The Cognitive Heuristics Behind Disclosure Decisions

Vincent Marmion, Felicity Bishop,
David E. Millard, and Sarah V. Stevenage

University of Southampton, Highcliff Campus,
Southampton. SO17 1BJ, UK

Abstract. *Despite regulatory efforts to protect personal data online, users knowingly consent to disclose more personal data than they intend, and they are also prone to disclose more than they know. We consider that a reliance on cognitive heuristics is key to explaining these aspects of users' disclosure decisions. Also, that the cues underpinning these heuristics can be exploited by organisations seeking to extract more data than is required. Through the lens of an existing credibility heuristic framework, we qualitatively analyse 23, one-to-one, semi-structured interviews. We identify six super-ordinate classes of heuristics that users rely upon during disclosures: PROMINENCE, NETWORK, RELIABILITY, ACCORDANCE, NARRATIVE, MODALITY, and a seventh non-heuristics TRADE class. Our results suggest that regulatory efforts seeking to increase the autonomy of the informed user are inapt. Instead the key to supporting users during disclosure decisions could be to positively nudge users through the cues underpinning these simple heuristics.*

Keywords: Cognitive Heuristics; Privacy Paradox; Informed Consent

1 Introduction

Disclosing personal data is self-managed. It is users that decide whether to consent to disclosure requests or whether to withhold their data. This consent-based model aligns appealingly with the ideals of information self-determinism [48]. However, in practice these ideals are not being met. Although regulations such as the UK Data Protection Act (1998) stipulate that organisations inform users of the operation of data processing, the explanation of risk is left to organisational discretion, making it incumbent on users to make the necessary risk calculation [40]. If we adopt the long-standing definition regarding consent in a medical context [49], the reality of consent in disclosure is akin to simple rather than informed consent. Simple consent involves a brief explanation of operations, followed by a trust-based agreement or refusal; to elevate this to informed consent, a detailed discussion of risks is also required. Regardless of likelihood, high risk necessitates informed consent. Because the consequences of data misuse are increasingly high and decreasingly rare [18], simple consent is unsatisfactory.

Unfortunately, increasing the autonomy of informed users may not be sufficient. This is because disclosure decisions are inherently uncertain, and when data is stored indefinitely there is no means of accounting for future uses or future capabilities, so the 'data controller' may also be ignorant of the risk [47]. Even if organisations were able to explain the risk, users rarely read privacy policies [45], and when there is an attempt, they lack the time and/or the capacity for the required uncertainty calculus ('privacy calculus') to comprehend them [13, 27]. Users are essentially left to trust that the data controller will behave in an expected and innocuous manner [38, 33], and to make heuristic judgements (i.e. using 'rules of thumb') about whether to disclose [44].

We share the view of [41, 14] in considering cognitive heuristics as key to understanding these decisions, and illuminating the problem of the privacy paradox - the tendency of user's to disclose more in their actual behaviour than in their previously stated intention [37].

In this paper we present a qualitative study to understand the heuristics that people use when making disclosure decisions. Using established credibility heuristics as a starting point we present an analysis of 23 semi-structured interviews, with the aim of exploring whether the heuristics related to credibility judgements are a general enough framework to also apply to disclosure decisions. We also seek to identify superordinate classes of similarly themed heuristics, and to explore the importance and limitations of heuristics within those classes, and therefore within the disclosure decision process as whole.

2 Background

When faced with difficult, uncertain, or intractable problems, using heuristics can be rational [17] as they fit with observations of decision-making *in the wild.* Nonetheless, they are prone to misjudgement and bias [44] and with that are prone to manipulation, or 'nudges'. While these nudges might be used to limit over-disclosure, they can also be used to encourage users to disclose more than they might otherwise be comfortable with [2, 19]. In this section we will look at exactly what is meant by heuristics, how they have been linked to credibility, and used when exploring disclosure decisions.

2.1 Heuristics and Credibility

Heuristics are used to reduce difficult decisions to solvable simple decisions [15]. For example, whether to invest in company A or B is a difficult decision, which depends on many complex factors. Whereas, an associated heuristic may be that size is related to success. Then the heuristic's decision variable, i.e., the cue, could be the number of service users or stock price. So, whether $A(cue) > B(cue)$ substitutes for whether A is a better investment than B.

Early research around activities such as *phishing* and *fake news* showed the impact that online cues have on users' trust of web sites, and their judgements of the credibility and legitimacy of those sites [12, 42]. Fogg [12] finds that rather than users seeking a particular cue, prominent cues affect users, and that with new digital interfaces comes changes in the prominent cues (e.g., the padlock

Table 1. Heuristic Approaches to Credibility Evaluation Online

Heuristic	Description
Authority	An official or primary authority
Recognition	Familiarity, even in name only
Reputation*	A prestigious service would not knowingly be wrong
Endorsement*	Recommendation from known others
Bandwagon	Recommendations or perceived actions of unknown others
Consistency*	Agreement with another source or procedure
Consensus	Agreement between many sources or procedures
Self-Confirmation*	Alignment with a pre-existing belief
Coolness	New modalities of the technology
Novelty	New encounters with the technology
Expectancy Violation*	Inferior site design, errors, poor visual appearance
Persuasive Intent*	A feeling of bias or being pushed

Note: Twelve (* six prominent) credibility heuristics collated from [34].

shown on a search bar). Understanding how to manipulate these effects is an active research agenda for behavioural economists [8, 2, 6].

