UNIVERSITY OF SOUTHAMPTON

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING

ELECTRONICS AND COMPUTER SCIENCE

Volume 1 of 1

Values in Technology and Practice:
Using Activity Theory to consider the role of values and technology in everyday activities

by

Richard C. Gomer

Thesis for the qualification of Doctor of Philosophy

April 2018

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING

ELECTRONICS AND COMPUTER SCIENCE

Thesis for the qualification of Doctor of Philosophy

**VALUES IN TECHNOLOGY AND PRACTICE:
USING ACTIVITY THEORY TO CONSIDER THE ROLE OF VALUES
AND TECHNOLOGY IN EVERYDAY ACTIVITIES**

Richard Charles Gomer

We have incorporated technology in many of our everyday activities like working, socialising, eating and exercising. These technologies alter how we behave, and often do so quite deliberately in service of aspirations such as fitness or weight loss. Moreover, they frequently impinge upon the things that we think are important in our lives, our values; things like family, intimacy, privacy or the environment.

A design methodology called Value Sensitive Design [68] aims to account for values in technology design. However, VSD often focusses on the level of technological artefacts [49,140] rather than the broader 'context of use' that a piece, or pieces, of technology might be deployed in. This limits the kind of values that it can consider, and the influence that artefacts have on the broader human activities in which they are embedded. To make it easier to consider values in contexts of use, it could be helpful to combine VSD with theories about how human activity is structured. One such theory, often used within HCI, is called Activity Theory [41,100]. However, values are not explicitly accounted for within Activity Theory, which complicate its applicability to VSD.

In this thesis, I report on investigations in two domains: privacy, and wellbeing; both of which bring values and everyday behaviours together. Through these investigations, I identify ways that Activity Theory can account for the observed practices, propose ways in which values can be made explicit within Activity Theory, and consider some design implications that arise. In doing so, I contribute a theoretical account of human activity in which values are explicit, which may be used as a model for conducting VSD at a 'context of use' level.

Contents

Index of Tables

Index of Figures

Author's Declaration

I, Richard Gomer, declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

Values in Technology and Practice: Using Activity Theory to consider the role of values and technology in everyday activities

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:

R. Gomer, E. M. Rodrigues, N. Milic-Frayling, and m.c. schraefel, "Network Analysis of Third Party Tracking: User Exposure to Tracking Cookies through Search," in 2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT), 2013, pp. 549–556.

Richard Gomer, m.c. schraefel, and Enrico Gerding. 2014. Consenting agents. In Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing Adjunct Publication - UbiComp '14 Adjunct, 653–658. https://doi.org/10.1145/2638728.2641682

H. Marreiros, R. Gomer, and M. Tonin, "Exploring user perceptions of online privacy disclosures," in Proceedings of 14th International Conference on WWW/INTERNET 2015, 2015.

Richard Gomer and m.c. schraefel. 2016. Consentful Surveillance : Supporting User Understanding and Control. Workshop Paper for Everyday Surveillance Workshop, CHI 2016. Retrieved from https://eprints.soton.ac.uk/389835/

m.c. schraefel, Richard Gomer, Alper Alan, Enrico Gerding, and Carsten Maple. 2017. The internet of things: interaction challenges to meaningful consent at scale. interactions 24, 6: 26–33. https://doi.org/10.1145/3149025

Signed:

Date:

1. Introduction

People's values are important. Values are the bases upon which we form preferences; they are the lenses that shape our view of the world, they determine the outcomes that we seek, and shape the actions that we take to reach those outcomes. As technology has encompassed a wider and wider range of activities in our lives, including healthcare [92], socialisation [187], shopping [176], exercise [45,116,148,191] and leisure [8,38], it has changed our behaviour and societies, often in ways that we did not predict [81], but in ways that unquestionably impinge on a whole host of diverse values like fairness, equality, privacy, independence, and friendship.

The far-ranging consequences of an increasingly technologically-mediated society demand that we consider the political and ethical assumptions – and the values – that are 'folded into' technologies [87]. Or, from a more synthetic design perspective, that we deliberately consider the ways in which we embed values and power dynamics into technology as it is designed and engineered [49,68,147,181]. The fundamental importance of human values within Human-Computer Interaction (HCI) as a discipline is already recognised (e.g. [81] cited in [152]); and that HCI's *"outward trajectory"* [77] (cited in [103]) should eventually come to encompass human values is not surprising; after all, Rokeach, whose early work established the importance of considering values within the social sciences, wrote:

> *"It is difficult for me to conceive of any problem social scientists might be interested in that would not deeply implicate human values,"* [153].

Current controversies demonstrate the importance of considering how technology interacts with the things we care about at a very high level. Problems like "fake news" [7] and its apparent impact on things like democracy and broad social values show a link between factors such as economic models (online publishing driven by advertising revenue), information sharing that removes traditional editorial and fact checking roles, and values like democracy or liberalism that we consider foundational to our society. At an individual level, a lot of attention is given to values like privacy [142,154], which is enshrined in law going back over a century (and which continues to evolve today) [168,173], but which is challenged by the 'creepy' [171] or difficult-to-understand [127,150] technologies that we use. The technology that we build could look different if

we asked "*How can we ensure the accuracy of news media?*" rather than "*How can we let people share things quickly with their friends?*"; or "*How can we protect the privacy of people who are reading about sensitive topics?*" instead of "*How can we generate the revenue to sustain an online publication?*". The relative prominence that we give to those different values during design and engineering shapes the technology that we build, the society that we live in, and the behaviours that we engage in as individuals.

In light of the importance of human values, they are given central prominence within a design approach called Value Sensitive Design (VSD) that provides "*an overarching theoretical and methodological framework for which to handle the value dimensions of design work*" [68]. VSD has been used to study and design technologies in a whole host of areas; from the design of technical artefacts like web browsers [67] to socio-technical systems like wind farms [140].

However, considering a technical artefact in isolation is different to considering it as part of a broader practice, something that Le Dantec *et al.* note in their critique of VSD [49]. They argue that while VSD has a valuable role to play in designing technology to respect values of 'ethical import', a focus on those aspects can be reductive and detract from values that are implicated when the artefact is embedded in human activity. Oosterlaken similarly notes that VSD tends to be constrained to the design of technical artefacts [140] rather than the broader 'contexts of use' that proponents of VSD such as Friedman, Kahn and Borning [68] argue that it could be applied to. Rogers [152], writing in 2012, similarly suggested that value-based analyses and "context" rather than "users" would become increasingly important concerns for HCI.

The distinction between considering an artefact and a broader context of use is perhaps illustrated by Epstein *et al.*'s [59] work on the sharing of fine-grained pedometer data; although the investigation uncovers values such as privacy, honesty, support and accountability; there is little sense of the broader activity that the sensed activities are situated within. For instance, although their scenarios embed values within contexts, it is not clear why (for instance) Ellen wants to be supported in her walking practice. Is her walking practice part of an activity in service of weight-loss, or in preparation for a walking holiday? Why has Hunter gone for a run? Nardi's analysis of blog use [134] suggests that considering the overall motives behind an activity can uncover important nuance about the meaning that stands behind people's behaviour: transforming blogging from an apparent extension of a private activity like keeping a diary, to a social activity intended (for instance) to promote discourse.

VSD's deliberate and methodical focus on values in technology design remains compelling; even in light of concerns that it has, to date, often focussed on artefacts to the

detriment of context [49,140] or *"situatedness"* [181]. Expanding VSD's consideration of broader contexts of use could support the design and analysis of technology that takes into account a wider range of values and so help the individuals and societies that adopt those technologies to enact behaviours that are consistent with what they consider to be important.

However, VSD is a design methodology and not in itself an account of how human practice is structured and how technologies fit within a broader context. I argue that it would be easier to apply VSD to broader 'contexts of use' if it were coupled with a theoretical account of what those contexts contain and how they are constructed.

Other areas of HCI theory *do* provide theoretical accounts of the contexts in which technologies are used, as Nardi explains in her comparison [132] of Activity Theory, Situated Action and Distributed Cognition, and as Qin *et al.* discuss – in terms of the same three theories – in their more recent work on context-awareness [148]. In bridging the gap between a value sensitive approach, and the broader context in which technology is adopted and used, there should be value in combining a theory that provides an account of those contexts, with VSD's focus on, and methodology for considering, values. In discussing the role of theory, Kaptelinin & Nardi [97] cite Barthelmess & Anderson [21]:

> *The value of any theory is not "whether the theory or framework provides an objective representation of reality" but rather how well a theory can shape an object of study, highlighting relevant issues.*

Which, of Activity Theory, Situated Action and Distributed Cognition might be best suited to highlighting the role of values within a broader context of use, then?

There is an inherent tension with a top-down notion, like the role of *values* in human behaviour, with Situated Action's bottom-up approach, a focus on what Rogers [152] describes as "*the actual details of a work practice,*" that leaves it uncompelling as a candidate for considering high-level values and the context of technology use together.

Distributed Cognition continues in the spirit of first-wave HCI theory's cognitive approach to understanding human-computer interaction, but reframes the unit of study from being an individual user to groups of people and artefacts involved in collective tasks [132,152]. In the case of Distributed Cognition, analysis is typically based on extensive fieldwork and grounded in the interactions between individuals and artefacts. Rogers [152] and Nardi [132] point out that it lacks a set of core concepts to explain and

model the systems that are studied, rendering it more of a methodology or approach than a theory-proper.

In contrast, Activity Theory *"offers a set of perspectives on human activity and a set of concepts for describing that activity"* [133]; with a conceptual breadth that spans (at the most granular) the individual operations that make up human activity, and (at the most abstract) the needs and motives that direct that activity, and the tools that mediate people's relationships with the world.

In her comparison of Activity Theory, Situated Action and Distributed Cognition, Nardi concludes that "*Activity theory seems the richest framework for studies of context in its comprehensiveness and engagement with difficult issues of consciousness, intentionality, and history"* [132]. Activity Theory, then, seems to provide a conceptual model that could complement VSD's design methodology with a set of inter-related concepts that can explain the behavioural contexts in which technologies are situated. In particular, Activity Theory's account of how needs and motives are translated into an Object of purposeful activity seem conceptually close to the notion of value as beliefs about conduct; and a view of tools as mediational means. It is, therefore, Activity Theory that I have adopted as a basis for considering how context might be modelled in support of Value Sensitive Design.

Notably, there has already been a small amount of work that considers the role of values within Activity Theory. Barr *et al.[20]* adopt activity theory in their discussion of how values are enabled and enacted through the interface of the computer game Fable; drawing on Rokeach's original definition of a value as an *"enduring belief that a specific mode of conduct or end-state of existence is personally or socially preferable to an opposite or converse mode of conduct or end-state of existence" [153].* They draw particular attention to the role of values as linkages between specific conditions (the presence of an enemy) and modes of in-game conduct such as swinging a sword, but also hint at the possibility of considering broader values and motives that might lead to playing Fable as an activity in itself.

The aim of the work presented here is to situate values within Activity Theory; to provide a more explicit account of how they interact with Activity Theory's existing conceptual account of purposeful human activity in order to highlight their role in the overall context that technology is situated within. In doing so, the goal is to make Activity Theory more directly applicable as means of undertaking Value Sensitive Design at the level of what Friedman *et al. [68]* call *"contexts of use,"* and which Le Dantec *et al. [49]* call a *"value-sensitive analysis of the design space"*.

Introduction

In order to examine the role of values within activity, two very distinct domains that bring together values and human behaviour are studied. First, Privacy – itself a value – in the context of online advertising. Second, Wellbeing – and in particular health behaviour change – as an area that is both rich in values and motivations and which has received considerable attention from designers – both in research and commercially. The aim of choosing to study how values are implicated in these two domains is to develop constructs that generalise beyond the immediate empirical results [83]; the contrasting accounts provided by considering multiple domains provide multiple perspectives to temper the abstractions that are drawn from the results.

Privacy can itself be framed as a value [130,154,179] – it's a commonly held belief that maintaining a degree of privacy through one's actions is preferable to not doing so – and privacy is frequently invoked in discussions of technology and often at odds with how technology is designed and used [4,142,171]. The design challenge that relates privacy to our everyday activities is best framed as an attempt to resolve the "privacy paradox" - the empirical finding that, despite valuing privacy, people frequently fail to protect it through their actual behaviours [17,138,166]. To examine privacy as a value in context, it is considered in relation to a very widespread online practice called third-party tracking that supports targeted online advertising. By drawing on empirical analysis of the advertising ecosystem, and qualitative research via focus groups, and a design probe called the Web Mirror, I provide an account of third-party tracking as a multi-stakeholder activity system and use this to reframe consent – an important concept in data protection law – through Activity Theory's notion of focus shift, breakdown and Object construction.

Conversely, Wellbeing is a domain in which objective and subjective factors are considered in relation to people's overall sense of 'being well' [52,85]. Existing work in HCI typically considers wellbeing in relation to health behaviour change [82,83,129] - frequently the design of devices to promote exercise[42,182]. Wellbeing provides a domain in which a set of activities – typically exercise – are involved with a potentially diverse range of values. As Crossley [48] notes, exercise can form part of activities directed at motives as diverse as recovering from injury, losing weight, ogling other gym users or responding to a recent health 'scare'. In this domain, I draw on qualitative research with people who are involved in fitness practices to understand the values that they draw on in shaping it, and then use a speculative design exercise to re-frame fitness activities as relating to values other than health and fitness. Doing so uncovers value conflicts between what we might call the "instrumental motivation" provided by persuasive technology, and the motives and needs that ultimately direct and excite the activity.

Together, these two domains provide insight to two sides of what appears to be a many-many relationship between values and actions: Investigation of privacy situated within third-party tracking provides an account of a single value embedded in many online activities. Considering wellbeing, in particular health behaviour change, provides a fairly constrained set of activities (like running, walking or cycling) that are manifestations of diverse values and motives.

In considering the results from these studies in isolation and in comparison with one another, the contribution is to provide an account of how values, as an explicit concept that is foundational to Value Sensitive Design efforts, fit within Activity Theory, which provides a framework for considering the contexts of use that VSD could productively be used to design for. Effectively bridging the gap between what are, currently, two distinct bodies of HCI theory.

Practically, my work has used a mix of grounded theory [163] – building theory from concepts that are apparent in the data that I have collected – and theory-led analysis that has drawn on Activity Theory as a lens through which empirical data can be understood and explained.

Taking this dual approach has allowed me to consider how Activity Theory can be used to think about Privacy, Wellbeing, and Values in a general sense, but also to draw upon concepts that were present in the empirical data even if they were not already explicit within Activity Theory. In doing so, I have been able to introduce new concepts into Activity Theory which are grounded in the data that I have collected. My overall approach resembles what Goldkuhl and Cronholm call "Multi-Grounded Theory" [71] in which theory is developed based on a combination of "*theoretical, empirical, and internal grounding.*"

My conclusions are therefore two-fold: First, domain-specific accounts of how values play out in Third-Party Tracking and Wellbeing, grounded in the concepts that I encountered in the data. Second, Activity Theory-led analyses of those grounded accounts, and – where necessary – the introduction of new concepts into Activity Theory to account for observations that did not fit into the existing framework that Activity Theory provides.

The rest of this document is arranged into five parts; the first outlines the **general background** of values research both in the social sciences and HCI, as well as the theoretical and methodological approaches that I have used. In particular, I draw on literature from behaviour change and persuasive technology to consider the ethical imperative of being sensitive to values when designing technology that will be situated within people's everyday activities.

The second and third cover **privacy** and **wellbeing** in isolation, respectively; situating findings with Activity Theory and suggesting theoretical implications for Activity Theory that arise for each domain.

The fourth draws generalisable lessons from those two domains into a more **unified view** of how values can fit into Activity Theory's account of human activity, explicitly situating values in Activity Theory in a general theoretical sense. In particular, I expand on the link between values and needs, and consider how values act as constraints on how activities are constructed.

Finally, I conclude with a summary of the conclusions in the domains of Privacy and Wellbeing, and the general theoretical implications for Activity Theory.

2. General Background

In order to situate Value Sensitive Design within HCI, this chapter discusses the background, and existing work related to technology design, behaviour change, human values, and theoretical approaches from HCI, most notably Activity Theory.

I begin with a discussion of what we mean by values, and work to date on incorporating them more explicitly into technology design through Value-Sensitive Design (VSD). Understanding precisely what a value *is* and how, to date, they are used within VSD provides an account of what exactly it is that, conceptually, needs to be situated within Activity Theory.

I then consider work on behaviour change and persuasive technology. At first glance, this work appears to be more closely related to Wellbeing than the more general topic of VSD However, these topics provide an important explication of the ethical challenges at the intersection between design and behaviour that are generalisable beyond attempts to deliberately influence behaviour and therefore provides an account of why being value sensitive at a context of use level is important.

In keeping with the overall structure of this document, and to aid clarity, the domain-specific background and related work for Privacy and Wellbeing is presented in those sections, respectively.

2.1 Values & Value-Sensitive Design

Colloquially, discussion of values is common in political discourse – for instance in the amorphous "British Values" that all political parties seem to lay claim to – and there is an intuitive sense that values – perhaps things like tolerance, or tradition, or internationalism – are important concepts that can be expressed through the machinery of government. The academic literature paints a picture of values that doesn't differ too much in essence from that colloquial sense of the word.

Milton Rokeach made important contributions to the definition and study of values in his 1973 book "The Nature of Human Values" [153]. Rokeach identifies five properties of values:

1. A value is enduring; they are not fixed, but change slowly over time both on an individual and societal level. Notably, Rokeach points to the necessity of balancing values against one another, often in situations in which multiple values are implicated.
2. A value is a belief; in particular, a value is a prescriptive or proscriptive belief in which "*some means or end of action is judged to be desirable or undesirable*."
3. A value refers to a mode of conduct or end-state of existence; called instrumental and terminal values, respectively. Rokeach suggests that this distinction is important because although there is a functional relationship between the two types, there is not necessarily a one-to-one relationship between the two types. There are multiple sub types, but Rokeach specifically points to inter versus intra-personal (societally oriented versus personally oriented) terminal values; and two types of instrumental value: moral and competence – the former typically relating to interactions between people (such as being polite or honest), with competence values focussed on one 's own actions (being imaginative or logical).
4. A value is a preference as well as "*conception of the preferable*"; a value determines what we will find desirable and therefore which, of two scenarios we will prefer. Possessing a hierarchy of values implies preference between different scenarios even when they relate to different values.
5. A value is a conception of something that is personally or socially preferable; it can be conceived of as applying to a differing extent to different people, to oneself or to others.

Rokeach also differentiates between values and needs. A human being needs to eat, and to move, in order to survive; but they may also eat to mark a social occasion or enact a social convention, or move to enjoy their surroundings and compete with others. It is these latter concerns – of leisure, socialisation, family, competition and improvement that constitute values and which motivate activity beyond what is necessary simply to survive.

More recently, Shalom H. Schwartz developed a Theory of Basic Human Values based on cross-cultural research [159]. Schwartz proposes a set of ten basic values, divided into four higher-order groups. These values seem to exist across all human cultures, hence their designation as "basic". Schwartz contradicts some of Rokeach's earlier suggestions about values. In particular, he downplays the functional importance of a distinction between instrumental and terminal values.

Schwartz defines the nature of values in six points:

1. Values are beliefs, linked to affect; congruent with Rokeach's suggestion that values are beliefs about what is desirable.
2. Values refer to desirable goals, and motivate action. In this sense, Schwartz sense of values is closer to the idea of terminal values expressed by Rokeach.
3. Values transcend specific actions and situations. This is itself implicit in Rokeach's writing.
4. Values serve as standards or criteria that guide selection and evaluation of things like actions, policies or people. This is similar to the idea of values as a conception of the preferable.
5. Values are ordered by importance; which reflects Rokeach's notion of values being hierarchical.
6. The relative importance of values guides action.

The set of ten basic values proposed by Schwartz is not merely descriptive, the theory also posits relations between them (reinforcing or antagonistic), and suggests distinct social roles that the values might play.

2.1.1 Value-Sensitive Design

In a more specifically design-oriented context (as opposed to the social science context of Rokeach and Schwartz), Values are a key element, as the name suggests, of Value Sensitive Design (VSD). VSD is, according to Friedman, "*an overarching theoretical and methodological framework for which to handle the value dimensions of design work*" [68], a systematic approach to accounting for human values throughout the design process.

VSD defines values in a loose sense, as being "what a person or group of people consider important in life" [68]. Methodologically, VSD is described as an iterative process that combines, in stages, conceptual, empirical and technical investigations in a so-called "*tripartite integrative methodology.*"

Of particular interest, given the topics covered in this thesis, are efforts to apply VSD to the design of privacy mechanisms within Web browsers, notably Friedman's own work on the visualisation of browser cookies [67] which involved, among other things, a value-sensitive explication of what informed consent means in the context of the Web. More recently, but in a similar vein, Xu *et al.* used VSD to propose a suite of Privacy-Enhancing Tools [179]. Privacy is a value that's considered – among others such as accuracy – in Epstein at al's [59] application of VSD to the sharing of fine-grained pedometer data.

Boyd *et al.* [34] consider privacy in a work activity-tracking system, alongside values such as autonomy, intimacy, ownership and security; concepts that occur frequently elsewhere in the privacy literature (e.g. [25,154]).

Other examples of VSD include Kuznetsov's [104] study of the values that are important to Wikipedia contributors, which found values – such as community, reputation and autonomy – are "*deeply embedded in the technology*." Friedman & Kahn [69] consider the role of values in augmented reality technology, drawing attention to physical and psychological well-being, privacy, deception, informed consent, ownership and property, and trust. De Greef *et al.* [75] build on notions of ethical import and VSD to create a design methodology, sCEthics, within a Cognitive Engineering methodology used for the design of autonomous robotic systems; noting, for instance, the ethical dilemmas that result from arming robots with weapons.

VSD's theoretical bases, and claims of utility, are not without criticism, though. Notably, Borning & Muller [30] offer critique on two fronts; that a stance on the universality of human values is unnecessary, and that the roles of participants' and researchers' voices are important when conducting and reporting on VSD. To the first point, they suggest that claims from within the VSD literature concerning the universality of human values are unnecessary; that VSD can accommodate a range of viewpoints on *where* values come from. To the second, they draw attention to the need to strengthen participants' own voices in VSD, and to more explicitly acknowledging the voice of the researcher. An important element of this critique is the notion of power, specifically the power relationships between participants and researchers, an issue that is also raised by Zaman & Jafari Naimi [181] in their account of an attempt to apply VSD. In fact, given the cultural importance of values and their role in human society, it seems that power – in particular in so much as it pertains to *whose* values are considered most important – is a fundamental issue for design to consider.

Le Dantec *et al.* [49] provide a critique of the VSD methodology, and offer some suggestions for improvement. For instance, they suggest that an initial technical investigation can more effectively uncover values than beginning with a conceptual one, and confronting researcher bias and power through the choice of elicitation methodology. Notably, they criticise a focus on the twelve "values of ethical import" that are often used in VSD, on the grounds that *"this heuristic privileges a discursive definition of values over values that may be discovered or encountered through investigation, and produces systems that are aligned with these twelve values rather than those aligned with values expressed in the context of design.*"

Barn *et al. [19]* use experiences from integrating technology into youth offending teams to propose a model for incorporating VSD into a software engineering process more formally, in order to continue to focus on values after the design phase.

Reading the literature around VSD, one can ascertain three particularly salient points:

1. That design can take place at multiple levels – from individuals artefacts, through to more systems-oriented socio-technical structures – and that values can be implicated at any (or all) of those levels [68,140,181].

2. That design has, by necessity, to resolve value conflicts that arise within or between stakeholders, and so should consider the power relationships that exist between them [49,181].

3. That precisely what a value is, and whether they are universal, is contested; but that VSD can accommodate – and be used productively within – multiple viewpoints [30].

VSD is not constrained to the design of ICTs, it has also been applied in domains as diverse as the design of wind power projects [140] and cardiac implants [51]. Nonetheless, the majority of VSD work remains, as Oosterlaken notes, "*at the level of technical artifacts,"* rather than at the level of socio-technical systems. Friedman *et al.* explicitly suggests that a "context of use" could be the starting point for a VSD enquiry [68] and suggest that Borning *et al.*'s application of VSD to the use of software called UrbanSim in an urban planning consultation [31] is an example of doing so. Nathan's [135] analysis of technology use within an eco-village considers the community itself as a source of values, and examines how they interact with IT use. Durrant *et al.* [57] also consider a context of use in their analysis of how values are manifested through technology in memorialisation of the Rwandan Genocide.

Considering the role of values in relation to everyday practices like walking, running, losing weight or using social media suggests that these contexts are in themselves important; a sentiment echoed in Le Dantec's critique [49]. Intuitively, use of social media (for instance) is part of a broader social practice that might also involve face-to-face interactions, shared experiences, and physical contact. The diverse contexts in which a single tool or platform might be used is something acknowledged by Marwick & boyd [123] in their work on context collapse: the observation that using the same platform to communicate with multiple social groups or for multiple purposes can bring previously distinct contexts together. There is a sense in their work that these contexts might have different values associated with them; that the contexts themselves determine how the technology will or won't, should or shouldn't, be used.

Context seems to be, therefore, an important determinant of how values relate to our everyday practices, and hence how values are relevant to the design technical artefacts that we build to facilitate and support those behaviours. Incorporating theories about what a context of use *is*, and how such contexts operate, could therefore be a promising contribution to enabling the application of VSD beyond the level of technical artefacts.

One broad area area of research that brings values, everyday practices and technology together is the idea of behaviour *change;* the recognition that our behaviours sometimes fall short with respect to the outcomes we (or others) would like to see and the use of tools and technology to try address those shortcomings. It's behaviour change that I turn my attention to next.

2.2 Behaviour Change

The questions raised in this thesis about how values are reflect in our everyday activities, and how we can design in response, are closely related to the notion of behaviour change; the deliberate use of interventions that aim to alter people's behaviour, often (though not always) to change outcomes related to health [82,83,129]. Behaviour change is an area of research where technology is deliberately designed to have an effect on a larger behaviour, like eating or exercising. The wearable devices being deployed in health and fitness scenarios usually have the explicit aim of encouraging a change in daily practice for health reasons. For instance, encouraging more physical activity. Intuitively, the very notion of intervening to alter behaviour is based on the premise that some behaviours we engage in lead to outcomes that, evaluated against some set of criteria – often (but by no means always) an individuals own conscious preferences or values – are considered negative, or at least far from ideal.

No discussion of designing tools that are embedded within human activities can be entirely separated from concerns about how human behaviour will be impacted by them, though. That is, fundamentally, why being value sensitive in their design is important; and it is a direct implication of the view within Activity Theory that meaningful activity is the result not just of cognitive processes within an individual, but embodied through individuals' purposeful action interactions with the world [98,101,133]. This means two things: First, that in concerning ourselves with how values are enacted within purposeful activities that involve tools, and asking how we can design those tools in response, we are implicitly concerned with how the activity will be influenced or transformed as a result; for instance by constructing an activity that has less of an effect on our privacy. Secondly, that as designers who are making judgements about how to resolve the tensions between

different values, or different stakeholders in a multi-stakeholder activity system, we are placed in the same ethical position as designers working on what are deliberately framed as "behaviour change" devices; for instance, when we choose (or, more commonly, choose not) to make privacy a core value in design, we are altering the horizon of possible actions that are available to potential users, and frequently – because of the bounded rationality [10,89] that characterises how humans make decisions – impacting the outcomes of activities in ways that many users will not even notice, l*et al.* ne fully appreciate.

Three bodies of behaviour change literature are of particular relevance: First, the theoretical accounts of *how* behaviour can be changed, which inform many behaviour change interventions. Second, the notion of behavioural "nudges," as proposed by Thaler & Sunstein [167], which has been widely adopted in domains such as public health promotion and so provides an interesting account of the ethical and theoretical challenges in behaviour change. And third, in the delivery of "persuasive" (ie behaviour altering) techniques specifically via technical means, so-called persuasive technology.

Before discussing those two areas, though, it is necessary to briefly visit the psychological explanations of why our behaviour sometimes needs to be changed in the first place.

2.2.1 Bounded Rationality & Internalities

Human beings make rational decisions, up to a point. In some circumstances, human behaviour can be predicted based on the *utility* of a particular decision. This rational view is the basis for classical economics. In many circumstances, though, humans do not make the choices that the model of a purely rational "*homo economicus*" would predict. In essence, human beings are rational only within certain, fairly predictable, bounds [12,167]

Of particular interest is Daniel Kahnemann and Amos Tversky's work on psychological biases and heuristics [169], for which Kahnemann was later awarded a Nobel prize. Their collaboration produced much of the seminal work in behavioural economics [84] and is closely associated with dual-process theory, an account of the two parallel but interacting systems that are involved in human decision making. System I, the fast automatic system that leads to intuition, and system II, the more reasoned conscious system typically associated with homo economicus [89]. Other accounts that draw such a distinction include the Elaboration Likelihood Model, which differentiates between a thoughtful *central route* and an inferential *peripheral route* for processing stimuli [145].

The distinction between two systems, or at least two broad kinds of information processing, is a theme that's frequently revisited in the literature about nudging,

persuasion and self-control. Alós-Ferrer and Strack [10] review theories from psychology and economics, and present a "*multiple selves model*" which aims to bridge similar theories from the two disciplines into a more unified account. We'll return to these notions of dual-process, and dual selves, later.

As a result, human beings often make decisions that seem to be completely at odds with their own interests, as they themselves would define them. While dual-process psychology and its accounts of the biases and heuristics that underlie our choices can explain many of the reasons for those bounds on our rationality, behavioural economics provides a conceptual model that helps to account for the costs that these "bad" decisions impose upon the individual, the notion of *internality* [8,76,177].

In a similar way that externalities are a form of market failure, internalities can be viewed as a failure of an individual to make optimal decisions with regard to their own welfare. Allcott and Sunstein explain internalities as "*costs we impose on ourselves by taking actions that are not in our own best interest*" [8]. The concept of internality has been applied to describe health-related activities such as smoking [76].

The central tenet of internality – the idea of *negative* side-effects on oneself – is clearly an important principle when considering technologies that are notionally grounded in individual values. After all, whether something is 'negative' is a question of preference, and preference is value-based [153].

When applied in a public policy context (as much of the work on internalities from behavioural economics *is*) the regulation of internalities raises moral questions over the role of the state, and whether top-down interventions to minimise self-control problems are a legitimate function of government, in the same way that regulating to control externalities is [177]. From a personal technology design perspective, in which we assume voluntary adoption of an intervention, these concerns are avoided though, perhaps providing a motivation to pursue self-directed avenues for internality reduction in addition to, or *instead of*, top-down imposition of behaviour change interventions.

In some cases, the distinction between an internality and externality may not be clear cut. For instance, many of the negative effects of smoking are borne by a smoker's family as a result of second-hand smoke. Are those effects externalities, or are they more properly considered an internality, given that they're effects which should factor into the smoker's own decision making [76]? Values prove conceptually useful in this scenario: I contend that side effects that are considered negative vis-a-vis an actor's own values might be considered internalities. Thus, the effects of smoking on those around them might be considered an internality with respect to somebody who valued the health of their family

or community, while being considered as purely an externality in somebody who had no such community spirit. Likewise, the environmental effects of driving a petrol-powered vehicle are conventionally labelled externalities, since they inflict a cost on broader society, but they might also be considered, as a result of a personal belief in and commitment to environmentalism, as an internality.

This is an important distinction. If an activity leads to internalities, an individual might (rationally, at least) have some interest in altering their behaviour to reduce those internalities. If an activity leads *purely* to externalities, why should a rational individual care? The implication is that internalities are *in-scope* for self-adopted behaviour change, while externalities are not. This is a particularly interesting proposition in the case of activities that produce both internalities and externalities, since it raises the possibility of leveraging internal effects to motivate behaviour change in a way that is also beneficial to those around us. Helping people to find those internalities – where their practice conflicts with the things that they care about – is one area that VSD, deployed at a context of use, could be used to positive effect.

There are, then, a number of factors that lead to behaviours that might be considered sub-optimal, or worth changing. There are also theories, often adopted from health psychology, that aim to explain how behaviours can be influenced through deliberate interventions. I consider those next.

2.2.2 Behavioural Theory

There are a number of theories that explain behaviour change, and which cover different conceptual "levels" [83]: from highly generalisable (but non-specific) meta-models to highly specific (but non-generalisable) empirical findings. Hekler *et al.* note that:

> *"Behavioral theories vary widely in which behaviors they describe and how these behaviors are described. Some theories focus on one behavior (e.g., smoking), others describe the specific process (e.g., relapse prevention), and still others describe dynamics between behaviors and other constructs (e.g., theory of planned behavior). As a consequence, behavioral theories can be categorized in a variety of ways. One common distinction, for instance, is between behavioral theories that describe determinants of behavior (e.g., the health belief model) versus the process of change." – Hekler et al. [83]*

Theories commonly adopted as the basis for behaviour change interventions include Ajzen's Theory of Planned Behaviour [6] and the Transtheoretical Model of Behaviour Change [146].

Ajzen describes the Theory of Planned Behaviour as *"a theory designed to predict and explain human behaviour in specific contexts"* [6]. It describes three main determinants of behaviour: *Attitude towards the behaviour, Subjective norms, and perceived behavioural control.* The theory does not account for where these attitudes, norms and perceived control come from though, nor the broader context in which the behaviour is occurring or, from a change perspective, *why* someone might be changing their behaviour at all.

The Transtheoretical Model [146] (TTM) attempts to divide the process of behaviour change into multiple stages, as a means to assess readiness for change and scaffold interventions. However, it has been criticised on several grounds – including a lack of predictive power, internal inconsistencies and lack of supporting evidence – and West [175] argues that it should be abandoned entirely. These criticisms aside, the TTM is essentially orthogonal to a broader context; it describes a process through which a particular behaviour is changed, rather than providing a framework for situating that behaviour within a broader system of motives (including values), other behaviours, and factors such as culture and social influences.

These theories, while undoubtedly useful in designing behaviour change interventions, lack the conceptual breadth of modelling, more broadly, the notion of *context of use* that a Value Sensitive Design exercise might wish to engage with. There is a difference in kind between a theory that helps to explain *how* to change behaviour, and a theory that can explain why we might *want* to change behaviour, and how behaviour change might relate to values-oriented concerns like family, community or achievement.

One area in which behaviour change has been considered within a broader context is in Thaler & Sunstein's "Nudge"; which offers not only a discussion of how bounded rationality might shape how we present decisions to people, but also offers a political and ethical framework that is helpful is explicating and considering these broader questions about why we might want to alter behaviour, and the ethical implications of doing so.

2.2.3 Nudging

"Nudging" is one public policy response to citizens' bounded rationality that aims to present choices in ways that promote positive outcomes. In comparison to a "sin tax," which aims to increase the economic cost of an option that has inter- (or exter-) nalities and hence to affect consumption by appealing to a person's rational economic reasoning,

a nudge intervenes in a choice *without* altering the economic incentives, and instead uses the predictable but often *irrational* ways that people make decisions to effect behavioural change.

Simply, a "nudge" is a change in a choice architecture (the way that a choice is presented) that uses some aspect of human psychology other than mere economic rationality to predictably alter the decisions that people make. Richard Thaler and Cass Sunstein, who first proposed the idea in their book *Nudge* describe them as *"any aspect of the choice architecture that alters people's behaviour in a predictable way without forbidding any options or significantly changing their economic incentives"* [167].

Nudges are of interest less in their guise as a policy mechanism (since that is not the focus of the thesis), and more as a mechanism that deliberately leverages our predictable irrationality [12] to effect behaviour change. Techniques and interventions that could legitimately be considered "nudges" could clearly be implemented by organisations and individuals other than governments, and the prospect that an individual might nudge themselves (for instance, as in the example given by Thaler & Sunstein (p.48), by purchasing "Clocky") is of interest to people designing technologies that assist personal behaviour change.

The notion of nudging is not uncontroversial, or without potential drawbacks, though. An oft-cited critique is the tension that sits at the heart of Thaler & Sunstein's "liberal paternalism". As Thaler & Sunstein themselves acknowledge, this seems on the face of it to be an oxymoron. Despite the good intent behind many nudges, the fact remains that nudges deliberately work to alter behaviour in favour of a particular set of values, which are not *necessarily* those of the people who are being nudged [32]. Selinger & Whyte [162] go further, arguing that nudges can alter the *meaning* of an action, moving them from interventions at a subconscious automatic level, to ones that can change the values that are implicated in a particular choice. They revisit Thaler & Sunstein's example of a fly painted on a urinal (to reduce spillage [167]) to consider how such a nudge might be received in a culture that revered all life.

Technology has a potential role to play in resolving some of these conflicts by identifying and delivering behavioural interventions that are consistent with an individual's *own* values and attitudes, as well as, of course, being a medium through which nudges can be delivered.

2.2.4 Persuasive Technology

The idea of using technology as a means of altering behaviour (through nudges, or other means) is well established within HCI, specifically in "persuasive technology" or "captology" [63]. Work in this area tackles a number of related areas; for instance Fogg provides a taxonomy (the "functional triad") of *how* technology can persuade people [63], a behavioural model that explains the psychological precedents necessary for action [64] and a design process for persuasive technologies [65].

Persuasive techniques have been used in a wide range of technologies, particular those related to health-behaviour change such as diet change or exercise promotion. Notable examples include Fish'n'steps [116], which used progress in a computer game to motivate real-world physical activity, and wearable devices such as the Fitbit that combine techniques such as competition, goal-setting and (virtual) rewards to encourage physical activity.

However, like non-technological nudges, these persuasive technologies are typically paternalistic in the sense that "positive" behaviours are identified externally and then 'baked-in' to the finished interventions during the design process. This raises an ethical question (again similar to that of nudging) about whose will is being enacted via the behaviour that persuasive devices can encourage [24]. Purpura *et al.* use the design fiction "Fit4Life" [147] to explore these issues. Fit4Life is described as "*a thought experiment to stimulate reflection on the social, ethical, and political issues that technology raises.*" They raise three critical issues: 1) who should be in control of individual behaviour, and the (in)distinction between encouragement, persuasion and coercion; 2) reinforcement of the trend towards control, rationalisation and 'optimization' (my quotes) of human behaviour; 3) that technological control raises issues of surveillance, and the process of quantification "discards the value of personal experiences and emotions".

Berdichevsky & Neuenschwander identify what they call "the Golden Rule of Persuasion," "*never seek to persuade a person or persons of something they themselves would not consent to be persuaded to do.*" Their framework makes a distinction between the persuader and the persuaded and points to an important aspect of what we might consider *personal* behaviour-change devices; the persuader (in the case that persuasion is involved) is the same as the persuaded person. That is to say, the motivations for seeking behaviour change belong to the target of the intervention themselves. This addresses, to some extent, Purpura *et al.'s* first concern.

Critique of the motives and means that are built into persuasive technologies echoes that of technology more generally. Behavioural interventions are perhaps an extreme example

of a more general principle, that what on the face of it are 'mere tools' can in fact embody a range of assumptions about means and ways of doing, and conceptions of what is preferable – values. Introna [87] provides a more generalised critique of several apolitical views of technology that minimise the concern which is given to non-user agency in discourse around technology. Noting, for instance, that *"Folded into – or enclosed in – the ongoing co-constitutive horizon or nexus of human and technology relationships are (un)intentions, (im)possibilities, (dis)functions, affordances/prohibitions that renders possible some ways of being and not others, that serves the (il)legitimate interests of some and not others"* he suggests that, ethically, it's important to be able able to reverse those foldings – to be transparent about the intentions, possibilities, functions and affordances that technology encloses. Engendered within this observation is the implication that persuasive technologies should be open about what they are trying to persuade people to do, that those assumptions, values and intentions should be open to scrutiny, as well as somehow aligned with those of the end-user as Purpura, Berdichevsky, Neuenschwander *et al.* suggest.

The behavioural theories and persuasive techniques that I have discussed are important to consider from a Value Sensitive Design perspective for two reasons:

First, because there are clearly good reasons to try and influence people's behaviour to promote outcomes such as obesity reduction, and these techniques offer ways to do that. Attempts to do so are not value-*less* in fact quite the opposite, they are driven by ideas of health, quality of life, and wellbeing.

Second, because they provide insight into the ethical and contextual issues that arise when designing to deliberately change or support human behaviour. The role of power, and of personal empowerment, is crucial; and is reminiscent of the discussion of power relationships in discussion of VSD itself. Designing for behaviour change forces us to ask in whose interests we are designing, and how our interventions will interact with concepts like agency, empowerment and fairness.

From the perspective of understanding how values might be implicated in the process of behaviour change (or other technology-supported human practices, for that matter) we need a theoretical account of human activity that is conceptually broader than how individual behaviours are determined and changed, though. For this, I turn to Activity Theory.

2.3 Activity Theory

As I have already explained, in considering how we might model the context of use that Friedman *et al.* suggest VSD could be applied to [68], Activity Theory seems to provide a promising framework for exploring and modelling those contexts.

Before I move on to explaining Activity Theory itself in more detail, a quick word on theory in general: As Reeves *et al. [149]* explain, *"a theory is an organised, coherent, and systematic articulation of a set of issues that are communicated as a meaningful whole,"* theories *"give researchers different 'lenses' through which to look at complicated problems and social issues, focusing their attention on different aspects of the data and providing a framework within which to conduct their analysis."*

The benefit of adopting theory within a Value Sensitive Design process should be similar to the benefits of adopting theory anywhere else. It provides a set of validated concepts, and an account of the relationships between them, that can be built upon; in this case as a means of guiding the process of identifying and accounting for values within the broad context of use that is of interest.

Unlike the behavioural theories discussed previously, Activity Theory provides a much broader account of human activity that links the ultimate high-level motives with the low-level atomic operations that human behaviour is ultimately composed of. It is, therefore, Activity Theory that I have adopted as a basis for considering how context might be modelled in support of Value Sensitive Design.

Activity theory provides a very thorough account of objective-oriented activity and has a long history of use (and development) within Human-Computer Interaction [100]. Most importantly, given the subject at hand, it is conceptually broad enough to provide an account of how abstract needs and motives translate into the individual actions that compose human behaviour.

2.3.1 Basic Concepts

Activity Theory (AT) is a theoretical approach that considers how human behaviour is constructed as 'activities'. It is relatively well-established as a theoretical approach within HCI, and has a number of conceptual elements that make it well-suited to the issues described in this research. This chapter gives a brief overview of the key features of activity theory, and explains why I adopted it as a basis for the discussion of my own research.

Essentially, Activity Theory frames the behaviour of purposeful actors as *Activities,* and takes those activities as the basic unit of analysis. An *Activity* is a "purposeful interaction with the world" (undertaken by a subject) that works towards a desired outcome (an 'Object'). Activities are often long-lived and consist of complex series of more basic actions and goals. The idea that actions can only be understood properly in the context of the activity that they are part of is foundational to Activity Theory – the act of one individual fashioning spears is given meaning in the broader context of a group that hunts animals for food. Moreover, AT allows us to move the study of human-computer interaction away from a focus on how a user operates a system (a 'user-system view') and expand our inquiry to consider aspects such as "higher-level, meaningful tasks" [96] (p.36). On the face of it, this seems well-suited to the investigation of the role of values in design, since intuitively it's those meaningful tasks where values are implicated.

Central to the concept of Activity is the notion of an ***Object***. An Object (or objectified motive) is the ultimate aim of an activity. The object encapsulates the *motives* for acting and directs the activity. Understanding the Object of activity is crucial to understanding the activity itself, as apparently similar activities may have different Objects that subtly alter the way activity is understood by those involved, and alter the way that an activity plays out. This is particularly important where an activity involves multiple parties, who could conceivably see the same activity as contributing to different Objects. Despite the linguistic similarity, it is important to distinguish the Object of Activity from physical objects that may be involved in an activity; for instance tools, and even physical outputs from actions within the activity. By way of example, Hardman [80] recounts a classroom situation in which students produce written reports using computers. These written reports may be thought of as objects (albeit virtual ones), but their production is tangential to the Object of the activity, which is the production of technically literate students.

There are complementary conceptualisations of Object provided by different Activity Theory communities, that reflect slightly different approaches to the use of Activity Theory. Kaptelinin discusses the distinction between the concepts of Object used by Dmitry Leontiev (and adopted in the more HCI-centric AT of Nardi and Kaptelinin) with a more group-activity oriented conception adopted by Engeström [93]. The main differences between the two perspectives are summarised in Table 1 [93]. Unless otherwise stated, I'll use the conceptualisation provided by Leontiev, and adopted by Kaptelinin and Nardi, throughout.

Facets of Activity	Leontiev	Engeström
Activities are carried out by:	Individuals	Communities
Activities are performed:	Individually and collectively	Collectively
The object of activity is related to:	Motivation, need ("the true motive")	Production (what is being transformed into the outcome)
Application domain:	Psychology	Organizational change

Table 1: Two perspectives on the Object of Activity, from [93]

Objects are always formed in response to ***Motives***. As Leontiev explains, *"Behind the Object, there always stands a need or desire, to which [the activity] always answers,"* [96] (p.73). Kaptelinin explores the distinction between motives and the Object that is constructed in response to them [93], to conclude that Objects are also formed to take into account *constraints,* including *"available means, other potentially related activities, and other actors involved, each with their own motives and objects."* Hence, Objects may be reconstructed as motives and constraints change. The motives that are objectified in an Object are potentially diverse, and it seems intuitive that some of those motives might be related to, or are in themselves, *values*. Activity Theory's consideration of *motive* as the underlying driver of actions and activity – something that distinguishes it from, for instance, Situated Action Models [132] – makes it particularly appropriate as a theoretical framework for considering the role of values – by definition conscious but abstract arbiters of preference – in technology design.

Activities needn't arise in answer to a single motive, though, they may be *polymotivated* – that is, they may be shaped by more than one motive. Different relationships between motive, object and activity could be used to model polymotivation [96] (p146) although Kaptelinin and Nardi conclude that the most promising model is that *N motives* are objectified into *one object* which shapes *one activity* ("Model C, N:1:1" p.149).

Activity itself is composed of a series of ***actions***, which are themselves made up of ***operations***. The notion of action accounts for intermediate steps that are involved in achieving the Object of activity, such as fashioning a spear in order to hunt animals for food. Operations are distinguished from Actions in that they are, from the subject's perspective, atomic. Any complexity in an operation has been automated to the point where the subject needn't consciously break an operation down into more granular tasks in order to carry it out. Importantly, actions can be internalised to become operations,

through a process of learning, or as the result of adopting a tool that embodies that operation [156] (p.108).

This three-level object/action/operation hierarchy forms the basis of most Activity Theory research; but González *et al.* [73] point to a conceptual gap in between actions – which are very short in duration, and objects, which are slow to change and long-lived. They suggest the need for an intermediate concept that can represent series of thematically-related actions – such as preparing a report – that are bigger than a single action, but are not an object in their own right. They call this construct an *ensemble* and define ensembles as *"sets of actions thematically connected, oriented towards a particular purpose, and framed within a particular object-related activity"*. I have found the notion of an ensemble very useful in constructing accounts of the activity that I have observed through my research, and so I adopt it as a concept from the beginning.

