UNIVERSITÀ DEGLI STUDI DI PADOVA

Head Office: Università degli Studi di Padova

Department of General Psychology

Ph.D. COURSE IN: Brain, Mind and Computer Science
CURRICULUM: Neuroscience, Technology and Society
XXXV SERIES

THE PRINCIPLE OF TRANSPARENCY AS A DESIGN GOAL AND AS A DISCOURSE TOPIC: TWO CASE STUDIES

Coordinator: Prof. Anna Spagnolli
Supervisor: Prof. Anna Spagnolli
Co-Supervisor: Prof. GianAntonio Susto

Ph.D. student: Mariavittoria Masotina

ABSTRACT

The principle of transparency was first proposed as a safeguard with reference to state organization; it later became the advocated protection of society as a whole, and is now incorporated into laws and regulations to minimize the risks related to information technology. Transparency derives its metaphorical meaning from the signifier of *something that can be seen through,* thereby promising to display and to understand. Despite the immediacy of this visual metaphor, however, transparency is not visibility per se but a medium to enhance visibility. As such, it can influence people's interpretations and understanding of what they see.

This dissertation embraces a user-centered approach when designing for transparency in information systems and digital interfaces. This approach focuses on the recipient of the information (i.e., the user), their language, and comprehension, to make sure that transparency is genuine and not a mere legal compliance.

My first three studies focus on improving comprehension by using the information already implicit in the context when designing privacy notices. They define context in a novel way, as provided by those elements spatially and temporally surrounding the action at stake. The reason is that action, according to ethnomethodology and conversation analysis, creates a background against which the subsequent events are interpreted. The guiding hypothesis of the first three studies is, therefore, that the action performed by the user on the website immediately preceding the appearance of the privacy notice can affect its comprehension. In the first two studies (N = 132, 128), following a between-participants design, I manipulated the *consecutiveness* of a cookie notification presented in an ad-hoc website by either preserving the sequential connection between its appearance and the users' action triggering it or broking it with a delay. I also manipulate the notice's *explicitness* by explicitly mentioning the trigger in its title or omitting this information. Through a final survey, I measured the participants' comprehension (topic and cause) and experience of comprehension (perceived comprehension, clarity of the notice, and sense of control). Behavioral aspects (notice acceptance and the time to respond) were collected through the ad-hoc website. In the third study (N = 91), I followed the same rationale, adding the interpretation of the notice as a dependent variable and

investigating the effect of different contexts (generic action – i.e., entering the website – or specific action – i.e., downloading). Overall, the results of statistical analysis suggest that the action preceding the notice affects the identification of its cause and the interpretation of its content, whereas the explicit text of the notice does not. The variables did not influence the acceptance of the notice, as would be foreseen by the transparency paradox. These results show the explanatory power of good contextualization: considering the sequential context in which the notice appears seems an effective design practice to achieve genuine comprehension.

In a fourth study I turned my attention to methodology, to explore a method to find how people discuss transparency and transparency-related concerns in spontaneously expressed, real-life discourses. This endeavor still pursues a user-centered approach to transparency, because it aims to access how citizens talk about transparency and then create a shared ground between designers and users. The method consists of collecting in a corpus the discourses of interest and applying qualitative analysis and natural language processing techniques to: 1) assess the relevance of a certain topic in a corpus by checking the overlap between the corpus's keywords and some target passages related to the object of investigation; 2) identify the terminology the document's authors used to refer to the object of investigation. The method is applied to a corpus of newspaper articles as a case study. The analysis shows that even though the newspapers articles might not directly use the term ‘transparency’ to a great extent, transparency related concerns are pervasive, and relate to the core arguments of the corpus as expressed by its keywords (*human, ai, gpt, machine, intelligence*).

CONTENTS

GENERAL INTRODUCTION

This work focuses on the principle of transparency applied to information technology and tackles it from a user-centered perspective. The concept is attracting increasing attention also due to its inclusion of some major legal regulations produced by governments and federations. The General Data Protection Regulation (GDPR; European Union, 2016) in the European Union has a full article devoted to it (art. 5). According to GDPR, "*The principle of transparency requires that any information and communication relating to the processing of those personal data be easily accessible and easy to understand, and that clear and plain language be used*." (GDPR, recital 39).

Before GDPR, transparency was already referred to as such or as "openness" in other codes and regulations as described by Barth et al. (2022) - e.g., Australian Privacy Act (2001); CSA's Model Code for the Protection of Personal Information (Canadian Standards Association, 2001; 2014); Global Privacy Standard (GPS; Cavoukian, 2006); ISTPA's Privacy Framework (Sabo, 2007); ISO29100 Privacy Framework (Technical Committee ISO/IEC, 2011); OECD's Privacy Principles (OECD, 2013). After GDPR, transparency as recurrently been used in other EU regulatory acts such as the Digital Markets Act (September 2022) or the Digital Services Act (October 2022).

Outside information technology, transparency emerges as a taken-for-granted ideal and explanation of how society should work; the call for transparency has become a distinctive aspect of our era – the "transparentocene" (Alloa, 2018). It was born in the political and legal framework of democracy as safeguard for protecting the citizens by giving them a control on political power. Its boundaries have been expanded and its narrative has become ubiquitous, through a process that has been defined "discoursification" (Koivisto, 2022). As a consequence, its meaning has become broader and broader, its normative attractiveness extreme, its conceptual opposites blurred, and its potential market global.

As a result, nowadays transparency is a *magic concept* (Alloa, 2018) of high rhetoric power and context-dependent acceptations (Koivisto, 2022). According to Alloa (2018), transparency has been declined as: *accessibility* (all citizens should access to information – right to know); *accountability* (stakeholders are assumed to develop shared responsibility by making their decisions available to the public); *asymmetry reduction* (if secrecy privileges a few, transparency can rebalance the power); *authenticity* (when nothing is concealed, things can be true to themselves); *moralization* (constant visibility drives people to act virtuously); *procedural fairness* (everyone should be involved in process as a safeguard); *public good* (self-interest temptation is reduced when actions are placed under public scrutiny); *rationalization* (decision-makers improve their rational behavior standards if they have to justify their actions); *self-knowledge* (knowing own selves is the prerequisite for knowing own selves accountability); *truth-making* (falseness is banned by forcing people to speak out). In the variety of declinations transparency has been associated with, its appeal derives from the promise of seeing (a transparent object is something that can be seen through) and understanding (e.g., *I see* as *I understand*). After the datafication of society, its call for has been emphasized with regard to users' data (privacy) and to the algorithms used to elaborate them (Artificial Intelligence; AI). These concepts will be elaborated on in **Chapter 1**.

This dissertation embraces a user-centered approach to transparency in the assumption that focusing on the recipient of information, their language, and comprehension is a necessary step to implement transparency as a principle genuinely. User-centeredness is approached in two separate sections, each one guided by a specific research question inspired by the call for transparency.

The first section is driven by the research question "*How can privacy notices be designed to be genuinely transparent – i.e., comprehensible - to their users?*". This section tests a user-centered strategy called contextualization to make privacy notices usable, thereby improving their understandability. It presents three experimental studies. The studies follow a between-participants design and involve 351 real internet users in total.

The second section explores a method to find the way in which transparency is treated as a topic in real life discourse. This section embraces a user-centric perspective as it seeks to understand how individuals discuss transparency in spontaneously expressed, real-life discourses and establish common ground between designers and users. The methodology explored is proposed to answer the research question "*How can we detect any references to transparency in real-life discourses?*". The study reported in this section applies qualitative analysis and natural language processing techniques to analyze a corpus of newspapers articles about a specific AI (GPT-3).

Chapter 12 synthetizes the present dissertation and provides its conclusions.

CHAPTER 1.
THE PRINCIPLE OF TRANSPARENCY

This chapter elaborates on the different acceptations of transparency as a concept and as a principle. First, it defines transparency as a concept to display how its literal and figurative meaning is related to its promise of being a safeguard. Following that, it briefly overviews the evolution of transparency as a principle from its birth until the contemporary era.

1.1 Transparency as a concept: *seeing is understanding*

The Oxford English dictionary defines *transparency* as “the quality or condition of being transparent” and “that which is transparent”; “a transparent object or medium”; *transparent* is defined as “Having the property of transmitting light, so as to render bodies lying beyond completely visible”; “that can be seen through”. According to this definition, transparency as a concept bears a material and a symbolic meaning. From a material point of view, transparency is a physical property. Primarily, it allows seeing. This material connotation roots its metaphorical meaning (Koivisto, 2022). We tend to prefer visual information over verbal information to retrieve accurate information about reality: describing something by words implies the speaker's interpretation and decoding of the receiver; verbal information acquires authority from the speaker. Visual information seems to derive their authority from reality itself (Morris, 2011). Transparency allows the gaze to reach its target, and *seeing is understanding* (e.g., *I see what you mean*; Christensen & Cornelissen, 2015). Consequently, transparency promises visibility, clarity, and – for analogy - understanding.

Second, transparency allows *seeing through.* It is not visibility. It is a medium enhancing visibility (Koivisto, 2022). Metaphorically it has been associated with a window and with a flashing light. Depending on the metaphor, the implications of being a medium change (Hansen, Christensen, & Flyverbom, 2015). The external reality is concealed from us without a window (a transparent object). By looking through a window, external reality becomes visible. This metaphor presents transparency as a neutral medium, letting reality emerge as it is. The variant of flashing lights introduces the idea of intentionality, temporality, and partiality. In this acceptation, transparency is a tool for illuminating. It allows us to see by directing a beam of light toward certain areas, leaving other areas in the dark (Teurlings & Stauff, 2014). The metaphor implies that a) there is somebody that directs the beam of light, presupposing an intentional and purpose-driven act, b) it enhances the visibility of the illuminated reality, not of the reality in its whole, and c) being flashing, it is also temporal (Flyverbom & Albu, 2019). Regardless of the metaphor, transparency is consistently a medium. According to Birchall (2011), an inconspicuous one: "it is seen as not having a particular meaning in itself but as, rather, merely the invisible medium through which the content is brought to our attention, into the visible realm." Still, not paying attention to the medium doesn't mean its effects on the content are irrelevant: "A medium implies the conscious sharing or constructing of its objects: it implies selection, highlighting, omission and following the conventions of the medium. (…) When we see with the aid of a tool, the tool becomes part of what we see." (Koivisto, 2022).

Third, transparency, as being a physical property, is also an affordance – i.e., a call to action, the physical quality of an object that suggests to a human being the appropriate actions to manipulate it (Gibson, 1996; Norman, 1998). To be realized, affordances need an addressee. Otherwise, they will remain a possibility. In this light, transparency as created visibility needs someone who addresses the affordance, it needs a gazer. And by enhancing the gaze, it empowers the gaze beholder. Following the categorization of Koivisto (2022), in line with the metaphor of *seeing is understanding,* a gaze can not only help us navigate the world by spontaneously retrieving information about our surroundings (*observatory gaze*), but also change our perception of reality, revealing that things are not as they seem (*revelatory gaze*). The revelation assumes the object to have no time to react if it's gifted with intentionality. Beyond understanding, a gaze can create its object, exercising a form of power on it. On the one hand, individuals can modify their behavior in response to the awareness of being observed (e.g., Hawthorne effect; Franke & Kaul, 1978). On the other hand, we can be guided by our beliefs while we look at something (Kahneman, 2017), forming our interpretations of it accordingly (*believing is* *seeing;* Morris, 2001). Thus, neither the observer nor the object itself – when aware - can be considered mechanical participants in the inner logic of transparency.

1.2 Transparency as a safeguard: the moral, political and legal project

The call for transparency has had a long run. The emergence of transparency as an ideal dates back to Aristotle and to Augustine in the Middle Ages. In this context, transparency was a moral project, referring to the purity of heart or soul (Koivisto, 2022). With the Enlightenment, the call for transparency emerged in political discourses, being advocated as a moral, legal, and political project (Baume & Papadopoulos, 2018). It started to be one of the very cornerstones of democracy and the mechanism through which power could be legitimized and controlled. In a social contract-style democracy, the governed are the ultimate wielders of political power. The governor exercises power on their behalf and is accountable to the governed for his use of power (Baume, 2018). Enhancing the citizens "seeing" the operation of their governor, they will be able to know how the latter is using the power they conceded. Based on the information retrieved, they can sustain the governors or substitute them. In contemporary democracy, the rationale changes slightly: it is a legal and political authorization that puts the governor in the position for exercising the power and transparency practices (specified by the legal framework) are additional procedure guarantees against the misuse of that power. They function as tool for giving information about the ways and the reasons why the power is being used on citizens. In both cases, transparency is connected to the understandability of the power, which becomes a necessary precondition for legitimacy (Koivisto, 2022).

In this context, transparency works through preventive and justificatory mechanisms. The preventive mechanism acts by making a misuse of power as little fascinating as possible. It builds on the analogy of how transparency works in social life (Ringel, 2018). In social life, (intentional) transparency is a way to manage the impression we want to project to our audience. In interacting with other people, we show some aspects of our selves, in a way that is in line with its their real or imagined expectation for us (Goffman, 2006). The goal is to give a positive impression. When we lose control over our self-presentation, our audience will project on the our selves elements we wish it would not (e.g., socially inappropriate behaviors or thoughts). The consequence will be shame and possibly the loss of our social status (Goffman, 1982). In the same way, the governors have to manage their self-presentation. Being aware that their actions are surveilled and documented, they will act by the citizens' expectations, or risk being humiliated and losing their charge and status. The public opinion will serve as a deterrent for illegit actions and a nudge for exercising power fairly (transformative gaze).

The justification mechanism works in a post-ante way (Koivisto, 2022). The governors will justify their actions, creating a plausible narrative of how they used the power. The assumption is that by demonstrating to have followed legitimate procedures, their power's legitimation will be further confirmed. The presence of the preventive mechanism will make their actions retraceable, limiting the opacity of their narratives. In this framework, transparency as a law project represents the complement of the political project. The law will be the technocratic realization of the political project, regulating the practices in the production of the documentation and the citizens' access to the government-held documentation, balancing the right of the citizens to know and the importance of keeping some information protected.

1.3 The growing popularity of transparency as a safeguard in the contemporary era

Born to legitimize and control the power in the optics of state and popular sovereignty, transparency as a safeguard has transcended this limit and penetrated the social tissue in the last decades. From being a safeguard for the state organization during the birth of the democratic society, it grew into protection for the entire society in the contemporary era – keeping the mechanisms guaranteeing its defensive capabilities described in transparency as a political project. The increasing demand for transparency, and its growing popularity, can be traced back to the late 1980s when the World Bank introduced the concept of good governance as a criterium for granting aid to the countries asking them (Koivisto, 2022). If governance was defined as "the manner at which the power is exercised in the management of a country's economic and social resources for development", transparency was identified as both the condition and the objective for development (World Bank, 1989). The concept became popular, and more and more global governance institutions started applying it (e.g., International Monetary Fund, United Nations Development Program, and European Union; Koivisto, 2014). Transparency was normalized and associated with a wider concept of a good and functioning society – and uncontestable concepts such as human rights, democracy, and participation.

In 2009, after the financial crisis, the growing power of global institutions was raising a growing concern for their legitimacy (Peters, 2013). As a response, these institutions started adopting transparency policies to demolish the perceived institutional-informational asymmetry using transparency "*as if* it were a legal concept and *as if* the policies worked as binding legal regulations" (Koivisto, 2020). Their adoption became so common that a *transparency turn* was claimed (Gupta, 2008). Along with this phenomenon, transparency became more widely known. In the 2000s, it was institutionalized in other domains (e.g., ethics and economics), and its popularity grew to the point that its spread assumed the characteristics of a *pandemic* (Pollit & Hupe, 2011). At the same time, with the rise of the internet and the growing use of digital platforms, private companies strategically started using transparency to enhance their credibility. These entities assumed their own transparency narrative by publishing transparency policies and declarations on their websites (Koivisto, 2022).

The advent of the internet associated with the growing computer power, data analytics, and employment of Artificial Intelligence (AI) has meant a drastic change in our society. It turned digitization (i.e., the conversion of analogical contents - e.g., books - in digital ones) into datafication (i.e., turning all the aspects of users' life in data). The datafied information can be used to generate new value, impacting the economy and the balance of power (Cukier & Mayer-Schoenberger, 2013). With people relying on digital services for pursuing their tasks during their daily life, the few private companies offering the most popular services – and collecting their users' data at the same time - have gained a considerable power (Mäihäniemi, 2020). According to Koivisto (2022), instead of governments or intergovernmental organizations, the primary power holders of the digital society are the monopolies owning of these services. Similarly to what happened for global institutions, the growing power these entities can exercise on society – and on their users - has been accompanied with a call for transparency in their practices. This call has been twofold, deriving from the two main enablers of their power: data and AI.

1.3.1 The call for transparency for data protection

Users' data represent the fuel of the datafied economy. The footprints people leak online are gathered by the companies providing the digital services we use. On the one hand, users can put efforts and attention into curing their digital image (e.g., using social networks). On the other hand, we – as users - seldom consider ourselves when performing online (e.g., inquiries on Google, time spent watching videos, webpages visited). Through our behavior, we reveal aspects of ourselves that we usually would not show to other people. We act as if no one is watching, as we are only giving information to a machine, as our involuntary disclosure will not be seen or used. "But we know that's not true" (Harcourt, 2015). In addition, users do not know what they are suggesting through their behavior and to whom (Koivisto, 2022). This inequality in power between the digital platforms and their users has been recognized, and those companies are increasingly required to be transparent (e.g., GDPR).

1.3.2 The call for transparency for AI systems

Along with the growth of computational power – and the proliferation of data – AI systems have become increasingly employed (e.g., Algorithm Watch, 2019). Their ability to generate human-like products (e.g., writing, pictures, deep fakes), and their functioning as black boxes have raised concerns about their potential drawbacks. The spread of post-truth politics, alternative news, and fake news has been considered a clue of the reality principle loss of its privileged position in public opinion (Koivisto, 2022). The advent of verisimilar products representing something never happened (e.g., images of people never existed, deepfakes) has been argued to exacerbate this trend (Kline & Holland, 2020). On the other hand, users rely on decisions guided by these systems. The results during an online search or the contents shown on social media are an example of these decisions. Decisions that the system makes on its users' behalf. And we seek to understand how power is wielded over us, regardless of whether the decision-maker is a person, organization, or machine (Koivisto, 2022). The call for transparency has been advocated as solution for limiting both these issues as a variation from the traditional one.

In this case the object of the call is not only a legal entity (person or company), but it targets the AI itself.

1.4 Transparency as a safeguard of the (digital) society

Transparency is so deeply woven into the fabric of our society that we consider it a taken-for-granted principle. Nonetheless, the concept has come a long way, gradually exiting from its initial boundaries to become the legitimization of societal and new power-related dynamics. As presented in the previous sections, as a consequence of the datafication of society process, the global power constellation is changing, fueled by human-users' data. After having been extended from its initial legal and political framework, transparency has been presented as a solution for responding to the derived new societal challenges. In its conceptual definition, transparency promises the possibility of seeing, understanding, and consequently controlling the way the wielders of power exercise it. Nonetheless, despite the immediacy that the visual metaphor suggests, transparency is not visibility. It is a medium enhancing visibility. As a medium, the way it is built influences how the people will interpret the object they are observing and their ability to understand it – and understanding a pre-requisite to exercising their civil and existential rights.

SECTION I

Testing A Strategy To Improve The Understandability (Transparency) Of Privacy Notices[1]

[1] The content of this section has been published in Masotina, M., & Spagnolli, A. (2022). Transparency of privacy notices and contextualisation: effectively conveying information without words. *Behaviour & Information Technology*, 1-31. The data are available via my institution data archive at http://researchdata.cab.unipd.it/637/.

Section I

CHAPTER 2.
INTRODUCTION

Regulations such as the General Data Protection Regulation in Europe (GDPR, 2016) and the California Consumer Privacy Act (CCPA, 2018) require websites to ask for the visitor's informed consent before collecting any data on their browsing activity in order to grant the right to privacy (i.e., the protection of personal or private information from misuse or unauthorized disclosure). They also require that such informed consent implements a principle of transparency according to which "*any information and communication relating to the processing of those personal data be easily accessible and easy to understand, and that clear and plain language be used*" (GDPR, 2016). The premise is that people will not be able to grant any *informed* consent - and possibly exercise their right to privacy - if they cannot understand what they are consenting to. In other words, consent forms should not only contain information, but should also be designed in such a way that average users can easily understand them. Contrary to this principle, however, privacy notifications and security warnings seem to do little to empower users to protect their online privacy and security. A survey in the European Union in 2019 found that 84% of the interviewees (N = 23.106 European citizens) do not read the privacy notices (European Commission, 2019). Visitors prefer using shortcut buttons instead of reading privacy policies (Obar & Oeldorf-Hirsch, 2018), and miscomprehend even very simple and intuitive privacy policy explanations (Korunovska et al., 2020).

How can privacy notices be designed to be genuinely transparent to the user and then being fully compliant *with the spirit of* current privacy regulations? A user-centered approach, which considers the users' psychology in designing security or privacy notices, can help to address these issues (Birge, 2009; Schaub et al., 2015; Rossi & Lenzini, 2020). In particular, we can exploit a principle of usability that was elaborated well before privacy and security became objects of investigation in HCI; the usability expert Donald

Norman calls this principle 'offloading the task'[2], i.e., reducing the cognitive load required to carry out a cognitive task by exploiting familiar pieces of information that the users can recognize and work as an external memory. Icons or familiar color codes conveying information about the severity of a danger (Habib et al., 2021; Felt et al., 2015; Silic & Cyr, 2016; Yang et al., 2017) apply this usability principle. My approach is in the same vein and aims to offload transparency notices by providing privacy information without adding such information explicitly to the notice; I exploit, when possible, the information already available in the context. This is in line with recent recommendations to transfer the burden of understanding privacy information from the user onto the system without reducing the users' control and power (Acquisti et al., 2017; Spagnolli et al., 2017).

The research question "*How can privacy notices be designed to be genuinely transparent – i.e., comprehensible – to their users?*" guides this section. The proposal made here is: *by carefully designing the notice context of appearance*. The goal is not to provide an example of how privacy notices should be designed to fully comply with the guidelines for enacting GDPR at the time this dissertation is written – since they have been updated and ameliorated as time went by. Here the focus is on contextualization as a user-centric solution. This strategy can be applied regardless of the specific privacy notice and the updates foreseen in the guidelines. The suggested solution is tested through three consecutive experimental studies. The studies investigate the effect of the context of appearance of a cookie notice – defined according to conversation analysis and ethnomethodology - on participants' comprehension of its cause and content, their interpretation of the notice, their experience of comprehension (perceived comprehension, clarity of the notice and sense of control), and response to the notice. The experiments follow a between-participants experimental design and reproduce the experience users live in their daily life: they enter a website to pursue a task and encounter a cookie privacy notice.

The section is organized as follows:

[2] https://www.nngroup.com/articles/minimize-cognitive-load/

Chapter 3 introduces the related works, providing a background based on the scientific literature.

Chapter 4 gives an overview of the studies and of their rationale, starting with the identification of the research gap, deepening the definition of context I relied on, and presenting the hypotheses.

Chapter 5 provides a technical and legislative framework of the privacy notice presented to participants during the studies and describes the studies' experimental materials, procedure, and ethics.

Chapter 6, Chapter 7, and Chapter 8 describe Study 1, Study 2, and Study 3, respectively. Each chapter explicates the specific hypotheses of the study, provides specific information about the materials and methods, and reports the results and conclusion.

Chapter 9 presents the conclusions of Section I, discussing the hypotheses and drawing the design implications of the findings.

Section I

CHAPTER 3.
RELATED WORKS

The studies presented in this section will test contextualization as a strategy for improving users' comprehension. This chapter introduces the factors identified in the scientific literature as limiting the comprehension of privacy and security notices and preliminary evidence about the importance of context in influencing users' interpretations and decisions.

3.1 Obstacles preventing the comprehension of privacy and security notices

The ineffectiveness of online privacy and security notices has been explained in terms of habituation, convenience, and bad design. Habituation is the "decreased response due to repeated stimulation" (Groves & Thompson, 1970), i.e., a decreased attention to the warnings due to frequent exposure to them. Several studies have pointed out the role played by habituation on the response to warnings and alerts (Egelman et al., 2008; Wogalter & Vigilante, 2006; Akhawe, & Felt, 2013; Egelman & Schechter, 2013). Egelman et al. (2008), for instance, found a correlation between recognizing a warning and ignoring it. Evidence of habituation has been also found at a neurological level since the activation of the visual processing centers of the brain seems to decrease during the interaction with advice to which participants have been already exposed (Anderson et al., 2015).

A second explanation is that users rely on shortcuts and heuristic criteria to make privacy or security decisions. A recent survey with 6000 users reports that the response to security warnings depends on the extent to which the visitors trust the website or the browser (Reeder et al., 2018), also confirmed by Almuhimedi et al. (2014). Visitors' decision is an overall act of trust instead of an informed choice following a genuine comprehension of the security risks they are running. Users' reliance on peripheral cues and cognitive shortcuts can be deliberately exploited by so-called dark design patterns to increase their acceptance to disclose personal data (Soe et al., 2020; Waldman, 2020).

For instance, participants in a study by Utz and colleagues (2019) accepted to have their data treated by third parties if acceptance was the pre-selected option on the notice; the same option was selected more often when colored in blue instead of grey. Chang et al. (2016) found that the participants were more likely to disclose sensitive information about themselves (e.g., about their sexual experiences) if they were exposed to provocative profile pictures on a fictitious social network than participants who were exposed to less provocative images.

Disconcertingly, transparency itself can work as a heuristic, creating a transparency paradox. Indeed, solutions designed to increase users' awareness and comprehension might backfire, decreasing users' privacy concerns, increasing users' trust in websites and services, and encouraging incautious behaviors (Acquisti et al., 2013). For example, being transparent on the data collection aims and methods can reassure users and make them perceive that the risks are lower. In Oulasvirta et al. (2014), participants were asked to report their privacy concerns in nine hypothetical surveillance scenarios. Despite the intrusiveness of the scenarios remained unchanged, respondents provided with intentions about data usage reported lower concerns for their privacy than participants who did not receive this information. In Masotina et al. (2019), participants perceived the sensitivity of data asked in an experimental study depending on the transparency of the method used for collecting them: the participants who self-reported the information (transparent method) found them less sensitive than the participants who believed an eye tracker (opaque method) had collected the same information. The results in Xiong et al. (2020) suggest that also perceived comprehension could participate in the transparency paradox: the version of privacy policy participants reported to have understood better reached also the numerically largest disclosure rates of highly sensitive information. Furthermore, the mere presence of a privacy policy can suggest to users that they are protected regardless of its contents. For example, in Earp and Baumer's survey (2003), only 54% of participants reported they would read privacy policies, but 66% of the sample indicated the presence of privacy policy as increasing their confidence in the website. The authors suggested that users were reassured by the provision of privacy policies, but less concerned about the specific information contained in them. Hoofnagle & Urban (2014) came to the same conclusion when the mere presence of a privacy policy was incorrectly considered by 62% of survey respondents as a warranty that a website could not share their personal information.

To be blamed is not simply the users' laziness or gullibility, but the obscurity of privacy and security notices. The difficulty in understanding the language of a notice or the number of steps required to shield one's device might seem to exceed the benefits (Herley, 2009). The users' lack of comprehension of privacy and security threats is well documented. In a study by Felt et al. (2015), less than half of the sample was able to correctly identify the source of the threat after having seen an SSL warning, and less than 20% could detect which pieces of data could be at risk. The comprehension of the specific content of warnings and notices is low (Reeder et al., 2018), the recall rate of the warning content is poor (Malkin et al., 2017), and the content might be altogether misunderstood (Milne & Culnan, 2004). Only 31% of participants correctly answered that phishing scams attempt to steal personal information in a study by Egelman & Schechter (2013).

Not only actual comprehension is poor, but also users report their understanding to be low. For example, Milne and Culnan (2004), Vail et al. (2008), and Reidenberg et al. (2015) measured their participants' perceived comprehension by asking them if they felt confident in their understanding of privacy notices using 5-point Likert scales. In all the cases, typical users did not reach the agreement threshold. Notably, there are contrasting results for the relationship between perceived actual understanding. If Kühtreiber et al. (2022) reported a correlation between the two scores, Vail et al. (2008) found that participants' perceived comprehension did not always reflect their actual comprehension. The authors assessed the actual and perceived comprehension of four versions of the same policy, and participants reported higher confidence in their understanding of some versions of the policy they performed worse in the objective comprehension score. Milne and Culnan (2004) pointed out that low perceived understanding was associated with a decreased willingness to read notifications in the future.

Obscure language, confusing terms, and technical jargon reduce the likelihood of having a clear assessment of the warning implications (Zaaba et al., 2014). It has been estimated that to read the privacy policies would take an average of 30 minutes (Obar & Oeldorf-Hirsch, 2018) and that complex privacy policies drain resources from the decision itself (Franz et al., 2021).