The MAIN model [42] structures ten years of psychological research into the cues that have been empirically shown to affect users. This model assembles the cues in terms of four technological affordances of digital media: Modality, Agency, Interactivity, and Navigability. The result is an extensive array of cues and associated heuristics, providing a framework for more applied research. Drawing from this model, Metzger et al. [35] conducted 11 focus groups, and found five prominent credibility heuristics: *Reputation, Endorsement, Consistency, Expectancy Violation and Persuasive Intent*; later adding *Self-confirmation* as a sixth [34]. For instance, determining the credibility of a website is a difficult decision. An associated heuristic may be the Expectancy Violation heuristic that illegitimate websites appear unprofessional. So, whether the website has spelling mistakes can substitute for whether the website is credible.

The value of [35, 34] is in representing an expansive array of cues and decision variables, as in the MAIN model, into something simple and coherent (as summarised in Table 1). In addition, this table contains the six additional heuristics that [34] discussed as relevant, yet omitted for their purposes[1]. A concern is that the original six may be overly fitted to credibility judgements.

2.2 The Role of Heuristics in Disclosure

Credibility and disclosure decisions are both trust-based decisions [33, 41], and heuristics are somewhat abstracted from the *weeds* of a problem [26]. In fact, is it predicted that in some circumstances one simple heuristic can take an individual through a cycle of disclosure through related yet independent decisions, flowing

[1] Metzger et al., found six heuristics through a process of reduction, i.e., Recognition is subsumed under Reputation, as to perceive reputation involves a prior recognition.

from a credibility judgement regarding the legitimacy of a service, through a judgement to determine a service's trustworthiness as a data controller, and finally to whether the individual is willing to disclose a particular item of data.

In similar work, [14] used grounded theory, on data from eight focus groups, to reveal eight heuristics that underpin disclosure; four of which promote disclosure (Gatekeeping, Safety-net, Bubble and Ephemerality), and four that inhibit it (Fuzzy-boundary, Intrusiveness, Uncertainty, Mobility). Where the work of Metzger et al. provided the motivation for this work, the findings here have been refactored in light of these eight newly revealed heuristics.

There is a question however, as to the value of producing an ever-expansive set of heuristics separated only by subtleties. For instance the Intrusiveness heuristic, i.e. unsolicited communications inhibiting disclosure willingness from [14], shares similarities with the Persuasive Intent heuristic i.e. a feeling of bias or being pushed that inhibits the willingness to disclose [34]. A balance is required to avoid returning to the extensive set of decision cues such as in the MAIN model, and remaining in the scope of the simple findings bespoke to credibility. Hence our focus on developing superordinate classes of similarly themed heuristics that might accommodate emerging work.

3 Method

A series of semi-structured, one-to-one, face-to-face interviews were conducted. This approach was consistent with the qualitative nature of the focus groups in [35] and [14], but had an advantage of providing for a deeper focus on individual experience. The interview structure followed that of a cognitive walk-through, this was chosen as it is productively used for heuristic evaluations within human computer interaction (HCI) studies [36, 22]. However, instead of a specific target system as in HCI studies, the interviewee were asked to recall an interaction with an online service. This meant that no one system had to be contrived, and also the focus on what had already occurred avoided talk of ideals that misalign with actual behaviour (i.e., the privacy paradox) [37].

Furthermore, focusing on interviewee interpretation (the heuristic) means that the actual system, and the content of the cues was of less importance than the type of cues. For instance, one person may look for a NUS seal of approval, whereas another looks for a Royal Warrant mark, either way they can both be using the Endorsement heuristic. This meant that the required sample were those who regularly engaged with online actives and services, and who could also reason and articulate about these engagements. With this in mind we recruited from the Psychology department's participant pool, resulting in 23 Interviewees, all of whom were 18 to 25 years old.

The interviews were all under an hour (43 to 57 minutes), and were conducted on campus. Interviewees were briefed, and given the opportunity for questions before signing a consent form, the audio recording then commenced for the duration of the interview, and at the end they were debriefed. Participants were paid a small amount of compensation for their time. Finally, the audio files were transcribed verbatim by the interviewer, and imported into NVivo for analysis.

The interviews comprised three stages. Stage one (approx. 5 minutes), involved simple questions that established the interviewee's general digital engagement, whilst also easing them into the interview process.

In stage two (approx. 40 minutes), each interviewee was asked to recall a recent instance whereby they registered with an online service. Then, interviewees were primed to think within one of two possible self-regulatory mindsets [21]; twelve were *promotion primed* to recall a system relating to social activities, entertainment or freebies (in the analysis these are denoted with S, for social, i.e., S1, S2,...S12), and eleven were *prevention primed* to recall a system relating to responsibilities or financial and commercial transactions (denoted with T, for transaction, i.e., T1, T2,...T11).

Once a relevant situation had been recalled they were asked to discuss the process from first considering the service, through to completing the registration or transaction. They were allowed to speak freely around the task, however, when recall became disjointed a prompt sheet of short questions that moved chronologically through the process was referred to. While the interviewer steered the conversation around the contexts of the original primer, some contextual cross-over was unavoidable, and in many cases these instances were insightful.

Finally, stage three (approx. 5 minutes), included questions regarding the general concept of identity and privacy in the media. As well as a winding down exercise, this section allowed the interviewees to express any related thoughts or concerns that may have occurred during the interview.

3.1 Analysis

The analysis was undertaken in three parts.

The first part was ensuring familiarity with the data. The first author conducted the interviews, and transcribed the audio recordings, which ensured a base level of familiarity with the data. The data was then divided into; A) the data relating to disclosure decisions, and B) the data not relating to disclosure decisions. Before being set aside, data set B was examined to provide context and validity as to the interviewees' suitability for, and engagement with the process.

The second part involved coding the interview data into distinct categories. To do this, the transcripts were examined using an interpretivist approach [23, 10], wherein the interviewee constructs the theoretical connection between cue and decision, the analyst is then left to categorise the self-reported 'rules of thumb' [42]. An analysis challenge stems from people combining heuristics or interweaving heuristic and non-heuristic interpretations to inform decisions [16].