Figure 1 shows the relationship between Motive, Object, Activity, Actions and Operations, and associated concepts. It is based on a diagram by Wilson [178] but updated to include Motives, based on [93], and Ensembles, based on [73].

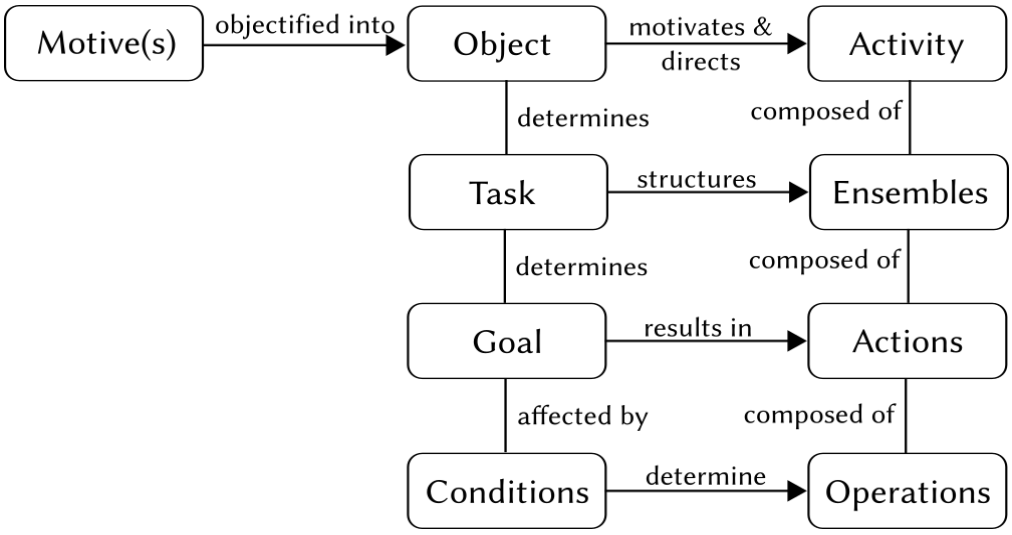

Figure 1: Levels of Activity and the relationship to Motives and Object. Adapted from Wilson [178] to incorporate motives [93] and ensembles [73].

2.3.2 Applications of Activity Theory in HCI

Activity Theory can be used in a number of ways in HCI research, something that Clemmensen *et al.* found in their 2016 review [41]. At a basic level, there is a distinction between attempts to use Activity Theory to support the design of new tools or systems, and attempts to use it in the analysis of how existing systems are used.

As an example of the use of AT to inform design, Kaptelinin's UMEA system [92] attempted to apply concepts from Activity Theory to the design of a virtual work environment by providing explicit support for "project contexts" that encompass file storage and personal information management facilities. Lightweight aggregation of activity-related resources was based on observation of actions undertaken in other tools such as a word processor – so-called "interaction histories".

Bødker & Klokmose [29] apply AT to explain how the relationships between artefacts in "artefact ecologies" evolve over time as people learn more about devices' functionality and affordances; specifically in the appropriation of iPhones by twelve users. They identified a three-stage process of *anticipating, exploration* and *equilibrium.* They further suggest that artefact ecologies themselves move between *unsatisfactory*, *excited* and *stable* states. Most interestingly, they relate these stages to the "web of activities" that the device is situated within; and the primary role of AT within their analysis is to provide this broader notion of activity beyond specific functionalities of the devices within the ecology.

As touched on previously, González at al [73] applied Activity Theory to the study of two IT departments in order to understand how their activities were structured; this study provide empirical data to inform the addition of *ensembles* to AT, and the authors conclude that such a concept is particularly useful in HCI where previous accounts have struggled to *"establish adequate borders for the notions of action and activity as proposed by the hierarchical framework."*

Döweling *et al. [53]* also extend Activity Theory, by presenting a model *"that takes both the physical/technical and the human/social context explicitly into account and provides an integrated perspective on the different levels of analysis in AT."* In doing so, they explain how AT can help to understand "context-of-use" within the ISO 9241-210 user-centered design process, and expand Engeström's model of Activity Theory to more explicitly account for physical and human context. Their notion of context seems to encompass some aspects of the *Conditions* construct provided in other accounts of AT (e.g. [90,178]).

Almalki *et al. [9]* use AT to structure analysis and discussion of quantified-self activity, and conclude that these activities encompasses two distinct types of action: Health Management (HM), and Data Management (DM) in pursuit of health outcomes. This is probably the closest application of AT to the wellbeing domain; although the quantified-self community is not necessarily representative of fitness-device users more generally, and the proposed framework considers only health outcomes rather than broader (or values-based) motives that might be associated with more subjective notions of wellbeing.

Nardi *et al.* [134] draw on AT's concept of Object to help understand blogging as an activity. They identified five Objects among their participants: updating others, expressing opinions, seeking opinions and feedback, thinking by writing, and releasing emotional tension. They debunk the idea of "blog-as-diary" in part by considering these Objects, despite the apparent similarity between blog posts and diary entries, and make design recommendations; illustrating the value of considering the Object of Activity and not just the artefacts that are involved.

Of particular interest is the use of Activity Theory by Barr *et al.* [20] to situate findings from their investigation of how values are expressed and interpreted in the role-playing computer game *Fable*. Based on Rokeach's definition [153] of values as pertaining to modes of conduct, they suggest that *"we might consider 'scoring points' and 'completing the game' as being specifically game-related values"*. They point out that the values which can be expressed through game conduct are influenced by the control system; noting that *"the player can certainly kill a bandit to express their valuing of justice and also of achievement, but they cannot, for instance, try to talk the bandit out of their wicked ways."* They identify two important values: *progress* and *amusement;* and identify in-game actions relating to each. These observations are then grounded in Activity Theory, leading to a brief discussion of how the overall Activity of playing Fable could be framed. Although they speculate on the Object and motives behind doing so, they do not identify specific elements in the collected data. The values of *progress* and *amusement* that are important in the game are not necessarily the values that motivate the Activity itself.

Activity Theory's approach, in which tools are viewed as mediating means between a subject and the Object of activity has also been applied to reground HCI's notion of affordance; which is discussed in detail in the next section.

2.3.3 AT & Affordance

Affordance is an important concept in HCI, and relevant to discussion and exploration of how people adopt and use technologies for different purposes. Kaptelinin and Nardi argue

[94] that the sociocultural approach of activity theory can improve our conceptualisation of affordance and make it more directly applicable to the challenges in HCI than the original Gibsonian definition. In doing so, they provide an account of affordance, and contribute to an activity theoretical model of tools that proves useful in a discussion of values, technology and practice.

Gibson originally defined affordances in the context of the environment (for he was an environmental scientist), saying that *"the affordances of the environment are what it provides or furnishes, either for good or ill"* [52]. Although broadly this ecological definition seems to encompass the human-made environment, and by extension individual artefacts within it, it is fairly non-specific and has been supplemented, in HCI, by additional definitions that are more design-centric. Robert St. Amant provides a formalised account of affordances along with a description of the broad conceptualisations that have followed the original definition. He suggests that affordances have been conceptualised as *"relationships or properties of relationships," "actions," "perceived properties,"* and "*mental constructs"* [11]. Reframing affordances as perceived properties is particularly interesting, as it suggests that affordance is, to some degree, in the eye of the beholder; and not merely a property solely of the environment or artefact itself; although design could, of course, aid in creating such perceptions.

Kaptelinin and Nardi consider the compatibility of contemporary HCI notions of affordance with the original Gibsonian definition and argue for a "mediated action" perspective that considers the three-way interaction *"between the person, mediational means and environment"* [94]. In doing so, they suggest regrounding (not merely extending) the Gibsonian notion of affordance, away from the *natural* environment, towards an explicitly sociocultural context containing deliberately (and socially) designed *tools* and technology. One implication of this regrounding is that an object's affordances are no longer independent of an agent's needs, but instead emerge from them, *"the mediated action perspective points that the dynamic nature of tools make it possible to adjust a tool to a situational need."*

This leads to the development of several more specific forms of affordance. Of particular interest for our purposes is the notion of *instrumental affordance.* An instrumental affordance is a possibility for acting *through* a tool or technology on an external object. An instrumental affordance can itself be decomposed into an *effector affordance* (what change a tool can be used to effect in the world) and a *handling affordance* (the ways in which the tool itself can be operated). This is, essentially, an elegant theoretical account of why (for instance) hammers have two ends – a handle and a head – that, in *combination*, allow a user to effect change in the world around them. We interact with tools, whether

simple physical devices or complex ICTs, as a means of interacting indirectly with something else.

In total, Kaptelinin and Nardi identify six types of affordance, which I have summarised in Table 2. These different types of affordance are referred to later.

Affordance	Description
Instrumental	A way in which a person can act *through* a tool to bring about change in the world, by mediating between a user and an object. An instrumental affordance can be decomposed into a handling and effector affordance.
Handling	A handling affordance is a way in which a person can interact with a tool itself; for instance, a scroll bar provides a means to interact with many user interfaces.
Effecter	An effecter affordance is a way in which a tool can bring about change in an object; for instance by cutting, or sending a message.
Aggregation	Some artefacts can be combined with others to form a larger-scale compound mediator; such as the combination of mobile phone with headphones to form a personal music player.
Learning	Learning (itself a second-order mediation) is required in order to access the other affordances of a tool. A learning affordance provides the opportunity to learn, by providing information. Examples include tool tips, or the USB icon that informs users about the possibility of inserting a USB device.
Maintenance	Technology often requires maintenance, and so must provide affordances for doing so. Examples could include a mobile phone with the facility to change the battery or sim card.

Table 2: Six types of affordance identified by Kaptelinin and Nardi in [94]

2.4 Summary of Background & Research Questions

Prior work on values [153,159,160] emphasises their importance in understanding how humans form preferences about outcomes and conduct, and therefore suggests their central importance in design. Values' fundamental role in design is reinforced by Dunne & Raby when they describe their collaboration with Michael Anastassiades as designing *"a collection of electronic products that that intentionally embodied values at odds with those we would expect from products today" [55].* Barr *et al. [20]* demonstrate how values can be used as a lens for analysing and understanding action in video games, and how value expression is enabled (or not) through a user interface.

The importance of considering values is underlined by work from persuasive technology, where ethical issues such as power imbalance [24,147], semantic variance [162] and the issue of whose interests are being served through persuasion become particularly apparent. Introna's work on the ethics of technology [87] reinforces that values are "*folded into*" the technology that we design, build, and use.

Value Sensitive Design already provides a design approach that considers values and seems promising in promoting ethical technology design that is conscious of, and deliberately attentive to, the values that design takes into account; but as Oosterlaken [140] notes, these efforts are often restricted to the design of technical artefacts rather than the broader "contexts of use" that proponents such as Friedman suggest it could be applied to [68]. Activity Theory, as a theory that attempts to explain purposeful human activity has some promise (as Barr *et al.* demonstrated) in considering these broader contexts. As Qin notes, "*Activity theory proposes that the activity itself is the context'* [148]; and AT's application in numerous HCI studies demonstrates its value in making sense of how tools and artefacts – like blogs [134], or health logging [9] – are used within purposeful human activities.

I propose that, by combining Activity Theory with Value Sensitive Design, the latter might be more easily applied to the broader "contexts of use" or socio-technical systems that could benefit from a values-first approach to design. In this spirit, the goal of the work presented in this dissertation is to situate values within Activity Theory, to understand where and how they are relevant to human activity, in order to support the use of Activity Theory within Value Sensitive Design.

To this end, I pose three specific research questions:

Summary of Background & Research Questions

1. **Where, and how, do values fit within Activity Theory's model of human activity?**

   Intuitively, values seem to have a role as motives; but as Barr *et al.* show through their work on Fable [20], values may have a role within an activity (progression) that is not the ultimate driver of the activity (such as enjoyment, socialisation). Understanding how values interact with different concepts from Activity Theory will help designers who are applying AT in VSD exercises to more quickly identify the role of values within the overall activity system that they are designing for.

2. **How does technology, used within human activity, interact with values?**

   I am particularly interested in the ways that technology itself – as a mediational means – interacts with the values that are implicated at different parts of the activity system; which leads to question three:

3. **What are the design opportunities for technology to better respect, or leverage, values as part of broader activity?**

   Knowing how technology relates to values at different points in an activity system, are there design strategies or guidelines that might be inferred to aid VSD?

Answering these questions will situate values within Activity Theory, particularly with respect to how they relate to the technology that is embedded in activity. By doing so, it is hoped that Activity Theory may provide a rich account of context that can be used, in conjunction with VSD's value-oriented design methodology, to consider values at a broader context of use level.

The rest of this dissertation is split into three parts. The first considers Privacy – a Value in itself – and in particular the way that Privacy is involved with a very common online practice called Third Party Tracking. The second considers Wellbeing, and the activities that people engage in with respect to their wellbeing – most often physical health and fitness – as a broader domain in which multiple values might be at play. Finally, based on the primary studies presented, I draw together findings to answer the three questions posed above; and then conclude with a summary of the contributions.

3. Exploring Privacy in Third-Party Tracking

Privacy is, basically, the idea of being left alone. It can be conceptualised as a *value*, and as a legal or moral *right*. In the 21st century (at least thus far), most conversations about privacy revolve around what we might call *informational privacy* – the right (or desire) to control who knows what about us.

Crucially, privacy can also be framed as a *value*. Privacy is not a need in the biological sense, but is clearly a property that many people feel is important in itself. In many jurisdictions, the importance of privacy is – as many values are – enshrined in citizens' fundamental rights.

The general approach of this section is to take, prima facie, Privacy as a value and to consider the ways that it is (or frequently isn't) embodied in the technology around us, how that is reflected in users' concerns, beliefs and behaviours, and how we might be more respectful of privacy, as a value, when we design and build new technology.

Privacy has become a major policy area for governments; concerns about privacy range from commercial misuse of personal information through to concerns about the rights of the individual with regards to the state, and (for instance) the legality and desirability of mass interception for the purposes of national security. Privacy concerns may form the basis of regulation that is designed to protect fundamental rights, or (as in the case of the European Union's original Data Protection Directive) to provide a solid legal basis through which data can be used and shared for the benefit of businesses, while at the same time providing the reassurance and rights that are required to empower the data subjects themselves. To this end, successive policy interventions, in numerous jurisdictions, such as the EU's e-privacy directive and the upcoming General Data Protection Regulation (GDPR) have attempted to regulate the use of personal information, with varying degrees of success.

It could be argued that privacy is a job for technologists or regulators; that we should build Web browsers which – in response to widespread consumer concern – "turn off" the now-pervasive tracking of individuals by advertisers, or pass laws that prevent, for instance, employers from using Google to conduct background checks on new employees. Under the European data protection regime, though, citizens are given specific rights to

access and correct the data held about them, but they may also give permission to just about any use imaginable if they so choose. There seems to be little political will to regulate specific data uses, and it is doubtful that such an approach could keep up with technological developments even if policy makers tried. Likewise, the myriad ways to track an individual user on the Web make purely technological responses fragile, and once data has been willingly shared with an organisation it is hard to conceive of a technological approach to ensuring that the data is only used for a specific set of purposes. Instead, the European policy approach requires, legislatively, that data processors comply with certain principles, such as consent, but essentially leaves the decision about exactly which data practices to permit with individuals and broader European society. This approach is reinforced through an increased reliance on consent in the GDPR; empowering users through choice, but putting a greater emphasis on their own understanding of the process, intent and potential consequences of the data processing that they consent to; and, in the absence of direct regulation, an awareness of the steps that can be taken to limit practices such as third party tracking which fall (for jurisdictional or other reasons) outside of the requirement for data subject consent.

In short, the European model empowers citizens through rights, but also, to some extent, requires that they take responsibility for their own data. It provides a regulatory environment in which individuals are given the opportunity, and responsibility, of protecting and enacting their own values through the data sharing decisions that they make.

Privacy, as a domain, is particularly interesting in that – unlike wellbeing – the challenges that we face have been brought about largely, maybe entirely, by human-made systems.

Despite efforts within HCI to improve the understanding and control that individuals have over their own data, user concern over privacy is still rife and, as I will expand upon through the research presented, understanding of practices such as third party tracking is low despite its ubiquity and the concerns that are expressed by users when it is explained to them.

i Third-Party Tracking

The studies presented in this thesis are each fairly small scale, but cover a number of different aspects of a common online practice that I refer to throughout as “tracking”. Essentially, this is a practice whereby websites use persistent identifiers to aggregate information about individual web users’ browsing activity. On a small scale, this technique can be used by a website to identify repeat visitors – an important metric when understanding audience and website performance. On a larger scale, though, it can be

used by "third-party" websites such as advertising networks to aggregate information about a user's browsing over a number of different first-party websites, potentially allowing those third parties to collect detailed logs of what an individual looks at on the Web, and to use those logs to infer things about their demographics, income, interests or even health.

This third-party tracking (TPT) is a complex network of stakeholders, comprising advertisers, content providers (who display ads in return for profit), the brokers that link the advertisers with the advertising space and – of course – the Web users themselves, who are tracked and profiled in order to optimise the brokerage process, getting the right advert in front of the right person.

Third-party Tracking raises many of the issues that are present in the broader privacy debate, making it a good choice as a case study for exploring notions of privacy. It raises particular issues around visibility (the extent to which it can be perceived or noticed by users), transparency (the extent to which it can be understood and interrogated by users) and control (the extent to which choice can be exercised by users).

The pervasiveness of tracking, and the advertising model that it supports, also makes it an appealing case study. It occurs on the majority of websites, and is something that virtually every Web user is subject to, if not actually aware of, and, to some extent, takes place "in the open". Unlike the private networks of Facebook or Google, many aspects of third-party tracking on what we might think of as the "open" or "classic" web can be *detected* through the appropriate instrumentation of a Web browser. The HTTP requests to a third party website are made by a user's own Web browser, and the cookies that are commonly used to set persistent identifiers are stored within it. This is in stark contrast to a system like Facebook, where data is shared with third parties "behind the scenes," and where a persistent identifier is provided not through a pseudonymous property of the device (like a unique number in a cookie) but by virtue of the user being logged in to a long-lived account that's associated with detailed information about them.

ii Structure of this Section

The studies here examine third-party tracking 'vertically' through investigations that cover the technical aspects of tracking as well as peoples' attitudes towards it. After an initial exploratory study, the Cookie Visualiser, which helped to identify some key features of the domain, the studies consider several aspects of the tracking ecosystem. From the structure of the tracking infrastructure itself, through the interfaces that have been deployed (as a result of legislation) to inform users about tracking practices, and the understanding and concerns that users actually have when confronted with tracking.

Finally, a research probe called the Web Mirror, is used to consider how we can help Web users to make *sense* of the tracking that they are subject to, and how they respond to options that would allow them to exercise some control over it.

3.2 Privacy Background

There are, broadly, two key areas in the existing privacy literatures relevant to the work presented here: 1. Definitions of privacy, individual and societal expectations of privacy and discussions of its importance, 2. Studies on the relationship between technology and privacy.

Given the huge amount that has been written about privacy and the Web, I restrict this section to a discussion of the broad strokes of privacy as a concept – in law and in the abstract – and how privacy might be framed as a value. I then move on to consider in more detail *online tracking* itself, which serves as a useful case study for considering privacy.

Privacy is, typically, framed as a right that individuals have to control their personal space or communications. Privacy in a non-digital world intuitively has a spatial component – the right to exclude others from what we might consider "private space" such as our home, or to choose who is party to a particular conversation. These spatial concepts both point towards an informational concept of privacy, in which we protect not only space but information or knowledge. We would generally consider it to be an invasion of our privacy if a stranger were to take photographs through the windows of our home even if they did no physical damage and even if we were not within the home at the time those photographs were taken. Photographs convey information about our home, and therefore about us, to the outside world; they potentially move *information* from a private place out into the public.

The concept of a legal right to privacy in US law is generally traced to Warren and Brandeis who, writing in 1890 [173] in response to intrusions by the press into individuals' privacy, declared that *"For years there has been a feeling that the law must afford some remedy for the unauthorized circulation of portraits of private persons; and the evil of the invasion of privacy by the newspapers, long keenly felt, has been recently discussed", "of the desirability – indeed the necessity – of some such protection, there can, it is believed, be no doubt."* Warren and Brandeis deliberately distance the right to privacy from the concepts of libel or slander that deal with reputational harm and instead draw upon the fact that *"the common law secures to each individual the right of determining, ordinarily, to what extent his thoughts, sentiments, and emotions shall be communicated to others."* Fundamentally, Warren and Brandeis' argument is that privacy itself is worthy of protection regardless of whether any other subsequent harm arises from its violation.

A detailed account of the legal development of privacy is not within the scope of this work, but the value that privacy has within our society is clearly evidenced by the legal principles that have developed around it, not least the inclusion of a "Right to respect for private and family life" as Article 8 in the European Convention on Human Rights [44].

In his 2014 book "Internet Privacy Rights" [25] Paul Bernal frames privacy as a means for protecting individual autonomy and sets out four key rights that "represent principles that could be reflected in laws, and understood as the basis of laws." These rights are: 1. a right to roam the internet with privacy, 2. a right to monitor those who monitor us, 3. a right to delete personal data, 4. a right to an online identity.

Although a traditional view of privacy effectively defines two spheres – the public and the private – Andreas Birkbak argues that we need to consider *"the public and the private as something that is continuously connected*" [26]. Intuitively, the existence of many spheres to which an individual discloses different information or affords different types of access is a more accurate assessment than merely "public" or "private". A sphere might be as small as a single person or as big as everyone on the internet. A single private sphere (for instance, one's doctor or paramour) might permit knowledge of a small amount of highly personal information (about a person's health or particular sexual preferences) as well as intimate physical access, but not be party to one's more general private home life. Indeed, multiple private spheres might be entirely distinct from one another, and it may be quite important to keep them that way.

The existence of multiple private spheres can lead to a so-called "context collapse", in which two spheres are brought together – often inadvertently. Context collapse – and ways of preventing or managing it – is discussed in the work of Stefanie Duguay around LGBTQ young people [54] and earlier work by, for instance, danah boyd [33][123]. As boyd points out, *"the value of imagining the audience or public is to adjust one's behaviour and self-presentation to fit the norms of that collective."* Thus privacy, despite often framed as a value in itself, is often a means to deal with the divergent values of different groups or audiences – to present oneself according to the values held by one's audience.

3.2.1 Third Party Tracking

As mentioned previously, *Third-party Tracking* (TPT) refers to the practice of third-party organisations tracking user visits to another website, the first-party. TPT is a mechanism through which advertising networks, market analytics or other businesses can obtain information about a particular web user by partnering with multiple first-party websites and using a common identifier (such as an HTTP cookie) across those first parties. It is

commonly facilitated by devices such as visible content (advertising, or "widgets" such as Facebook or Twitter's sharing buttons) or hidden "web bugs" [122]. TPT is highly prevalent on the Web today, and is instrumental to the online advertising industry which uses the collected data to target adverts based on inferences about users and therefore increase the effectiveness of online advertising [22,61].

TPT is particularly interesting from a values perspective in that it is typically largely invisible to users, but, nonetheless, users express discomfort at the practice when it is explained to them. It represents, therefore, a value-relevant practice that users can exercise some control over. Indeed, Value Sensitive Design has previously been applied to Web browsers in order to devise interfaces that are sensitive to the privacy concerns that relate to the use of cookies [67].

TPT occurs in numerous scenarios. Web cookies are a common tracking mechanism, but users can also be tracked by means of data stored in their Web browser cache, HTML5 local-storage, E-Tag data, Flash locally-stored objects (LSOs) [16], or through the long-lived unique IDs provided by many mobile devices [9]. These techniques are *stateful*, in that they rely on data (state) that's stored on a user's device. Still more *stateless* tracking mechanisms exist, which do not rely on state held on the user's device but instead "fingerprint" a device based on, for instance, clock skew, installed fonts or IP address [51].

In many cases, TPT is facilitated by co-operation between first and third party sites (for instance websites that receive income from a third-party advertising network) although security vulnerabilities such as cross site scripting (XSS) or cross site request forgery (CSRF) can allow third-parties to track user visits to first party websites without the consent or co-operation of the first party website [124].

Responses to third-party tracking can be broadly categorised as either technical or legal. Technical responses include tools designed to help users understand or manage the TPT that they are subject to. Legal responses include data protection and privacy regulations.

Technical responses often take the form of Web browser extensions. Some such add-ons are focussed exclusively on third-party tracking (for instance Ghostery [60] or TACO [3]) while others focus on related concerns such as JavaScript blocking (NoScript [120]) or blocking advertisements (Adblock Plus [141]). However, such tools do not enjoy universal or even majority use, and can be difficult for users to configure or understand [46].

Recently, two high-profile policy responses to TPT have emerged. The first, "Do Not Track" (DNT) [172], provides a technical mechanism by which users can signal their wish not to be tracked. However, unlike the browser extensions mentioned previously, DNT does not disable the tracking mechanisms themselves but relies on third-parties

complying with the users' wishes. At present, such compliance is voluntary but future compliance could be mandated through legislation.

The second is the European Union's so-called "e-Privacy" directive [184], which targets "stateful" tracking by requiring that site operators get informed consent from users before storing or retrieving data (such as cookies or Flash LSOs) from a user's device. While the implementation varies between EU member states, websites in the UK now routinely seek user consent to cookie use.

3.2.2 Regulation and Policy

Privacy and data protection are societal issues that have received a great deal of attention from governments and policy makers. As mentioned previously, the European Union and the United States both have regulations around the use of personal data in some circumstances, although they take fundamentally different approaches. The United States takes a sectoral approach, in which specific data uses – such as healthcare [2], education [36] or video rentals [1] are regulated directly. Europe, as discussed, takes an umbrella approach in which all personal data is governed by the same data protection principles and under which citizens have rights to access and correct the data held about them, as well as recourse to national regulatory bodies when data is held without justification or rights such as the fair processing or security of the data are not upheld by data controllers. The European Union does have some sectoral legislation, for instance the ePrivacy directives of 2002, 2006 and 2009 [184] that regulate the electronic communications sector and – in particular – create the requirement for explicit opt-in consent to receiving electronic marketing mail as well as the much maligned "cookie law" that requires consent from end-users before non-essential cookies are set by websites.

Globally, privacy and data protection regulations typically build on principles from the OECD's 1970 report on privacy [139], including a transparency principle that requires data controllers to provide information to data subjects about how and why data is collected and processed.

This requirement for transparency is manifested as a requirement that data controllers (to use the EU term) make available privacy notices – or privacy policies – to inform people about how they process personal data. In the EU, Data Protection Authorities such as the UK's ICO provide extensive guidance on what such notices should contain [86].

The effectiveness of these notices has been called into question by research, though. For instance, many privacy policies are complex and require a high degree of literacy to interpret [88] and Reidenberg *et al.* found that users' interpretations of privacy policies

often differ substantively from those of legal experts [150]. This challenge seems to be about more than just the inaccessibility of long strings of legalese, though; McDonald *et al.* found that a range of alternative privacy policy modalities did not markedly increase comprehension or acceptability of privacy notices to users [126]. More generalized critique of this "notice and consent" model comes from, for instance [117], fundamentally asking whether a focus on transparency or self-management of privacy is a workable approach to privacy regulation.

More recently, legislative requirements for "Privacy by Design" [155] aim to mitigate privacy concerns at the point that technologies are designed, rather than leaving it to users to make decisions about later on.

3.2.3 Privacy & Bounded Rationality

As in other areas of behaviour, bounded rationality has effects on how people behave with respect to their privacy. This is most commonly expressed in terms of a "privacy paradox" – the observation that people will reveal more information when asked to do so, than they indicate they *would* reveal when asked to predict their own response [138]. More generally, people seem to engage in practices that risk their privacy, and fail to take actions to protect it, even though they say that they are concerned about their privacy. Taddicken finds, for instance, that privacy concerns have only a small negative effect on German internet users self-disclosure on social media [166].

Acquisti *et al. [4]* provide a fairly comprehensive review of the privacy paradox and offers three themes that connect insights from across behavioural science: uncertainty, context-dependence, and malleability of privacy concerns. They note the cultural nature of privacy concerns, but also the high degree of variability within individuals between different contexts and over time. As a result of these factors, they suggest that policy approaches based solely on transparency and control are likely to be ineffective.

As Baek [18] notes, the "privacy paradox" observed in many polls and surveys seems to arise because, for many respondents, opinions are an instantaneous response to a questionnaire rather than being based on more thorough consideration of the decisions that they are being asked about. Thus, many respondents in research that exhibits the privacy paradox have unstable preferences, failing to take into account the benefits of sharing personal data when asked, in the abstract, about sharing it; but conversely failing to take into account privacy implications when offered a benefit in return for personal data. Alternative measures such as deliberative polls provide more stable assessments of privacy attitude; a critique that echoes that of Lipman [117].

The privacy paradox has methodological implications for research (as Baek notes), but also reminds us that the behaviours we engage in can be counter to our own beliefs about how we should, or how we think we should, behave; and draw our attention to the extent that values like privacy do *not* shape our activities in the ways we would like or expect.

3.2.4 Summary

Privacy is established culturally and legally as a value worthy of protection, in itself, and as a means of achieving other ideals such as self-direction, and as a means of avoiding harms like discrimination.

While technology often fails to respect people's notions of privacy, and despite efforts to redress the balance through policy measures, the end result seems to be veneer of accountability and legitimacy with little meaningful impact upon individuals' capacity to make decisions about technology use based on their values.

As in other areas of behaviour, bounded rationality seems to play a role in privacy-related activity, leading to a disconnection between people's conscious and deliberate attitudes towards privacy and the ways that they do, or don't, preserve it through their actions.

Finally, third-party tracking – as the immediate context of interest – is a largely invisible and technically complex practice that Web users become entangled in, often unknowingly, as a result of using the Web. It provides an interesting space for exploration: the practice is very widespread, happens without deliberate user intervention or even their knowledge, yet raises privacy concerns when it is explained to them. Those factors provide an interesting design space, and an ideal scenario to investigate how privacy – as a value – is, or isn't, manifested through the actions of users that become involved with the tracking system.

3.3 Exposing Tracking to Users through the 'Cookie Visualiser'

One of the more obvious (and perhaps objectionable) properties of third-party tracking is the extent to which it is largely invisible. It's quite possible to use the Web, be tracked by many different organisations, and to have no knowledge that it's happening. Much has been made of this invisibility, and tools such as Mozilla Lightbeam aim to make it more visible to users. The underlying assumption being, apparently, that visibility of these practices might afford users greater control over them. The purpose of this first study was to explore what happens when, simply, we implemented a mechanism to confront people with tracking.

To do so, I developed an add-on for the Firefox browser, (imaginatively) called Cookie Visualiser (CV). CV shows users a graphical representation of the third-parties that are associated with the Web page they're viewing. It uses a simple learning heuristic to flag commonly encountered third-parties that therefore have greater access to the user's web browsing information.

The goal of CV was essentially to act as a research probe to begin exploring several aspects of users relationship with tracking. Specifically, we wanted to investigate; 1) participants' existing awareness of, and aversion to, being tracked online; 2) their responses and attitudes towards tracking when it is revealed to them; 3) how users think tracking works and what they think are the possible implications of online tracking; and 4) what design constraints might exist on tools that aim to make tracking more visible and controllable to end users.

> *This exploratory study was conducted in the Summer of 2012 at Microsoft Research, Cambridge; in collaboration with Natasa Milic-Frayling. Some parts of this chapter are taken from an unpublished manuscript that was submitted to CHI 2013.*

3.3.1 Hypotheses

We framed elements of the research questions as hypotheses that would allow us to test the assumption that making tracking more visible would encourage and enable users to control tracking.

At the most basic level, we expected CV to increase our participants' awareness of tracking.

H1: **Using cookie visualiser will increase user *awareness* of tracking and cookie use.**

Given the existing literature that suggests many users find tracking creepy, or otherwise objectionable, we also hypothesised that using Cookie Visualiser, which exposes the *extent* of tracking to users would increase their aversion to it.

H2: **Using cookie visualiser will increase user *aversion* to tracking and cookie use.**

In order to assess the impact of the add-on designs with regards to these hypotheses, we focussed on three aspects and attempted to measure them using pre- and post-study questionnaires:

(1) Cookie Knowledge: We used 16 true/false questions to assess the objective cookie knowledge of each participant. The responses to these questions were summed to assign "cookie knowledge" score to each participant.

(2) Awareness of and (3) Aversion to Third Party Tracking: Participants were asked to indicate their agreement to statements relating to their *perceived extent of* and *attitude towards* third-party tracking. The answers were used to assign scores to each participant indicating the scale of third-party tracking that they perceived (awareness) and their averseness to being tracked. These scores relate to H1 and H2, i.e. the perception of the level of "collusion" between first-party websites and third-party trackers.

3.3.2 Qualitative Research Questions

In addition to testable hypotheses, we had a series of qualitative research questions that could not be framed as simple rejectable hypotheses. These research questions were investigated during interviews with some of the participants at the end of the study period.

The qualitative research questions cover two angles. First, is the Cookie Visualiser as a tool with instrumental affordances for participants:

- Do participants understand what CV is showing them?
- Do participants use CV to try and control tracking by deleting cookies?

Secondly, is CV as a research probe to catalyse a discussion of tracking itself.

- What do participants' understand about tracking?
- What are participants' attitudes towards tracking?

3.3.3 Methodology

We used a mixed-methods between-participants protocol to test the qualitative and quantitative effects of receiving real-time information about tracking on participants. Participants were divided into three groups C, V & G. Group **C**, the control, used the standard Mozilla Firefox browser with default settings, group **G** installed an add-on called Ghostery, and group **V** installed the Cookie Visualizer. All the participants also installed a logging add-on that would collect information about the Websites that the participants visited during the study, and the third parties that were engaged by each one. CV and Ghostery are described in greater detail shortly.

The participants were provided with a short explanatory document about cookies and their potential use as a tracking mechanism, and a document about their respective browser add-on. Participants in group **C** were provided with information about Firefox's built-in cookie manager.

The participants were asked to use Mozilla Firefox with the respective browser add-on for a 10-day period. Where possible, participants were asked to use the computer that they use at home, even if they use that computer at both home and work. We chose to study Web use in the home since we assumed that users are more likely to engage in privacy-sensitive Web activities such as social networking, e-commerce or seeking health information at home.

Finally, we conducted semi-structured interviews with members of group **V** to obtain a qualitative account of their experiences using CV as well as more general attitudes towards tracking and their conceptualisations of privacy.

3.3.4 Apparatus

We used two interventions during the study, the Cookie Visualiser and Ghostery.

Cookie Visualiser: The Cookie Visualiser (CV) displays a series of coloured vertical bars in the bottom left hand corner of the browser viewport, superimposed on top of the Website itself (Figure 2). Each bar represents a distinct domain that was involved in the currently displayed Web page – including the "first party" domain (the domain of the URL that the user is viewing) as well as any "third party" domains, such as advertising networks, analytics providers or other third-party content.

The height of each bar represents the number of cookies set by the domain and the colour of the bar changes to indicate the number of times the domain has acted as a third party. Thus, a common third party, such as large ad network, eventually becomes indicated by a

red bar while a less common third party such as a site-related content delivery network remains green. Very frequently observed third parties are also indicated by an exclamation mark.

Ghostery: The ability of CV to learn which third-parties the user has encountered most often, and to change the colour of each bar accordingly, differentiates it from the Ghostery add-on. While Ghostery displays a list of the third-parties that have been detected, it does not provide an indication of how often each of those third-parties has been encountered on other websites. Unlike CV, Ghostery provides a global list of cookies that have been accumulated by the browser and easy access to additional information about third parties that could help to inform the user's attitude and response to the tracking.

3.3.5 Results

Participants covered a range of age groups and levels of experience with computers and the Web. Of 26 recruited participants, 18 completed the study. All the participants were new to the use of the cookie/tracking visualisation add-ons.

i Awareness and Aversion

We scored the pre and post-study questionnaires to produce an *awareness* score, and an *aversion* score for each participant before they used the intervention, and afterwards; shown in Table 3. These scores were computed as the sum of the chosen answers on 3 and 7 of the questionnaire statements, respectively; the awareness score is therefore between 3 and 18, and the aversion score is between 7 and 42, for each participant.

	Awareness			**Aversion**		
	Pre	**Post**	**Δ**	**Pre**	**Post**	**Δ**
Control n=6	15.2 (1.9)	14.5 (2.7)	-0.7 p=0.51	34.2 (7.5)	34.3 (7.7)	+0.3 p=0.84
Cookie Visualiser n=4	12.5 (3.7)	12.8 (3.2)	+0.3 p=0.91	32.0 (9.8)	34.8 (5.0)	+2.6 p=0.44
Ghostery n=5	14.2 (3.6)	15.0 (1.4)	+0.8 p=0.60	29.4 (7.3)	31.0 (6.7)	+1.6 p=0.20

Table 3: Mean pre- and post-intervention awareness and aversion scores by group, with (S.D.)

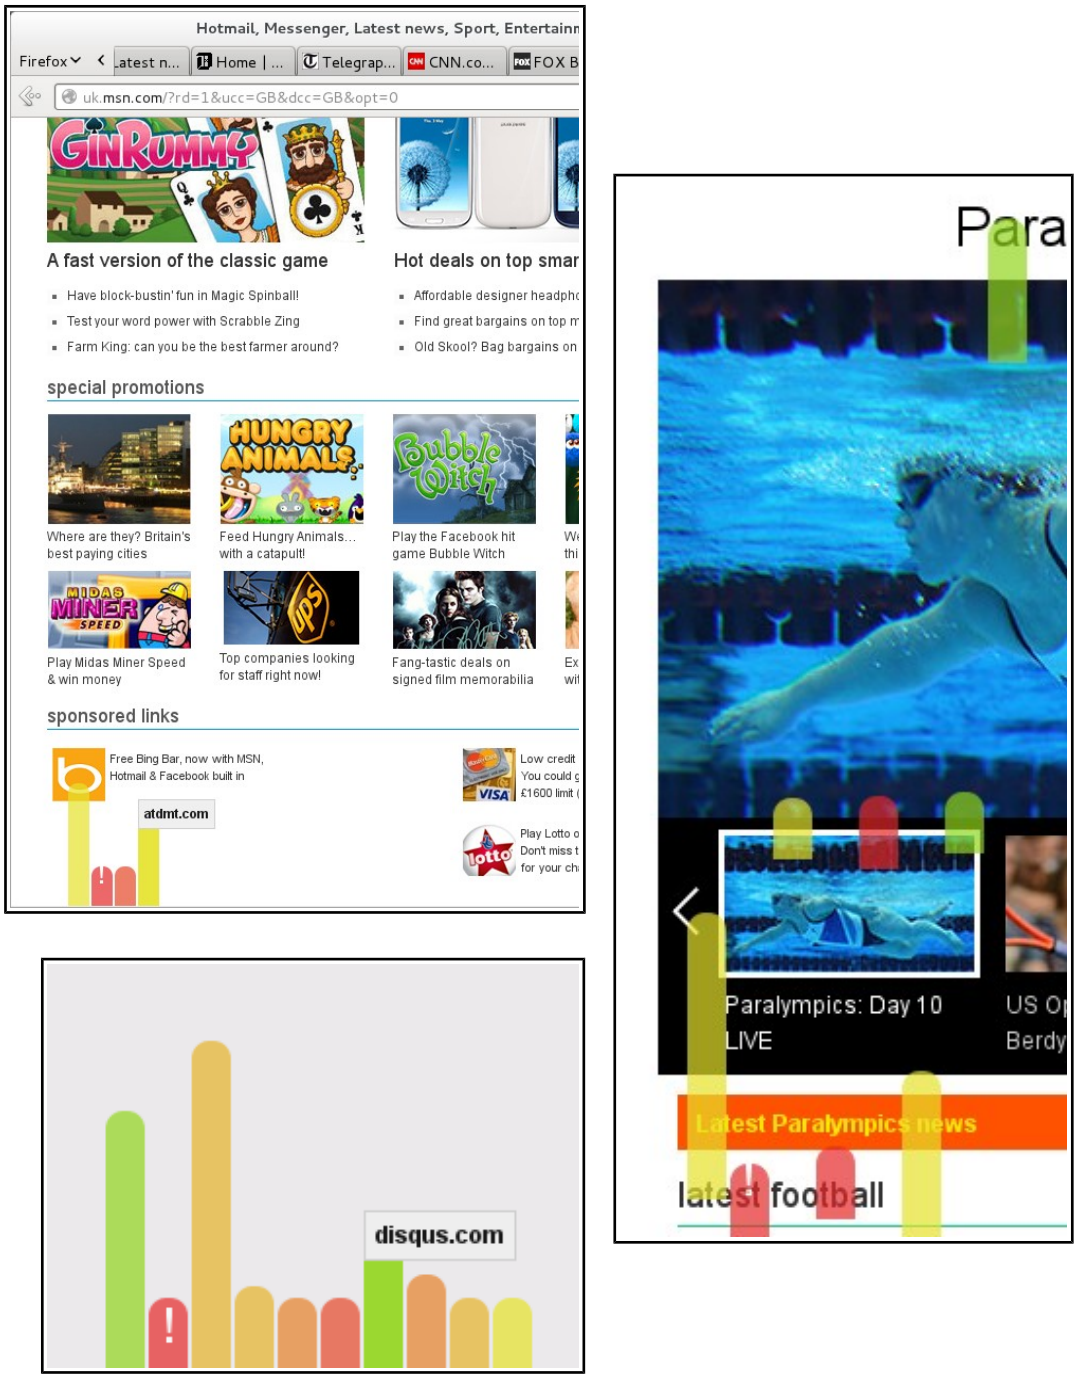

Figure 2: The Cookie Visualiser is shown in the top bottom left hand corner of the browser viewport (top left) and consists of a set of coloured bars (bottom left) that represent third party domains. To attract attention, bars drop from the top of the viewport once the page has loaded (right).

Aversion and awareness increased for all groups, except for the control group where Awareness decreased. We conducted paired-samples two-tailed T-tests for each pair of pre-post scores; none of the changes were significant even prior to correcting for multiple tests, and so no correction coefficient was applied.

ii Cookie Knowledge

We observed no change in participants' cookie knowledge in the pre- and post-study questionnaires, as shown in Table 4. This is surprising, as we had included a description of what cookies are how they are used in the information that was given to participants, after the pre-study questionnaire. This points, perhaps, to the difficulty of conveying technical concepts in a manner that leads to a persistent understanding of their possible uses and effects. Significance was, again, tested using paired-samples two-tailed T-tests.

	Cookie Knowledge		
	Pre	**Post**	**Δ**
Control n=6	13.3 (2.0)	13.2 (0.81)	-0.1 p=0.81
Cookie Visualiser n=4	12.5 (2.4)	12.0 (1.8)	-0.5 p=0.50
Ghostery n=5	13.6 (1.5)	13.2 (1.6)	-0.4 p=0.62

Table 4: Mean pre- and post-intervention cookie knowledge scores by group, with (S.D.)

3.3.6 Interview Results

The quantitative results from the questionnaires indicated no change in the users' understanding of cookies, nor in their perceptions of or averseness to TPT. The lack of change in understanding is somewhat surprising since the participants were provided with information about the use of cookies. This raises further questions; first, whether the information was adequate and comprehensible, and, if so, whether the user can reason about the implications and whether those implications are a concern. The finding also suggests that empowering users to control their own exposure to surveillance may not be as simple as providing educational resources.

We did find a statistically significant correlation between the amount of TPT perceived by participants and their aversion to TPT (Table 1). This invites further explorations of these two user attributes. Given that neither CV nor Ghostery affected understanding,

perception or averseness, we have focussed further investigation on the qualitative data about user experiences with CV obtained through interviews. We focused on three key aspects of CV:

Did the participant recall noticing CV at any point during the study?

Participants indicated that they were generally aware of CV during the study, with some commenting on the "dropping" behaviour as attracting their attention. No participants indicated that they did not recall noticing CV at all.

Did the participant recall interacting with CV in any way, such as hovering or clicking on it, for instance to find out more information or close CV?

Responses were generally negative, with some participants indicating that they tended to focus on their current task rather than look at CV. The majority of participants indicated that they always used the Web with a task in mind. Some participants indicated that they hovered over CV to find out the domain name, but that they had not clicked the bars for more information. Participants that hovered over the bars indicated that knowing the domain names was not very useful as they did not recognize the companies involved. This suggests that the information provided by Ghostery, which includes the company name (and other information) may be more relevant to users. One participant indicated that they did not interact with CV because they had assumed it was "just part of the study," and may have interacted otherwise.

Did the participant feel like they understood what CV was displaying through the colour, size and number of bars?

Some participants were unable to explain the purpose of the colours or height of the bars, although others alluded to the level of tracking by various sites. Participants generally understood that more bars indicated a greater level of tracking on a particular website. In general, the height of the bars seemed to confuse participants.

These qualitative findings confirm that the type of information provided through add-ons, and CV in particular, is not adequate to affect user understanding. Furthermore, the number of cookies per domain is not necessarily a good indicator of the level of TPT on a particular website. The height of the bars could instead be used as a second indicator of how frequently a third party has been observed.

3.3.7 Discussion

The study revealed a couple of critical design issues. Task focus, that the participants' focus on the Web task they had in mind deflected from giving attention to their concerns about privacy and CV itself; and Scale, that seeing the scale of tracking through aggregate visualisations elicits a strong reaction from people.

i Task Focus

While the majority indicated that privacy was a concern, the focus on the task deflected attention away from information provided by the add-ons. We hypothesize that, unlike, for instance, receiving an invalid security certificate warning, which one might encounter once a week, the visualization and warnings of third-party cookies occur constantly and are not *necessarily* indicative of surveillance or threat to privacy. In addition, privacy concerns around TPT are, perhaps, better described as "chronic" concerns – a constant, long-term concern – that is rarely, if ever, sufficiently acute enough to interrupt the user 's immediate task. TPT therefore demands a different notification approach to acute events such as invalid security certificates or hazardous situations in the physical world.

Models such as the Elaboration Likelihood Model [145] predict that preoccupation (task focus), and lack of urgency (from being a constant low-level concern) decrease the likelihood of deliberate "central route" engagement. As a result, the possibility of deferring reflection and decision-making about tracking – and associated aspects such as cookies or Web browser settings – might hold more promise than a real-time or just-in-time approach.

ii Scale

The invisibility of the cumulative effects of surveillance and the sheer *scale* of the TPT networks pose another challenge. When asked, users clearly indicated their concern about privacy and the information that others hold about them. Yet, the lack of transparency around TPT seems to dampen the appreciation of the impact that Web surveillance can have. Thus, making explicit the connection between TPT and the user's own privacy is a key challenge that CV does not appear to have addressed adequately. The aggregate visualisations that we presented to participants during the interviews *did* elicit more of a reaction, and seemed to prompt a greater sense of aversion.

It seems from the results that while CV was able to demonstrate ubiquity, it did not convey scale to participants. Being able to infer and appreciate the scale or structure of the tracking network by seeing the connections to individual websites – as shown in Ghostery and CV – is non-trivial.