The meta-analysis conducted by Argo et al. (2004) on warnings suggests that the strategies to increase the vividness of the notice such as font size, colors, chunking, and the use of pictures and symbols affect attention but do not significantly impact the users' comprehension and recall. Therefore, specific strategies to improve comprehension must be found, so that the principle of transparency advocated by privacy regulations is genuinely applied. Solutions are explored to simplify the text of security and privacy notices (e.g., by mimicking nutrition and energy labels; Kelley et al., 2009) and make it more readable.

3.2 Offloading comprehension via contextualization

Context can be defined as "a set of situational elements in which the object being processed is included" (Bazire & Brézillon, 2005), namely the environment in which that object is placed; that environment contributes to the process through which such element is made sense (e.g., Bonner & Epstein, 2021). The role of context in effecting the perception of stimuli has also been observed via neuroimaging techniques (Willems & Peelen, 2021); likewise, the different responses to actions embedded in congruent, incongruent or ambiguous context have observed via transcranial magnetic stimulation (Amoruso et al., 2018). The role of context in leading interpretation has also been described in the studies of persuasion; they show that the environment in which a persuasive message is received affects the recipients' appreciation of its relevance and ultimately its effectiveness; this phenomenon is called by Cialdini pre-suasion and refers to persuasive elements of the context of a persuasive the message (Cialdini, 2016). In designing security and privacy notices, a good contextualization exploits the information already present in the environment to understand the meaning of a notice, while at the same avoiding any contextual cues that can distort such meaning.

The context of security and privacy notices has been considered in terms of the conspicuity of the notice against their perceptual background, for instance by varying their position, interactivity and color (e.g., Utz et al., 2019). The context is also taken into account in studies varying the temporal relation between the appearance of the privacy notice and the activity at stake: active warnings (i.e., warnings that interrupt the task of the users) have been proven to be more effective than passive ones (Wu et al., 2006) so that they are now a standard for security aims (Akhawe & Felt, 2013). Sometimes, however, users are annoyed by active warnings interrupting their tasks (e.g., Sunshine et al., 2009; Schaub et al., 2015). Egelman and colleagues (2009) found that consumers would pay a privacy premium when interested in sensitive goods online and that this willingness is affected by the timing at which the privacy level of the website is displayed.

The relation between the notice's context of appearance and their comprehension has been less investigated. Bolchini et al. (2004) suggest allowing direct access to that portion of the privacy policy relevant to the visitors' current action on the website. Relevance is operationalized as the affinity between the subject of the privacy policy and the action the user is performing on the website, providing the context to understand the privacy policy. A similar notion is advocated a decade later by Schaub et al. (2015), who suggest - among other solutions – to provide smaller pieces of privacy information specific to the transaction at stake. Both Bolchini et al. (2004) and Schaub et al. (2015) describe this solution but do not test it. A direct test is provided by Klumpe et al. (2020), who manipulated the triggering action for authorization requests to collect the users' location information; they assumed that a push trigger when entering the app versus a pull trigger when performing specific location requests on the map would affect the notice's interpretation. In the former case, it would look more intrusive because it would clear the collection of users' behavior on the app at large, and then lead to less acceptance. They check the effect of their assumption on the risk vs benefits tradeoff of privacy decisions but do not test the notice interpretation directly.

Section I

CHAPTER 4. THE PRESENT STUDIES

This chapter reports the research gap identified in the previous studies, deepens the definition of context followed in designing contextualization in the experiments, and presents the hypotheses driving them.

4.1 Research gap

As presented in Chapter 3, the context of privacy and security alerts was primarily considered in previous studies for its effect in influencing users' *attention* and *response* to privacy and security alerts (e.g., Utz et al., 2019; Wu et al., 2006; Akhawe & Felt, 2013; Sunshine et al., 2009; Schaub et al., 2015). Solutions for increasing the understanding of the information provided in notices mainly focused on improving the clarity of their contents by simplifying the text presented (e.g., Kelley et al., 2009). Less attention has been given to the relationship between privacy notice's context of appearance and their *comprehension*. Bolchini et al. (2004) and Schaub et al. (2015) proposed providing users a context to understand privacy policy to increase their comprehension, but no direct evidence is available to support this proposal. The present studies aim to fill this gap and test the effect of contextualization on the comprehension of the privacy notice. It will also be considered whether it is necessary to put all information on notice since exploiting the context can make some of this information superfluous. Indeed, the assumption is that contextualization exploits the information *already available* in the context.

Concerning the derivation of the hypotheses, as previous research did, the studies will consider perceived comprehension (Milne and Culnan, 2004; Vail et al., 2008; Reidenberg et al., 2015), clarity of the text (e.g., Zaaba et al., 2014), sense of control (e.g., Brandimarte et al., 2013), and actual comprehension (e.g., Milne & Culnan, 2004, Felt et al., 2015; Malkin et al., 2017; Reeder et al., 2018). Regarding the latter, previous studies mainly focused on the *contents* of notice. Here, the comprehension of its *cause* will be taken into account as well. In addition, prior research highlighted that a *transparency paradox* may offset the strives made to improve users' awareness on the information they provide. Despite not being the main interest of the present studies, this aspect will be taken into consideration too, contributing to the available literature in this regard.

The studies presented here are specifically focused on *contextualization* as a user-centric solution. Previous works provided little information about the theorical background guiding the operationalization of the *context*. Here, I will rely on a definition of context following ethnomethodology and conversation analysis.

4.2 Definition of context in my studies

The notion of context can be fuzzy to define since it might include everything the whole environment in which the user is situated. For instance, the context can be the situation in which the data transaction occurs and related expectations on the kind of data flow that is appropriate there (Nissenbaum, 2011). This approach is close to studies such as Spagnolli et al. (2015) or Ebert et al. (2020), who found that the willingness to share personal data differs according to the party collecting the data and the nature of the transaction (e.g., medical, commercial). Here I operationalize the context following ethnomethodology and conversation analysis, which uses action as the center around which to sort out which elements of the environment are relevant and represent then the type of context to take into account. In this approach, the context is primarily provided by those elements spatially and temporally surrounding the action at stake that are consequential to such action (Goodwin & Duranti, 1992). In particular, I focus on the temporal context; actions create a background against which the following events are interpreted (sequential implications; Schegloff, 2007). For instance, Schmidt et al proved that users having previously authorized the tracking of their browsing behavior attribute to that choice the cause of subsequent price changes in the services, which in turn drives their evaluation of the price fairness (Schmidt et al., 2020). So, I define context as the *action performed by the user on the website immediately preceding the appearance of the privacy notice*.

4.3 Hypotheses

According to Bruner (2005), we understand events in narrative form, by linking one event to another through case-effect relationship. Sloman and Lagnado (2015) aligned with the idea that we unfold events over time in a narrative-like way. In their review, they highlighted how our perception, reasoning, decision making, judgements, and attributions are guided by causal reasoning, arriving to claim that "the human cognitive system is built to see causation as governing how events unfold" (p.224). Not being able to retrace an adequate causal explanation for an event can impair individuals' understanding and, consequently, appreciating the cause of an event allows a fuller appreciation of its motivation and meaning. Therefore, the aspects of the notice investigated during my studies as object of comprehension are its content (henceforth "topic") and the trigger causing its appearance (henceforth "cause"). Also, transparency aims not only at achieving an actual improvement in users' knowledge but also to empower the user so they feel to be in control (Kools et al., 2006); so, I assume that the users' subjective experience, i.e., perceived comprehension, clarity, and sense of control on the data might be affected as well. I hypothesize that:

H1 – Effect of consecutiveness on comprehension: the event preceding the appearance of the notice affects the comprehension of the notice's cause and topic, and the experience of comprehension of the notice (perceived comprehension, clarity, sense of control).

Bolchini et al. (2004) and Schaub et al. (2015) also contend that contextualization is stronger than a mere mention of the context in the text of a notice. They state that privacy policies mention the situations to which they are applicable instead of appearing in those situations; being read under circumstances removed from those that would make information actionable, privacy information fails to become useful and to inform decisions. It can then be expected that a mere mention to the event motivating notice would not provide as effective an interpretive frame as appearing after that event. On the other hand, the verbal categories used to describe the message do affect its apprehension, a process referred to as framing (e.g., Adjerid, et al., 2013; Gluck et al., 2016). Therefore, I also varied the presence in the notice of the mention to the event motivating its appearance and measured its effect on comprehension. I then hypothesize that:

H2 - Effect of explicitness on comprehension: explicitly mentioning the trigger affects the comprehension of the notice's cause and topic, and the experience of comprehension of the notice (perceived comprehension, clarity, sense of control).

The literature points at the possibility that the transparency of a website or service can encourage incautious behaviors in the users, by making them feel that they can trust the service (Acquisti et al., 2013; Oulasvirta et al., 2014; Masotina et al., 2019; Paunov et al., 2019). Therefore, I will also consider the participant's response to the notice. It is possible that my manipulation has the undesirable effect of increasing the acceptance of the notice's request. My third hypothesis is that:

H3 - Transparency paradox: the factors increasing transparency also increase the acceptance of the request to track the users' activity.

The hypotheses are visually summarized in Figure 1. Consecutiveness and explicitness are expected to influence participants' comprehension of the cause and of the topic of the notice, experience of comprehension (i.e., perceived comprehension, clarity, and sense of control), and their response to the notice, increasing the acceptance to its request to accept cookies. As a behavioral measure, the time elapsed from the notice appearance to the participants' decisions will be collected as well, to show how long the participants took to decide. Each of the three studies reported in the rest of the section declines the three core hypotheses of this project more specifically based on the manipulations attempted there.

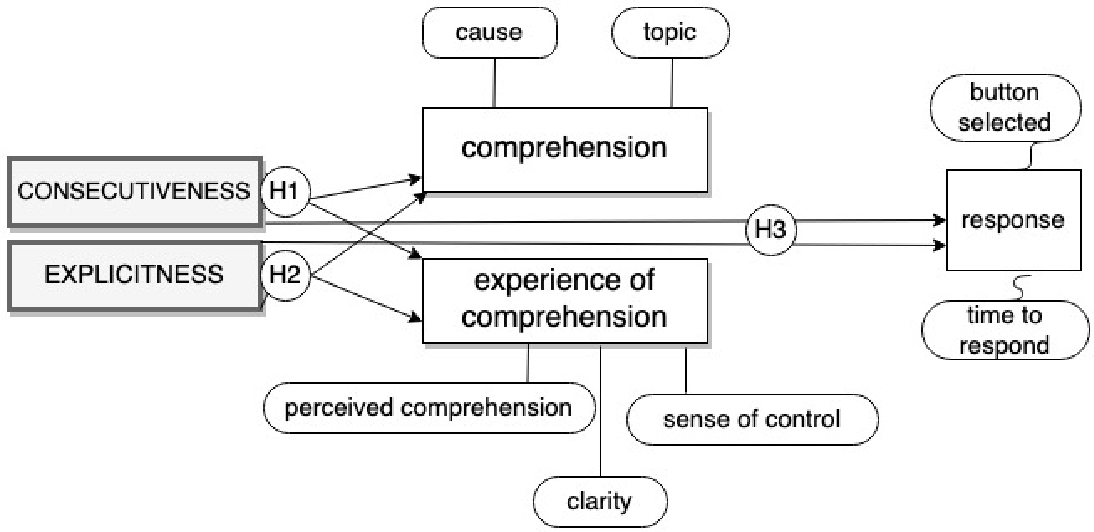

Figure 1: Visual summary of the hypotheses.

Section I

CHAPTER 5.
METHOD

During the experiments, cookie notices were considered as a case study for privacy notices. This chapter introduces the technical and legislative framework of cookie notice before providing information about the three studies' materials, procedures, and ethics. Since slight variations were introduced in the materials across the studies, this chapter provides an overview of the features all the studies had in common. More detailed information about the specific materials used in each study will be provided in the relevant chapters (Chapter 6 for Study1, Chapter 7 for Study 2, and Chapter 8 for Study 3).

5.1 The cookie notices

Cookies are small blocks of data generated by a webserver while a user is browsing a website. These blocks are stored on the users' devices by their web browser. Generally, cookies are classified according to three different aspects: their duration, their provenience, and the purpose they serve (e.g., https://gdpr.eu/cookies/). Concerning their *duration*, cookies can be divided in session cookies - which expire when the user closes their browser - or persistent cookies. The latter will last until the users or the browser actively erase them or until their expiration date. Regarding their *provenance*, cookies are characterized as first-party cookies (i.e., the website the user is visiting generates them) or third-party cookies (i.e., the cookies are not placed on the device by the website the user is on but by a third party – e.g., an advertiser or an analytic system). According to their *purpose*, cookies can be considered as: strictly necessary if their purpose is to make it possible to browse the website and use its features; preference cookies if, by using them, a website can retain users' previous preferences and selections (e.g., the language or name and password to automatically log-in); statistics cookies, which collect info about how the website is used (e.g., visited pages) for the sole purpose of improving the website functions; marketing cookies, that track the users' activity online and can share information with third parties. The latter are usually persistent and fairly often third-party.

Currently, cookies are commonly utilized to monitor users' online behavior. Tracking users' behavior is made possible thanks to the association of a unique identifier to the user's browser. This identifier allows the server to recognize the browser and access information associated with it. If a user logs in to a website, the server can identify the logged user through the cookie. Data stored in cookies can be considered personal if they allow identifying the user and, as a result, they fall under the jurisdiction of the GDPR (Recital 30 of GDPR). GDPR is supplemented by the ePrivacy Directive (EDP, 2009), which addresses crucial aspects of Internet based tracking. According to the current legislative framework, users should: 1) provide an informed consent before the website uses cookies (except the strictly necessary ones); 2) receive accurate and specific information about data tracked by each cookie and its purpose before the consent is provided; 3) being able to access the service regardless of their refusal to allow certain cookies; 4) be provided with an easy way for withdrawing their consent. Cookies notices are then thought for asking the consent to their users at their entrance on the website. In this context, it is important to underline that being under the GDPR, the consent should be informed and implement transparency as a principle of understandability (GDPR, Article 12).

The rules governing cookie policies are still on set. The EDP was supposed to be replaced by the ePrivacy Regulation (EPR; Gonzalez, et al., 2020) at the same time GDPR came into force. Nonetheless, after five years, the EPR is still a proposal, and an enactment before 2025 is considered unlikely (Kretschmer et al., 2021). At the same time, cookie notices design and the granularity of control they offer is in continuous evolution. If nowadays a state-of-the-art cookies notice will provide the users the option not only to accept but also for rejecting the cookies directly in the banner (e.g., the cookie banner in the UK government website allows to *accept additional cookies*, *reject additional cookies* or *view cookies*; https://www.gov.uk/), at the time of the first two studies not having this option was the norm (e.g., Figure 2). Here, I used as baseline a cookie notice inspired by the design, the text, and the options of typical cookies notice users encountered during their daily life. As for warnings and other privacy notices, patterns of bad and dark design have been found in cookies notices (e.g.,Utz et al., 2019; Kretschmer et al., 2021), users show insufficient knowledge about their contents, report to be worried that their personal identifiable information will be attached to tracking activities – thereby affecting their right to privacy - and admitted to make compromises when choosing between effort, awareness, and control (e.g., Ha et al., 2006; Smit et al., 2014).

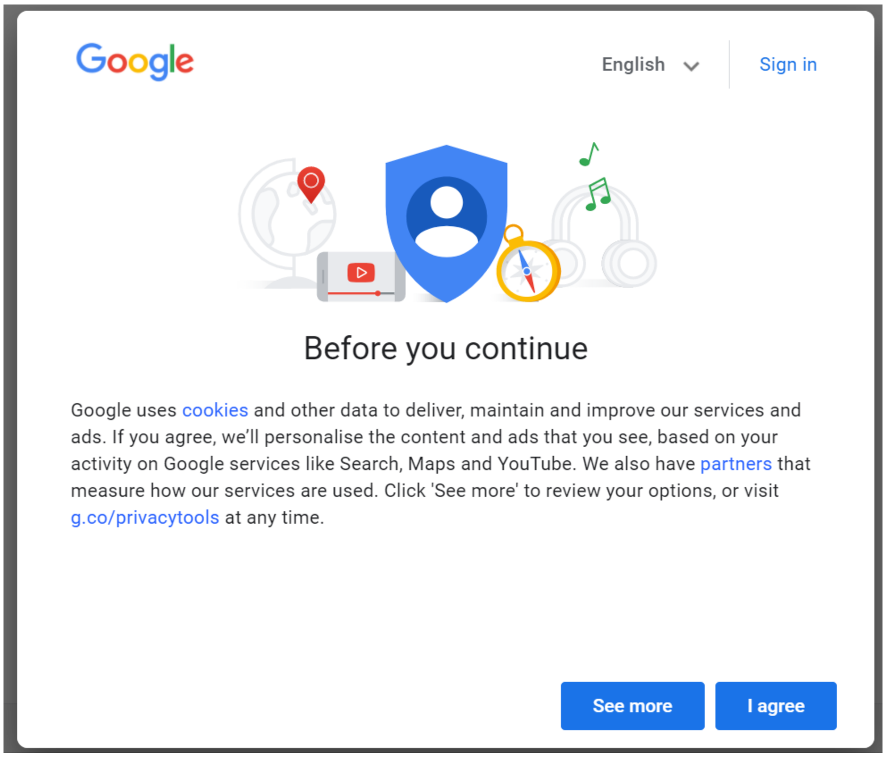

Figure 2. Google's cookie notice at the time of Study 1.

5.2 Material

I designed an *ad hoc* webpage representing the homepage of a website specialized in food recipes called Foodit. If I used existing websites, the participant might have visited them already, and have their cookie installed onto their device; this would prevent the cookie notice to appear during the study. In all the studies, the website included a logo, the images of some dishes, and an element the participants could interact with (a search bar in Study 1 and Study 2, download buttons in Study 3). The link to the website was provided in a starting window after collecting the consent to participate. Immediately after landing on this webpage or a few seconds later, a notice appeared, preventing any other action, and containing two buttons. The buttons, modeling cookie notices that were common at the time of the studies, allowed either to accept the request in the notice or to personalize the privacy settings (Study 1 and Study 2) or to deny the consent (Study 3). The webpage was built in PsychoPy3 (https://www.psychopy.org/) and presented through Pavlovia (https://pavlovia.org/).

To manipulate contextualization, the notice was programmed to appear immediately after the participants' arrival on the website or after having let them the possibility to explore it. To manipulate explicitness, one version of the cookies notice mentioned the triggering action.

The button selected by the participant and the time taken to select the button (time to respond, in seconds) were automatically recorded and stored as part of the data collected. Other data were collected via a survey, self-reported by the participant. To do so, I needed some features such as skip logic and the randomization of the answer options that were not available in Pavlovia. On the other hand, exiting latter to connect to a different online survey platform required a few seconds wait to save the data. So, I included in Pavlovia the few items testing the participants' comprehension of the notice, to display them right after the participant clicked on the notice itself. The rest of the items (perceived comprehension, age, gender, level of education, and privacy concerns) were answered via SurveyMonkey, where the participant was automatically transferred.

5.3 Measures

The measures collected in the studies included the dependent variables of interest, namely the comprehension of the cookie notice, the experience of comprehension and the response to the notice. In addition, a set of control variables was collected, i.e., age, gender, level of education, and privacy concerns.

5.3.1 Comprehension

Comprehension, i.e., the users' ability to discern appropriate meaning (Smith & Taffler, 1992), was assessed by administering two multiple-choice questions, and then counting the rate of correct answers; this strategy is very common when assessing comprehension (e.g., Balebako et al., 2015). Since the text of the privacy notice was very short, I only administered one multiple-choice test related to the topic of the notice. In addition, I administered a multiple-choice item related to the cause of the notice.

Cause comprehension was assessed with the item: "What caused the alert's appearance?" and three answer options, "Entering the website" (correct), "Something that I did on the website" or "None of the above". The participant was asked to select one option. *Topic comprehension* was tested with the item: "What was the alert trying to communicate?" and three answer options, "Foodit website will show advertising", "Foodit website needs permission to use cookies" (correct), "Foodit website collects user's data during the navigation". The participant was asked to select one option. The results of these two items were not consolidated in one sole comprehension score, but were kept distinct, because I was interested in studying how each of the two aspects was affected by the manipulations.

5.3.2 Perceived comprehension, clarity, and control

Perceived comprehension is the subjective experience of having a good grasp of the content of the notice; it might not be related to the actual comprehension (e.g., Kobsa & Teltzrow, 2004). Inspired by similar items in Knijnenburg & Cherry (2016), I measured the perceived comprehension with two items: "I would feel confident answering questions about the contents of Foodit alert" and "After reading Foodit alert, I understand the implications of visiting Foodit website". The participants answered by expressing their agreement on a 5-point Likert scale (from "Strongly disagree" to "Strongly agree"). These items were prefaced by a multiple-choice question asking participants whether they read the notice (“Did you read the content of Foodit alert?” (“Yes”, “Approximately”, “No”, “I don’t know”). The participants answering “No” or “I don’t know" were not administered the perceived comprehension items.

Clarity is the easiness of understanding and depends on the sentence structure and language (Hargis, 2000). Clarity is central for correctly processing information and is a pre-requirement of the understandability of a notice (Antòn et al., 2004). It was measured by collecting the agreement with three statements adapted from Knijnenburg & Cherry (2016) on a 5-point Likert scale (from "Strongly disagree" to "Strongly agree"): "The language of Foodit alert was clear", "Foodit alert was difficult to understand" and "I feel that the information in Foodit alert was explained clearly". These items were presented only to the participants who answered "yes" or "approximately" to the filter item described above.

Sense of control is the extent to which the participants feels that they control the way the website uses their data; a transparent notice that is truly understood should be accompanied by this experience (Tsai & Brusilovsky, 2021). The sense of control was measured with an item adapted from Xu et al. (2009): "I believe I have control over how personal information is used by Foodit website". The level of agreement was measured on a 5-point Likert scale (from "Strongly disagree" to "Strongly agree").

5.3.3 Response to the notice

The response to the notice was measured by recording *which button* the participant selected on the notice, whether “accept” or “set options”. The participants could only select one button, since the pressure of the button led them directly to the survey. I also measured the *time* elapsed, in seconds, since the appearance of the notice till the pressure of the button.

5.3.4 Control variables

The control variables in this study were privacy concerns, some demographic variables, and the familiarity with cookie notices. The items investigating the level of *privacy concerns* were selected and adapted from Antòn et al. (2010) and used a 5 points Likert scale (from "Strongly disagree" to "Strongly agree"):

1. I mind when a website uses cookies to customize my browsing experience.
2. I want a website to disclose how my information will be used.
3. I mind when the information about myself is shared with third parties.
4. I mind when a website that I visit collects (without my consent) information about my browsing patterns.
5. I am concerned about unauthorized employees getting access to my information.

Demographic variables. As in Milne & Culnan (2004), Klumpe et al. (2020), and Herbert et al. (2021), I asked participants for their gender, age, and education, as they are typical demographic variables, although they do not always have a direct effect on privacy-related behavior; Biselli and Reuter (2021), for instance, found no effect of these variables on privacy behavior. This demographic information was collected through 2 multiple-choice items (G*ender*: female/male/other; *Level of education*: "Which level of education have you reached?" Primary School/High School/Bachelor's degree/Master's degree/Ph.D.) and 1 open-answer item (*Age*: "How old are you?").

Familiarity. To check if the participants behaved similarly to a daily life situation, I introduced two multiple-choice items focused on the *habits* with similar notices and on the estimation of the *novelty of the notice*: "Did you notice any difference between that Foodit alert and any alert you usually find on the Internet?" (Yes/No), and "In my real life, what do you do when you see a similar alert?" ("I click on Accept"; "I ignore it if it is possible"; "I enter the options").

5.4 Procedure

All participants were recruited online through a platform called Prolific (https://www.prolific.co). The invitation included the link to enter the study platform. Separate links were sent based on the participants' gender, to reach an equal number of male and female participants.

The inclusion criteria, checked with the filters provided by Prolific, were to reside in Europe or the United Kingdom, where GDPR or GDPR-compliant data protection policies are adopted; to be English native speakers, since the content was in English; to accept studies with deception (considering that people who agreed with this last option may receive invitations to studies with or without deception, this filter did not provide any clue about the nature of the study). As is recommended in online studies to improve the quality of the sample (Oppenheimer et al., 2009; Shamon & Berning, 2019), I inserted two *attention checks* in the questionnaires; failure to correctly reply to these attention checks suggested that the participant was just clicking randomly on the response scales without reading the items, and was then excluded from the dataset. I modeled the attention checks to be fair and in line with Prolific guidelines, i.e., making it possible for respondents who were paying attention to pass the check without needing to remember any information. The attention checks were merged into the rest of the questions and read as follows: "It is important that you pay attention to this survey. Please, tick

"Disagree". Participants were alerted of the presence of these attention checks in the initial informed consent for ethical reasons, since the exclusion from the dataset involved the exclusion from the compensation.

The participants were serially assigned to one of the four experimental conditions (e.g., the first participant to the explicit-consecutive condition, the second one to the explicit-delayed condition, etc.). After entering the online study platform, the experimental session started (Figure 3). The participants were first displayed the information notice and asked their consent to participate. If they accepted, they were given the task instructions in a subsequent screen, i.e. reach the Foodit website, look for the "Carbonara" recipe, and find out when the eggs are supposed to be added. This task was a cover task, allowing to expose them to the cookie notice. They would click on the 'Go to the website' button embedded in the instructions screen and reach the Foodit webpage and be shown the cookie notice. Upon clicking on either button on the notice, they were informed that the search task was completed as far as the study was concerned and that they were about to be led to the survey. At the end of the survey, a second informed consent was displayed disclosing the goal of the study.

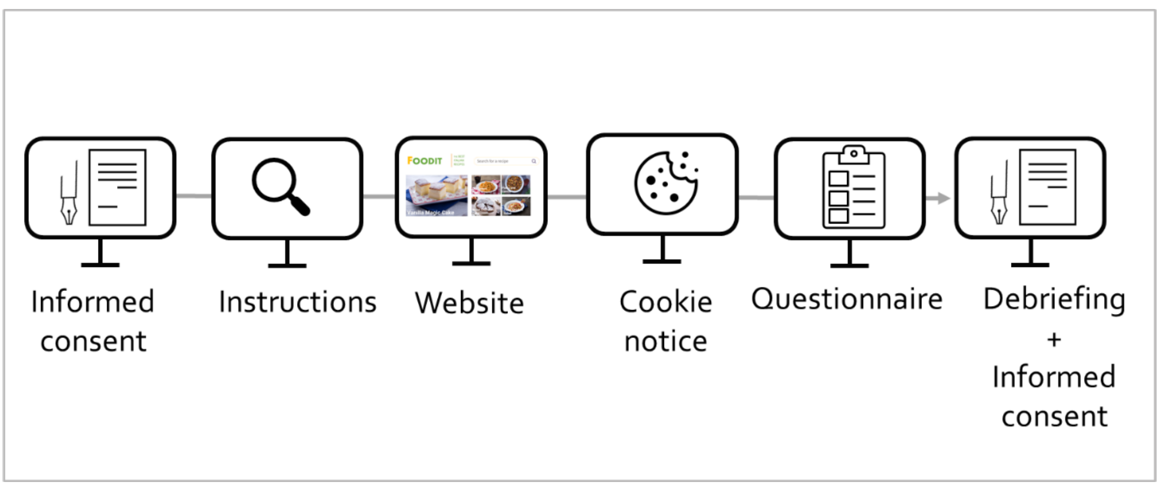

Figure 3: Outline of the procedure.

5.5 Ethics

The studies were approved by the HIT ethical committee of the University of Padova (2021_99R, 2021_99R replica, 2021_99R replica2). The participants were asked two informed consents, one before starting the experimental session and one after it. The initial consent provided all possible information except a full disclosure of the goal of the study, to prevent influencing their behavior. They were told that the study wanted their opinion on the usability of a website. The participants received monetary compensation in line with Prolific fair standards (7.50£/hour: https://www.prolific.co/pricing/); the only constraint for receiving the reward was the completion of the questionnaire and to successfully pass the attention checks. The participants were free to withdraw from the study by closing the browser window; in that case, the monetary compensation would be lost. After completing the surveys at the end of the study, I asked for their consent again through a second informed consent disclosing the real goal of the study. If the consent was denied, I checked if they passed the attention checks to proceed with the payment; then, their data were deleted. The only identifying information I collected was the participants' Prolific ID (the alphanumeric code that identifies participants on Prolific and that ensures their anonymity) to be able to pay them after the experiment and exclude participants who already participated in a study from each subsequent study. After the payment, this information was replaced in the dataset with a progressive code. The participants who failed the attention check or did not accept to participate or who abandoned the survey were led to a thank-you page explaining what was entailed in terms of compensation.