To help address this data set A was first categorised into; A1) heuristic-based, and A2) non-heuristic decisions. Then by matching the language used and the cues mentioned with those outlined in the literature, data set A1 were deductively categorised to align with the heuristics in Table 1 as described in [34]. Data not aligned with Table 1 were then inductively coded, as described in [39], to reveal additional heuristics. Data set A2 provided a set of non-heuristic decisions, that coded inductively acted as deviant cases to counterbalance any confirmatory-bias residing in the efforts to explore heuristics.

While part two teased the results apart, part three of the analysis serves to recombine them by identifying super-ordinate classes that encapsulate the results from part two. For example, while the Reputation and Recognition heuristics involve different interpretations and reasoning, they seem to be based upon similar underlying cues (for instance the size of the organisation) therefore, they are discussed within a parent *prominence* class. This allows the extension of the original credibility framework whilst maintaining its richness, and also helps to keep the overall results concise and pragmatic.

4 Interviewee Context and Validity

Interviewees reported habitual engagement with digital living, with typical comments such as; “Oh, literally, all the time”, and “when I wake in the morning, I just look in bed on Facebook.” However, we remain mindful of the temporal nature of the responses [28], as summed up by interviewee T6 when they said “I was 15, now I’m 19, so I have different interests.”

The Interviewees were open and candid within the interviews, this was exemplified by a common revelation about having only two or three passwords across all online systems, a finding that mirrors those of [29]. For example, one participant admitted “when I am creating a password they say you need a capital or a punctuation so it might vary, but generally I use one of three.” In some cases the participants even admitted to writing them down for ease, “I just created a word document to remember all of my passwords.”

While the similar age and educational status of the participants naturally scopes our findings to a particular demographic, it is a key demographic for the problem of disclosure. Our analysis of the first stage of the interview shows that our participants were engaged with the problem of disclosure in their everyday lives, and prepared to give rich answers to the interview questions.

5 Findings

Table 2 summarises the results of the second and third parts of our analysis (coding, and identifying super-ordinate classes), and contains sample dialogue to illustrate the coding process. Six super-ordinate classes are discussed; PROMINENCE, NETWORK, RELIABILITY, ACCORDANCE, NARRATIVE, and MODALITY. A seventh non-heuristics TRADE class was also identified. Within the following sections, the interview extracts address the flow of decisions through the cycle of disclosure, leading an interviewee from an initial assessment of service legitimacy, through the assessment of the service as a trustworthy data controller, and finally to the assessment as to whether the interviewee actually disclosed an item of PII.

Table 2. Superordinate Classes and Heuristic Coding Reference

CLASS	**Heuristic:** Description	Example Extract -*Interviewee*
PROMINENCE	**Reputation:** Prestigious services would not knowingly do wrong	"FIFA is well-known and probably not evil."$_{S3}$
	Recognition: A familiarity with a service, even in name only	"it is just a very small app that I have not heard much about, I think I wouldn't put my information on it."$_{S4}$
NETWORK	**Endorsement:** A recommendation from known others	"My brother has been telling me it is more secure, it is easier, better and safer."$_{S1}$
	Authority: A recommendation from official or primary authority	"He was a journalist, so he knows a lot of those sort of things."$_{T7}$
	Bandwagon: Perceiving the actions of unknown peers or general population	"I was quite influenced by what everyone was doing."$_{T7}$
RELIABILITY	**Consistency:** Interacting with a familiar process	"I tried another website, and also they ask for the same thing. The same questions. You have to sign up first, and it was the same thing. So I signed up."$_{S1}$
	Consensus: A normative or standardised process	"Just the normal, name, date of birth and the important one is the mobile number to create an account."$_{S10}$
	Expectancy:[1] Inferior site design, errors, poor visual appearance	"this looked fashionable and genuine."$_{S6}$
ACCORDANCE	**Intent:**[2] A feeling of bias or being pushed	"they wanted all my details to tell me how much it would cost. So I provided false details."$_{T3}$
	Self-confirmation: Feeling a consistency with pre-existing beliefs	"Why do you need ID? I'm only buying make-up."$_{T11}$
MODALITY	**Coolness:** Gratifying features of a technology	I like the effect on the photos, I only did it for that, I don't like the privacy really, but I don't really use it very often, it is just on there in case."$_{S6}$
	Novelty: An new encounter with a technology	"When I actually started, I was so happy about it, that I completed absolutely everything."$_{S3}$
NARRATIVE[3]	**Availability:** The ability to recall similar instances	"there has never been a dodgy situation when I don't want to give [my location data] because it is harmless games like Flappy Bird"$_{S3}$
	Coherence: The ability to envisage the result of an action	"if someone hacked my Twitter account I honestly wouldn't care because it is utter rubbish, nothing important, it is only entertainment"$_{S2}$

1:Violation aspect removed from the Expectancy heuristic to provide neutral label, [See Section 5.3]

2: Persuasive aspect removed from the Intent heuristic to provide neutral label. The Intent heuristic also incorporate the Intrusiveness heuristic from [14] [See Section 5.4]

3: Not part of the original Credibility heuristics framework as described in Table 1

5.1 PROMINENCE: Recognition and Reputation

Many of the participants expressed terms aligning with the **Reputation** heuristic, by implying that a prestigious service would not knowingly do wrong, as described in [34]. When asked for location data, Interviewee S3 says; "I mean FIFA is well-known and probably not evil." This is the first example of a simple heuristic, reputation, being sufficient to complete the cycle of disclosure. We know that reputation relates to credibility judgements regarding an organisation's legitimacy, also, within this exchange with S3 it also provided an implicit trustworthiness of FIFA as a data controller, and then specifically in regards to their willingness to disclose their location data.