3.3.8 Conclusions

This study confirmed that people dislike the idea of being tracked, but that designing tools to provide awareness of tracking, either for its own sake or as a precursor to enabling action, needs to contend with at least two key challenges. First, that web users are engaged in other tasks, which are often mentally-demanding or being done in the course of employment. Second, that the scale of the tracking ecosystem is what prompts most concern in people, but is very difficult – or impossible – to appreciate when information is shown on a site-by-site basis.

These two design challenges point towards using aggregate visualisations during dedicated time. Techniques such as slow technology, which deliberately create time for reflection, could be a more promising approach than real-time interventions which constantly struggle against the user's main task for attention. Happily, the idea of using an aggregated view of the tracking network is entirely consistent with such a non-real-time approach; data can be collected during browsing, then aggregated and displayed in a more reflective period.

The two studies presented next – of the Tracking Networks themselves, and of a reflective tool called the Web Mirror, pick up on these two aspects. First, by developing a better understanding of tracking as a network; and second, by developing a reflective tool called the "Web Mirror" that embodies on the findings from CV in an aggregate non-realtime reflective tool.

3.4 Analysing Third-Party Tracking Networks

A key finding from the Cookie Visualiser study was that using an aggregate view of third-party tracking seems to make the potential privacy impact more appreciable. Furthermore, the use of a network visualisation, that shows the degree of interconnection between first and third-party websites seems to foreground these concerns in a way that is easily understood.

In fact, the network aspects of TPT seem to be fundamental to the concerns that participants have about privacy. This is apparent from the interviews that we conducted, and seems to be intuitively true as well. Informational privacy is undermined when data is shared about us, and an invisible network of websites that collects and shares information about our web browsing seems to embody these concerns.

Whilst the cookie visualiser study identified subjective concerns among our participants, it provided little in the way of objective, empirical insight into the tracking networks themselves. The purpose of this next study was to understand more about the structure and extent of the tracking networks, with a view to better understanding, at an objective level, how they might interact with users, and users' personal data, to give rise to subjective concerns relating to informational privacy.

It makes little sense to think about values in isolation from the world around us. If values form the basis of our preferences about things then what actually *exists* around us, the environment in which our activity takes place deserves some interrogation, too. How will it respond when we operate within it? What externalities will arise as we do so? The purpose of the study presented next was to understand more about the structure of the tracking ecosystem, with a view to better understanding the possible implications for people who use the Web.

> *This chapter reports research that was conducted in collaboration with, and is based on a paper written with, Natasa Milic-Frayling, Eduarda Mendes-Rodrigues and m.c. schraefel [72]. Some parts are included verbatim from that manuscript, while others have been modified to align with the specific focus of this thesis.*

3.4.1 Methodology

Search engines are essential for accessing relevant content on the Web. Commercial search services such as Google (google.com), Bing (bing.com), and Baidu (baidu.com), process millions of queries and serve millions of customers each day. In this study, we used search engines as a mechanism to collect a sample of websites for inclusion in a later analysis of the tracking networks associated with each one. We considered search results collected from Google, Bing, and Baidu based on a series of search terms.

i Search Queries

For search queries we considered a subset of broadly accessible data from the KDD Cup 2005 Challenge [183]. The KDD Cup task involved automated categorization of queries and the organizers published a set of 800 training queries that had been categorized by three human labellers. We selected a subset of the labelled queries, aiming for a consistency in the assignment of queries across the categories and the label accuracy. We excluded labels with no agreement among the labellers and discarded the corresponding queries. The resulting list contains 662 labelled queries. The categories comprise a two level hierarchy with top categories: Computers (8 sub-categories), Entertainment (9 sub-categories), Information (8 sub-categories), Living (18 sub-categories), Online Community (6 sub-categories), Shopping (6 sub-categories), and Sports (11 sub-categories). Only one second level category had a further sub-category. Each query was assigned between one and four labels from a set of 67 distinct second and third level categories. A list of labels associated with 30 or more queries is shown in Table 5. We hypothesize that search results related to the queries in different categories will lead to the retrieval of different types of Web sites which, in turn, may have different marketing practices and involve different third party entities.

ii Search Results

For each of 662 queries we gathered the top ten search results from both Google and Bing. We used the Bing search API from the Azure Marketplace to obtain Bing search results while for Google we extracted URLs from the search result pages. The process was carried out four times to collect results for four English-language search markets: the United States, the United Kingdom, South Africa, and India. Search within a specific market was conducted by providing the appropriate location identifier for the Bing API and by altering the search engine domain for Google (see Table 6).

Some search queries returned zero results for one or more of the search engine-market combinations and were excluded from the set for the sake of consistency. That resulted in

the final set of 659 search queries, with ten search results for each of the eight search engine-market combinations—a total of 5,272 sets of search results comprising 9,776 unique Web domains.

In addition to the English-language results, we collected search results from Baidu in the Chinese market by considering popular queries, as published weekly by baidu.com. We compiled a set of 98 queries and then collected the top 10 search results for each of those queries. Although this data set is smaller, it enabled us to analyse the characteristics of the tracking practices in a non-English language market.

iii Crawling

In order to collect information about cookies and third parties associated with each URL in the search results, we used the Selenium Web browser automation framework to automate visits to each set of search results. We used the Firefox browser v.16.0 on multiple Linux machines to access Web sites in parallel, and each browser instance was controlled by a Python script. A new browser instance was spawned every 15-30 seconds, based on the real-time resources that were available on each computer. After an initial pilot phase, we applied heuristics to manage the crawling processes—browser instances that did not complete their crawl within 5 minutes were automatically terminated.

A new Firefox "profile" was created to visit each set of 10 search results, in order to provide a clean browsing environment. We installed a custom logging add-on to record the referrer header of all HTTP requests, and to log cookie creation, updates, and deletions. No other add-ons or plug-ins were enabled within the Firefox browser and its "do not track" feature was not enabled. Each URL from each set of search results was loaded by the browser in the ranked order, starting with the first result. At the end of the crawling session, the browser was directed to a web service to deposit, via means of a web form, a log file from the add-on that contained cookie and HTTP referrer information.

Category Label	Number of Search Queries
Shopping\Stores & Products	101
Information\Local & Regional	95
Information\Companies & Industries	60
Living\Health & Fitness	49
Living\Car & Garage	41
Information\Law & Politics	40
Living\Travel & Vacation	39
Living\Fashion & Apparel	37
Information\Science & Tech	36
Living\Finance & Investment	34
Living\Food & Cooking	33
Information\Education	30

Table 5: Distribution of search queries across the top 12 categories, ranked by the number of queries used in the experiments.

Market	Bing API Identifier	Google Search Domains
India	en-IN	www.google.co.in
South Africa	en-ZA	www.google.co.za
United Kingdom	en-UK	www.google.co.uk
United States	en-US	www.google.com

Table 6: Bing market identifiers and Google search domains used during crawling

iv Analysis Methods

We conducted two types of analysis on the results from the crawling:

1. An analysis of the distribution of tracking domains across search results and topic categories, and
2. an analysis of the referral network involving Web sites retrieved during search and the third party domains associated with them.

To that effect, for each set of search set results (9 sets in total), we analysed the referrer headers of the browser's HTTP requests as the search result pages were loaded into the Firefox browser. We collected a list of all the Web domains that appear as search results and those that are referred to as third parties. We noticed that some Web domains, such as facebook.com and twitter.com, appear as both the Web sites in the search results and third party domains referred to by other Web sites. That motivated us to differentiate among four types of Web domains:

Web sites—Web domains whose pages appear among search results and are not referred to by other sites. They are thought of as the first party only domains.

Third party only—Web domains that are referred to by Web sites or other third party domains and never appear among search results nor refer to other domains. Such are, for example, googleanalytics.com or ad services that place ads directly on the Web pages.

Dual role—Web domains that appear as both first party and third party domains. An example of which is facebook.com which appears among search results and is referred to by sites that include Facebook integrations such as "like" buttons.

Ad Exchange Service—Web domains that appear only as third parties, i.e., do not appear in search results, and refer to other third party domains. They are intermediary third parties that provide a bridge between Web sites and other third parties involved in ad bidding.

From the referral header information we produced an edge-list representing the 'refer to' relationship and created directed network graphs in which A→B corresponds to the fact that the domain B is 'referred by' domain A when the page from A is loaded into the browser. As noted above, in many instances, the domain B is the ad network or ad exchange service that delivers the adverts to the Web page. In other instances, the referral is due to the fact that the Web site is running a script that refers to a monitoring service,

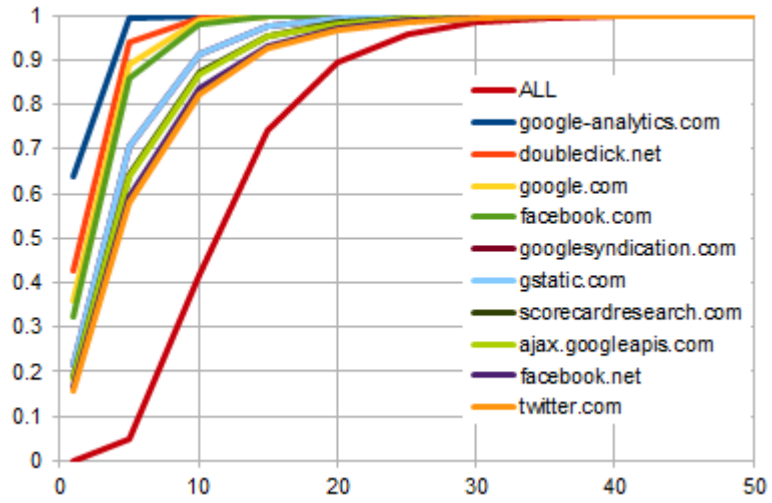

*Probability of encountering tracking domains while browsing results of **Google** in the US Search market*

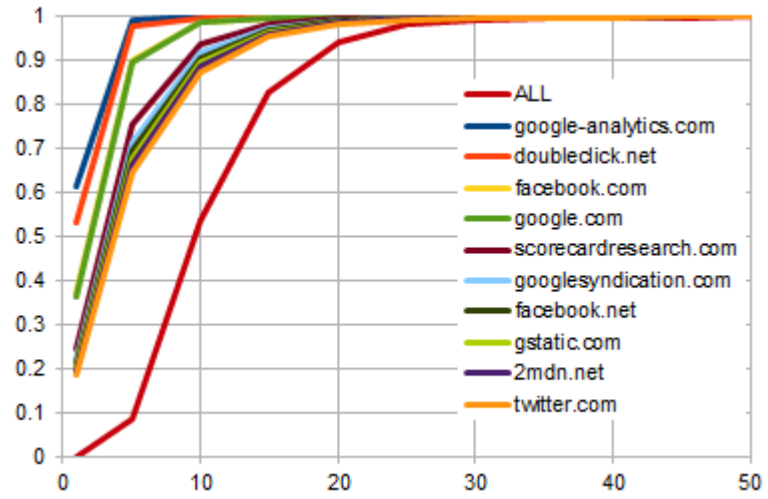

*Probability of encountering tracking domains while browsing the results **Bing** in the US market*

Figure 3: When browsing search results, the user is exposed to tracking domains. We calculate the probability that the user encounters top tracking domains while visiting a random set of search result pages for both Google-US (top) and Bing-US (bottom)

collecting statistics of the Web site usage. We loaded the edge and node lists into R[1] and NodeXL[2] in order to calculate graph metrics and visualize the referral networks.

3.4.2 Analysis

Three areas of analysis are of particular interest (more are reported in the full paper). First, we analysed "exposure" to tracking – the likelihood of a user being tracked as they visit search results. Second, we analysed how tracking varied between the original search query topic. Finally, we analysed the structure of the observed tracking networks.

i Exposure to Tracking

In order to estimate the rate at which users are exposed to third parties, we estimate the probability P(T) that a search result exposes the user to a third-party T by calculating the proportion of search results that refer to T. We rank third parties based on P(T) and, for the top 10, determine the likelihood that the user will encounter each of these parties after accessing a number of retrieved search results. We make two simplifying assumptions. First, we assume that any Web page from a given Web site is exposing the user to the same set of trackers. Second, we expect that the user's choice to visit a search result is independent from the previously seen pages. Based on this model we observe the probabilities that a user would have encountered all top ten third parties. We find that after visiting just 30 search results, the probability of getting cookies from all top 10 third party domains is 99.5%. Figure 3 shows the probabilities of encountering the trackers when using Bing and Google search engines in the US Market. Our analysis of the Chinese search market indicates that the user exposure rate may be even higher than in the English markets. However, the Chinese dataset is smaller and not directly comparable.

1 R: http://www.r-project.org/
2 NodeXL: http://nodexl.codeplex.com/

ii Influence of Search Topics

Categories associated with the search queries enable us to observe the users' exposure to tracking while accessing information within a specific area of interest. We analysed queries associated with the top ten most commonly applied labels to our query set (see Table 7). In particular, we compute the expected exposure to tracking when the user selects the top ranked search result for a given query. Although basic, this model is instructive. Most of the topic labels led to similar levels of tracking. However, two labels (Information\Science & Tech., and Information\Law & Politics) show lower numbers of both third parties that set cookies and third parties that do not.

Label	Num. of Logs	TPs w/ > 0 Cookies	TPs w/ 0 Cookies
Shopping\Stores & Products	785	2.79 (3.57)	3.65 (3.37)
Information\Local & Regional	726	2.16 (4.26)	3.44 (4.27)
Info\Companies & Industries	459	2.88 (4.02)	3.79 (4.28)
Living\Health & Fitness	362	2.33 (3.54)	3.42 (3.42)
Living\Car & Garage	286	3.11 (4.05)	3.84 (3.98)
Information\Law & Politics	298	0.44 (1.26)	1.23 (1.42)
Living\Travel & Vacation	301	3.10 (4.85)	3.19 (2.93)
Living\Fashion & Apparel	289	3.37 (3.87)	3.91 (3.31)
Information\Scienc e & Tech..	271	1.77 (3.16)	2.07 (2.35)
Living\Finance & Investment	245	3.08 (3.68)	3.99 (4.28)

Table 7: When browsing search results , the user is exposed to tracking domains. We calculate the probability that the user encounters top tracking domains while visiting a random set of search result pages.

iii Network Structure

We analyse the tracking network that was observed in each of the search markets, as well as the aggregate global network. We computed global graph metrics that highlight the similarity of the tracking networks across the search engines and the search markets.

All the networks include a dominant network component that, in some instances, includes more than 92% of nodes and 99% of edges. We compared the tracking networks with synthetic networks based on the Watt-Strogatz random model [136,137]. For each network we generated a graph with the same number of nodes and the same average degree of the nodes. For example, a synthetic network for the global tracking network uses 8.91 for the average degree of a node. We then compare the synthetic and real networks based on their average path and the clustering coefficient. As Figure 4 illustrates, the global tracking network closely follows a small world network model with a rewiring probability of 0.2. We get similar results for the tracking networks affiliated with each individual search market.

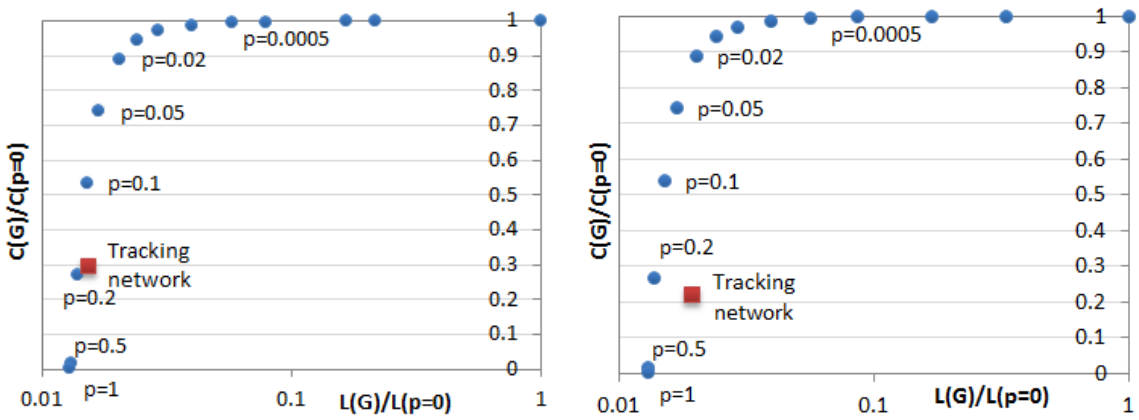

Figure 4: Watts-Strogatz random model, plotted against the average path length L(G) and the clustering coefficient C(G) of left: the global tracking network from English search markets; right: the global tracking network after removing doubleclick.net

As Latora and Marchiori [106] have shown, small world networks have high local and global efficiency in supporting exchanges of information. Thus, the detected tracking networks are well equipped to support a range of processes: gathering and disseminating contextual information about the user, real time processing and aggregating information, and bidding and delivery of adverts. We use the same analysis to assess the robustness of the network. Considering the dominant role of the ad exchange domains such as DoubleClick.net, we ask how properties of the network would be affected should such a node be excluded from the network. As Figure 4 shows, removal of the DoubleClick node would increase randomness due to the longer paths and lower clusterability. This, in turn, would reduce the efficacy of the transactions within the network.

3.4.3 Conclusions

The results confirm that tracking is a widespread practice on the Web. Moreover, they allow us to estimate user exposure and to conclude that a typical user has a 99.5% chance of being observed at least once by each of top ten tracking domains that we found, after visiting just 30 search results

From a design perspective, the near ubiquity of tracking is a key element to consider. As we saw in the Cookie Visualiser study, the inevitability of tracking makes bringing it to users ' attention *more* difficult, because it makes the appropriate timing of such information non-obvious. A notification each time that tracking is observed is quickly normalised by users, and becomes a "chronic" background concern.

The relatively uniform frequency of trackers observed across search query categories is also noteworthy, since it includes categories such as "Health and Fitness" that relate to areas that many people would consider to be sensitive. Indeed, information about a person 's health or physical condition is classed as sensitive personal data by European Union data protection laws, and research such as Google 's Flu Trends (despite poor prediction results at a population level [108,109]) suggests that people do disclose health information to search engines via their search behaviour.

Finally, the network structure itself, which supports efficient propagation of information (in this case, data collected about web users) and which is dominated by a few large nodes such as Doubleclick, has some relevance to user concerns around privacy. A single omnipotent surveilleur seems, intuitively, to be more problematic than a large number of surveilleurs with only restricted access to one 's activity.

3.5 Surveying Cookie Notices in the Wild

The tracking network analysis showed, empirically, the extent and structure of tracking on the Web, and provides a useful model to help understand the potential privacy implications of using the Web to find information. However, it says nothing about how this tracking is currently made visible or explained to Web users. A finding from the Cookie Visualiser study was that the way information about tracking is presented is important; despite being based on relatively simple principles, the emergent tracking ecosystem, and the inferences that can be drawn from that data that is collected, are complex and hard to appreciate. Thus, the way that tracking practices are framed to web users is an important factor in whether they are able to foresee how their use of a website or online service might impact upon their privacy.

Driven by the well-established principle of transparency in privacy and data protection regulation, in 2009 the European Union adopted the so-called "Cookie Law" [184], which introduced a new requirement for services to obtain consent before storing or retrieving information on or from a user's device, except where technically necessary to deliver the requested service. In effect (and by intention) this law introduced the requirement for users in the EU to be notified when cookies are set in their Web Browser for purposes such as analytics or behavioural profiling.

Although the interpretation of the new rules differed between member states, in the UK this led to the widespread use of "cookie notices" to inform Web users about the use of cookies. These notices – unlike privacy policies – are fairly hard to avoid and so the cookie law became perhaps one of the first industry-wide measures to inform users that cookies are in use for various purposes. The purpose of this study was to understand, qualitatively, how that information was being presented visually and as a choice architecture as a precursor to further qualitative research with web users.

3.5.1 Methodology

I used a survey methodology, which involved surveying 300 UK websites and recording information about the "cookie consent" elements that appeared (or did not appear) on each one.

I obtained a random sample of 5385 domain names from Nominet, the UK domain name registrar. These had been selected from a subset of the .UK domain pool, which included only those domains associated with a website. Because the survey was conducted manually, to make the sample size tractable I selected, at random, 150 of those domain

names to include in the survey. The random sample of UK domain names provided a broad picture of UK websites, but it was apparent very early on that the measures taken by these (mostly small business) websites were not necessarily the same as those taken by higher-profile sites. While fewer in number, high-traffic websites account for a large proportion of the interactions that people have with the Web, and so I wanted to ensure that they were represented in the survey. To do so, I also included the Alexa top 150 websites in the sample; giving a total of 300 different websites.

I visited each website in turn, using the Google Chrome Web browser. Installed in the browser was an add-on that I developed (shown in Figure 5) that allowed me to select the "consent element" on the page (if one was present), and which saved a snapshot of the page to a web service. This snapshot included information about the cookies on the page, a screen shot in PNG format, the HTML source code of the page, and the screen co-ordinates of the selected consent element. Conditions such as domain name "holding pages" and other errors were also recorded, so that all websites were accounted for in the final results.

I then took a grounded theory approach to analyse the collected consent element samples, treating them as qualitative data but considering aspects such as their position on the page, and visual aspects, as well as the language that was used. Although qualitative analysis is often thought of with respect to textual data – such as interview transcripts – it can also be applied (and frequently is in disciplines such as media studies) to visual content such as images or video [23]. Applying it to analyse the visual and interactive elements of computer interfaces, as I do here, is not much of a stretch from those other visual forms.

The aim of the qualitative analysis was to identify themes and categories in terms of how consent elements are presented to users, the content they contain, and the choices (or instrumental affordances) that they provide. Having examined the collected samples to identify common themes and categories, I devised a lightweight codebook of categories (described in the results) and then classified the full sample to produce the quantitative results.

As will quickly become apparent from the results; this analysis exercise was fairly lightweight – largely because of the low variation that was observed within the surveyed consent elements. The results of that qualitative analysis are presented next.

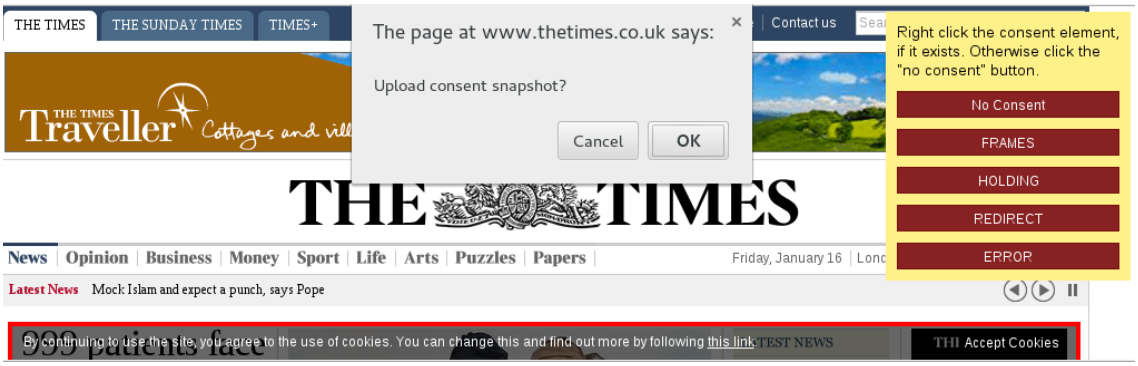

Figure 5: The Chrome extension being used to select the consent notice on thetimes.co.uk

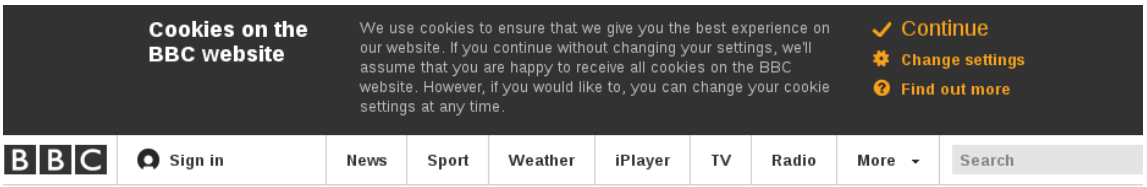

Figure 6: A cookie notice on the BBC website (not included the sample studied) that includes a method for the user to opt out of cookies.

3.5.2 Results

Of the 300 websites in the sample, 290 were surveyed during January 2014. The remaining sites could not be surveyed for some reason, generally due to compatibility problems with the browser tool (such as pages consisting entirely of frames) or failure to connect to a functional website.

In total, 89 sites in the sample exhibited a "cookie consent" feature of some kind (76 of the Alexa top 100 out of 129 that used cookies and 12 out of 74 of the .co.uk sample, excluding domain holding pages and error pages). The number of sites in each category, split by those that use cookies and consent, is shown in Table 8.

The position that elements take on the screen has been shown as an important determinant of attention [165]. Looking at the collected samples, I noticed four categories that describe where on the page the cookie consent feature has been placed. They are described below:

1. **Header**: The consent feature was placed towards the top of the page in an area above what might be considered the main content of the page.

2. **Fold Float**: The feature was placed so that it was attached to the bottom of the browser viewport, so that it remained there even as the page scrolled.

3. **Footer**: The feature was placed towards the bottom of the page, below what would be considered the main content of the page. Typically this was close to links for the privacy policy, and terms of use as well as the copyright notice.

4. **Other**: A small number of features were placed in other locations, usually within the content itself. For instance, one website in the sample placed a link to the cookie policy underneath the new-user registration form on the homepage.

I also considered the form that the elements took on the page, and the way that the choice was presented to users. Three clear categories emerged, covering all of the collected samples.

1. **A Policy Link:** The simplest of the features, consisting just of a link to another page, usually with the words "Cookies", "Cookie Policy" or "Privacy and Cookie Policy". This was only considered to be the consent feature if no other cookie consent features were present.

2. **Ultimatum**: This includes all features to the effect that "This website uses cookies, and by using this website you agree to their use", including features that just said "This website uses cookies".

3. **Justification**: This includes all features where the use of cookies is acknowledged, but some reason for their use is provided. In the overwhelming majority of cases, this justification was vague, for instance "This website uses cookies to improve your experience."

Combining the position categories with form categories produced a three-by-four matrix of 12 subcategories; 11 of which were represented in the collected sample (Table 9).

	Alexa UK 100				**Nominet co.uk Sample**											
Status	OK				OK				Holding Page				Error			
Cookies	Y		N		Y		N		Y		N		Y		N	
Freq	127		2		74		49		10		21		2		4	
Consent	Y	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y	N
Freq	**76**	51	0	**2**	**12**	62	1	**48**	0	10	0	21	0	2	0	4

Table 8: Overview of survey results; numbers that represent a "correct" behaviour are shown in bold

	Policy Link	**Ultimatum**	**Justification**	
Header	6	11	24	**41**
Fold Float	0	2	6	**8**
Footer	31	3	3	**37**
Other	1	1	1	**3**
	38	**17**	**34**	**89**

Table 9: Notice breakdown by theme and position

3.5.3 Discussion

The results show relatively little variation between UK websites, with all falling into a few common interaction patterns. In analysis of these common patterns, I draw on the notions of affordance as mediated action proposed by Kaptelinin & Nardi [94].

In almost all of the surveyed cases, the consent elements had no *instrumental* affordances with respect to cookies themselves. None of them provide a direct means for the user to prevent or remove non-essential cookies. This is why I refer to one group of elements as the "ultimatum" group – the use of cookies is stated, but no means of control (other than leaving the Website outright) is mentioned.

There are some high-profile examples – such as the BBC (figure 12) – of cookie consent elements that *do* have instrumental affordances, but their absence on the surveyed websites suggests that they are few and far between. Of the sites in the survey, only bt.co.uk provided direct controls, within the consent element itself (figure 7); and although the handling affordance is reasonably clear (in the form of a "change cookie settings" link) it's not clear what the associated icons mean, suggesting a missing learning affordance.

A handful of websites – including natwest.co.uk and barclays.co.uk – provided controls within their privacy policies, which are themselves linked to from the consent element. Overall, these compound interactions which consist of an initial informative banner, and then more detailed information with controls, exhibit instrumental affordances over the cookies that are set by their respective websites. Taken in isolation, though, the initial banners provide only learning and handling affordances; alerting the user to the use of cookies, and providing a means to obtain more detailed information and controls, but not providing any control in themselves.

All of the surveyed websites provide a handling affordance, and a subsequent learning affordance via their links to more detailed information. Many of the linked policies provide information – directly, or via yet another link – about how to control cookies via a Web browser's built-in tools.

Linguistically, the consent elements displayed a high degree of ambiguity; as evidenced by the distinction between the "ultimatum" and "justification" groups. In the case of the ultimatum group, no explanation for the use of cookies was given. In the case of the "justification" group, though, some attempt to justify the use of cookies was included. For instance, "We use cookies on this site to enhance your user experience". Other common justifications included statements such as "we use cookies to help make this website better." Aside from a linguistic justification for the use of cookies, the ultimatum and justification groups are very similar. There is an obvious lack of clarity over what phrases like "improve your experience" might mean; it could be construed as reference to the

collection of analytics that allow decisions over the Website at an editorial or structural level, or it could refer to specific personalisation features, or even to functionality that is technically dependent on cookies, although strictly those cookies fall outside the scope of the legal consent requirement created by the ePrivacy directive.

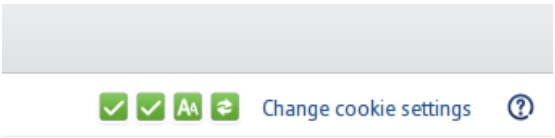

Figure 7: The cookie controls provided on bt.co.uk have instrumental affordances, although they lack learning affordances that would allow a user to understand what they do.

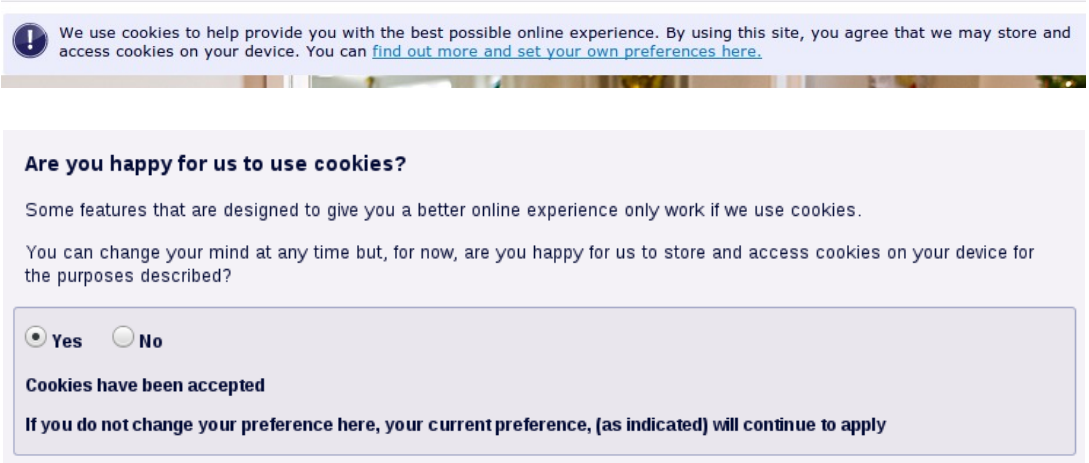

Figure 8: A few sites, including natwest.co.uk (pictured), used consent elements (top) that linked to a more detailed policy which itself included instrumental affordances over cookies (bottom).

3.5.4 Conclusions

The survey found relatively little variation in how consent to cookie use is sought on UK websites, with all of them falling into a narrow range of interaction patterns. Those patterns typically provide no instrumental affordance with respect to cookies, instead being limited to learning affordance and the opportunity to find more information. Where control is offered, in all but one or two cases (BT and the BBC) that control is provided on a separate page.

In many cases, the purpose of the cookie use is unspecified. In the cases where some purpose *is* specified, it is specified in very vague language, in terms such as "improve the user experience" that do not clearly convey whether, for instance, disabling cookies will damage an individual user's immediate experience, or simply prevent the collection of statistical data that might be used to inform changes to the Website for all users at a later date.

3.6 Notice Focus Groups: Understanding People's Perceptions of 'Cookie Notices'

Having identified the scale of the tracking networks themselves, and surveyed the information elements that are being used to inform Web users about tracking, the next step was to understand more about how individual Web users interpret the information that they're shown, and how they reason about online tracking. To do so, we used a series of focus groups to explore, qualitatively, our participants knowledge, beliefs and feelings towards online tracking, utilising a series of statements taken from privacy policies and cookie notices as talking points.

Our goals were to understand more about:

a) what the participants felt that they understood by a series of statements taken from online privacy policies and whether they felt positively or negatively about the statements, particularly the vague justifications such as "improve your experience" that were found in the previous study;

b) whether the participants had noticed cookie consent boxes while using the Web, and what the participants understood by the wording used in some of consent notices that had been found in the previous survey,

c) whether the participants linked the information from privacy and cookie notices with their own sense of online privacy, and

d) the sort of information that participants thought online advertisers might have about them and where they thought this information was obtained from – i.e. their ability to correctly reason about the privacy implications of third-party tracking.

> *Some of these focus groups were conducted and used as part of a mixed-methods study with my colleague Helia Marreiros, published in [121], although their inception and design is my own.*

3.6.1 Methodology

Participants for the focus groups were recruited from interns and postgraduates in the School of Computer Science and the School of Economics by means of mailing lists and personal invitation. We provided participants with pizza during each session to encourage attendance and to create a more relaxed atmosphere for conversation.

Participants were seated around a table with the investigators. The sessions were structured through the use of a set of slides that were projected on to a screen. The slides had three sections:

1. A series of statements taken from online privacy policies. This section was conducted primarily to inform a secondary study that is not reported here.
2. Screenshots of some cookie consent notices from UK websites. We asked participants to explain the reasons that they thought the Website was displaying the notice, what the notice meant the Website would do and what they thought other parts of the notice (including phrases such as "improve your experience") might mean.
3. Two exercises; in which participants were asked to imagine what information a) Facebook and b) a behavioural advertising company, like DoubleClick, might know about them.

We kept detailed notes during the focus groups, as well as audio recording them. We used thematic analysis to identify and group together similar attitudes and beliefs into a set of key themes that had emerged from our conversations, and these are discussed in the results section, below.

Participants spanned a range of academic backgrounds, including social science, computer science and humanities. In total we had 21 participants; 12 male and 9 female.

3.6.2 Results

The results of the focus groups are reported as qualitative accounts of user concerns. Themes were identified by two investigators based on their participation in the focus groups as well as audio recordings. For smaller focus groups (2-3 participants) we also used paper handouts for the later exercises, providing a physical record of the data that participants thought companies hold about typical users.

i Privacy Policy Statements

Participants were mixed in their responses to the individual privacy statements. Given the extent of the services provided by Facebook and Google, participants were unsure which features would involve the different types of data processing. For instance, one participant asked whether particular statements were applicable if they only used Google to search. A similar sentiment was apparent in participants ' responses to statements regarding the transfer of their personal data to the "family of companies" that constitute

Facebook. Participants were unsure which companies this included, and some interpreted it as very broad range of affiliates. Some mentioned companies such as Snapchat, Instagram and WhatsApp.

Participants were generally negative about the idea of processing their personal data in foreign countries. This was seen as unnecessary and potentially risky, some participants commented that they might have less legal protection if their data was transferred abroad.

When we asked participants to comment on purely technical statements, such as explanations of cookies and pixels, they were generally less negative but felt that the purpose of their use was important.

ii Cookie Notices

All but one of the participants indicated that they were familiar with seeing cookie consent notices on websites that they visited. However, all admitted that they did not pay attention to the notices and the majority did not understand what the function of cookies was.

One participant (a computer science undergraduate), gave a fairly complete explanation of how cookies could be used to track a user across multiple visits to a website and another (who had previously been employed by an advertising company) was aware that cookies could be used to find out demographic information about a person, but was unsure about the technical mechanism for doing so.

When asked what they thought the notices meant, participants were often unable to suggest how cookies could fulfil a purpose such as "make this website better" or "improve your experience". Typically, though, they interpreted this as personalization, for instance by remembering previously visited pages to personalize navigation. A few expressed that the intent was to collect analytics through which the Website could be improved in general rather than made to work better specifically for them, but those participants were in a minority.

Participants were confused about the difference between cookies and browser features such as auto-complete and browser history. A number of participants spoke about the "private browsing" feature of modern Web browsers, as a way to avoid being tracked if they wanted to do so, although it was unclear to what extent they made use of this feature themselves.

Two participants knew that websites were required to disclose cookie use by law, although none mentioned the specific EU directive or UK legislation.

A number of participants spoke about the "private browsing" feature of modern Web browsers, although it was unclear to what extent they made use of this feature themselves. Interestingly, the simple description of non-persistent browser state as "private browsing" foregrounds the feature's value-relevance in contrast to the typical technology-centric terminology used by other Web browser features and typical explanations of cookies.

Most participants erroneously believed that a behavioural advertising companies such as DoubleClick would have data such as email address, address and postcode via, for instance, orders that people placed online, and did not raise the use of cookies as a way to obtain partial browsing history and potentially infer demographics or interests from that information.

The consent notices that we showed to our participants do not appear to have conveyed the purposes of cookie use. Participants were unable to explain how cookies worked, the real purposes for their use (for instance, how the use of cookies might improve their experience on the Website) and were unable to relate the use of cookies to their own privacy.

When prompted to consider the potential privacy aspects of cookies by a notice message that specifically mentioned privacy, participants questioned how cookies might relate to their privacy, but again the understanding was incomplete.

Once notions of privacy or demographic or behavioural targeting had been introduced into the discussion, either by participants themselves or the investigators, there was still no evidence that many participants understood the link between those practices and the use of cookies. Furthermore, when asked, specifically, to identify the types of information that an online advertiser might have about a web user, participants nearly all failed to mention information gleaned through the use of persistent third party cookies.

Some participants felt it was unfair to declare the use of cookies but provide no means to opt-out. In the words of one participant, it is "undemocratic" to provide no means to use a website without being able to reject the cookies. This sentiment does not necessarily seem to be driven by a particular concern over the use of cookies in general, rather a response to the lack of choice in itself.

iii General Findings

In both scenarios – privacy policies and cookie notices – participants seemed to take into account the purpose of the request when articulating their assessment of the desirability of the practice. For instance, cookies that ostensibly "improve" experience are seen more

positively, particularly when participants were able to articulate methods through which they believed this improvement could be brought about, such as personalization.

Cultural context, or specific privacy concerns, seem to be taken into account by participants. This was particularly common among the foreign participants who often related it the cultural context of their home countries. For instance, a Turkish participant spoke of how they felt there were limits on their freedom to criticize the government; a Mexican participant linked posting photos of his dog to the risk of it being stolen for ransom – a phenomenon of which he was familiar from his home town. Another participant spoke of the different risks that male and female internet users might face because of the cultural expectations around gender in his home country.

At the end of the focus groups, and following a debrief session during which we answered any questions that participants had about the issues we'd touched in the focus group, many participants admitted to feeling more concerned about their privacy than before taking part in the focus group. This sentiment was not universal, though. Some of the participants, particularly those who were very relaxed towards the privacy notices that we showed them did not appear to have changed their attitude. Some of those participants expressed that had learned new things about the mechanisms or extent of, for instance, third party tracking, but still did not feel it was a problem to them personally.

3.6.3 Discussion

Participants covered a range of attitudes towards privacy, consistent with previous work showing diverse attitudes (e.g. [102,176]). There is a notable distinction between those who are concerned with privacy in itself, and those who are concerned with privacy as a means to another end. Although this interpretation requires some caution – consequentialist arguments could be a post-hoc rationalisation of why privacy is important to a participant – this poses an interesting issue when considering privacy from a values perspective: Is it more appropriate to design tools from a consequentialist perspective, or from a deontological one? Or, is it possible to do *both*?

However, in both the deontological case and the consequentialist one, there is a degree of reasoning and inference required on the part of users in order to assess the impact of a given activity. In the consequentialist case, participants tried to reason over how simple actions like sharing a photograph or accepting cookies could lead undesirable outcomes, such as the theft of a pictured possession. In the deontological case, where privacy is valued, participants attempt to reason over how an action impacts upon privacy. The use

of a cookie, for instance, might lead to a loss of privacy if it were used for profiling purposes, but most participants were unaware of this possibility, and confused about the technical mechanisms that are involved.

I propose a three-tier model of how our participants' concerns relate to the data processing practices of data controllers, such as websites.

The bottom, most objective level is a *technical* explanation of a practice, such as the use of cookies.

A middle layer, *privacy*, represents privacy as an interim construct; essentially the state of controlling information about oneself.

Finally, a *consequential* layer encapsulates concerns that stem from privacy effects – such as the result of some knowledge becoming public – but which are not in and of themselves privacy concerns; such as having a possession stolen, or being discriminated against. These concerns are highly subjective in the sense that they are often dependent on personal context, though not *necessarily* subjective in the sense of being open to disagreement – most people would agree that being the victim of theft is a bad thing, but not everyone is in a position where loss of privacy would conceivably lead to a theft.

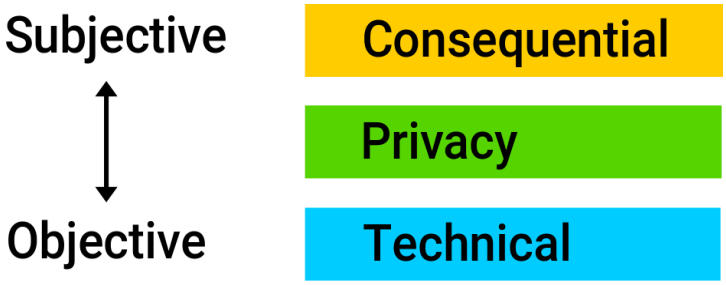

Figure 9: A three-tier model of privacy reasoning

The most important finding from these focus groups, though, is probably that transparency about Technical implementation is not, by itself, enough to allow Web users to make reasoned decisions about privacy – even assuming that they had the time, will and rationality to do so. Using the information that's provided often requires a level of technical expertise and literacy beyond what most people possess. Even *beyond* understanding the words and explanations that are provided, which is in itself a challenge [131], being able to infer the impact of a cookie, or a tracking pixel, on something abstract like privacy, or even to predict a possible concrete harm, is simply too complex. While privacy implications, and possible consequences, might logically follow from a particular

configuration of technology, web users cannot reasonably be expected to make the connection between technical explanations, privacy as an abstract concept, and specific consequences that could occur.

Treating Privacy as an intermediate concept that sits somewhere between a technical explanation of a particular practice and an assessment of that practice 's potential consequences helps to illustrate why considering privacy in and of itself might be useful. Predicting the likely, or possible, consequences of an action can be difficult, if not outright impossible. There is a huge degree of uncertainty associated with most privacy-related decisions both in terms of how the technology and services operate, and also in terms of which events might – or might not – actually occur. Security breaches (such as recent high-profile cases involving Yahoo, TalkTalk, Ashley Madison or Equifax) can have potentially major privacy implications, but service users cannot predict if or when such a breach will occur.

Privacy as an intermediate concept helps to address some of this uncertainty, *and* removes some of the necessity of even trying to predict consequences; by grounding concerns in *Privacy* it is no longer strictly necessary to identify other specific harms. Valuing (and protecting) privacy in and of itself seems to provide (whether the en user is conscious of it or not) some protection against harms that *could* arise from the disclosure of personal information.

The framing of privacy as an abstract intermediate construct between concrete activity and concrete harms is congruent with work, such as that by Rossler [154], that frames the value of privacy as being about individual control and self-determination. It also reinforces privacy's status as a *value;* recall that Rokeach describes values as a conception of the preferable, and Schwartz describes them as criteria for evaluating the selection and evaluation of actions. As a concept that simplifies the otherwise difficult (or impossible) task of predicting the possible consequences of sharing data about oneself; privacy acts as a value against which actions can be judged. Where privacy is valued, actions that preserve privacy should be preferred to those that do not.

3.6.4 Conclusions

The focus groups provided insight into how information about personal-data processing is perceived and reasoned about by users. Of particular note was how participants attempted to reason about the information that we showed them, and the distinction between attempts to reason about privacy (or related values such as choice) and more

grounded attempts to identify specific consequences that might arise from disclosing particular pieces of information.

Adopting privacy as an intermediate construct, and valuing it in and of itself – that is, conceptualising of privacy as a *value* – seems to be helpful in making decisions about sharing personal data less complex, by reducing the need to foresee specific possible consequences.

3.7 The Web Mirror

This chapter reports on a design probe called "the Web Mirror". The Web Mirror is an interactive tool that combines elements of the cookie visualiser and tracking analysis studies, and combines them to create an online visualisation of someone's web browsing history, and the associated tracking networks. It builds on the interaction findings from the Cookie Visualiser study, and the tracking network analysis methods developed in the Tracking Network study.

There were two primary motivations behind building the Web Mirror. The first motivation was largely practical: I wanted to be able to explore the use of visualization with participants in lab studies, or short online experiments (as well as in public outreach activities, or in demonstrations) but existing add-ons that visualise web tracking – like Mozilla's Lightbeam tool, or the Cookie Visualiser tool described earlier – analyse web browsing as it happens; as a result, they have to be used for a period of days or weeks before they yield detailed aggregate information about tracking. The Web Mirror achieves similar results in a matter of minutes; making it far more amenable to the intended uses. The **second** motivation was to experiment with explanations and visualisations of web tracking that conveyed the profiling and targeting aspects as well as the tracking network itself; building on the findings from the earlier Cookie Visualiser study, where my conclusions were that aggregated display of information, in a dedicated reflective activity, would be more effective than intervening in the course of another activity to surface privacy concerns. As explained shortly, the Web Mirror goes beyond the visualisation offered by any existing tools in the way that it analyses and presents tracking information to the user.

I begin with a technical description of the Web Mirror itself, before moving on to discuss how it was used as the basis for a mixed-methods study, and what can be discerned from those results.

