The Willingness of Crowds: Cohort Disclosure Preferences for Personally Identifying Information

Vincent Marmion, David E Millard, Enrico H Gerding, Sarah V Stevenage
Electronics and Computer Science, University of Southampton, Highfield Campus, SO17 1BJ
v.marmion@soton.ac.uk, dem@soton.ac.uk, eg@ecs.soton.ac.uk, S.V.Stevenage@soton.ac.uk

Abstract

Exploiting personal identifying information (PII) is critical for secure access to digital and web-based systems, it is also a significant element of the online social media business model. However, how this exploitation relates to users' valuation of their PII is poorly understood as an individual's willingness to disclose items of PII in different situations is unknown. For instance, an individual may delight in accessing their smartphone using facial recognition, yet they may hesitate when accessing banking services or vice versa. Moreover, the actual cost of disclosure gets obfuscated within dense and lengthy policies in a manner designed to exploit additional data. Thus, an individual may not understand that systems such as facial recognition can be a gateway to infer further PII.

Even with respectful intentions, identity-dependent technologies face a myriad of challenges to transparently balance users' sensitivities with their own need for high veracity PII. In a novel application of the ELO ranking algorithm, we detail a frugal and scalable method of capturing and combining some of these sensitivities. The design involves a set of 33 items of PII, and a cohort ($N = 115$) divided into three contexts: expression (35), transaction (40) and submission (40). The results indicate that while individuals may have many differences, as a cohort the personal utility of PII still collates and forms distinct clusters of PII that relate within and across contexts. This result means that technologies that treat PII as one amorphous group, and those transferring PII across contexts, risk failing to adhere to the sensitivities of the user. However, by working with these cohort-based clusters in mind, it is plausible that system designers and policymakers may better appropriate system needs with the wants of the individual.

Introduction

The remits of identity dependent technologies range from securing access to nations to delivering advertisements. These remits involve processing personal identifying information (PII) to reach the desired confidence that a person is as they claim or is whom they are believed to be (Beres et al. 2007). Ideally, these remits balance two requirements. On the one side, they need to exploit the identifying utility of PII, and on the other, they must respect the personal interests of users concerning privacy (Williams 2008). Unfortunately, efforts to meet these requirements are currently imbalanced. This imbalance is because, due to direct security and financial incentives (Clarke 1988; Black et al. 2012), exploiting PII's identifying utility is a clear and longstanding goal for both nation states and commercial organisations. Whereas, attempts at respecting the privacy interests of users lack cohesion and maturity (Nissenbaum 2004; Solove 2007), and have mostly indirect, hard to measure benefits such as goodwill (Toubiana et al. 2010).

Despite this imbalance, there is growing interest in personal privacy, with organisations recognising the market potential in addressing user concerns in the expanding Internet of Things (IoT) (Böhme, Koble, and Dresden 2007; Toubiana et al. 2010; Brunton and Nissenbaum 2015). Likewise, the recent EU general data protection act (GDPR) has added timely impetus to this agenda. So there remains urgency in steering systems designers, researchers, and policymakers towards identity-based technologies that can meet both requirements together (Cavoukian 2011).

We contribute to understanding the personal sensitivities of users when disclosing PII. Perhaps counter-intuitively, we do this using a cohort approach that aggregates the partial preferences of individuals. We use a cohort approach as our method is scalable and adaptable in terms of how the preference indication work is divided. Then, concluding the results, we return from that basis to show that this method can estimate individual preferences. To do this we explore a set of PII for indication of an underlying structure of personal utility, and we compare the results across three contexts. Also, in the method section we present our novel application of Elo's ranking algorithm (1978), and later show that this method not only ranks subjective data types, but that also reveals relative distances between the ranks.

The paper is structured as follows. The background section looks at some of the difficulties that veil the disclosure decisions that individuals face, along with some proposed technological solutions. The methodology section details a novel application of an existing algorithm. Then, the results section explores the aggregation of partial sets of individual preferences into complete sets of cohort preferences, and then back to more complete personal preferences. The results also examine the effect of changing the context on the results. A discussion section follows, before our conclusion.

Background

Research in the area of behavioural economics has focused on the monetary utility of PII from a personal perspective (Acquisti 2004), wherein individuals often attribute a low price in exchange for their PII (Grossklags, Hall, and Acquisti 2007). However, work regarding the privacy paradox repeatedly tells us that relying on actual behaviours, while initially seeming to be a sound approach, actually misrepresent the personal intentions and ideals of disclosure (Norberg, Horne, and Horne 2007). Because, in practice, disclosure ideals get outweighed by a combination of immediate, often small, gratifications (Acquisti 2004; Grossklags, Hall, and Acquisti 2007), perceived security benefits (Udo 2001), social norms (Zafeiropoulou et al. 2013), and other such heuristic-based decisions (Marmion et al. 2017a). Equally, a disclosure decision is often context dependent. Whether a user is in dealings with a government, in the process of self-expression or perhaps merely shopping, can affect their willingness to disclose (Van Zoonen and Turner 2014), it can also alter their perceptions of risk (Higgins 1998). Likewise, a disclosure is subject to interpersonal sensitivities and personality traits (Quercia et al. 2012). This subjectivity means that what one person considers sensitive and requiring of anonymity another may not, a finding that includes demographic inconsistencies such as older individuals being more sensitive about relational and taboo disclosures (Correa et al. 2015). Moreover, actual disclosures are complicated further by an information asymmetry between informed organisations that control data extraction and the users that poorly understand the extent, the purpose or even the identifying power of each, even innocuous, extraction (Preibusch, Krol, and Beresford 2013; Acquisti 2014; Perez, Musolesi, and Stringhini 2018). Finally, these decisions are also uncertain, as, once third parties are involved, the distribution paths can continually splinter as the technology evolves (Weitzner et al. 2006).

Despite the discord between users' intention, knowledge, and behaviours, due to the self-deterministic mechanism of consent, regulations place the consequences of disclosure with the discloser (Solove 2012). Also, as regulations do not distinguish online contexts in the way that users do, organisations such as data brokers often transfer PII outside of a user's intended boundaries (Tsesis 2014). With users readily disclosing, and retaining the risk, this serves only to embolden organisations in extracting and trading more and more PII (Vila, Greenstadt, and Molnar 2003; Tene and Polonetsky 2013b). This extraction may be for legitimate security purposes, such as Apple Inc. adopting fingerprint enabled access for its iPhone 5 in 2012. Alternatively, it may be for non-security purposes, such as Rovio Inc. extracting location data via the Angry Birds game. Either way, despite the hacking of the iPhone 5 on day one (Goodman 2015), fingerprint enabled access (and increasingly face recognition) is now a standard feature of modern phones, and despite unease at Rovio's tactics, it is one of a growing list of games and apps now extracting PII as a norm (Balebako et al. 2013).