Interpreting cues related to size, being low-key, or being a known brand were typical. Such as when Interviewee S5 suggests that "Twitter is such a big company you assume they would not [...] pass your information on." In a related tone, Interviewee T11 associates size and risk "because smaller companies don't have as many resources for security." Meaning that an organisation would want to protect their reputation, and thus protect the user, and the bigger the organisation the better the protection. There was an overall sense that if something has gained prominence then it must be doing something right, whereas lacking prominence suggests otherwise.

This sentiment is also reflected in the **Recognition** heuristic; trust occurring due to a basic familiarity with an entity [17, 35]. T6 makes the connection from the prominence of a high-street presence, and thus being "well-known", and them "not trying to scam me". This may seem similar to reputation, but there is value in the distinction. Reputation seems to involve other people, as in "well-known" compared to an inwards reflection, as in "I have not heard".

Perhaps the most notable difference between recognition and reputation, is that reputation extends beyond the original entity towards subsidiaries. Gambino et al. [14] refer to this as a safety-net heuristic, exemplified by Interviewee T10 when they state "with independent people, you need a barrier". This is a repeated factor in disclosure decisions, yet it is still a factor of reputation because reputation is acting as a form of collateral for the data exchange, as something "to live up to", T1. For instance, online organisations acting as a trust intermediary for other associated organisations, because the parent "company image is that valuable", S11. This protection is also inherited by other service users, as S4 finds it "really dodgy" being young and female on "a site that isn't well-known", yet on Twittter they "wouldn't worry too much".

Seemingly, being of prominence provides organisations with many cues interpreted towards trust and willingness to disclose. We have a scenario where credible organisations are "not trying to scam", which attracts users, which in turn adds to them being "so well-known and so big you can trust it", and this trustworthiness is reinforced by having the "resources for security" against outside threats, and being "probably not evil" to cause inside threats. This cocktail of credibility and trustworthiness leading to willing to disclosure is a prime example of the simplicity in the cycle of disclosure. However, from Interviewee T2's perspective, "it is not really about the reputation it is about the price", reminding us that although these are simple heuristics, the decision cues remain diverse, and moreover, that when finance is involved it changes the decision further.

5.2 NETWORK: Endorsement, Authority and Bandwagon

Evident from the interviews was that an individual's interpersonal network has a considerable influence on disclosure decisions. This was also reflected in [34] through the Endorsement and Bandwagon heuristics. These heuristics are similar to the heuristics in the Prominence class, the difference being that the Prominence class regards a service's place in the world, i.e., a high-street presence, whereas the network class has a personal characteristic, i.e., my friends were doing it.

Focusing first on the **Endorsement** heuristic; testimonial by known others, Interviewee S2 found "that two of my housemates were already there made it seem more comfortable". These findings mirror that of [35], suggesting that in some cases individuals prefer recommendations to their own decision. Interviewee S1 reflects delegation to others to make decision for them, when admitting; "I am more affected by what people tell me as I am not really an IT person". This type of sentiment, by an educated individual with habitual use of technology, runs counter to the notion that disclosure is to be self-managed.

The **Authority** heuristic is when trust stems from expert or official authority endorsements [42]. When Interviewee T7's father convinced T7 to allow electoral roll information to be traded, he was the authority but not as a parent, "he was a journalist, so he knows a lot of those sort of things". In effect it is an endorsement from an individual with reputation, but is not passive like the reputation heuristic. This feature was seldom present in the interviews.

Similarly, the **Bandwagon** heuristic involves recommendations and often shares decision cues with the Endorsement heuristic [35]. However, the findings here agree with [42], that the two are meaningfully different. Instead of a personal endorsement from friends and family, in the Bandwagon heuristic the recommendations can be from unknown others via less personal factors such as aggregated testimonials or star ratings embedded within the interface. This places the Bandwagon heuristic conceptually close to the Prominence class, illustrated by "many thousands of people have downloaded them they can't be that bad", S4. Yet, it is also has a socially compelling aspect to it, as T7 explains, "I wasn't 100% satisfied with the [privacy policy]" but, "everyone is doing it" or as T10 reflects "I thought everyone else was. I assumed you had to fill it in".

Throughout this Network class of heuristics, there is a degree of delegating the decisions within the cycle of disclosure, through direct council and endorsement or indirectly through the assumed behaviours of peers. There is a free-riding aspect, whereby there is an expectation of others doing the risk discovery [45]. But this is a self-fulfilling 'social proof' [9], whereby a herd mentality can follow without due consideration of the circumstances [5].

5.3 RELIABILITY: Expectancy, Consistency and Consensus

The **Expectancy Violation** heuristic has negative connotations surrounding poor design, central to which is an expectation of professionalism [34]. In this regard, on numerous occasions interviewees were cued by presentation details, with features such as poor layout, inferior design or errors impacting on perceptions of service trustworthiness. To this end T6 trusted their judgement "by looking at their website or social networking site, whether it looks professional or not", or for T7 it was that, "something in my mind saying it is not right". We have chosen to remove the 'violation' label and simply use Expectancy, because the cues can work both ways, in that "you sort of get the feeling that it is not right, but this looked fashionable and genuine" (S6), this relabelling reflects that writing inconsistency, non-consensus, or disreputable would be unproductive.

This Reliability class also contains the **Consistency** heuristic; trust based on the agreement between two independent sources [35]. When S1 says "I didn't want to sign up, so I tried another website, and also they ask for the same thing. The same questions. You have to sign up first, and it was the same thing. So I signed up.", we say they are using the Consistency heuristic. In this case, S1 expected to engage without registering, however, when it became apparent that the seemingly non-standard requirement for registration was a consistent requirement across similar TV services, the user became willing to disclose.