> *The inception and development of the Web Mirror are my own, but this chapter reports research that was conducted in collaboration with, and is based on an unpublished manuscript co-authored with, Helia Marreiros, Dorota Filipczuk and m.c. schraefel.*
>
> *Some parts of this chapter are included verbatim, while others have been modified. I am particularly grateful to Helia for her knowledge of statistics. Analysis, discussion and conclusions as they pertain to Activity Theory are entirely my own contribution.*

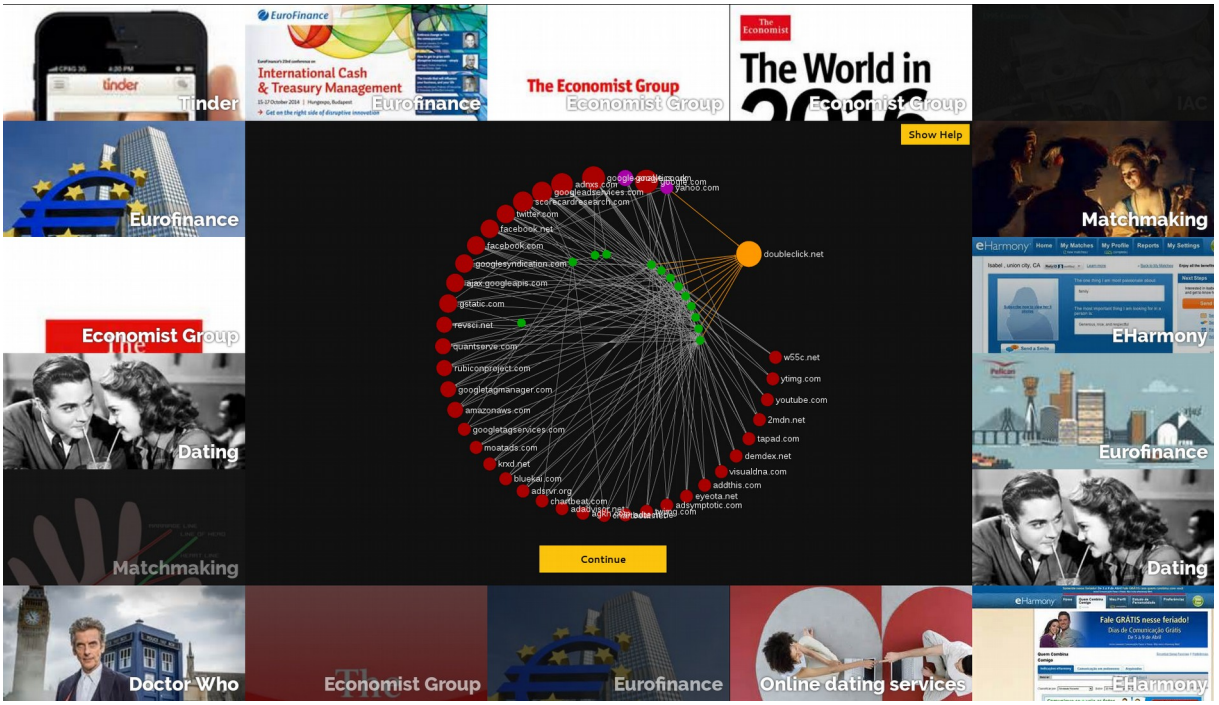

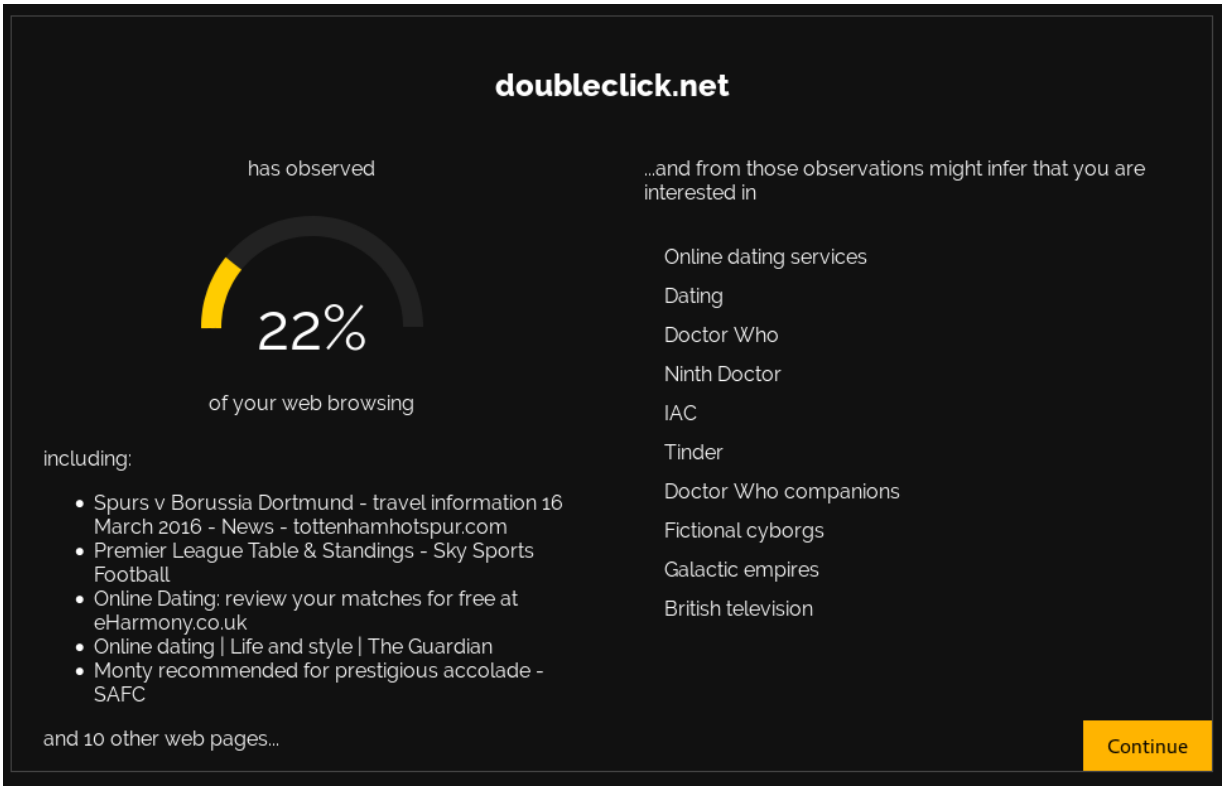

Figure 10: The Web Mirror displays a network diagram, depicting the tracking network detected on the participants' uploaded browsing history, and topics related to the visited pages are shown around the outside of the page (top); detailed information about a particular tracking profile (bottom),

3.7.1 Description of the Web Mirror

The Web Mirror is an interactive tool to explore user understanding, attitudes and reactions to online tracking practices. Recent work by Schaub *et al. [150]* shows that, when reflecting on existing Transparency components that are part of tools like Ghostery and Disconnect, participants are still confused about how tracking works. To address this confusion, inspired by the use of non-technical metaphors like the desktop to represent technical processes, like a file system, we designed the Web Mirror to deliberately de-emphasise the technical representations of tracking such as number of cookies and instead foreground both a more familiar and more personalized, less abstract metaphor of a personal "web reflection". It is designed to promote reflection on the link between web use, tracking, identity, and privacy; in the spirit of Hallnäs and Redström's "slow technology".

That is, by using a person's browser history, the Web Mirror lets a person see:
(1) how trackers are connected to visited websites;
(2) what they may have observed, and
(3) based on these observations, how each tracker might "see" the user.

Figure 10, shows an example of the Web Mirror in use. The Web Mirror is displaying that *adnxs.com* has observed the user visit (among others) *doctorwho-online.co.uk*. This is presented by highlighting the node in the network graph, and then displaying the detailed information (shown in the bottom right image) about the total amount of browsing observed, examples of pages tracked, and the inferences drawn. Interests associated with the profile are shown around the outside of the page as text and images (for instance, "Doctor Who", "Tinder" and "Economist Group").

Another advantage of the Web Mirror is that feedback is almost immediate: unlike other tools such as Lightbeam or F-Secure's Freedome Tracking Mapper[3], the Web Mirror is not implemented as a browser add-on and does not need to be used for a period of time (typically days or weeks) before results about tracking patterns to become available. Instead, it analyses a sample of browsing history that is provided by the participant by copying and pasting from their Web browser's built-in history facility. This property makes the Web Mirror suitable for use in the kinds of lab and survey studies we describe in the next section, without installation and use of additional software. There are three main stages in using the Web Mirror:

3 https://freedome.f-secure.com/whoistrackingme/tracker-mapper.html

1. *Upload*. The first stage is to collect a sample of web browsing history from the user's device. On desktop, this is done by copying and pasting data from the browser history feature. On Android devices, this is accomplished using an app that accesses the browser history API.

2. *Crawling*. We then use an instrumented Web browser to visit each page in the uploaded history, recording the third parties that are connected to each page (via the HTTP referer: header) and the page content itself. This results in a graph of the tracking network itself, and allows us to extract key topics from each page using the OpenCalais[4] service to build a profile of topics for each visited web page.

3. *Analysis and Display*. Using the network graph, we can build a "profile" of the user on behalf of each detected website, including both first- and third- parties. The profile consists of the URL of each visit that the third party observed during the crawl, and the topics that were extracted from that page. Each of these profiles represents an individual "reflection" of the user – our guess of what they look like to each node in the tracking graph. This information is visualized to the user in a desktop Web browser (mobile users are given a code that allows them to access their results).

Prior to the studies developed below, we developed the Web Mirror iteratively based on feedback from users in step-up demonstrations, small focus group sessions and a school outreach program run by the University of Southampton. Participants suggested that the network graph itself was useful in helping to contextualize the individual tracking profiles – corroborating a similar finding from the earlier Cookie Visualiser study. The Web Mirror therefore begins with that network (and an explanation of it), and animates the diagram to highlight which reflections are being displayed. We found that integrating several aspects of each reflection into a single information pane helped to reinforce the link between tracking that had been undertaken by a particular domain, and the interest-inferences that might have been drawn from it. Previously, we had presented statistics about surveillance coverage, pages observed and inferred interests sequentially, but changed this to the parallel presentation shown in Figure 10 which generated more engagement in subsequent studies. Finally, participants suggested that they were unsure of what they could do about tracking. This was manifest either as direct questions, or cynicism that "there's nothing I can do about [surveillance]". As a result, we produced simple countermeasure information, similar to the descriptions given in Figure 11. This finding itself provides motivation for the study presented next; which considers the relative importance of transparency – provided by something like the Web Mirror – and actionable information – like information about available countermeasures.

4 http://www.opencalais.com/

Our preliminary results suggested that with the Web Mirror, participants felt they had gained more understanding about being tracked and were interested in exploring countermeasures. From this preliminary work, we determined the Web Mirror was sufficiently clear to use in a more formal, experimental study, which is presented next.

Browser Extension (eg. Ghostery, Disconnect.me)
Browser extensions can be installed in your web browser (Firefox, Chrome, Internet Explorer). **How it works:** these extensions block the 'requests' (connections) that your browser makes to known tracking websites, preventing them from tracking you. **Pros:** fairly simple to install; blocks most tracking. **Cons:** not available for older browsers; might not be installable on company or school equipment; sometimes need to be configured properly to block tracking (for instance; by default, Ghostery only reports tracking rather than blocking it). **Resources:** http://www.disconnect.me/; http://www.ghostery.com/

Figure 11: A sample of the countermeasure information that was provided to participants during development of the Web Mirror, and in Experiment 1

3.7.2 Formulating the Research Question

The Web Mirror had demonstrated usefulness in schools outreach, and as a catalyst for discussing issues around tracking – similar to the discussions that I'd conducted with focus groups based on the cookie consent notices and privacy policies earlier. However, I – and colleagues – wanted to test whether tools like the Web Mirror, which aim to make tracking more transparent, have a role to play in helping Web users to understand practices like tracking, and to opt-out of those practices if they wish to do so. In framing a question in this way, we draw a distinction a between two types of information that can be provided to people: *transparency information,* and *actionable information;* the distinction between which I'll discuss next.

Tools such as the Web Mirror, the Cookie Visualiser, or Mozilla Lightbeam provide *transparency* – they aim to show people how tracking works, and there is an implicit assumption – including in my own work – that somehow this will be valuable, perhaps ethically, but perhaps also in helping people to make more informed decisions, or to control practices (like tracking) that they find objectionable. Transparency also has an important *legal* role; in EU data protection law, data controllers are required to be transparent about how they collect and process personal data. As research has shown, though, transparency (typically in the form of privacy policies) has largely been

unsuccessful in furnishing individuals with an understanding of, or control over, their personal data.

As one of the participants during the development of the Web Mirror remarked, despite being informed about tracking *"there's nothing I can do about it."* While transparency can help people to understand tracking, it does not necessarily contribute to people's ability to control the things that we are being transparent about. Instead, there exists *actionable information* that describes things that can be done; like installing an ad-blocker, or clearing browser cookies, that provides a plan of action for a person to take.

While transparency might help to build motivation and understanding, it is (prima facie) *actionable information* that suggests an activity (or action within a broader activity) in answer to that motive.

In developing some *actionable information* – in our case a set of information sheets that explain different tracking *countermeasures* – we observed that actionable information needs, for reasons of narrative clarity, to give some reference to the need that that action meets. So, for instance, the ability for an ad-blocker to prevent some tracking from taking place requires a brief mention of tracking. This is only superficial information, nowhere near as detailed or personalised as the information displayed in tools like Lightbeam or the Web Mirror – but it begs the question of just how much transparency is actually required. Do we need tools that make tracking more visible, or could we just tell everyone that ad-blockers exist? More broadly, what is the functional role of *transparency* in enabling the sort of privacy self-management that is required in the EU (and US) regulatory regime, and can tools like the Web Mirror contribute to a critique of transparency as a regulatory concept?

In this study, we tested whether the Web Mirror – as an example of detailed *transparency* – increased (or not) people's intention and desire to take *action* in the form of using an anti-tracking *countermeasure*. We conducted two experiments, with two specific hypotheses related to this research question, which are described next.

3.7.3 Methodology & Hypotheses

Using a between-participants experimental set-up, we investigated whether the Web Mirror had an effect on self-reported *intention to adopt* privacy-preserving tracking countermeasures, and self-reported *willingness to pay* for privacy preserving countermeasures, versus just describing the countermeasures themselves.

Self-reported *adoption* of privacy measures and self-reported *willingness to pay* have been used to study individuals' behaviour in prior research (e.g. [5,17] and [157], respectively). Although they have less ecological validity than measuring *actual* use of or payment for countermeasures, they are considerably easier to measure. Furthermore, they overcome confounds – such as technical ability *to* use countermeasures – that would be introduced by measuring actual adoption or payment. For the purpose of evaluating experimentally whether the Web Mirror leads to additional *motivation* to adopt privacy-preserving tracking countermeasures, these measurements allow two groups to be compared without the overhead or additional confounds that measuring actual behaviour would introduce.

In total, we tested two hypotheses. The rationale for each is described in the experimental design section, below.

H1: Transparency information (in the form of the Web Mirror) plus actionable information will lead to greater adoption of tracking countermeasures than just actionable information by itself.

H2: Transparency information will result in a greater willingness to pay for a tracking countermeasure than among participants who do not receive transparency information.

To test our hypotheses, we designed and conducted two similar randomized experiments in two parts (Experiment 1 and Experiment 2).

3.7.4 Experimental Design

Both experiments were conducted in two phases, where each iteration was a refinement of the main protocol based on the results of the previous version. This approach helped to inform a robust set of results, described below.

The experiments had a between-subjects design, where participants were randomly allocated to one of two treatments, Control (C) or Transparency and Countermeasures (TC) in experiment 1, and Control or Transparency and WebCloak (TW) in experiment 2. In the TC and in the TW, participants were asked to use the Web Mirror and then asked to complete an online survey, described below. In the control treatments (CI in experiment 1 and CII in experiment 2), participants only completed the online survey.

Phase 1 of experiment 1 consisted of a laboratory study followed by focus groups with students from our university. The purpose of this phase was to pilot the experimental design in a controlled environment, and to test aspects such as the resilience and comprehensibility of the Web Mirror prototype, the timing and complexity of the survey instrument, and initial measurement of treatment effects. We led focus groups after each

session to check participants' understanding of the information provided and to understand their views qualitatively.

The laboratory experiment was designed to give robustness to subsequent online experiments, as it is a more controlled environment. Online experiments were used to reach more participants and increase the data set. Therefore, the remaining phases (phase 2 of experiment 1, and both phases of experiment 2) were conducted online.

The online studies were orchestrated through a custom-built website that assigned participants to conditions and administered the study components accordingly. We recruited participants through Prolific Academic[5], a UK-based crowd-sourcing service that recruits participants for academic purposes.

3.7.5 Experiment 1: Countermeasure Information

Experiment 1 had two components: the Web Mirror and a survey about tracking countermeasures. the Web Mirror presented transparency information. The survey (available in the supplementary materials) contained questions about participants' privacy attitudes, demographics, their current use of different countermeasures and about their intent to adopt countermeasures. We presented detailed information about each countermeasure as Actionable information – i.e. information that describes a possible action, rather than information that describes an extant data-processing practice.

The survey had five pages. On the first page of the survey, we asked participants about their online behaviour (such as, how often they browse the Web and click on ads they see), their beliefs about their own knowledge of key concepts (such as, "I know what online tracking is") and their existing attitudes ("I would like to do something to prevent some or all online tracking (even if I don't know how)"). Next, we asked them to indicate which of the suggested countermeasures they presently use. This list of countermeasures covers a range of tools and manual behaviours. At this point we did not provide any information about any of the countermeasures. Notably, we excluded the private browsing feature of Web browsers, as it effectively combines other entries in the list such as clearing cookies and browser history. Also, we included "clearing browser history" as it was a technique that participants had frequently mentioned to us in earlier work, even though it does not actually counteract tracking, and we wanted to check how widespread this practice already was among our participants. We were clear about its ineffectiveness in the provided information.

5 http://www.prolific.ac

The third page presented detailed information about how each of the countermeasures worked, how it is used and/or installed, and the pros and cons of each countermeasure. Figure 11 shows an example of the information provided for Browser Extension, such as, Ghostery and Disconnect.me.

Participants were then asked to indicate which of the nine countermeasures they intended to use. They chose from a five-point scale: "Definitely going to use"; "Going to use"; "Might use"; "Not going to use" and "Definitely not going to use". We wanted to capture undecided participants in a dedicated category, to get an indication of how strong their intent was, and hence chose five points over four.

We also asked them to report anything that they had learned during the study. This question was to evaluate whether they felt they had learned about tracking and/or about options to avoid the tracking i.e. the countermeasures.

The fourth page contained questions to evaluate participants' level of concern about privacy. We used an instrument proposed by Chellappa and Sin [37], which evaluates participants' desire for privacy against the value derived from personalisation, and so overcomes some of the philosophical problems associated with instruments that attempt to measure privacy by itself, without reference to the value that might be derived in return for *losing* some privacy. Finally, we asked demographic questions and if participants felt more concerned about their privacy after participating in the study.

The manipulation condition was the presentation (or not) of the Web Mirror. In the CI treatment, participants only completed the survey, without seeing the Web Mirror at all. In the TC condition, participants were required to provide a sample of their web browsing history and to interact with the Web Mirror for at *least* three minutes.

i Phase 1: Lab Experiment + Focus Group

Participants were invited to take part in the study via email and social media. During online registration, they were required to upload a sample of their web browsing history (at least 70 unique URLs), which was pre-processed and used to populate the Web Mirror. During the study session, each participant was allocated to a computer terminal. After reading the information sheet and signing a consent form, the experiment started. Participants were allocated, at random, to one treatment. To protect our participants' privacy, we did not store the uploaded history along with any identifying information, and deleted it after 72 hours.

Those in the TC treatment were first presented with their Web Mirror reflection, and then asked to complete the survey. Those in the CI treatment completed the survey first, and

were presented with their Web Mirror reflection afterwards, to enable them to participate in the focus group that followed.

After the individual study, participants were taken to another room where pizza was provided and the investigators led a semi-structured discussion about the information presented in the Web Mirror, as well as their concerns and attitudes towards tracking and countermeasures.

ii Phase 1 Participants and Results

We recruited a total of 25 students from the University of Southampton. Ages varied from 18 to 34. 48% identified as female and the remaining 52% as male. 35% were computer science students and the remaining 65% were from different disciplines, such as social sciences, natural sciences and engineering. The completion time was 13 minutes.

The results for the laboratory experiment indicate that 80% of the participants used between zero and three countermeasures (8% did not use countermeasures; 8% used only 1; 36% used 2; 28% used 3; 16% used 4 and 4% used 7). The most commonly used countermeasures were manually clearing cookies (56%) and browser history (52%) and using ad-blockers (44%).

After receiving detailed information about each countermeasure, 64% of the participants indicated willingness to use at least five countermeasures, 24% at least four countermeasures and 32% all of the presented countermeasures.

We next created an index variable – Premeasures – sum of all the measures already used, and an index variable – Intent – sum of the willingness to use the countermeasures.

We found that the countermeasures information produced a significant increase in intention to use the countermeasures versus previously used measures (Premeasures = Intent – Wilcoxon test p=0.000). We define willingness to use including “Might use” and “Definitely going to use” or “Going to use”, however we also tested a more narrow definition of intent to use only those saying “Definitely Going to use” or “Going to use”. We conducted the same test with the narrower definition, and the results are similar.

However, we found no significant difference between the two treatment groups (Man-Whitney test: All: p=0.543; sub-sample: p=0.2324). The Web Mirror does not increase intention to use countermeasures beyond just seeing the information about countermeasures by itself. This initial result on the lack of treatment effects was surprising.

In order to better understand reasons behind participants' answers to multiple-choice questions, we performed a thematic analysis of their responses to open-ended questions, identified common aspects and developed a code book to apply to the full data set. For each of the questions, we first grouped similar answers, which helped us to distinguish the major themes. We then assigned categories to each of the groups, counted the number of items in groups and calculated the percentage of all responses in each category.

The results have shown that the main reasons why people would not use countermeasures were the fact that they may break sites or make it more difficult to browse (11% – 8% in C1 and 15% in TC), slow down a computer/browsing (11% – 13% in C1 and 9% in TC), installing and/or using them causes too much hassle and/or requires too much effort (9% – 8% in C1 and 9% in TC) and that the participants cannot install any new software on their corporate computers (9% – 11% in C1 and 7% in TC). 17% did not identify any reasons (15% in C1 and 20% in TC). When asked what they learnt in this study, 25% said that they learnt about the existence of countermeasures in general (28% in C1 and 21% in TC), 15% that they learnt about Disconnect.me and Ghostery (24% in C1 and 5% in TC) and 14% that they learnt about tracking and/or the scale of tracking (0% in C1 and 32% in TC).

The focus groups in this phase provide additional qualitative data to help discuss our quantitative results. During the focus groups, most of the participants said they learned about tracking by using the Web Mirror and that the information about countermeasures was informative and would help them to protect themselves from tracking. We found that a majority of the participants expressed surprise at the personalized information received though the Web Mirror. They felt they had learned about tracking and that the Web Mirror helped them to correct beliefs they had about it. For example, participant 14 said: *"It's incredible, the amount of data they can get of you from just browsing the Internet*". The visualization tool worked as an incentive to look for more information about it, with participant 18 explaining that *"I'm actually surprised, because I hear that there is stuff like tracking, something about that, but who cares, I don't care. But with this one, we know something big is going on, that I don't know about. […] I think I'm gonna do more research to understand it".*

We also found that, qualitatively, the combination of the Web Mirror and the countermeasure information seemed to increase some participants' intent to act. For example, participant 11 said: *"I suppose I always kind of suspect that there's lots of people tracking […] but I think when you just take [your data] and you see it on a display, you kind of confront with it, you kind of go – yeah, I should do something about that. So, I'll probably download one of these things".* Participants felt that the generic information on

countermeasures was important in order that they 'd know how to protect themselves against tracking and protect their privacy. For example, participant 11 explained *"I think I should try some of them, I found some good information here"*, and participant 12 said *"Maybe I will download one of these blockers, kind of to protect my privacy, I don't want to be a target"*.

The results of the laboratory experiment gave us confidence in the protocol and study instruments. We therefore proceeded to phase 2 of the experiment to increase the subject pool and provide more robust results.

iii Phase 2: Online Experiment

Phase 2 applied the same experimental design, but was conducted online with participants recruited via Prolific. The aim of phase 2 was to provide a larger sample for statistical analysis.

All participants saw a brief study description, which informed them that the study was about online privacy, that data collection was subject to the Data Protection Act, and that our institution 's ethics committee had approved the study. After reading the briefing information, participants were asked to confirm that they consented to taking part in the study.

Following their consent, participants in the TC treatment were directed to upload a sample of their web browsing data to the Web Mirror, and following a brief waiting time while their data was analysed, were shown their reflection.

While the data was analysed, participants completed a short privacy-attitudes instrument consisting of three statements:

> *1. Consumers have lost all control over how personal information is collected and used by companies.*
>
> *2. Most businesses handle the personal information they collect about consumers in a proper and confidential way.*
>
> *3. Existing laws and organizational practices provide a reasonable level of protection for consumer privacy today.*

These statements were used to categorize the participants by privacy type, prior to their exposure to the Web Mirror and survey. This Privacy Segmentation Index is based on the

long-running Westin privacy surveys [102]. It classifies participants into three categories: *Fundamentalists*, who are the most protective of their privacy, *Unconcerned*, who do not know what the 'privacy fuss' is all about, and *Pragmatists*, who weigh the potential pros and cons of sharing information.

Participants were asked to choose an option from a four-point scale: "Strongly disagree"; "Somewhat disagree"; "Somewhat agree" and "Strongly agree". Those agreeing with statement 1 and disagreeing with statement 2 and 3 were classified as fundamentalists; those disagreeing with statement 1 and agreeing with statement 2 and 3 were classified as unconcerned and the rest of the participants were classified as pragmatists.

Participants then interacted with the Web Mirror for at least three minutes, after which a "Continue" button allowed them to continue to the survey. This time was determined based on our observations of participants during the laboratory experiment.

In the control treatment (CI), participants were immediately directed to the survey after deciding to take part in the study and did not see the Web Mirror at any point.

iv Phase 2 Participants and Results

We recruited a total of 117 participants. All participants were native English speakers born in the US, UK, Ireland, Canada or Australia.

The majority of the participants were from the UK (50%) and 44% were born in the US. The average age was 28, 70% of the participants identified as male and 50% had college education, 30% preferred not to say, 11% had completed high-school. The average time to complete the study was fifteen minutes. 43% participated in the TC treatment and 57% participated in the control treatment. Participants were paid £6 per hour independently of the treatment allocation.

Figure 12 shows the percentages of participants who indicated, during the survey, that they intended to use (or not use) each countermeasure, and the percentage who already used each one. The majority of participants said they were going or definitely going to clear cookies, clear their browser history and use ad-blockers. This is partly because participants were already using those three measures. 52% of the participants said they would not use Tor browser and 30% said they might use it. 17% said they were going to use it, even though the information provided highlighted the cons of this countermeasure. Around 50% of the participants said they might use Ghostery and Disconnect.me and more than 20% said they would use it. Most participants said they will use Opt-out and DNT (around 42%).

Table 11 shows the percentage of participants already using each countermeasure (Used) and their intent to use that countermeasure after receiving the detailed information ("Intent to use"). The percentage of intent to use included those answering "Going to use" and "Definitely going to use". Most of participants were already clearing cookies and browser history, and using ad-blockers. This is not surprising, as these behaviours are well known. Therefore, the main increase is not on the three most used countermeasures. This finding seems to confirm previous work [125] that found cookie deletion is a well-known, if widely misunderstood, action, and that many users installed ad-blockers for reasons other than privacy [58]. As an example, in the focus groups we asked participant 8: "*Do you remember why you started [using ad blockers]?"* to which they responded: *"Because it's annoying, seriously so many ads".* Participant 9 said: *"I think that an ad blocker is useful for me, because I always am annoyed by the ads".*

The increase in use of each of the nine countermeasures is significant at 1% level (Wilcoxon test $p<0.0074$). We conducted the same "Intent to use" test with both the broad (includes "Might use", "Going to use", "Definitely going to use") and the more narrow ("Going to use" and "Definitely going to use") groups, as in the pilot. The results are similar (significant at 1% level (Wilcoxon test $p<0.005$)).

Seeing information about countermeasures has a slight effect on participants' intent to use countermeasures, which can be seen in Figure 12. This effect is stronger on Disconnect.me and Ghostery; which were relatively unused by participants prior to taking part in the study. The Premeasures index is significantly different from the Intent index at 1% level (Wilcoxon test $p=0.000$). Notwithstanding this difference, the effect of the transparency information (the Web Mirror) on action (intent to use countermeasures) is not statistically significant (Mann-Whitney test $p= 0.6069$). Therefore, our hypothesis 1 (H1) is not supported. Transparency information (the Web Mirror) plus countermeasure information does **not** lead to greater adoption of tracking countermeasures than just countermeasure information by itself.

Table 10 reports a random effects OLS regression on the willingness to use countermeasures. The variable Treatment is a dummy variable that indicates that the participant saw the Web Mirror. Use Measures is a dummy variable that takes the value of 1 if participants already use one or more countermeasures. We controlled for participant's level of privacy concern with the dummy variables: Fundamentalists (highly concerned); Concerned or Unconcerned, and whether participants answered that they would like to prevent tracking in some way (Prevent).

The results are consistent with the non-parametric tests. We find no treatment effects and a significant tendency towards using countermeasures from those that already use at least

one of them. In the regression analysis, we controlled for privacy attitudes. 35% were classified as fundamentalist, 8% as unconcerned and 57% as pragmatists. These variables do not have an effect on willingness to use countermeasures. We also controlled for privacy concerns based on the Chellappa and Sin [37] instrument. Participants that had a score higher than four in questions 7-10, were classified as concerned. Those concerned about their privacy were more willing to use the countermeasures, as well as those stating that they would like to do something to prevent tracking. Controlling for knowledge about tracking and demographic characteristics did not change the results.

	[1]	[2]	[3]
Saw Web Mirror	0.114 (0.34)	0.244 (0.33)	0.112 (0.35)
Has used CMs in past	0.457 *** (0.09)	0.442 *** (0.09)	0.422 *** (0.11)
Is Fundamentalist	0.128 (0.36)		
Is Concerned		1.183 *** (0.46)	
Wants to prevent tracking		0.324 * (0.19)	
Intercept	5.274 *** (0.43)	4.296 *** (0.55)	3.367 ** (1.44)
Knowledge	No	No	Yes
Demographics	No	No	Yes
Observations	117	117	117
Rsq	0.200	0.241	0.242

*Table 10: OLS regression analysis on Intent to use countermeasures (** significant at p=0.05; *** trend towards significant, p<0.10)*

	Used (%)	Def Not	Not	Might	Will	Def Will
Clearing Cookies	84	0.8	5.1	13.7	21.4	59
Clearing Browser History	82	2.5	5.2	12	22.4	57.8
Tor Browser	12	12	40.2	30.8	10.3	6.8
Disconnect.me	10	3.4	27.3	47	9.4	12.8
Ghostery	14	5.2	18	53.4	7.8	14.7
Disabling 3rd Party Cookies	37	0.8	15.4	38.5	24.8	20.5
Ad Blocker	79	0	7.7	6.8	16.2	69.2
Enabling Opt-Out Cookies	46	2.6	18.8	36.7	17.9	23.9
Do Not Track	30	2.6	15.5	40.5	17.2	24.1

Table 11: For each each countermeasure , the percentage of participants who had used, or intended to use after receiving countermeasure information sheets (both treatments, phase 2)

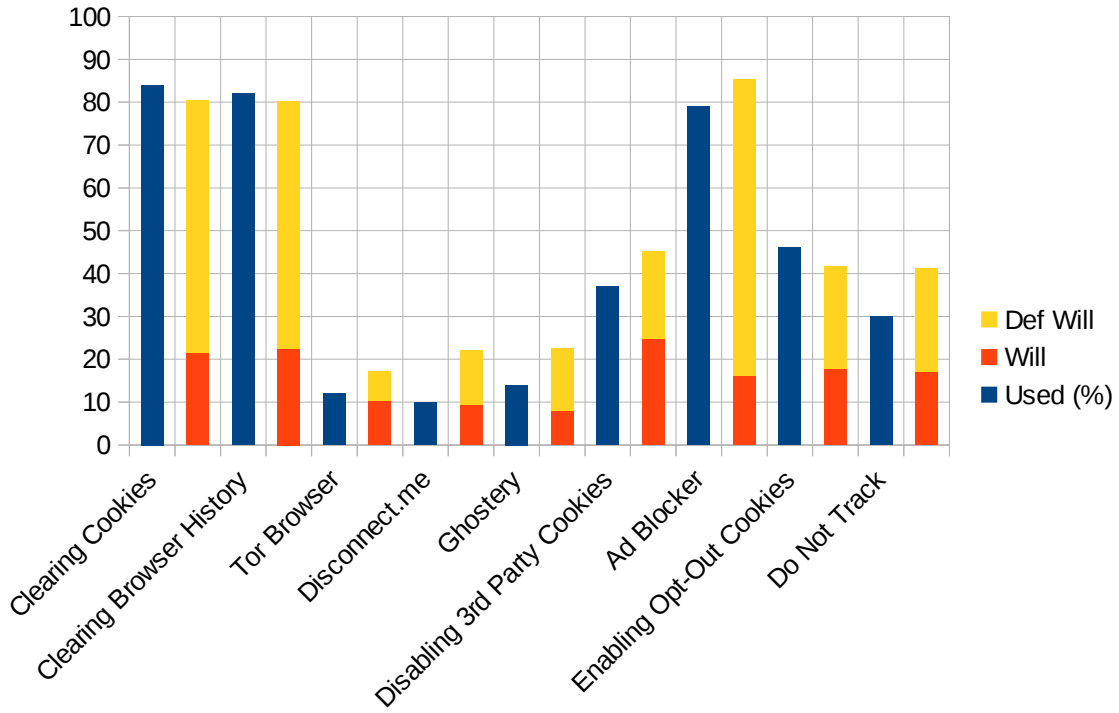

Figure 12: Participants who had used, and indicated intention to use (definitely will use / will use), each countermeasure (both treatments, phase 2)

3.7.6 Experiment 2: Web Cloak

In experiment 1 we found that transparency (the Web Mirror) did not increase participants' intent to adopt additional countermeasures beyond information about the countermeasures themselves. We next designed experiment 2, an online experiment similar to phase 2 of experiment 1, to test whether the lack of treatment effects was because the Web Mirror had no effect on participants' attitude to being tracked, or because the detailed countermeasure information was sufficient to motivate action and was crowding out the effect of the Web Mirror.

The main design change from experiment 1 was on page three of the survey: instead of providing information about the countermeasures and asking participants about their intention to adopt each one, we asked about their willingness to pay for a fictional countermeasure called WebCloak.

This experiment was designed to control for a limitation of Experiment 1, and tested the effect of Transparency in combination with *far less detailed* information about countermeasures on participants' intent to act. In Experiment 1, our survey instrument meant we could not ask about intent to adopt countermeasures without defining what they were, thus providing no opportunity to test the effect of just transparency information. To overcome this limitation, in Experiment 2 we asked about willingness to pay for a single more loosely-defined countermeasure, thus communicating far less through the Countermeasure information itself.

The text read as follows:

> *We are developing a tool, called Web Cloak, that prevents websites from tracking you as you browse the Web. Web Cloak uses advanced analysis techniques to detect and block a range of different tracking types, without breaking the functionality of any websites that you use. WebCloak comes in three tiers; Gold, Silver and Bronze. Gold is the most effective at blocking tracking, but is more expensive.*

Table 12 shows how effective each tier was said to be, on average, at blocking tracking. In the first instrument, participants were shown a table containing columns 1 and 2 (tier name and tracking blocked) and were asked to state the maximum amount that they would be prepared to pay for an annual subscription to each tier: bronze, silver and gold.

In the second instrument, they were presented with the full table, including the prices shown in Table 12, and asked to choose which tier (or none) they would prefer to subscribe to, based on the stated prices and effectiveness.

Tier	Tracking Blocked	Monthly Cost
Bronze	76%	£1.00 / €1.20 / $1.33
Silver	87%	£2.00 / €2.40 / $2.66
Gold	98%	£3.00 / €3.60 / $4.00

Table 12: Information that was provided to participants about the fictional 'Web Cloak' countermeasure.

i Phase 1 Participants and Results

We recruited 82 participants. The majority of whom were from the UK (55%) and 38% were born in the US. The average age was 28. 48% of the participants identified as female and 52% as male. 50% had a college education and 34% had completed high-school. The average time taken to complete the study was fifteen minutes. 80% of the participants were already using ad-blockers. 42% of the participants were allocated to the TW treatment and 58% were allocated to the CII treatment. As in experiment 1, participants were paid £6 per hour independently of the treatment allocation.

Table 13 reports our participants' willingness to pay for each tier (columns 2 and 3) and the percentage of the participants that chose each WebCloak subscription tier, if any (columns 4 and 5). Participants that had seen the Web Mirror were willing to pay more than those who had not. However, the treatment differences are not significant (Man-Whitney test, overall: p=0.1135, Bronze: p= 0.1909, Silver: p=0.1043, Gold: p= 0.0907). We also observe that seeing the Web Mirror had an effect on whether participants were willing to pay for any of the offered tiers, i.e. more participants in the TW condition were willing to *pay* for either gold, silver or bronze. Nonetheless, this difference is also not significant (Man-Whitney test – Overall: p=0.6299; Chi2 -Bronze: p=0.599, Silver: p=0.10, Gold: p=0.631, None: p=0.249). These results indicate that H2 is not supported; transparency information does **not** increase willingness to pay for countermeasures.

ii Phase 2

One possible explanation for the small significant differences in willingness to pay for Web Cloak, or the chosen subscription level, between participants who had seen the Web Mirror and participants that had not, is that some participants were already aware of tracking. The willingness to pay of those who were not already using ad-blockers appeared to be higher (although not significantly) in the TW group than in CII. We therefore wanted to apply a simple filtering heuristic to limit our analysis to the effect that the interventions had on the less tracking-aware participants. However, there was an insufficient number of such participants in the existing data to conduct that analysis (only 20%, resulting in a Man-Whitney test result of p=0.0620 on overall willingness to pay). Therefore, we conducted Phase 2, to specifically recruit additional participants that had **never used** an ad blocker.

In Phase 2, to balance the sample size, we recruited 82 additional participants (the same number of participants in phase 1), 50% of whom participated in the TW treatment and 50% of whom participated in the CII treatment. Table 13 presents the overall results. The results are similar to those of Phase I. Participants paid more for the WebCloak after the Web Mirror. However, the results were again not significant. Therefore, our hypothesis that the Web Mirror would have a strong effect on those that who had not previously used ad-blockers as a countermeasure was not confirmed.

The willingness to pay of UK and US citizens was quite similar (£20 and $18 on average, respectively). Those more concerned about their privacy and those using less than three countermeasures were willing to pay about 5 currency units more for the WebCloak than those less concerned or who were already using more than three countermeasures.

The results of the analysis of qualitative responses, performed in the same way as in the first study, have shown that 39% of the participants (39% in CII and 38% in TW) have chosen the bronze WebCloak tier, simply because this was the cheapest option. 20% (20% and 22% respectively) selected the gold one, as it provided the maximum protection. Although a significantly high number (48% – 38% in CII and 59% in TW) of them did not indicate any reason why they could not use the WebCloak, 16% (20% in CII and 12% in TW) said that what would stop them from using it was the cost. Surprisingly, 7% (12% in CII and 0% in TW) did not trust the WebCloak – even though they believed that it was to prevent tracking, they wondered whether the tool itself might perform tracking. As the main learning outcomes from the survey 35% (51% in CI and 15% in TW) of the participants reported the fact that it is possible to prevent tracking and 26% (12% in CII and 44% in TW) generally that tracking exists.

	Phase 1				Phase 1 + Phase 2			
	Payment Value *		Chosen by %		Payment Value *		Chosen by %	
	CII	TW	CII	TW	CII	TW	CII	TW
Bronze	11.21	13.72	18.75	23.53	11.50	13.21	22.73	25.33
Silver	16.75	19.32	8.33	20.59	16.88	19.06	12.50	17.33
Gold	22.19	28.7	37.5	32.35	23.20	27.84	34.09	34.67
None	-	-	35.42	23.53	-	-	30.68	22.67
Total	17.17	20.54	100	100	17.44	19.92	100	100
Obs	48	34	48	34	89	75	89	75

Table 13: Mean Value of willingness to pay (in currency units), and percentage of participants choosing each level of WebCloak, in Phase 1 and Phases 1 + 2*

3.7.7 Discussion

In this study, we tested the effect of providing both transparency and countermeasure information towards intention to act to limit tracking. Our results suggest three key findings: first, that detailed, personalized transparency information in the form of the Web Mirror does not seem to increase intended countermeasure adoption beyond that which generic information about the countermeasures does. Second, that most participants had a desire to adopt surveillance countermeasures after being told about them. Third, that the lack of increased intent to adopt countermeasures is in spite of participants reporting that the Web Mirror had brought aspects of online tracking – such as the scale – to their attention.

i Desire to Adopt Countermeasures

Our results show that, in all treatments, there is a willingness to adopt additional countermeasures (71% of the participants intend to use most of the suggested countermeasures). This is consistent with previous research [112,119] that shows many users have a negative attitude to being tracked, when told about it. In experiment 1, our results show that the range of countermeasures that participants considered adopting was not increased by the use of personalised transparency information. In experiment 2, our results show that the strength of intention to adopt a countermeasure (measured via

willingness to pay for such a countermeasure) was also not significantly increased by the transparency information provided by the Web Mirror.

From the qualitative analysis of the focus groups, we find that most people have some knowledge about tracking and although the Web Mirror has clarified this idea, the most important factor was to know what to do to stop this tracking. For example, participant 22 said: *"It raised my question about online tracking (...) I want to make some private space which I can only see on the Web and maybe those tools will help me to do so"*.

The provision of information about countermeasures themselves seems to provide adequate stimulus to motivate intent to adopt countermeasures and any effect from the transparency information is rendered inconsequential. Crucially, from a practical perspective there is a greater degree of freedom about whether or not to supply transparency information than there is about whether or not to provide information about what control mechanisms are available. It is not realistic to expect users to be able to act to control data processing activities in the absence of information about how to do so, although such information might be expressed through interactive features rather than static prose.

ii Motivation & Comprehension

The qualitative results from the survey do suggest that the Web Mirror has an effect on participants' understanding of the tracking ecosystem. For instance, 33% of them (in both experiments) mentioned some aspect of tracking when asked what they'd learned by taking part in the study – typically referring to the scale of the tracking network. This is consistent with our earlier work developing the Web Mirror, in which we consistently found participants express surprise about regard to the amount of tracking that takes place, and a relatively low initial understanding of how advertising companies obtain information such as interests and demographics.

In the context of the overall results, though, this more detailed understanding does not appear to translate into additional intent to act. From an empowerment perspective, this raises an interesting question – do we inform citizens about data processing practices purely to enable them to make decisions, or are we informing them because transparency is valuable in and of itself? If the former, then identifying the minimum amount of information that's necessary for citizens to reach a "stable" decision (i.e., one that is not affected by being given further information) appears, prima facie, to have some interaction benefits in terms of lowering cognitive effort and decreasing the amount of time required to engage with privacy-protection mechanisms.

In experiment 2 we expected that personalized transparency would lead to a significant Action effect in each experiment, based on the responses to the Web Mirror in our focus groups. Instead, our results suggest that while the Web Mirror seems undoubtedly to affect comprehension of tracking, that kind of comprehension is not required in order to be motivated to act.

iii Design Implications

The failure of the transparency that's provided by the Web Mirror to motivate additional intent to adopt, or willingness to pay for, anti-tracking countermeasures is in some senses surprising. The results during the Cookie Visualiser study suggested that seeing the scale of tracking was surprising to participants, and objectionable, to people; and the qualitative results from this study seem to confirm that. Nonetheless, that does not translate into increased intent to actually *act* at least in terms of the instruments (intent to use countermeasures, and willingness to pay for a countermeasure) that we used here.

Broadly, I suggest that this finding challenges the assumption in privacy and data protection regulation that detailed transparency is necessarily helpful to individuals when making decisions about their engagement with services; and clarifies the scope of tools such as Mozilla Lightbeam that, while offering interesting insight into the track ecosystem for interested users, may not provide great value to general users.

However, this result is not necessarily negative; in fact, our participants' willingness to consider, and even pay, for anti-tracking countermeasures when those measures are described to them suggests that providing readily-available *means of acting* with respect to privacy could be effective even without finding ways to make detailed information easily consumable.

These findings are revisited, along with results from the previous studies, in the next section.

3.7.8 Limitations

The current study considered intent to use countermeasures, rather than their actual use. This was a deliberate decision, as the ease of installing or using countermeasures could be a major confounding factor in their actual adoption. This raises two immediate questions for further exploration: 1) Are countermeasures usable enough to be adopted by the people who said they'd like to use them? 2) Is there a disconnect between reported intent to use surveillance countermeasures and actual use, in the same way that the "privacy paradox" describes a disconnect between reported intent to disclose personal data and

actual disclosure? By extension, and given the strong qualitative response to the Web Mirror, but the limited significant effect on intent to use countermeasures, it would be interesting to come back to participants over time to see if their behaviours have changed, and how their practices correlate with treatment condition. One might hypothesise that the additional comprehension facilitated by the Web Mirror (or other transparency information) sustains concern or action over a longer period of time, or that longer-term adoption of countermeasures is increased having had more time to think about the information presented by the Web Mirror.

3.7.9 Conclusions

In this study, we drew a distinction between information that increases the transparency of online surveillance and information about actions that Web users can take to counter online surveillance. We showed that, contrary to our initial expectations and the apparent assumptions that underlie many tools and regulations, detailed transparency information itself may be of little value in shaping the intentions of users with regard to the adoption or use of Countermeasures.

The surprising absence of additional intention to adopt surveillance countermeasures as a result of being shown more detailed Transparency information has implications in a range of spheres. In data protection policy, it calls into question the regulatory focus on transparency – such as the EU's much maligned "cookie policies" – and instead points to opportunities around articulating the available control mechanisms. In HCI, our research helps to understand the contexts in which information visualization may be useful, helping researchers and practitioners to more confidently identify which types of intervention (transparency information, or control mechanism) may be useful in a particular scenario.

In the next section, I revisit results from the Web Mirror, and the previous privacy studies, from an Activity Theory perspective.

3.8 Discussion: An Activity Theory View on Privacy

We observed in the Cookie Visualiser study that privacy is, by and large, a longer term concern that is overshadowed in the moment by the task at hand. That, while to some extent privacy management is an activity in its own right, privacy is also something that 's affected by a wide range of *other* activities; privacy is frequently an internality.

In an activity theory sense, in the course of carrying out many of these other activities people seem to fail to adequately take privacy into account – at least in as much as the outcomes are often incongruent with people 's stated privacy attitudes.

If one were to summarise the role of privacy in web-related activity systems, it would probably be best described as a factor that is missing in the construction of complex multi-stakeholder objects; one which users must struggle to introduce, rather than one which is reflected in how the activities of first-party websites and the rest of the tracking ecosystem are shaped.

I begin this section with a discussion of how privacy relates to the idea of "constraints" that Objects are constructed in response to; I then consider how privacy relates to the Internal Plane of Action, the overall structure of third-party tracking as an activity system, and finally possible design implications arising from those prior observations.

3.8.1 Privacy as a restrictive constraint on activity

As designers who are interested in helping users to make more consistent decisions, AT seems to offer a helpful framework to explain and reason about values such as privacy. Essentially, our design challenge is to help shape activities in a way that reflects privacy attitudes, but which still fulfils the original object. On the face of it, values such as privacy do not seem to necessarily motivate *new* activities, but may act as constraints that shape other activities.