Effectively, in vast numbers, users are either engaging without consideration, without understanding or with acclimatisation to PII requests (Tene and Polonetsky 2013a). Then, consciously, reluctantly or unwittingly, users engaging with one generation of technologies drive on the next generation (Flanagin, Flanagin, and Flanagin 2010; Das et al. 2015). Thus, with each generation of technologies, unwilling users to face a diminishing choice of non-disclosure or non-engagement (Staddon et al. 2012).

From a macro view, these actions combine towards situations whereby common systems introduce PII of increasing identifying utility, which in turn, normalises further everyday uses, a concerning and potentially unsustainable escalation. Subsequently, actual behaviours belie the intrinsic value of PII, so using observations to assess the privacy interests of users only favours those advantaged by disclosure. This situation motivates our investigation into the personal value of PII to individuals, not regarding what people disclose in a given circumstance, but on what they prefer to reveal.

It is from this ideals-based standpoint that technological solutions should position. For instance, future users may delegate the disclosure of PII to automated negotiation agents, whereby the agents can manage the multitude of permissions required in the IoT (Baarslag et al. 2017). One barrier to the success of such methods is the initial training period for preferences, yet we present a way to aggregate the partial preferences of individuals into a cohort result, then it is possible to start such an agent with this baseline training.

Methodology

We want to know users' perception of PII concerning personal disclosure preferences. We base our methodology on a participant experiment in the form of a pairwise comparison task. In each round participants choose between two items of PII within a particular context. After the completion of all comparisons, we then use a ranking algorithm to combine the individual comparisons into an ordered list for the group. The following sections give a detailed description of the experimental design and the process with which we collected and analysed the data.

Experimental Design

Pairwise Competition We use pairwise comparisons to mirror the process of a competitive match between two players. This feature is harnessed here, replacing opponents (or the players) with types of PII, and the win condition with a participant judgement. In each match, the participant chooses from two items of PII, e.g., work email and personal email, that which they are most willing to disclose in a given situation. The winner is the chosen item, and the other is the loser, the format does not permit a draw. Then, if there exists a preference in the judge's willingness to disclose then, some items will win significantly more than others, and therefore, it is possible to produce a ranking on a set of PII. Then repeat the process with a cohort of judges, and it is possible to produce an aggregated ranking.

Set of PII (The Players) For the 'players', we adopted a set of 33 items of PII (Table 1) from a report on personal data by Rose and Kalapesi (2012). This set aims for familiarity

Table 1: List of 33 Personal Data Identifiers.

Knowledge and Token				Biometric	
Biographic	*Demographic*		*Knowledge*	*Physical*	*Behavioural*
First Name	First pet	Age	Password	Fingerprint	Signature
Surname	Postcode	Education	PIN code	DNA	Geo-location
DoB	First School	Nationality	Username	Iris Image	Swipe Dynamics
Work Email	Phone Number	Home Address	Favourite Colour	Face Image	Keystroke Dynamics
Personal Email	Ethnicity		Account Number	Hand Pattern	Voice Signal
Mother's Maiden Name	Sex			Ear Pattern	

Adopted from a "Rethinking personal data" report (Rose and Kalapesi 2012), and adapted through colleague discussion. DNA and favourite colour provide sense test bounds to the set. One would expect these to be outliers in the results. The categorisation of these identifiers is for illustration only; the participants did not receive this information.

and a balance of knowledge, token and biometric identifiers. Future work would benefit from extending this list, for now, this work focused on the method, and what the results can tell us about the characteristics of PII's personal utility.

Experiment Scope Assuming that the order of presentation is not important, i.e., Postcode vs Fingerprint = Fingerprint vs Postcode, the 33 items of PII produce 538 possible combinations. During an informal pilot, individuals made 50 pairwise comparisons, taking 5-8 minutes, without any apparent attention fatigue. We decided to present the pairs randomly to participants as this simple design feature is frugal and scales consistently. The random presentation means that adding extra items of PII to be judged or adding participants to do the judging does not influence the experience of the participants. The point of distributing the work is so that the individual does not even need to see the entire list to contribute to the list in the same way a chess player need not play an opponent yet have a similar ranking.

Due to the work distribution, the number of judges is less critical than each item of PII having ample opportunity to play. For the algorithm detailed in the next section, a heuristic of fewer than 30 pairings implies a provisional player (Coulom 2007). A mock simulation of this random process determined that each of the three cohorts of 40 participants each making 50 pairwise comparisons ensure over 75 pairing for each item. The randomness does, however, impose variation in the pairings. Table 2 provides an overview of the actual outcome. For instance, Cohort 1, had fewer than 40 participants after removals meaning one item only played 73 matches, while another played 120.

Table 2: PII Pair Presentations

	S	**C**	**G**
Total Pairs Presented	1750	2000	2000
Min Presented Per Item	73	101	99
Max Presented Per Item	120	149	151
Range	47	48	52
SD	9.64	12.04	11.23

Contextual Variety To abide by evidence of the contextual nature of disclosure (Knijnenburg and Kobsa 2013), the ecological validity of this enquiry was helped by situating the task across three context, these are a social (networking) context, to account for activities involving 'expression', a commercial context (banking), to account for 'transactional' disclosure, and a government (passport) context that aligns with 'submission' based disclosure (Van Zoonen and Turner 2014). In a between-subjects design, random allocation placed participants in one of three distinct cohorts (see Table 3 for details of the scenarios presented to participants). Consequently, the outcome is three separate personal identifier ranks, one for each cohort, representing the relative value of PII within that particular scenario of the disclosure.

The Algorithm

Elo's (1978) ranking algorithm is the base mechanism to achieve the aim of quantifying willingness to disclose PII. The competition aspect of Elo's algorithm is mirrored in the process of pairwise comparisons (Sarma et al. 2010), and can be used more generally outside of sports to rank subjective-based entities, i.e., photographs (Coulom 2007). In our instance, these entities, or 'players', are items of PII.