Similarly, the Reliability class includes the **Consensus** heuristic; a normalised and general agreement [42]. Consensus has a broader application than the Consistency heuristic. These situations are exemplified by the use of normative terms, such as Interviewee S4 noting that "obviously name, email address" and "obviously it wanted a photo" to describe an interaction with a social networking site. Likewise, Interview S10 with "Just the normal, name, date of birth", and T6 with "obviously name and email", however, the data within these normal and obvious requests did tend to differ. The result is that "it almost bypasses you because you have done it so many times, but if something unexpected came up like a page you have not seen before that would make you doubt it. [...]. If it is the same process as usual I would assume it was fine".

The three heuristics are linked by the idea that if something is broken, has mistakes, or if something changes, it can cue users against disclosure, whereas a professional and as-expected interaction goes unnoticed. Problematically, this manner of thinking could incentivise service providers to request more information than is currently required, because in waiting to do so at a later occasion, the service risks disturbing the user's sense of routine and invoke questions such as S6's when a TV service started asking for information; "why are you doing that? It used to be different. They didn't use to ask for details".

5.4 ACCORDANCE: Self-confirmation and Persuasion

When Interviewee S6 expressed that a TV service "didn't use to ask for details" they were disrupted as a factor of the consistency heuristic, however, when they then reflected on "why are you doing that" this is closer to the **Self-Confirmation** heuristic; when something aligns with one's prior belief [35]. S6 later went on to reflect on why the BBC iPlayer "need[s] to check up on me, and my full name? To see what I'm watching?". Then as a result, "I just put my initials in, because I'm just watching TV".

The Accordance class differs from the Reliability class, in that it refers to beliefs and understanding rather than process or interface. Also, the Self-confirmation heuristics does not require a norm to the request, as long as there is an understanding that the request "comes up for good reasons" such as a store requesting a delivery address from S3. Whereas, when asked for ID for a birthday promotion, T11 refrains because "that is not a good reason to give my ID, especially when I just want to buy make-up. I wasn't happy, so I didn't sign up". They could not justify the disclosure when told the reason, although in contrast T11 did give their ID to a storage company when told it was in case "something happened" despite being vague and not particularly compelling.

Also in the Accordance class is the **Persuasive Intent** heuristic, the underlying principle of which is that perceived manipulation leads to negative judgements [35]. For instance, pop-up messages have been shown to produce a negative psychological effect [11, 46]. Throughout the interviews, such instances related to unsettling aspects of an interaction being "too violent, in your face" or annoying features that "as you try and get a page and they are flashing up at you". Gambino et al. [14] recognise such aspects as being an *intrusiveness* heuristics, leading users to "question the integrity" of the service. Removing the 'persuasion' part of the name of this heuristic to leave it labelled simply as 'intent' serves the purpose of being close to the 'integrity' element in Gambino et al., whilst maintaining a neutral description. Intent better describes the grey area between it being "quite helpful if they have picked up on what you are trying to find" and "it seems to be everywhere, [...]. It is annoying and unnecessary. I suppose there is two sides to it".

We learn more about this class when Interviewee T7 implies that paying for prominence on a search result was something to be "wary about", as if it was not in the spirit of things, compared to those who achieve prominence through merit of popularity. Or when Interviewee T2 was deterred by an insurance company because "they wanted all my details to tell me how much it would cost". In this instance, T2 realised it was a 'consistent' process for insurance companies to request this, yet the feeling of being pushed meant they "provided false details".

5.5 MODALITY: Coolness and Novelty

Sundar [42] associates the **Coolness** heuristic with new technological features, or the bells and whistles of existing technologies, with positive credibility evaluations. For instance, Interviewee S6 consents [Instagram] access to all their photos, despite that they "don't particularly like to, but you can't download it without giving that permission", and they "like the effect on the photos".

The **Novelty** heuristic is subtly distinct from the Coolness heuristic as it is invoked by a user's initial experience with a technology [42]. S3 describes two instances of “when I first started Facebook I think I got a bit carried away”, and “when I actually started [Deviant Art] I was so happy about it, that I completed absolutely everything”. However, that early exuberance waned and “looking back on my profile I used to disclose more information than I do now”.

In Sundar’s MAIN model [42], these heuristics are seen as a factor of modality. Instances of these heuristics were sparse, and limited to social and entertainment situations, perhaps aligned with the explanation that in these instances individuals are mostly concerned with gains and immediate gratification [21, 1].

5.6 NARRATIVE: Availability and Coherence

The framework in Table 1 was insufficient to explain all of what the interviewees described. This is mainly because the credibility framework referred to individuals establishing trust, it does not account for individuals considering risk. Instead of “why are you doing that” type questions, interviewees would engage with past examples, and/or hypothetical situations, asking themselves “why would they be interested in me”, or more pertinent, “if I had been affected” type reflections.

To frame these instances of introspection, we refer to the description by [43, p. 15] of the **Availability** heuristic; a judgement of the likelihood of an event based on the ‘ease with which relevant instances come to mind’. A example is S4’s work insight meaning they “would not sign up for anything like [comparison websites], because I worked in insurance and basically if anyone put information on GoCompare it would come straight to us”.

Interviewee S4’s experience was not typical in relation to those less aware, such as Interviewee S3’s lack of risk availability in that “there has never been a dodgy situation when I don’t want to give [my location data] because it is harmless in games like Flappy Bird”. There were many similarities between the perspectives of the interviewees here and [7] as when the extent of data leakages were revealed to their participants they were ‘very surprised’ by the frequency and the destination of data leakage from mobile games.

The overriding difference is that in [7] the full extent of data disclosures as a result of playing a game was demonstrated, which in turn, allowed their participants to envisage a list of possible negative outcome, thus completing a disclosure narrative. In the end, these participants stated a desire to change future behaviours, and one participant even changed from perceiving disclosure as useful for customisation, to later referring to the game as being ‘slime’.