Section I

CHAPTER 6.
STUDY 1

Cookie consent notices inform the visitors about the type and purpose of the data collected and ask permission to process it. Entrance is the usual circumstance triggering the display of a cookie notice, to collect the users' data since their earliest interactions with the website. I compared a condition in which the notice appears right at the entrance on the website (consecutive condition), to a condition in which the notice appears a few seconds later when the content of the website has already been displayed (delayed condition). In the former condition, a sequential tie is established between the entrance and the notice, whereas in the latter condition such tie is broken by interposing events such as the appearance of all the content of the website and, possibly, the users' first exploration of it. I also manipulated the explicit reference to the triggering action in the notice; in the explicit condition, the title of the notice stated the cause of its appearance, in the no-mention condition it did not. The specific hypotheses of this study are the following:

Effect of consecutiveness on comprehension: compared with a delayed notice, a notice appearing right after its trigger improves the users' comprehension of its cause (**H1a**), topic (**H1b**), and the users' sense of comprehension and control (**H1c**).

Effect of explicitness on comprehension: mentioning the trigger in the notice improves the users' comprehension of its cause (**H2a**), topic (**H2b,**), and the users' sense of comprehension and control (**H2c**).

Transparency paradox: the acceptance of the notice is more frequent in the consecutive condition (**H3a**), in the explicit condition (**H3b**), and with a higher perceived comprehension (**H3c**).

6.1 Material

During this study, the homepage of Foodit website included a logo, the images of some dishes, and a search bar (Figure 4, top left). Immediately after landing on this webpage or a few seconds later, a notice appeared, preventing any other action, and containing two buttons. The buttons, modeling cookie notices that were common at the time of the study, allowed either to accept the request in the notice or to personalize the privacy settings (Figure 4, top right).

To manipulate contextualization, the notice was programmed to appear either 0.6 seconds (consecutive conditions) or 5.5 seconds (delayed conditions) after the participant's arrival on the website. To manipulate explicitness, one version of the cookies notice mentioned the triggering action (explicit conditions, Figure 4, down-right), and a second version did not (no-mention condition, Figure 4, down-left).

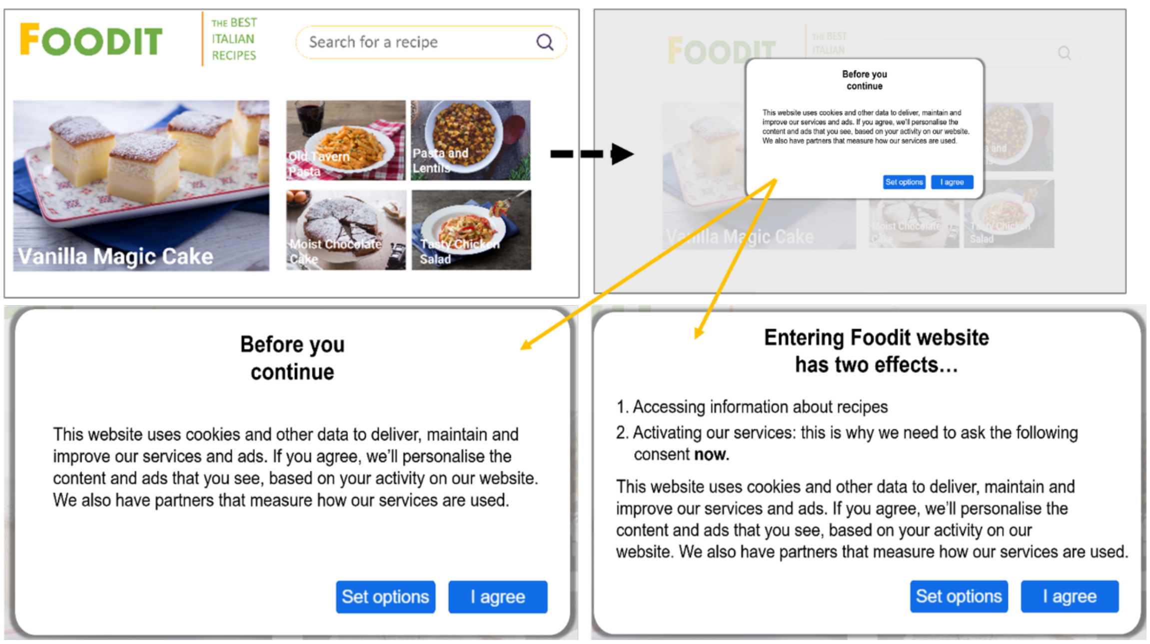

Figure 4: The website at the participants' entrance (top left) and when the cookies notice appeared (top right); the notice in the no-mention conditions (bottom left) and in the explicit mention conditions (bottom right).

6.2 Participants

The initial sample was composed of 159 people. I excluded from the analysis 27 participants: 15 for inappropriate completion of the questionnaire (i.e., response patterns, failure at the attention checks), 11 encountered technical issues during the study, and 1 for answering "Other" to the gender item and was then not possible to run a separate analysis for this gender category since its numerosity was too low for being included as a variable level in the regressions built for analyzing the data. The final sample after removing invalid responses consisted of 132 people aged 18 to 76 years (M = 30.99; SD = 11.02; 66 males, 66 females), all residing in the EU or the UK.

6.3 Data and Analysis

The answers to the items measuring the actual comprehension of the cause and the topic of the notice were re-coded as "Correct" (cause: *N* = 80; topic: *N* = 98) and "Incorrect" (cause: *N* = 52; topic: *N* = 34); age was re-coded into "Under 30" (*N* = 65) and "Over 30" (N = 67) with reference to the median of the sample; the level of education was re-coded into "High School" (*N* = 56) and "University" ("Bachelor' degree", "Master's degree", "Ph.D.", *N* = 76); during the studies, no participant selected "Primary school".

All analyses were run in RStudio (v. 1.4.1717). Since my interest was to understand the effect of the independent variables (consecutiveness and explicitness) and their interaction, considering at the same time the effect of the control variables (age, gender, education, privacy concerns), I opted for regression analyses. This method enables the assessment of the effect of one variable considering the effect of the other variables included in the model; it is appropriate for testing interactions and allows to deal with independent and dependent variables measured at a continuous, ordinal or nominal level. Regression models were built (Model 1) for every dependent variable (comprehension of the cause of appearance, comprehension of the topic, perceived comprehension, clarity, sense of control, response), inserting the independent variables, and the control variables as predictors. Then, following a stepwise procedure, I built a second model (Model 2) adding the interaction between explicitness and consecutiveness. To interpret the direction of the coefficients in the models, I set the following references for the contrasts: *delayed* for consecutiveness, *no-mention* for explicitness, *female* for gender, *under 30* for age, *university* for education, *not correct* for comprehension, and *set options* for response. Each Model 1 was then compared with the corresponding Model 0 (null model) employing goodness of fit tests. This comparison allowed to verify whether the inclusion of the predictors in Model 1 explained the data variability better than mere chance. The results of the significant goodness of the fit tests are reported in the Appendix.

Since the dependent variables were measured on different scales, I employed logistic, ordinal, or linear regressions based on the nature of the dependent variable considered. Logistic regressions were employed for binary outcomes (i.e., comprehension of the cause of the notice appearance, comprehension of the topic of the notice, response to the notice). Logistic regressions were also used to assess the influence of privacy concerns (based on the average score of the items, since the alpha of the scale was acceptable: $\alpha = .74$) and of the perceived comprehension on the participants' response to the notice. Linear regressions were run for continuous dependent variables, i.e., the time to respond to the notice (in seconds) and the notice's clarity. Since clarity was a multi-item scale, after having reversed the scores for the item "Foodit alert was difficult to understand", I calculated the alpha of the scale. The alpha was good ($\alpha = .87$), so I ran the analyses on the average score of the scale. Given the ordinal nature of the Likert-scale items when considered individually, ordinal regressions were employed for the sense of control and the perceived comprehension, considered separately. The regression results (Model1 and Model 2 coefficients) for models that could not explain the variability of the data better than the chance are reported in the Appendix.

I checked if the assumptions of the different regressions were met. First, since the presence of strong correlations among the predictors (multicollinearity) reduces the possibility to discriminate their effect on the outcome, I assessed multicollinearity through the Variance Inflation Factor (VIF) before running every regression. If the VIF was less than 2.5, I excluded collinearity. No variable in my studies showed a VIF greater than 2.5. Second, I checked the linearity, normality, homoscedasticity assumptions, and the absence of influential outliers in the linear regressions through a visual inspection of the residuals. Third, I verified the proportional odds assumption for ordinal regressions with the Brant test: were the test significant, the assumption would be violated. No evidence of violation of the proportional odds assumption was found during my studies.

The effect size of the model and the predictors was assessed considering the proportion of variance explained (R^2) by them. For single predictors, I included the predictor alone and checked the R^2 resulting from the comparison with Model 0. Different estimates can be more or less suitable depending on the kind of regression employed. So, I will report the adjusted R^2 for linear regressions, the Hosmer and Lemeshow's R^2 for logistic regressions, and the McFadden pseudo R^2 for the ordinal ones.

The sample size for each predictor is reported in Table 1. The subsample (N = 108) composed of participants declaring to have read the notice (called 'readers') was used to test the effect of the predictors on the notice's perceived comprehension, clarity, and on the time to respond.

Table 1: Sample size by the predictor's levels for the complete sample and the readers' subsample

Variable	Level	Complete sample (N)	Readers' subsample (N)
Consecutiveness	Consecutive	65	47
	Delayed	67	61
Explicitness	Explicit mention	65	54
	No mention	67	54
Gender	Female	66	53
	Male	66	55
Education	High School	56	44
	University	76	64
Age	Under 30	67	50
	Over 30	65	58

6.4 Results

Table 2 provides an overview of the descriptive statistics of Study 1. In the following sections, I report the results of the statistical analyses organized by groups of dependent variables.

Table 2: Descriptive statistics of the outcomes depending on the predictors.

Variable	Level	Cause	Topic	Perceived Comp. (item1)		Perceived Comp. (item2)		Clarity		Sense of control		Time to respond (sec)		"I agree" Response
		% correct	% correct	M	SD	M	SD	M	SD	M	SD	M	SD	%
Consecutiveness	Consecutive	89.23	76.92	3.66	0.87	3.11	1.03	3.80	0.77	3.25	1.05	4.58	3.50	95.39

Variable	Level	Cause	Topic	Perceived Comp. (item1)		Perceived Comp. (item2)		Clarity		Sense of control		Time to respond (sec)		"I agree" Response
		% correct	% correct	M	SD	M	SD	M	SD	M	SD	M	SD	%
	Delayed	32.84	71.64	3.41	0.86	3.08	1.04	3.68	0.73	2.87	1.03	7.43	8.08	83.58
Explicitness	Explicit	52.31	66.15	3.57	0.84	3.22	0.99	3.67	0.72	3.06	1.01	8.21	8.67	90.77
	No mention	68.67	82.09	3.46	0.91	2.96	1.06	3.80	0.77	3.05	1.09	4.17	2.27	88.06
Gender	Female	57.58	74.24	3.45	0.91	3.00	1.04	3.67	0.77	2.97	0.99	6.39	6.29	95.46
	Male	63.64	74.24	3.58	0.83	3.18	1.02	3.79	0.72	3.14	1.11	6.00	6.98	83.33
Education	High School	62.50	73.21	3.57	0.89	3.14	1.07	3.77	0.79	3.05	1.00	6.02	7.77	91.07
	University	59.21	75	3.48	0.85	3.06	1.01	3.70	0.72	3.05	1.09	6.31	5.77	88.16
Age	Under 30	55.22	76.12	3.45	0.86	2.98	1.00	3.71	0.70	3.08	1.11	4.59	2.91	86.57
	Over 30	66.15	72.31	3.60	0.88	3.22	1.06	3.76	0.80	3.03	1.00	8.04	8.91	92.31

6.4.1 Effect of consecutiveness and explicitness on comprehension (H1a,b, and H2a,b)

I hypothesized that consecutive and explicit notices increased the users' comprehension of the cause and topic of the notice. Figure 5 shows the results of the comprehension of the cause (left) and of the topic of the notice (right).

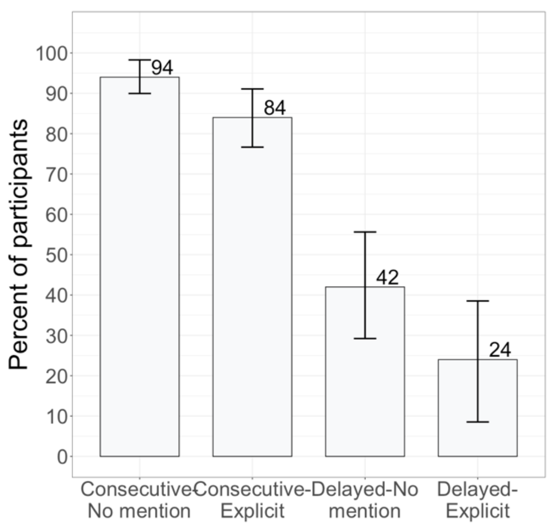

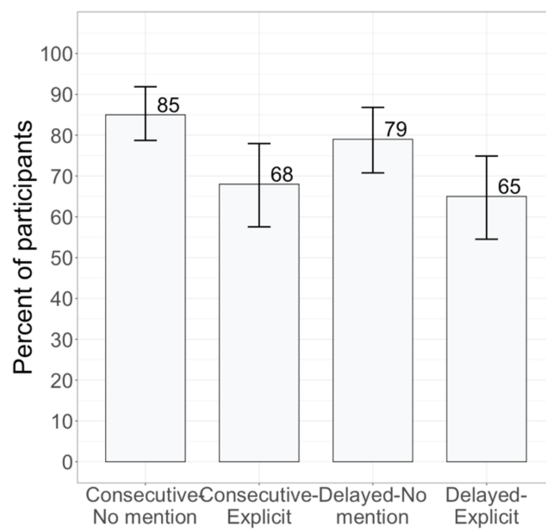

Figure 5. Percentage of participants correctly identifying the action triggering the appearance of the notice in the four conditions (Cause Comprehension - left) and the topic of the notice (Topic Comprehension – right). The error bars represent the standard error.

I tested whether there was any statistically significant difference between the four conditions in the comprehension of the cause of the notice, following the procedure already described in Section 6.3. The VIF of the predictors included in Model 1 was less than 5, excluding the presence of multicollinearity (Explicitness = 1.11; Consecutiveness = 1.08; Gender = 1.06, Education = 1.05; Age = 1.09), so the absence of multicollinearity assumption held. The inclusion of the predictors in Model 1 was useful for explaining the variability in the data; overall, Model 1 could explain the 32% of data variability (R^2 = .32). The coefficients of Model 1 (Table 3) showed a main effect of consecutiveness in identifying the cause of the appearance of the notice: being in the consecutive conditions increased the odds that the participants selected the correct cause of the notice, i.e., their entering the website. This predictor alone explained 27% of the variability in the data (R^2 =.27). Model 1 also showed an effect of explicitness in identifying the cause of the appearance of the notice in the direction of reducing the likelihood of a correct response. However, the model in which only explicitness was specified as predictor (Model 1.2) was not better than Model 0 in explaining the variability of data, $X^2(1) = 3.71$, p = .05, R^2 =.02. Hence, the overall effect of explicitness was not confirmed. Model 2 did not reveal an effect of the interaction between consecutiveness and explicitness. So, I could find evidence for H1a, but I could not refuse the null hypothesis for H2a.

Table 3: Regression results for Cause Comprehension. *p <.05, **p<.01,***p<.001

		b	Std. Error	z	Odds Ratio	p
	(Intercept)	-0.92	0.54	-1.70	0.40	.09.
	Explicitness	-1.07	0.48	-2.24	0.34	.03 *
MODEL1	Consecutiveness	2.98	0.51	5.84	19.69	<.001***
direct effects	Gender	0.62	0.47	1.34	1.87	.18
	Education	0.13	0.47	0.26	1.13	.79
	Age	0.66	0.47	1.40	1.94	.16
MODEL1.1	(Intercept)	-0.72	0.26	-2.75	0.49	.006**
effect of consecutiveness	Consecutiveness	2.83	0.48	5.93	16.95	<.001***
MODEL1.2	(Intercept)	0.78	0.26	2.98	2.19	.003**
effect of explicitness	Explicitness	-0.69	0.36	-1.91	0.50	.06
	(Intercept)	-1.01	0.58	-1.76	0.36	.08
	Explicitness	-0.93	0.56	-1.66	0.40	.10
MODEL 2	Consecutiveness	3.29	0.84	3.94	26.93	<.001***
direct effects	Gender	0.64	0.47	1.374	1.89	.17
+	Education	0.14	0.48	0.30	1.15	.77
interaction	Age	0.69	0.48	1.46	2.00	.15
	Explicitness: Consecutiveness	-0.52	1.05	-0.50	0.59	.62

I then considered the comprehension of the *topic* of the notice (Figure 5). The VIF of the predictors included in Model 1 was less than 5, excluding the presence of multicollinearity (Explicitness = 1.09; Consecutiveness =1.02; Gender = 1.00, Education = 1.05; Age = 1.09). Model 1 was not different from Model 0 at the likelihood ratio test ($X^2(5) = 5.25$, $p = .39$, $R^2 = .03$) but its coefficients showed a significant effect of explicitness (Table 4). I built a model in which I specified only explicitness as predictor (Model 1.1) which confirmed the presence of an effect. This effect was the opposite of what was hypothesized: being in an explicit condition reduced the odds of correctly identifying the topic of the notice. Nonetheless, this effect was small ($R^2 = .03$). Model 2 was not different from Model 0 at the likelihood ratio test, $X^2(6) = 5.37$, $p = .50$, $R^2 = .04$, and revealed no significant influence of the predictors on the comprehension of the topic of the notice. In conclusion, I could not refuse the null hypothesis for H1b and H2b.

Table 4: Regression Results for Topic Comprehension. *p <.05, **p<.01,***p<.001

		b	Std. Error	z	Odds Ratio	p
	(Intercept)	1.54	0.51	3.05	4.66	.002**
	Explicitness	-0.88	0.43	-2.05	0.41	.04*
MODEL 1	Consecutiveness	0.27	0.41	0.66	0.27	.51
direct effects	Gender	0.02	0.41	0.06	1.02	.95
	Education	-0.27	0.42	-0.63	0.77	.53
	Age	-0.06	0.42	-0.13	0.94	.89
MODEL 1.1	(Intercept)	1.52	0.32	4.78	4.58	<.001***
effect of explicitness	Explicitness	-0.85	0.41	-2.07	0.43	.04*
	(Intercept)	1.49	0.57	2.56	4.26	.01*
	Explicitness	-0.75	0.57	-1.33	0.47	.18
MODEL 2	Consecutiveness	0.44	0.65	0.69	1.56	.49
direct effects	Gender	0.03	0.41	0.08	1.03	.94
+	Education	-0.26	0.42	-0.62	0.77	.54
interaction	Age	-0.04	0.43	-0.08	0.97	.93
	Explicitness: Consecutiveness	-0.29	0.84	-0.35	0.75	.73

6.4.2 Effect of consecutiveness and explicitness on Perceived Comprehension and Control (H1c and H2c)

The scores of the users' perceived comprehension, clarity, and sense of control are shown in Figure 6 and Figure 7 below. The tests for perceived comprehension and sense of control were based on ordinal regressions, while the test for clarity was based on linear regression. The VIF for the predictors described in this section is reported in Table 5. The results of the Brant test for the ordinal regressions showed that the proportional odds assumption was respected since no test was significant (Table 6).

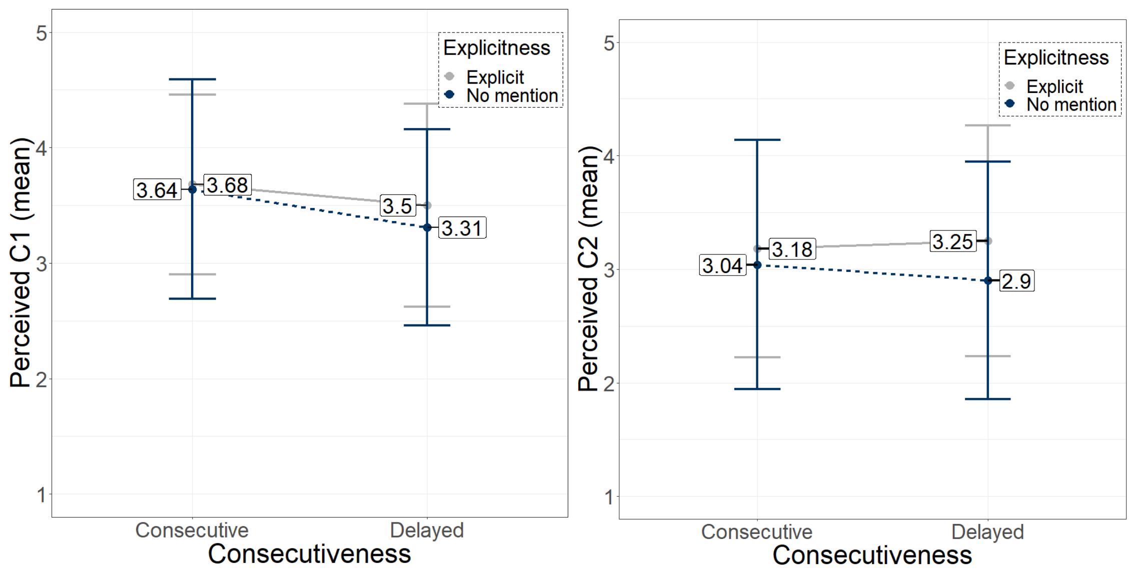

Figure 6: Effect of consecutiveness and explicitness on the scores of the two items measuring the perceived comprehension (1: I would feel confident answering questions about the contents of Foodit alert; 2: After reading Foodit alert, I understand the implications of visiting Foodit website).

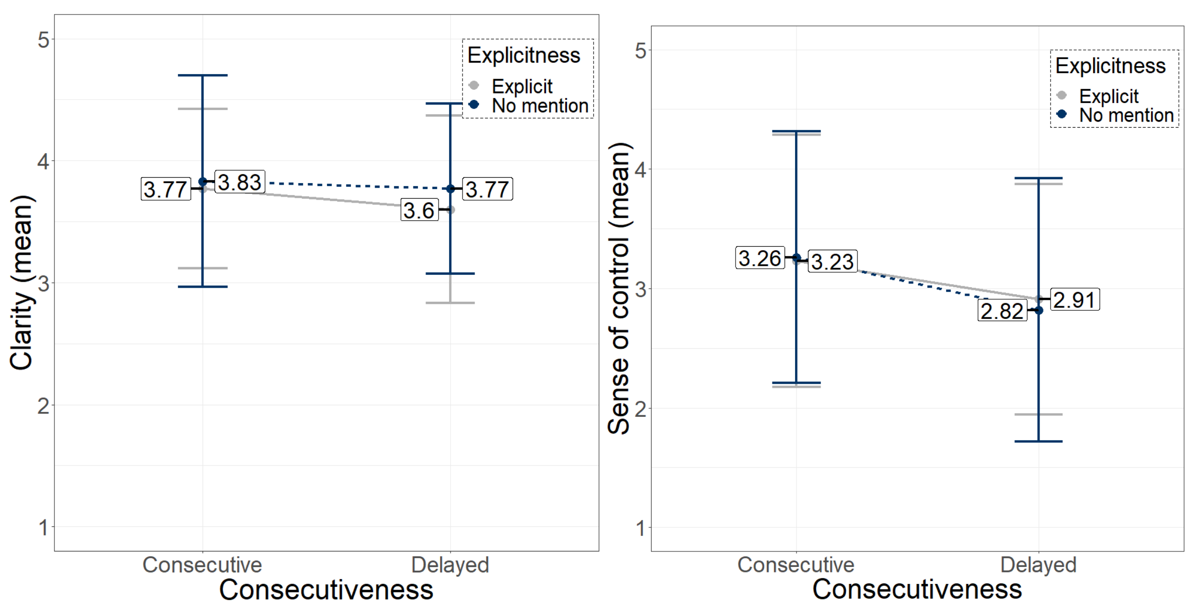

Figure 7: Effect of consecutiveness and explicitness on clarity (average score of the three items “The language of Foodit alert was clear", "Foodit alert was difficult to understand (reverse)" and "I feel that the

information in Foodit alert was explained in a clear manner") and sense of control ("I believe I have control over how personal information is used by the Foodit website").

Table 5: VIF of the predictors depending on the outcome.

	Perceived Comp. 1	Perceived Comp.2	Clarity	Sense of control
Explicitness	1.09	1.09	1.09	1.09
Consecutiveness	1.01	1.01	1.01	1.01
Gender	1.02	1.02	1.02	1.00
Education	1.05	1.05	1.05	1.04
Age	1.07	1.07	1.07	1.08

Table 6: Results of the Brant test depending on the outcome.

	Perceived Comp. 1			Perceived Comp.2			Sense of control		
	X^2	df	p	X^2	df	p	X^2	df	p
Omnibus	8.53	12	.74	7.38	18	.99	17.17	18	.51
Explicitness	0.73	2	.69	0.12	3	.99	0.19	3	.98
Consecutiveness	3.76	2	.15	0.33	3	.96	2.61	3	.46
Gender	1.26	2	.53	0.89	3	.83	1.82	3	.61
Education	0.16	2	.92	1.17	3	.76	3.68	3	.30
Age	1.83	2	.40	1.32	3	.72	2.97	3	.40
Consecutiveness: Explicitness	1.48	2	.48	2.01	3	.57	0.21	3	.98

Model 1 (i.e., the model in which the independent variables – consecutiveness and explicitness, and the control variables – gender, age and education - were specified as predictors) and Model 2 (i.e., Model 1 with the addition of the interaction between timeliness and explicitness) could not explain the variability of data better then Model 0 at the likelihood ratio test neither for the *perceived comprehension* (Model 1 for item1: $X^2(5) = 5.54$, $p = .35$, $R^2 = .02$; Model 1 for item2: $X^2(5) = 4.18$, $p = .52$, $R^2 = .01$; Model 2 for item1: $X^2(6) = 6.09$, $p = .41$, $R^2 = .03$; Model 2 for item2: $X^2(6) = 4.68$, $p = .59$, $R^2 = .02$), nor for the *sense of control* (Model1: $X^2(5) = 6.06$, $p = .30$, $R^2 = .02$; Model 2: $X^2(6) = 6.29$, $p = .39$, $R^2 = .02$). Nonetheless, the results of Model1 for sense of control showed a significant effect of consecutiveness (Table 7). So, I built a model (Model 1.1) in which I specified only consecutiveness as predictor. This model showed that the participants included in the consecutive conditions significantly felt to be more in control of their data compared to the participants who were in the delayed conditions. Even if this finding was significant, only 2% of the variability in the data could be explained by this factor alone

(R^2 = .02). Model 2 showed no effect of the interaction between explicitness and consecutiveness.

Table 7: Regression results for Sense of Control. *p <.05, **p<.01,***p<.001

		b	Std. Error	z	p
MODEL 1	Explicitness	0.15	0.33	0.45	.65
direct effects	Consecutiveness	0.73	0.33	2.25	.02*
	Gender	0.32	0.32	1	.32
	Education	-0.02	0.32	-0.05	.96
	Age	-0.11	0.33	-0.34	.74
MODEL 1.1 effect of consecutiveness	Consecutiveness	0.69	0.32	2.16	.03*
MODEL 2	Explicitness	0.30	0.46	0.66	.51
direct effects	Consecutiveness	0.89	0.46	1.94	.05
+	Gender	0.34	0.32	1.04	.30
interaction	Education	-0.01	0.33	-0.02	.98
	Age	-0.09	0.34	-0.26	.80
	Explicitness: Consecutiveness	-0.32	0.65	-0.49	.63

The overall goodness of fit test showed that the manipulations did not have any effect on *clarity* (Model1: F (5, 102) = 0.49, p = .78, R^2 = -.03; Model2: F (6, 101) = 0.41, p = .87, R^2 = -.03). While H2c was not supported by the data, H1c was partially confirmed for the sense of control. In conclusion, the consecutiveness seemed to slightly increase the sense of control; none of the contextualization strategies affected the level of perceived comprehension and clarity.

6.4.3 Participants' Response to the Notice (H3a,b,c, H4, H5)

The distribution across conditions of the two possible responses to the notice and the time to respond to the notices are reported in Figure 8. First, I tested the effect of the independent variables (consecutiveness and explicitness) and of the control variables (gender, age, education, and privacy concerns) on the participants' *response* to the notice (I agree/Set options). The VIF of the predictors was less than 5 (Explicitness = 1.11, Consecutiveness = 1.08, Gender = 1.04, Education = 1.08, Age = 1.21, Privacy concerns = 1.12), so the multicollinearity assumption was met. The inclusion of the predictors in Model 1 was useful for explaining the variability in data; overall, Model 1 could explain the 24% of data variability (R^2 = .24). Model 1 showed a main effect of privacy concerns (R^2 = .11) (Table 8): when the level of privacy concerns increased, the odds of selecting the "Set Options" button increased as an effect of this control variable. The addition of the interaction between explicitness and consecutiveness in Model 2 increased the variability explained by the model (R^2 = .27) and revealed an effect of consecutiveness (R^2 = .06), robust at the likelihood ratio test, $X^2(1) = 5.13$, p = .02, R^2 = .06. Being in a consecutive condition increased the odds of selecting "Accept", thereby supporting the transparency paradox (H3a); being in an explicit condition had no effect on the response to the notice (H3b). Nonetheless, privacy concerns and consecutiveness alone explained little of the total variability in the data. Finally, to test the effect of the perceived comprehension on the response to the notice (H3c), I built Model 3 introducing only the perceived comprehension item 1 ("I would feel confident answering questions about the contents of Foodit alert") as a predictor, and the response as the outcome. The introduction of the predictor did not explain the variability in the data better than chance alone, $X^2(1)$ = 1.12, p = .29, R^2 = .01. H3c was then not supported.