The algorithm produces an expected value for whether player A will win against player B, and uses that value to update both players' scores depending on the actual outcome. For example, players A and B are opponents, each with an associated pre-game score, A_i, B_i. If A wins, then $A_{i+1} = A_i + x$, and $B_{i+1} = B_i - x$, i.e. it is a point exchange of the value of x. This way, incorrectly scored players can rise or fall through the ranks, while correctly scored items remain relatively stable. When repeated with many players and many matches an overall ranking stabilises with the best players on top. This process is achieved with equation (1),

$$EA_{i+1} = \frac{1}{1 + 10^{(B_i - A_i/400)}} \tag{1a}$$

$$EB_{i+1} = 1 - EA_{i+1} \tag{1b}$$

$$A_{i+1} = A_i + x \tag{1c}$$

$$x = K(A_i - EA_{i+1}) \tag{1d}$$

where, x is the number of points exchanged between A and B. The value of x is calculated using the relative magnitudes of A_i, B_i, to determine the expectancy of player A winning EA_{i+1}. Players with higher scores are expected to win, and

Table 3: Group Conditions and Tasks

S. Social Network	**C. Commercial Bank**	**G. Passport Office**
Your friends are discussing a new social networking site and suggest that you should join. You are keen and determined to join. So you visit the website and begin registration. However, this site is trying something different with their security.	Your bank calls to say they are upgrading their security methods. However, they are trying something different. You have been with the bank for a long time and wish to comply.	You have had your passport for ten years, and it now needs renewing. You need a passport so you begin the registration process. However, the Government are trying something different with their security.
Task: We take security very seriously and want to improve how we secure our service for your use. Instead of a single identifier, we intend to use a broader set of identifying information. However, we want to give you some control. Therefore, below is a pair of identifiers; what we want is for you to click on the identifier that is most appropriate for the security of this service. Think about this in terms of security and what you are comfortable giving to this service.		

an expected win yields a smaller x than an upset victory. Referred to as the K-factor, K is a constant that controls the maximum magnitude of x.

Algorithm Parameters To set the parameters, we followed Hvattum and Arntzen (2010). Each item of PII receives an initial score of 1500. Modifier values of 10 and 400, meaning that if $A_i - B_i = 400$, then the expected score of A winning is ten times that of B. These parameters have little impact other than to set the scale for the scores that underpin the final ranking. However, the K-factor has a more nuanced role. Typically this number ranges from 10 to 40 in general applications, with a higher number resulting in more sensitivity in score exchanges and also, therefore, in rank positions (Coulom 2007). This can be useful for situations where new players are entering stable ranks, for instance, the K-factor can be tuned, i.e., higher for novices, this way a new player may affect the ranking more initially while they close in on a real rank, then after a certain number of games, they are assigned the same K-factor as others. Likewise, a lower K-factor for 'master' players means that they do not affect the scoring as much when playing lesser opponents. In our application, there is no concept of novice and master, we, therefore, choose a constant 32 for our K-factor, and while this may be considered as high, we explain in the next section why this is less of an issue in our method.

Systematic Error An unexpected feature of ELO's algorithm is that the match order along with an overly sensitive K-factor can affect the final ranking. For example, using arbitrary results between five fictional players, and randomising the order in which we process the results produces different scores, then, from these scores differing ranks form (Figure 1). To eliminate this systematic error, we use a combination of Monte Carlo simulation and rolling scores. At T_0 all scores are equal (1500), then randomising the order of the participants and producing a score and rank, at $T_{i=1}$. Then, repeating the randomisation of participant order, and using the resulting scores from T_i as the base score for T_{i+1}, each item approaches its correct position (seen later in Figure 2). Each turn aims to produce less fluctuation until order sensitivity becomes insignificant. Adopting this technique also eliminates the hazard of setting the K-factor too high.

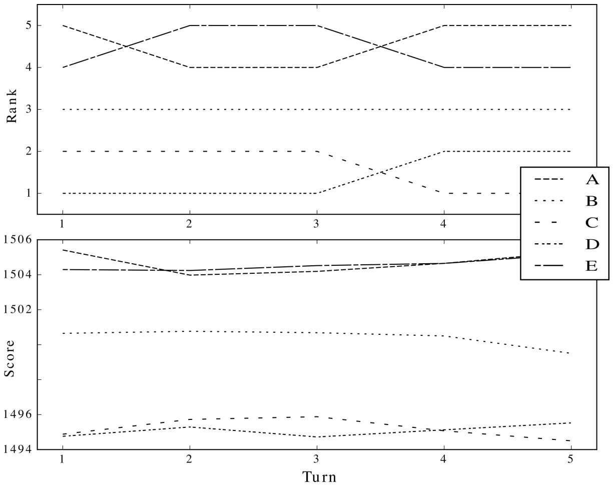

Figure 1: Using the ELO ranking algorithm on a random list of results between five mock players [A:E], produces slightly different rank [top] and scores [bottom] depending on the order that the results of passed through the algorithm. Each turn is independent.

Advantages of ELO Using Elo's algorithm is advantageous due to the way the scoring occurs. For instance, in Condorcet voting, e.g., Copeland's method, it is the total votes that indicate rank position (Saari 1999). Likewise, Likert type systems impose a rating scale on the individual. In Elo's method no scale is imposed and the number of 'votes' is not essential. Instead, it is who or what was in opposition that matters, the magnitude of a victory is primary. These magnitudes not only produce the ranking they also provide a relative value, meaning that we may also interpret any score separations as more meaningful scale. This method would be a worthwhile addition to other methodologies that aim to classify and order subjective data, especially those relying on human moderators rating all elements, such as examining the 'anonymity sensitivity' of social media posts (Correa et al. 2015). With our method, their participants would be able to cover a more massive repository of data without increasing workload.

Additionally, in using this ELO approach, not all items have to be paired an equal number of times or paired at all. Likewise, new items added to the original set will eventually find their place without the need to repeat the entire process. As PII is a continuously growing set (Black et al. 2012), the feature to add new items and or participant is a significant advantage. Later, we discuss the need for togetherness due to the interdependencies of PII disclosure, and it is from this perspective that we may perceive a snowball of engagement, with new information continuously added, and new participants joining the process, all without disrupting the core methodology. That sort of data set would be a powerful negotiating tool against a system that holds all the cards. An added advantage is that, given the differences in how people make disclosure decisions (Quercia et al. 2012; Van Zoonen and Turner 2014; Knijnenburg 2013), if we repeat the process with diverse cohorts, the results can be aggregated over time, and or kept separate. This flexibility means that groups that share the same ideals and motivation could form forcible coalitions.

Procedure

Participants responded to an online advertisement. They followed a hyperlink to an online survey system and completed the required consent to continue. During the consent process, we briefed the participants on the procedure. They then completed three questionnaires (as part of a study reported elsewhere) before following a link to a bespoke webpage for the pairwise comparison task. The site presented one scenario to each participant, then the first of the 50 identifier pairs in the form of two hyperlinks, the clicked link signified the winner, and the next pairing was displayed. The system randomly allocated the situational context from the three described in Table 3. Cohort (S) considered disclosure of information to sign up to a new social media site, cohort (C) considered disclosure of information to their bank, and cohort (G) considered disclosure of information to a government-based site for a passport renewal.