In essence, privacy (at least in most scenarios) does not seem to be a motive in quite the same sense as "obtain food" might be. In the latter case, the motive clearly energizes and directs an Object by itself. Privacy, though, typically only becomes relevant in the context of another activity (such as using online services); rather than being something that could be pursued directly and in isolation.

We live, socially and legally, in a society where our privacy tends to remain intact unless we undermine it in the course of doing something. By and large, we are not subject to informational privacy violations in the course of sitting at home or even in the conduct of our general lives. To some extent this is because only we have access to much of what we consider private – like our preferences, thoughts, and personal information. It is hard for someone to obtain your phone number unless you give it out in the first place. Even where privacy can be violated by observation and inference, there has to be an action to observe, or to infer *from.* Legally, there is a presumption of privacy. In most cases, organisations cannot process data about us unless we voluntarily engage with them as customers or service users, or by otherwise giving them our consent. Socially, well-developed notions of physical and personal privacy often serve to protect our informational privacy, too.

Privacy does become important within activities, though. Participants in the focus groups have privacy concerns, and do try to adjust their actions in response to those concerns. As discussed earlier, this is sometimes grounded in a belief that privacy is valuable in itself, and in others a more consequentialist concern about particular negative outcomes that could arise from a loss of information privacy. Despite not being a need to which those participants' activity ultimately answers, privacy *does* shape their activity, and there is a strong sense that there is a desire for privacy to shape many participants' activities to a greater extent than it currently does; particularly when faced with information about practices such as third-party tracking and the privacy impact that it engenders.

Kaptelinin discusses in [93] how Objects are dynamically created in response to *constraints* such as motives, available means and other actors. Privacy certainly has a role in shaping the Object of activity; but because it does not serve the same motivating role for the activity that a true motive would, I suggest the term *restrictive constraint* to describe it. A restrictive constraint is a factor that shapes the Object of activity, but does not excite that activity in its own right.

The idea of privacy – a value – as a restrictive constraint is consistent with the definition of a value provided by Rokeach, *"some means or end of action is judged to be desirable or undesirable," [153]* in this case, Privacy is an important determinant of whether particular means of action within an activity – and the overall Object itself – are desirable. It restricts which, of all possible actions and Objects that could be directed towards the motives for activity, are desirable, and which are preferable to one another.

3.8.2 Privacy and the Internal Plane of Action

A defining feature of my participants' relationship with online services is their inability to articulate how those services work, particularly with regard to how the services collect and process personal data. Activity Theory's concept of the Internal Plane of Action can be used to explain the relevance of this result. Kaptelinin explains [91] that *"the internal plane of action is a concept developed within activity theory that refers to the human ability to perform manipulations with an internal representation of external objects before starting actions with these objects in reality."*

Participants frequently engage in web use with simplistic or even incorrect models of how interacting with services will impact upon their privacy, and this can be seen quite clearly in the qualitative results obtained during the Web Mirror and the Consent Notice studies. Moreover, privacy research highlights the limitations on the extent to which the IPA can be brought to bear in everyday decisions. Bounded rationality, task focus and limits on (for instance) the number of factors that can be considered simultaneously mean that making predictions about the impact of an action can be fraught even when the subject has relatively good understanding. To some extent, (as discussed in chapter 3.6) privacy as a construct for dealing with uncertainty about the consequences of an action is itself helpful in this regard. Thinking about an abstract notion of privacy may be more tractable than foreseeing specific consequences.

In the case of online tracking, though, even the privacy implications of just browsing the Web are not immediately apparent, and are hard to infer or understand even with a working knowledge of how technology likes cookies, or the advertising ecosystem, works. Tools that draw a more direct link between technologies like cookies, the underlying tracking ecosystem, and values like identity and privacy may have a role in simplifying the model, or automating the inferences that would otherwise have to be made by a user themselves. They can do this partly as tools with learning affordances, allowing users to develop their knowledge and enriching the internal plane of action; but can also act to supplement the IPA, by doing the processing and drawing inferences that humans might otherwise struggle with.

As the Web Mirror showed, though, a deep understanding of the implications of something like tracking is not necessarily required in order to formulate a response. The simple explanations that we provided of available actions – installing or using different countermeasures – were sufficient to motivate a similar level of motivation to *use* those countermeasures as the very detailed transparency that was provided via the Web Mirror. From an interaction design perspective, this is a powerful result: values and motives can

be leveraged as a source of motivation, and the link between an action and values to which it is relevant, can be communicated without detailed models and information. In contrast to the over-simplified user interactions that often drive concerns like privacy *out* of people's awareness and decision-making, the evidence suggests that we can use relatively simple explanations to bring those values back *in* to the overall activity by creating simple actions, like installing anti-tracking countermeasures, that can take care of these concerns outside of the task-focussed cognitively-limited contexts in which the initial internalities arise.

3.8.3 TPT as a Multi-Stakeholder Activity

As a case study in privacy, third-party tracking is interesting to consider as an activity system. As the tracking network study showed, there are at least three types of party within the ecosystem itself: first party websites, brokers, and advertisers; plus the user themselves. I'll begin by considering the apparent Objects of the different stakeholders: advertisers, ad brokers, first-party websites, and the Web user.

Of these, the activities of ad brokers and advertisers are fairly well aligned, and could be described as a single coherent activity; albeit with slightly different motives between the two different types of actor. The Object of the activity undertaken by these parties is to match screen space on first party websites with adverts that are likely to be relevant to the user. The ad broker is motivated by the revenue that this activity generates (either from clicks – "pay per click" – or total advert views – "pay per impression"). The advertiser is motivated, in most cases, by a desire to market a product or service to users who might be interested in it. Unlike contextual advertising, where adverts are chosen based on the content of the first party web page itself, behavioural advertising tracks the user across multiple pages in order to infer which market segments they might be part of, such as a particular demographic or interest group [118]. Moreover, given that this ecosystem is largely automated, it can be conceived of as exhibiting agency in itself, on behalf of the companies that build and maintain it [99].

The first-party website's activity is harder to define as a particular Object; the range of first parties is broad, and there is no motive common to all of them. What they all have in common, though, is a need for the revenue that the advertising system provides. For some first party websites, this revenue could be a motive in itself; for others, it may be an economic necessity in service of another non-commercial motive. In all cases, though, the first-party incorporates the activity of the ad brokers, and advertisers, by making arrangements for adverts delivered by those parties to be incorporated into their web pages; and for the tracking that supports the targeting process to take place.

As with first-party websites, the activity of users is hard to define exactly. Using ad-supported websites can be part of many different activities; from conducting research for education, to browsing the Web for leisure, purchasing items through an e-commerce site, undertaking paid employment, or researching a health condition. At first glance, all of these activities appear to be largely orthogonal to the activity of the advertising ecosystem; to the extent that it doesn't really make sense to think of them as voluntary participants in that broader activity at all. It is through the actions and choices of the first party website that users become involved in the tracking activity system.

To illustrate how a first-party website constructs an activity system, let's consider the example of an online newspaper like the Guardian. The Guardian is owned by Scott Trust, a body that aims to "*secure the financial and editorial independence of the Guardian in perpetuity and to safeguard the journalistic freedom and liberal values of the Guardian free from commercial or political interference" [185].* The Object of activity of the Guardian is fairly clear; to publish journalism, in order to promote liberal values in public discourse (rather than commercial gain). The Guardian website allows readers, without whom the journalism would have no effect, to access the content. In order to fund this activity – that is, to remain a sustainable organisation that can pay for human effort and other resources – the Guardian makes use of behavioural advertising and in doing so brings the (diverse) activities of its readers, its own publishing activity, and the ad-delivery activity of the ad brokers and advertisers into an inter-related activity system.

In doing so, a restrictive constraint of the Guardian's readers – privacy – is brought into conflict with the actions of the behavioural advertising networks – which, most people agree, fail to preserve privacy because of the way that their services operate.

In their work on Ajaxe (a pharmaceutical company) [95] Kaptelinin and Nardi note how, in a multi-stakeholder activity, different motivations between those multiple stakeholders can lead conflicts in the pursuit of a shared object, even though each stakeholder has a good reason to pursue that shared object. The different stakeholders bring different motives, essentially resulting in a shared polymotivated Object; although some motives may be shared, it is not necessarily the case that all of them are. In the case of Ajaxe, stakeholders involved in biotechnology research used *curation* to selectively reveal information to others, and so to influence how the activity unfolded in practice.

Notably, the conflicts between the different stakeholders arose "*from the motives in relation to one another, not the motives in relation to the object,*" "*the object itself was not contested; it was the instantiation of the object that led to tensions*". I argue that this is not the case in the third-party tracking example, and that despite the coherence of the activity system as seen from the perspective of a first-party website, the Objects constructed by

web users (which commonly includes privacy as a restrictive constraint) and the Object constructed by ad brokers (which, as judged by most users, does not) are essentially disjoint; the revenue generated by advertising is tangentially related to the activity that web users undertake through tools like Facebook or online news sources, and those users exercise little to no control over how those tools derive the revenue that's necessary to sustain themselves. The power relationships – and imbalances – between the different stakeholders shape how the conflicts between the different stakeholders are resolved in the collective object; in this case, largely relegating privacy entirely.

So, are there means for users to influence the overall activity that the first-party website has constructed? And, situated within this activity, what is the role of design vis-a-vis privacy?

3.8.4 Privacy Breakdowns in the TPT Activity System

The results from the Web Mirror demonstrate that participants were willing to incorporate actions into their activity in response to a privacy motive, even when those actions imposed some cost.

Conceptually, there appear to be two stages involved in a user's decision to install anti-tracking countermeasures; or to make use of the controls provided by a first party. First, that a link between an action – like using a website – and a value that is relevant to that action – like privacy – is established. Second, that an action in response to that motive is offered. This is a combination of an initial learning affordance – that enriches the user's internal plane of action to link web use, tracking, and privacy together – and provision of an instrumental affordance that allows that relationship to be altered.

From my discussions with participants, an explanation of *tracking* seems to be more easily understood than a discussion of *cookies;* and the information that we provided about countermeasures during the Web Mirror study was sufficient to draw an association between web use and privacy; and to provide a set of possible actions that users could take in response.

Applying this basic model – of establishing a link between and action and a value, and then providing an instrumental affordance to mediate it – to the results of the Cookie Notice survey, we can observe that not only are instrumental affordances missing, but the learning affordances that would allow a user to relate the proposition (accepting cookies) to a relevant value (privacy) are also missing. Most websites effectively require users to reconfigure their Web browser, or leave the website entirely, in order to avoid tracking cookies; with only a minority offering more fine-grained control.

Conceptually, the conflict between *privacy* – a value held by the user – and the tracking activity conducted by a functional organ composed of the first-party website and ad broker (FP-AB), has the potential to create a breakdown [27,101] in the activity. When it is brought to the users attention, the unanticipated behaviour (the collection of browsing history data) of the tool (the FP-AB, and the web browser) vis-a-vis privacy causes the tools themselves to become the focus of the users attention. At this point, privacy becomes a *need* and the new motive of preserving privacy necessitates a reconstruction of the Object of Activity [93].

In the case of third-party tracking, different actions can resolve this breakdown in different ways. Leaving the first party website entirely removes the user's own Object from the emergent Object of the activity system that the first party website had incorporated it into; but this course of action demands a potentially major reconfiguration – either finding an alternative website, or (more likely, given the prevalence of tracking on the Web) seeking a non-Web alternative. If the user instead changes their Web browser settings, they can resolve the unwanted behaviour by breaking the ad broker's activity instead, preventing their automated activity from taking place; but, in the process, potentially breaking some aspects of the first party website itself.

In practice, the actions people take in response to concerns over tracking – such as the adoption of anti-tracking countermeasures – are different to the actions suggested by websites themselves. Many of the tracking countermeasures that exist, and which we considered in the Web Mirror study, interfere with tracking practices, and – like reconfiguring a Web browser to reject third party cookies – interfere with the ad broker's activity while leaving the user's own Object in essentially its original construction.

Considered in this way, elements of an interaction that bring unanticipated behaviour to the user's attention have an important role in enabling the user to reconfigure their activity to protect their privacy. One special case of creating deliberate breakdowns for this purpose is *consent*, and I will discuss that next.

3.8.5 Consent as Focus Shift

As touched on in the background section, the concept of consent is well-established in many areas of law and ethics. Intuitively, consent provides people with the opportunity to exercise control over important aspects of their life or personhood, for instance when undergoing medical procedures. Consent is now recognised as an important element of privacy and data protection, notably in the European Union's General Data Protection Regulation (GDPR) [168].

Consent is a particularly interesting aspect to consider in the context of VSD and AT. From a VSD perspective, consent can itself be construed as a value, and – as work by Friedman *et al. [67]* shows – decomposed into values relating to how it should be employed in practice. The idea of "giving consent" also makes consenting an action, that can be situated within an activity system. In the GDPR, consent is considered to be *"a clear affirmative act establishing a freely given, specific, informed and unambiguous indication of the data subject's agreement"* (Recital 32) [168].

The effect of the law is that, where applicable, consent requirements create a point at which processing becomes contingent on a deliberate and informed action of the data subject that signifies their consent.

That is to say that, in the context of a multi-party activity system, which involves (at least) a data subject and data controller, consent becomes a mechanism that first establishes (in the data subject) an understanding of the activity that is proposed by the data controller, and second, waits for a voluntary and affirmative signal that the data subject wishes that activity to take place.

From an AT perspective, the consent process deliberately aims to expose aspects of the multi-party activity system to a person. This action must, in itself, become what Bødker [28] calls a focus shift; it shifts the focus of the user towards aspects of the activity – in the case of TPT, the tracking that is operationalised through the tools that the user is working with – that were previously unarticulated.

Moreover, the goal of that focus shift is (as Friedman's six factors explicate) for the user to develop a deeper understanding of aspects of the activity, and then to deliberately re-instantiate that activity through signalling their voluntary agreement. In the event that the user does n*ot* consent to the activity, the focus shift becomes a breakdown; and the user must somehow reconfigure their activity.

Framing consent as a breakdown in the user's activity system draws attention to the necessity of re-constructing the activity following such a breakdown. Voluntariness is one aspect of consent that Friedman et identified [67] and refers to the viability of a user declining consent; or, in an Activity Theory sense, the viability of re-constructing their activity following the breakdown that a consent interaction has induced. The implication is that voluntariness could be increased if that re-construction were facilitated more directly; for instance by offering means of resolving the value tension that caused the breakdown, as we did in the Web Mirror study.

Some of these themes are revisited later in part 4, Unification, where I draw them together with findings from the next section, which considers values and technology in the context of wellbeing.

4. Exploring Values in Wellbeing Activity

Wellbeing is a complex topic that combines aspects of a person's physical and mental health, as well as more subjective notions such as personal fulfilment and happiness [14,35]. The distinction between objective indicators of wellbeing, such as physical fitness, and subjective factors, such as feeling fulfilled, is a major feature of the domain. Research into improving the health of human populations is often motivated by a desire to reduce the cost of morbidity to public health service such as the NHS. For an individual, though, wellbeing is often an intensely subjective concept. Few would argue that they indulge in chocolate, alcohol, or even unprotected sex because of the objective health benefits, but there is an acknowledgement that all of these things can contribute to our subjective experiences. It's not realistic to expect to die in perfect health, and so in some ways, at some times, objectively unhealthy things can make our lives *better* from the complex subjective position that we experience them from.

Nonetheless, historic levels of Obesity and associated morbidities such as diabetes have been identified as major public health concerns in many countries, including the United Kingdom. Despite large-scale public health campaigns and regulatory efforts such as changes to food labelling, obesity prevalence has risen steadily and, in 2012, stood at around 25% for both men and women in England, with 34% of men and 45% of women having a 'raised' waist circumference [115]. In addition, many individuals demonstrate difficulty in maintaining what would be considered objectively healthy lifestyle practices, such as regular exercise or a healthy diet. Even if copious amounts of chocolate could improve life in the short term, there is little subjective wellbeing to be had in an early grave. Subjective wellbeing is not entirely divorced from objective physical health and finding an acceptable balance between short-term enjoyment and long-term wellbeing is often non-trivial.

As a result of concerns over the health implications of widespread obesity and sedentary lifestyle, there has been significant interest in the role of technology in supporting and encouraging behaviour change from both academic and commercial actors. In particular, the HCI community has published heavily on the use of "persuasive" devices to encourage health behaviour change. Commercially, apps such as MyFitnessPal allow individuals to record the exercise that they perform and the food they consume, and

wearable devices such as the Fitbit or JawBone Up provide, among other features, step tracking to quantify everyday physical activity. The persuasive techniques first employed in HCI research – such as gamification and social support – are now often present in commercial offerings. Fitness tracking is a key feature of both the Apple Watch and Android Wear devices. However, the abandonment rate for wearables is relatively high; a survey released by Endeavour Partners [110] showed that in 2013, only 43% of consumers were still using a wearable device 12 months after acquiring it, and that despite improvements to the devices, in 2014 a third of consumers had abandoned their wearable device after six months. While this attrition could have various explanations, it does appear that they fail to deliver sufficient value – to improve the lives of their end users *enough* to outweigh the hassle of wearing them.

In this section, I first consider background and related work in Wellbeing – in particular health behaviour change – and then report on three studies that I conducted: A survey of wellbeing practices, interviews with runners, cyclists and Fitbit users, and finally a speculative design exercise in which I attempted to reframe fitness activities using alternative values. Finally, I draw these findings together to consider how they can be explained using Activity Theory.

4.1 Background

As previously mentioned, conceptions of wellbeing fall broadly into two groups: the *objective theories*, which conceptualize wellbeing as a property of an individual that is independent of that individual's preferences or values, and *subjective theories*, which conceptualize an individual's wellbeing as contingent on their adopted values and attitudes [13,14,35,52]. Typically, the focus of public health and technological interventions is on practices that will influence attributes of objective wellbeing such as obesity or blood pressure through behaviours such as exercise, the promotion of fruit and vegetables and limiting salt, sugar or fat consumption. However, the practical implications of the distinction between subjective and objective wellbeing do not demand allegiance to either viewpoint. Human beings share largely identical physical embodiments, and therefore the same basic needs. Thus, while cultural context, such as a place, time, demographic, and socio-economic conditions vary between us, the subjective determination of 'good' related to one's physical embodiment is likely to have similarities.

The traditional approach to health and wellbeing, typically exhibited within medicine, is to consider health as the absence of any disease or health problem. As Felicia Huppert notes in a review commissioned by the UK Government, "*Future developments in the*

science of well-being and its application require a fresh approach—beyond targeting the alleviation of disorder to a focus on personal and interpersonal flourishing", though [85]. There is considerable scope for *proactive* efforts to improve wellbeing, even if the absence of specific diseases and beyond just *avoiding future disease,* too.

A similar, though distinct, point is made by Michele Crossley in "Rethinking Health Psychology" [47]. In particular, when critiquing mainstream health psychology she comments that *"insufficient attention is paid to the fact that human beings are often motivated by 'irrational', unconscious forces and emotions. In addition, health-related behaviours such as as eating smoking, drug-use and exercise embody latent social, cultural and value-laden meanings that individuals incorporate into their thinking but are not necessarily consciously aware of." (p.38)*

Crossley continues, drawing on the work of numerous other authors, to discuss, for instance, the role of "trust, intimacy and love" (p.46) – undoubtedly factors that most people would consider relevant to their subjective wellbeing – as motivations for gay men to engage in unprotected sex, putting them at greater risk of HIV transmission – a factor that most people would consider detrimental to their wellbeing. Wellbeing, then, is not so simple as fulfilling certain needs, it is a complex process of balancing potentially competing factors based on physical, cultural and individual factors and often with incomplete understanding of the risks or implications of a particular action.

The overall picture is of health as existing on a continuum between objectivity and subjectivity. Aspects of our shared human physiology provide a common basis for objective measurement of some aspects of wellbeing – for instance the likelihood of stroke or heart attack. Shared, but not universal, factors from our culture provide a more subjective, but not entirely individual aspect to our wellbeing, and our own beliefs, knowledge and experiences provide subjective and sometimes unique aspects to our own self-perceived wellbeing.

I next turn my attention to existing work in HCI that has considered wellbeing, in particular "fitness", which perhaps most directly corresponds to the major health (and hence wellbeing) challenges facing populations in the UK and many other countries.

4.1.1 Fitness Devices

HCI has studied – and produced – a wide range of ICT devices for use in fitness activities. These devices are hugely diverse, both in terms of their capabilities and how the fit into an overall activity. This section doesn't aim (nor could it hope) to exhaustively review all

of HCI's work on fitness-related ICTs; rather it attempts to synthesise salient aspects of the application of digital technology to fitness activities.

Consolvo *et al.*'s '*Houston*' [43] is a relatively early example of a mobile fitness technology; motivated the potential of physical exercise such as walking to counteract an upward trend in obesity, Houston used a combination of step-counting and social sharing to promote *opportunistic* exercise into users' daily routines.

Gamification – the introduction of game mechanics into activity-promoting devices – as a means of persuasion is a major area of study. Two broad approaches have been examined: incorporating game mechanics such as leader boards or achievements into fitness-logging technology [182], and incorporating physical activity into existing games; the latter often being referred to as "exergames". Examples of exergames include Fish'n'Steps [116], and SwimTrain [39], which extends the concept to groups of people.

Chen *et al.* [38] explored how the framing of an "exergame" – a game designed to incorporate physical exercise – would affect exercise duration. They found that participants primed to think of their involvement as exercise spent considerably longer (49 minutes vs 39) playing the game than those primed to think of it as enjoyment. They found that displaying health information such as calories burned did not impact upon exercise duration, but did elicit increased positive affect.

Zhao *et al.* [182] looked at long-term exergame use (70 days) and deliberately tested the effects of gradually introducing new features on retention and overall activity. They found that gradually adding new features reduced the decrease in user activity over time, but that activity levels in all cohorts in the study declined over time. While promising in some respects, the results point to a more general trend of declining effectiveness over time for many different fitness devices (e.g. [74][164]). A notable exception to this trend was observed by Meyer at al [128], who found no notable decrease in usage among people using Fitbits as part of recovering from heart attack; suggesting that usage context – and perhaps motives in particular – might be an important factor in long-term usage of devices, and engagement with the related physical exercise.

Gouveia *et al.* [74], despite finding decreased engagement over time with the *Habito* system, note that scaffolding technologies, which make up for gaps in expertise, knowledge or motivation, might be expected to be used less over time, and to eventually be abandoned, as the skills that they scaffold are internalised by the user. Conversely, Fritz *et al.*'s study [70] of long-term Fitbit users (although not in itself a longitudinal study) found that a *failure* to abandon devices can lead to negative affect among users – feeling that steps are 'wasted' when the Fitbit is not worn – and engagement in inefficient

activities – for instance pacing around a room when the time could be better spent on a more intense physical activity.

Clawson *et al.* [40] examined the active abandonment of fitness technology by studying secondary sales of those devices on Craigslist. They found a wide range of reasons for abandonment; some positive – like achieving fitness goals or upgrading to a different device – and others negative – like finding it less useful than they had anticipated. Around 26% of the ads were coded as demonstrating an *"*internal decrease*"* in motivation, either to use the device or to engage in physical activity in general; suggesting that the device itself did not create or sustain sufficient motivation in the user. As Gouveia at al, the authors point to the phenomenon of positive abandonment as both a signal of success and an event to design for. They conclude that there is a need for *"new ways to conceptualize and theorize about individuals' health behaviours that take into account new streams of information, dynamically changing contexts of use, and rich messy practices that emerge around technologies for health"* and suggest that self-monitoring must be designed as a means to diverse ends, rather than seeing self-monitoring as a goal in itself; the sort of contextual account of activity that activity theory is well-suited to provide.

Li *et al.* [113] studied the *information* needs of health-technology users, and found two distinct phases of use – discovery and maintenance – each posing distinct questions of the logged data. In the discovery phase, users are trying to understand their behaviour, understand what influences it, and identify goals. In the maintenance phase they are concerned with their progress towards their goals, and any discrepancies. The authors point to the need to identity and tailor to the appropriate phase, and to support the transitions between them.

Of course, devices designed specifically for use in fitness activities are not the only technology that is *used* in that context. Patel & O'Kane [143] studied technology use within a gym setting, including fitness devices and other technology such as televisions and music players. They concluded that a whole range of aspects influence how technology is used in gyms, including the user's values and type of motivation. They also found that different actions within the gym led to very different uses of technology; with, for instance, music players being a welcome distraction during aerobic activities like running, but interfering with the need to count reps when lifting weights.

The overall picture of ICT use within physical exercise is promising, but points to many challenges. In the long-term, fitness devices can fail because they do not sustain engagement, or they can fail because they sustain engagement in suboptimal activities or create feelings of dependence. In a more immediate sense, a music player can be a

beneficial distraction in the context of a repetitive activity, but becomes intrusive when the focus shifts to counting repetitions of a lift.

On the face of it, Activity Theory seems to offer some structure to account for these different concerns. It provides a broad object-level context that describes the ultimate goal of an activity – improving health, for instance – that allows for consideration of why, ultimately, somebody has engaged with a device at all. It also provides the granularity to talk about the conditions necessary for an individual operation to take place, and hence to conceptualise the distraction from a music player. In the series of investigations presented next, I have attempted to shed light on the motive-object end of this activity theory spectrum; asking, basically: How do devices align (or not) with the ultimate motives of the people who do (or might) use them? In the hope that understanding more of how ICTs can be situated within an overall activity in the AT sense (that is, motivated towards an object) and not just how they fit with engagement in particular exercise (like walking or dancing) might influence their future design.

Two concepts stand out in existing HCI work on fitness devices; **motivation** – and the necessity of motivating people to adopt particular behaviours, and **reflection** – helping people to make sense of past actions as a means of changing and improving future behaviour. I'll briefly review these two specific concepts in more detail, as they are both broadly applicable concepts, with roots outside of the HCI literature.

4.1.2 Motivation

Motivation is identified a key factor in behaviour change [129], and in behaviour more generally. It's therefore a major feature in most theories of behaviour change, and is commonly cited – in academic literature and in conversation – as an important factor in wellbeing-related activity such as dieting for weight loss, or taking part in sport. Given its importance, motivation is well studied in cognitive psychology and in applied disciplines like education, where understanding motivation helps to explain engagement – or disengagement – with the learning process.

Most of the literature on motivation distinguishes between intrinsic and extrinsic motivation. The former refers to "doing something because it is inherently interesting or enjoyable", while the latter refers to "doing something because it leads to a separable outcome" [67]. Deci *et al.* describe intrinsic motivation as an "innate psychological need for competence and self-determination" – leading to a sub-theory of self-determination theory, called Cognitive Evaluation Theory [50], that uses perceived competence and autonomy to explain intrinsic motivation, and in particular the observation that extrinsic

motivation – such as rewards – seem to decrease intrinsic motivation. The antagonism between intrinsic and extrinsic motivation is challenged by Covington & Mueller [45], who suggest that the observed loss of motivation may be explained by the *scarcity* condition under which extrinsic motivations are often given – particularly in classroom situations – instead arguing that extrinsic rewards, in an appropriate reward system, can reinforce intrinsic motivation. They use a qudrapolar model to frame motivation as a combination of *approach* (seeking success) and *avoidance* (avoiding failure).

The concept of motivation seems to be closely linked to Activity Theory's notion of Motive, and the idea that Motives – via the object of activity – *motivate* as well as *shape* an activity. There seems to be no literature discussing the conceptual implications of a distinction between intrinsic and extrinsic motivation on activity theory, though, and an account of approach vs avoidance aspects of motivation is similarly absent. This absence is not too surprising, given Activity Theory's focus on purposeful, thoughtful human activity; but the extent to which issues of motivation, and other facets of bounded-rationality, seem to dominate theories of health behaviour change suggests that an effort to incorporate these aspects into Activity Theory would be necessary in order to successfully apply it as a means of thinking about or designing wellbeing technology.

In Engeströmian Activity Theory, where activity systems are construed as inherently societal and carried out communally [93], the motivations of the individual may not align with the motives of the activity system as a whole [174]. Even from a more typically HCI perspective, in which activity may be carried out by a single individual, technology *users* are typically not the technology *creators*, and so a similar misalignment may occur; the motives that are assumed, or deployed, by a technology's designers are not necessarily the motives that led the user to adopt it. To the extent that behaviour change technologies can imbue an activity with motives from a third-party, they have the potential to bring the social dimension of "individual" ICT-mediated activity into sharp relief.

4.1.3 Reflection

Reflection – the process of critically considering one's past actions – is, quite intuitively, part of deliberate personal development. Reflection also forms a major component of many technological behaviour change interventions, such as the Apple Watch, Fitbit or Jawbone devices. Indeed, a whole industry now exists around the idea of "fitness tracking" – the apparent assumption is that, armed with data, individuals will be able to reflect on their physical activities and improve their fitness.

Hallnäs and Redström propose an approach called "slow technology" that is intended to support user reflection [79], and which focuses on "moments of mental rest rather than efficiency in performance." They contrast this active exploitation of calm moments with entirely passive approaches like ambient displays, and with "fast" technologies which are designed to make time "disappear", suggesting that deliberately building-in opportunities to reflect and think by providing the necessary time. Their examples are primarily in-place installations, but it is an intriguing concept to consider in the context of personal (often wearable) fitness and wellbeing devices.

Lee at al [111] explore reflection in the context of personalization, arguing that the "freedom to interpret technologies" allows those technologies to be appropriated to the user's own specific needs. They created a planning tool, "Fitbit Plan" that prompted users to reflect on why their walking goals were important to them and showed that this increased daily walking, and elicited a very diverse range of motives.

Reflection has been deliberately applied in support of behaviour change in systems such as Nutriflect [151], which invited reflection of household shopping patterns against personal nutrition goals. In both of the above cases – Fitbit Plan and Nutriflect – there is a notion of some *ideal,* a sense of what one aspires too; such as long-term physical or mental fitness or weight loss.

Reflective Practice, a technique originally proposed by Schön, is a technique that is used to, for instance, help clinical or education students develop their skills [158]. Research has shown that these reflective skills can be developed through, for instance, the preparation of skills portfolios. Li *et al.* apply self-reflection to physical activity, and demonstrate that incorporating contextual information helps users to understand factors that influence their physical activity, and to identify opportunities to increase it [114].

Existing work seems to confirm that reflection is a useful design principle for wellbeing technology, and that reflection can be promoted through the interactions that the technology affords and creates.

4.1.4 Summary

Although conceptions of wellbeing exist that are broader than objective models of health and fitness, they are not well represented in existing behaviour change devices. Nor do these devices seem to give much attention to the broader motives for engaging in practices such as exercise; the very diverse motives and values that Crossley identifies among gym goers [48] seem to be missing from discourse about how to motivate people, through technology, to be more active.

In this section, I consider broader activity-level descriptions of people's motivations for engaging in behaviour change, and the role that techniques such as persuasion play within that activity, subordinate to the purposeful and deliberate motives that initially encouraged them to use a device like a FitBit. Moreover, I reflect on the role that devices like a FitBit take within that broader activity.

4.2 Wellbeing Survey – Values and Goals

I began investigations into wellbeing using a large online survey that looked at factors that theory suggests might be related to wellbeing. This investigation deliberately straddles the distinction between subjective and objective models of wellbeing; we asked people to self-assess their own sense of wellbeing, but also to comment on the things that matter to them.

The survey had a dual purpose; first, to collect quantitative data on some hypotheses that I, with colleagues, had about some objective elements of wellbeing, such as the extent to which things like body mass index (BMI) are linked to self-assessed (subjective) wellbeing. Confirming that such a link exists would justify the inclusion of interventions aimed at physical health and fitness within a broader category of devices aimed at wellbeing.

Second, to collect qualitative data on broader aspects of wellbeing-oriented activity, such as the goals that people have and the sources of support that they draw upon when trying to reach them, and the things that people think are important to them in a very general sense.

The study presented here, does two things: First, it establishes, empirically, that wellbeing is a construct which encompasses, but which is broader than, physical heath and fitness. Second, it begins to unpack the non-health factors that are involved in wellbeing, and situates wellbeing-oriented practice within a social environment that provides a very wide range of support in the form of motivation, information and tools.

> *This work was undertaken in the Summer of 2013 while working as an intern at Microsoft Research, Cambridge. In collaboration with m.c. schraefel, Natasa Milic-Frayling (Microsoft Research) and Matthew Kay (University of Washington). This chapter is based on a paper prepared for submission to CHI 2014.*

4.2.1 The Big Five

For the purposes of this study, we conceptualised objective wellbeing as a five-factor model, based on existing research into wellbeing. The Big Five model proposes that five areas of activity are crucial for human survival: Movement. Food, Cognitive Challenge, Socialisation and Sleep.

The Big Five provides a simple conceptual model for evaluating aspects of objective wellbeing. It firstly allowed us to assess how effectively our participants' everyday practices are likely to support their objective wellbeing, without needing to ask and interpret potentially invasive questions about their health. Rather than asking about disease and discomfort, we can ask about sleep quality and level of physical activity, for instance. Secondly, it allowed us to make such an assessment even in the absence of specific health conditions, which is beneficial as poor dietary or exercise practices could take years to give rise to any specific diagnosed health "problem".

4.2.2 Research Questions & Methodology

There were four specific research questions that we aimed to address through this study:

1. Do objective factors of wellbeing, like BMI, influence self-reported subjective wellbeing?

We wanted to test for a link between reported, subjective wellbeing, and objective measures of wellbeing like BMI. This question relates to the earlier assumption that, because of our shared physiology, subjective human wellbeing will have some commonality, determined by our physiology.

2. Does awareness of the "big five" model of physical exercise, sleep, eating, socialisation, and challenging cognitive activity help to predict self-assessed wellbeing?

Demonstrating higher subjective wellbeing among that people who already consider all of these five factors to be important would provide some evidence that they are worthwhile areas of practice for design to focus on in order to improve wellbeing.

3. What are the factors that people identify as being important to their wellbeing?

We are not solely interested in designing for objective wellbeing, and wanted to understand more about the *subjective* factors that people identify as being related to their wellbeing.

4. What sources of information do people draw on when trying to change aspects of practice relating to their wellbeing?

Sources of information form an important part of behaviour change, and of activity systems more generally. We wanted to see which sources of information or support participants were already making use of when trying to alter their behaviour.

To answer these questions, we conducted an online survey and asked participants questions about their general sense of wellbeing, the things that are important to them (an open-ended values question) and whether they have specific goals. We then asked, in greater detail, about their past experiences, current practices and future goals relating to two out of the big five (randomised between participants).

Factor	Coefficient	SD	$F_{1,356}$	
importance of Big 5	0.302	0.100	9.17	**
age	0.018	0.006	8.38	**
male	0.413	0.158	6.78	**
personal trainer	0.367	0.198	3.46	†
BMI	-0.059	0.016	13.36	***

† p<.1 (trend toward significance) *p<.05 **p<.01 ***p<.001

Table 14: Correlation of five factors: the importance of the Big Five categories, age, gender, whether a person is a personal trainer, and BMI with the self-assessment of wellbeing (from 1–10).

We recruited participants through a range of international sources, including personal social networks and colleagues, but also by reaching out through contacts with health and wellbeing related blogs and social media groups. In total, more than 1100 participants began the survey. After removing responses that were substantially incomplete, 682 responses remained and it is from these 682 responses that the statistics in this paper are drawn.

4.2.3 Results

Participant age ranged from 17 to 75, with a median value of 40 (SD 12.28). 353 participants identified themselves as male and 324 as female; 5 participants indicated that they thought of themselves in another way. Because we solicited input from lists where people may professionally coach others around wellbeing, we asked people to say whether or not they were trainers. 133 respondents (18%) indicated they there were coaches or personal trainers.

i Objective Wellbeing Factors

We wanted to see if current physical state had any correlation with self-reported wellbeing (research question 1), so we asked for self-report of two standard markers of health: Body Mass Index (BMI) and Activity Rating. BMI is a score based on weight, height and gender and is used as an assessment of obesity. The mean BMI of the survey participants, at 25.6 with SD of 4.84, averages as what is considered “overweight” but only just. As a population, the BMI distribution of the group maps almost exactly onto the 1990 distribution for general US population with mean 25.44 and SD 4.88 [144]. By 2010 it is

28.7 [62]. Trainers' BMI averages significantly lower (24.44), but just on the border of normal and overweight.

We used the mean of two 10-point Likert items (where 10 is "very good") to assess subjective wellbeing. We then conducted an ANOVA to evaluate the effects of various factors on subjective wellbeing (summarized in Table 14). Men typically rated their wellbeing as higher than women (Men: 7.26 SD: 1.48 Women: 6.98 SD: 1.68); older participants also had a higher wellbeing rating; and those with a higher BMI tended to rate their wellbeing as lower.

We also wanted to see whether people who rated the Big Five as important reported higher wellbeing (research question 2) The importance of the Big Five categories to each respondent was measured by asking them to consider their own goals and different scenarios in which a friend or a colleague wish to improve some aspect of wellbeing (e.g., concentrating on work), and to rate how important they felt each of the big 5 factors was to each scenario on a 6-point Likert item (where 6 was "very important"). We then took the mean of these ratings as an overall assessment of importance.

We found that rating aspects of the Big Five as more important to achieving their own goals, and as important to the presented scenarios, was associated with increased self-reported wellbeing.

ii Wellbeing Factors & Goals

For the qualitative aspects of the study, I undertook a clustering exercise to identify common themes in the text answers that participants had provided. Because of the large number of responses, which made coding of the full dataset impractical, this analysis was performed on a random sample of 198 responses. These themes were assigned codes, which were used to classify each answer into the themes that described it. Some answers were allocated multiple codes, as they mentioned multiple themes.

This qualitative analysis covered three areas: the things respondents indicated

GOAL	*N*
BODY	145
CAREER	129
FAMILY	92
LEISURE	79
FINANCIAL	78
EDUCATION	53
MOVE	20

IMPORTANT	*N*
OTHERS	149
HEALTH	115
RELATIONSHIPS	47
FEELING	30
GOD	29
CAREER	26
FINANCIAL	23
LEISURE	21

Table 15: GOAL and IMPORTANT themes, and the number of participants that mentioned each.

were important in their life (IMPORTANT), the goals that participants had for the next 12 months and 5 years (GOALS), and the resources that participants had drawn (or would consider drawing) upon to reach those goals. Given the large number of responses that were collected, we extracted a random sample of 198 responses for thematic analysis and coding. In the survey we asked a range of quantitative and qualitative questions focusing on participants' in-bodied wellbeing practices, their goals and values, and a range of demographic questions.

We asked participants to list, starting with the most relevant, three things that were important to them, and three goals that they had in the coming year. When asked about their goals, the top three areas were body 73% career 65% and family 46%.

Themes with at least 20 responses in are shown in table 15. Each theme in the Goal and Importance sets is briefly described, below.

Goals

The themes here relate to the Goals that people had set for themselves in the next twelve months and five years. These questions were designed to elicit areas that people felt were important to their subjective wellbeing, and which they were actively pursuing.

Body: This theme included references to body composition and strength.

Career: References to career progression or performance at work.

Family: References to family members or the quality of interpersonal family relationships.

Leisure: References to non-work relaxation and leisure time, including specific leisure activities and references to things like "quality time".

Financial: References to income and financial stability; this code was differentiated from *Career* by explicit reference to salary in the context of a job.

Education: References to learning new skills or educational attainment; for instance completing a degree or qualification. Conceivably, this theme could overlap with career if a qualification was in support of a promotion, although no edge cases were encountered in practice and so fine differentiation was not necessary.

Move: References to movement or physical activity. We found that references to exercise were framed in terms of body goals (the Body theme) or with respect to movement itself, such as running a particular time, or achieving a particular weight lifting goal.

Important Factors

We also asked participants "What would you say are the three things in life that most matter to you?" This question was designed to elicit general factors that they consider important to their subjective wellbeing. In particular, the framing was designed to capture factors that they may not have had specific *goals* related to. Conceptually, this is an important distinction as many important factors may be in a state of maintenance rather than active goal-oriented improvement. The themes with at least 20 responses were as follows:

Others: References to other people's wellbeing; including one's own role – such as being a part of a community, or supporting other people – or the general wellbeing of others.

Health: References to one's own health and fitness.

Relationships: References to interpersonal relationships; such as having supportive or good relationships with friends and family. This code was applied only to responses that mentioned relationships specifically, to differentiate it from the *Others* code that refers to the wellbeing of others. This distinction was drawn because one could conceivably work to improve the wellbeing of others without having a direct relationship with them; through community volunteering, for instance.

Feeling: Some participants referred to a sense of wellbeing, or feeling good, as being important in itself. That is to say, they gave a sense of subjective wellbeing as being important in itself. This is a little bit meta, and points towards a model in which personal wellbeing is distinct from some other sense of duty or responsibility.

God: References to god, religion and spirituality.

Career: References to one's career or job.

Financial: References to financial matters, typically financial security. This included things like owning a house.

Leisure: References to leisure time; including phrases such as "work-life balance" that allude to having an appropriate degree of leisure (or at least non-work) time.

iii Support Resources

The analysis of people's goals, and things that they think are important in life, yielded a range of factors; many of which might be considered values. We also wanted to understand, with respect to the goals that people had, how they might find support and information to reach that goal. Conceptually, while these things are not valued in

themselves, they are things that seem to be *practically* important to people; and so begin to inform our enquiry into ways that wellbeing can be supported through tools and technology. We refer to these things as *Resources*.

We asked participants in an open question what kind of support or information they had sought or might seek in future to help them achieve one of their 12-month and one of their 5-year goals. We received a range of replies, the top themes from which are listed in table 16 and discussed below. Notably, not mentioned at all in the analysed sample were any kind of fitness application or trackers, although technology-related resources such as the internet and DVDs *were* mentioned. We had specific hypotheses about the Trainers in our sample, and found that they were significantly *more* likely to seek help from other experts ($X^2 = 21.72$, $p < .001$) and significantly *less* likely to seek help from family/friends ($X^2 = 11.30$, $p < .001$).

Resources

Expert: References to seeking information or advice from an expert or professional as a one-off or intermittent activity.

Books: References to using printed books, or a library, as a source of information.

Internet: References to using the Web or Internet to find information.

RESOURCE	N	RESOURCE	N
EXPERT	52	EXPERT_ONGOING	22
BOOKS	52	DVD	18
INTERNET	49	COMMUNITY	15
FRIENDS	35	NOTHING	14

Table 16: RESOURCE themes, and the number of participants that mentioned each in the n=198 subsample.

Friends: References to seeking support and advice from friends and family.

Expert_Ongoing: References to the support from an ongoing expert relationship, such as a coach or personal trainer.

DVD: Some participants mentioned DVDs and video as a source of information; for instance video-delivered fitness programmes.

Community: References to community support, such as weight-loss groups.

Nothing: Some participants indicated that they had no particular need for additional support or information to reach their goal.

iv Success in Change

We were interested to see how success in implementing *change* varied across aspects of the Big Five.

Aspects of Big 5	Coefficient	SD	t	
MENTAL – FOOD	-1.16	0.41	-2.835	*
SOCIAL – FOOD	-1.27	0.41	-3.021	*
MOVEMENT – MENTAL	1.11	0.41	2.702	†
SOCIAL – MOVEMENT	-1.22	0.42	-2.895	*

†p<.1 (trend toward significance) *p<.05

Table 17: Relative success at change between aspects of Big Five. All pairwise contrasts were conducted using Tukey's HSD; we report only the significant results here.

Each participant was asked whether they had ever tried to make a change to their practice in relation to one of the Big Five (each participant was randomly assigned only one of the aspects of the Big Five). They were asked to rate how successful they were at this change on a 10-point Likert scale (where 10 was "very successful"). We used an ANCOVA to evaluate whether participants who were asked about different aspects (varied between subjects) reported different levels of success. We included BMI, gender, trainer, and age as covariates in this model.

The aspect of the Big Five considered by the participant had a significant effect on their success ($F_{4,357}=4.90$, p<0.001); post-hoc pairwise contrasts are summarized in Table 17. BMI also had a significant negative relationship with success ($F_{1,357}=11.86$, p<0.001), and males reported higher success on average ($F_{1,357}=10.79$, p<0.01).

4.2.4 Discussion

I now consider the results from the perspective of the four stated research questions, as well as discussing some interesting aspects that emerged during analysis. I begin with some observations about the group of participants that were recruited.

i Participants

While our respondents have an atypical BMI for 2013, thus being considered more healthy than the current norm, they are still interested in body oriented goals; health is important to them, usually only after family. While trainers are more aware of the role of the Big Five in affecting other goals and aspirations than non trainers, all demographics have shown that they have tried to make a change to each area of the Big Five. Critically, the majority reflect that having tried to improve each of the Big Five, they have found food and movement related practices easier to change than cognitive performance, social engagement or sleep.

ii Objective Wellbeing Predictors (RQ 1 & 2)

To answer research question 1, *Do objective factors of wellbeing, like BMI, influence self-reported subjective wellbeing?*, we considered several aspects of objective wellbeing. As reported in the results, we found that BMI was negatively correlated with self-reported subjective wellbeing. This is not surprising, as higher BMI is often associated with health problems such as diabetes and cardiovascular disease. However, this result is also interesting in context of the qualitative results. We found that many participants had goals around health, fitness and body composition. It is not much of a stretch to hypothesise that BMI also impacts on people's own judgement on their physical fitness, particularly given that many reported goals relating to weight loss. The results suggest that BMI does not impact on subjective wellbeing merely through physiological pathways, but by interacting with values such as a cultural pressure to be lean.

Also in relation to research question 1, and in direct answer to research question 2 (*Does awareness of the big five model help to predict self-assessed wellbeing?*), we tested whether rating elements of the Big Five as important was related to self-reported subjective wellbeing. As shown in the results, those participants who rate the Big Five highly also report higher subjective wellbeing. This suggests that an awareness of the Big Five factors promotes subjective wellbeing.

There are differences between how successful participants were in achieving change in each of the Big Five areas, though. The pairwise comparisons shown in Table 17 indicate that our participants had greatest success in changing aspects of their practice relating to Food, followed by Movement, Cognitive Challenge and Socialisation. None of the pairwise comparisons against Sleep were significant, indicating that it is likely to be in the middle of the ranking. Relative rankings are visualised in Figure 13.

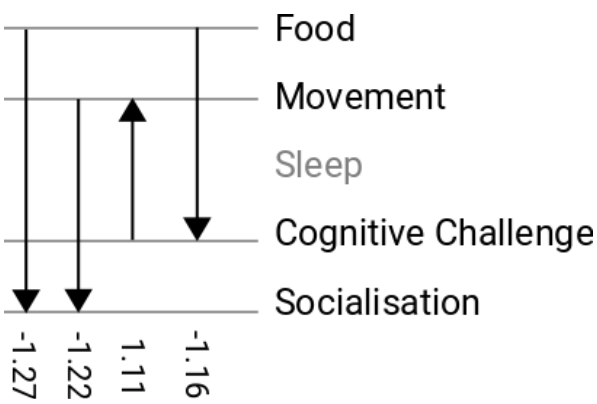

Figure 13: Pairwise comparisons between success in change in different elements of the Big Five provides a ranking of which elements participants reported greatest success in

Together, the results demonstrate a link between objective measures of wellbeing – like BMI – and an objective model of wellbeing practice (the Big Five) with *subjective* wellbeing; how "well"

people *say* they feel. They also demonstrate significant differences in how easy it is to effect change in each of those areas, though.

iii Subjective Wellbeing Components (RQ 3)

Despite the results confirming that there are objective components to self-reported wellbeing – such as BMI, and an appreciation of the Big Five area, the qualitative results also show a great degree of heterogeneity among participants when it comes to other aspects of their lives that they say are important. Although health was a frequently cited value, it came second to "Others" – The welfare of family, friends or the community. Other common values include the strength of interpersonal relationships or having a general sense of wellbeing.