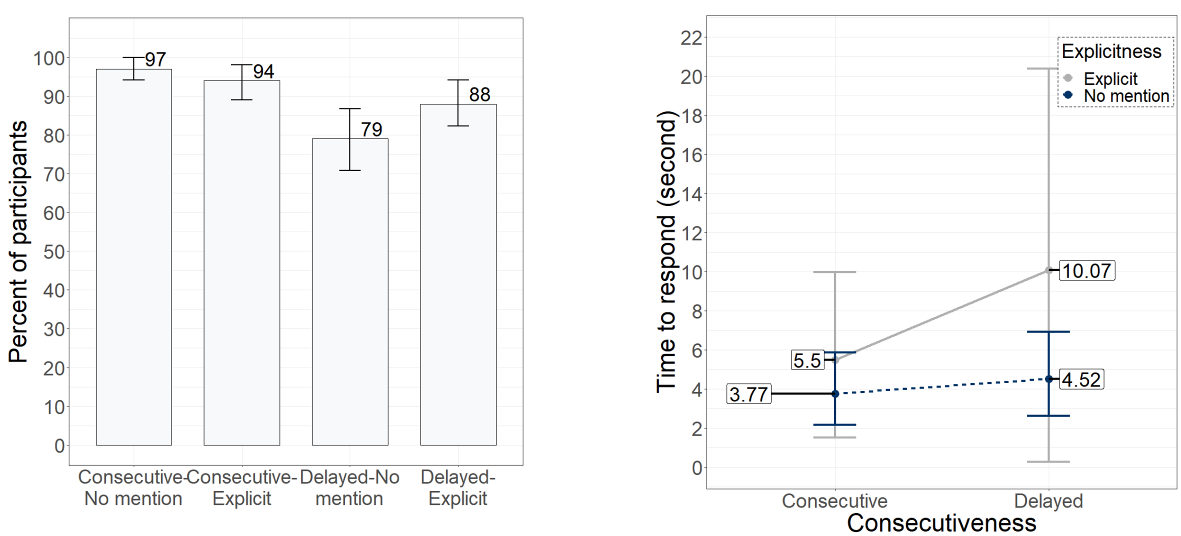

Figure 8: Distribution of the ‘accept' response to the notice by experimental condition (the error bars represent the standard error - left) and effect of the main variables on the response time (right).

Table 8: Regression results for Response. *p <.05, **p<.01,***p<.001

		b	Std. Error	z	Odds Ratio	p
MODEL 1	(Intercept)	9.68	3.02	3.21	16037.89	.001**
main effects	Explicitness	0.12	0.67	0.18	1.13	.86
	Consecutiveness	1.40	0.75	1.88	4.07	.06

		b	Std. Error	z	Odds Ratio	p
	Gender	-1.23	0.72	-1.70	0.29	.09
	Education	0.40	0.68	0.59	1.49	.56
	Age	0.90	0.72	1.25	2.45	.21
	Privacy concerns	-1.92	0.70	-2.74	0.15	.006**
MODEL 1.1	(Intercept)	8.72	2.54	3.43	6110.75	<.001***
effect of privacy concerns	Privacy concerns	-1.61	0.59	-2.73	0.20	.006**
MODEL 2	(Intercept)	12.14	3.66	3.32	4775.58	.001**
main effects	Explicitness	0.77	0.79	0.97	2.15	.33
+	Consecutiveness	2.77	1.24	2.24	16.03	.03 *
interaction	Gender	-1.26	0.75	-1.69	0.28	.09
	Education	0.51	0.70	0.73	1.67	.46
	Age	-1.36	0.80	-1.71	3.91	.09
	Privacy concerns	-2.29	0.79	-2.90	0.10	.004**
	Explicitness: Consecutiveness	-2.81	1.66	-1.70	0.06	.09
MODEL 2.1	(Intercept)	1.63	0.33	4.94	5.09	<.001***
effect of consecutiveness	Consecutiveness	1.40	0.68	2.07	4.05	.04*
MODEL 3	(Intercept)	0.74	1.09	0.68	2.10	.50
effect of PerceivedC1	PerceivedC1	0.34	0.32	1.07	1.40	.29

Second, I tested the time to respond to the notice. The inclusion of the predictors in Model 1 was useful for explaining the variability in the *time to respond*; overall, Model 1 could explain the 14% of data variability (R^2 = .14). The VIF of the predictors was less than 5, excluding the presence of multicollinearity (Explicitness = 1.09, Consecutiveness = 1.01, Gender = 1.02, Education = 1.05, Age = 1.07), confirming this assumption. The model showed a main effect of explicitness (R^2 = .09), consecutiveness (R^2 = .04), and age (R^2 = .06): the participants in the explicit conditions, in the delayed conditions, and more than 30 years old took longer to select a button (Table 9). Nonetheless, the data showed great variability and little was explained by these three factors alone. Model 2 did not show any effect of the interaction between consecutiveness and explicitness.

Table 9: Regression results for Time to respond. *p <.05, **p<.01,***p<.001

		b	Std. Error	z	p
MODEL 1	(Intercept)	4.21	1.42	2.97	.004**
direct effects	Explicitness	3.39	1.23	2.75	.007**
	Consecutiveness	-2.85	1.20	-2.37	.02*
	Gender	-0.32	1.19	-0.27	.79
	Education	0.85	1.23	0.69	.49

		b	Std. Error	z	p
	Age	2.90	1.23	2.37	.02*
MODEL 1.1	(Intercept)	4.17	0.86	4.84	<.001***
effect of explicitness	Explicitness	4.04	1.22	3.31	.001**
MODEL 1.2	(Intercept)	7.43	0.83	8.93	<.001***
effect of consecutiveness	Consecutiveness	-2.85	1.26	-2.26	.03*
MODEL 1.3	(Intercept)	-0.44	1.72	-0.26	.80
effect of age	Age	0.22	0.05	4.12	<.001***
MODEL 2	(Intercept)	2.98	1.54	1.94	.06
direct effects	Explicitness	5.30	1.58	3.36	.001**
+	Consecutiveness	-0.61	1.67	-0.37	.72
interaction	Gender	-0.13	1.18	-0.11	.91
	Education	0.83	1.22	0.68	.50
	Age	3.20	1.22	2.62	.01*
	Explicitness: Consecutiveness	-4.54	2.38	-1.91	.06

6.5 Conclusions of Study 1

The results of Study 1 support H1a and provide a weak and partial evidence for H1c; in other words, if a notice appears right after its triggering action the participant will be more able to identify the cause of the notice (H1a) and their sense of control will increase (H1c), compared with the conditions in which the notice is delayed. I also found that – contrary to H2b – an explicit mention to the triggering action in the title of the notice did not affect or – in some models - decreased the participant’s ability to identify the topic of the notice. No other effect on comprehension (H1b, H2a), perceived comprehension, clarity or control was found (H1c, H2c). Regarding the response to the notice, the participants seemed to act according to their habits, selecting the 'accept' button (89.39%) like they reported to do in their daily life (82.58%). The consecutiveness of the notice, in line with the transparency paradox, seemed to increase the possibility of accepting the notice (H3a) and to decrease the time taken to make such decision. Privacy concerns were the only factor reverting the preference for the acceptance button. I could not find any effect of explicitness (H3b) or perceived comprehension (H3c) on the participants’ response.

These results mean that effect of contextualization on comprehension was found only for consecutiveness and with respect to the cause comprehension and sense of control. A possible explanation for the lack of effect on the topic comprehension and on perceived comprehension is that the notice in this study was designed after the typical cookie notices that can be found while navigating on the Internet. The participants themselves reported no difference with the notice they were familiar with (81.06%). Such familiarity might explain the lack of effect on the topic comprehension and perceived comprehension, which were good in all conditions. The cause of the notice, instead, is not commonly mentioned in privacy notices, so the manipulation was able to affect this knowledge.

The results also show that contextualization via explicitness had no effect or, sometimes, detrimental effects on comprehension. Before concluding that explicitness is counterproductive or ineffective in increasing the understanding of the notice, however, the possibility that the message in the explicit notice was somehow unclear should be considered. Indeed, the part of the notice referring to the triggering event might have been to indirect in identifying such trigger ("Entering Foodit websites has two effects…"). Also, the text of the notice was longer in the explicit conditions, which could partly explain the longer response time. Thus, I decided to replicate the study addressing these two flaws in the manipulation of explicitness.

Section I

CHAPTER 7.
STUDY 2

Study 2 was a replica of Study 1 with a more straightforward reference to the triggering action in the explicit notice, and a similar text length in the two notice variants (Figure 9). The explicit reference to the triggering action was obtained by changing the notice's title, while the rest of the notice remained identical across conditions.

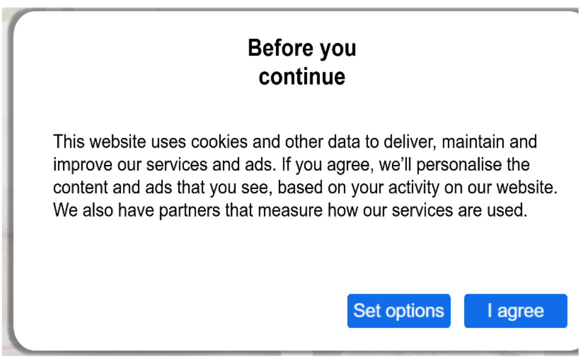

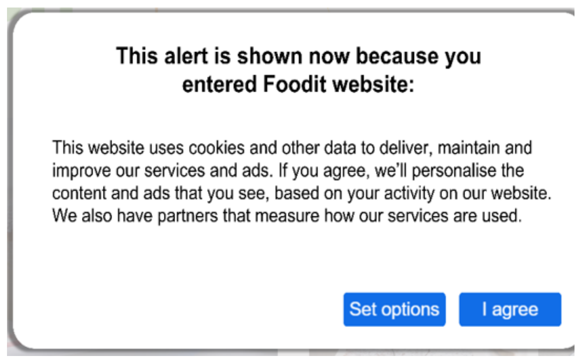

Figure 9: The cookies notice in no-mention (left) and explicit (right) conditions.

I also made some minor rewording in the questionnaire: I added the proper name Foodit to the generic name 'website' to make the referent clearer ("*My* entering *Foodit* website", "Something that I did *when I was on* Foodit website", "None of the above"); and removed any reference to cookies in the answer options to the topic comprehension, which might have been too suggestive of which option was the correct one ("Foodit website declines responsibility for mistakes in finding the appropriate recipe", "Foodit website asks permission to track my activity on the website" (correct) and "Foodit websites credits the authors of its content and pictures"). Except for these adjustments, the materials and the method were the same as in Study 1. The hypotheses of this study were also the same as in Study 1:

Effect of consecutiveness on comprehension: compared with a delayed notice, a notice appearing right after its trigger improves the users' comprehension of its cause (**H1a**), topic (**H1b**), and the users' sense of comprehension and control (**H1c**).

Effect of explicitness on comprehension: mentioning the trigger in the notice improves the users' comprehension of its cause (**H2a**), topic (**H2b**), and the users' sense of comprehension and control (**H2c**).

Transparency paradox: the acceptance of the notice is more frequent in the consecutive condition (**H3a**), in the explicit condition (**H3b**), and with a higher perceived comprehension (**H3c**).

7.1 Participants

The initial sample was composed of 158 people. I excluded from the analysis 30 people: 13 for inappropriate completion of the questionnaire (i.e., response pattern, failure to the attention checks), 16 because of technical issues, and one who did not provide the second consent. The final sample then consisted of 128 participants, aged 18 to 63 years (M = 27.44; SD = 9.01; 64 males, 64 females), all residing in the EU or UK. None of them participated in Study 1, thanks to a filter automatically available in the recruitment platform.

7.2 Data and Analysis

The data treatment and the testing procedure described in Section 6.3 for Study 1 also holds for this second study. The sample size for every predictor in the complete sample and in readers' subsample is reported in Table 10.

Table 10: Sample size by the predictor's levels for both the complete sample and the readers' subsample.

Predictor	Level	Complete sample (N)	Readers' subsample (N)
Consecutiveness	Consecutive	61	47
	Delayed	67	48
Explicitness	Explicit mention	70	53
	No mention	58	42

Predictor	Level	Complete sample (N)	Readers' subsample (N)
Gender	Female	64	53
	Male	64	42
Education	High School	40	29
	University	88	66
Age	Over 30	39	30
	Under 30	89	65

7.3 Results

Table 11 provides an overview of the descriptive statistics of Study 2. The results of the tests are reported in separated sections.

Table 11: Descriptive statistics of the outcomes depending on the predictors.

Predictor	Level	Cause	Topic	Perceived Comp. (item1)		Perceived Comp. (item2)		Clarity		Sense of control		Time to respond (sec)		"I agree" Response
		% correct	% correct	M	SD	M	SD	M	SD	M	SD	M	SD	%
Consecutiveness	Consecutive	83.61	93.44	3.55	0.80	3.06	0.90	3.83	0.52	3.00	0.97	8.07	6.64	95.08
	Delayed	41.79	95.52	3.48	0.68	2.94	0.93	3.81	0.59	3.21	0.99	6.79	3.61	92.54
Explicitness	Explicit	60	95.71	3.45	0.64	2.98	0.93	3.77	0.53	3.03	1.01	8.42	6.38	94.29
	No-mention	63.79	93.10	3.60	0.86	3.02	0.90	3.87	0.59	3.21	0.95	6.16	3.28	93.10
Gender	Female	65.63	96.88	3.47	0.82	2.96	0.96	3.80	0.60	3.28	0.98	7.30	4.77	95.31
	Male	57.81	92.19	3.57	0.63	3.05	0.85	3.84	0.49	2.94	0.96	7.57	6.03	92.19
Education	High School	77.50	95	3.59	0.73	3.04	0.94	3.93	0.59	3.00	0.93	8.39	5.49	95
	University	54.54	94.32	3.49	0.75	2.99	0.90	3.77	0.54	3.16	1.00	7.00	5.26	93.18
Age	Over 30	64.10	97.44	3.53	0.63	2.93	0.91	3.72	0.39	3.15	1.04	7.32	3.72	92.31
	Under 30	60.67	93.26	3.51	0.79	3.03	0.92	3.86	0.61	3.09	0.96	7.47	5.96	94.38

7.3.1 Effect of consecutiveness and explicitness on comprehension (H1a,b, and H2a,b)

The *comprehension of the cause of the notice* is shown in Figure 10 (left). The VIF of the predictors included in Model 1 was less than 5, excluding the presence of multicollinearity (Explicitness = 1.01; Consecutiveness =1.02; Gender = 1.06, Education = 1.03; Age = 1.02) and confirming that this assumption held. The inclusion of the predictors was helpful to explain the variability in data. Overall, Model1 could explain the 18% of this variability. The coefficients of Model 1 are reported in Table 12. Model 1 showed a main effect of consecutiveness in identifying the cause of the appearance of the notice.

Specifically, being in the consecutive conditions increased the odds of the participants correctly answering that the notice appeared because they entered the website. This predictor alone explained the 15% of the variability in data (R^2 =.15). Explicitness instead did not affect this variable (R^2 =.001). In Model 2, the interaction between consecutiveness and explicitness was added. No effect of this interaction was found. So, I could find evidence for H1a, but I could not refuse the null hypothesis for H2a.

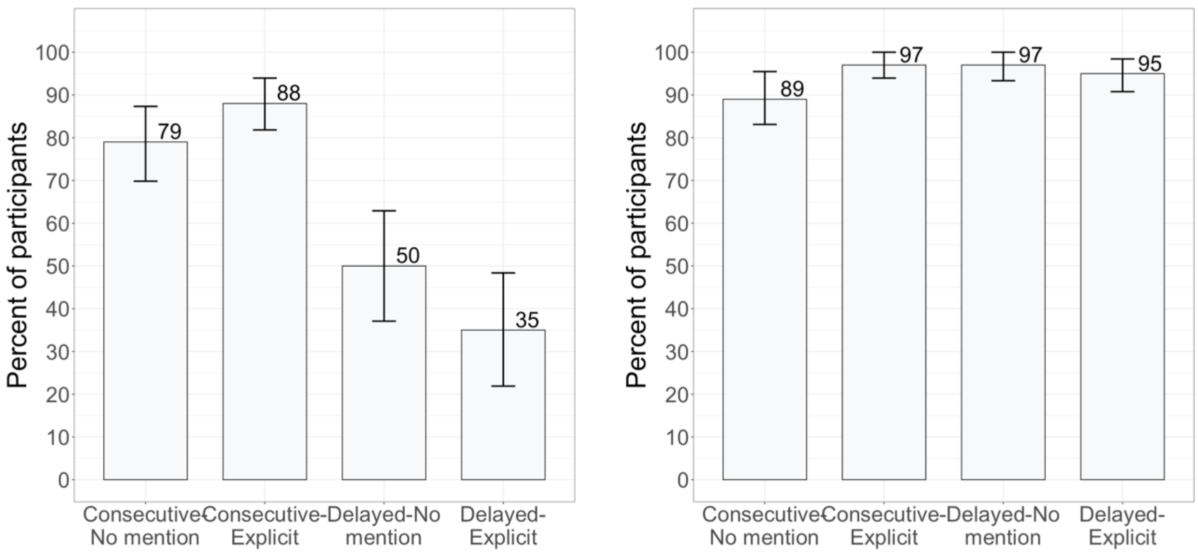

Figure 10: Comprehension of the cause of the notice's appearance (left) and of the topic of the notice (right) depending on the experimental condition (percentage frequency of correct answers). The error bars represent the standard error.

Table 12: Regression results for Cause Comprehension. *p <.05, **p<.01,***p<.001

		b	Std. Error	z	Odds Ratio	p
MODEL 1	(Intercept)	-0.38	0.44	-0.87	0.69	.39
direct effects	Explicitness	-0.12	0.42	-0.30	0.88	.77
	Consecutiveness	1.96	0.44	4.47	7.07	<.001***
	Gender	-0.39	0.43	-0.92	0.68	.36
	Education	0.89	0.48	1.86	2.44	.06
	Age	0.21	0.45	0.47	1.24	.64
MODEL 1.1	(Intercept)	-0.33	0.25	-1.34	0.72	.18
effect of consecutiveness	Consecutiveness	1.96	0.43	4.61	7.10	<.001***
MODEL 2	(Intercept)	-0.14	0.46	-0.29	0.87	.77
direct effects	Explicitness	-0.58	0.52	-1.13	0.56	.26
+	Consecutiveness	1.28	0.60	2.12	3.59	.03*
interaction	Gender	-0.39	0.43	-0.90	0.68	.37
	Education	0.91	0.48	1.88	2.48	.06
	Age	0.20	0.46	0.44	1.22	.65

		b	Std. Error	z	Odds Ratio	p
	Explicitness: Consecutiveness	1.32	0.88	1.50	3.45	.13

Regarding the comprehension of the *topic* of the notice, the VIF of the predictors included in Model 1 was less than 5, excluding the presence of multicollinearity (Explicitness = 1.04; Consecutiveness = 1.06; Gender = 1.06, Education = 1.08; Age = 1.05). Neither Model 1 nor Model 2 differed from Model 0 at the likelihood ratio test (Model1: $X^2(5) = 3.66$, $p = .60$, $R^2 = .07$; Model2: $X^2(6) = 4.91$, $p = .56$, $R^2 = .09$), so there was no significant influence of the predictors in comprehending the topic of the notice. I could not refuse the null hypothesis for H1b and H2b. Indeed, almost the total of the sample answered the item correctly (Figure 10).

7.3.2 Effect of consecutiveness and explicitness on Perceived Comprehension and Control (H1c, H2c)

The scores of perceived comprehension, clarity, and sense of control are shown in Figure 11 and Figure 12 below. Neither Model 1 (i.e., the model in which the independent variables – consecutiveness and explicitness, and the control variables – gender, age and education – were included as predictors) nor Model 2 (the model in which the interaction between explicitness and consecutiveness was added) could explain the variability of the data better than Model 0 at the likelihood ratio test for *perceived comprehension* (Model1 – item 1: $X^2(5) = 3.04$, $p = .69$, $R^2 = .02$;Model 2 - item1: $X^2(6) = 3.10$, $p = .80$, $R^2 = .02$; Model1 – item2: $X^2(5) = 0.76$, $p = .98$, $R^2 = .003$; Model2 - item2: $X^2(6) = 1.01$, $p = .99$, $R^2 = .005$) and *sense of control* (Model1: $X^2(5) = 7.93$, $p = .16$, $R^2 = .03$; Model2: $X^2(6) = 9.27$, $p = .16$, $R^2 = .04$). The overall goodness of fit test showed the same results for *clarity* (Model1: $F(5, 89) = 0.75$, $p = .59$, $R^2 = -.03$; Model2: $F(6, 88) = 0.64$, $p = .70$, $R^2 = -.02$), so no manipulation had any effect. The VIF for the predictors described in this section shows that the absence of multicollinearity assumption held (Table 13); the results of the Brant test indicate that the proportional odds assumption was confirmed in ordinal regressions since no test returned a statistical significance (Table 14).

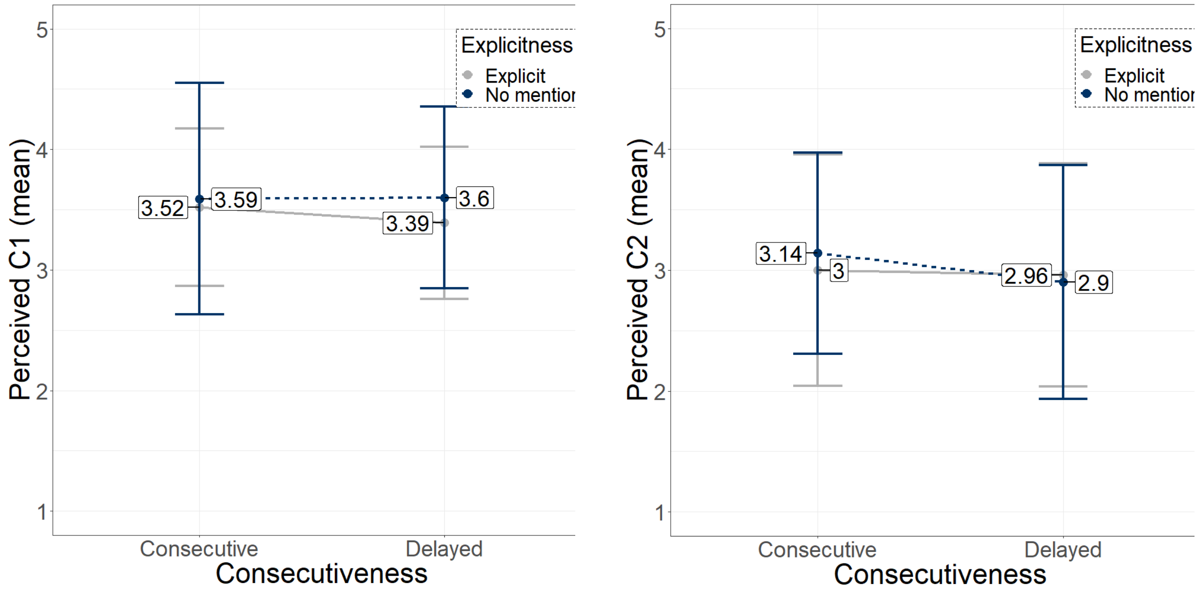

Figure 11: Effect of consecutiveness and explicitness on the scores in the two items measuring the perceived comprehension (1: I would feel confident answering questions about the contents of Foodit alert; 2: After reading Foodit alert, I understand the implications of visiting the Foodit website).

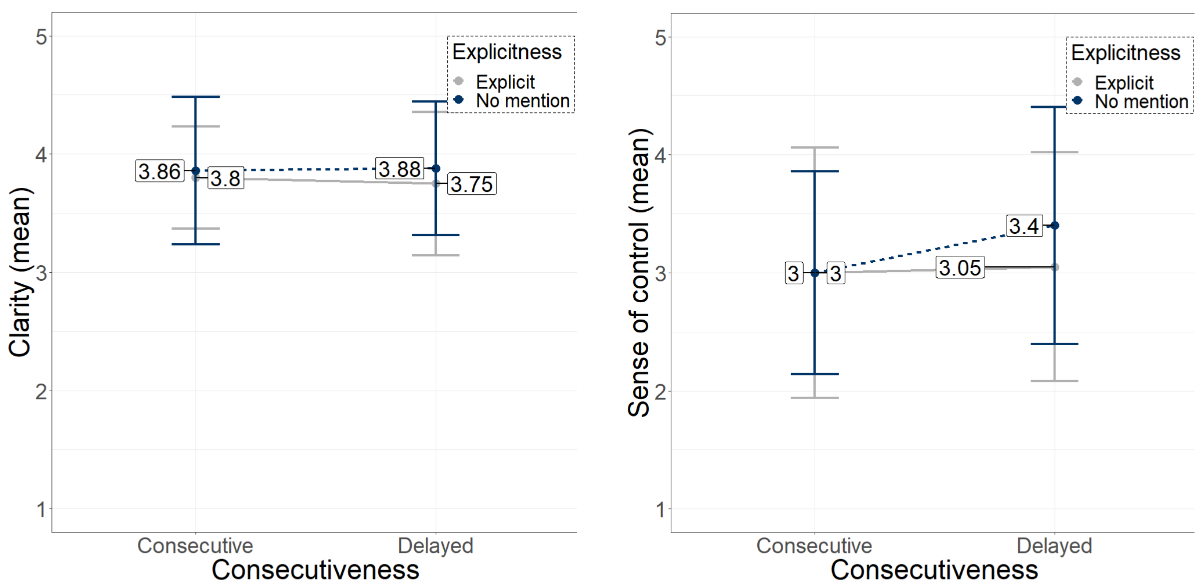

Figure 12: Effect of consecutiveness and explicitness on clarity (average score of three items: “The language of Foodit alert was clear", "Foodit alert was difficult to understand (reverse)" and "I feel that the information in Foodit alert was explained in a clear manner") and sense of control (“I believe I have control over how personal information is used by the Foodit website”).

Table 13: VIF of the predictors depending on the outcome.

	Perceived Comp. 1	Perceived Comp.2	Clarity	Sense of control
Explicitness	1.02	1.02	1.02	1.01
Consecutiveness	1.05	1.05	1.05	1.02
Gender	1.05	1.05	1.05	1.04
Education	1.08	1.08	1.08	1.06
Age	1.05	1.05	1.05	1.01

Table 14: Results of the Brant test depending on the outcome.

	Perceived Comp.1			Perceived Comp.2			Sense of control		
	X^2	df	p	X^2	df	p	X^2	df	p
Omnibus	8.68	18	.97	7.5	12	.82	5.38	18	1
Explicitness	1.03	3	.79	0.04	2	.98	0.13	3	.99
Consecutiveness	1.1	3	.78	0.16	2	.92	0.62	3	.89
Gender	1.95	3	.58	0.96	2	.62	0.15	3	.98
Education	2.16	3	.54	0.7	2	.71	0.17	3	.98
Age	0.17	3	.98	1.53	2	.47	3.55	3	.31
Consecutiveness: Explicitness	0.03	3	1	1.88	2	.39	0.29	3	.96

7.3.3 Response to the Notice (H3a,b,c)

The distribution of the “I agree” response to the request to install cookies (agree/set options) across conditions and the time taken to respond are reported in Figure 13.