Participants

Differences could be introduced by variety in participant demographics, as described in (Knijnenburg 2013). This work focuses on the influence of context, yet for alternative questions, the context could instead be factors such as age, education or gender. It would just be a case of controlling the recruitment for these factors. However, these differences can become overly granular whereas we consider PII and identity assurance to be a collective matter (Fairfield and Engel 2015; Marmion et al. 2017b), so constraints may be of diminishing value. For our purposes, participants self-selected from within three UK Universities, based on an advert seeking those aged 18 to 26, of UK or Irish nationality. Table 4 summarises the recruitment and selection. Recruitment stopped when each of the three cohorts reached 40 participants, in total 125 participants took part (93f / 32m, mean age 20.21, SD 1.78). As part of a partner project, participants also completed three questionnaires, one of which, a two-factor regulatory mode (Kruglanski et al. 2000), included a 6-item lie index to assess the credibility of engagement. As a precaution, we excluded 10 participants from the results due to scoring over 1.5 standard deviations from the mean lie scale, or for leaving more than ten answers blank over the three questionnaires.

Table 4: Participant Recruitment Overview

		Age		
(Cohort) Context	**M - F**	$\bar{x}$	$_{min,max}$	**N/ N***
(S) Social Network	7 - 33	19.7	18 - 26	40 / **35**
(C) Bank	10 - 32	20.5	18 - 25	42 / **40**
(G) Passport Office	15 - 28	20.4	18 - 25	43 / **40**

N* is the participant numbers after eliminations

Results

The results comprise three sections. The first section examines the results solely from Cohort S; the PII requests from a social network service. However, these results typify each cohort. The second section of results will incorporate Cohorts C & G to provide a three-way concordance analysis. The third then shows how it is possible to use these cohort preferences to bootstrap the efforts of the individual.

Ranks and Clusters

Within Cohort S, there were a total of 1750 matches from 35 participant judges each making 50 pairwise comparisons. Figure 2 [top] shows the ranking of the 33 items of PII, while Figure 2 [bottom] are the corresponding scores. These images are for illustration of how the data settles. Initial observations show that predictably DNA sample and favourite colour occupy opposite ends of the scale concerning willingness to disclose; this is consistent for each cohort. Note that in practice these participants are potentially willing to disclose all of these items; this is a relative willingness. Higher scores indicate more willing to disclose. This expected coherence somewhat validates that participants cooperated by not randomly clicking through the pairwise comparisons.

As seen in Figure 1, Figure 2 [top] illustrates how ranking can fluctuate depending on the process order of the matches. The closeness in the scores in Figure 2 [bottom] are what are causing these rank fluctuations. Lowering the K-Factor would produce fewer fluctuations, however using Monte Carlo simulations with scores at T_i used for T_{i+1} these initial fluctuations dissipate, and although they persist, by T_{20} most positions have stabilised.

What is already apparent by T_{20} is that the scores have separated into clusters. Taking the results at T_{1000}, Figure 3 is an alternative view of these clusters as score box plots. Ordered by final ranking, this figure illustrates how these clusters form what appears to be step-levels of PII, with items such as age and sex in a cluster of PII relatively more inclined to be disclosed by participants, whilst face and iris image reside within another cluster of PII that the participants were collectively less willing to disclose. Due to the variation and overlap within these scores, it is prudent to consider rank and score together as rank alone does not tell the whole picture. For example, Ethnicity, Nationality, and

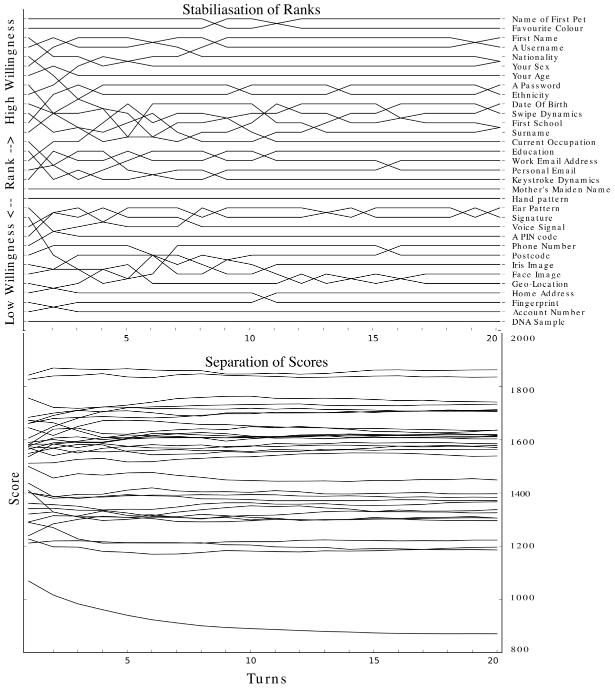

Figure 2: **Rank and Score Stabilising**. The purpose of these plots are to illustrate the shape of the data as it stabilises while being processed. [Top] The rank positions fluctuate at the beginning but visibly settle by T_{20}. [Bottom] The corresponding scores indicate a clustering of some items of PII that are sensitive to small fluctuations. K-Factor 32 was used, a lower factor would produce less fluctuation. Scores at T_i are used as base for T_{i+1}.

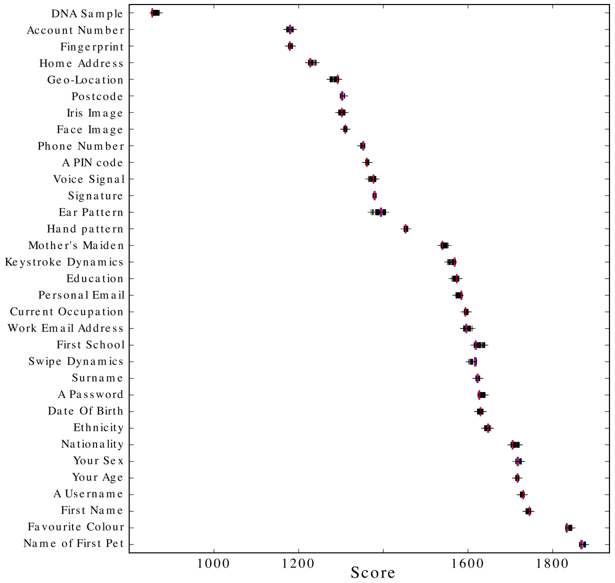

Figure 3: **Score Variation Boxplots**. Although individual scores changed over the T_{1000}, the spread of scores was narrow.