These diverse values are reflected in the diversity of goals that participants reported. Although, as we'd expect given the surveyed population, many participants had body related goals, many also had career or family-oriented goals. Some had health related goals (distinct from body-oriented goals) and others specific goals about their movement practice.

The fact that these aspects are cited by participants as being important to them would justify their description as *values* for many purposes, but many of them also have striking similarities to the values in Schwartz' work on universal values. Health and Financial security are elements of Schwartz' *Security* value, for instance, Career encompasses aspects of *Achievement* and *Power*.

These results suggest that any form of technological wellbeing support that attempts to move past objective models of wellbeing into more subjective conceptions will need to be flexible to these very different values, motives and goals. However, they also suggest that there are a wide range of motives within the population that could be leveraged to motivate behaviours that promote objective wellbeing. One participant felt that their family, dogs, and health were the most important things in their life. In that example, there is a clear opportunity to leverage time spent walking dogs as an opportunity to improve or maintain physical health, and to involve family in that activity. Another participant wanted to see interesting parts of the world, be active, and enjoy good food; again a set of diverse motives that while important in themselves, also have intriguing intersections with one another, as well as with objective wellbeing.

iv Sources of Information

Participants drew on a range of information sources in the course of pursuing their goals. Seeking support from other people – often experts – was very common. Intuitively, the ability for another person to offer guidance and to apply knowledge to specific circumstances provides benefits over and above *passive* information sources such as books or most web pages. One participant explained they would *"work with a strength coach to lay out a better plan"*, pointing specifically to a coach's ability to synthesis knowledge into a plan, rather than just providing information for the participant themselves to apply.

A significant minority of participants noted that they didn't need any information resources to help them achieve their stated goal (the NOTHING code). For instance, one participant noted that *"I know what I need to do.*" It's difficult to ascertain, from the responses given, whether these participants have the necessary knowledge, or don't know that they have missing knowledge. The same participant had drawn on information in the past, suggesting that they'd already developed some understanding and formulated a plan towards their goal. Another participant listed some sources of information, but caveated it with the comment *"I actually know what to do and how to do it, I just need to get back into a routine."* Another participant had a more negative experience of trying to find information, saying that *"I cannot take one more piece of information into account for my recovery. There is so much conflicting information on nutrition that … sometimes I just get frozen. … I no longer go to expensive coaches or peers since I know what I need to do and often I know I will not comply with whatever work they give me."* This participant felt they had enough information, but had trouble resolving conflicting advice, and applying that knowledge to their everyday practice. The experience of *not* seeking information is clearly diverse, and appears to contain both people who feel they have enough information and skills, but also those who do *not*. In the latter case, the synthesis that could be provided by expert advice – such as a coach – could be beneficial; particularly in terms of learning techniques to improve adherence, which the participant suggested was a problem.

In the case of both participants who seek expert assistance, and in the case of those have found that knowledge can be hard to translate into practice, there is a sense that two types of knowledge are required. First, an understanding of the domain of change itself (for instance nutrition) but also a knowledge of how to build a sustainable practice. This is reminiscent of the distinction that Thaler & Sunstein point to between "the planner" and "the doer" [167]. It is all very well for the planner to have a set of goals and aspirations; but – like having an alarm clock that runs away – the planner needs strategies for ensuring compliance with an identified behaviour – like getting out of bed – later on. The planner has several jobs to do: Plan actions towards an Object(ive), plan for practicalities

(like packing a gym bag in advance), and planning for the ways in which the doer might later fail to carry out the plan.

v Goals

We observed significant diversity in the *goals* that our participants reported. The diversity of goals is interesting, as it indicates that participants with similar practices may have different narratives surrounding it, motivated by different values. For instance, we observed that some participants had goals to *lift a certain mass* during weight lifting; others had goals to become *stronger*. These two narratives, one focussed on the activity and one focussed on the self may well refer to the same activity, but are perhaps motivated by different values. Whether these subtly different ways of expressing goals are a reliable indication of the underlying values, and whether these narratives should be supported in subtly different ways through the use of technological support, is an open question.

These diverse framings of ostensibly similar activity are interesting for another reason, though: they have parallels with the ways in which devices frame activity, too. Compare, for instance, a scale with a Fitbit. The scale reports on personal state, one's current mass. The Fitbit reports on *process*, the number of steps that have been taken. We saw examples of goals that are framed in relation to both: some participants want to lose weight, something that could be confirmed by a scale, others wanted to ensure that they got 10,000 steps per day, something that could be confirmed with a Fitbit. Both of those processes, though, can be motivated by a sense that health is important and needs to be improved. As Fritz [70] reports, although these framings are superficially similar, and in practice may motivate very similar practices, in some cases a focus on a goal related to steps, rather than a goal related to (for instance) intrinsic physical fitness, can lead to undesirable outcomes such as distress when steps are "missed" from the count, or a missed opportunities to try different physical activities that could better meet the underlying motives.

To some extent, the different framings that we observe in the results of this study could be a methodological limitation of the survey methodology of specific questions that were used. Colloquially, usage of the word "goal" can refer to a specific quantifiable target – like reaching a step count – or to a more abstract aspiration like improved fitness. It is important, therefore, not to overstate these differences on the basis of this study.

4.2.5 Conclusions

The results confirm that objective measures of wellbeing, like BMI, have some correlation to subjective wellbeing. This justifies the inclusion of objective wellbeing promotion within attempts to improve subjective wellbeing, even beyond other worthwhile goals such as reducing pressure on health care services. Objective aspects of wellbeing seem to have an impact on subjective wellbeing through values-based as well as physiological routes; weight loss (for instance) can be motivated by self-esteem or body image, as well as by health concerns, and we observed a wide range of goals and motives among participants.

People who rank the Big Five areas of practice as being more important generally have higher subjective wellbeing. This suggests that having a greater appreciation of the Big Five areas improves wellbeing, and that the Big Five is promising as a model for wellbeing promotion.

Participants identify a very diverse set of elements as part of their subjective wellbeing. Aspects such as family, relationships, career progression and financial stability interact with concerns about work-life balance and health concerns. This diversity, and the ways in which those aspects interact suggests that finding synergies between different aspects could be an efficient way to create activities that contribute to multiple aspects of wellbeing at the same time – for instance by combining family time with inclusive physical exercise like walking.

Participants draw on a range of information sources when seeking to change their everyday practices such as what they eat, or how they exercise; but the most popular form of information is *other people*. I hypothesise that this is because other people – often experts – can not only provide information on request, but can apply knowledge to specific circumstances, and provide broader guidance. Notably, none of the participants mentioned apps or fitness devices as sources of information.

There are distinct *types* of knowledge required by participants. Those types are (at least) formulation of an Object of activity that's linked to their motives, based on some kind of domain-specific model, like a knowledge of nutrition; a knowledge of how to undertake the actions required; and knowledge of how to plan for their own adherence.

The survey had some limitations; it predated my decision to ground my analysis and investigation in Activity Theory and so was not deliberately constructed around the concepts that AT provides. A survey also provides no opportunity to clarify understanding or to explore particular participants' responses in more detail. For

instance, it is unclear why some participants were not seeking additional information, or the significance of the ways in which they had framed their goals. The next study builds on these findings, and attempts to overcome some of their shortcomings, by using semi-structured interviews.

4.3 Understanding the Values and Technology involved in Running, Cycling & Walking

The online survey reported in the previous chapter provided a broad snapshot of what our participants felt was important in their lives, the goals that they had, and their experiences of trying to reach those goals. However, as I noted previously, it also had some limitations; most notably the inability of surveys to explore particular answers in greater depth, and to clarify what participants mean when their answers are unclear. In this chapter, I report on a series of semi-structured interviews that I carried out in order to understand the experiences of two diverse groups of people: Those who have an established physical practice such as running or cycling; and people who are currently using a device such as a Fitbit.

Existing work on fitness device use has looked at aspects such as accuracy [180], adoption and abandonment [40,107] and experiences of long-term use [70]. Fritz *et al.*'s findings [70] are particularly interesting, as they point to limitations in the way that these devices cause users to think about their activity, that in some cases lead to negative experiences, despite their apparent success in building long-term physical practices like frequent walking. It seems, at least to some extent, that some of those negative effects are related to the way that the device had become embedded in the activity; feelings of "wasted" steps when the device was not worn, for instance.

Of course, fitness practices certainly pre-date the use of fitness trackers, and there is no doubt that many people build sustainable, long-term practices without these devices. I wanted to understand more about those people's activities in order to be able to compare the activities that they have, with the experiences of people who have a fitness practice that involves a fitness tracker such as a Fitbit.

The aim of this study, and this chapter, is to unpack whether the Activities engaged in by the established runners and cyclists are comparable to the Activities that are built around a device like a Fitbit, or whether there are qualitative differences to explore. And, if there are differences, can we learn anything from the former group that might help in designing technology in support of physical activity?

In this chapter, I present the results of these exploratory interviews. I highlight three areas of contrast between these two groups, and discuss how these differences may inform device design, and an Activity Theoretical view of device use within physical practices.

4.3.1 Methodology and Participants

Our methodology has been first to explore how participants explain their engagement with fitness practices. We carried out this stage of work with formative interviews with half a dozen adults from a range of backgrounds who identified primarily either as "device users" (Fitbit, jawbone, fuel) or "people with a fitness practice" like running, cycling, walking rowing, weight lifting etc.

The initial formative interviews were conducted, often informally, by myself and my supervisor, m.c. schraefel; and formed the basis for many discussions of physical practice over a period of several months. These discussions were an important source of sense-making during the research, and where I refer to "we" in the discussion that follows, it is indicative of the importance of these sessions of shared thought and discussion.

> *This chapter is based on an unpublished manuscript that was authored in collaboration with my supervisor Prof m.c. schraefel. I personally conducted all of the interviews with participants included here, and their analysis is my own; not withstanding the afore-mentioned importance of the discussions and opportunities for sense-making afforded by m.c.'s keen interest and collaboration in this work.*

From these preliminary open ended interviews, based on discussions and reflection, I developed questions to guide semi-structured interviews. users) differed; in part because knowing that the Fitbit users all had devices allowed the inclusion of specific questions about those devices, and partly because the Fitbit interviews took place, chronologically, after the Runner/Cyclist interviews, and I was able to refocus these later interviews on specific areas that were of interest. This allowed the later interviews to be considerably more targeted, a shorter in duration.

Interviews with the R/C group covered five main areas. These areas were not typically conducted in order, as conversations took different courses and different aspects came up in conversation in different orders; although I always started with area 1.

1. Aspects of life that are important to the participant; broadly framed.
2. Whether the participant had any plans relating to their physical practice, and how those plans were formulated.
3. How the participant measured and reflected on their practice; including the use of any devices.

4. Any resources, such as social support, equipment or locations that are important to the participant's practice.
5. The reasons for engaging in the activity (initially and at currently), and whether the activity serves any other purposes; like cycling for transport purposes as well as exercise.

In the case of the F group, the interviews covered four main areas. Again, the order varied between interviews based on the course that the conversation took; but I always began with area 1.

1. The device that's used, and any past device use.
2. The activity (or activities) that the device was used as part of. I was conscious that initial questions about the device itself might create an artificial focus on the device, by priming participants. To overcome this, and encourage them to think about their practice more broadly, I asked them about ways in which their activity is (or would be) different when *not using* their device.
3. The participant's reasons for starting physical activity, and any ways in which they felt their device was helpful with respect to those reasons. Again, I tried to make participants consider this broadly, and not just to think about their adoption of the device itself.
4. What the device measures, and how (if at all) those measurements are used; which measurements the participant pays most attention to. Whether there are ways of reflecting on progress other than the device itself.

These questions, not surprisingly, are similar to those used by Fritz & colleagues to explore properties of engagement for longer term Fitbit users [4]: we are all keen to understand what scaffolding helps whom, when, to stay engaged with their fitness.

We recruited participants from open calls posted to social networks and email invitations to mailing lists that reached across multiple organizations. We also deliberately invited participants we knew had health practices either because they talked about their Fitbit or their wearing them was obvious, or whom we knew went to the gym, ran, cycled and so on. In total, we completed 15 30-90 minute interviews. These interviews were recorded and thematically coded. In the end, participants came from three groups: those who identified primarily as cyclists (C), as runners (R) or as Fitbit users (F). Participants are described in Table 18, numbered to identify the group they are from.

P#	Main Activity / Device	Gender
R1	Running	m
R8	Triathlon	m
R11	Running	m
C6	Cycling	m
C7	Cycling	m
C9	Cycling	m
C10	Cycling	m
F1	Fitbit Charge	f
F2	Fitbit Charge HR	f
F3	Fitbit Charge HR	f
F4	Charge	m
F5	Fitbit Blaze (previously Charge HR and Apple Watch)	f
F6	Fitbit Charge	f
F7	Fitbit Charge HR	f
F8	Fitbit Charge HR	f

Table 18: Participant Information

4.3.2 Findings

In this section I present the main themes that emerged from the interviews that were common across groups. We had anticipated that there would be areas unique to each group, but that was not the case. Instead, what emerges are different ways of engaging with these concerns between the R/C and F groups.

i Health Relationship

All participants, across groups, were motivated to engage with their activities for a variety of reasons, but all but one participant shared a common motivator for engaging in a fitness activity: concern for their health. In that unique case, health still played a role, perceived to be a positive side-effect and enabler of/for an enjoyable activity.

Health was usually the first reason for participants to engage in a fitness practice. Several participants, for instance, started running or cycling because of a health concern. For

most of these participants, this was expressed as a desire to lose weight. C10, who stopped exercising following an injury, was motivated to *"take it more seriously"* again when he noticed that he was *"getting out of shape"* and *"feeling heavy and lethargic"*. R1 described how he had *"reached a point where I couldn't go on,"* from a combination of stress, health concerns, being overweight and poor quality interpersonal relationships.

Two Fitbit participants (F5, F6) specifically mentioned obesity and diabetes as health problems that regular walking would help them to avoid. F5 explained that *"getting ten thousand steps means I won't get fat, get diabetes and die."* Weight loss was also mentioned by F1, F2, F4 and F7. F8 had a more acute medical motivation, and was using the Fitbit to encourage walking as part of her recovery from a hip replacement.

However, health was not the initial motivation for all participants and nor did all participants cite the health benefits of the activity as being particularly important to them. One participant (C1) expressed that although they did feel that cycling (and swimming) had a fitness benefit, that was primarily only important as fitness was necessary to do those activities. For this participant, fitness was not an end merely an enabler for these activities that provided enjoyment: "*The only reason I want to be fit is so that I can do these things, so it's not really a priority for me that it does keep me fit it just means that I can ride my bike quite well."*

ii Everyday Context: Synergies & Antagonism

All participants also mentioned how their fitness practices interacted with other activities in their daily lives. This interaction was manifest in two ways: (1) getting a double benefit from one action and (2) the opposite: having conflict between two seemingly competing activities

A synergistic benefit was noted when riding a bike to get to work for cyclists, or walking to work "to get steps" for Fitbit'ers. These actions enabled two needs to be realized: engaging with an activity for their health, and getting to work. Participant C6 explained that time spent commuting by bike was not "wasted" in the same way that time spent in a car or bus would be: cycling in fact provided value to him in multiple ways, both as a mean of getting from A to B but also promoting health and sustainability and modelling good practices for his children. Interestingly, sometimes these fitness activities begin not as a health activity but as a functional activity: one participant C7 started cycling because it was "a pragmatic way to get around" when his wife was using the car.

Sometimes instead of multiplying benefits, fitness activities are presented by the participants as antagonists, in conflict with other values. R1 reported that his growing

ability to go for longer runs meant that he was away from his family more on weekends – the time not at work when he felt he should be with his family. He suggested that his family were making sacrifices to support his health activity.

Both R1 and C7 talked about resolving antagonistic demands by "negotiating" time for their activity with members of their family. This engagement represents a deliberate attempt not only to synergistically and opportunistically carry out their fitness activity when the context permitted, but to deliberately alter their context and *constraints* to accommodate their activity.

While the runners/cyclists often reported how they created these synergies or had to address antagonist activities, we more often saw the Fitbit participants defeated by antagonistic interactions that, to us, were sometime surprising.

Like the cyclists and runners, Fitbit participants stack activities into synergies, often using their devices to record activity "opportunistically". In the case of F1, the device was used to record steps taken while walking to work. This synergistic strategy, however, was interrupted by moving house; her walk was shortened and so she achieved less physical activity; she did not seek out new opportunities to replace these 'lost' steps. During a formative interview, one participant mentioned that she'd moved flats. The new flat had fewer steps, ending the high degree of stair climbing she'd previously achieved each day. In this case, the person noticed this change in both energy loss and weight gain, but, she felt, was not able to create an alternative that created the same results in the same seemingly effortless way. These examples seem to demonstrate a high reliance for Fitbit'ers in particular between the activity that the device captures (steps, stair climbing) opportunistic synergies in their current daily practice. That is, while it was common for Fitbit'ers to embrace "more" steps or more stairs as and when available, these participants seem less resilient at maintaining these levels of engagement if their physical circumstances changed, compared with the runners / cyclists. We offer this not as a judgement, but an observation.

iii Phases of Activity

Complementary to synergies between a fitness activity and a transportation activity are Phases of activity. Not all our participants spoke of this, but it came up sufficiently to be of interest. In particular, among the runners and cyclists who also raced, they would deliberately focus on their training in the run up to a race. Their use of monitoring devices seemed to be connected with that locus of attention.

C3, a participant who raced regularly, for instance referred explicitly to a "fitness phase" when a heart rate monitor would be added to his tool use. R8 's idea of phases was also apparent: *"At the moment, it [using the technology to create a log] is quite infrequent. Most of the time I do use my Garmin [a running wrist computer] to log, but because I'm not really training for anything, and am slightly injured, I'm just not bothering at the moment. I would definitely use the log for a race or a functional threshold session."* C9 mentioned a "fitness phase" where fitness goals such as waking heart rate were important measures. Devices were typically used more deliberately when preparation for a particular event became important.

F3, who had been using a Fitbit at that point for 3 months, also shared that she was not particularly engaged in her Fitbit practice, but expected that to change over time and that she anticipated she would begin to use measures such as sleep more pro-actively, rather than just observing them. She explained that work pressures were preventing action to tackle observed problems in sleep and exercise, and that thus she was currently in "damage control mode".

iv Device Use

All participants used, at some point, a device or piece of technology as part of their activity. We found that participants in the Device group tended to be Device-centric, fitting aspects of their activity to the features of the device, whereas participants in the Activity group were led by their activity, but reached out to a range of tools based on their current goals and focus.

Across the participants, we found that devices supported several aspects of activity: *Measuring and quantifying* (for instance counting steps or measuring heart rate), *logging* (creating artefacts such as GPS route maps), *sharing* (for instance posting route maps on social media or telling other people about their activity), *persuading/motivating* (by providing feedback or goals), and *seeking information* (such as routes). The degree to which our two groups seem to rely on these different aspects of devices is different, though.

By and large the Fitbit users, with the partial exception of F2, used the Fitbit to measure their steps, measure "active minutes", to track progress towards daily goals and, implicitly, to suggest which physical activities they should engage in. The measures that are built into the Fitbit – steps taken and stairs climbed – implicitly suggest that walking and stair-climbing are beneficial activities. Several participants reported the ten thousand step per day target as being *"what everyone says you should get"* (F4). Other activities, cycling for instance, would likely not contribute to those specific measures. As noted by

Fritz *et al.* [4], Fitbit users often frame their activity in terms of progress towards their daily step target rather than the more abstract health motives that initially prompted adoption of the device. This was the case with our participants even where additional measures such as waist circumference and weight were being used in parallel to the Fitbit.

The runners and cyclists, instead of using devices to suggest physical activities, incorporated devices into their activity as required to achieve particular ends. For instance, C7 spoke of the importance of having GPS tracking in order to share routes with others. R8 mentioned the use of a heart rate monitor when training in conjunction with a coach in order to facilitate advice and support, but also noted that it *"takes all the fun out of [running] when it starts beeping*" and so he usually runs without it. Rather than shaping the whole activity, these devices are adopted (and discarded) in response to shifting needs, goals and priorities. The overall direction of the participants' physical activity determines which devices are in use at a particular time.

A focus on walking and stair climbing is not necessarily a result of adopting the Fitbit, and seems to be reflective more of the sophistication of the practice itself. We note that F2 had an established physical activity that centred around working out at the gym. The Fitbit was adopted primarily because of concern about heart attack and the belief that monitoring heart rate would help to identify heart problems in advance, although opportunistic stair climbing while at work had been adopted as a result of using the Fitbit. In this case, the reassurance provided by the Fitbit (even if medically dubious) was an enabler, allowing the participant to more confidently and vigorously engage in her practice.

v Side Effects and Experiences

Several participants mentioned that time spent running or cycling had a beneficial effect on the level of stress and anxiety that they felt. One runner expressed that running provided time to reflect on work and helped him to do his job better. Two participants mentioned the physiological basis for feeling better, referring to the effects of neurochemicals released during exercise. No participants cited these psycho-physiological effects as reasons that they started their practice and none suggested that they anticipated these effects. However, for at least one participant (R1) knowing that running would provide time for reflection and confidence that it would be effective at eliciting a general feeling of wellness appeared to be very important qualities of the practice.

For C6, C7, C9 and C10 time spent with a family member was perceived as valuable in itself, or because it was seen as promoting a healthy practice in that person. For instance,

C6 and C7 hoped that towing their respective children in trailers behind their bikes would encourage their children to adopt cycling later in life.

R1 expressed that running, climbing and training in the gym gave him a chance to mix with people he would not typically encounter during his job, as well as providing an opportunity to spend time with colleagues who were also interested in these activities, saying that *"another benefit, that surprised me, was I gained a wider circle of friends. So socially it has been really beneficial."*

vi Environment & Surroundings

The physical context for an activity was raised by both groups, but in different ways. Participants in the Activity group were typically more concerned with the qualitative effects of the environment – for instance the experience of being in a place or the effect of the weather. Device participants typically referred to the environment in terms of affordance towards specific goals, usually the availability of stairs to climb.

For C6 and C7, being outside in "*inspirational*" or "*beautiful*" countryside was a positive aspect of the activity. Both these participants mentioned that living in areas with nice countryside was a positive thing that made a difference to their cycling experience. C10 talked about how races provide an opportunity to visit places, often foreign countries. For him these trips provided an opportunity to appreciate "*different aspects"* of both cycling and the place itself. He recalled feeling emotional when imagining the *"hell on earth"* of WWI when cycling through Flanders in Belgium and visiting a war cemetery.

C7 and C9 mentioned how the weather impacted their experience of cycling, and how they could prepare for environmental conditions by, for instance, taking adequate water in hot weather.

F5 explained that she had found new routes between work and home that included additional steps and incorporate steps that helped contribute to her targets by affording opportunities to incorporate more exercise into her commute, and F2 referred to the stairs in her place of work that were convenient for increasing the number of flights she'd climbed.

vii "In the Moment" Motivation

A frequent subject of discussion among the Fitbit participants was how the device itself provided "motivation" to engage in fitness activity. For instance, their progress towards their step goal, and the knowledge that the device would turn (as F3 put it) *"smiley and green"* spurred them to engage in the activity 'in the moment'.

F5 found the social competition feature motivating, as she was motivated to "*beat*" her friends and family. Similarly, F7 was a member of a Facebook group that encouraged teams of Fitbit users to compete against one another, which she found motivated her to achieve more steps or active minutes.

There is a difference between this source of motivation and the ultimate goal behind the use of the Fitbit. F7 had not obtained the Fitbit in order to be able to compete with people on Facebook, but the motivation that the competition provides is nonetheless important, in an instrumental sense, to achieving her ultimate health and fitness goals by stimulating her engagement with the fitness activity.

On the other hand, despite using tools like Strava that feature persuasive elements like leader boards and achievements, the more experienced participants did not seem to rely on that external motivation.

C9 was the only one of the runners and cyclists that mentioned a specific motivational technique, saying: *"when I'm on my bike, I am happy, I'm frustrated that I don't do it more … I need to find strategies to actually do it."* C9's technique was to hold himself to account by blogging in advance about his intent to cycle.

4.3.3 General Discussion

We've observed that there are a number of differences between the practices and perceptions reported by the runners and cyclists, and the Fitbit participants. Broadly, the activity of the runners and cyclists seems to be situated in a wider range of motivations, and are potentially more resilient as a result. That is not to say runners/cyclists are more resilient than Fitbit users, nor that these robust approaches to a practice are innate: participants shared how they developed their practice often by trial and error over time. The lessons learned from the groups' engagement with their practice, however, seems to be different. Fundamentally, the more practice-focused groups – the runners and cyclists – seemed to create opportunities for their activities, whereas the device-centric group seemed more at the mercy of synergies of activity (steps, stairs) and opportunistic contexts to achieve these goals. To some extent this could be down to the relatively greater experience of the runners and cyclists, but crucially there was little evidence from the Fitbit participants that they had begun to view their practice in separation from the device itself, except for F2 who had adopted a Fitbit *into* an established gym routine.

Those participants who were exposed to multiple ways of viewing their practice seemed to have a better appreciation of what is possible, and how they might adapt their practice. For instance, not all runners or cyclists raced. Those who did foregrounded knowledge

and experience of how to vary both the intensity of their activity and their use of technology to support it, deliberately: it is not just and only 10k steps every day: some days its more or less, faster or slower – deliberately, depending on a larger context for that activity. From an Activity Theory perspective, the "modes" or "phases" that some participants referred to – primarily runners and cyclists, although also F3 – seem to correspond to distinct Objects of activity. The relative importance of the different motives changes, and the way in which devices and information are used shifts to reflect the shift in goals. This is perhaps most clearly illustrated by the distinction that several participants drew between recreational practice and racing; in the run up to a race, participants focus more on their performance – drawing on devices and expertise that are performance-oriented.

The runners/cyclists seemed more able to create opportunities for their activity, even if this meant having to negotiate time from other important activities. We saw earlier that the Fitbit participants seemed to be more willing to wait on circumstances to change either to allow more engagement, or for that matter to restrict and reduce engagement: they perceived no other options. This lack of imagining alternatives may be both the strength and weakness of the way fitness devices are currently designed: based on persuasive technology [3] they focus on attempting to habituate a very clear specific repeatable activity, like "10k steps". As such, another key aspect of these automated activities is that they are largely context dependent [7], and application specific, and are not therefore designed to create skills transfer. If the context changes such that fewer stairs are available to climb as part of one's daily activities, fewer stairs get climbed.

We also know that, despite manufactures claims about growth of these devices entering the market, not everyone picks up a fitness tracker, and among those that do abandonment is relatively high [110]. That said, not everyone who starts to run, sticks with running. Notoriously, not everyone who takes out a gym membership in a fit of post-Christmas enthusiasm keeps going to the gym. Our small study focusing on successful longer term fitness engagement seems to suggest that there is value in ultimately foregrounding an activity rather than relying on external persuasion; there is value in learning to internalize what a practice feels like, rather than staying reliant on the lens provided by a device.

We suggest that, in building sustainable long-term fitness engagement, there is scope to design in support for transitions from the characteristics we observed in less automated or habituated externalized actions like counting steps to the more durable, deliberate actions like going for a walk. Based on our results, we propose three transitions that technologies could aim to instigate in the people who use them:

i From Framing to Integrating

This transition reflects the difference between having a device and finding activities that satisfy its measures and targets, and having an activity and finding devices that provide useful insights into it.

While having persuasive device that provides clear goals and cues to specific activities makes adoption easy, it also constrains possible activity and potentially discourages the adoption of other activities that could be more beneficial (in terms of health outcomes) and more intrinsically enjoyable.

We observed, among the Fitbit users, that the device tends to frame their physical activity by limiting the types of activity that they engage in and the ways that they reflect on that activity. The device encourages users to think in terms of steps taken and floors climbed, rather than how the activity relates to their underlying motives such as health improvement, or aspects such as enjoyment.

In contrast, the experienced runners and cyclists seemed to approach their activities in reverse: they first set their goals and objectives in line with their overall motives and based on their previous experience, and then adopt devices that met the needs that those activities gave rise to. Their overall practice is not dependent on a particular device, but they make use of measuring and logging when it makes sense to them to support a particular insight or aspiration.

We suggest interactive technological fitness ecosystems and devices support a transition from framing the activity through the 'lens' provided by a counter/tracker, to an activity focus or aspiration focus, supporting their users to deliberately choose measures, goals and devices that align with their own aspirations for their quality of life. This transition incorporates learning challenges, including potentially how to incorporate other measures like heart rate or sleep cycles or nutrients that are more nuanced than "must get steps." As we see from our racers who have a more sophisticated view of the phases of their activity, this learning offers more variety, more options for achieving or maintaining health.

ii From External Persuasion to Internal Motivation

Devices like the Fitbit use persuasive techniques like goal setting and achievements to provide a sense of motivation to users. However, the more experienced participants did not rely on this external encouragement and instead seemed to have internalized a sense of how their physical activity related a range of things – like socialization, enjoyment and performance – that they found motivating.

Although the Fitbit seems to have helped the Fitbit participants to increase their physical activity, this easy-but-constrained activity does seem to have negative side-effects over the long term such as a dependence on the device or misalignment between the provided measures and desired health outcomes. Essentially, users that rely on external persuasion may stop exercising when the device is no longer present, and may focus on the persuasive elements to the detriment of their desired health outcomes.

We were struck by the extent to which the more experience runners and cyclists could relate their physical activity to a broad range of motivations. Being able to account for how, for instance, running contributes to a sense of wellbeing, improves interpersonal relationships or reduces stress seems important in motivating their continued engagement in the activities.

We therefore propose that building an appreciation of the wider value that a physical activity provides, and internalizing persuasive techniques (for instance as in Cognitive Behavioural Therapy, [2]) is what we term as the transition from external persuasion – provided by a device, or perhaps personal trainer – to internal motivations such as the confidence that going cycling will be enjoyable and stress-relieving.

iii From Quantitative Measures to Internalized Sensing

The more experienced participants were able to explain how their fitness activity made them feel, rather than relying on an external tool to measure their activity for them. They are able to judge aspects of their performance and progress intuitively, based on past experience.

There are some benefits to being able to do so. As previous work (and our R1) have found, when relying on the external measures there can be a sense that effort is "wasted" if the device is not present and, as discussed above, there can be a tendency for external measures to alter how participants think and feel about their activity.

The more experienced runners and cyclists often spoke of qualitative feelings associated with their physical activity. In some cases (for instance R1), there was a sense that running reduced stress and anxiety. C10 explained that, when cycling he "can't stop smiling" and that exercise provides time to *"feel good*". C7 referred to the *"buzz of endorphins*" associated with exercise. These qualitative feelings of wellbeing seem to encourage activity.

F7 explained that without the Fitbit it would be hard to keep a track of (for instance) how many flights of stairs she had climbed. This difficulty is understandable, but it might be possible for a device to teach somebody what "enough exercise" feels like, and to help

them validate those feelings by comparing them to objective measures as they learn. Internalizing notions of performance and success – or adopting "point in time" measures like weight or waist circumference, could avoid the feeling that activity undertaken without the device are "wasted" – something that F1 recounted and which is also apparent in earlier work. Foregrounding interplay between feeling first and validating second – more deliberately.

There was a sense that some Fitbit participants had started to translate the measures used by the device into more device-independent measures. For instance, F5 explained that one thousand steps corresponded to about ten minutes of walking. C7 spoke about the sense of achievement associated with a "challenging ride", reflecting an internalized sense of the degree of challenge involved in a particular route.

There is undoubtedly a balance to be struck between internalized sensing and external measures. There is a need for any internal sense of performance, achievement or wellbeing to be objectively grounded and trustworthy, but balanced against the previously-discussed problems with entirely external measures.

4.3.4 An Activity Theory Discussion

The discussion so far has been grounded primarily in concepts that emerged from the data, and the transitions that those themes suggest could be possible. It is also enlightening to view the emergent themes, and the accounts of our participants, through the lens of activity theory. From a methodological point of view, I have found it useful to take this dual approach; first taking an almost grounded-theory approach to analysing the data, unconstrained by trying to directly fit it to Activity Theory, followed by a more integrative period of reflecting on how the original data – and the concepts that emerged – might relate to the existing framework offered by Activity Theory. A more unified Activity Theory account across the three wellbeing studies that I carried out is provided later, but there are three ways in Activity Theory is particularly relevant to this specific study, which I will briefly discuss now: First, how multiple distinct Objects of Activity are present within some participants' accounts of their practice in the form of "phases"; second, the different ways in which device use relates to, and suggests, Objects of Activity; and finally, exploring a distinction between *motive* and *motivation.*

i "Phases" as distinct Objects of Activity

One interesting area to consider are the different "phases" of activity that several participants mentioned. These different phases seem to correspond correspond to distinct Objects of Activity, and a reconfiguration of many aspects of the Activity's structure and

focus. There are two reasons that I think these phases deserve to be modelled as shifts in the Object of Activity, rather than being considered as part of a single broader Activity: First, is the extent to which these phases correspond with changes in how information and data are sought and used; and second is the very definite change in the outcomes that the participants are seeking.

The *"damage limitation mode"*, mentioned by F2 is a good example of a shift in intended outcome. In this "mode", F2 sought *reassurance* from having access to her sleep data, but was not actively attempting to use it to inform or measure deliberate change. There was a strong sense that, for this participant, the Object of Activity would eventually shift to a more active attempt to *improve* her sleep and exercise, when other factors in her life become less of a constraint.

A common example from the runners and cyclists of a change in the Object of Activity is the distinction between a "training" phase – in preparation for a race – and a "leisure" phase. These phases lead to marked changes in both the information that the participants sought, and the outcomes that they envisaged: In the run up to a competition, runners and cyclists seek external expertise (from coaches in the case of R1 and R2) and often make greater use of technology to help them gauge their improvement and performance – even though, in the case of R2, they found that technology could detract from the overall enjoyment of running. We observed a shift from enjoyment and socialisation, to (in a race phase) aspects such as performance and competition – in fact, several participants valued races precisely because they provide a controlled and repeatable measure of performance, and an exciting environment that encourages them to push themselves harder and to perform better.

It should be noted that the overall mix of motivations behind these activities do not appear to change; only the extent to which those different motives come to fore. It would not be fair to suggest that socialisation, stress-relief or enjoyment were not important to participants during a race phase, just relatively less influential.

The shifts between these different Objects are often cyclical; frequently based around an annual calendar. Taking part in the same race each year was relatively common, with less intense more leisure-oriented periods in between. Repeating phases allowed participants to increase the sophistication of their practice; for instance by seeking more specialist expertise (in the case of one participant, a research unit dedicated to marathon running), or by entering increasingly more challenging races.

ii Where do Objects come from?

The second observation I'd like to draw with respect to Activity Theory concerns the ways that technology use relates to the Object of Activity. As explained in the first transition, "From Framing to Integrating," Fitbit participants often seemed to construct their activity in response to the measures provided by the device, in contrast to the Runners and Cyclists (and to some extent F2) who integrated technology into their practice in support of particular goals or phases. From an AT perspective, the Fitbit seems to provide an Object of Activity to many of the Fitbit participants; strongly suggesting stair climbing and walking in answer to their health and fitness motives. This contrasts with other available technology, such as Cancer Research UK's "Active App" that can be used in preparation for participation in their fund-raising races, and which therefore orients towards a '*racing for charity*' Object.

The fact that the Fitbit (and other tech) seems to suggest a particular Object so strongly begs a broader question for design: How do users identify an Object of Activity in answer to their motives? And how can we design to facilitate that process? Building on the notion of changing Objects discussed previously, there could be a design opportunity to deliberately review and alter wellbeing-oriented Objects, not only in terms of degree (like setting step goals) but also deliberately suggesting phases of activity, identified based on a more nuanced understanding of the user's motives; something that is largely missing from current generation devices, which provide affordances for weight loss, stamina building and post-operative recovery, but which lack many of the learning affordances that could help users to make efficient use of those different affordances at different points in their developing fitness practices. The data from Fitbit participants indicates, for instance, a high degree of uncertainty over the significance of Ten Thousand daily steps, and how (or even if) step count might relate to Objects such as weight loss. Some participants were under the impression that reaching this goal would lead to weight loss (F1), others felt it would help them to maintain their weight, and others framed it more generally as a way of avoiding health problems later in life. In essence, there is often a lack of clarity about how the Goals presented by the device – 10k steps – are related to the Object of Activity – e.g. walking to achieve weight loss – and the user's underlying motives – such as feeling better about their appearance.

iii Motive vs Motivation

Finally, the accounts provided by participants indicate a subtle distinction between *motives* – the things that shape and motivate their fitness Activities in a deliberative sense – and *motivation* – the desire to engage in the activity "in the moment".

F7 illustrates this distinction; her fitness activity is oriented towards a weight-loss motive, tightly bound to health and body image motives. Nonetheless, she suggested that having the Fitbit was *"massively motivation, and keeps me stepping"* through the step goal and social competition features. It is quite possible that motives such as socialisation and competition could become additional motives in their own right, much as they did for R1, but there is also a strong sense that the motivation the Fitbit provides is in some sense *instrumental*; that is to say that the intrinsic motivation provided by those features is valuable in so far as it helps F7 to maintain her practice regularly in service of more abstract, deliberative, motives like health and weight-loss.

C9's use of a blog to publicise his training intentions and to hold himself to account seems to perform a similar function. The posts create a sort of personal contract, enforced in part by the sense that his readers will judge him if he does not keep to his stated plans, and so provide motivation to carry those plans out. This a deliberately contrived situation, though, which itself exists only because of other more fundamental motives like fitness.

This is not, on the face of it, *quite* the same as a distinction between intrinsic and extrinsic motivation. Motives like health are not external to F7, and have not been imposed by some other actor; but becoming healthier is a long-term activity: it happens slowly, it is hard to measure (especially as the ultimate goal is often to avoid *developing* a health condition), and the pay-off may be far in the future. A health motive may provide little intrinsic motivation, despite being the ultimate need to which the Activity is oriented; and conversely reaching a step target for the day may have little long-term relevance, but provides the impetus necessary to engage in an otherwise potentially tedious activity like walking.

4.3.5 Conclusions

The fitness practices of the two quite different groups of people explored in this study raise several interesting contrasts that contribute both to the design of fitness technology directly, and to enriching an Activity Theory account of fitness practice and the role that technology can play in it.

I observed three main contrasts between groups of participants (typically between R/C and F, although with some exceptions).

1. There is a difference between participants who constructed a practice based on a device, like a Fitbit, and participants who had a practice and incorporated technology use into it as their goals dictated.

2. Some participants relied on externalised motivation, typically provided by the 'persuasive' features of a Fitbit, while others had internalised their motivations to a greater degree, and were able to derive motivation to engage in their fitness practice without these features.
3. Some participants were more oriented towards quantitative measures (like step count or floors climbed), while others used qualitative measures (like noticing how stressed one feels, or recognising the "buzz of endorphins" during exercise). Qualitative measures are more internalised, are less reliant on a device, and perhaps help in the process of internalising motivation mentioned in contrast with point 2 (above).

Based on those contrasts or transitions, and the themes in the data, I identified three implications for an Activity Theory-based account of fitness practice.

1. The Object of Activity in fitness practice can change, and this is reflected as different "phases"; such as the distinction between preparing for a race and running for leisure. Shifts between objects may be cyclical, such as training phases before races with leisure-oriented phases in between.
2. An Object of Activity can be suggested by a tool, through the lens that it provides on a fitness practice. Counting steps and flights of stairs strongly suggest walking and stair-climbing, instead of entering and training for a race, for instance. Obtaining and using a Fitbit can, in *itself*, be the initial construction of an Object that answers to a fitness motive.
3. There is a distinction between the ultimate motives behind an activity (like health) and things that provide the *instrumental motivation* that's required for someone to get up and do that activity.

We'll revisit these ideas later on, but I wanted to explore some of them in greater detail. In particular, I wanted to unpack the different motives that fitness activities can answer to, and the different *instrumental motivations* that can be leveraged to help sustain those activities. The next chapter reports on a Speculative Design exercise that focuses on the idea of motive and motivation more specifically.

4.4 Speculatively Designing Fitness Trackers for Alternative Values

The two previously reported studies showed that fitness practices can be motivated by a whole range of things, many of which seem to fit the definition of a value. Participants were able to relate their exercise not only with their health and fitness, but also with things like the quality of their family and social life, climate change, and setting a good example to people around them. The second study in particular began to unpack some ways in which these fitness practices can be modelled using Activity Theory; in particular the ways that these devices can implicitly provide Objects of Activity in answer to a fitness motive, and how the persuasive features seem to provide a kind of *instrumental motivation* that helps to sustain the practice, but which is otherwise unrelated to the underlying *motives*.

I wanted to explore, in greater detail, the possibility of anchoring fitness technology in non-health motives like enjoyment and socialisation; building in particular on values that seemed to motivate the runners and cyclists that I'd interviewed. There are two reasons why, on the face of it, such devices could be useful. First, they could engage with people who are not motivated to act by purely health-based narratives. Recall from Crossley's work in critical health psychology [47] that people sometimes make decisions based on non-health factors which lead them to what we might consider objectively *negative* health outcomes (like smoking or having unprotected sex). Grounding an objectively healthy activity, such as running or walking, in a non-health motive such as enjoyment or socialisation could provide a way to engage a new population in those healthy activities by providing alternative motives. Secondly, it was clear from the interviews that although health reasons had motivated some participants to start their running or cycling practice, it was not necessarily health factors that drove the continuation or expansion of that practice. Instead, new aspects such as achievement, enjoyment, socialisation, competition and leadership came to the fore. Devices that are able to engage these other factors might stand a better chance of maintaining long-term practice. After all, even once more acute health concerns have been addressed, most people would benefit from a more active lifestyle.

Designing, building and evaluating such devices would be at least another thesis in itself, and therefore doing so was not an option. Instead, I adopted an approach called "speculative design". Speculative design, as described by Auger [15], and building upon notions of design as a means of social critique and enquiry proposed by Dunne & Raby

[56], involves the crafting of fictional but believable designs that can be used to elicit feedback from potential users. In this study, I combined semi-structured interviews and an online questionnaire to gather feedback on designs that I produced for a series of wearable fitness devices. The goal was to understand how people would think about those designs, how they imagined they would interact with them, and the factors that seemed to be important to them in evaluating each device.

As a secondary goal, I wanted to determine whether the Portrait Values Questionnaire (PVQ) [161], an instrument designed to measure the relative importance of different values to a person, had any predictive power in determining which devices most appealed to people. For instance, would someone who the PVQ rated as valuing hedonism highly be more likely to pick a device that ostensibly aimed at maximising their enjoyment?

I first describe the speculative designs that I produced, followed by the research methodology, results and discussion.

4.4.1 The Designs

I produced designs for five different wearable devices. All five had several things in common, including their form and general range of features. Each one was designed to appeal to a motive that had been identified in the previous research. The devices are described below (descriptions shown are the final descriptions, as shown in the questionnaire, following some refinement during the interview stage).

The Vitality

The vitality adds new health-related features. As well as tracking heart-rate and steps, it can measure your heat rhythm to detect abnormalities, your blood sugar to help you tune your food intake, and blood oxygen to monitor your cardio and pulmonary health.

The Vitality just wants you to be – and feel – healthy! It asks how you 're feeling, and can help you to find which exercises and activities work for you, taking into account any injuries that you might have.

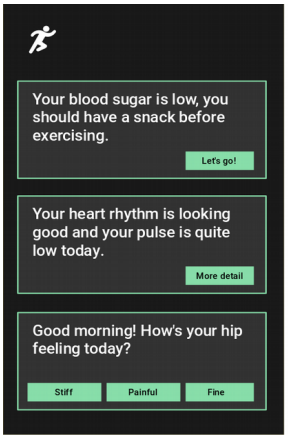

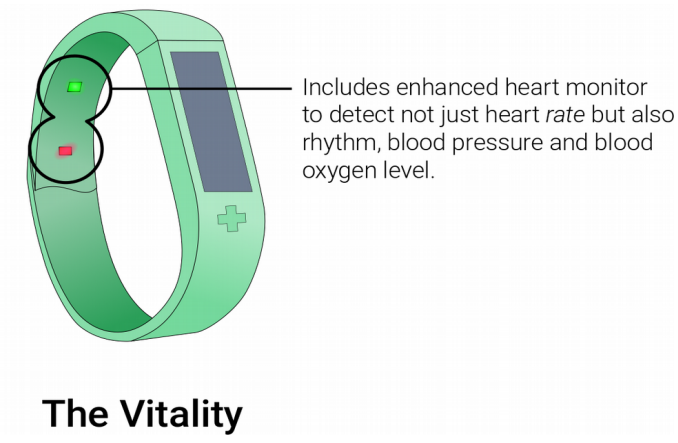

The Giggle

The Giggle just wants you to have fun! It's constantly looking for opportunities to play a game, or to find time for an activity that you enjoy. The Giggle won't ask you how much water you've drunk, or remind you to take more steps, but it might suggest that you go dancing or have a stroll outside in the sunshine! By asking you rate how happy you're feeling, and how much you've enjoyed its suggestions, the Giggle works out what you think is fun and finds new activities that you might be into.

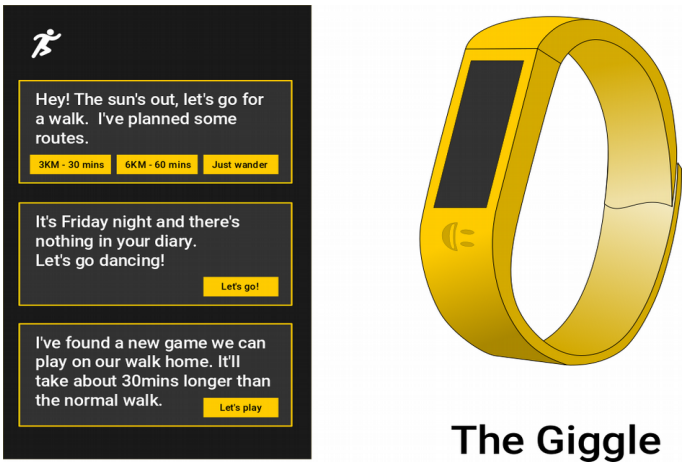

The Trend

The Trend helps you to share your fitness story with the people around you. Whether you're trying to encourage friends and loved ones to get a bit more active, or just need everyone to know about your amazing gains the gym, the Pioneer helps you to tell the world!