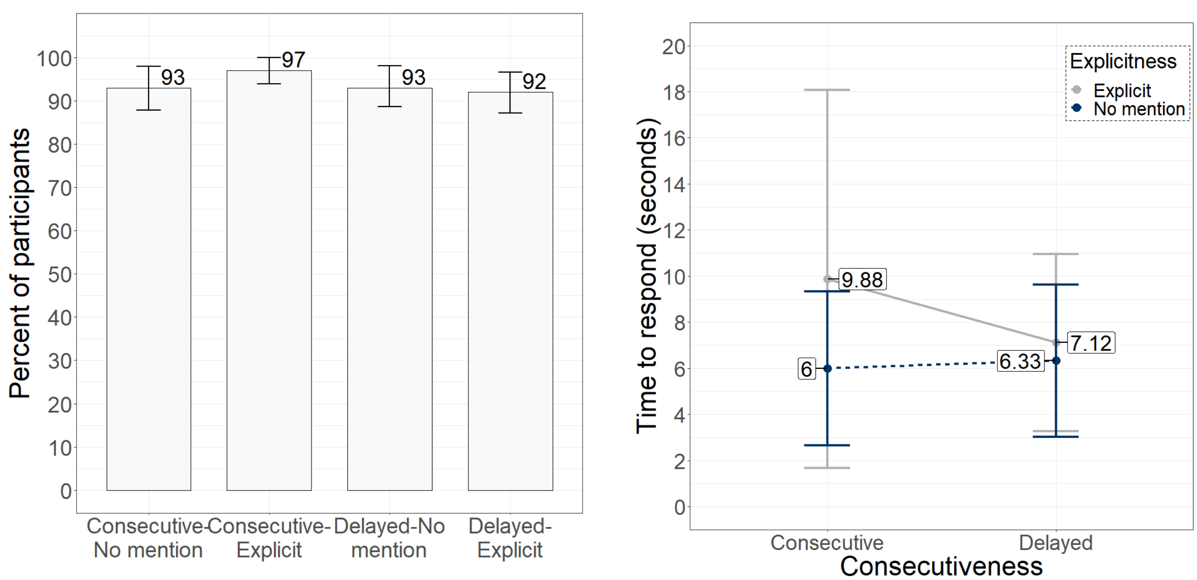

Figure 13: Distribution of the ‘I agree' response to the notice by experimental condition (the error bars represent the standard error - left) and effect of the main variables on the response time (right).

Regarding the effect of the predictors on the participants' *response* to the notice, the VIF of the predictors was less than 5, excluding the presence of multicollinearity (Explicitness = 1.01, Consecutiveness = 1.04, Gender = 1.10, Education = 1.14, Age = 1.03, Privacy concerns = 1.03). Neither Model 1's fit nor Model 2's was better than Model 0 at the likelihood ratio test (Model1: $X^2(6) = 5.55$, $p = .48$, $R^2 = .09$; Model2: $X^2(7) = 6.39$, $p = .50$, $R^2 = .11$), so no predictor had any effect (H3a, H3b). By including privacy concerns only in Model 3, I found a tendency of significance for the effect of privacy concerns ($b = 1.56$, $z = 1.97$, $p = .05$, $R^2 = .08$).

I then checked whether there was any effect of the perceived comprehension scores on the response (H3c). So, I built Model 4 introducing only the perceived comprehension item 1 ("I would feel confident answering questions about the contents of Foodit alert") as a predictor, and the response as the outcome. I found that perceived comprehension did not explain the variability in the data better than chance alone ($X^2(1) = 0.19$, $p = .66$, $R^2 = .004$).

Finally, I considered the *time to respond*. From the visual inspection of residuals, the assumptions of linearity, homoscedasticity, and normality were not respected in Model 1 for time to respond. Furthermore, the data were not normally distributed for the levels of my independent and control variables at the Shapiro-Wilk test (Table 15), violating the normality assumptions of t-test and ANOVA. So, I run separate Mann-Whitney U tests to verify the effect of the independent (consecutiveness and explicitness) and control variables (age, gender, and education) on the time to respond; I employed the Vargha and Delaney's A as a measure of effect size. To assess the effect of the interaction between Consecutiveness and Explicitness I employed a Kruskal-Wallis test using the experimental condition as the independent variable. In this case, I will report the E^2 as the index of the effect size.

Table 15: Time to respond: Results of the Shapiro-Wilk test on the levels of the independent and control variables. *p<.05, **p<.01, ***p<.001

Predictor	Level	W	p
Consecutiveness	Consecutive	0.69	<.001***
	Delayed	0.86	<.001***
Explicitness	Explicit	0.71	<.001***
	No mention	0.82	<.001***

Predictor	Level	W	p
Consecutiveness: Explicitness	Consecutive - no mention	0.80	<.001***
	Consecutive - explicit	0.72	<.001***
	Delayed - no mention	0.82	<.001***
	Delayed - explicit	0.81	<.001***
Gender	Female	0.74	<.001***
	Male	0.53	<.001***
Education	High School	0.66	<.001***
	University	0.57	<.001***
Age	Over 30	0.51	<.001***
	Under 30	0.63	<.001***

No differences were found between: the experimental conditions ($\chi^2(3) = 7.14$, $p = .07$, $E^2 = .08$.; Explicit-Consecutive: *Md* = 6.23; Explicit-Delayed: *Md* = 6.51; Implicit-Consecutive: *Md* = 4.70; Implicit-Delayed: *Md* = 4.79); consecutive (*Md* = 5.54) and delayed (*Md* = 5.93) conditions (*W* = 1113, *p* = .92, *A* = .49); females and males (*W* = 1077, *p* = 0.79, *A* = .48; Female: *Md* = 5.32; Male: *Md* = 5.90); high school and university (*W* = 1105, *p* = .23, *A* = .58; high school: *Md* = 6.62; university: *Md* = 5.51); under and over 30 (*W* = 838, *p* = .27, *A* = .43; under 30: *Md* = 5.31; over 30: *Md* = 6.32). I only found a difference between explicit (*Md* = 6.41) and implicit (*Md* = 4.70) conditions (*W* = 1450, *p* = .01). The size of this effect was small (*A* = .35), given that the variability in data was high (Table 11).

7.4 Conclusions of Study 2

Study 2 supports the main findings of Study 1. In both studies, having the cookie notice appear immediately after the triggering action (i.e., entering the website) improved the participants' ability to identify the notice's *cause* (H1a) compared to the conditions in which the appearance was delayed. Both studies found no effect of explicitness on the cause comprehension (H2a), while the negative effect of explicitness on topic comprehension was found only in the first study. In both studies I could not find evidence for an effect of my independent variables on perceived comprehension and clarity (H1c, H2c). Study 1 partially supported H1c for sense of control, showing a positive yet weak (R2 = .04) influence of consecutiveness, and during Study 2 did not find this effect at all. Finally, Study 1 and 2 found that the participants' privacy concerns affected their response to the notice (i.e., the lower the concerns the higher the acceptance rate); the significance was higher in the first study, but the effect size was similar (.11 and .08) and the effect was in the same direction. They both did not provide evidence for the effect of explicitness (H3b) and perceived comprehension (H3c) on the response. The weak effect of consecutiveness was not replicated in Study 2 (H3a). The results about the time to respond were partially, although weakly, confirmed in Study 2: participants included in the explicit mention conditions took more time than participants in no-mention conditions to respond, even after having reduced the differences in the length of the text and simplified the way in which I explicitly mentioned the cause of appearance. Neither the effect of age nor the one of consecutiveness found in Study 1 was replicated here, perhaps due to the younger age of the sample - the median was 30 in Study 1 (M = 30.99; SD = 11.02) and 24 in Study 1 (M = 27.44; SD = 9.01), - and for the weakness of the effect of consecutiveness during Study1 (R^2 = .04).

In conclusion, consecutiveness affected the comprehension of the cause of the notice but not the comprehension of the topic or any other variable measuring the experience of comprehension. I then stand with the explanation proposed after Study 1: the participants were already familiar with the general gist of a cookie notice and could successfully answer the item requiring such general knowledge. Indeed, most participants reported no difference with the notice they were familiar with (81.06% Study 1, 88.28% Study 2), and selected the 'accept' button like they declared to do in real life (82.58% Study 1, 84.38% Study 2). The trigger of the notice, instead, is not commonly described in cookie notices, and was more easily affected by my manipulations. I also can be more confident in suggesting that adding some information in the notice might not be an effective contextualization strategy, since it does not obtain any improvement in the participants' comprehension of that specific piece of information.

7.4.1 Open questions for further investigation

Study 1 and Study 2 left two open tasks. One regards the effect of contextualization on the comprehension of uncommon notices. Therefore, I decided to run a third study to measure comprehension on more layers, a generic one related to common knowledge easily obtained by navigating websites vs a specific layer, related to information conveyed by a given notice in particular. Another task that remains open consists of testing the effect of different contexts: in the delayed conditions of Studies 1 and 2 I broke the consecutiveness between entering the website and seeing the cookie notice, but I did not know what the context of the notice was. The participant could have just waited, looked around, started doing something on the website, and this might have affected their interpretation of the notice. Therefore, in the third study, I decided to control the context of the notice' appearance and test its effect on the participants' interpretation of the notice.

Section I

CHAPTER 8.
STUDY 3

In Study 3 I used a different version of the Foodit website, which implemented a push vs pull paradigm of consent acquisition (Klumpe et al., 2020). In both conditions, the notice had the same text, asking the participant's permission to track their activities on the website. In one condition, the notice appeared as the user clicked on the website link and the Foodit website started loading; the context of the notice (i.e., the action immediately preceding it) was then the act of entering the website. In the second condition, the notice appeared after the user clicked on one of the download buttons on the Foodit website; the context of the notice was then a specific operation (downloading) performed on the Foodit website.

The two conditions reflect two models that are found in real life: asking permission to collect the users' data as they enter a website (push-based strategy) or as they require a specific service/content for which permission is relevant (pull-based strategy; Klumpe et al., 2020). In the push model, the context is a very preliminary stage of the website usage (generic condition): the notice asks permission when the user has not identified any specific content of interest on the website yet, and refers to future and un-formed courses of action, whose nature and value is not clear. In the pull model (also called 'just-in-time'; Schaub et al., 2015), instead, the course of action to which the notice refers has already started; the user has already identified some content of interest and tried to obtain it. The context of the notice in the pull model is then a very specific action. This makes the scenario of the request more vivid and present to the users. In online marketing and advertising, a push model has already proven to raise users' privacy concerns (Klumpe et al., 2020; Kobsa & Teltzrow, 2004) and be vulnerable to financial compensations (Xu et al., 2009). In terms of comprehension, Balebako et al. (2015) found that the notice's content was better recognized if shown during the app use compared with showing it early in the app store. Here I study the effect on the comprehension and interpretation of the notice using open-ended and multiple-choice questions.

In continuity with study 1 and 2, I also tested the effectiveness of providing contextual information explicitly in the notice, by mentioning or not a specific in action in the title of the notice (“before you continue” vs “before you download”). The study followed a 2x2 between-participant design, with consecutiveness and explicitness as the two variables. The hypotheses for this study were the following:

Effect of consecutiveness on comprehension: displaying the notice after a specific action increases the identification of that action as the notice’s cause **(H1a),** and the generic comprehension of the topic of the notice **(H1b)**. The context also affects the interpretation of the specific topic of the notice, in a direction that I cannot predicted and intend to explore **(RQ1).**

Effect of explicitness on comprehension: explicitly mentioning a specific action in the notice increases the identification of that action as the notice’s cause (**H2a**) and improves the generic comprehension of the topic of the notice **(H2b)**. It also affects the interpretation of the specific topic of the notice, in a direction that I cannot predicted and intend to explore **(RQ2).**

Transparency paradox: the acceptance of the notice is more frequent when it is shown after a specific action **(H3a)**, and in the explicit conditions **(H3b)**.

8.1 Material

I slightly modified the website used in the two previous studies to obtain more control over the actions immediately preceding the notice. The webpage designed for this study included the Foodit logo and the three recipes represented by an image, a label, and a download button (Figure 14, top left). The notice appeared immediately after landing on the webpage in the generic condition, when the content was not yet visible, or after clicking on the only interactive elements on the webpage, the download button, in the specific condition. The notice prevented any other action and contained two buttons, "Deny" or "I agree" (Figure 14, top right). By pressing on either button, the participant was led to the questionnaire. The text of the notice asked the permission to track the users' activity on the website. In the explicit condition, there was also a reference to a specific activity to be tracked (downloading), while in the generic condition there was no reference to it (Figure 14, down-left and down-right). As in the two previous studies, the webpage was built in PsychoPy3 (https://www.psychopy.org/) and presented through Pavlovia (https://pavlovia.org/).

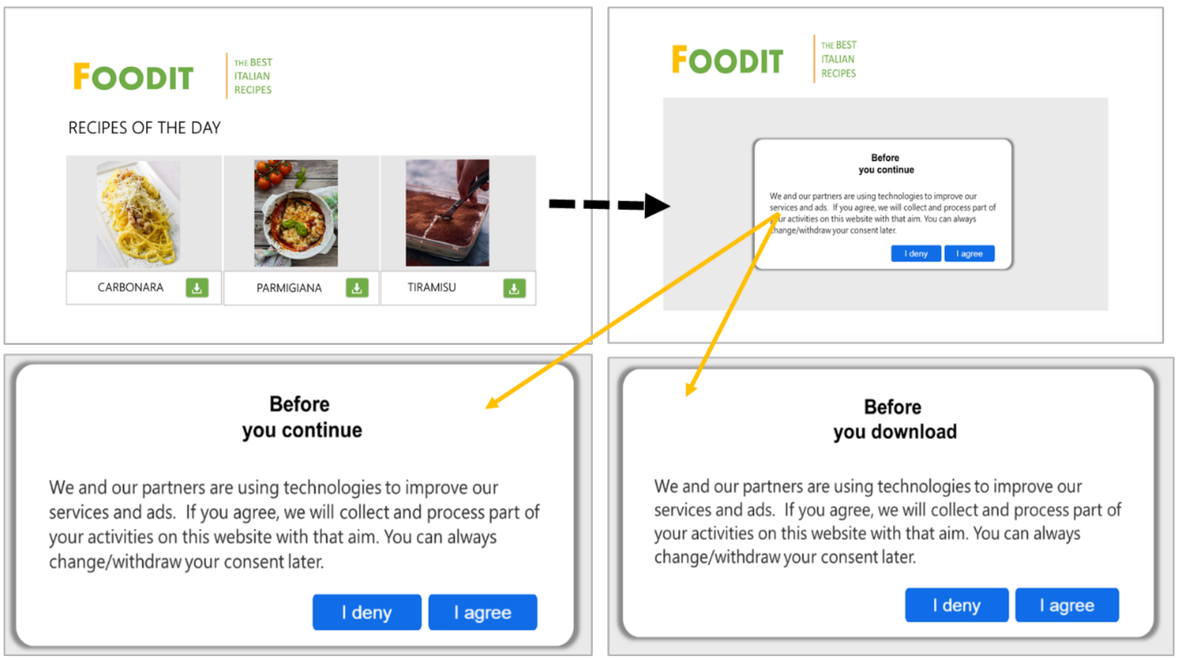

Figure 14: The website at the participants' entrance (top left) and when the Cookies notice appeared (top right); the notice in the no-mention conditions (bottom left) and in the explicit mention conditions (bottom right).

8.2 Measures

In this study, I measured:

- The comprehension of the *cause* of the notice with a multiple-choice question (*What caused the alert appearance?)* with three answer options*:* "My entering the website", "Something that I did when I was on the website" or "I don't know", like the one I used in the first two studies.

- The comprehension of the generic topic of the notice with an open-ended question ("*Consider the alert displayed by Foodit website. In the box below, try to describe as accurately as possible in your own words what that alert said")* allowing participants to use their own words (similarly to Korunovska et al., 2020). The answer was analyzed to see how many participants provided a depiction of the topic compatible with the actual topic of the notice, i.e., the permission to track the activity on the website.
- The interpretation of the specific topic of the notice with a multiple-choice question (*Which of the following information will the Foodit website collect if the visitor agrees with the request in the alert?)* with four options: "The visitor's personal information, such as the IP address", "The website that linked to Foodit (inbound links)", "The recipes the visitor downloads", "I don't know"). This item tested the users' comprehension of the notice at a more specific level than the open question. The options were inspired by the ones used in Utz et al. (2019) in their study. In this case, the interest was not in the correctness of the option proposed to the participants, but in the *interpretation* they would give to its specific topic.

In the questionnaire, the open question on the generic topic of the notice was asked first, in order not to influence the participants' spontaneous recollection. Like in Study 1 and Study 2, I collected the participants' *response to the notice* ("I agree"/ "Deny") and the *time taken to respond to the notice*, as described in Section 5.3.3. Finally, I measured the participants' privacy concerns, some demographic variables (gender, age, and education), and the participants' familiarity with similar notices as *control variables*. The description of these variables is provided in Section 5.3.4.

8.3 Procedure

I followed the same procedure described in Study 1 and Study 2. The only change was in the instructions, which in Study 3 read as follows: *use as many functions as possible in 15 seconds*. These instructions did not hint to any specific action available on the website because this would have interfered with the genericity/specificity manipulation at the core of this third study; the instructions also introduced some time pressure, so the participant could focus on the task of using the website and avoid dwelling on the notice.

8.4 Participants

The initial sample was composed of 158 people. I excluded from the analysis 67 people: 63 for inappropriate completion of the questionnaire (i.e., a response pattern, failure at the attention checks, use of incompatible devices) and 4 because of technical issues. The final sample consisted of 91 participants, aged 19 to 73 years (*M* = 36.71; *SD* = 12.78, 40 males, 51 females), all residing in the EU or UK. None of them participated in Study 1 or Study 2, thanks to a filter automatically available in the recruitment platform.

8.5 Results

8.5.1 Comprehension of the cause of the notice

The frequency at which each option was selected in each condition is reported in Figure 15. The data seems to show that the participants identified the cause of the notice in line with the hypothesized effect: in the conditions with a generic context, the participants attributed the appearance of the notice to their entering the website; in the conditions with a specific context, the participants attributed the appearance of the alert to something they did. Instead, like in Studies 1 and 2, the presence of a different title on the notice did not seem to direct the participants' choice. To analyse these data, I followed the same procedure as Study 1 and Study 2, described in Section 6.3. I first built a logistic model (Model1) where the independent variables (consecutiveness and explicitness) and the control variables (age, gender, and education) were specified as predictors. I excluded from this analysis the participants who answered "I don't know" (N = 2). The inclusion of the predictors was useful to explain the variability of data (R^2 = .62). The results are reported in Table 16. Model 1 showed a main effect of consecutiveness (R^2 = .61) in line with my hypotheses: the action preceding the notice influenced the participants' interpretation of its cause (H1a). In particular, being in a generic context increased the odds of participants answering that the notice appeared because they entered the website; being in a specific condition increased the odds of them answering that the notice appeared because of something they did on the website. No effect was found for explicitness (H2a) and the interaction between consecutiveness and explicitness (Model 2). Thus, I could find evidence for H1a, but I could not reject the null hypothesis for H2.

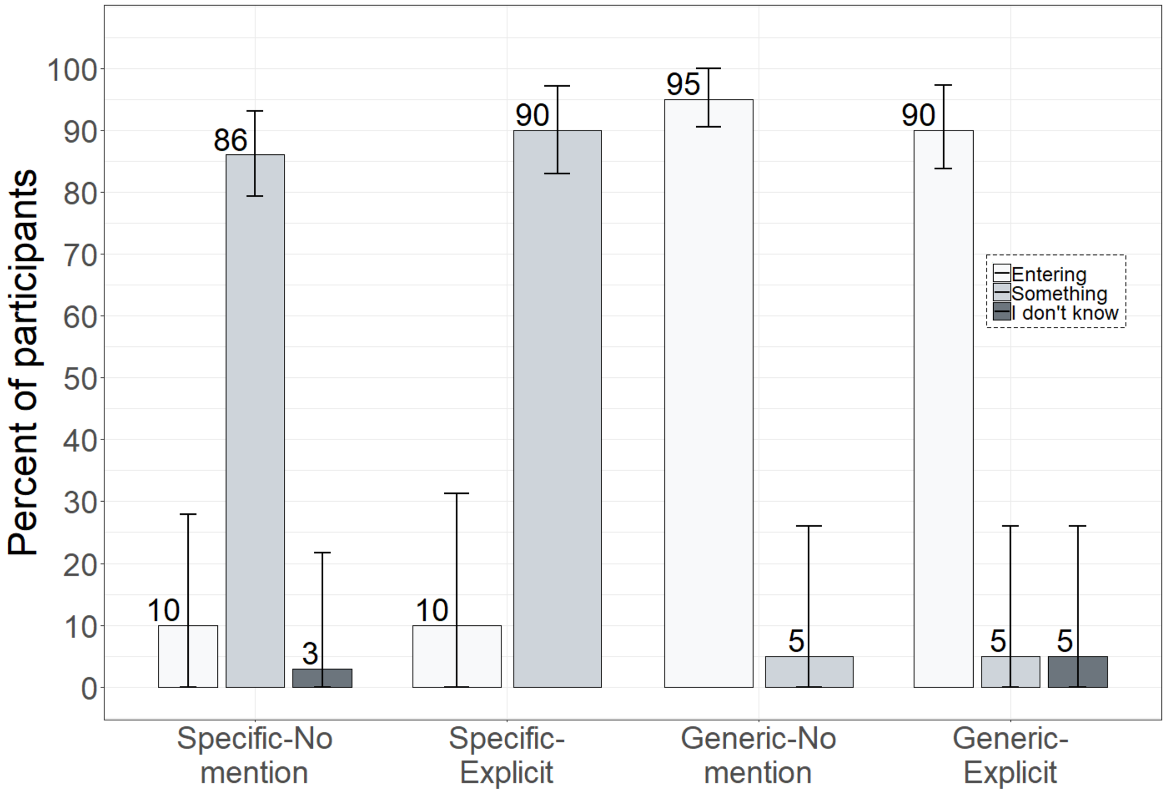

Figure 15: Comprehension of the cause of the notice depending on the experimental condition. The error bars represent the standard error.

Table 16: Regression Results for Cause Comprehension. *p <.05, **p<.01,***p<.001.

		b	Std. Error	z	Odds Ratio	p
MODEL 1 direct effects	(Intercept)	-3.32	1.21	-2.74	0.04	.006**
	Explicitness	0.03	0.81	0.03	1.03	.98
	Consecutiveness	5.43	0.99	5.51	227.81	<.001***
	Gender	0.55	0.83	0.66	1.73	.51
	Education	0.72	0.90	0.80	2.06	.42
	Age	-0.62	0.93	-0.66	0.54	.51
MODEL 1.1 effect of consecutiveness	(Intercept)	-2.97	0.73	-4.10	0.05	<.001***
	Consecutiveness	5.12	0.87	5.92	167.70	<.001***
MODEL 2 direct effects + interactions	(Intercept)	-3.27	1.43	-2.29	0.04	.02*
	Explicitness	-0.05	1.49	-0.04	0.95	.97
	Consecutiveness	5.38	1.29	4.18	216.34	<.001***
	Gender	0.55	0.84	0.66	1.73	.51
	Education	0.71	0.90	0.79	2.04	.43
	Age	-0.63	0.95	-0.66	0.54	.51

	b	Std. Error	z	Odds Ratio	p
Explicitness: Consecutiveness	0.11	1.80	0.06	1.12	.95

8.5.2 Comprehension of the general topic of the notice

In analyzing the answer to the open question on the topic of the notice, I identified the answers that were compatible with the topic of the notice (e.g., reference to cookies, data collection, sharing data with third parties) and those that were not. The number of compatible answers is reported in Figure 16. What did the incompatible answers refer to? Eleven participants admitted having no idea. A couple of answers mentioned topics related to food but not to the notice ("The alert would be about allergies" and "Identifies what industry and businesses can voluntarily use to decrease the risk of intentional contamination of their food products"). Three participants stated that the notice was about the terms and conditions "Accept the terms and condition of the use of Foodit website", or recalled bits of the notice unrelated to collecting data, for instance, "It said something about if I gave permission or not. I was then able to select if I agreed or I denied at the bottom". Some others in the explicit condition tried to connect creatively some words they might have noticed, such as "downloading" and "agree"; for instance, "that they were going to download information that I placed into the website" or "Something about agreeing or disagreeing with a download". Finally, in the specific explicit condition, six more participants mentioned 'downloading' and nothing more, as if they only read the title of the notice. The subset of participants who were unable to describe the topic of the notice seemed to have retained some elements of the notice and to fill the rest based on their previous knowledge. For instance, twenty participants reported that the notice concerned cookies even though the term "cookies" was not used in the notice.

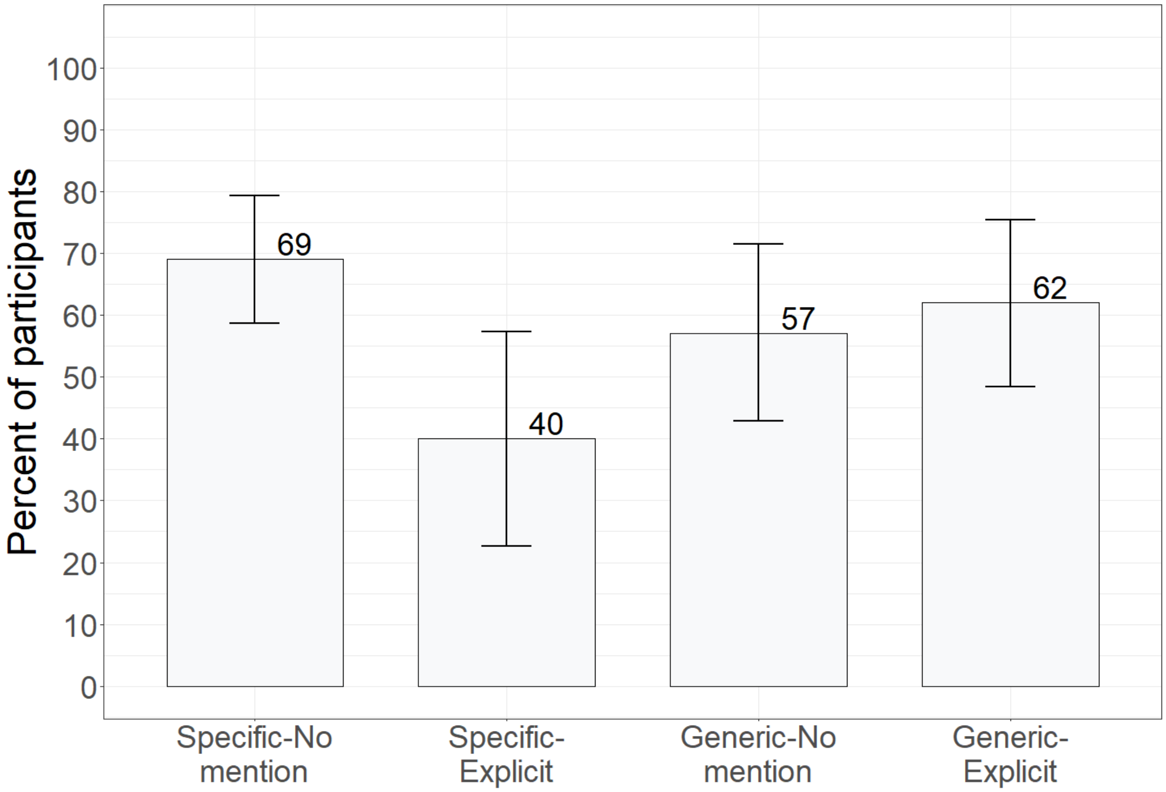

Figure 16: Percentage of compatible answers to the item measuring the general comprehension of the topic of the notice, broken down by the experimental condition. The error bars represent the standard error.

To verify whether the compatibility of the answers differed depending on my independent (consecutiveness and explicitness) variables also considering the effect of my control variables (gender, age, and education), I followed the same procedure described in Section 6.3 (H1b, H2b). So, I built a logistic model (Model 1) specifying the independent and control variables as predictors. Then, I built a second model (Model 2) adding the interaction between explicitness and consecutiveness. Neither model was better than Model 0 in explaining the variability of the data (Model1: $X^2(5) = 6.66$, $p = .25$, $R^2 = .05$; Model2: $X^2(6) = 8.21$, $p = .22$, $R^2 = .07$). So, no variable influenced this outcome. The table reporting the results of the logistic models is included in the Appendix.

8.5.3 Interpretation of the specific topic of the notice

A multiple-choice question asked the participant to choose which specific topic was covered by the notice. The frequency with which each answer option was selected in the four conditions is reported in Figure 17 and in Table 17 (observed proportions).