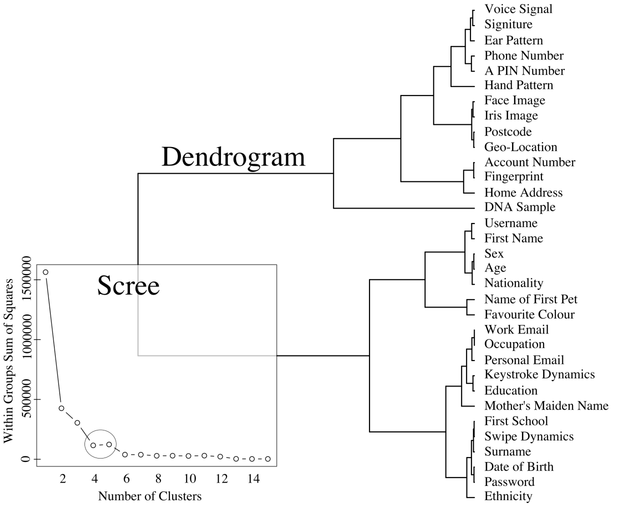

Figure 4: **Clusters**. The scree plot (bottom left) and dendrogram (right) for the social condition. Here 4-5 groups can be visually deduced from the scree type plot, at 6 there is no added benefit. The dendrogram signifies groupings due to mean distances, not in rank order. Technique: Hclust in R's stats package using ward.D2 (Murtagh and Legendre 2014)

Sex, rank 26, 27, and 28th respectively, yet the corresponding scores are 1648, 1704 (+56), and 1716 (+12).

To determine how many clusters, we use R's K-means function, as described in Murtagh and Legendre (2014), to reveal four to five clusters, see the scree plot in Figure 4. The corresponding dendrogram [right] shows the list divided into four; the length of the horizontal lines is a representation of the relative distance between the scores.

Contextual Concordance

Table 5 contains the final ranks after T_{1000}. The table is sorted based on the average rank of an item. Included is a σ score indicating volatile or stable items within the rankings, with a lower number indicating greater stability across contexts, i.e., phone number $\sigma = 9.64$, home address $\sigma = 1.52$.

Figure 5 depicts each of the three scenarios reaching stability. This stability resulted from comparing the ranking in Turn T_i to ranking in T_{i-1} using Kendall's Tau, as described by Legendre (2005). This figure illustrates a reduction of the systematic error within the algorithm by T_{300}, yet some persist beyond T_{700}. With some items the score grouping was such that some flux could not be entirely eradicated using the initial parameters and data set, e.g., in cohort S, sex (1716.15) $\approx$ age (1716.40) making the ranking susceptible to small movements in score even by T_{2000}.

To investigate the effect context has on the willingness to disclose, the final ranking from cohorts S, C and G, although comprised of 35+ judges each, are here considered as ranking from three independent judges. The context becomes a meta judge, and we can compare the meta judges' responses. Following Legendre (2005), Kendall's coefficient of concordance (W) measures the agreement between > 2 judges. The result $(W = 0.81, p < .0001)$ suggests an over-

Table 5: Final Ranks and Volatility

PII Label	S	C	G	σ
Username	30	33	29	2.08
Nationality	27	27	33	3.46
Surname	24	30	32	4.16
Favourite Colour	32	31	21	6.08
Password	23	32	27	4.50
First Name	31	29	22	4.72
Name of First Pet	33	25	24	4.93
Age	29	18	28	6.08
Sex	28	15	31	8.50
Mum's Maiden Name	15	28	30	8.14
Date of Birth	25	20	26	3.21
First School	21	22	23	1.00
Current Occupation	19	19	25	3.46
Work Email Address	20	26	13	6.50
Personal Email	18	21	18	1.73
Education	17	23	17	3.46
Swipe Dynamics	22	12	20	5.29
Ethnicity	26	11	16	7.63
Postcode	6	16	19	6.80
Signature	12	13	14	1.00
Phone Number	9	24	6	9.64
PIN code	10	17	11	3.78
Ear Pattern	13	6	15	4.72
Keystroke Dynamics	16	5	7	5.85
Voice Signal	11	8	8	1.73
Hand pattern	14	3	10	5.56
Fingerprint	3	9	12	4.58
Face Image	8	4	9	2.64
Account Number	2	14	4	6.42
Geo-Location	5	10	2	4.04
Home Address	4	7	5	1.52
Iris Image	7	2	3	2.64
DNA Sample	1	1	1	0.00

Sorted by the average rank over the three contexts (S)ocial, (C)ommercial, and (G)overnment. The standard deviation (σ) provides a reference to how volatile an item was across the cohorts. T_{1000}

arching similarity in rankings across the three contexts. We see from items such as phone number or ethnicity that there is sufficient discord between the contexts. When examining the conditions in pairs, Table 6 indicates that this seemingly high value of W is due to the close similarity of two of the conditions; the passport and social cohorts $(\tau) = 0.62)$.

Table 6: Rank Concordance, Kendall's Tau (τ)

	Social	Bank
Bank	0.44	-
Passport	0.62	0.52

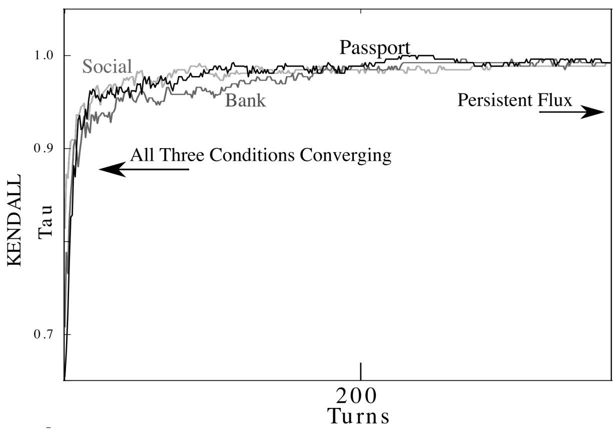

Figure 5: Illustrating Stabilisation. Kendall's Tau (τ) Rank Coefficient (Legendre 2005), shows rank stabilisation up to T_{400}. The small flux persisted.

Thus far, the contextual differences have been the focus. However, the similarities between the contexts resulting in $W = 0.81$ also tell a story. Figure 6, illustrates that distinct clusters exist across the contexts. Excluding the outlier (DNA) our K-means clustering arrives at five clusters of PII separated mainly regarding willingness to disclose (PC2).

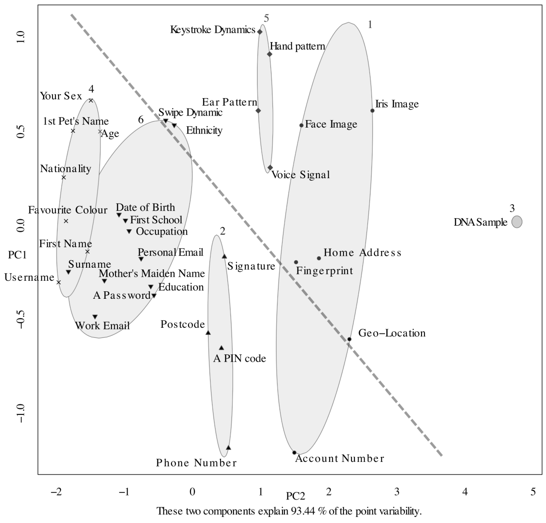

Figure 6: Cluster Plot. This plot uses the K-means clustering (max.iteration = 1000) function in R. Input was the scores of all three cohorts (T_{1000}). It illustrates that despite contextual differences, distinct groups exist within the data. PC1 associates with volatility between the cohorts, e.g., personal email ranked consistently thus holds the centre. PC2 associates with the general willingness to disclose.