A small integrated camera allows you to take photos on the go, just hold up your wrist and tap your finger to instantly share what you're doing, what you're eating, or where you've been. The screen glows brighter for each 'like' that your photos receive!

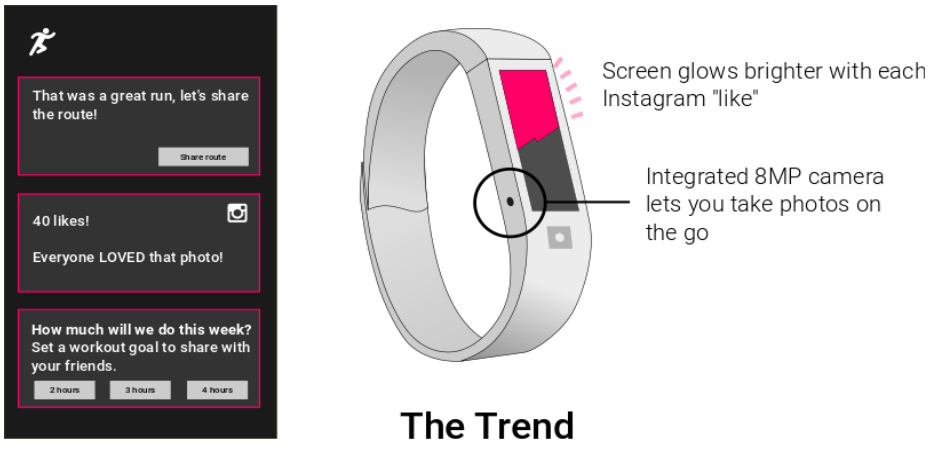

The Enforcer

The Enforcer is on your side, helping you to achieve what you want to achieve, for YOU. It knows that sometimes you don't feel like pushing yourself, but it also knows you'll feel better if you do! It includes a loud alarm clock function to wake you up in the morning, and can only be silenced once your start moving around. Armed with small electrodes on the inside of the band, it will (gently) remind you to get some exercise when you're lounging around, or encourage you to push a bit harder when you're taking it too easy.

The Butterfly

The Butterfly loves to be sociable. It tries to find activities that you can take part in with other people. It knows what your friends are up to, and will vibrate to tell you if they're exercising nearby; sometimes it will suggest inviting a friend or loved one to cycle in the forest, or go to go for a walk along the river. Other times it might suggest a group class at a local gym, or a club for an activity that you seem to like. If all that fails, and the Butterfly is feeling outgoing, it might even try to match you with a stranger you have things in common with!

4.4.2 Methodology & Hypotheses

I combined qualitative semi-structured interviews with a more quantitative online survey to evaluate the designs. The purpose of the interviews was to identify areas of interest and specific questions for use in the survey, as well as gleaning a series of more in-depth answers than are typically possible via a questionnaire. Interviews, a two-way dialogue, provide much greater scope to correct misunderstandings, check comprehension and clarify participants' statements than a questionnaire does.

The online survey gathered a much larger dataset, allowing hypotheses around the PVQ to be tested, as well as a more generalisable comparison of the popularity of the different approaches.

i Interviews

I began the study by carrying out semi-structured interviews with people who were interested in starting new physical activities, but who had not necessarily done so. I chose this demographic as broadly reflective of the sort of person who might consider purchasing a wearable fitness device. These interviews ran through each device, and asked the participants to consider which was most suited to them. I asked them to explain any aspects that they felt were particularly positive or negative, as well as asking them to rank the devices on various criteria.

Each interview was audio recorded and transcribed. I also made notes during each interview, including the answers to each of the ranking questions. This was particularly useful where participants pointed to devices on the paper materials, rather than verbalising the name.

I revised the interview protocol during the course of conducting the interviews based on my experience with the participants. For instance, introducing new rankings and removing questions that did not seem to elicit much relevant discussion. These changes are discussed in more detail later.

ii Questionnaire

The questionnaire had two main goals: First, to gather rankings and short qualitative reflections from a much wider range of participants than would be possible via interviews. Secondly, I wanted to compare the chosen devices with the results from a values instrument called the Portrait Values Questionnaire.

The online questionnaire consisted of four main sections: 1) A PVQ based on Schwartz' 40-item version, 2) an IPAQ instrument [78], which provides a measure of physical activity, 3) a set of ranking questions that asked participants to choose between the different devices and 4) a set of short qualitative answers asking participants what they felt was positive and negative about the different devices. Information about the five different devices was presented to participants after the IPAQ instrument.

Since values theoretically provide a means for assessing preferences, my hypothesis was that particular values in the PVQ might predict which device most appealed to different participants. A list of values from the PVQ, and the device that I hypothesised might most closely correspond to them, is reported in Table 22 in the results section.

I had found in the earlier studies that some values are mentioned relatively commonly by participants in relation to fitness practice, but are not explicitly covered by Schwartz' model of universal values: Health, Socialisation and Performance. I introduced some additional portraits into the PVQ instrument in order to provide a sense of how important these three more specific values were to each participant; those portraits are shown in Table19.

Value	**Portrait Statement**
Health	They want to avoid illness and are careful to avoid things that might make them unwell in the future.
	It is important to them to look after their health. They don't want to be ill.
Socialisation	It is important to them to be around other people. They like to spend time with friends and loved ones.
	They like to share experiences with others. They would rather do things with others than alone.
Performance	It is important to them that they push themselves. They want to achieve as much as possible.
	It is important to them that they reach their goals, even if it means some discomfort on the way.

Table 19: Additional values and portrait statements introduced into the portrait values questionnaire (PVQ)

I discovered during the interview phase that participants interpreted ranking questions slightly differently when they were asked about the usefulness, effectiveness (at encouraging more exercise) and overall likeability of the different devices. I wanted to explore these concepts, and the factors that participants related to each one in more detail

through the online survey. I therefore asked three different ranking questions in the online survey, along with free-text questions that allowed participants' to justify their choices. The three ranking tasks were as follows:

1. Please rank the devices in order of how **useful you think they would be.**
2. Please rank the devices in order of how effective you think they would be at **getting you to do enough physical activity**.
3. Please rank the devices in order of how much **you would like to own them.**

I used a combination of quantitative and qualitative analysis to test specific hypotheses about associations between values (measured in the PVQ) and device choice, and to explore the concepts and factors that affected participants opinions of the different designs.

The results and implications from the two phases – interviews and online survey – are reported next.

4.4.3 Interview Phase

In total, I interviewed nine participants (6 female, 3 male). My main interest during the interviews was to check understanding of the designs and to test questions prior to their use in the online survey.

i Suitability of Participants

I wanted to establish whether the participants I interviewed were within what might be considered the target audience for casual fitness devices like the ones I had designed. I therefore asked whether were currently engaged in any kind of deliberate physical practice to help contextualise their responses to the rest of the interview, and to check their level of expertise. All bar one of the participants was currently engaged in at least one regular physical practice, as summarised in Table 20.

All but two of the participants wanted to increase the amount of physical activity that they engaged in. Participant 3 did not necessarily want to increase the amount of time that she spent engaged in physical activity, but did have goals around stamina and body image that suggested she wanted to make some qualitative changes to the types and intensity of the activities that she undertook. SD5 was happy with her level of physical activity.

All of the participants had used some kind of fitness technology; either a wearable device or a smartphone app in the past, but only two participants (SD3 and SD5) still used one regularly. Broadly, all of the participants could be characterised as open to the idea of incorporating technology into their physical activity, but as being non-users or casual users of such technology at present. All of the participants were therefore within what might reasonably considered the target audience for the devices that I was evaluating in this study, and therefore all of them are included in the subsequent analysis.

Participant	Motives	Activity	Tech
SD1	Getting to work, environmentalism	Cycling to work	None
SD2	Lose weight & improve fitness	Gym & Yoga	None
SD3	Socialising, working off food, Improve Stamina, Tone body	Hockey, Yoga, Dance, Rounders, Walking, Running	Map My Run
SD4	- No deliberate physical activity -		
SD5	Enjoyment, Medical Condition is managed through strength exercise	Gym, Bouldering	HR Monitor (Gym)
SD6	Weight loss	Walking to work (in conjunction with diet plan)	None
SD7	Health; establishing a healthy lifestyle while young	Swimming, Volleyball	None
SD8	Improving body appearance	Swimming	None
SD9	Health	Walking twice a week	None

Table 20: Participants current motives, physical activities and technology use

Themes

I conducted a thematic analysis of the interviews to identify common areas of discussion. Conceptually, two areas of discussion stood out to me: *Purpose* – what the participants thought the devices would be for; and *Challenges* – things that participants felt the devices would need to take into account in order to be useful.

Participants identified two main purposes for the devices: Providing *information*, and providing *motivation*. They identified three challenges: Life Constraints (working around other activities), Adherence (avoiding abandonment the device), and Relevance (ensuring a fit between the device and the activity that they are or would be engaged in).

Purpose: Information

Participants were split over what they felt the main purpose of using one of the devices would be. They fell, broadly, into two camps: those who saw the devices as a source of *information*, and those who saw them as a more direct source of *motivation*.

SD3, in particular, articulated this division while deliberating over which device she would rank as her favourite; she was unsure whether being given suggestions about activities she could do (information) would more important than something that could somehow 'make' her exercise:

Figure 14: Additional information introduced during the online study to clarify features common across the devices.

> *The science side of [appeals to me], 'okay, my blood sugar is low, I should deal with that' rather than going 'let's go dancing' I want to go dancing, but if I'm low on blood sugar I'd still struggle. So, I like that side of it. I also like it's sort of making you get up and do the exercise, so the alarm goes off, so it makes you do something, whereas the other one is more sort of fun suggestions but you don't have to do, you don't get electrocuted. – SD3*

SD4 was fairly information-oriented, and ranked the Vitality as his preferred device. However, he also suggested that having information could be motivating in itself, as well being useful in other ways.

> *I think if I can see that sort of information in real time it's more of a motivator for me to improve, and you can also see if anything's wrong. Things like the blood sugar is god because I've got a lot of diabetes in my family, so you know that's an added benefit for that from a, I guess a medical perspective, which the others you wouldn't get that. – SD4*

A similar focus on information was apparent in his previous experience of fitness technology, too, where he explained that:

> *I [had] wanted to improve my fitness by going running, and [RunKeeper] seemed like a good way to get a little bit more out of the run. To tell you exactly how far you'd gone and how long you've been going. So, I guess overall because it gave you more information about running. – SD4*

SD1 twice referred to the device as a source of reminders to do things, a form of *context-driven information.*

> *You kid yourself that you've done more than you have, so a little reminder that, um, that you haven't been doing it hard enough is probably a good thing – SD1*
>
> *I do find that, you know, when I've forgotten to eat and things that it does affect me so sort of little reminders like [those from the Vitality] might no be a bad thing. – SD1*

Some participants were concerned about how they might ***interpret*** information that was provided by the device, in particular the *health* information from the Vitality.

> *Um, I kind of like the idea of measuring things like that, I just wonder whether it would just worry me if it told me that things weren't normal or something. – SD1*
>
> *I guess I were looking to buy [the Vitality], what I'd want to reassure myself … is how I would, as a lay person, be able to interpret the information. – SD8*

Other participants were concerned about the ***accuracy*** of any reported data:

> *I'd worry about the accuracy of those measures [on the Vitality], how accurate would it be when it was measuring those things? But, if it was accurate enough then I think that could be a good idea, but I wonder if people who had problems with their blood sugar might try and use it as a replacement for their medical devices that are designed specifically for that. That could be potentially dangerous. – SD7*

Finally, information about the *environment* around her appealed to SD7, who found locating information about (for instance) classes and swimming times to be a chore, but who had found her ability to swim was limited by public swimming times at her local pool:

> *A lot of what puts me off of trying to find something to do is that it's quite tedious to go through all the Websites and find the times and work out when and where. So, if something did that for you, that would be quite handy. – SD7*
>
> *if it could look up swimming in other swimming pools and tell me how far away they were, that'd be handy. – SD7*

Purpose: Motivation

Being able to derive motivation by using a device was a major factor for most participants; and there was a strong sense that a lack of motivation was the reason that

they don't engage in more regular exercise. SD5 and SD8 good examples of a motivation-oriented participants. SD5 explained that:

> *The main thing I would want out of one of these devices is the motivation and encouragement to actually do [exercise]; over information on what's happening with my body, because I already exercise without that information. So if I was to have more motivation in terms of the enforcer I think, yeah, it would be my favourite. – SD5*

SD8 saw the primary purpose of a fitness device as being to help him get fit, and felt this would be achieved through motivation – either his own intrinsic motivation, marshalled by the enforcer – or through *additional* motivation provided by the butterfly's social features.

> *Um, my goal is to get fit, therefore the one that is most useful to me has to the one that is most likely to get me fit, and that's the enforcer – SD8*
>
> *I would be motivated by having kind of a 'real time' 'other people are doing things', but ... it's quite easy for me to ignore that. Whereas if it's sending electric shocks to me, it's more of a prompt. – SD8*
>
> *I think that [the Butterfly] could be a good alternative source of motivation because I don't think you can have that internal motivation, unless you're an Olympian, all of the time. – SD8*
>
> *I like that it's sociable. Generally I find exercise more enjoyable and more fun if it's with other people – SD8*

There was a sense in P5's comments that motivation could be derived from introducing additional motives into her activity through the device. Since she already enjoyed activity, and felt that her health was already good, those were factors that could not provide any additional motivation.

> *This [the vitality] would be super informative, but I don't feel like it would motivate me at all. I'd just be lying down all the time, like, "everything's fine, I'll just not get up". – SD5*

> *I don't think [the Giggle] would offer me any more motivation other than the fact that I enjoy going climbing and working on exercising. Like, it would probably be fun for a week, but then 'd be like 'this isn't really offering me anything now' – SD5*

Challenge: Life Constraints

When thinking about the devices that they might use, participants considered how those devices would fit into their lives, around the other activities that they carry out. These constraints emerged in several ways. Most commonly, it was fitting physical exercise around other activities. For instance, SD1 talked about how she'd not want to be interrupted during other tasks:

> *I wouldn't want it interrupting me in the middle of things ... I would probably resent having a device that kept interrupting me and stopping me from getting things done.- SD1*

While P9 identified fitting exercise around work as a more general barrier to engaging in it – with a device or without:

> *I'd say it's motivation, well, partly motivation. Partly the level of exhaustion, fitting it in after work. So yeah, time. – SD9*

P1 had experienced something similar, but saw a device (particularly the Giggle) as an opportunity to make better use of her time, opportunistically:

> *just doing something for twenty minutes, fitting it in, just doing something before I start the next thing – SD1*

P5 identified the device itself as being a constraint on her activity. The proposed form of the device – a wristband – would interfere with her primary physical activity, bouldering:

Factors	Spearman's R	p
Effectiveness – Likeability	0.707	1.39e-75 *
Usefulness – Likeability	0.755	1.590e-91 *
Usefulness – Effectiveness	0.744	2.074e-87 *

Table 21: Correlations between rankings for usefulness, effectiveness and likeability; all correlations are highly significant.

Device	Value	p	Explanation of Hypothesis
Enforcer	Performance	0.107	The Enforcer embodies a focus on pushing oneself and achieving the goals it sets.
	Achievement	0.407	
	Stimulation	0.187	St. includes valuing "challenge"
	Self-Direction	0.110	Tested as a *negative* predictor, as the enforcer provides little flexibility or self-direction.
Vitality	Health †	0.160	The vitality is health-centric
	Security	0.118	Security encompasses physical safety and health
	Fitness †	0.156	Health and fitness are conceptually related, and often conflated
Trend	Social †	0.165	The trend encourages sharing of photographs on social media and engaging with others.
	Share †	0.251	
	Benevolence	0.313	Some runners and cyclists reported that they share in order to encourage others, a sense of benevolence
Giggle	Hedonism	0.261	The Giggle focuses on enjoyment, conceptually linked to hedonism and stimulation.
	Stimulation	0.353	
Butterfly	Social †	0.00320 *	The butterfly encourages social activity.
	Benevolence	0.00293 *	Benevolence and universalism are socially-oriented values concerning the welfare of others.
	Universalism	0.032	

Table 22: Values that I hypothesised would be associated with 'liking' each device in the survey and p values from Mann Whitney U test (= significant at p=0.05, after applying a Bonferroni correction for repeated tests to give p = 0.05/15 ≈ 0.00333). † denotes values that were added to the PVQ and are not part of Schwartz' universal values.*

> *The main downside of the activity that I do in terms of climbing …is I would take this wristband off. I wouldn't want to climb with it on my wrist. So, it measuring all these things wouldn't to a lot for me while I was actually working out, if it was just going to be in my bag. – SD5*

Challenge: Adherence

Some participants tried to predict which devices they thought they'd stick with. SD2 felt that the information provided by the Vitality would be interesting, but that she probably wouldn't use it for a long time.

Concerns about adherence were most commonly expressed in relation to the Enforcer, with several participants expressing that they'd probably take it off.

> *I like the alarm clock you can't silence unless you're moving. Um, I do think I'd just take it off and put it in a drawer if it was annoying me too much, trying to be too motivating. – SD7*

P1 was unsure about how she'd react to being regularly reminded to do exercise, but predicted that she might dislike the device and stop wearing it:

> *You kid yourself that you've done more than you have, so a little reminder that, um, that you haven't been doing it hard enough is probably a good thing but I'm not sure how I'd want that incorporated into my life and whether I'd just , um, throw it away. If it kept telling me I hadn't done much, yeah, I think I'd start to hate it quite soon, um, see it would have to know whether I'm just actually bored and would jump up and do something, and when I'm actually in the middle of something really important and don't wanna get up. – SD1*

Challenge: Relevance

P5, who already had an established activity (bouldering) was curious about whether the device could be used more specifically within that activity; for instance by recognising the type of climbing she was doing:

> *can I set targets in this that are specific to anything, it doesn't have to be just running? 'cause that would be really useful, and then keeping track of like, oh, if it could recognise what kind of climbing I was doing, like if I could input 'oh I've done this grade today' and all that kind of thing. – SD5*

iii Implications of Interview Data

Based on the interviews, I identified several implications for the online survey. First, exploring the distinction between effective, useful and likeable devices; second, revisions to the Trend device; and finally, revisions to the explanation of commonality between the devices.

Effective, Usable, Likeable

I'd asked participants to rank devices based on how much they liked them, how useful they thought they were, and how effective they thought they'd be at encouraging physical activity. This distinction was prompted by observations in some of the earlier interviews that some participants were drawing a distinction between devices they'd enjoy owning, and which they thought were more "useful". The interviews showed that these concepts were slightly different, and that usefulness potentially reflected more on informational affordances than motivational ones; but I wanted to collect a larger set of qualitative data for each of these concepts to better understand how they related to one another. This observation led to the set of three questions presented earlier, in the methodology section.

Revising the Trend

I also revised the Trend device during the course of the interviews. Initially, it was called "The Selfie" and focussed more on the act of posting images to social media. Participants in the interviews were universally negative about the device. Notwithstanding that none of the interview participants were By renaming it to the "Trend", and emphasising other social features (such as sharing routes) I hoped to give greater emphasis to the idea of *being a role model* that some participants in earlier interviews had mentioned to me.

Common Functionality

In the survey, I introduced the idea that devices would synchronise with the user's calendar to avoid interrupting them in the middle of other activities.

I found during the interviews that some participants chose devices over one another for supposed differences in their ability to do things like plan around an existing schedule, or find information about local facilities. I had not intended for there to be any major difference in those features, and wanted instead to focus on the relative motives embedded in each device. I therefore drafted an explanation of the common features between the designs to include in the survey, that aimed to minimise perceived differences around some of those aspects. The revised information is shown in Figure 14 and was presented before the designs themselves.

Having made these changes, and settled on a final set of questions, I proceeded with the online survey.

4.4.4 Online Survey Phase

102 participants were recruited via Prolific Academic and completed (or substantially completed) the online survey. 51 participants identified as male, 50 as female, and 1 participant did not answer this question. Participants ranged from 18 to 40 years old (median age 29).

i Usefulness, Effectiveness & Likeability

Each device was ranked three times by each participant for predicted effectiveness at encouraging them to do physical activity, usefulness and likeability.

I had observed some differences between these ranking questions during the interview phase, and wanted to test how strongly they are associated with one another. Although the rankings are not continuous, they are ordinal, and so I used Spearman's Test to check correlation between the rankings provided for each device. All three pairwise comparisons are shown in Table 21. All three constructs correlate strongly, but not completely, with one another, and these correlations are *highly* significant.

Theme	Effectiveness	Usefulness	Likeability
Find Activities to take part in	**Yes**		**Yes**
Remind me to exercise/eat	**Yes**	**Yes**	
Motivate me	**Yes**	**Yes**	**Yes**
- by shaming me or making me guilty	**Yes**		
- by forcing or pushing me	**Yes**	**Yes**	**Yes**
- by making it fun	**Yes**	**Yes**	**Yes**
- by letting me share on social media	**Yes**	**Yes**	**Yes**
- by making it social (in person)	**Yes**	**Yes**	**Yes**
- by providing health data	**Yes**	No – Health data was mentioned as being useful, but never explicitly as *motivation*	
- by tracking my progress	**Yes**		
Providing targets/goals	**Yes**	**Yes**	
Overcoming constraints (time, people, facilities)	**Yes**		
Giving information about my health	(No – Always mentioned as motivational, above)	**Yes**	**Yes**
Encouraging healthy eating		**Yes**	
Validating the effect of my exercise		**Yes**	
Improving my health		**Yes**	
Improving my fitness		**Yes**	
Helping me be sociable		**Yes**	**Yes**
Making me more active		**Yes**	

Table 23: Themes containing at least three responses, on the effective and useful responses

The data collected for the *likeability* responses is somewhat different in character to that collected for *usefulness* and *effectiveness*; in part because the positive likeability question was framed with relation to specific features rather than the device more generally. In hindsight, this choice of wording is limiting and may be why the themes that I identified

are more limited than those for *usefulness* and *effectiveness*. The positive likeability data seems to be similar in scope to the usefulness data, although with less total themes.

The responses to the negative likeability question were again very different, and so I discuss those separately in the next section.

ii Dis-likeability

The data on dis likeability was quite device-specific. One or two themes encompassed multiple devices, notably the sense that the device had no purpose or would not be effective at improving health or fitness.

The **Vitality** was the only device that nobody mentioned disliking. This is similar to the interviews, where it was fairly well received by all participants.

There were no specific features of the **Giggle** that participants disliked, but it was the device most often referred to as not having a purpose, or being pointless.

Concerns about meeting strangers put some participants off of the **Butterfly**, and others suggested that relying on friends being available was risky; suggesting that this would be a constraint on their activity.

The **Enforcer** was most often disliked because of how demanding it is perceived to be. It was described as "cruel", "demanding", "extreme", "brutal" and "terrifying". Participants felt that it would ineffective because they would not wear it for a sustained period of time.

The **Trend** was by far the most disliked device, attracting about half of the total comments. Like the Giggle, it was often cited as being pointless. A substantial number of participants noted that they do not use social media, rendering the device useless for them. Others disliked the experience that they'd had of *other* people posting fitness information on social media, with one commenting that *"the Trend sounds like more social media induced egoism. Society has enough of this already in my opinion."* Perhaps most interesting, though, are a group of responses that suggest a values-conflict between the actions embodied by the Trend and what they perceived fitness activities to be about. Most commonly, they suggested that fitness should be a personal activity; one explained that *"if you are exercising and trying to get fitter it should be for you not because you want likes on social media"*, another said "*it focuses more on body image ... and links workout success to impressing others."* Others framed their objection in terms of privacy – *"I'm more private and like to keep my activity to myself".*

iii Values and Device Choice

I wanted to see whether specific values were predictive of particular devices being adopted; to test this, I devised fifteen specific hypotheses about values (measured by the PVQ) that might be associated with specific devices. Those hypotheses are summarised – along with results – in Table 22.

To test these hypotheses, I grouped participants into those who liked, disliked or were neutral towards each device, based on how they had ranked the devices in order of how much they liked them. Ranking a device first or second was treated as liking that device, ranking it fourth or fifth was treated as disliking it, and ranking it third was considered being neutral. This categorisation provided five variables, "LIKES_X" (one per device) with three values, "YES", "NO" and "NEUTRAL". I chose likeability for this test as, conceptually, it seems to account for a wider range of factors than either effectiveness or usefulness; the very high correlation between the three factors (above) suggests that results for the others would be similar, though. To avoid introducing Category I errors (or reducing the power of p to an unrealistically low level through correction) I did not perform hypotheses tests against effectiveness and usefulness.

I conducted a Mann-Whitney U Test to compare the PVQ scores for the hypothesised value between those who *like* and those who *do not like* each device. The results of these tests are shown in Table 22. After controlling for multiple tests by applying a Bonferroni correction to *p*, two of the tests showed significant differences in values between those who liked and disliked a particular device; specifically, I found a difference in Sociability and Benevolence between those who like and dislike the Butterfly.

4.4.5 Discussion

Several areas of the results warrant further discussion, in particular the relative roles of motivation and information within individuals' wellbeing practices, the values conflicts that the Trend seemed to elicit,

i Motivation vs Information

There is a distinction evident from participants' responses between the devices as *informational* tools, and as *motivational* tools. In the first case, the device is seen as a means of obtaining information to be acted upon; this was most commonly associated with the Vitality where health information was seen as valuable. In the latter case, the device is seen as a source of motivation to engage in an activity. Only a small minority of participants suggested that the information was in itself motivational in itself, and P5

event suggested that knowing she was in good health could actually be *de*-motivational in the sense that it would undermine the motivation that she derived from feeling that she was improving or maintaining her health.

Although most participants focussed on information about their health, a few referred to information about *how* to engage in an activity, such as the route maps provided by the Giggle or by using a device to help locate facilities like swimming pools. The accounts provided by participants suggest that information is incorporated at different levels of the activity hierarchy; sometimes to aid understanding, sometimes motivating, and sometimes dealing with conditions (like the availability of a facility).

The perceived ability of the devices to provide motivation was a key factor in many participants' decisions about which they preferred, and this was observed in both the interviews and the online survey. I suggest that this indicate there is, quite commonly, a *second-order volition* about physical exercise: people *want to want* to exercise, and see devices as one way of fulfilling that desire by making engagement in physical activities more enjoyable and appealing. I also suggest that there is a conceptual difference between the motives behind the second-order volition itself ("I'd like to be leaner"), and the motivation that can be found in answer to it ("This device motivates me to run"). This latter motivation is *instrumental* and forms an important part of the activity from a practical perspective, but it is not related to a motive that the activity ultimately answers to.

ii Semantic Conflict

Participants' reactions to the Trend demonstrated some limits on the notion of instrumental motivation in the form of semantic conflict, the sense that instrumental motives can be in conflict with the overall meaning of the Activity. This was most clearly illustrated by several survey participants who felt that incorporating social media into fitness activity was incompatible with their view that fitness should be oriented towards the individual.

> *The Trend annoyed me – to me if you are exercising and trying to get fitter it should be for you and not because you want more likes on social media.*
>
> *I don't like all the social media features on the devices, it is not what a fitness device should be about.*
>
> *Fitness should be for yourself, not publicising to others.*

There is a strong sense in these comments that, even if sharing can be motivating, it is essentially incompatible with fitness activity. There is an interesting contrast here between these participants, and some of the more established runners and cyclists in the previous study who felt that sharing information about their activity – even cultivating an identity as a runner – was (at least in part) a benevolent act, aimed at encouraging others and indirectly improving the wellbeing of those around them.

This conflict warrants further investigation, and is undoubtedly a relevant design consideration for persuasive or instrumentally motivating technology more generally. It is reminiscent of a concept proposed by Selinger & Whyte *semantic variance,* which they advance in their critique of Nudge techniques by arguing that "*nudge designers really understand how different people re-interpret the meaning of situations after a nudge has been introduced into the situations"* [162]. Despite being (in part) instrumental (much like a nudge would be), the idea of sharing of fitness activity led these participants to re-interpret the meaning of the overall activity in a way that they felt was incompatible with the original motives.

iii Affordances

As an observation, many participants focus on the affordances of specific features, rather than the motives that each tool seems to embody. Those motives had been made explicit in the descriptions, though, for instance: "*The Giggle just wants you to have fun*"; but aspects like route planning, and having access to a camera, were still valued.

Some participants also seemed to imagine the use of these devices in activities oriented towards non-fitness Objects. For instance, several participants noted that the Butterfly would allow them to be more sociable, apparently a valuable Object in itself and not merely an instrumentally motivating factor.

It's worth bearing in mind that even devices designed around specific values – and which are therefore presented quite teleologically – still afford a range of possibilities to people.

iv Value Influence on Technology Choice

The quantitative results from the online survey suggest a weak link between people's personal values and the technology that appeals to them. The results show a significant link between Benevolence (a socially-oriented value) and the Butterfly device. This partially confirms the hypothesis that values would influence device preference, based on the established role of values in preference formation (e.g. [153,160]), but the considerably

number of hypotheses that were not confirmed points to the complexity of this relationship.

The fact that many participants (as noted above) seem to consider the affordance of the features of each device, as well as the more teleological explanations provided about their purpose provides some explanation for this; one participant with an interest in photography chose the Trend for the built-in camera, rather than the social sharing narrative that it was presented with, for instance. It is also partly explained, perhaps, by the distinction between instrumental motivation and the ultimate motives behind a fitness activity. Participants demonstrated some awareness of what they felt would motivate them to engage in activity – such as finding enjoyable activities, or working towards goals – that seem to be distinct from their more deliberative conscious motives such as weight loss or body image.

4.4.6 Conclusions

The Speculative Design study was instructive on a number of fronts. It established a link between values and device preference – albeit incomplete and apparently complex. It strongly suggests a role for *instrumental motivation* that motivates instantiation of an activity, in response to a second-order volition, but which is not necessarily grounded in the motives that the Activity ultimately oriented towards. It also demonstrates the need to ensure that this motivation does not alter the overall meaning of the Activity in a way that people would find objectionable, which is reminiscent of of Selinger & Whyte's *Semantic Variance* [162].

These findings are interesting to consider in relation to the results from the previous wellbeing studies, as I have touched on already in the discussion, above. The next chapter aims to draw these three studies together more deliberately and to discuss wellbeing practices, and technology, from an Activity Theory perspective.

4.5 An Activity Theory View on Health Behaviour

In this chapter, I reflect on the results of the previous studies and discuss some of thine findings through an Activity Theory lens. Further discussion of the issues identified below, and others, is provided in section 5, Unification.

I being with a discussion of instrumental motivation – the ways in which participants seek motivation to realise second-order volitions – then go on to discuss how the Object of activity can be made more resilient through deliberate polymotivation, and finally discuss how bounded rationality related to self-control and motivation could, conceptually, be incorporated into Activity Theory.

4.5.1 Instrumental Motivation

What's striking in the results of the speculative design exercise is the extent to which choosing a device is seen by participants as a way to obtain *motivation*. The focus on obtaining motivation seems to be indicative of a second-order volition, the participants *want to want* to engage in physical activity. On the face of it, second-order volition is a defining feature of many health-related behaviours like smoking cessation, healthy eating, or exercise.

There is a lack of consistency not only between what medical experts tell us is good for us in an objective sense and our actual behaviour, but also between behaviour that would help to achieve particular aspirations (like being physically fitter) and what we actually do. There is an important temporal aspect to this, which reflects Thaler & Sunstein's [167] notion of "the planner" and "the doer". Psychological effects like hyperbolic discounting [105] mean that what we want to do in the future is usually different to what we want to do right now, in the moment. In the future, we typically care about very conscious concepts like our health, wellbeing, wealth or achievement – things we might call values. In the moment, we are guided by how tired we feel, or how hot, cold, hungry, lonely, excited, or sad; things that we might think of, broadly, as emotions. Crucially, the latter are experienced rather than being deliberately thought about. They seem to surface from somewhere unconscious, rather than being the result of conscious thought. Nonetheless, one's very conscious plan to eat salad can be derailed by a bad day at work when, it transpires, ice cream is more emotionally appealing (even though one is under no delusion about its compatibility with the aspiration). Essentially, there is often a really big

difference between what we want now, what we want in the future, and (consequently) what we *want to want* right now.

In the speculative design interviews, several participants talked about how the devices could help them to be more active. This took several forms: By reducing barriers (such as finding a convenient time), by constraining their choices (in the case of the enforcer), or by combining physical activity with other motivations such as socialising. As P5 said:

> *The main thing I would want out of one of these devices is the motivation and encouragement to actually do [exercise] – SD5*

During the discussion of the Speculative Design exercise, I referred to this additional motivation as *instrumental motivation*; in the sense that it is subordinate to the motives that are encapsulated within the Object of activity itself. It was hugely significant to participants in the speculative design exercise, and also to the Fitbit users who took part in the comparative interviews.

4.5.2 Object Hardening

However, instrumental motivation is not the only motivation-building process that seems to be at play; F1's account of how her practice had changed over time shows how she had combined multiple motives into a single polymotivated object – by walking to the train station, she was able to fit steps into her routine while also accomplishing part of her daily commute. The combination of these two motives allowed her health and fitness motive to be incorporated into an activity that was routinely instantiated; the polymotivated Object was more resilient than one motivated only by health and fitness. Indeed, when her circumstances changed she no longer needed to commute, and she no longer found time to fit the steps into her routine. F5 also used walking as a means of getting from home to work, but built on that to deliberately find longer routes that increased her step count. In her case, the two motives (getting to work, and health) both shaped the construction of the Object, but the deliberate incorporation of additional steps seems to reflect a greater role of health and fitness in motivating the activity than in F1's case.

Polymotivation was hugely important to the runners and cyclists who I interviewed; they'd all, over time, found that their practice tied into a wide range of motives and were able to articulate that quite clearly. R1 felt that the motives he had discovered since starting to run were now so important that he wasn't sure if he would continue to run

without them. As a result, the runners and cyclists sought out ensembles that would maximise these non-health motives; such as R1 who found the social aspect of races to be important, and C10 who used races in particular locations to explore those locations and explore their historical significance.

The polymotivated Objects of activity that underlie these activities lead to a more robust activity *because* they are polymotivated. By finding activities that are enjoyable, or sociable, or practical, *as well as* fitness-improving, participants demonstrate the construction of an Object of activity that will motivate them at a point later in time, when the focus is no longer on planning an activity, but on actually doing what one has planned to do.

I'll refer to this process of identifying a conscious aspiration, and combining it with other more immediate motivations, as "*Object hardening*"; a process (deliberate or serendipitous) of incorporating multiple motives into an Object of activity so that the associated activity becomes more resilient to fluctuations in how motivated we might feel to pursue a particular motive in the moment.

These two approaches to motivating an activity – through *instrumental motivation* or *object hardening* – both seem promising from a design perspective; in fact, the latter seems to encapsulate many aspects of persuasive technology, where motivation is built through features that are motivating but often not particularly meaningful in their own sense. In fact, they appear to be complementary: the fact that the more established runners and cyclists had all discovered ways in which their practice could contribute to motives beyond health and fitness suggests that a motive which was originally incorporated as an instrumental motivation – like running with friends – could be incorporated into the Object as it began to provide more meaning to the activity and became an aspect of the activity that the participant seeks out, rather than being important only to the extent that it serves another motive like fitness.

4.5.3 Bounded-Rationality in Object Instantiation

The notions of *object hardening* and *instrumental motivation* have some additional implications for Activity Theory itself. Recall Kaptelinin's 2005 n:1:1 model of polymotivation; in the process of accounting for why object hardening works at all, that model can be revised to account for another important aspect of activity, one which is pervasive in the fitness domain: whether or not an activity takes place *at all.*

As explained earlier, the notion of polymotivation is not new to Activity Theory, and Kaptelinin proposes a model of polymotivation in which a single object combines

multiple needs and motives [93]. However, while the model proposed (n:1:1) allows for the relative importance of the different motives to alter how an activity is shaped at any particular time, it does not in itself directly account for whether or not a particular activity even takes place. In the case of physical practices such as walking, running and cycling, this *non*-instantiation was often mentioned by participants as a problem to be overcome. Despite participants' intentions and aspirations, and knowledge that (for instance) going for a run could be beneficial to them, finding the motivation in the moment to get on a bike, or go for a run, is a challenge, and a factor that needs to be modelled through theory in order to provide a full account of how these activities play out.

Figure 15 shows Kaptelinin's n:1:1 model incorporated into a larger overall model that allows for multiple polymotivated objects to be held simultaneously. The new model reintroduces the notion of a hierarchy of motives as a means for choosing which object (and hence activity) is instantiated at any particular time. I refer to it as the n:n:n:1 model: multiple needs, multiple motives, multiple *constructed* objects but a single *instantiated* activity.

In contrast to the n:1:1 model, the 1:1:1 model that Kaptelinin also discusses, based on Leontiev's earlier writing, uses a *hierarchy of motives* (HoM) to explain how – out of the multiple objects and associated activities that a given person may hold – one is actually instantiated at any given time. An ability to model multiple activities (as well as a single polymotivated activity) is a useful property of the theory to retain in itself. Nonetheless, an n:n:n:1 model maintains other features of the n:1:1 model, such as multiple motives being incorporated into a single object simultaneously.

Of course, reintroducing the hierarchy of motives as a theoretical device also reintroduces Kaptelinin's critique of it; notably the claim that the HoM does not provide much predictive power, and does not adequately deal with factors such as task *urgency* which may alter priorities in the moment. In contrast, the HoM that I propose here is conceptually highly dynamic, representing the relative importance of different motives at a point in time. It is, notionally, the place where a stressful day at work causes the instantiation of *eating for pleasure* instead of *eating for weight loss.*

A lack of predictive power within the HoM – whether in the generalised Leontiev'ian sense, or the more momentary construct proposed here – is a criticism that bears further scrutiny. Given the well researched difficulties that people have in introspecting their own decisions, and the predictable irrationality with which they behave, the observational and ethnographic research methods often associated with Activity Theory do seem to be poorly suited to better understanding the influence of factors like self-

control, intrinsic motivation and bounded rationality on how Objects translate into actual activity. Rather than a limitation of Activity Theory, though, the HoM may be seen as an *opportunity* to begin bridging the gap between two sides of human psychology; the rational, deliberative and slow System 2, with the fast, emotional and momentary System 1. In essence, a highly dynamic and momentary hierarchy of motives is not a black box for factors that cannot be explained, but a suggestion for where Activity Theory should consider *integrating* the theory and predictive models of Behavioural Economics that deal with the less purposeful side of our behaviour that AT typically excludes.

In providing a conceptual space for these models, the HoM allows AT to account for the observable incongruence between the *aspirations* people have about their future outcomes and the behaviour that they know or believe is necessary to reach them, and the behaviour that *actually* transpires. Furthermore, it explains how *object hardening (*altering the Object of activity to incorporate additional motives that consistently or regularly feature more highly in the HoM) and *instrumental motivation* (motivation derived from factors that are not motives in themselves) increase the likelihood that an activity is actually undertaken.

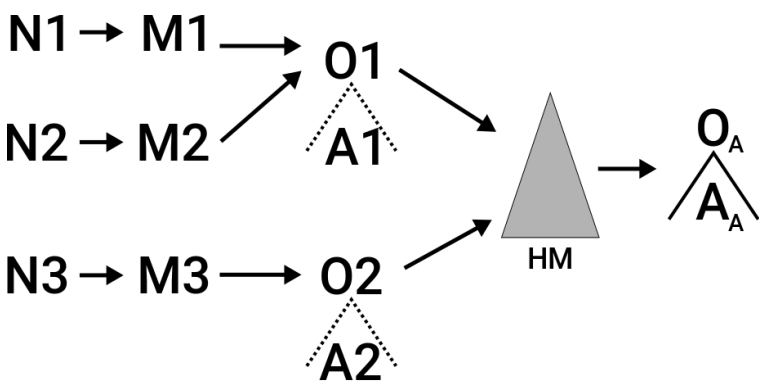

Figure 15: An N:N:N:1 model of needs, motives, objects and activities

Given the importance of System 1 factors in domains such as behaviour change, which are purposefully initiated and directed towards conscious needs and motives but often complicated by people's relationship to factors like their own motivation, this is a conceptually important bridge to build.

These factors – object hardening, instrumental motivation, and a means of explaining them through the incorporation of System 1 factors in Activity Theory – provide in themselves insights that may help to apply to Activity Theory to domains, like health behaviour change, where bounded rationality and fast decision making are implicated within more purposeful activity. In the final section of this document, I reflect on how findings from privacy can be combined with findings from health behaviour change in order to situate values more generally within Activity Theory, and to identify design implications for Value Sensitive Design.

5. Unification

The two domains explored previously, privacy and wellbeing, each reveal different aspects of the role that values play in people's activities; and I have already considered their implications for Activity Theory in isolation. Building on the previous discussion, in this chapter I propose a more unified account of how values can be considered through Activity Theory, and discuss design opportunities and challenges that the findings suggest.

5.1.1 Needs and Values

Rokeach's early work on values draws a distinction between *needs* and *values*. This distinction is not made in Activity Theory, but I will argue that introducing such a distinction is necessary in order to maintain AT's current definition of needs and to adequately account for conceptual differences between the two.

Rokeach's original definition of a value is *"an enduring belief that a specific mode of conduct or end-state of existence is personally or socially preferable to an opposite or converse mode of conduct or end-state of existence" [153].* Two aspects of what constitutes a value are particularly germane: First, that a value is a belief, and second, that values are the basis for preferences.

In contrast, Kaptelinin and Nardi describe needs as *"a directedness of activities toward the world, toward bringing about a desirable change in the environment" [98]* (p60). In this sense, that needs are formed based on desires—preferences—there is a clear role for values, as arbiters of what is preferable, in the formation of some (if not all) needs. This in turn suggests that a need is a function of both a value and the environment as experienced by a person. Accordingly, we can conclude that values give rise to needs when current circumstances – either states, or activities – do not adequately meet the value's conception of what is preferable.

A value such as "wealth" provides a good example of this distinction in practice. A person who values wealth but has little money or few assets might experience a need for more; while a person who values wealth but is very well off would not do so. Conversely, there are communities and cultures that reject wealth as a value and instead seek values such as spirituality. Placed in the circumstances above, those people might experience a need in the latter – high-wealth – circumstances, but not necessarily in the low-wealth ones.

That is not to say that all needs are formed in response to values; Rokeach suggests that although animals experience needs, we would not consider them to possess values in the same way that human beings do. Some physiological experiences such as hunger are intuitively experienced as a need, even in the absence of a particular *belief* that hunger is a negative state.

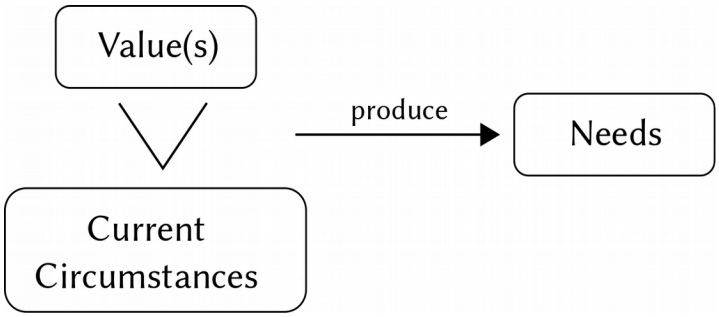

Figure 16: Values are applied to current circumstances and produce needs.

Drawing a distinction between needs and values places values within the activity hierarchy at a level above needs. Values give rise – when used to judge current circumstances – to needs, which in turn become motives from which an Object can be constructed, as shown in Figure 16. This framing is compelling for two reasons. First, it resolves conceptual differences between values, which are enduring, and needs, which may be transient; and clarifies that a value can be held without necessarily giving rise to a need. Second, it provides a conceptual space for values to be situated, explicitly, within Activity Theory.

5.1.2 Values and the Internal Plane of Action

Values, held as beliefs, are independent of immediate circumstances. It is when a value is applied, in mind, to current circumstances that it gives rise to a need. However, values can also be applied to possible scenarios, and possible actions (modes of conduct). Unlike actions on external objects, which can have unpredictable effects as a result of incomplete models of how the objects relate to and interact with one another, the relationship between a perception of objects and resulting value-based judgements is always entirely in mind; the act of considering current circumstances vis-a-vis a particular value is no different, conceptually, than considering an accurate prediction of possible circumstances vis-a-vis the same value.

In Activity Theory, the process of predicting how actions will affect the world "in mind" takes place in a construct called the Internal Plane of Action (IPA). Kaptelinin explains that "t*he IPA is a concept developed within activity theory that refers to the human ability to perform manipulations with an internal representation of external objects before starting actions with these objects in reality,"* [91].

Given that values exist only in mind, and do not directly correspond to real-world objects, they could be said to exist *only* in the IPA, and it therefore follows that values have a role as restrictions on possible actions beyond those imposed by objects as they exist in reality. Values become additional conceptual objects, in mind, that through their interaction with our understanding of external objects (as currently understood, or predicted) create needs and motives that can cause us to seek change, or to avoid particular courses of action.

A model of values as objects in the IPA adequately describes how privacy, which I previously described as a restrictive constraint, influences activity. Although abstract, and not directly present in the world, privacy, when applied to judge the preferability of current circumstances can give rise to a feeling of need; and when applied to possible actions and their outcomes by way of a motive such as "maintain privacy" restricts the scope of possible actions beyond those that would answer to the needs of another directing motive, such as "locate health information"; in doing so, the value can result in a material change to the way an activity or action is carried out, for instance motivating the use of private browsing mode in a Web browser. The crucial factor in being able to apply the value of privacy to inform action is an understanding of how the objects involved in that activity system – such as the records kept by ad brokers – will be transformed by particular actions, and hence to have an accurate prediction of possible future circumstances that can be considered vis-a-vis privacy.

However, as I observed in the cookie notice focus groups, there are two conceptual stages in the successful application of a value to a set of circumstances. First, there must be an appreciation that a particular value is somehow *relevant* to the situation. This could be the result of learning that a link exists in advance, or because a link is articulated somehow, such as through a consent interaction. The relevance of a particular value must be *apparent.* Second, the subject must be able to predict how their actions will alter the objects in reality and apply the identified value to consider the preferability of the resulting circumstances; the impact that the action will have vis-a-vis the value must be *appreciable.* As the Web Mirror showed, this appreciability could be as simple as making the link that "browsing the Web is bad for privacy" (introducing a need in the predicted

circumstances and an additional motive to the activity); and hence motivate the installation of an ad blocker.

5.1.3 Values in Context

In the model proposed above, values are superordinate to needs; occupying a conceptually higher and more abstract level in the activity hierarchy. The possibility of these 'global' values influencing activities is confirmed empirically in the speculative design exercise as the observed links between (some of) the participants' values as measured by the PVQ (values) instrument and their device preference. However, as Barr *et al.* note in their work on Fable [20], the definition of a value as an enduring belief about the preferability of a mode of action or end state can equally apply to more contextual beliefs, such as the notion that it is good to make progress in a video game, or that *"it is always good to swing my sword at bandits"*. In the privacy domain, the idea of contextual values are implicit in the notion of context collapse [54,123]

We can see evidence of contextual values in the primary research; most notably among survey participants in the speculative design exercise who felt that exercise should be a personal activity, rather than one which is shared on social media. Although for some of those participants, a general aversion to social media use in general was obvious, for others this appears to be a *contextual* value; not a general aversion to social media, or to exercise, but a result of a belief about what exercise *should* be.