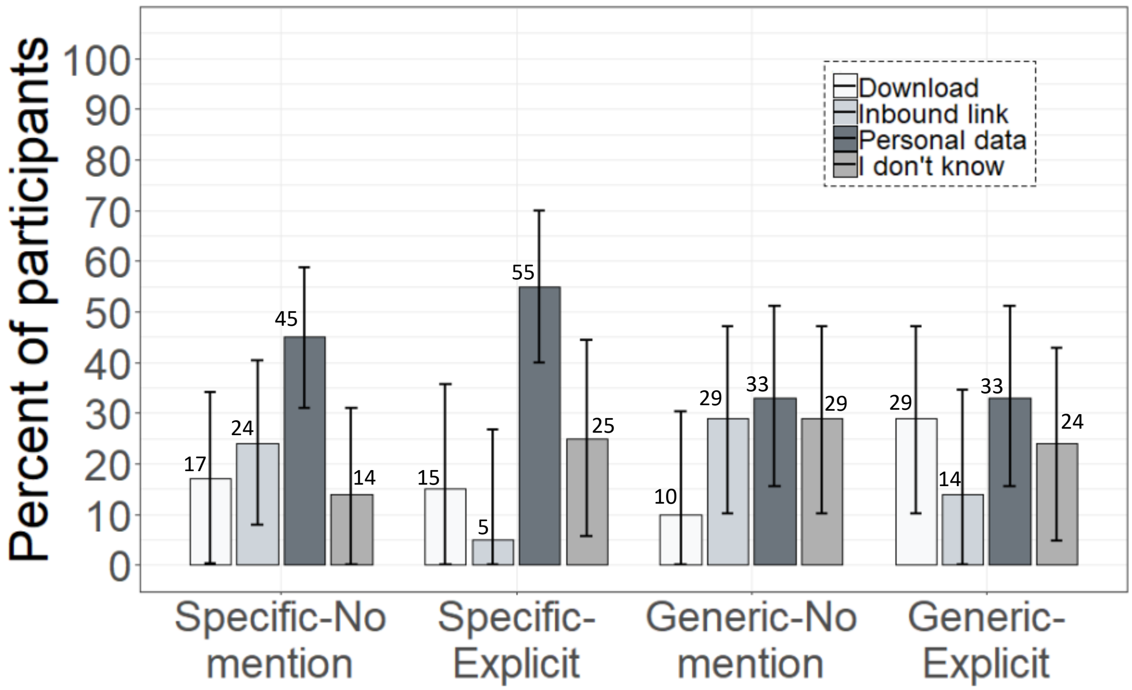

Figure 17: Answers for the specific topic item by the experimental condition (percentage frequency). The error bars represent the standard error.

I wanted to test whether the probability of selecting one option differed from chance inside each condition (RQ1 and RQ2). By running logistic regressions, I would be limited to dichotomous outcomes (e.g., right/wrong), while my interest was in the participants' *interpretations* of the notice and in how they differed across condition. Thus, I ran a series of exact binomial tests both for my experimental conditions and for the levels of my control variables (Table 17 and Table 18). My null hypothesis (H0) was that each response was equiprobable (expected proportion = .25). The results show that the only conditions in which one option had a higher probability of being selected than mere chance were those with a specific context (pull strategy) (Table 17). In such conditions, the answer selected was 'personal data'. In other words, a notice appearing right after the user had attempted a specific action on the website increased the possibility that they associated the notice with the collection of personal data. The control variables did not affect the chances of selecting a different answer (Table 18): the answer referring to the collection of personal data had higher odds of being selected across gender, age, or level of education.

Table 17. Exact binomial tests results for Topic Interpretation. p <.05, **p<.01,***p<.001.

	Specific - no mention (N = 29)		Specific – explicit (N = 20)		Generic - no mention (N = 21)		Generic – explicit (N = 21)	
	Observed proportion	p	Observed proportion	p	Observed proportion	p	Observed proportion	p
Download	.17	.40	.15	.44	.10	.30	.29	.80
Inbound link	.24	1	.05	.03*	.29	.80	.14	.32
Personal data	.45	.02*	.55	.004**	.33	.45	.33	.45
I don't know	.14	.20	.25	1	.29	.80	.24	1

Table 18. Exact binomial tests results for the item about the Specific Topic Interpretation (control variables). p <.05, **p<.01, ***p<.001.

	Gender: Female (N = 40)		Gender: Male (N = 51)		Age: Under 30 (N = 29)		Age: Over 30 (N = 62)		Education: High School (N = 34)		Education: University (N = 57)	
	Obs. prop.	p	Obs. prop.	p	Obs. prop.	p	Obs. prop.	p	Obs. prop	p	Obs. prop	p
Download	.13	.07	.22	.63	.14	.20	.19	.38	.18	.43	.18	.22
Inbound link	.23	.86	.16	.15	.17	.40	.19	.38	.06	.01**	.26	.88
Personal data	.45	.01**	.40	.02*	.41	.05	.42	.003**	.47	.01**	.38	.02*
I don't know	.20	.59	.24	1	.28	.83	.19	.38	.29	.55	.18	.22

For the sake of accuracy, I run again these exact binomial tests on the subset of participants who understood the generic topic of the notice. The answers selected by this subsample are reported in Figure 18 and the results of the binomial test are reported in Table 19. The results of this subsample confirm the result of the whole sample, according to which notices appearing as the user attempted some action (specific contextualization) resulted in answers different from mere chance and referring to personal data. The results of this subsample confirm the result of the whole sample, according to which notices appearing as the user attempted some action (specific contextualization) resulted in answers different from mere chance and referring to personal data. Likewise, the results regarding the control variables in this subsample (Table 20) were also in line with the results of the complete sample. Even though the option about the personal data did not reach the significant level in the under 30's subsample and was tendent to significance in the male's subsample, it was still the most selected one, independently from gender, age, or the level of education.

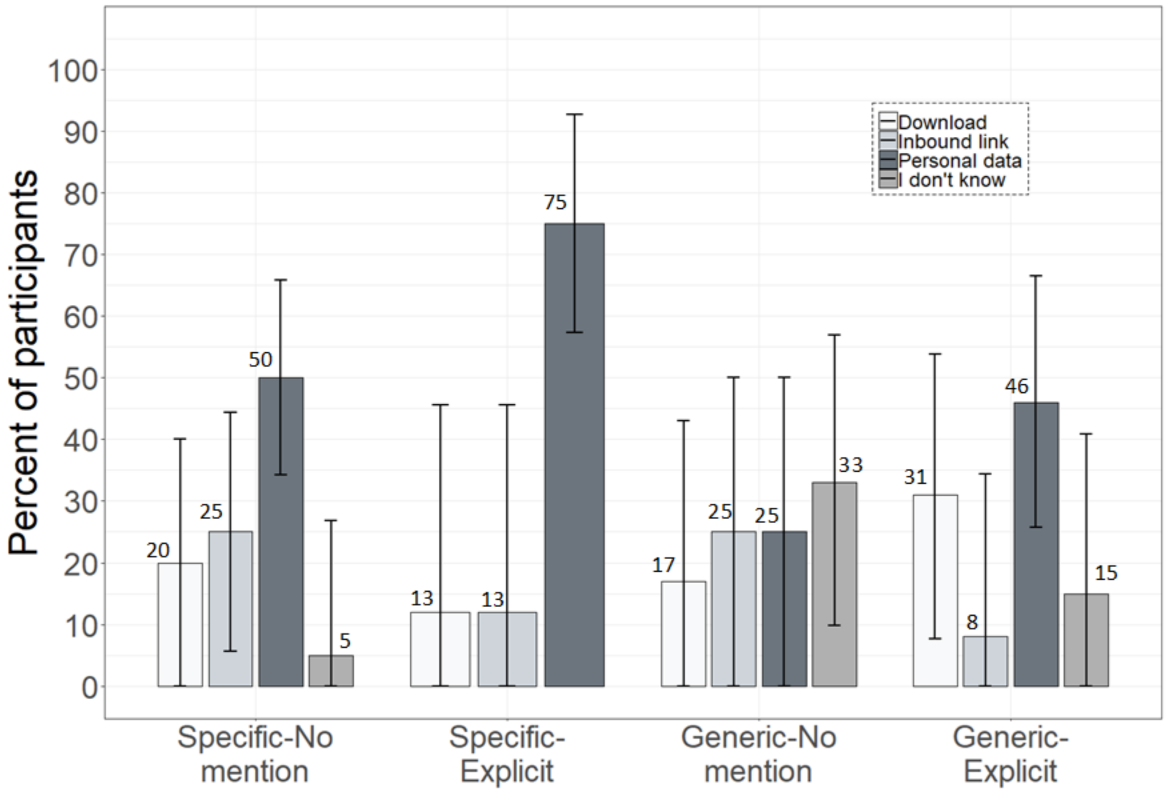

Figure 18: Topic of the notice by the experimental condition (compatible answers subsample). The error bars represent the standard error.

Table 19. Exact binomial tests results for item about the specific topic interpretation (compatible answers' subsample). *p <.05, **p<.01,***p<.001.

	Specific - no mention (N = 20)		Specific – explicit (N = 8)		Generic - no mention (N = 12)		Generic – explicit (N = 13)	
	Observed proportion	p	Observed proportion	p	Observed proportion	p	Observed proportion	p
Download	.20	.80	.13	.69	.17	.74	.31	.75
Inbound link	.25	1	.13	.69	.25	1	.08	.21
Personal data	.50	.02*	.75	.004*	.25	1	.46	.10
I don't know	.05	.04*	0	.21	.33	.51	.15	.54

Table 20. Exact binomial tests results for the item checking the specific topic interpretation (control variables - compatible answers' subsample). *p <.05, **p<.01, ***p<.001.

	Gender: Female (N = 21)		Gender: Male (N = 32)		Age: Under 30 (N = 20)		Age: Over 30 (N = 33)		Education: High School (N = 18)		Education: University (N = 35)	
	Obs. prop.	p	Obs. prop.	p	Obs. prop.	p	Obs. prop.	p	Obs. prop	p	Obs. prop	p
Download	.24	1	.19	.54	.15	.44	.24	1	.28	.79	.17	.33
Inbound link	.24	1	.16	.31	.20	.80	.18	.43	.06	.06	.26	1
Personal data	.48	.02*	.41	.05	.45	.07	.49	.004**	.50	.03*	.46	.01**
I don't know	.28	.83	.28	.83	.20	.80	.09	.04*	.17	.59	.11	.08

8.5.4 Response to the notice

The responses to the notice across conditions and the time taken to the participants to respond are reported in Figure 19. Regarding the effect of the predictors on the participants' *response* to the notice, the VIF of the predictors was less than 5, excluding the presence of multicollinearity (Explicitness = 1.13, Consecutiveness = 1.09, Gender = 1.18, Education = 1.05, Age = 1.21, Privacy concerns = 1.08). Neither Model1 (the model in which the independent variables – explicitness and timeliness – and the control variables – gender, age, education and privacy concerns – were specified as predictors) nor Model 2 (Model 1 with the addition of the interaction between timeliness and explicitness) were better than Model 0 in explaining the variability of data (Model 1: $X^2(6) = 9.99$, $p = .12$, $R^2 = .14$; Model2: $X^2(7) = 12.48$, $p = .09$, $R^2 = .18$). The coefficients of the model showed a significant effect of privacy concerns (Table 21), so I built a model where only privacy concerns were inserted as predictor (Model1.1). Model 1.1 was better than Model0 at the likelihood ratio test ($X^2(1) = 8.30$, $p = .004$, $R^2 = .12$). This result suggested that when the level of privacy concerns increased, the likelihood to accept the request of the notice decreased. Since no effect of consecutiveness or explicitness was found, I could not reject the null hypothesis for H3a and H3b.

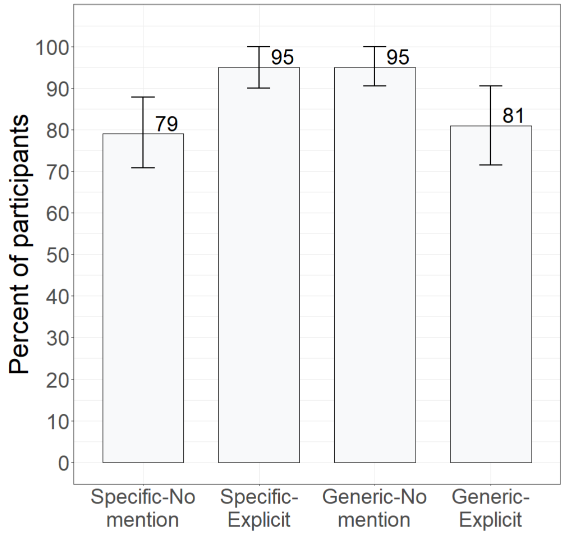

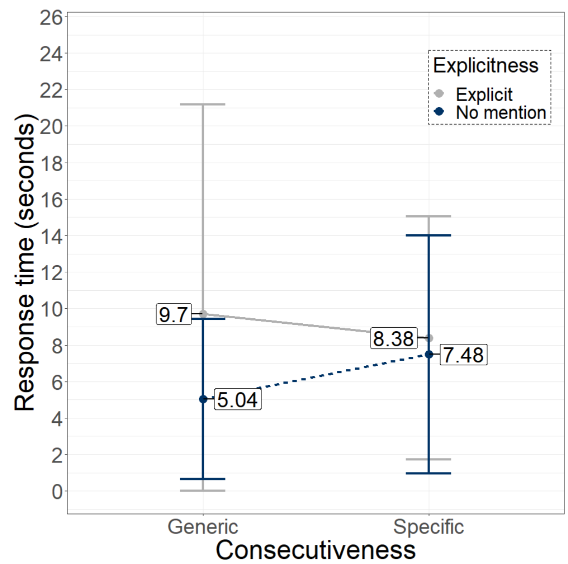

Figure 19: Distribution of the ‘accept' response to the notice by experimental condition (the error bars represent the standard error - left) and effect of the main variables on the response time (right).

Table 21: Regression Results for Response. *p <.05, **p<.01,***p<.001

		b	Std. Error	z	Odds Ratio	p
MODEL 1 direct effects	(Intercept)	7.93	2.85	2.78	2767.31	.005**
	Explicitness	0.53	0.71	0.75	1.70	.45
	Consecutiveness	0.05	0.70	0.07	1.05	.95
	Gender	-0.31	0.72	-0.43	0.73	.67
	Education	-0.20	0.69	-0.28	0.82	.78
	Age	0.77	0.74	1.04	2.16	.30
	Privacy concerns	-1.15	0.46	-2.50	0.32	.01*
MODEL 1.1 effect of privacy concerns	(Intercept)	8.01	2.68	2.99	3000.34	.003**
	Privacy concerns	-1.07	0.44	-2.42	.34	.02*
MODEL 2 direct effects + interaction	(Intercept)	8.12	2.85	2.85	3362.46	.004**
	Explicitness	-0.86	1.25	-0.69	0.42	.49
	Consecutiveness	-1.13	1.17	-0.93	0.32	.34
	Gender	-0.22	0.73	-0.31	0.80	.76
	Education	-0.18	0.70	-0.26	0.83	.79
	Age	0.58	0.76	0.76	1.78	.45
	Privacy concerns	-1.02	0.45	-2.24	0.36	.03*
	Explicitness: Consecutiveness	2.48	1.69	1.47	11.89	.14

Finally, I considered the *time to respond*. As in Study 1 and Study 2, I run this analysis on the readers' subsample. The distribution of the time to respond depending on my main variables is shown in Figure 19. The VIF of the predictors was less than 5, excluding the presence of multicollinearity (Explicitness = 1.07, Consecutiveness = 1.05, Gender = 1.01, Education = 1.01, Age = 1.02). Neither Model 1 nor Model 2 were better than Model 0 at the overall goodness of fit test (Model1: $F(5,48) = 1.18$, $p = .34$, $R^2 = .02$; Model 2: $F(6,47) = 1.17$, $p = .34$, R2 = .02. No significant contribution of my predictors was found.

8.6 Conclusions of Study 3

The effect of consecutiveness on identifying the cause of the notice confirms to be strong (H1a); in addition, I found that it increased the odds that the notice was interpreted as a request for collecting personal data (RQ1). The comprehension of the generic topic of the notice (H1b) instead was not affected, like in Studies 1 and 2. Regarding the effect of explicitness, it did not seem to affect comprehension, like in Study 1 and 2 (H2a, RQ2).

The association of the notice with a request for collecting personal data deserves some comments. It would have been plausible that seeing the notice while performing a specific action directed the participants to select that action as a better description of the specific topic of the notice (i.e., downloading). Instead, the participants selected "personal data", and did so more often than the participants who saw the notice at the entrance of the website. Why? That question followed the open question asking to describe the topic of the notice. Most participants answered the open question by using their general knowledge of privacy notices. Indeed, they judged the notice similar to the notices they saw in the past (84.62%, 77/91) and included in their answer several elements compatible with privacy notices but not present in the text of my notice (e.g., the term "cookies"). When they were asked to choose one option to describe the topic of the notice, the participants in the generic condition (push strategy) had no specific cue since they just entered the website. Their answers were evenly distributed across the four different options. The participants in the specific condition (pull strategy), instead, saw the notice amid their action and this might have led them to think that the notice related to them personally (collection of personal data).

Regarding the participants' responses, they were not affected by my manipulations, and hence by any different comprehension of the notice it might have caused. It was only affected by the participants' privacy concerns. The participants behaved as they declared to usually do (19.78% reported ignoring the notice if it is possible, 58.24% was used to agree, and 21.98% usually click on "Deny"). The transparency paradox seemed not to influence my participants' disclosure behavior (H3a, H3b).

Section I

CHAPTER 9.
CONCLUSIONS OF SECTION I

Section I tested contextualization as a strategy to implement the principle of transparency in privacy notices, with the end of contributing to answering the research question *how can transparency notices be designed to be genuinely transparent – i.e., understandable – for their users?*. Following ethnomethodology and conversation analysis (Schegloff, 2017), the context was defined as the *action performed by the user on the website immediately preceding the appearance of the privacy notice*. During the studies, the *consecutiveness* – i.e., the sequential connection between the notice appearance and the participant's action triggering it – and the *explicitness* – i.e., the mention of the trigger in the text of the notice - of a cookie notification were manipulated and three core hypotheses were tested:

H1 – Effect of consecutiveness on comprehension: the event preceding the appearance of the notice affects the comprehension of the notice's cause **(H1a)** and topic **(H1b),** and the experience of comprehension of the notice (perceived comprehension, clarity, sense of control; **H1c)**.

H2 - Effect of explicitness on comprehension: explicitly mentioning the trigger of the notice in the text affects the comprehension of the notice's cause **(H2a)** and topic **(H2b),** and the experience of comprehension of the notice (perceived comprehension, clarity, sense of control; **H2c)**.

H3 - Transparency paradox: the factors increasing transparency also increase the acceptance of the request to track the users' activity. In particular, it was expected to observe an effect of consecutiveness **(H3a)**, explicitness **(H3b),** and perceived comprehension **(H3c).**

The third study also explored two research questions:

RQ1 – Effect of consecutiveness on the interpretation of the notice: displaying the notice after a specific action influences the interpretation of the specific topic of the notice (i.e., what data the notice is asking to collect).

RQ2 – Effect of consecutiveness on the interpretation of the notice: explicitly mentioning a specific action in the notice influences the interpretation of the specific topic of the notice (i.e., what data the notice is asking to collect).

The specific hypotheses of each study, as reported in the relevant chapters of the dissertation (Chapter 6 for Study 1, Chapter 7 for Study 2, and Chapter 8 for Study 3), were a declination based on the specific manipulations attempted there. In particular, in the first two studies, *consecutiveness* was manipulated by either preserving the sequential connection between the notice appearance and the users' action triggering it (consecutive condition) or broking it with a delay (delayed condition). During the third study, the notice appearance could be consecutive to a generic or a specific action performed by the participant (entering the website or clicking on a download button, respectively). In all the studies, *explicitness* was manipulated by explicitly mentioning the trigger in its title or omitting this information.

In *Study 1 and Study 2* it was expected that presenting a notice appearing right after its trigger (H1)/ mentioning the trigger in the notice (H2) would improve the users' comprehension of the cause (H1a, H2a) and topic of the notice (H1b, H2b), as well as the participants' experience of comprehension (H1c, H2c). *In Study 3* it was expected that displaying the notice after a specific action (H1)/ explicitly mentioning a specific action in the notice (H2) would increase the identification of that action as the notice's cause (H1a, H2a) and improve the generic comprehension of the topic of the notice (H1b, H2b). Concerning the *transparency paradox (H3),* it was expected that the acceptance of the notice request to track the participants' activity would be more frequent in the consecutive (Study 1 and Study 2) / specific conditions (H3a), and in the explicit conditions (H3b) in all the Studies, and with a higher perceived comprehension (H3c) in Study 1 and Study 2.

Regarding the *effect of consecutiveness on comprehension and interpretation of the notice (H1 and RQ1)*, the studies showed that consecutiveness consistently affected the comprehension of the cause of the notice across the studies as hypothesized (H1a) but, contrarily to the hypothesis, did not influence the comprehension of the topic (H1b) or any other variable measuring the experience of comprehension – perceived comprehension, clarity and sense of control (H1c). These findings may be due to the high familiarity reported by participants and to habituation. When the experiments were run, cookie notices were shown every time users entered a website for the first time. The studies also suggested that participants had retained some elements of the notice and filled the rest based on their previous knowledge. This phenomenon retraces the pattern demonstrated by users exposed multiple times to other online warnings and alerts (e.g., Egelman et al., 2008; Wogalter & Vigilante, 2006; Akhawe, & Felt, 2013; Egelman & Schechter, 2013). The trigger of the notice, instead, was not commonly described in cookie notices participants could find in their daily life experience and was more easily affected by my manipulations. Additionally, Study 3 revealed that displaying the notice after a specific action had an impact on how participants interpreted the specific topic of the notice i.e., the type of data being requested (RQ1). In particular, participants who encountered the notice after performing a specific action (i.e., downloading) were more likely to associate it with a request for personal data. On the other hand, participants in the generic condition did not have any specific contextual cue since they had just entered the website, and therefore distributed their responses evenly among the proposed options. This suggests that participants in the specific condition may have linked the notice to themselves personally since it amid their action. These findings align with the suggestions of Bolchini et al. (2004) and Schaub et al. (2015) of providing users a context to understand privacy policy.

Concerning the *effect of explicitness on comprehension (H2 and RQ2)*, contrarily to the hypotheses, explicitness did not show any consistent effect on participants' comprehension (H2a, H2b) or experience of comprehension (H2c) across the studies, and it did not affect the interpretation of the notice (RQ2). As Bolchini et al. (2004) and Schaub et al. (2015) contended, contextualization proved to be a stronger strategy than the mere mention of the context in the text of the notice.

With regard to the *transparency paradox (H3)*, no consistent differences were found in the acceptance of the request to track the participants' activity depending on consecutiveness (H3a), explicitness (H3b), or perceived comprehension (H3c). In contrast to previous literature (e.g., Holland et al., 2018; Masotina et al., 2019; Paunov et el., 2019), our manipulation of transparency did not lead to the paradoxical result of increasing riskier behaviors.

The findings clarify how contextualization a strategy can be implemented. The sequential adjacency of the notice and its triggering action ("consecutiveness") improved the users' ability to correctly identify such trigger; conversely, disrupting the sequential adjacency between trigger and notice led the users to identify incorrect triggers. The findings also suggest that including the trigger as a piece of information inside the text of the notice did not improve the users' identification of the trigger. Far from implying that the content of the notice is irrelevant, these results show the explanatory power of a good contextualization. Bad contextualization, on the opposite, might lead to wrong inferences.

9.1 Design implications

The design solutions adopted in the three studies have implications on the strategies adopted in real websites to ask for the users' consent. Such strategies follow a push model, where an encompassing consent is asked at the entrance of a website as an application is downloaded, or a pull model, where a specific consent is asked about a service once it is requested by the user (Klumpe et al., 2020). The findings of my studies suggest a general principle of *pursuing contextuality with the users' actions when first displaying privacy notices relevant to such actions.* So, while at the website entrance a general policy presentation is appropriate, specific implications related to specific actions should be displayed when such actions are performed. Ignoring the context in design, does not make it disappear: any notice has a context, which will guide the interpretation of the notice's meaning in a direction that, if unguarded, can be even misleading. The advantage of exploiting contextualization is to exploit the pieces of information already relevant in the context, without the need to add them in the notice and make it longer. Such additions to the notice might be useless, as shown in my three studies, and adds nothing to its informativeness.

One can object that the contextual information is a vague concept, especially because it is unclear which configuration such information will have for each specific user. To address this difficulty, I proposed a user-centered perspective, which defines context with respect to the user's current action and finds such context in the action prior to that: the users interpret the events in the light of what has been just done, in a temporal context that is sequentially organized. This way of conceptualizing context is borrowed from the study of conversation (Schegloff, 2007) and applies well to disambiguate the events during the interaction with an interface, which unfolds as a series of turns between the user and the machine.

My results also allow to advance our understanding of the transparency heuristic, where transparency is used as a cue to trust a website; if transparency always had this effect, any interventions to improve it would always risk of achieving the undesirable outcome of increasing careless privacy-related decisions (Holland et al., 2018; Masotina et al., 2019; Paunov et el., 2019). My results are encouraging in this sense; they suggest that even if the level of transparency increases, by creating notices that are more comprehensible, the users' response to the notice is not directly affected; in my studies, such acceptance of the notice seemed affected by familiarity and privacy concerns (as in Norber et al., 2007 and Soe et al., 2020), regardless of the changing levels of comprehension. Transparency might be more influential a criterion when users face unfamiliar transactions, or websites where sensitive data is collected.

9.2 Limits and future work

While the studies presented in this section demonstrated the effectiveness of contextualization as a user-centric solution, they also present some limits.

Response to the notice. The appearance of the privacy notice was part of a study, and the content of the website and the data allegedly collected were not sensitive. Therefore, the time taken to click the button on the notice as well as the rate of acceptance of the privacy policy is only to be appreciated in its possible variation across conditions; they are not to be taken as an absolute estimation generalizable to real-world situations. The privacy decision was only a secondary interest in this project, for their relation to transparency (transparency paradox).

Exclusion of non-readers. In studies 1 and 2, I excluded non-readers from the part of the survey testing comprehension and all the parts following it. These excluded non-readers also from the analysis of perceived comprehension, clarity, and control; but non-readers made a privacy decision, so it would have been interesting to measure their perceived comprehension and control. To do so, I should have filtered out non-readers at the time of the analysis, and not of the data collection.

Context in the delayed conditions. It would have been interesting to check which action was performed by the users during the condition with delays. Having this information would allow creating a subset of users who interacted with the website during the delay, and for whom the sequentiality of trigger and the notice was more remarkably broken. I recommend logging this data in case the study is replicated; I remedied this in the third study, where the action preceding the appearance of the notice was controlled.

Comprehension. The measures of comprehension used in the first and second study were too easy and allowed the participant to rely on any previous experience with similar privacy notices to pass it. This was rectified in the third study.

In general, the notice in my studies was recognized as familiar to participants and indeed was designed to resemble a common privacy notice. It would have been interesting to know whether an unusually looking notice would have been affected by my manipulation. At the same time, it would be interesting to continue the investigation of the sensemaking process and the heuristics used by visitors who face a privacy notice while using a website or a digital device: which cue they recognize as relevant, and which inferences they make based on these cues.

SECTION II

Transparency Concerns in Real-Life Discourses

Section II

CHAPTER 10.
STUDY 4

The goal of this study is to explore a method to capture how transparency is voiced in spontaneous discourses. This goal is consistent with the user-centered approach to transparency embraced in this dissertation in the sense of accessing the way transparency is talked about by citizens. Getting familiar with these concerns and terminology is a mean to reach out for real-life relevancy when designing for transparency and to create a symmetric ground between designers and users. Being guided by the research question "*How can we detect any references to transparency in real-life discourses?*", this study is explorative in nature; the discourses analyzed in this study are newspaper articles in the UK press related to a specific type of artificial intelligence, GPT-3 (Generative Pre-trained Transformer 3). GPT-3 is an autoregressive language model that uses deep learning to produce human-like text (Brown et al., 2020). It was released on 11/06/2020, and, for the first time, the quality of the text generated was high enough to make it difficult to determine whether a human wrote it. According to Elkins & Chun (2020) it "writes better than some humans". Furthermore, its applications are broad, since with a simple prompt written in natural language it can, e.g., translate, summarize, comment and write in every style. It has been defined as the "biggest transformation of the writing process since the word processor" (Floridi & Chiriatti, 2020).

The methodology, based on a mixed quantitative and qualitative textual analysis, thereby showcased can be transferred and applied to other corpora (e.g., discussions on Internet fora), other technologies, other authors differing in expertise, nationality and other relevant variables. In the rest of this section, I will describe the methodology of the study and the results.

10.1 Corpus

A corpus was created – i.e., a collection of natural language textual data (texts and/or transcription of speech or signs) constructed with a specific purpose (Björkenstam, 2013) – composed by the articles published in the UK offshoot of one of the most popular magazines about technology – Wired[3] – and the United Kingdom national newspapers. The documents were retrieved through the Wired UK website (https://www.wired.co.uk/) and LexisNexis+ news database[4] (https://www.lexisnexis.com/uk/legal/news). The search included the UK national newspapers listed in the database: The Daily Mail (London), The Daily Mirror, The Daily Telegraph (London), Daily Star Online, Mail on Sunday (London), The Guardian (London), The European, The Express, The Independent (United Kingdom), The Observer (London), The People, The Sun (England), The Sunday Express, The Sunday Mirror, The Sunday Telegraph (London), The Sunday Times (London), and The Times (London). The keywords searched were "*'gpt-3'* OR '*gpt3'*" (the search engines were not case-sensitive); the search retrieved every article citing it. The period covered by the search was from 11/06/2020 (i.e., the date of GPT-3 release) to 20/10/2022.