The Willingness of Crowds

It would take 538 pairings for an individual to rank 33 items using pairwise comparisons of each combination. However, in Table 7 it is shown that the partial sets of 50 individual pairwise comparisons, aggregated into cohort-based rankings, is a reasonable estimate of the individual's preferences.

Table 7: N-1 Individual Estimates

	S	C	G
Correct order, > 5	59%	52%	49%
Correct order, ≤ 5	18%	15%	15%
Incorrect order, < 5	13%	20%	24%
Incorrect order, > 5	10%	13%	12%

Percentage of correct and incorrect ranking (within and outside of 5 ranking places) when compared to the base pairwise comparison matches.

Table 7 forms from comparing each individuals opening matches, i.e., fingerprint vs retina, to the final $N - 1$ rankings. The percentages indicate that a new individual using the cohort ranking will have a good (77%, 67%, 64%) chance that is in the correct order, rising (90%, 87%, 88%) when considering close, i.e., (< 5) yet incorrectly ordered items. Then, with this bootstrapping, it would is possible to keep 'turning' the system using only the individual's 50 preferences or by providing the occasional manual override to bring raise these percentages further to rectify the list towards their preferences.

Discussion

Overall, the results suggest that this is a useful and economical method to work with subjective data. By using a modest number of participants, for under 10 minutes each, we analysed the relationship between 33 items of PII. The method is scalable, as increasing the list of PII requires more individuals, rather than increasing the load on an individual. Also, the method is adaptable, as adding more individuals or more items is possible without discarding any data.

It was unsurprising that items such as DNA faired differently from details such as email address, but also, the results show the clustering of similarly perceived PII. Together, this means that along with a ranking, there is a relative distance between certain groups of PII that could be meaningful to identity-based systems that wish to collect PII yet also want to avoid overreach in their PII requests, at least from the perception of users.

Repeating the process, yielded similar yet distinct results across the contexts, suggesting that this method is sensitive enough to reveal subtle inter-contextual differences. The results of the three contexts were mixed, in some instances setting made little difference, e.g., personal email ranking at S:18, C:21, and G: 18 in others, for example, phone number ranking at S:9, C:23, and G:6, suggests that context is significant for a subset of PII. This finding echoes the suggestion that the online world is not one undifferentiated space for identity matters (Emanuel et al. 2014).

This ranking of PII from a personal utility perspective could provide system designers with new insight into the data they may consider requesting. A designer of a social site may wish to have users' phone numbers, but not understand that there are potentially over 20 items of PII that users may find more agreeable. Perhaps, said designer could seek a combination of these secondary items, forgoing the phone number. In such a circumstance the user and the extractor may prosper, as combinatorial identification has proved to be powerful (Black et al. 2012), also the user does not face the dilemma of disclosing a highly valued item of PII or retracting engagement. This latter dilemma has further importance because of the gap of knowledge that exists as designers, and PII extractors may never know why lack of engagement has occurred, or whether users are resentful for the disclosure.

Understanding such nuances could be advantageous for the efficacy of personal disclosure recommender systems (Balebako et al. 2011; Knijnenburg and Kobsa 2013) or automated disclosure negotiation agents (Baarslag et al. 2017). For instance, instead of a user training a system from scratch, all users can share the workload to arrive at a cohort ranking that has relies on the 'wisdom of crowds', or more apt, the willingness of crowds. Then, using the results, i.e., those in Table 5, to pre-populate a disclosure aid system, it would be relatively simple for a user to tune the results to their preference by conducting the occasional sense check on the ranking via a simple pairwise comparison. Likewise, new entries can reach their position by conducting pairwise comparisons in a simple binary search pattern. Future users may then avoid the quagmire of disclosure decisions required in the IoT and instead delegate the disclosure of PII to autonomous negotiating agents (Baarslag et al. 2017).

In Table 1, we nominally separated the set of PII into broad categories: knowledge and token, and biometric identifiers, the results here suggest that this conceptual model aligns with the psychological model of PII held by our cohorts. We draw a line in Figure 6 to indicate a separation of the biometric (clusters 1 and 5) from most other biographical or demographical PII, and this separation was stable across each context. The implication of this is that it shows a discrepancy in the practice of processing PII. For instance, biometric PII is widley considered of greater identifying utility than biographical PII (Rathgeb and Uhl 2011), from our results it also has a higher personal utility, yet are treated the same when it comes to extracting, processing, storing and trading such information. Therefore, there is an externality to PII disclosures, as the benefit of extraction increases and so does the 'cost' of the disclosure. In light of our results, one way to change this situation would be to regulate PII by assigning special status to items of low disclosure willingness. Then, it would be plausible to rebalance disclosures by exerting a cost on the organisation using these 'special' items of PII. Whether it remains at the discretion of organisations to show restraint in their extraction of PII or becomes a power of regulators to restrain the use of PII on behalf of the user, the results here practically inform these options.

The demographics of our sample pose a limitation. For example, it is possible that results that place home address consistently low on the willingness to disclose rank regardless of context reflect that the participants comprise mostly students and there is an element of being displaced, and temporariness to student living arrangements. Unknown in these results are any latent variables within the data, that is, whether users envisaged an actual disclosure method and incorporated a willingness to engage with the usability or novelty of such a system. Equally, items of PII such as swipe dynamics would likely have been interpreted differently by individuals as some may not be even aware of such methods, so some explanations or demonstrations would enhance future implementations of this study, but only if done promptly and succinctly. However, these limitations lessen because social drivers and interdependence not only help promote and develop technologies in the extraction of PII (Fairfield and Engel 2015; Marmion et al. 2017b), but they also provide social proof for individual and groups to establish disclosure norms (Das et al. 2015). That is because, correlations, inconsistencies, misunderstandings, and even maliciousness are all thrown into the mix with this method, just as they are in the real world. The strength of the process, however, is that outliers are not abandoned instead consensus merely counterweights them. Another point arising from the limitation posed by narrow demographics is the extent an individual may wish to exploit the aggregated cohort data from a self-reflecting homogeneous sample instead of seeking consensus from a diverse group. Either way, the final result can be trained to the individual, but how much training would be required could be a factor of such a decision, or perhaps the decision would be to take the consensus as is, thus delegating to a particular crowd. These questions require a qualitative investigation that falls beyond our scope.

