Convergence*, *28*, 1841–1860. https://doi.org/10.1177/13548565211042234

Murray, D., Fussey, P., Hove, K., Wakabi, W., Kimumwe, P., Saki, O., & Stevens, A. (2024). The chilling effects of surveillance and human rights. *Journal of Human Rights Practice*, *16*, 397–412. https://doi.org/10.1093/jhuman/huad020

Noelle-Neumann, E. (1974). The spiral of silence: A theory of public opinion. *Journal of Communication*, *24*, 43–51. https://doi.org/10.1111/j.1460-2466.1974.tb00367.x

Oz, M., & Yanik, A. (2024). Fear of surveillance. *Mediterranean Politics*, *29*, 1–25. https://doi.org/10.1080/13629395.2022.2046911

Penney, J. (2016). Chilling effects: Online surveillance and Wikipedia use. *Berkeley Technology Law Journal*, *31*, 117–182. https://doi.org/10.15779/Z38SS13

Penney, J. (2021). Understanding chilling effects. *SSRN Electronic Journal*, 1451–1530. https://dx.doi.org/10.2139/ssrn.3855619

Penney, J. (2025). *Chilling effects: Repression, conformity, and power in the digital age*. Cambridge University Press. https://doi.org/10.1017/9781108641784

Ratcliff, C. L. (2021). Characterizing reactance in communication research. *Communication Research*, *48*, 1033–1058. https://doi.org/10.1177/0093650219872126

Robinson, L., Cotten, S. R., Ono, H., Quan-Haase, A., Mesch, G., Chen, W., … Stern, M. J. (2015). Digital inequalities and why they matter. *Information, Communication & Society*, *18*, 569–582. https://doi.org/10.1080/1369118X.2015.1012532

Rosso, M., Nasir, A., & Farhadloo, M. (2020). Chilling effects and the stock market response to the Snowden revelations. *New Media & Society*, *22*, 1976–1995. https://doi.org/10.1177/1461444820924619

Sakariassen, H., & Meijer, I. C. (2021). Why so quiet? Exploring inhibition in digital public spaces. *European Journal of Communication*, *36*, 494–510. https://doi.org/10.1177/02673231211017346

Scheerder, A., van Deursen, A., & van Dijk, J. (2017). Determinants of internet skills, uses and outcomes. A systematic review of the second- and third-level digital divide. *Telematics and Informatics*, *34*, 1607–1624. https://doi.org/10.1016/j.tele.2017.07.007

Segijn, C. M., Opree, S. J., & Ooijen, I. V (2022). The validation of the Perceived Surveillance Scale. *Cyberpsychology*, *16*, Article 9. https://doi.org/10.5817/CP2022-3-9

Solove, D. (2006). A taxonomy of privacy. *University of Pennsylvania Law Review*, *154*, 477.

Steiner, T. (2023). Zwischen Autonomie und Angleichung. In M. Widmer (Ed.), *Datenschutz* (pp. 51–138). DIKE Verlag. https://doi.org/10.3256/978-3-03929-027-7

Stoycheff, E. (2016). Under surveillance: Examining Facebook's spiral of silence effects in the wake of NSA internet monitoring. *Journalism & Mass Communication Quarterly*, *93*, 296–311. https://doi.org/10.1177/1077699016630255

Stoycheff, E. (2023). Cookies and content moderation: Affective chilling effects of internet surveillance and censorship. *Journal of Information Technology & Politics*, *20*, 113–124. https://doi.org/10.1080/19331681.2022.2063215

Stoycheff, E., Liu, J., Xu, K., & Wibowo, K. (2019). Privacy and the panopticon. *New Media & Society*, *21*, 602–619. https://doi.org/10.1177/1461444818801317

Strycharz, J., Kim, E., & Segijn, C. M. (2022). Why people would (not) change their media use in response to perceived corporate surveillance. *Telematics and Informatics*, *71*, 101838. https://doi.org/10.1016/j.tele.2022.101838

Strycharz, J., & Segijn, C. M. (2023). Consumer differences in chilling effects. In A. Vignolles & M. K. J. Waiguny (Eds.), *Advances in advertising research (Vol. XII)* (pp. 107–120). Springer Fachmedien. https://doi.org/10.1007/978-3-658-40429-1_8

Strycharz, J., & Segijn, C. M. (2024). Ethical side-effect of dataveillance in advertising. *Journal of Business Research*, *173*, 114490. https://doi.org/10.1016/j.jbusres.2023.114490

Stubenvoll, M., & Binder, A. (2024). Is knowledge power? *Computers in Human Behavior*, *150*(2024), 107949. https://doi.org/10.1016/j.chb.2023.107949

Theocharis, Y. (2015). The conceptualization of digitally networked participation. *Social Media + Society*, *1*, 1–14. https://doi.org/10.1177/2056305115610140

Thomas, F., Shehata, A., Otto, L. P., Möller, J., & Prestele, E. (2021). How to capture reciprocal communication dynamics. *Journal of Communication*, *71*, 187–219. https://doi.org/10.1093/joc/jqab003

Valkenburg, P. M., & Peter, J. (2013). The differential susceptibility to media effects model. *Journal of Communication*, *63*, 221–243. https://doi.org/10.1111/jcom.12024

van Dijck, J. (2014). Datafication, dataism and dataveillance. *Surveillance & Society*, *12*, 197–208. https://doi.org/10.24908/ss.v12i2.4776

Véliz, C. (2020). *Privacy is power*. Penguin (Bantam Press).

White, G. L., & Zimbardo, P. G. (1975). *The chilling effects of surveillance.* https://apps.dtic.mil/dtic/tr/fulltext/u2/a013230.pdf

Journal of Communication, 2026, 00, 1–13
https://doi.org/10.1093/joc/jqag004
Original Article