The search produced a match in 8 of the 18 journals considered, totaling 101 articles. After a first screening, 22 were excluded, either because the keyword was not present in the article content (e.g, it was in the "related articles" section, N = 7) or for being duplicates (N = 15). The final corpus was composed of 79 articles (Table 22) written by: 53 different journalists (2 articles were the result of the collaboration between two journalists), GPT-3 (N = 1), and the editorial board (N = 2).

[3] Wired is a monthly American magazine focused on how emerging technologies affect culture, economy, and politics. According to its Publisher's records, it reaches more than 53.1 million people monthly (https://www.condenast.com/brands/wired).

[4] NexisLexis contains more than 4 billion searchable documents from 36,000 global business and legal sources, including full-text articles from UK National and regional newspapers.

Table 22. Search results for GPT-3. Results = initial results; Not GPT-3 = articles excluded because the keywork was not present; Duplicates = articles excluded because the article was a duplicate; TOT = final number of articles included in the corpus.

Source	Results	Not GPT-3	Duplicates	TOT
Daily Star Online	9	0	0	9
The Guardian (London)	26	4	0	22
The Independent (UK)	15	1	2	12
The Observer (London)	9	2	0	7
The Sunday Times (London)	13	0	6	7
The Times (London)	14	0	7	7
The Daily Telegraph (London)	4	0	0	4
Wired UK	11	0	0	11
tot	101	7	15	79

All the articles full texts were copied and pasted in .txt format, reporting their date, title, author, highlights (when present), and content. Links and references to other journals' contents, as well as images, were eliminated.

10.2 Analysis

The analysis followed a mixed-method approach. As a preliminary exploration, it was checked whether the word "transparency" emerged as a keyword after analyzing the corpus with natural language processing techniques. In parallel, the content of the articles was deepened by investigating whether and how the concept of transparency was present in the articles through qualitative analysis. Finally, it was checked if the keywords were present in the most frequent words in the transparency-related sentences – identified during the qualitative analysis.

10.2.1 Keywords extraction

For retrieving the keywords, the NLTK library in Python was employed to pre-process the files in the corpus and identify the most frequent words as a preliminary step. Then, those were compared with a reference corpus to determine if they could be considered keywords using SketchEngine. Data were pre-processed by:

1. Performing a lower-case transformation (i.e., transforming upper-case letters in lower case to avoid the same word not being recognized as equal because of case sensitiveness),
2. Eliminating the stopwords through the nltk.stopwords function (i.e., common words like e.g., "the", "and" that are very frequent in text and don't convey insights)
3. Part of speech tagging through nltk.pos_tag function (i.e., associating the linguistics function -e.g., verb - to the word; this passage increases the precision of the lemmatization)
4. Lemmatizing the text through WordNet Lemmatizer (i.e., reduction of the inflectional forms to their lemma)
5. Tokenizing the strings at a word level through the word.tokenize package (tokenization converts characters in strings; in this way, the system can recognize the words).

After pre-processing the data, the words in the corpus were reduced from 98540 to 45087.

Then, a wordlist was created (i.e., a list of tokens with their associated frequencies) and the 30 most frequent tokens were extracted. All the instances presenting those tokens were read in their context through the NLTK concordance function.

The 30 most frequent tokens frequencies found were then compared to a reference corpus to check if they were more frequent in my corpus than in the reference corpus. This method allows to identify keywords, i.e., those nouns that appear more frequently in a corpus than in general language (https://www.sketchengine.eu/quick-start-guide/keywords-and-terms-lesson/). This passage permits to appreciate the peculiarity of a corpus. This part of the analysis was performed by exploiting the keyword function provided by SketchEngine software (https://www.sketchengine.eu/). The English Web 2020 (enTenTen20) corpus was used as a reference corpus. It comprises texts collected from the Internet, downloaded between 2019 and 2021, consisting of 36 billion words. The quality of this corpus is guaranteed since it was checked for poor-quality text and spam (https://www.sketchengine.eu/ententen-english-corpus/). The software compared the frequencies of all the tokens in my corpus with those in the reference corpus and returned the list of the keywords accomplished with their keyness score (the software calculates this score using the simple math method; https://www.sketchengine.eu/documentation/simple-maths/). Then, it was checked if the tokens found in the top 30 were present in the list. Only the tokens having this characteristic were considered keywords.

10.2.2 Manual coding

The document in the corpus were all coded using a coding scheme. UAMCorpusTool v.6.2e was used for manually annotating the texts (http://www.corpustool.com/). The focus of the annotation were the utterances connected to the concept of Transparent A.I. To identify these utterances, I relied on the definition provided in the guidelines for A.I. transparency (European Commission, 2018) in its three declinations: the datasets and processes that yield the A.I. decision should be documented (traceability); A.I. should communicate meaningful insights about its decision-making processes to their users (explicability); users should be aware of the presence of an A.I., about its nature, and must be informed of the system's capabilities and limitations (communication). For every utterance, it was annotated which of those three declinations was present (transparency acceptation). It was also annotated the way in which transparency was referred to (transparency presence): an "explicit" reference was a statement about the importance of A.I. being traceable, explicable, or communicative, according to the definition provided above. An example is in TG11: *Humans are stumbling into an era when the more powerful the A.I. system, the harder it is to explain its actions. How can we tell if a machine is acting on our behalf and not acting contrary to our interests?* An implicit reference was an utterance referred to the transparency definitions without stating its importance. For example, in TI3 we can read *Bot posing as human fooled people on Reddit for an entire week;* in this case, the reference is to transparency as communication. Finally, transparency can be mentioned by referring to the way A.I. works to reach its result (description), implying that people are usually unaware of it (e.g., TG17: *Neither LaMDA nor any of its cousins (GPT-3) are remotely intelligent. All they do is match* *patterns, draw from massive statistical databases of human language. The patterns might be cool, but language these systems utter doesn't actually mean anything at all*).

10.2.3 Transparency and corpus keywords

The utterances singled out manually were pre-processed, and a frequency list was created as described in Section 10.2.1. The list was then compared to the corpus keywords to identify the overlap between the two.

10.3 Results

10.3.1 Results of the keywords' extraction

The list of the most frequent 30 tokens in my corpus (Table 23) is composed of nouns (N = 14), adjectives (N = 2), verbs (2), auxiliary verbs (N = 1), phrasal verbs (N = 3), and adverbs (N = 3). From a focus on the nouns, we can observe their relationship with the system (*artificial, ai, gpt, machine, system, model, intelligence, technology*), with the stakeholders involved in the system use (*human, one, people, company, world*), and with the concept of time and novelty (*time, year, new*). The check of the concordances suggests that the verb *to write* can be mainly related to actions performable by the system, the verb *to work* is present in the dual acceptation of *doing work* and *functioning,* while the high frequency of the verb *to say* is mainly due to the presence of reported speeches. The most frequent auxiliary verb (*could*) was used primarily in its hypothetical acceptation. Finally, the verbs *to go, to make,* and *to get* were used primarily in their phrasal function, assuming different meanings depending on the entity they were associated with.

Six of the 30 most frequent words resulted in being also in the list of keywords produced by SketchEngine. Ordered by relevance, they were *gpt, ai, artificial, intelligence, machine,* and *human*.

Table 23. Top 30 tokens in the corpus. The words identified as keywords are reported in bold. NA = not available in the keywords.

TOP 30			
Rank	Token	N	Keyness Score
1	ai	596	188.32
2	say	395	Na
3	human	348	11.72

TOP 30			
Rank	Token	N	Keyness Score
4	gpt	294	2729.08
5	use	283	Na
6	like	255	Na
7	one	253	Na
8	write	252	Na
9	make	240	Na
10	machine	203	13.48
11	could	200	Na
12	system	200	Na
13	work	197	Na
14	model	193	Na
15	people	192	Na
16	company	190	Na
17	year	179	Na
18	would	174	Na
19	new	170	Na
20	intelligence	154	28.90
21	language	152	Na
22	world	145	Na
23	text	145	Na
24	technology	142	Na
25	get	140	Na
26	time	140	Na
27	even	139	Na
28	also	138	Na
29	go	137	Na
30	artificial	132	59.70

10.3.2 Results of the manual coding

In my corpus, 22 utterances explicitly referred to transparency were identified. This number increased when considering an implicit reference to this concept (N = 61) or the provision of a description of its functioning (N = 45). In total, 137 utterances referred to transparency were identified across 55 articles (69.62% of the articles in the corpus). The following paragraphs will deepen the results for each acceptation.

Across the corpus, the speakers referred to the acceptation of transparency as *communication* in 101 instances across 53 articles. Table 24 synthetizes the frequencies it was referred to explicitly, implicitly, or as providing the reader a description, reporting an example for each.

Table 24. Results for transparency – communication.

Presence	Instances (N)	Articles (N)	E.g.,
Explicit	9	8	TST1: You might think people would understand that these machines merely simulate compassion. But, as Peter (who was evidently not senile) told Easton: "We like to believe that empathy is a human trait but, troubling though it might be, it seems that some robots are more caring than some people."
Implicit	47	31	TST7: We have all read stories that have been written by software. They are just not labelled as such.
Description	45	31	TG5: "By collecting a historic repository of human-made speech, GPT-3 can map out patterns in how we communicate, using those rules to create new content. […] These machines do not have will, they do not have originality and they cannot claim authorship. In fact, this week's "robot authored" op-ed was a human affair from beginning to end. It was human beings who selected the prompt for the piece […]. Then, it was told who it was, what humans thought about it, what we feared and what we wanted. [...] Again, it was human authors who selected how to pare down the resulting material, discarding nearly 90% of what GPT-3 created. Without these human decisions - both the inputs and edits - there would have been no essay. GPT-3 would have nothing to contribute to the public discourse, as it has no thoughts of its own".

Across the corpus, the speakers referred to the acceptation of transparency as *explicability* in 22 instances across 17 articles. Table 25 synthetizes the frequencies it was referred to explicitly and implicitly and provides an example for each.

Table 25. Results for transparency – explicability

Presence	Instances (N)	Articles (N)	E.g.,
Explicit	10	8	WIUK1: (…) I've been saying this for ages - there's a broader cultural point around how important it is to create a legibility of AI - creating a way for people to understand how AI even works.
Implicit	12	11	TT8: GPT-3 will also remain in a limited beta for academics to test the capabilities and limitations of the model.

Across the corpus, the speakers referred to the acceptation of transparency as *traceability* in 5 instances across 3 articles. Table 26 synthetizes the frequencies it was referred to explicitly and implicitly, and provides an example for each.

Table 26. Results for transparency – traceability

Presence	Instances (N)	Articles (N)	E.g.,
Explicit	3	1	TI13: Wolf said that opening up the datasets used for language models helps humans better understand their biases.
Implicit	2	2	TG9: AI algorithms rely on data, and if that data is coming from a single demographic, it will reflect that group's biases and blind spots.

Table 27 reports a synthesis of the frequencies of the transparency acceptations and their presence across the articles.

Table 27. Results for the transparency acceptations.

Acceptation	Explicit	Implicit	Description	Instances (TOT)	Articles (TOT)
Communication	9	47	45	101	53
Explicability	10	12	-	22	17
Traceability	3	2	-	5	3

10.3.3 Manually annotated references and keywords

Table 28 reports the 30 most frequent tokens in the manually annotated utterances referring to transparency (the total token frequency after the pre-processing was 2822).

Table 28. Top 30 words in the transparency-related utterances. In bold the words identified as keywords in the corpus; for these the % on the corpus total frequency, and an example of their use are provided as well.

Top 30 words in the transparency-related utterances

Rank	Token	N	% on corpus tot	e.g.,
1	human	69	19.83%	WIUK3: we need to be careful about how such AI is used to generate articles, speeches, or any other text, read by the Queen or not - and make sure a **human** is accountable and responsible.
2	ai	53	8.89%	TG9: The emphasis on diverse data is important, because it highlights a misconception about **AI**: that it is somehow objective because it is the result of computation.
3	gpt	40	13.61%	TG12: The argument given is that building these systems helps us understand risks and develop solutions, but what did we learn between GPT-2 and **GPT-3**? It's just a bigger model with bigger problems.
4	make	26	-	-
5	machine	25	12.32%	TST3: (…) measures should be put in place to make it clear when text has been made by a **machine**, and laws should be drafted against passing off AI generated text as human. As AI improves, creating fake personas will only get easier.
6	write	25	-	-
7	model	24	-	-
8	say	23	-	-

Top 30 words in the transparency-related utterances

Rank	Token	N	% on corpus tot	e.g.,
9	text	22	-	-
10	system	21	-	-
11	like	20	-	-
12	understand	20	-	-
13	data	20	-	-
14	language	19	-	-
15	learn	18	-	-
16	word	17	-	-
17	think	17	-	-
18	lamda	14	-	-
19	way	14	-	-
20	people	13	-	-
21	intelligence	13	8.44%	TI17: (…) One of Google's (former) ethics experts doesn't understand the difference between sentience (AKA subjectivity, experience**),** **intelligence**, and self-knowledge
22	one	13	-	-
23	ability	12	-	-
24	create	12	-	-
25	know	12	-	
26	work	12	-	-
27	image	11	-	-
28	get	11	-	-
29	google	11	-	
30	tell	11	-	-

Five out of the six keywords of the corpus are also present in the manually annotated utterances: *human, ai, gpt, machine, intelligence*. In particular, *human* represents the most frequent token. Its frequency in this subsample represents almost 1/5 of the total occurrences of this token in the corpus.

Section II

CHAPTER 11.
CONCLUSIONS OF SECTION II

In sum, the results of this section show that even though the documents in the corpus might not directly use the term ‘transparency’ to a great extent, transparency related concerns are nonetheless present. They are represented by the occurrences in the three acceptations of transparency, namely traceability (the datasets and processes that yield the A.I. decision should be documented), explainability (A.I. should communicate meaningful insights about its decision-making processes to their users), and communication (users should be aware of the presence of an A.I., about its nature, and must be informed of the system’s capabilities and limitations).

Moreover, the stretches of the document referred to transparency share most keywords with the entire corpus (*human, ai, gpt, machine, intelligence*), showing that transparency-related passages in the documents are not peripheral to the core arguments conveying by the articles, but actually part of them.

My method, substantially, allows to:

- assessing the relevance of a certain topic in a corpus by checking the overlap between the corpus’s keywords and some target passages in a corpus, related to the object of investigation.
- identifying the terminology used by the document’s authors to refer to the object of investigation, after singling out the portions of text addressing that object. Such terminology can be used in the design phase of privacy-related information in order to make it more recognizable and appreciable by a specific class of user.

As such, the method can be adopted to investigate the way in which transparency is referred to in other corpora and the relevance transparency has in each specific corpora, without relying exclusively on a qualitative or quantitative procedures.

CHAPTER 12.
GENERAL CONCLUSIONS

The present dissertation has focused on implementing the principle of transparency in the realm of digital technologies. In this context, transparency has been called for to shield users from the risks of information technology and emphasized regarding the users' data and AI's employment. Indeed, transparency promises to act as a safeguard by allowing users to see and understand, and thus to have control over their choices, decisions, and rights. Stepping back to its definition, however, makes it apparent that transparency is a medium, *something that can be seen through.* As such, it becomes part of the object it makes visible, influencing the way the latter is seen and thereby understood. Consequently, drawing attention to how transparency is built becomes pivotal for enabling the fulfillment of its promise.

This thesis has embraced a user-centered perspective and approached transparency as a design goal and as a discourse topic. Transparency as a design goal has been sought in Section I, where a user-centered solution - namely contextualization - was tested to increase the transparency of privacy notices. Transparency as a discourse topic has been dealt with in Section II, by exploring a method to find how transparency is treated in real life discourses and applying it in the context of AI as a case study. The user-centered perspective stresses the importance of focusing on the users during the design process. The present work aligns with this perspective, in the assumption that prioritizing the recipient of information, their language, and their understanding is a crucial prerequisite for achieving the goal of designing transparency in a way that can effectively empower the desired target, i.e., the final users.

The contribution of this dissertation is twofold. First, this work provides evidence about how contextualization can be applied to design privacy notices that are genuinely transparent – i.e., understandable - to their users. Finding new strategies for improving users' comprehension is necessary to make privacy notices truly compliant with *the spirit* of current regulations. The latter protects the users' right to privacy by requiring digital platforms to ask for the users' *informed* consent before collecting any users' data. It also specifies that such informed consent has to implement a principle of transparency, according to which the information provided must be clear and easy to comprehend. Merely overwhelming the users with information they will barely read is an easy way for being superficially compliant with the regulations, but has been proven to be ineffective to achieve that empowerment that transparency promises to grant. Section I demonstrates that it is possible to exploit the *context of appearance* of privacy notices to guide the users' interpretation and understanding of their contents. Borrowing the definition of context from ethnomethodology and conversation analysis, it encourages designers to carefully consider the *action performed by the user on the website immediately preceding the appearance of the privacy notice* and to *pursue contextuality with the users' actions when first displaying privacy notices relevant to such actions.* Aligning the information with the users' action at stake provides the users with a context able to increase their understanding - without the need to decrease the usability of the notice itself by adding this piece of information in its text. At the same time, ignoring the context can lead users to misinterpret the notice contents and draw the wrong conclusions. As time passes, it is likely to foresee regulation and guideline updates. Still, the proposed strategy remains relevant since, regardless of the type of privacy notice or possible updates, every notice will come with a context.

The second contribution of this dissertation is methodological. Section II has explored a method to detect references to transparency in real-life discourses. Understanding how people discuss transparency is a way to access their language, thereby establishing a common ground between designers and users and reducing the asymmetries between the former and their stakeholders. The method proposed here relies on qualitative and natural language processing techniques. It enables the evaluation of the importance of a particular subject in a body of text and facilitates the identification of the language used by the informers to refer to the topic. This language can be utilized when creating privacy-related information to make it more recognizable and understandable to a particular group of users. Here, the methodology has been applied to the topic of transparency in AI. This aspect was one of the calls for transparency emphasized in the digital society. As a case study, the documents included newspaper articles published in the UK about GPT-3. Nonetheless, the possible applications for exploring transparency as a discourse are broader. The process starts with the creation of the corpus. As long as discourses can be represented by textual data, a corpus can be created by selecting *who* the informants of interest are*, where* they are talking about technology, and *which technology* is the object of the talk. In this light, to name a few examples, the analysis may target other Countries – e.g., newspaper articles published in *Italy* about GPT-3. It can consider other technologies - e.g., newspaper articles published in the UK about *Dall-E* (the "cousin" of GPT-3 able to generate images). It can also focus on other “places” where the discourse of transparency can emerge – e.g., *tweets* published on *Twitter* about GPT-3. As a further step, the same methodology can be applied to more than one corpus differing in one of the features above or in the period of time considered. In this way, it is possible to appreciate the evolution of the topic over time and comparing how different users talk about transparency depending on their characteristics (e.g., age, nationality, profession), the "place" where they are speaking (e.g., specific social networks, traditional media) and the technology they are talking about.

In conclusion, the present dissertation represents a contribution towards achieving genuine transparency in digital technologies. It provides applicative insights for implementing a user-centered transparency, in the belief that if transparency is meant to empower the users, then the users have to be the focus in its design.

REFERENCES

Acquisti, A., Adjerid, I., & Brandimarte, L. (2013). Gone in 15 seconds: The limits of privacy transparency and control. *IEEE Security & Privacy, 11*(4), 72-74.

Acquisti, A., Adjerid, I., Balebako, R., Brandimarte, L., Cranor, L. F., Komanduri, S., ... & Wang, Y. (2017). Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR), 50*(3), 44.

Adjerid, I., Acquisti, A., Brandimarte, L., & Loewenstein, G. (2013, July). Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proceedings of the ninth symposium on usable privacy and security* (pp. 1-11).

Agozie, D. Q., & Kaya, T. (2021). Discerning the effect of privacy information transparency on privacy fatigue in e-government. *Government Information Quarterly*, 101601.

Akhawe, D., & Felt, A. P. (2013, August). Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness. In *USENIX security symposium* (Vol. 13).

Albu, O. B., & Flyverbom, M. (2019). Organizational transparency: Conceptualizations, conditions, and consequences. *Business & Society, 58*(2), 268-297.

Algorithm Watch (2019). Automating Society – Taking Stock of Automated Decision-Making in the EU.

Alloa, E. (2018). Transparency: A magic concept of modernity. In *Transparency, society and subjectivity* (pp. 21-55). Palgrave Macmillan, Cham.

Almuhimedi, H., Porter Felt, A., Reeder, R. W., & Consolvo, S. (2014, July). Your Reputation Precedes You: History, Reputation, and the Chrome Malware Warning. Proceeding presented at *Symposium On Usable Privacy and Security,* Menlo Park, CA.

Amoruso, L., Finisguerra, A., & Urgesi, C. (2018). Contextualizing action observation in the predictive brain: Causal contributions of prefrontal and middle temporal areas. *Neuroimage, 177*, 68-78.

Anderson, B. B., Kirwan, C.B., Jenkins, J. L., Eargle, D., Howard, S., & Vance, A. (2015, April). How Polymorphic Warnings Reduce Habituation in the Brain: Insights from an fMRI Study. Proceeding presented at *33rd Annual ACM Conference on Human Factors in Computing Systems*, Seoul, Repubblica di Corea.

Antón, A. I., Earp, J. B., He, Q., Stufflebeam, W., Bolchini, D., & Jensen, C. (2004). Financial privacy policies and the need for standardization. *IEEE Security & privacy, 2*(2), 36-45.

Antón, A. I., Earp, J. B., & Young, J. D. (2010). How internet users' privacy concerns have evolved since 2002. *IEEE Security & Privacy, 8*(1), 21-27

Argo, J. J., & Main, K. J. (2004). Meta-analyses of the effectiveness of warning labels. *Journal of Public Policy & Marketing, 23*(2), 193-208.

Balebako, R., Schaub, F., Adjerid, I., Acquisti, A., & Cranor, L. (2015, October). The impact of timing on the salience of smartphone app privacy notices. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices* (pp. 63-74).

Barth, S., Ionita, D., & Hartel, P. (2022). Understanding Online Privacy—A Systematic Review of Privacy Visualizations and Privacy by Design Guidelines. *ACM Computing Surveys (CSUR), 55*(3), 1-37.

Baume, S. (2018). Publicity and transparency: the itinerary of a subtle distinction. In *Transparency, Society and Subjectivity (*pp. 203-224). Palgrave Macmillan, Cham.

Baume, S., & Papadopoulos, Y. (2018). Transparency: from Bentham's inventory of virtuous effects to contemporary evidence-based scepticism. *Critical Review of International Social and Political Philosophy, 21*(2), 169-192.

Bazire, M., & Brézillon, P. (2005, July). Understanding context before using it. In *International and Interdisciplinary Conference on Modeling and Using Context* (pp. 29-40). Springer, Berlin, Heidelberg.

Birchall, C. (2011). Introduction to 'Secrecy and Transparency' The Politics of Opacity and Openness. *Theory, Culture & Society, 28*(7-8), 7-25.

Birge, C. (2009, October). Enhancing research into usable privacy and security. In *Proceedings of the 27th ACM international conference on Design of communication* (pp. 221-226).

Biselli, Tom and Reuter, Christian, "On the Relationship between IT Privacy and Security Behavior: A Survey among German Private Users" (2021). *Wirtschaftsinformatik 2021 Proceed*ings.

Bolchini, D., He, Q., Antón, A. I., & Stufflebeam, W. (2004, July). "I Need It Now": Improving Website Usability by Contextualizing Privacy Policies. In *International Conference on Web Engineering* (pp. 31-44). Springer, Berlin, Heidelberg.

Bonner, M. F., & Epstein, R. A. (2021). Object representations in the human brain reflect the co-occurrence statistics of vision and language. *Nature Communications, 12*(1), 1-16.

Bravo-Lillo, C., Cranor, L., & Komanduri, S. (2014, July). Harder to Ignore?: Revisiting Pop-Up Fatigue and Approaches to Prevent It. Proceeding presented at *Symposium On Usable Privacy and Security*, Menlo Park, California.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, *33*, 1877-1901.)

Bruner, J. (1990). *Acts of meaning*. Harvard University Press.

California Consumer Privacy Act, 2018 Cal. Legis. Serv. Ch. 55 (A.B. 375) (WEST)

Canadian Standards Association (CSA Group). 2014. Model Code for the Protection of Personal Information of 23 December 2014, First Published in March 1996; Reaffirmed 2001. CAN/CSA-Q830-96. (December 2014).

Cavoukian, A. (2006). *Creation of a Global Privacy Standard. Report*. Information and Privacy Commissioner of Ontario, Ontario, Canada.

Chang D, Krupka E.L., Adar E, Acquisti A. (2016). Engineering information disclosure: norm shaping designs. *Proceedings of the Conference on Human Factors in Computing Systems* (pp. 587 – 597), ACM.

Christensen, L. T., & Cornelissen, J. (2015). Organizational transparency as myth and metaphor. *European Journal of Social Theory, 18*(2), 132-149.

Cialdini, R. (2016). Pre-suasion: A revolutionary way to influence and persuade. Simon and Schuster.

Cukier, K., & Mayer-Schoenberger, V. (2013). The rise of big data: How it's changing the way we think about the world. *Foreign Aff., 92*, 28.

Digital Markets Act, 2022. Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828.

Digital Services Act, 2022. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC.

Dhamija, R., Tygar, J. D., & Hearst, M. (2006). Why phishing works. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems (pp. 581–590)*. ACM Press: New York.

Earp, J. B., & Baumer, D. (2003). Innovative web use to learn about consumer behavior and online privacy. *Communications of the ACM*, 46(4), 81-83.

Ebert, N., Ackermann, K. A., & Heinrich, P. (2020, April). Does context in privacy communication really matter?—A survey on consumer concerns and preferences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1-11).

Egelman, S., Cranor, L. F., & Hong, J. (2008, April). You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1065-1074). ACM.

Egelman, S., & Schechter, S. (2013, April). The importance of being earnest [in security warnings]. In *International Conference on Financial Cryptography and Data Security* (pp. 52-59). Springer, Berlin, Heidelberg.

Egelman, S., Tsai, J., Faith Cranor, L., & Acquisti, A. (2009, April). Timing is everything?: the effects of timing and placement of online privacy indicators. Proceeding presented at *SIGCHI Conference on Human Factors in Computing Systems*, Boston, MA, USA.

Elkins, K., & Chun, J. (2020). Can GPT-3 pass a writer's Turing Test? *Journal of Cultural Analytics,* 2371, 4549.

European Commission (2018), Ethics Guidelines For Trustworthy A.I.

European Commission. (2019). Special Eurobarometer 487a: The general data protection regulation. Brussels: European Commission.

European Union (2016). General Data Protection Regulation 2016/679.

Felt, A. P., Ainslie, A., Reeder, R. W., Consolvo, S., Thyagaraja, S., Bettes, A., … Grimes, J. (2015). Improving SSL warnings: Comprehension and adherence. In *Conference on Human Factors in Computing Systems - Proceedings* (Vol. 2015-April, pp. 2893–2902). Association for Computing Machinery.

Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, *30*(4), 681-694.

Franke, R. H., & Kaul, J. D. (1978). The Hawthorne experiments: First statistical interpretation. *American sociological review*, 623-643.

Franz, A., Zimmermann, V., Albrecht, G., Hartwig, K., Reuter, C., Benlian, A., & Vogt, J. (2021, August). SoK: Still Plenty of Phish in the Sea—A Taxonomy of User-Oriented Phishing Interventions and Avenues for Future Research. In *Seventeenth Symposium on Usable Privacy and Security ({SOUPS} 2021)* (pp. 339-358). USENIX Association.

Gibson, J.J. (1966). *The Senses Considered as Perceptual Systems*. Allen and Unwin, London.

Gluck, J., Schaub, F., Friedman, A., Habib, H., Sadeh, N., Cranor, L. F., & Agarwal, Y. (2016). How short is too short? Implications of length and framing on the effectiveness of privacy notices. In *Twelfth Symposium on Usable Privacy and Security ({SOUPS} 2016)* (pp. 321-340).

Goffman, E. (2006). The presentation of self. Life as theater: A dramaturgical sourcebook, 129-139.

Gonzalez, E. G., De Hert, P., & Papakonstantinou, V. (2020). The Proposed ePrivacy Regulation: The Commission‘s and the Parliament‘s Draft s at a Crossroads?. *Data Protection and Privacy. Data Protection and Democracy.* Hart Publishing, 267-298.

Goodwin, C. & Duranti, A. (1992) Rethinking context: An introduction. in Duranti, A., & Goodwin, C. (Eds) *Rethinking Context: Language as an Interactive Phenomenon* (pp. 1-42). Cambridge, UK: Cambridge University Press.

Groves, P. M., & Thompson, R. F. (1970). Habituation: a dual-process theory. *Psychological review, 77*(5), 419.

Gupta, A. (2008). Transparency under scrutiny: Information disclosure in global environmental governance. *Global environmental politics, 8*(2), 1-7.