Also outside the scope of this work, there are a few natural progressions to this research. The first is to consider expanding the list of PII to include other personal information such as credit score, sexuality, health records, and religious beliefs. While these are often disclosed or extracted through modern living already and can contribute to profiling methods (Kosinski, Stillwell, and Graepel 2013), we do not include them within our set of more formal identifiers. However, in the context of tomorrow, they may become valuable identifiers swept up along with metadata to provide additional ‘behavioural signals’ towards passive identification (Perez, Musolesi, and Stringhini 2018). Second, this method could track preferences over time, as new extraction methods become mainstream; for example, it would be fascinating to see how society’s reluctance erodes, or becomes entrenched, regarding biometric identification. Third, to use this method across cultures and age groups, as the ideal and norms of the UK may differ vastly internationally.

Conclusion

As users repeatedly face undesirable disclosure decisions, a lack of strategy, regarding robustness and social acceptability, hampers the long-term effectiveness of identity-dependent technologies. Our work premises that actual behaviour is often a watered down remnant of previous preferences. We arrive at this view due to the sheer amount of disclosure requests faced each day as people try to get from permissions to the task. Moreover, whether these requests are active or passive, our actions consent to accept the uncertainty and risk in PII disclosures. So, if observing the actions of individuals in the wild is unreliable, then what remains are the users’ intentions and ideals. It is prudent, therefore, to protect intentions and ideals now, before social norms adapt to what had initially been reluctant actions.

A contribution of this work is in developing a method enabling exploration of value in subjective data, in this case into PII’s relative personal value. Using this method across three contexts, submitting to government requests, transacting with commercial organisations, or self-expression on social media has provided a novel insight into the intentions and ideals of users’ willingness to disclose.

The results reflect a non-linear distribution within a set of 33 items of PII, indicating clusters of similarly valued PII, some that users seem relatively willing to disclose, but suggesting others that need more protection as they represent high personal value. We also show that, for many items of PII, the personal utility changes depending on context. Together, these findings indicate an inadequacy within current PII regulations which treat PII as a homogeneous set, whereas more nuanced regulations would perhaps help avoid any over-use of high-value items.

These findings indicate that different types of PII should warrant different levels of protection, where organisations perhaps should have to justify and bear some cost for the use of these higher value PII. Understanding these levels could promote regulation that transfers the value as a cost to those wishing to extract. Such control would be one way of tempering the escalation in PII extraction. Another would be to use autonomous consent-based agents that are less likely to succumb to the normalisation of disclosure or other cognitive bias that affect user disclosures.

The findings also raise questions regarding the transfer of PII across contextual boundaries, as the results suggest that it is possible to manipulate a user in such a way that disclosure sought in settings of high willingness to disclose get transferred to settings naturally of low willingness. Such an action would likely go against the ideals of the original discloser. Moreover, as the utility of the extracted PII increases for those now in possession, the original trade-off value, e.g., gratification, to the individual may not.

Finally, it is essential to protect PII for users across all societies, and not just from those at the head of the technological curve. For the reach of modern technologies means that regarding privacy and identity we need to think collectively, of PII as a public good, rather than individualistically (Fairfield and Engel 2015; Marmion et al. 2017b). For if others are increasingly sharing their PII, however personal their own motivations, then the pressure increases on us to share our own PII. The goalposts move, and increasingly we will be expected to share PII that we consider of high value. A public good indicates a common problem.

References

Acquisti, A. 2004. Privacy in electronic commerce and the economics of immediate gratification. In *5th ACM conference on Electronic Commerce*, 21. New York, New York, USA: ACM Press.

Acquisti, A. 2014. The Economics of Privacy. *The Economics of Privacy - Resources on financial privacy, economics, anonymity* 16.

Baarslag, T.; Alan, A. T.; Gomer, R.; Alam, M.; Perera, C.; Gerding, E. H.; and schraefel, m. 2017. An automated negotiation agent for permission management. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '17, 380–390. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Balebako, R.; Leon, P. G.; Almuhimedi, H.; Kelley, P. G.; Mugan, J.; Acquisti, A.; Cranor, L. F.; and Sadeh, N. 2011. Nudging users towards privacy on mobile devices. In *CEUR Workshop Proceedings*, volume 722, 23–26.

Balebako, R.; Jung, J.; Lu, W.; Cranor, L. F.; and Nguyen, C. 2013. "Little brothers watching you". In *Proceedings of the Ninth Symposium on Usable Privacy and Security - SOUPS '13*, 1. ACM.

Beres, Y.; Baldwin, A.; Mont, M. C.; and Shiu, S. 2007. On identity assurance in the presence of federated identity management systems. In *Proceedings of the 2007 ACM workshop on Digital identity management - DIM '07*, number 1, 27. New York, New York, USA: ACM Press.

Black, S. M.; Creese, S.; Guest, R. M.; Pike, B.; Saxby, S. J.; Fraser, D. S.; Stevenage, S. V.; Whitty, M. T.; Stanton Fraser, D.; Stevenage, S. V.; and Whittty, M. T. 2012. Superidentity: Fusion of identity across real and cyber domains. *ID360: Global Identity*.

Böhme, R.; Koble, S.; and Dresden, T. U. 2007. On the viability of privacy-enhancing technologies in a self-regulated business-to-consumer market: Will privacy remain a luxury good? In *WEIS*.

Brunton, F., and Nissenbaum, H. 2015. *Obfuscation: A User's Guide for Privacy and Protest*. Mit Press.

Cavoukian, A. 2011. Privacy by Design: Origins, Meaning, and Prospects. *Privacy Protection Measures and Technologies in Business Organizations: Aspects and Standards: Aspects and Standards* 170.

Clarke, R. 1988. Information technology and dataveillance. *Communications of the ACM* 31(5):498–512.

Correa, D.; Silva, L. A.; Mondal, M.; Benevenuto, F.; and Gummadi, K. P. 2015. The many shades of anonymity: Characterizing anonymous social media content. In *Proceedings of the Ninth International Conference on Web and Social Media, ICWSM 2015, University of Oxford, Oxford, UK, May 26-29, 2015*, 71–80. pub_id: 1310 Bibtex: DBLP:conf/icwsm/CorreaSMBG15 URL date: None.

Coulom, R. 2007. Computing elo ratings of move patterns in the game of go. In *Computer games workshop*.

Das, S.; Kramer, A. D.; Dabbish, L. A.; and Hong, J. I. 2015. The role of social influence in security feature adoption. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work &#38; Social Computing*, CSCW '15, 1416–1426. New York, NY, USA: ACM.

Elo, A. E. 1978. *The Rating of Chess Players, Past and Present*. Arco Pub.