It is unsatisfactory to wait for users’ negative experiences to instil a more cautious, considered approach to disclosure. Instead, it may be possible to inform users of disclosure risk through a relatable narrative. In this regard we refer to a **Coherence** heuristic; being able to envisage the result of a decision as a plausible outcome. For example, S11 does not profess to having been mugged, yet they can reason that “when you post a picture you can add your location then people in the area can look at the picture and they can find you and they could mug you or something like that”. For S2, they can envisage that it is possible to hack a Twitter account, yet “honestly wouldn’t care because it is utter rubbish, nothing important, it is only entertainment”. Seemingly, the interviewee does not have the available recall or imagination to see the potential negative results, such as those increasingly experienced by victims of facility takeovers [24, 25].

Norberg et al. [37] finds that people abstractly perceive risks in over-disclosure, yet when faced with a specific disclosure decision they most likely disclose. The evidence in this study contributes to that observation, and further suggests that the often missing narrative could play a significant role. In practice however, due to the consent model being 'simple', such narrative is rarely available to the user, and therefore, the resulting behaviour is similar to that observed in Norberg et al. Furthermore, there is an expectation of sorts that this narrative will be brought to them as noted by T1 who "assume[s] that if I had been affected [...] I would be contacted by eBay. [...]. Only at that point if that happened would I care about it a bit more". Or T3 suggesting that "[i]f there was a serious problem I'm sure it would be in the news".

5.7 TRADE: Gains and Worth

Despite the primary focus on heuristics, we examined the data for deviant cases to counter some confirmation bias. From this, it was evident that along with heuristics, interviewees were also weighing up their disclosures in terms of trading utility gains versus losses [4, 3]. In many ways the Modality class (coolness and novelty) reflects the notion of a trade. Interviewee S3 considering that "it doesn't seem like a good investment" to disclose location data to a poorly designed game, seemingly interpreted in a manner associated with the Effort Heuristic [32] in that lack of effort reduces utility. Likewise, S6 explains that it was "quite a lot of details, but I felt like I was getting something back with the [rail] voucher".

Trade-type behaviours were often imbalanced in favour of disclosure. For instance, S9 perceived a lack of real option "[w]hen Google linked Gmail and YouTube [...] I didn't have much of a choice, I would have had to close my YouTube account, and I didn't want to do that", because "I didn't want to lose" my "personal videos" and "amateur stuff". Likewise, S11 described how "they force you to have it on your phone", with the sentiment that "I sort of need it. I have 100-200 friends on there that I need to contact". S11 also reflected on "why would they be interested in me?" Conveying a common sentiment of insignificance around personal data [20], hence apportioning a low overall value to the data disclosed, compared to a clear understanding of Facebook's utility.

Within efforts to disclose in a more calculative manner, the variables underpinning the decision often remain heuristics based. When S2 explains a differences when disclosing "name, phone number, that I need to give, that is fine. But I wouldn't tell them where I'm working or what I study", they rely on the Coherence heuristic when envisaging someone turning up at their work. Equally, S1 can envisage the risk and therefore caution "in terms of card and bank details. That will put me off subscribing or buying online", but this is only a relative value as "mobile, or equally email, is the least worrying compared to card".

Whilst in some instances there was a relative value to single identifiers, i.e., name vs. place of work, in other instances, disclosing a combination of identifiers impacted the valuation. Interviewee S4 talked of less willingness to disclose their age once gender had already been disclosed. Seemingly, being a female was a satisfactory disclosure, but not in conjunction with being young. Not evident is whether age in isolation holds the same value as when in combination with gender, or how this value may change over time. In contrast, from S9's perspective there is a threshold effect;"I consider my phone number a pretty private thing to begin with it, so if someone has it, they already probably know my name."

6 Conclusion

The interviews that we have conducted have provided a rich qualitative account of users interaction with the disclosure decision points of online systems. Looking at the simplicity of such decisions through the lens of credibility heuristics, we find that our prediction that the heuristics are also being used for disclosure decisions to be valid. Also, we find disclosure heuristics outside of the credibility framework, mainly the importance of narrative in how users make disclosure decisions. These results were then encapsulated within superordinate groups (Table 2), revealing PROMINENCE, NETWORK, RELIABILITY, ACCORDANCE, NARRATIVE, MODALITY and a seventh non-heuristics TRADE class.

The main implication of this paper is that the self-managed model whereby self-informed individuals are responsible for consenting or withholding personal data, is idealistic. The evidence here is that users tend to make impoverished decisions. They evaluate trustworthiness from heuristics formed of prominence and social networks, using decision cues such as popularity, brand exposure or word of mouth, resulting in somewhat of a herd mentality. Alternatively, users evaluate trustworthiness from heuristics formed of accordance with beliefs and a sense of reliability, using cues such as familiarity and regularity. However, this reasoning has inductive pitfalls based on the idea that if nothing negative occurred before as a result of a disclosure, then future like-for-like instances are deemed safe. Our results agree with [41], and [14] that a reliance on such cognitive heuristics is key to understanding users knowingly consenting to give more than intended (i.e., privacy paradox). But also, we find this key to understanding users consenting to give more than they know (i.e., simple consent).

Norberg [37] calls it a privacy paradox when describing how users base their disclosure intentions on risk, yet base their disclosure behaviours on trust. The result is behaviour that favours disclosure because in disclosure environments there are many trust based cues yet scarce information about the risks. Therefore, on occasions when users attempt a considered approach to disclosure, qualitative accounts of what may happen are not adequately portrayed, and users find it difficult to complete a coherent narrative which diminishes their ability to adequately conduct the 'privacy calculus' [31] required for informed consent.

Our new super-ordinate set of heuristics for disclosure is envisaged to allow future research into the heuristics, and also to provide a place for emerging heuristics. Our hope is that these heuristics, and the implication of their susceptibility to bias and manipulation [30], may one day be harnessed so users may benefit from some form of positive nudge and thus mediation of the risks [8, 6].