Habib, H., Zou, Y., Yao, Y., Acquisti, A., Cranor, L., Reidenberg, J., ... & Schaub, F. (2021, May). Toggles, dollar signs, and triangles: How to (in) effectively convey privacy choices with icons and link texts. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (pp. 1-25).

Hansen, H. K., Christensen, L. T., & Flyverbom, M. (2015). Introduction: Logics of transparency in late modernity: Paradoxes, mediation and governance. *European Journal of Social Theory, 18*(2), 117-131.

Harcourt, B. E. (2015). *Exposed: Desire and disobedience in the digital age*. Harvard University Press.

Hargis, G. (2000). Readability and computer documentation. *ACM Journal of Computer Documentation (JCD), 24*(3), 122-131.

Herbert, F., Schmidbauer-Wolf, G. M., & Reuter, C. (2021). Who Should Get My Private Data in Which Case? Evidence in the Wild. In *Mensch und Computer 2021* (pp. 281-293).

Herley, C. (2009, September). So long, and no thanks for the externalities: the rational rejection of security advice by users. Proceeding presented at workshop on New security paradigms workshop, Oxford, Regno Unito.

Hsieh, A. Y., Lo, S. K., Chiu, Y. P., & Lie, T. (2020). Do not allow pop-up ads to appear too early: Internet users' browsing behavimy to pop-up ads. *Behaviour & Information Technology,* 1-10.

Holland, D., Krause, A., Provencher, J., & Seltzer, T. (2018). Transparency tested: The influence of message features on public perceptions of organizational transparency. *Public Relations Review, 44*(2), 256-264.

Hoofnagle, C. J., & Urban, J. M. (2014). *Alan Westin's privacy homo economicus*. Wake Forest L. Rev., 49, 261.

Kahneman, D. (2017). *Thinking fast and slow.* Farrar, Straus, and Giroux.

Kelley, P. G., Bresee, J., Cranor, L. F., & Reeder, R. W. (2009, July). A" nutrition label" for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security* (pp. 1-12).

Kline, K., & Holland, K. (2020). From Representation to Simulation. In *Jean Baudrillard and Radical Education Theory* (pp. 43-60). Brill.

Klumpe, J., Koch, O. F., & Benlian, A. (2020). How pull vs. push information delivery and social proof affect information disclosure in location based services. *Electronic Markets, 30*(3), 569-586.

Knijnenburg, B., & Cherry, D. (2016). Comics as a medium for privacy notices. In *Twelfth Symposium on Usable Privacy and Security ({SOUPS} 2016).*

Kobsa, A., & Teltzrow, M. (2004, May). Contextualized communication of privacy practices and personalization benefits: Impacts on users' data sharing and purchase behavior. In *International Workshop on Privacy Enhancing Technologies* (pp. 329-343). Springer, Berlin, Heidelberg.

Koivisto, I. (2014). Varieties of good governance: A suggestion of discursive plurality. *International Journal for the Semiotics of Law-Revue internationale de Sémiotique juridique, 27*(4), 587-611.

Koivisto, I. (2022). *The Transparency Paradox*. Oxford University Press.

Kools, M., van de Wiel, M. W., Ruiter, R. A., Crüts, A., & Kok, G. (2006). The effect of graphic organizers on subjective and objective comprehension of a health education text. *Health Education & Behavior, 33*(6), 760-772.

Korunovska, J., Kamleitner, B., & Spiekermann, S. (2020, June). The Challenges and Impact of Privacy Policy Comprehension. ECIS.

Kretschmer, M., Pennekamp, J., & Wehrle, K. (2021). Cookie banners and privacy policies: Measuring the impact of the GDPR on the web. *ACM Transactions on the Web (TWEB)*, *15*(4), 1-42.

Kühtreiber, P., Pak, V., & Reinhardt, D. (2022). Replication: The Effect of Differential Privacy Communication on German Users' Comprehension and Data Sharing Attitudes. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)* (pp. 117-134).

Mäihäniemi, B. (2020). Competition Law and Big Data: Imposing Access to Information in Digital Markets. Edward Elgar Publishing.

Malkin, N., Mathur, A., Harbach, M., & Egelman, S. (2017, April). Personalized Security Messaging: Nudges for Compliance with Browser Warnings. Proceeding presented at *2nd European Workshop on Usable Security*, Parigi, Francia.

Masotina, M., Pluchino, P., Freuli, F., Gamberini, L., & Spagnolli, A. (2019, September). Transparency Heuristic: Effect of Implicitness of Online Data Acquisition on Sensitivity Perception. In *IFIP Conference on Human-Computer Interaction* (pp. 676-679). Springer, Cham.

Milne, G. R., & Culnan, M. J. (2004). Strategies for reducing online privacy risks: Why consumers read (or don't read) online privacy notices. *Journal of interactive marketing, 18*(3), 15-29.

Moe, W. W. (2006). Should I wait to promote?: The effect of timing on response to pop-up promotions. In *Marketing Science Conference*, University of Maryland.

Morris, E. (2011). Believing is Seeing: Observations on the Mysteries of Photography Errol. New York, NY: Penguin Press.

Nissenbaum, H. (2011). A contextual approach to privacy online. *Daedalus, 140*(4) 32–48.

Norberg, P. A., Horne, D. R., & Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of consumer affairs, 41*(1), 100-126.

Norman, D. A. (1988). *The psychology of everyday things.* Basic books.

Obar, J. A., & Oeldorf-Hirsch, A. (2018). The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society*, 1-20.

OECD (2013). *The OECD Privacy Framework*. Booklet. Organisation for Economic Co-operation and Development, Paris, France.

Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of experimental social psychology, 45*(4), 867-872.

Oulasvirta, A., Suomalainen, T., Hamari, J., Lampinen, A., Karvonen, K.: Transparency of intentions decreases privacy concerns in ubiquitous surveillance. *Cyberpsychol. Behav. Soc. Netw. 17*(10), 633–638 (2014)

Paunov, Y., Wänke, M., & Vogel, T. (2019). Transparency effects on policy compliance: disclosing how defaults work can enhance their effectiveness. *Behavioural Public Policy, 3*(2), 187-208.

Peters, A. (2013). *Towards transparency as a global norm* (pp. 534-607). Cambridge University Press.

Pollitt, C., & Hupe, P. (2011). Talking about government: The role of magic concepts. *Public management review, 13*(5), 641-658.

Reeder, R. W., Felt, A. P., Consolvo, S., Malkin, N., Thompson, C., & Egelman, S. (2018, April). An experience sampling study of user reactions to browser warnings in the field. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (p. 512). ACM.

Reidenberg, J. R., Breaux, T., Cranor, L. F., French, B., Grannis, A., Graves, J. T., ... & Ramanath, R. (2015). *Disagreeable privacy policies: Mismatches between meaning and users' understanding.* Berkeley Tech. LJ, 30, 39.

Ringel, L. (2019). Unpacking the transparency-secrecy nexus: Frontstage and backstage behaviour in a political party. *Organization Studies, 40*(5), 705-723.

Rossi, A., & Lenzini, G. (2020). Transparency by design in data-informed research: a collection of information design patterns. *Computer Law & Security Review, 37,* 105402.

Schaub, F., Balebako, R., Durity, A. L., & Cranor, L. F. (2015). A design space for effective privacy notices. In *Eleventh Symposium On Usable Privacy and Security* (pp. 1-17).

Schegloff, E. A. (2007). Sequence organization in interaction: A primer in conversation analysis I (Vol. 1). Cambridge University Press.

Shamon, H., & Berning, C. (2019). Attention check items and instructions in online surveys with incentivized and non-incentivized samples: Boon or bane for data quality?. *Shamon, H., & Berning, CC (2020). Attention Check Items and Instructions in Online Surveys with Incentivized and Non-Incentivized Samples: Boon or Bane for Data Quality,* 55-77.

Schmidt, L., Bornschein, R., & Maier, E. (2020). The effect of privacy choice in cookie notices on consumers' perceived fairness of frequent price changes. *Psychology & Marketing, 37*(9), 1263-1276

Silic, M., & Cyr, D. (2016, July). Colmy arousal effect on users' decision-making processes in the warning message context. In *International Conference on HCI in Business, Government, and Organizations* (pp. 99-109). Springer, Cham.

Sloman, S. A., & Lagnado, D. (2015). Causality in thought. *Annual review of psychology, 66*, 223-247.

Smit, E. G., Van Noort, G., & Voorveld, H. A. (2014). Understanding online behavioural advertising: User knowledge, privacy concerns and online coping behaviour in Europe. *Computers in human behavior*, *32*, 15-22.

Smith, M., & Taffler, R. (1992). Readability and understandability: Different measures of the textual complexity of accounting narrative. *Accounting, Auditing & Accountability Journal, 5*(4).

Soe, T. H., Nordberg, O. E., Guribye, F., & Slavkovik, M. (2020, October). Circumvention by design-dark patterns in cookie consent for online news outlets. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society* (pp. 1-12).

Spagnolli, A., Frank, L. E., Haselager, P., & Kirsh, D. (2017, December). Transparency as an ethical safeguard. In *International Workshop on Symbiotic Interaction* (pp. 1-6). Springer, Cham.

Spagnolli, A., Guardigli, E., Orso, V., Varotto, A., & Gamberini, L. (2015, October). Measuring user acceptance of wearable symbiotic devices: validation study across application scenarios. In *International Workshop on Symbiotic Interaction* (pp. 87-98). Springer, Cham.

Sunshine, J., Egelman, S., Almuhimedi, H., Atri, N., & Cranor, L. F. (2009, August). Crying Wolf: An Empirical Study of SSL Warning Effectiveness. In *USENIX security symposium* (pp. 399-416).

Technical Committee ISO/IEC JTC 1/SC 27 (2011). ISO/IEC 29100:2011 Information Technology – Security Techniques – Privacy Framework. Standard ISO/IEC 29100:2011(E). *International Organization for Standardization and International Electrotechnical Commission*, Geneva, Switzerland.

Teurlings, J., & Stauff, M. (2014). Introduction: The transparency issue. Cultural Studies? *Critical Methodologies, 14*(1), 3-10.

Trabasso, T., & Van Den Broek, P. (1985). Causal thinking and the representation of narrative events. *Journal of memory and language*, 24(5), 612-630.

Tsai, C. H., & Brusilovsky, P. (2021). The effects of controllability and explainability in a social recommender system. *User Modeling and User-Adapted Interaction, 31*(3), 591-627.

Utz, C., Degeling, M., Fahl, S., Schaub, F., & Holz, T. (2019, November). (Un) informed Consent: Studying GDPR Consent Notices in the Field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (pp. 973-990).

Vail, M. W., Earp, J. B., & Antón, A. I. (2008). An empirical study of consumer perceptions and comprehension of web site privacy policies. *IEEE Transactions on Engineering Management*, *55*(3), 442-454.

Waldman, A. E. (2020). Cognitive biases, dark patterns, and the 'privacy paradox'. *Current Opinion in Psychology, 31*, 105-109.

Willems, R. M., & Peelen, M. V. (2021). How context changes the neural basis of perception and language. *Iscience, 24*(5), 102392.

Wogalter, M. S., & Vigilante, W. J. (2006). Attention switch and maintenance. *Handbook of warnings*, 245-265.

World Bank (1989). Sub-saharan Africa: from crisis to sustainable growth.

Wu, M., Miller, R. C., & Garfinkel, S. L. (2006, April). Do security toolbars actually prevent phishing attacks?. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (pp. 601-610). ACM.

Yang, W., Xiong, A., Chen, J., Proctor, R. W., & Li, N. (2017, April). Use of phishing training to improve security warning compliance: evidence from a field experiment. In *Proceedings of the hot topics in science of security: symposium and bootcamp* (pp. 52-61).

Xiong, A., Wang, T., Li, N., & Jha, S. (2020, May). Towards effective differential privacy communication for users' data sharing decision and comprehension. In *2020 IEEE Symposium on Security and Privacy (SP)* (pp. 392-410). IEEE.

Xu, H., Teo, H. H., Tan, B. C., & Agarwal, R. (2009). The role of push-pull technology in privacy calculus: the case of location-based services. *Journal of management information systems, 26*(3), 135-174.

Zaaba, Z. F., Furnell, S. M., & Dowland, P. S. (2014, November). A study on improving security warnings. In *The 5th International Conference on Information and Communication Technology for The Muslim World* (ICT4M) (pp. 1-5). IEEE

APPENDIX

A1 Goodness of fit test results

A1.1 Study 1

Table 29: Likelihood ratio test results for logistic models. *p <.05, **p<.01,***p<.001.

Outcome	Model	X^2	df	p	R^2
Cause Comprehension	Model 1 vs Model 0 direct effects	55.82	5	<.001***	.32
	Model 1.1 vs Model 0 effect of consecutiveness	47.77	1	<.001***	.27
	Model 1.2 vs Model0 effect of explicitness	3.71	1	.05	.02
	Model 2 vs Model1 direct effects + interaction	0.26	1	.61	.32
Topic Comprehension	Model 1 vs Model 0 direct effects	5.25	5	.39	.03
	Model 1.1 vs Model 0 effect of explicitness	4.43	1	.04*	.03
	Model 2 vs Model0 direct effects + interaction	5.37	6	.50	.04
Response	Model 1 vs Model 0 direct effects	21.23	6	.002**	.24
	Model 1.1 vs Model 0 effect of Privacy Concerns	9.48	1	.002**	.11
	Model 2 vs Model1 direct effects + interaction	3.14	1	.08	.27
	Model 2.1 vs Model 0 effect of consecutiveness	5.13	1	.02*	.06
	Model 3 vs Model 0 effect of PerceivedC1	1.12	1	.29	.01

Table 30: Likelihood ratio test results for ordinal models. *p <.05, **p<.01,***p<.001.

Outcome	Model	X^2	df	p	R^2
PerceivedC1	Model 1 vs Model 0 direct effects	5.54	5	.35	.02
	Model 2 vs Model 0 direct effects + interaction	6.09	6	.41	.03
PerceivedC2	Model 1 vs Model 0 direct effects	4.18	5	.52	.01
	Model 2 vs Model 0 direct effects + interaction	4.68	6	.59	.02
Sense of Control	Model 1 vs Model 0 direct effects	6.06	5	.30	.02

Outcome	Model	X^2	df	p	R^2
	Model 1.1 vs Model 0 effect of consecutiveness	4.74	1	.03*	.02
	Model 2 vs Model 0 direct effects + interaction	6.29	6	.39	.02

Table 31: Overall goodness of fit test results for linear models. *p <.05, **p<.01,***p<.001.

Outcome	Model	F	df	p	R^2
Clarity	Model 1 vs Model 0 direct effects	0.49	(5,102)	.78	-.03
	Model 2 vs Model 0 direct effects + interaction	0.41	(6, 101)	.87	-.03
Time to respond	Model 1 vs Model 0 direct effects	4.52	(5, 102)	<.001***	.14
	Model 1.1 effect of explicitness	10.98	(1,106)	.001**	.09
	Model 1.2 effect of consecutiveness	5.11	(1,106)	.03*	.04
	Model 1.3 effect of age	16.97	(1,106)	<.001***	.13
	Model 2 vs Model1 direct effects + interaction	3.63	1	.06	.16

A1.2 Study 2

Table 32: Likelihood ratio test results for logistic models. *p <.05, **p<.01,***p<.001.

Outcome	Model	X^2	df	p	R^2
Cause Comprehension	Model 1 vs Model 0 direct effects	30.40	5	<.001***	.18
	Model 1.1 vs Model 0 effect of consecutiveness	24.85	1	<.001***	.15
	Model 2 vs Model1 direct effects + interaction	2.29	1	.13	.19
Topic Comprehension	Model 1 vs Model 0 direct effects	3.66	5	.60	.07
	Model 2 vs Model 0 direct effects + interaction	4.91	6	.56	.09
Response	Model 1 vs Model 0 direct effects	5.55	6	.48	.09
	Model 2 vs Model0 direct effects + interaction	6.39	7	.50	.11
	Model 3 vs Model 0 effect of privacy concerns	4.66	1	.03*	.08
	Model 4 vs Model 0 effect of PerceivedC1	0.19	1	.66	.004

Table 33: Likelihood ratio test results for ordinal models. *p <.05, **p<.01,***p<.001.

Outcome	Model	X^2	df	p	R^2
PerceivedC1	Model 1 vs Model 0 direct effects	3.04	5	.69	.02
	Model 2 vs Model 0 direct effects + interaction	3.10	6	.80	.02
PerceivedC2	Model 1 vs Model 0 direct effects	0.76	5	.98	.003
	Model 2 vs Model 0 direct effects + interaction	1.01	6	.99	.005
Sense of Control	Model 1 vs Model 0 direct effects	7.93	5	.16	.03
	Model 1.1 vs Model 0 effect of gender	3.73	1	.05	.01
	Model 2 vs Model 0 direct effects + interaction	9.21	6	.16	.03

Table 34: Overall goodness of fit test results for linear models. *p <.05, **p<.01,***p<.001.

Outcome	Model	F	df	p	R^2
Clarity	Model 1 vs Model 0 direct effects	0.75	(5, 89)	.59	-.03
	Model 2 vs Model 0 direct effects + interaction	0.64	(6, 88)	.70	-.02

A1.3 Study 3

Table 35: Likelihood ratio test results for logistic models. *p <.05, **p<.01,***p<.001.

Outcome	Model	X^2	df	p	R^2
Cause Comprehension	Model 1 vs Model 0 direct effects	76.74	5	<.001***	.62
	Model 1.1 vs Model 0 effect of consecutiveness	75.31	1	<.001***	.61
	Model 2 vs Model1 direct effects + interaction	.004	1	.95	.62
Topic comprehension (generic)	Model 1 vs Model 0 direct effects	6.66	5	.25	.05
	Model 2 vs Model 0 direct effects + interaction	8.21	6	.22	.07
Response	Model 1 vs Model 0 direct effects	9.99	6	.12	.14
	Model 1.1 vs Model0 effect of Privacy Concerns	8.30	1	.004**	.12
	Model2 vs Model0 direct effects + interaction	12.48	7	.09	.18

Table 36: Overall goodness of fit test results for linear models. *p <.05, **p<.01,***p<.001.

Outcome	Model	F	df	p	R^2
Time to respond	Model 1 vs Model 0 direct effects	1.18	(5, 48)	.34	.02
	Model 2 vs Model 0 direct effects + interaction	1.17	(6, 47)	.34	.02

A2 Regression results for non-significant models at the goodness of fit test

A2.1 Study 1

Table 37: Regression Results for Perceived Comprehension (item 1). *p <.05, **p<.01,***p<.001.

		b	Std. Error	z	p
MODEL 1	Explicitness	0.26	0.39	0.67	.51
direct effects	Consecutiveness	0.69	0.39	1.75	.08
	Gender	0.41	0.38	1.07	.28
	Education	0.22	0.39	0.57	.57
	Age	0.33	0.40	0.83	.40
MODEL 2	Explicitness	0.48	0.47	0.98	.33
direct effects	Consecutiveness	0.97	0.55	1.76	.08
+	Gender	0.44	0.38	1.13	.26
interaction	Education	0.21	0.39	0.54	.59
	Age	0.38	0.40	0.94	.35
	Explicitness:Consecutiveness	-0.58	0.79	-0.74	.46

Table 38: Regression Results for Perceived Comprehension (item 2). *p <.05, **p<.01,***p<.001.

		b	Std. Error	z	p
MODEL 1	Explicitness	0.47	0.38	1.25	.21
direct effects	Consecutiveness	0.07	0.35	0.19	.85
	Gender	0.36	0.35	1.00	.32
	Education	0.23	0.37	0.62	.54
	Age	0.36	0.38	0.96	.34
MODEL 2	Explicitness	0.69	0.49	1.41	.16
direct effects	Consecutiveness	0.33	0.52	0.65	.52
+	Gender	0.39	0.36	1.09	.28
interaction	Education	0.25	0.37	0.66	.51
	Age	0.39	0.38	1.04	.30
	Explicitness:Consecutiveness	-0.51	0.72	-0.71	.48

Table 39: Regression Results for Clarity. *p <.05, **p<.01,***p<.001

		b	Std. Error	t	p
MODEL1	(Intercept)	3.62	0.17	20.80	<2e-16 ***
direct effects	Explicitness	-0.13	0.15	-0.84	.41
	Consecutiveness	0.12	0.15	0.78	.44
	Gender	0.13	0.15	0.90	.37
	Education	0.04	0.15	0.26	.80

		b	Std. Error	t	p
	Age	0.09	0.15	0.58	.56
MODEL 2	(Intercept)	3.64	0.19	18.90	<.001 ***
direct effects	Explicitness	-0.16	0.20	-0.80	.43
+	Consecutiveness	0.08	0.21	0.38	.71
interaction	Gender	0.13	0.15	0.88	.38
	Education	0.04	0.15	0.26	.80
	Age	0.08	0.15	0.54	.59
	Explicitness:Consecutiveness	0.07	0.30	0.25	.80

A2.2 Study 2

Table 40: Regression Results for Topic Comprehension. *p <.05, **p<.01,***p<.001.

		b	Std. Error	z	Odds Ratio	p
MODEL 1	(Intercept)	3.08	0.92	3.34	21.74	<.001***
direct effects	Explicitness	0.69	0.81	0.85	2	.40
	Consecutiveness	-0.39	0.82	-0.48	0.68	.63
	Gender	-1.13	0.89	-1.27	0.32	.20
	Education	0.17	0.91	0.19	1.19	.85
	Age	1.24	1.13	1.10	3.46	.27
MODEL 2	(Intercept)	3.64	1.20	3.05	38.24	.002**
direct effects	Explicitness	-0.32	1.27	-0.25	0.72	.80
+	Consecutiveness	-1.27	1.23	-1.04	0.28	.30
interaction	Gender	-1.16	0.90	-1.30	0.31	.20
	Education	0.24	0.93	0.26	1.28	.80
	Age	1.29	1.16	1.11	3.62	.27
	Explicitness:Consecutiveness	1.88	1.75	1.07	6.54	.28

Table 41: Regression Results for Perceived Comprehension (item1). *p <.05, **p<.01,***p<.001.

		b	Std. Error	z	p
MODEL 1	Explicitness	-0.59	0.42	-1.39	.16
direct effects	Consecutiveness	0.24	0.42	0.58	.56
	Gender	0.18	0.42	0.44	.66
	Education	0.30	0.46	0.67	.51
	Age	0.04	0.44	0.09	.93
MODEL 2	Explicitness	-0.69	0.58	-1.19	.23
direct effects	Consecutiveness	0.11	0.66	0.17	.87
+	Gender	0.19	0.42	0.45	.65
interaction	Education	0.30	0.46	0.66	.51
	Age	0.04	0.44	0.09	.93
	Explicitness:Consecutiveness	0.21	0.83	0.26	.80

Table 42: Regression Results for Perceived Comprehension (item2). *p <.05, **p<.01,***p<.001.

		b	Std. Error	z	p
MODEL1	Explicitness	-0.06	0.39	-0.14	.89
direct effects	Consecutiveness	0.20	0.39	0.52	.61
	Gender	0.15	0.38	0.39	.70
	Education	0.04	0.44	0.10	.92
	Age	-0.19	0.42	-0.45	.65
MODEL 2	Explicitness	0.14	0.55	0.25	.80

		b	Std. Error	z	p
direct effects	Consecutiveness	0.42	0.58	0.72	.47
+	Gender	0.15	0.39	0.38	.70
interaction	Education	0.03	0.44	0.08	.94
	Age	-0.20	0.42	-0.47	.63
	Explicitness:Consecutiveness	-0.38	0.77	-0.50	.62

Table 43: Regression Results for Clarity. *p <.05, **p<.01,***p<.001.

		b	Std. Error	t	p
MODEL 1	(Intercept)	3.83	0.12	31.64	<2e-16 ***
direct effects	Explicitness	-0.10	0.12	-0.86	.39
	Consecutiveness	-0.01	0.12	-0.07	.95
	Gender	0.08	0.12	0.70	.48
	Education	0.16	0.13	1.26	.21
	Age	-0.11	0.13	-0.90	.37
MODEL 2	(Intercept)	3.85	0.14	27.63	<.001***
direct effects	Explicitness	-0.14	0.17	-0.85	.40
+	Consecutiveness	-0.05	0.18	-0.30	.77
interaction	Gender	0.08	0.12	0.70	.48
	Education	0.16	0.13	1.25	.21
	Age	-0.12	0.13	-0.91	.37
	Explicitness:Consecutiveness	0.08	0.23	0.34	.73

Table 44: Regression Results Sense of Control. *p <.05, **p<.01,***p<.001. a = The model in which only gender was specified as predictor (Model1.1) was not better than chance alone in explaining the variability in data at the Likelihood ratio test (X2(1) = 3.73, p = .05).

		b	Std. Error	z	p
MODEL 1	Explicitness	-0.34	0.33	-1.02	.31
direct effects	Consecutiveness	-0.34	0.34	-1.01	.32
	Gender	-0.75	0.35	-2.16	.03*[a]
	Education	-0.46	0.37	-1.23	.22
	Age	0.19	0.36	0.54	.59
MODEL 1.1 effect of gender	Gender	-0.63	0.33	-1.92	.06
MODEL 2	Explicitness	-0.72	0.48	-1.51	.13
direct effects	Consecutiveness	-0.75	0.50	-1.51	.13
+	Gender	-0.75	0.35	-2.15	.03*[a]
interaction	Education	-0.46	0.37	-1.23	.22
	Age	0.21	0.36	0.58	.56
	Explicitness:Consecutiveness	0.75	0.67	1.13	.26

Table 45: Regression Results for Response. *p <.05, **p<.01,***p<.001.

		b	Std. Error	z	Odds ratio	p
MODEL 1	(Intercept)	8.79	3.49	2.52	6566.27	.01*
direct effects	Explicitness	0.15	0.75	0.20	1.16	.84
	Consecutiveness	0.42	0.79	0.53	1.52	.60
	Gender	-0.44	0.80	-0.55	0.64	.58
	Education	0.002	0.91	0.002	1	1
	Age	-0.41	0.78	-0.52	0.66	.60
	Privacy concerns	-1.51	0.80	-1.88	0.22	.06

		b	Std. Error	z	Odds ratio	p
MODEL 2	(Intercept)	9.41	3.64	2.58	12201.51	.001**
direct effects	Explicitness	-0.40	0.98	-0.41	0.67	.68
+	Consecutiveness	-0.28	1.08	-0.26	0.76	.80
interaction	Gender	-0.44	0.81	-0.54	0.65	.59
	Education	-0.01	0.92	-0.01	0.99	.99
	Age	-0.44	0.79	-0.56	0.64	.58
	Privacy concerns	-1.58	0.82	-1.93	0.21	.05
	Explicitness:Consecutiveness	1.45	1.62	0.90	4.25	.37
MODEL 3	(Intercept)	8.86	3.28	2.70	7012.23	.01**
effect of privacy concerns	Privacy Concerns	-1.56	0.79	-1.97	0.21	.05
MODEL 4	(Intercept)	4.14	2.27	1.82	62.61	.07
effect of perceived C1	Perceived C1	-0.45	0.60	-0.74	0.64	.46

A2.3 Study 3

Table 46: Regression Results for Topic Comprehension (generic). *p <.05, **p<.01,***p<.001.

		b	Std. Error	z	Odds ratio	p
MODEL 1	(Intercept)	1.15	0.61	1.91	3.17	.06
direct effects	Explicitness	-0.66	0.45	-1.48	0.52	.14
	Consecutiveness	-0.09	0.45	-0.21	0.91	.84
	Gender	0.58	0.45	1.27	1.78	.20
	Education	-0.43	0.46	-0.94	0.65	.35
	Age	-0.90	0.50	-1.78	0.41	.08
MODEL 2	(Intercept)	0.79	0.67	1.19	2.21	.23
direct effects	Explicitness	-0.06	0.66	-0.09	0.94	.93
+	Consecutiveness	0.43	0.61	0.70	1.53	.49
interaction	Gender	0.55	0.46	1.20	1.73	.23
	Education	-0.40	0.46	-0.85	0.67	.39
	Age	-0.80	0.51	-1.57	0.45	.12
	Explicitness:Consecutiveness	-1.12	0.90	-1.24	0.33	.21

Table 47: Regression Results for Time to Respond. *p <.05, **p<.01,***p<.001.

		b	Std. Error	t	p
MODEL 1 direct effects	(Intercept)	3.31	2.87	1.15	.26
	Explicitness	3.03	2.16	1.40	.16
	Consecutiveness	0.59	2.17	0.27	.79
	Gender	-0.01	2.10	-0.01	1
	Education	4.28	2.17	1.97	.05
	Age	1.80	2.20	0.82	.42
MODEL 2 direct effects + interaction	(Intercept)	1.57	3.30	0.47	.64
	Explicitness	5.81	3.40	1.71	.09
	Consecutiveness	2.94	3.09	0.95	.35
	Gender	-0.29	2.12	-0.14	.89
	Education	4.37	2.17	2.02	.05
	Age	2.14	2.22	0.96	.34
	Explicitness:Consecutiveness	-4.67	4.41	-1.06	.30