Emanuel, L.; Neil, G. J.; Bevan, C.; Fraser, D. S.; Stevenage, S. V.; Whitty, M. T.; and Jamison-Powell, S. 2014. Who am I? Representing the self offline and in different online contexts. *Computers in Human Behavior* 41:146–152.

Fairfield, J. A. T., and Engel, C. 2015. Privacy as a public good. *Duke LJ* 65:385.

Flanagin, A. J.; Flanagin, C.; and Flanagin, J. 2010. Technical code and the social construction of the internet. *New Media & Society* 12(2):179–196.

Goodman, M. 2015. *Future crimes: Everything is connected, everyone is vulnerable and what we can do about it*. Anchor.

Grossklags, J.; Hall, S.; and Acquisti, A. 2007. When 25 Cents is too much : An Experiment on Willingness-To-Sell and Willingness-To-Protect Personal Information. In *Information Security*, 7–8.

Higgins, E. 1998. Promotion and prevention. Regulatory focus as a motivational principle.pdf. *Advances in Experimental Social Psychology* 30:1–46.

Hvattum, L. M., and Arntzen, H. 2010. Using ELO ratings for match result prediction in association football. *International Journal of forecasting* 26(3):460–470.

Knijnenburg, B. P., and Kobsa, A. 2013. Making Decisions About Privacy: Information Disclosure in Context-Aware Recommender Systems. *ACM Trans. Interact. Intell. Syst.* 3(3):20:1—-20:23.

Knijnenburg, B. P. 2013. On The Dimensionality Of Information Disclosure Behavior in Social Networks. *International Journal of Human-Computer Studies* 71(12):1144–1162.

Kosinski, M.; Stillwell, D.; and Graepel, T. 2013. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences of the United States of America* 110(15):5802–5.

Kruglanski, a. W.; Thompson, E. P.; Higgins, E. T.; Atash, M. N.; Pierro, a.; Shah, J. Y.; and Spiegel, S. 2000. To "do the right thing" or to "just do it": locomotion and assessment as distinct self-regulatory imperatives. *Journal of personality and social psychology* 79(5):793–815.

Legendre, P. 2005. Species associations: the Kendall coefficient of concordance revisited. *Journal of agricultural, biological, and environmental statistics* 10(2):226–245.

Marmion, V.; Bishop, F.; Millard, D. E.; and Stevenage, S. V. 2017a. *The Cognitive Heuristics Behind Disclosure Decisions*. Cham: Springer International Publishing. 591–607.

Marmion, V.; Millard, D. E.; Gerding, E. H.; and Stevenage, S. V. 2017b. The tragedy of the identity assurance commons. In *Proceedings of the 2017 ACM on Web Science Conference*, WebSci '17, 397–398. New York, NY, USA: ACM.

Murtagh, F., and Legendre, P. 2014. Ward's hierarchical agglomerative clustering method: Which algorithms implement ward's criterion? *Journal of Classification* 31(3):274–295.

Nissenbaum, H. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79:119.

Norberg, P. A.; Horne, D. R.; and Horne, D. A. 2007. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* 41(1):100–126.

Perez, B.; Musolesi, M.; and Stringhini, G. 2018. You are your metadata: Identification and obfuscation of social media users using metadata information.

Preibusch, S.; Krol, K.; and Beresford, A. R. 2013. The privacy economics of voluntary over-disclosure in web forms. In *The Economics of Information Security and Privacy*. Springer. 183–209.

Quercia, D.; Casas, D. L.; Pesce, J. P.; Stillwell, D.; Kosinski, M.; Almeida, V.; and Crowcroft, J. 2012. Facebook and privacy: The balancing act of personality, gender, and relationship currency.

Rathgeb, C., and Uhl, A. 2011. A survey on biometric cryptosystems and cancelable biometrics. *EURASIP Journal on Information Security* 2011(1):3.

Rose, J., and Kalapesi, C. 2012. Rethinking personal data: strengthening trust. *BCG Perspectives* 16(05):2012.

Saari, D. G. 1999. Explaining all three-alternative voting outcomes. *Journal of Economic Theory* 87(2):313 – 355.

Sarma, A. D.; Sarma, A. D.; Gollapudi, S.; and Panigrahy, R. 2010. Ranking Mechanisms in Twitter-like Forums.

Solove, D. J. 2007. 'I've Got Nothing to Hide' and Other Misunderstandings of Privacy. *San Diego Law Review* 44(May):1–23.

Solove, D. J. 2012. Introduction: Privacy Self-Management and the Consent Dilemma. *Harvard Law Review* 126:1880–1903.

Staddon, J.; Huffaker, D.; Brown, L.; and Sedley, A. 2012. Are privacy concerns a turn-off?: engagement and privacy in social networks. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, 1–13. ACM.

Tene, O., and Polonetsky, J. 2013a. A Theory of Creepy: Technology, Privacy and Shifting Social Norms. *Yale Journal of Law & Technology (. . .* 16:1–32.

Tene, O., and Polonetsky, J. 2013b. Big data for all: Privacy and user control in the age of analytics. *Northwestern Journal of Technology and Intellectual Property Volume* 11(5):240–273.

Toubiana, V.; Narayanan, A.; Boneh, D.; Nissenbaum, H.; and Barocas, S. 2010. Adnostic : Privacy Preserving Targeted Advertising. In *Proceedings of the NDSS Symposium 2010*, 1–21.

Tsesis, A. 2014. Right to Erasure: Privacy, Data Brokers, and the Indefinite Retention of Data, The. *Wake Forest L. Rev.* 49:433.

Udo, G. J. 2001. Privacy and security concerns as major barriers for e-commerce: a survey study. *Information Management & Computer Security* 9(4):165–174.

Van Zoonen, L., and Turner, G. 2014. Exercising identity: agency and narrative in identity management. *Kybernetes* 43(6):935–946.

Vila, T.; Greenstadt, R.; and Molnar, D. 2003. 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*, 403–407. ACM.

Weitzner, D. J.; Abelson, H.; Hanson, C.; Hendler, J.; Mcguinness, D. L.; Jay, G.; Waterman, K. K.; Berners-lee, T.; Kagal, L.; and Sussman, G. J. 2006. Transparent Accountable Data Mining: New Strategies for Privacy Protection. 1–12.

Williams, M.-a. 2008. Privacy Management , The Law and Global Business Strategies : A Case for Privacy Driven Design. In *Computational Science and Engineering, 2009. CSE'09. International Conference on*, volume 3, 71–79. IEEE.

Zafeiropoulou, A. M.; Millard, D. E.; Webber, C.; and O'Hara, K. 2013. Unpicking the privacy paradox. In *Proceedings of the 5th Annual ACM Web Science Conference on - WebSci '13*, 463–472. ACM.