References

1. Acquisti, A.: Privacy and security of personal information. In: Economics of Information Security in Advances in Information Security, vol. 12, pp. 179–186. Springer (2004), http://link.springer.com/content/pdf/10.1007/1-4020-8090-5{_}14.pdf
2. Acquisti, A.: Nudging privacy: The behavioral economics of personal information. Digital Enlightenment Yearbook 2012 pp. 193–197 (2012)
3. Acquisti, A., Grossklags, J.: Privacy attitudes and privacy behavior. In: Economics of Information Security, pp. 1–15. Springer (2004), http://link.springer.com/content/pdf/10.1007/1-4020-8090-5{_}13.pdf
4. Acquisti, A., Grossklags, J.: What Can Behavioral Economics Teach Us about Privacy ? In: Digital Privacy, pp. 363–377. Auerbach Publications (dec 2007), http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.145.7609{\&}rep=rep1{\&}type=pdfhttp://www.crcnetbase.com/doi/abs/10.1201/9781420052183.ch18
5. Acquisti, A., John, L., Loewenstein, G.: The impact of relative standards on the propensity to disclose. Journal of Marketing Research 49(2), 160–174 (2012), http://journals.ama.org/doi/abs/10.1509/jmr.09.0215
6. Adjerid, I., Acquisti, A., Brandimarte, L., Loewenstein, G.: Sleights of privacy. In: Proceedings of the Ninth Symposium on Usable Privacy and Security - SOUPS '13. p. 1. ACM (2013), http://dl.acm.org/citation.cfm?id=2501604.2501613
7. Balebako, R., Jung, J., Lu, W., Cranor, L.F., Nguyen, C.: "Little brothers watching you". In: Proceedings of the Ninth Symposium on Usable Privacy and Security - SOUPS '13. p. 1. ACM (2013), http://dl.acm.org/citation.cfm?id=2501604.2501616{\%}5Cnhttp://www.scopus.com/inward/record.url?eid=2-s2.0-84883078013{\&}partnerID=tZOtx3y1
8. Balebako, R., Leon, P.G., Almuhimedi, H., Kelley, P.G., Mugan, J., Acquisti, A., Cranor, L.F., Sadeh, N.: Nudging users towards privacy on mobile devices. In: CEUR Workshop Proceedings. vol. 722, pp. 23–26 (2011)
9. Cialdini, R., Trost, M.: Social influence: Social norms, conformity and compliance. In: The Handbook of Social Psychology, Vol. 2, pp. 151–192 (1998), http://psycnet.apa.org/psycinfo/1998-07091-021
10. Fereday, J., Muir-Cochrane, E.: Demonstrating Rigor Using Thematic Analysis: A Hybrid Approach of Inductive and Deductive Coding and Theme Development. International Journal of Qualitative Methods 5(1), 80–92 (2006)
11. Fogg, B.J., Soohoo, C., Danielson, D.R., Marable, L., Stanford, J., Tauber, E.R.: How Do Users Evaluate the Credibility of Web Sites? A Study with Over 2,500 Participants. In: Proceedings of the 2003 conference on Designing for user experiences (DUX'03). pp. 1–15. ACM (2003), http://dl.acm.org/citation.cfm?id=997078.997097
12. Fogg, B.J.: Prominence-interpretation theory: Explaining how people assess credibility online. In: Conference on Human Factors in Computing Systems - Proceedings. pp. 722–723. ACM (2003), http://dl.acm.org/citation.cfm?id=765951{\%}5Cnhttp://www.scopus.com/inward/record.url?eid=2-s2.0-84869039673{\&}partnerID=40{\&}md5=f36a1afb8a3a649f12e97c7d6b38854a
13. Furnell, S., Phippen, A.: Online privacy: A matter of policy? Computer Fraud and Security 2012(8), 12–18 (2012), http://dx.doi.org/10.1016/S1361-3723(12)70083-0

14. Gambino, A., Kim, J., Sundar, S.S., Ge, J., Rosson, M.B.: User Disbelief in Privacy Paradox: Heuristics that determine Disclosure. In: Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems. pp. 2837–2843. ACM (2016)
15. Gigerenzer, G., Gaissmaier, W.: Heuristic decision making. Annual Review of Psychology 62, 451–482 (2011)
16. Gigerenzer, G., Hoffrage, U., Goldstein, D.G.: Fast and frugal heuristics are plausible models of cognition: reply to Dougherty, Franco-Watkins, and Thomas (2008). Psychological review 115(1), 230–239 (2008)
17. Gigerenzer, G., Todd, P.M.: Fast and frugal heuristics: The adaptive toolbox. In: Simple heuristics that make us smart, pp. 3–34. Oxford University Press (1999)
18. Goodman, M.: Future crimes: Everything is connected, everyone is vulnerable and what we can do about it. Anchor (2015)
19. Hansen, P.G., Jespersen, A.M.: Nudge and the Manipulation of Choice: A Framework for the Responsible Use of the Nudge Approach to Behaviour Change in Public Policy. European Journal of Risk Regulation 1, 3–28 (2013), `http://ssrn.com/abstract=2555337`
20. Heikkinen, A., Wickström, G., Leino-Kilpi, H.: Understanding Privacy in Occupational Health Services. Nursing Ethics 13(5), 515–530 (2006), `http://nej.sagepub.com/content/13/5/515.abstract`
21. Higgins, E.: Promotion and prevention. Regulatory focus as a motivational principle.pdf. Advances in Experimental Social Psychology 30, 1–46 (1998)
22. Hollingsed, T., Novick, D.G.: Usability inspection methods after 15 years of research and practice. In: Proceedings of the 25th annual ACM international conference on Design of communication. pp. 249–255. ACM (2007)
23. Holloway, I.: Basic concepts for qualitative research. Wiley-Blackwell (1997)
24. Hoofnagle, C.J.: Identity Theft: Making the Known Unknowns Known. Harvard Journal of Law & Technology 21, 98–122 (2007), `http://papers.ssrn.com/sol3/papers.cfm?abstract{\_}id=969441`
25. Kahn, C.M., Roberds, W.: Credit and identity theft. Journal of Monetary Economics 55(2), 251–264 (mar 2008), `http://linkinghub.elsevier.com/retrieve/pii/S0304393207001250`
26. Kahneman, D.: Thinking, fast and slow. Macmillan (2011)
27. Kehr, F., Wentzel, D., Mayer, P.: Rethinking the Privacy Calculus: On the Role of Dispositional Factors and Affect. The 34th International Conference on Information Systems (1), 1–10 (2013), `http://cocoa.ethz.ch/downloads/2013/11/1624{\_}kehr{\_}2013{\_}privacy{\_}icis.pdf`
28. Knijnenburg, B.P.: On The Dimensionality Of Information Disclosure Behavior in Social Networks. International Journal of Human-Computer Studies 71(12), 1144–1162 (2013)
29. Komanduri, S., Shay, R., Kelley, P.G., Mazurek, M.L., Bauer, L., Christin, N., Cranor, L.F., Egelman, S.: Of Passwords and People. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11). pp. 2595–2604. ACM (2011), `http://dl.acm.org/citation.cfm?doid=1978942.1979321`
30. Krasnova, H., Günther, O.: Privacy concerns and identity in online social networks. Identity in the Information Society 2(1), 39–63 (2009)
31. Krasnova, H., Spiekermann, S., Koroleva, K., Hildebrand, T.: Online social networks: Why we disclose. Journal of Information Technology 25(2), 109–125 (2010), `http://www.palgrave-journals.com/doifinder/10.1057/jit.2010.6`
32. Kruger, J., Wirtz, D., Van Boven, L., Altermatt, T.W.: The effort heuristic. Journal of Experimental Social Psychology 40(1), 91–98 (2004)

33. Metzger, M.J.: Privacy, Trust, and Disclosure: Exploring Barriers to Electronic Commerce. Journal of Computer-Mediated Communication 9(4), 1–29 (2006)
34. Metzger, M.J., Flanagin, A.J.: Credibility and trust of information in online environments: The use of cognitive heuristics. Journal of Pragmatics 59, 210–220 (2013), `http://www.sciencedirect.com/science/article/pii/S0378216613001768`
35. Metzger, M.J., Flanagin, A.J., Medders, R.B.: Social and Heuristics Approaches to Credibility Evaluation Online. Journal of Communication 60(3), 413–439 (2010)
36. Nielsen, J.: Usability inspection methods. In: Conference companion on Human factors in computing systems. pp. 413–414. ACM (1994)
37. Norberg, P.A., Horne, D.R., Horne, D.A.: The privacy paradox: Personal information disclosure intentions versus behaviors. Journal of Consumer Affairs 41(1), 100–126 (2007)
38. Olivero, N., Lunt, P.: Privacy versus willingness to disclose in e-commerce exchanges: The effect of risk awareness on the relative role of trust and control. Journal of Economic Psychology 25(2), 243–262 (2004)
39. Ryan, G.W., Bernard, H.R.: Techniques to Identify Themes. Field Methods 15(1), 85–109 (2003)
40. Solove, D.J.: Introduction: Privacy Self-Management and the Consent Dilemma. Harvard Law Review 126, 1880–1903 (2012), `http://papers.ssrn.com/abstract=2171018`
41. Sundar, S.S., Kang, H., Wu, M., Go, E., Zhang, B.: Unlocking the privacy paradox: do cognitive heuristics hold the key? CHI'13 Extended Abstracts on Human Factors in Computing Systems pp. 811–816 (2013)
42. Sundar, S.S.: The MAIN model: A heuristic approach to understanding technology effects on credibility. Digital Media, Youth, and Credibility pp. 73–100 (2008), `http://www.mitpressjournals.org/doi/abs/10.1162/dmal.9780262562324.073`
43. Tversky, a., Kahneman, D.: Availability: A Heuristic for Juudging Frequency and robability. Cognitive Psychology 5(2), 207–232 (1973), `http://www.sciencedirect.com/science/article/pii/0010028573900339`
44. Tversky, A., Kahneman, D.: Judgment under uncertainty: Heuristics and biases. In: Utility, probability, and human decision making, pp. 141–162. Springer (1975)
45. Vila, T., Greenstadt, R., Molnar, D.: Why we can't be bothered to read privacy policies models of privacy economics as a lemons market. In: Proceeding ICEC '03 Proceedings of the 5th International Conference on Electronic Commerce. pp. 403–407. ACM (2003), `http://dl.acm.org/citation.cfm?id=948057{\&}dl=ACM{\&}coll=DL{\&}CFID=304526782{\&}CFTOKEN=23143651`
46. Ward, R.: Physiological responses to different WEB page designs. International Journal of Human-Computer Studies 59(1-2), 199–212 (2003), `http://linkinghub.elsevier.com/retrieve/pii/S1071581903000193`
47. Weitzner, D.J., Abelson, H., Berners-Lee, T., Feigenbaum, J., Hendler, J., Sussman, G.J.: Information accountability. Communications of the ACM 51(6), 82–87 (2008), `http://dl.acm.org/ft{\_}gateway.cfm?id=1349043{\&}type=html`
48. Westin, A.F.: Social and political dimensions of privacy. Journal of social issues 59(2), 431–453 (2003)
49. Whitney, S., Mccullough, L.B.: A Typology of Shared Decision Making , Informed Consent , and Simple Consent. Annals of Internal Medicine 140(1), 54–59 (2004)