Likewise, many of the participants who used a Fitbit placed high value on the ten thousand step target. Meeting that target becomes a value in the context of the activity - "it is good to take ten thousand steps" – that can excite and direct the activity.

Importantly, while these values (like meeting a ten thousand step target) provide motivation and direction, they are still, ultimately, constraints on the overall activity. Logically, the motives that they translate into specialise the Object of activity further and reduce the horizon of possible Objects.

Contextual values seem to be an important source of instrumental motivation, in the sense that they can be used as vehicles for the intrinsically motivating experiences such as progression and achievement [45] that exergames and other persuasive techniques use to motivate participation in particular activities. In this role, contextual values become – as Barr *et al.* observe [20] – conceptual links between conditions and operations; beliefs about how to act within the context of the activity as it is is constructed and mediated by persuasive tools such as exergames.

5.2 Design Implications

As well as progressing the general understanding of how values fit within activity systems, these results suggest some concrete design implications. The observation that tools can create values, or help people to relate activities to values, suggests a deliberate focus on doing so; and there are also opportunities relating to the process of Object construction and re-construction following breakdowns.

5.2.1 Creating Contextual Values

The view that tools mediate subjects' interactions with objects is central to activity theory, and as I have already touched on several times above, tools can mediate the relationship between actions, objects and values. However, as the Fitbit illustrates, tools are also able to *create* values in the context of an activity. Typically, this seems to happen through the ways that tools conveys notions of performance and progress; either in the form of a target (like ten thousand steps), through notions such as 'higher is better' (especially for participants who felt that the ten thousand step target was too hard or too easy), or by tying measures to specific motives (such as the 'fat burn' zone suggested by the Fitbit's heart rate monitor).

These contextual values can be vehicles for creating intrinsic motivation, as they might in an exergame, or simplifying concepts in the same way that Privacy seems to simplify reasoning about complex and uncertain consequences of sharing personal data. In introducing contextual values, though, we need to consider how they could interact with other beliefs about that activity in order to avoid conflicts and resulting breakdowns; and consider the possibility that values which are created by tools could occlude the original values and needs to which the activity answers.

In order to avoid conflicts it is necessary to understand both the motives that direct the activity, and any existing contextual values. For instance, the fact that some people believe fitness activities should be inherently personal limits the incorporation of socially-oriented contextual values, such as social sharing or 'role modelling'.

Fritz *et al. [70]* found that some users were reluctant to walk without the Fitbit because their steps felt "wasted"; the ten thousand step goal become such an important value that it altered users' experience of walking, and effectively *occluded* the original motives that they had for walking. Understanding how occlusion occurs, and how to design to mitigate it would be an interesting area for further research. It's possible that building appreciability between contextual values – like a step count – and overall motives could

help; as could encouraging the internalisation of more qualitative measures that can be assessed without the aid of tools.

5.2.2 Relating to Values

As well as creating new contextual values tools can also help subjects to *relate* actions or circumstances to a value; either in the context of an activity (like counting towards that ten thousand step target) or to global values (as in the case of tools like the Web Mirror that help to explicate the link between web browsing, tracking, and privacy). In this sense, tools can assist in the reflective process of understanding current needs and learning how activity alters outcomes that are desirable (or not) vis-a-vis a particular value.

Tools as mediators that can *create* values and that can help subjects to *relate* actions and outcomes to values present design opportunities, in terms of facilitating those processes, that are obvious candidates for consideration in Value Sensitive Design exercises.

Conversely, by considering how a wider range of values relate to activities such as walking or running, we can identify new opportunities for designing devices that appeal to a wider audience of users. For instance, by designing tools that are able to make a link between a value like family and activity like exercise *apparent*, and *appreciable*, there may be opportunities to contribute to broader adoption of practices with pro-social outcomes such as reduced healthcare costs. The accounts provided by runners and cyclists suggest that appreciability needn't be operationalised through the tool itself; R1 (for instance) was able to relate his running practice to the quality of his interpersonal relationships. It may be enough to prompt this reflection in order to make relevant values salient to the subject.

In the case of consent, making value tensions apparent to users is important in inducing the breakdown that, notionally, must occur in the event that user is unable to consent to the effective activity system that is proposed by a service provider. Values, as determinants of what is preferable, are conceptually key in making such a determination; and so, from a design perspective, identifying the values most relevant to the proposal for which consent is sought should guide attempts to deliver meaningful consent interactions. This is in stark contrast to the consent mechanisms that were observed during the notice survey, which were grounded in explanations of how the activity system is operationalised through technical mechanisms like cookies, rather than the motives or value tensions that are implicated.

5.2.3 Intervening in Object Construction

We have seen how tools like the Fitbit can influence object construction by suggesting particular activities, such as walking, in answer to health motives; and how, in doing so, they can create values within those activities. We've also seen how important the "in the moment" motivation that these tools can provide – through the values they create, and the ways that they frame progress – can be crucial in constructing activities in response to second-order volitions about physical activity.

Kaptelinin compares the process of Object construction to design [93], a proposal that in itself suggests the possibility of computer-mediated Object construction; and a combination of Activity Theory and Value Sensitive Design suggests deliberate intervention in this process with respect to values. One opportunity in this respect would be to facilitate Object hardening (already discussed at length), as the deliberate incorporation of multiple motives into an Object to improve the likelihood that it is instantiated in the moment.

In the privacy domain, we have observed how consent mechanisms can cause a focus shift, and possibly trigger a breakdown in the activity. In this way, the goal of consent mechanisms is to test, and possibly to reformulate, the Object that a user holds. Exposing possible contradictions to give the user an opportunity to identify them ahead of time, and to react accordingly. We also observed that many current consent mechanisms fail to do this. Cookie notices do not facilitate the understanding that's necessary for the privacy contradiction to become apparent to people, and – even if they did – they do not offer much in the way of support to the user when it comes identifying ways that their Object could be reformulated. By understanding the ways that values can constrain activity – including, but not just, privacy – we could perhaps design consent tools that aim for *recovery* as well as breakdown; perhaps in a similar way to the anti-tracking countermeasures that we proposed to participants in the Web Mirror studies.

5.3 Conclusions

Values can be situated, explicitly, within Activity Theory in two distinct ways: First, as general arbiters of preferability and hence progenitors of needs. In this sense, beliefs about what is preferable – values – when applied to the circumstances in which one finds oneself can give rise to feelings of need; for instance desiring greater privacy, or to improve relationships with loved ones. Second, as beliefs about how more specific activities should be conducted. For instance, the belief that activities oriented towards fitness should be inherently personal, and so not shared on social media.

In both cases, values can act as constraints on activity even in the absence of a needs-based motive. This is frequently the case with privacy, where complete inaction is not experienced as an unmet need for privacy, but where other Objects of activity might be constructed in ways that maintain privacy instead of in ways that do not. This is consistent with Kaptelinin's observation that Objects can be constructed in response to various types of constraints [93] and reflects a difference in kind between needs and motives (which excite), the object (which excites and directs) and values (which, as constraints, direct).

Values play a role that is tightly connected with the Internal Plane of Action; suggesting that they should be situated, conceptually, within or in close proximity to, the IPA. As conceptions of what is preferable, values can be applied to one's understanding of current circumstances but are also frequently applied to the imagined or predicted outcomes and value-systems that result from manipulations within the IPA. This implies that the IPA must contain a rich enough representation of the objects that are to be transformed, and how those objects relate to specific values, to facilitate accurate predictions about the desirability of particular courses of action. In the case of privacy, this is often not the case, making appreciation of how (for instance) accessing a website will impact upon one's privacy.

6. Conclusions

In the work presented here, I have considered two complementary domains – privacy and wellbeing – as loci in which people's everyday practices, like browsing the Web or commuting to work, interact with their beliefs about what is preferable, and what is not – their values.

This has led to three distinct areas of contribution:

1. A description of how third-party tracking is constructed as a multi-stakeholder activity system, with implications for privacy and consent.
2. A description of how physical activities like running, cycling and walking are situated with polymotivated activities, with particular emphases on the relationship between deliberate, purposeful action and the effects of bounded rationality and limited self-control that pervade efforts at health behaviour change.
3. A generalised account of how values are situated within Activity Theory, that can support Value Sensitive Design efforts at a 'context of use' level.

I will briefly summarise my contributions each of these areas in turn, before suggesting avenues for future work to consider in greater detail. However, I begin my conclusions by reflecting on my experiences of using Activity Theory as a framework within which to situate values.

By adopting Activity Theory as a framework, I have been able to consider how the actions that make up these practices relate to underlying needs and motives, and to suggest how values – as a distinct concept – fit within the broader theoretical model that AT provides. Considering how values relate to activity as progenitors of needs, or as constraints, enables us to uncover opportunities to design in response to the values that are implicated in specific contexts of use; and to identify opportunities to deliberately facilitate an alignment not just with the actions that need to be carried out, but with the values that direct and excite the activity.

In the speculative design exercise, by considering how different values create different motives within ostensibly similar practices of "walking" I was able to design concepts that presented instrumental affordances with respect to a range of different values. Whereas Epstein *et al.*'s application of VSD to the sharing of step data [59] allowed values such as

support, accountability, honesty and accuracy to be identified, I found that considering the overall activity system around fitness activities uncovered values such as enjoyment, environmentalism or being a role-model, which present exciting opportunities for design in themselves.

6.1.1 Conclusions on Privacy

By considering privacy as a value, and situating it within activity, its role as a constraint on activity can be appreciated. Moreover, situating privacy within Activity Theory allows 'consent' – a value in itself – and consent mechanisms (concrete interactions that aim to deliver consent) to be examined as focus shifts, and potential breakdowns, in multi-stakeholder activities.

Third-party tracking, as an activity system, embeds a number of different of actors each with different motives and Objects of activity. The end result is often to subordinate individuals' own Objects to an emergent one that fails to account for values like privacy. When modelled through Activity Theory, these conflicts in the emergent Object as it is constructed by a first party website become readily apparent.

Moreover, the ineffectiveness of current "consent" mechanisms, which provide very limited instrumental or learning affordances is also demonstrated, leading to the observation that, in order to be effective as regulation intends – and to deliver the alignment between user's own wishes and those of the other parties involved – these mechanisms need to facilitate a focus shift, and quite possibly a deliberate breakdown in the user's activity as a result of the value tensions that arise. Framing consent in this way suggests that it may be fruitful to consider ways to help users recover from the breakdowns in their activity that arise when they cannot legitimately consent to what is proposed by a service provider. The relative importance of offering alternatives over increased information is shown through the results of the Web Mirror study. Simply, finding that a value tension can be resolved through the adoption of a simple browser tool may obviate the need to build a more detailed understanding of precisely how that value tension is operationalised.

6.1.2 Conclusions on Wellbeing

Applying Activity Theory to the health behaviour change domain provides opportunities to consider how two different bodies of theory – Activity Theory's model of purposeful human action, and behavioural economics' accounts of the two systems of human decision making – can be reconciled and used productively in relation to one another, for

instance by framing effects such as "in the moment" motivation as factors that affect Object instantiation. Moreover, by doing so in a value-sensitive way we can begin to unpick the complexities of instrumentally motivating, or persuasive, techniques within an already purposeful activity, and to consider how they can change the overall meaning of the activity by introducing new values and possibly occluding others.

Despite the importance of behavioural theory in behaviour change, and the dominance of techniques like persuasion in contemporary research, we should not lose sight of the fact that, ultimately, people's decisions to try and change their own behaviour are often thoughtful and purposeful, grounded in very diverse values, and the value derived from participation in these activities is often felt most keenly in terms of factors *besides* health and fitness outcomes. There is, as I have begun to explore, promise in understanding more about how theories relating to persuasion, behaviour change and bounded rationality more generally can be situated within Activity Theory, and the focus on purposeful tool-mediated interaction with the world. Intuitively, both sides are important in building activities that are not only successful in engaging people in actions that improve objective health outcomes, but which answer to their own values and motives in ways that are meaningful to them.

6.1.3 Conclusions on the Role of Values in Activity Theory

Values can be situated within Activity Theory at two levels: as global conceptions of what is preferable, and as more contextual beliefs about how particular activities should be undertaken.

In the case of global conceptions of what is preferable, values can be progenitors of needs; experienced as a need when current circumstances are judged as inadequate with respect to a particular value.

Values – global and contextual – can also apply directly to activities as constraints; not necessarily being experienced as a need state, but suggesting constraints on how Objects of activity are constructed and how goals within that activity are realised through actions.

Arising from these findings are design implications: first, that the way tools mediate between activities and values should be considered directly; how do we identify and communicate the aspects and outcomes of an activity that are relevant to the values that the user has? Second, that deliberately introducing contextual values can help users to construct Objects and to introduce instrumental motivation or simplifying concepts into what could otherwise be undesirable or complicated activities. But, that in doing so we

must be cognizant of the risks of changing the meaning of the overall activity, introducing tensions with values that are already held, and occluding the original motives in ways that become restrictive or reductive.

Value Sensitive Design, applied at this level – of context – can help to explore and make explicit the ethical challenges posed by designers who – working on technology that is framed as behaviour change or otherwise – are engaged in altering human behaviour in ways that may, or may not, be desirable vis-a-vis the values of the subjects. VSD can provide a framework within which aspects of power, asymmetry and conflicting motives become foci for design attention situated within the broader value-laden contexts of use in which they are deployed and adopted. And, through a focus on making values apparent and appreciable, could help help to 'unfold' [87] the values that are incorporated into both artefacts and the broader multi-stakeholder activity systems that are the result of our design processes.

6.2 Future Work

There are many areas left to explore; both with respect to how Activity Theory relates to the two domains that I have considered, and in terms of how it can be coupled with Value Sensitive Design more generally. While my focus here has been to ask how values fit within activity systems, there work to be done in considering specifically how that can be re-integrated into VSD, perhaps through tools and techniques such as envisioning cards [66,170].

The design implications that I have suggested themselves warrant further investigation. In what circumstances is it safe to create new values within an activity? And when do they occlude the original motives?

For consent – as an interaction design challenge and policy area – Activity Theory's notion of breakdown and focus shift provide a theoretical account of the design space; but translating that into new approaches still needs to be explored. One possibility that I have mentioned is supporting the re-construction of the Object of activity following a consent-induced breakdown, and this could potentially help to overcome some of the limitations with current consent mechanisms, such as the lack of meaningful choice (and hence voluntariness) that users often experience.

In moving the focus of design to Objects (of activity) rather than artefacts, in a value-sensitive way, there are opportunities to explore – as I have begun to do – the re-framing of pro-social behaviours like exercise or healthy eating through a more diverse range of values.

By building on Activity Theory, Value Sensitive Design, and the findings that I have presented here, there is the potential for design to situate designed artefacts within the value-motivated and value-constrained activity systems that make up our lives; to avoid some of the conflicts and contradictions that are currently built into pervasive systems like online tracking, and to build more resilient practices in support of our wellbeing.

7. Bibliography

1. 100th United States Congress. 1988. *Video Privacy Protection Act*. United States of America. Retrieved from http://www.gpo.gov/fdsys/pkg/STATUTE-102/pdf/STATUTE-102-Pg3195.pdf

2. 104th United States Congress. 1996. *Health Insurance Portability and Accountability Act*. United States. Retrieved from http://en.wikipedia.org/wiki/Health_Insurance_Portability_and_Accountability_Act

3. Abine Inc. 2012. Targetted Advertising Cookie Opt-Out (TACO). *Firefox Add-Ons*. Retrieved from https://addons.mozilla.org/en-US/firefox/addon/targeted-advertising-cookie-op/

4. Alessandro Acquisti, L. Brandimarte, and G. Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221: 509–514. https://doi.org/10.1126/science.aaa1465

5. Alfssandro Acquisti and Jens Grossklags. 2005. Privacy and rationality in individual decision making. *IEEE Security and Privacy* 3, 1: 26–33. https://doi.org/10.1109/MSP.2005.22

6. I Ajzen. 1991. The theory of planned behavior. *Organizational behavior and human decision ….* Retrieved September 8, 2014 from http://www.sciencedirect.com/science/article/pii/074959789190020T

7. Hunt Allcott and Matthew Gentzkow. 2017. Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives* 31, 2: 211–236. https://doi.org/10.1257/jep.31.2.211

8. Hunt Allcott and Cass R. Sunstein. 2015. Regulating Internalities. *Journal of Policy Analysis and Management* 34, 3: 698–705. https://doi.org/10.1002/pam.21843

9. Manal Almalki, Kathleen Gray, and Fernando Martin-Sanchez. 2016. Activity Theory as a Theoretical Framework for Health Self-Quantification: A Systematic Review of Empirical Studies. *Journal of Medical Internet Research* 18, 5: e131. https://doi.org/10.2196/jmir.5000

10. Carlos Alós-Ferrer and Fritz Strack. 2014. From dual processes to multiple selves: Implications for economic behavior. *Journal of Economic Psychology* 41: 1–11. https://doi.org/10.1016/j.joep.2013.12.005

11. Robert St. Amant. 1999. Planning and user interface affordances. *IUI '99 Proceedings of the 4th international conference on Intelligent user interfaces*: 135–142. https://doi.org/10.1145/291080.291103

12. Dan Ariely. 2008. *Predictably Irrational: The Hidden Forces That Shape Our Decisions*. HarperCollins. https://doi.org/10.3389/neuro.01.013.2010

13. Richard Arneson. 2006. Desire Formation and Human Good. *Royal Institute of Philosophy Supplement* 59: 9–32. https://doi.org/10.1017/S1358246106059029

14. Richard J. Arneson. 1999. Human Flourishing Versus Desire Satisfaction. *Social Philosophy and Policy* 16, 1: 113. https://doi.org/10.1017/S0265052500002272

15. James Auger. 2013. Speculative design: crafting the speculation. *Digital Creativity*, May: 37–41. https://doi.org/10.1080/14626268.2013.767276

16. Mika Ayenson, Dietrich James Wambach, Ashkan Soltani, Nathan Good, and Chris Jay Hoofnagle. 2011. Flash Cookies and Privacy II: Now with HTML5 and ETag Respawning. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.1898390

17. Young Min Baek. 2014. Solving the privacy paradox: A counter-argument experimental approach. *Computers in Human Behavior* 38: 33–42. https://doi.org/10.1016/j.chb.2014.05.006

18. Young Min Baek. 2014. Solving the privacy paradox: A counter-argument experimental approach. *Computers in Human Behavior* 38: 33–42. https://doi.org/10.1016/j.chb.2014.05.006

19. Balbir Barn, Ravinder Barn, and Franco Raimondi. 2015. On the Role of Value Sensitive Concerns in Software Engineering Practice. In *Proceedings - International Conference on Software Engineering*, 497–500. https://doi.org/10.1109/ICSE.2015.182

20. Pippin Barr, James Noble, Robert Biddle, and Rilla Khaled. 2006. From Pushing Buttons to Play and Progress: Value and Interaction in Fable. *Conferences in Research and Practice in Information Technology Series* 50: 51–58.

21. P Barthelmess and K M Anderson. 2002. A View of Software Development Environments Based on Activity Theory. *Computer Supported Cooperative Work: The Journal of Collaborative Computing* 11, 1/2: 13–37. https://doi.org/Article

22. Howard Beales. 2010. *The value of Behavioral Targeting*. Retrieved November 26, 2012 from http://corp1.ewr1.turn.com/sites/default/files/wp-content/uploads/2010/06/Beales_NAI_Study.pdf

23. Philip Bell. 2001. Content analysis of visual images. In *Handbook of visual analysis*. 10–34. https://doi.org/10.4135/9780857020062

24. Daniel Berdichevsky and Erik Neuenschwander. 1999. Toward an ethics of persuasive technology. *Communications of the ACM* 42, 5: 51–58. https://doi.org/10.1145/301353.301410

25. Paul Bernal. 2014. *Internet Privacy Rights*. Cambridge University Press, Cambridge.

26. Andreas Birkbak. 2013. From networked publics to issue publics. In *Proceedings of the 5th Annual ACM Web Science Conference on Web Science - WebSci '13*, 24–32. https://doi.org/10.1145/2464464.2464489

27. Susanne Bødker. 1996. Applying Activity Theory to Video Analysis: How to Make SEnse of Video Data in Human-Computer Interaction. In *Context and consciousness: Activity theory and human-computer interaction*. 147–174.

28. Susanne Bødker. 1996. Applying Activity Theory to Video Analysis: How to Make Sense of Video Data in Human-Computer Interaction. In *Context and consciousness: Activity theory and human-computer interaction*, Bonnie A. Nardi (ed.). MIT Press, 147–174.

29. Susanne Bødker and Clemens Nylandsted Klokmose. 2012. Dynamics in artifact ecologies. *Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design*: 448–457. https://doi.org/10.1145/2399016.2399085

30. Alan Borning and Michael Muller. 2012. Next steps for value sensitive design. *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12*: 1125. https://doi.org/10.1145/2207676.2208560

31. Alan Borning, Paul Waddell, and Ruth Forster. 2008. Urbansim: Using Simulation to Inform Public Deliberation and Decision-Making. *Digital Government. Integrated Series in Information Systems Vol 17*: 439–464. https://doi.org/10.1007/978-0-387-71611-4_22

32. Luc Bovens. 2009. The Ethics of Nudge. *Preference change: approaches from philosophy, economics and psychology*: 207–219. https://doi.org/10.1007/978-90-481-2593-7_10

33. danah m. Boyd. 2010. Social Network Sites as Networked Publics: Affordances, Dynamics, and Implications. In *Networked Self: Identity, Community, and Culture on*

Social Network Sites, Zizi Papacharissi (ed.). 39–58. https://doi.org/10.1162/dmal.9780262524834.119

34. Karen Boyd, Adam Rule, Aurélien Tabard, and Jim Hollan. 2016. Human Values in Tracking and Sharing Computer-Mediated Activity at Work. *Cwcs*: 233–236. https://doi.org/10.1145/2818052.2869119

35. Ben Bradley. 2013. *The Cambridge Companion to Utilitarianism*. Cambridge University Press, Cambridge. https://doi.org/10.1017/CCO9781139096737

36. James L. Buckley. 1974. *Family Educational Rights and Privacy Act*. United States of America.

37. Ramnath K. Chellappa and Raymond G. Sin. 2005. Personalization versus Privacy: An Empirical Examination of the Online Consumer's Dilemma. *Information Technology and Management* 6, 2–3: 181–202. https://doi.org/10.1007/s10799-005-5879-y

38. Frank X. Chen, Abby C. King, and Eric B. Hekler. 2014. "Healthifying" Exergames: Improving Health Outcomes through Intentional Priming. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*: 1855–1864. https://doi.org/10.1145/2556288.2557246

39. Woohyeok Choi, Jeongmin Oh, Darren Edge, Joohyun Kim, and Uichin Lee. 2016. SwimTrain: Exploring Exergame Design for Group Fitness Swimming. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*: 1692–1704. https://doi.org/10.1145/2858036.2858579

40. James Clawson, Jessica A Pater, Andrew D Miller, Elizabeth D Mynatt, and Lena Mamykina. 2015. No longer wearing. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '15*, 647–658. https://doi.org/10.1145/2750858.2807554

41. Torkil Clemmensen, Victor Kaptelinin, and Bonnie Nardi. 2016. Making HCI theory work: an analysis of the use of activity theory in HCI research. *Behaviour & Information Technology* 35, 8: 608–627. https://doi.org/10.1080/0144929X.2016.1175507

42. Caleb Conner. 2016. Correcting Exercise Form Using Body Tracking. *CHI Extended Abstracts on Human Factors in Computing Systems*: 3028–3034. https://doi.org/10.1145/2851581.2892519

43. Sunny Consolvo, Katherine Everitt, Ian Smith, and James a. Landay. 2006. Design requirements for technologies that encourage physical activity. *Proceedings of the SIGCHI*

conference on Human Factors in computing systems - CHI '06: 457–466. https://doi.org/10.1145/1124772.1124840

44. Council of Europe. 1953. *European Convention on Human Rights*.

45. Martin V Covington and J M Kimberly. 2001. Intrinsic Versus Extrinsic Motivation: An Approach / Avoidance Reformulation. *Educational Psychology Review* 13, 2: 157–176. https://doi.org/10.1023/A:1009009219144

46. LF Cranor. 2012. Can Users Control Online Behavioral Advertising Effectively? *Security & Privacy, IEEE*, April. Retrieved September 8, 2012 from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6173007

47. Michele L Crossley. 2000. *Rethinking Health Psychology*. Open University Press, Buckingham.

48. N. Crossley. 2006. In the Gym: Motives, Meaning and Moral Careers. *Body & Society* 12, 3: 23–50. https://doi.org/10.1177/1357034X06067154

49. Christopher a. Le Dantec, Erika Shehan E.S. Poole, and S.P. Susan P. Wyche. 2009. Values as lived experience: Evolving value sensitive design in support of value discovery. *Proceedings of the 27th international conference on Human factors in computing systems (CHI '09)*: 1141–1150. https://doi.org/10.1145/1518701.1518875

50. Edward L Deci, Richard Koestner, and Richard M Ryan. 2001. Extrinsic Rewards and Intrinsic Motivation in Education: Reconsidered Once Again. *Review of Educational Research Spring* 71, 1: 1–27. https://doi.org/10.3102/00346543071001001

51. T Denning, Db Kramer, B Friedman, Matthew R. Reynolds, Brian Gill, and Tadayoshi Kohno. 2014. CPS: beyond usability: applying value sensitive design based methods to investigate domain characteristics for security for implantable cardiac devices. *Proceedings of the 30th Annual Computer Security Applications Conference on ACSAC '14*: 426–435. Retrieved from http://www.cs.utah.edu/~tdenning/files/papers/denning-ACSAC-2014-IMDs-VSD.pdf

52. Ed Diener. 2000. Subjective well-being: The science of happiness and a proposal for a national index. *American Psychologist* 55, 1: 34–43. https://doi.org/10.1037//0003-066X.55.1.34

53. Sebastian Döweling, Benedikt Schmidt, and Andreas Göb. 2012. A model for the design of interactive systems based on activity theory. *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work - CSCW '12*: 539. https://doi.org/10.1145/2145204.2145287

54. S. Duguay. 2014. "He has a way gayer Facebook than I do": Investigating sexual identity disclosure and context collapse on a social networking site. *New Media & Society*. https://doi.org/10.1177/1461444814549930

55. Anthony Dunne and Fiona Raby. 2013. Design as Critiqiue. In *Speculative Everything*. 33–45.

56. Anthony Dunne and Fiona Raby. 2013. *Speculative Everything: Design, Fiction, and Social Dreaming (MIT Press)*. The MIT Press, London. Retrieved from https://www.amazon.com/Speculative-Everything-Design-Fiction-Dreaming-ebook/dp/B00H5HA5VU%3FSubscriptionId%3D0JYN1NVW651KCA56C102%26tag%3Dtechkie-20%26linkCode%3Dxm2%26camp%3D2025%26creative%3D165953%26creativeASIN%3DB00H5HA5VU

57. Abigail C. Durrant, David S. Kirk, and Stuart Reeves. 2014. Human values in curating a human rights media archive. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*: 2685–2694. https://doi.org/10.1145/2556288.2557196

58. eMarketer. 2016. Ad Blocking to Jump by Double Digits This Year. Retrieved January 13, 2017 from http://www.emarketer.com/article/us-ad-blocking-jump-by-double-digits-this-year/1014111

59. Daniel A Epstein, Alan Borning, and James Fogarty. 2013. Fine-grained sharing of sensed physical activity. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing - UbiComp '13*, 489. https://doi.org/10.1145/2493432.2493433

60. Evidon Inc. 2012. Ghostery. Retrieved from http://www.ghostery.com/

61. Ayman Farahat and Michael C. Bailey. 2012. How effective is targeted advertising? *Proceedings of the 21st international conference on World Wide Web - WWW '12*: 111. https://doi.org/10.1145/2187836.2187852

62. Katherine M. Flegal. 2012. Prevalence of Obesity and Trends in the Distribution of Body Mass Index Among US Adults, 1999-2010. *JAMA: The Journal of the American Medical Association 307*, 491. https://doi.org/10.1001/jama.2012.39

63. B.J. Fogg. 2003. *Persuasive Technology*. Morgan Kauffmann Publishers, San Francisco.

64. B.J. Fogg. 2009. A behavior model for persuasive design. In *Proceedings of the 4th International Conference on Persuasive Technology - Persuasive '09* (Persuasive '09), 1. https://doi.org/10.1145/1541948.1541999

65. B.J. Fogg. 2009. Creating persuasive technologies. In *Proceedings of the 4th International Conference on Persuasive Technology - Persuasive '09*, 1. https://doi.org/10.1145/1541948.1542005

66. Batya Friedman and David G. Hendry. 2012. The envisioning cards: A toolkit for catalyzing humanistic and technical imaginations. *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems - CHI '12*: 1145–1148. https://doi.org/10.1145/2207676.2208562

67. Batya Friedman, DC Howe, and Edward Felten. 2002. Informed consent in the Mozilla browser: Implementing value-sensitive design. In *Proc 35th Hawaii Intl Conference on Systems Sciences*, 1–10. Retrieved January 17, 2014 from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=994366

68. Batya Friedman, Peter H. Kahn, Alan Borning, and Alina Huldtgren. 2013. Value Sensitive Design and Information Systems. In *Human Computer Interaction and Management Information Systems: Foundations Advances in Management Information Systems*, Ping Zhang and Dennis Galletta (eds.). Armonk, NY, 55–95. https://doi.org/10.1007/978-94-007-7844-3_4

69. Batya Friedman and Peter H Kahn Jr. 2000. New Directions: A Value-Sensitive Design Approach to Augmented Reality. *Dare 2000*: 163–164. https://doi.org/10.1145/354666.354694

70. Thomas Fritz, Elaine M Huang, Gail C Murphy, and Thomas Zimmermann. 2014. Persuasive technology in the real world. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, 487–496. https://doi.org/10.1145/2556288.2557383

71. Göran Goldkuhl and Stefan Cronholm. 2010. Adding Theoretical Grounding to Grounded Theory: Toward Multi-Grounded Theory. *International Journal of Qualitative Methods* 9, 2: 187–205. https://doi.org/10.1177/160940691000900205

72. Richard Gomer, Eduarda Mendes Rodrigues, Natasa Milic-Frayling, and M.C. Schraefel. 2013. Network Analysis of Third Party Tracking: User Exposure to Tracking Cookies through Search. In *2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, 549–556. https://doi.org/10.1109/WI-IAT.2013.77

73. Víctor M. González, Bonnie Nardi, and Gloria Mark. 2009. Ensembles: understanding the instantiation of activities. *Information Technology & People* 22, 2: 109–131. https://doi.org/10.1108/09593840910962195

74. Rúben Gouveia, Evangelos Karapanos, and Marc Hassenzahl. 2015. How do we engage with activity trackers?: a longitudinal study of Habito. *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*: 1305–1316. https://doi.org/10.1145/2750858.2804290

75. Tjerk de Greef, Arvind Mohabir, Ibo van der Poel, and Mark Neerincx. 2013. sCEthics: Embedding Ethical Values in Cognitive Engineering. 4:1–4:8. https://doi.org/10.1145/2501907.2501935

76. Jonathan Gruber. 2002. Smoking's "Internalities." *Regulation*, 52–57. Retrieved from http://www.pierrekopp.com/downloads/Smokings Internalities.pdf

77. Jonathan Grudin. 1990. The Computer Reaches Out: The Historical Continuity of Interface Design. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Empowering People (CHI'90)*: 261–268. https://doi.org/http://doi.acm.org/10.1145/97243.97284

78. M. Hagstromer, P. Oja, and M. Sjostrom. 2006. The International Physical Activity Questionnaire (IPAQ): a study of concurrent and construct validity. *Public health nutrition* 9, 6: 755–762. https://doi.org/10.1079/phn2005898

79. Lars Hallnäs and Johan Redström. 2001. Slow technology - designing for reflection. *Personal and Ubiquitous Computing* 5, 3: 201–212. https://doi.org/10.1007/PL00000019

80. J Hardman. 2007. Making sense of the meaning maker: tracking the object of activity in a computer-based mathematics lesson using activity theory. *International Journal of Education and Development using ICT* 3, 4: 110–130.

81. Richard Harper, Tom Rodden, Yvonne Rogers, and Abigail Sellen. 2008. *Being Human - Human-Computer Interaction in the Year 2020*. https://doi.org/citeulike-article-id:3349429

82. Daniel Harrison, Nadia Berthouze, Paul Marshall, and Jon Bird. 2014. Tracking physical activity. *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing Adjunct Publication - UbiComp '14 Adjunct*: 699–702. https://doi.org/10.1145/2638728.2641320

83. Eric B. Hekler, Predrag Klasnja, Jon E. Froehlich, and Matthew P. Buman. 2013. Mind the theoretical gap. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13* 51, 4: 3307. https://doi.org/10.1145/2470654.2466452

84. Floris Heukelom. 2007. Kahneman and Tversky and the Origin of Behavioral Economics. *Tinbergen Institute Discussion Paper*: 1–25. https://doi.org/10.2139/ssrn.956887

85. Felicia a Huppert. 2009. Psychological well-being: evidence regarding its causes and consequences. *Applied Psychology: Health and Well-Being* 1, 2: 137–164. https://doi.org/10.1111/j.1758-0854.2009.01008.x

86. ICO (Information Commissioners Office). 2010. *Privacy notices code of practice*. Retrieved from http://www.ico.gov.uk/for_organisations/data_protection/topic_guides/~/media/documents/library/Data_Protection/Detailed_specialist_guides/PRIVACY_NOTICES_COP_FINAL.ashx

87. Lucas D. Introna. 2007. Maintaining the reversibility of foldings: Making the ethics (politics) of information technology visible. *Ethics and Information Technology* 9, 1: 11–25. https://doi.org/10.1007/s10676-006-9133-z

88. Carlos Jensen and Colin Potts. 2004. Privacy policies as decision-making tools. *Proceedings of the 2004 conference on Human factors in computing systems - CHI '04* 6, 1: 471–478. https://doi.org/10.1145/985692.985752

89. Daniel Kahneman. 2003. A perspective on judgment and choice: mapping bounded rationality. *The American psychologist* 58, 9: 697–720. https://doi.org/10.1037/0003-066X.58.9.697

90. Victor Kaptelinin. 1996. Activity Theory: Implications for Human-Computer Interaction. In *Context and Consciousness*. 103–116.

91. Victor Kaptelinin. 1996. Computer-Mediated Activity: Functional Organs in Social and Developmental Contexts. In *Context and Consciousness*, Bonnie A. Nardi (ed.). MIT Press, Cambridge, MA, 45–68.

92. Victor Kaptelinin. 2003. UMEA: Translating Interaction Histories into Project Contexts. *Proceedings of the conference on Human factors in computing systems - CHI '03* 5, 1: 353–360. https://doi.org/10.1145/642611.642673

93. Victor Kaptelinin. 2005. The object of activity: Making sense of the sense-maker. *Mind, Culture, and Activity* 12, 1: 4–18. https://doi.org/10.1207/s15327884mca1201

94. Victor Kaptelinin and Bonnie Nardi. 2012. Affordances in HCI: toward a mediated action perspective. *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems*: 967–976. https://doi.org/10.1145/2208516.2208541

95. Victor Kaptelinin and Bonnie A Nardi. 2003. Objects of Desire. In *Acting with Technology: Activity theory and interaction design*. The MIT Press, 153–172.

96. Victor Kaptelinin and Bonnie A Nardi. 2009. *Acting with Technology: Activity Theory and Interaction Design*. The MIT Press. Retrieved from https://www.amazon.com/Acting-Technology-Activity-Theory-Interaction-ebook/dp/B002RDDI0W?SubscriptionId=0JYN1NVW651KCA56C102&tag=techkie-20&linkCode=xm2&camp=2025&creative=165953&creativeASIN=B002RDDI0W

97. Victor Kaptelinin and Bonnie A. Nardi. 2006. Do We Need Theory in Interaction Design? In *Acting with Technology: Activity theory and interaction design*. MIT Press, Cambridge, MA, 15–28.

98. Victor Kaptelinin and Bonnie A. Nardi. 2006. Activity Theory in a Nutshell. In *Acting with Technology: Activity theory and interaction design*. MIT Press, Cambridge, MA, 29–72.

99. Victor Kaptelinin and Bonnie A. Nardi. 2006. Artifacts, Agency, and (A)symmetry. In *Acting with Technology: Activity theory and interaction design2*. 237–252.

100. Victor Kaptelinin and Bonnie A. Nardi. 2009. Interaction Design Informed by Activity Theory. In *Acting with Technology: Activity theory and interaction design*. 73–115.

101. Timothy Koschmann, Kari Kuutti, and Larry Hickman. 1998. The Concept of Breakdown in Heidegger , Leont ' ev, and Dewey and Its Implications for Education. *Mind Culture and Activity* 5, 1: 25–41. https://doi.org/10.1207/s15327884mca0501

102. Ponnurangam Kumaraguru and Lf Cranor. 2005. *Privacy indexes: A survey of westin's studies*. Retrieved from http://www.casos.cs.cmu.edu/publications/papers/CMU-ISRI-05-138.pdf%5Cnhttp://repository.cmu.edu/isr/856/

103. Kari Kuutti. 1996. Activity theory as a potential framework for human-computer interaction research. In *Context and consciousness: Activity theory and human-computer interaction*, Bonnie A. Nardi (ed.). MIT Press, Cambridge, MA, 17–44. https://doi.org/citeulike-article-id:634717

104. Stacey Kuznetsov. 2006. Motivations of contributors to Wikipedia. *ACM SIGCAS Computers and Society* 36, 2: 1–es. https://doi.org/10.1145/1215942.1215943

105. D. Laibson. 1997. Golden Eggs and Hyperbolic Discounting. *The Quarterly Journal of Economics* 112, 2: 443–478. https://doi.org/10.1162/003355397555253

106. V Latora and M Marchiori. 2001. Efficient behavior of small-world networks. *Physical review letters* 87, 19: 198701. https://doi.org/10.1103/PhysRevLett.87.198701

107. Amanda Lazar, Joshua Tanenbaum, Christian Koehler, and David H Nguyen. 2015. Why We Use and Abandon Smart Devices. *UbiComp*: 635–646. https://doi.org/10.1145/2750858.2804288

108. David Lazer, Ryan Kennedy, Gary King, and Alessandro Vespignani. 2014. Google Flu Trends Still Appears Sick: An Evaluation of the 2013-2014 Flu Season. *SSRN Electronic Journal*: 1–11. https://doi.org/10.2139/ssrn.2408560

109. David Lazer, Ryan Kennedy, Gary King, and Alessandro Vespignani. 2014. The Parable of Google Flu: Traps in Big Data Analysis. *Science* 343, 6167: 1203–1205. https://doi.org/10.1126/science.1248506

110. Dan Ledger. 2014. *Inside Wearables - Part 2*. Retrieved from http://endeavourpartners.net/assets/Endeavour-Partners-Inside-Wearables-Part-2-July-2014.pdf

111. Min Kyung Lee, Junsung Kim, Jodi Forlizzi, and Sara Kiesler. 2015. Personalization revisited. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '15*, 743–754. https://doi.org/10.1145/2750858.2807552

112. Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. 2012. Why Johnny can't opt out. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12*, 589. https://doi.org/10.1145/2207676.2207759

113. Ian Li, Anind Dey, and Jodi Forlizzi. 2011. Understanding My Data, Myself: Supporting Self-Reflection with Ubicomp Technologies. *Proceedings of the 13th international conference on Ubiquitous computing*: 405–414. https://doi.org/10.1145/2030112.2030166

114. Ian Li, Anind K. Dey, and Jodi Forlizzi. 2012. Using context to reveal factors that affect physical activity. *ACM Transactions on Computer-Human Interaction 19*, 1–21. https://doi.org/10.1145/2147783.2147790

115. Health and Social Care Information Centre Lifestyle Statistics Team. 2014. *Statistics on Obesity, Physical Activity and Diet - England, 2014*. Leeds. Retrieved from http://www.hscic.gov.uk/catalogue/PUB13648/Obes-phys-acti-diet-eng-2014-rep.pdf

116. James J. Lin, Lena Mamykina, Silvia Lindtner, Gregory Delajoux, and Henry B. Strub. 2006. Fish'n'Steps: Encouraging Physical Activity with an Interactive Computer Game. In *UbiComp 2006: Ubiquitous Computing* (Lecture Notes in Computer Science). https://doi.org/10.1007/11853565

117. Rebecca Lipman. 2016. Online Privacy and the Invisible Market for Our Data. *Penn State Law Review*: 1–34.

118. Bin Liu, Anmol Sheth, Udi Weinsberg, Jaideep Chandrashekar, and Ramesh Govindan. 2013. AdReveal: Improving Transparency Into Online Targeted Advertising. In *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks - HotNets-XII*, 1–7. https://doi.org/10.1145/2535771.2535783

119. Mary Madden and Lee Rainie. 2015. Americans' attitudes about privacy, security and surveillance. *Pew Research Center*: 47. Retrieved from http://www.pewinternet.org/2015/05/20/americans-attitudes-about-privacy-security-and-surveillance/

120. Giorgio Maone. 2012. NoScript. *Firefox Add-Ons2*. Retrieved from https://addons.mozilla.org/en-US/firefox/addon/noscript/

121. Helia Marreiros, Richard Gomer, and Mirco Tonin. 2015. Exploring user perceptions of online privacy disclosures. In *Proceedings of 14th International Conference on WWW/INTERNET 2015*. Retrieved from http://eprints.soton.ac.uk/381134/

122. David Martin, Hailin Wu, and Adil Alsaid. 2003. Hidden surveillance by Web sites: Web bugs in contemporary use. *Communications of the ACM* 46, 12: 258–264. Retrieved November 26, 2012 from http://dl.acm.org/citation.cfm?id=953509

123. Alice E. Marwick and danah boyd. 2011. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society* 13, 1: 114–133. https://doi.org/10.1177/1461444810365313

124. Jonathan R Mayer and John C Mitchell. 2012. Third-Party Web Tracking : Policy and Technology. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy*, 413–427. https://doi.org/10.1109/SP.2012.47

125. A McDonald and L Cranor. 2010. Beliefs and behaviors: Internet users' understanding of behavioral advertising. *38th Research Conference on Communication, Information and Internet Policy*: 1–31. Retrieved November 27, 2012 from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1989092

126. Aleecia M. Mcdonald, Robert W. Reeder, Patrick Gage Kelley, and Lorrie Faith Cranor. 2009. A Comparative Study of Online Privacy Policies and Formats. *PETS '09 Proceedings of the 9th International Symposium on Privacy Enhancing Technologies* 5672: 37–55. https://doi.org/10.1007/978-3-642-03168-7

127. Aleecia McDonald, Robert Reeder, Patrick Kelley, and Lorrie Cranor. 2009. A Comparative Study of Online Privacy Policies and Formats. *PETS '09: Proceedings of the 9th International Symposium on Privacy Enhancing Technologies*: 37–55. https://doi.org/10.1145/1572532.1572586

128. Jochen Meyer, Jochen Schnauber, Wilko Heuten, Harm Wienbergen, Rainer Hambrecht, Hans Jurgen Appelrath, and Susanne Boll. 2016. Exploring Longitudinal Use of Activity Trackers. *Proceedings - 2016 IEEE International Conference on Healthcare Informatics, ICHI 2016*: 198–206. https://doi.org/10.1109/ICHI.2016.29

129. Susan Michie, Maartje M van Stralen, and Robert West. 2011. The behaviour change wheel: A new method for characterising and designing behaviour change interventions. *Implementation Science* 6, 1: 42. https://doi.org/10.1186/1748-5908-6-42

130. J.K. Miller, B. Friedman, G. Jancke, and B. Gill. 2007. Value tensions in design: The value sensitive design, development, and appropriation of a corporation's groupware system. *Proceedings of the 2007 International ACM Conference on Supporting Group Work (GROUP '07)*: 281–290. https://doi.org/10.1145/1316624.1316668

131. Stuart Moran, Ewa Luger, and Tom Rodden. 2014. Literatin: Beyond Awareness of Readability in Terms and Conditions. In *UbiComp*, 641–646. https://doi.org/10.1145/2638728.2641684

132. Bonnie A Nardi. 1996. Studying Context: A Comparison of Activity Theory, Situated Action Models, and Distributed Cognition. In *Context and Consciousness*, Bonnie A Nardi (ed.). MIT Press, 69–102.

133. Bonnie A. Nardi. 1996. Activity Theory and Human-Computer Interaction. In *Context and consciousness: Activity theory and human-computer interaction*, Bonnie A. Nardi (ed.). MIT Press, Cambridge, MA, 7–16.

134. Bonnie A. Nardi, Diane J. Schiano, and Michelle Gumbrecht. 2004. Blogging as social activity, or, would you let 900 million people read your diary? *Proceedings of the 2004 ACM conference on Computer supported cooperative work - CSCW '04*: 222. https://doi.org/10.1145/1031607.1031643

135. Lisa P. Nathan. 2008. Ecovillages, values, and interactive technology. *Proceeding of the twenty-sixth annual CHI conference extended abstracts on Human factors in computing systems - CHI '08*: 3723. https://doi.org/10.1145/1358628.1358920

136. M.E.J. Newman and D.J. Watts. 1999. Renormalization group analysis of the small-world network model. *Physics Letters A* 263, 4–6: 341–346. https://doi.org/10.1016/S0375-9601(99)00757-4

137. M E Newman and D J Watts. 1999. Scaling and percolation in the small-world network model. *Physical review. E, Statistical physics, plasmas, fluids, and related interdisciplinary topics* 60, 6 Pt B: 7332–7342. https://doi.org/10.1103/PhysRevE.60.7332

138. Patricia A. Norberg, Daniel R. Horne, and David A. Horne. 2007. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* 41, 1: 100–126. https://doi.org/10.1111/j.1745-6606.2006.00070.x

139. OECD. 1980. *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*. Retrieved from http://www.oecd.org/internet/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm

140. Ilse Oosterlaken. 2014. Applying Value Sensitive Design (VSD) to Wind Turbines and Wind Parks: An Exploration. *Science and Engineering Ethics* 21, 2: 359–379. https://doi.org/10.1007/s11948-014-9536-x

141. Wladimir Palant. 2012. Adblock Plus. *Firefox Add-Ons*. Retrieved from https://addons.mozilla.org/en-US/firefox/addon/adblock-plus/

142. Leysia Palen and Paul Dourish. 2003. Unpacking "privacy" for a networked world. *Proceedings of the conference on Human factors in computing systems - CHI '03*, 5: 129. https://doi.org/10.1145/642633.642635

143. Mitesh S Patel, David a Asch, and Kevin G Volpp. 2015. Wearable Devices as Facilitators , Not Drivers , of Health Behavior Change. *JAMA : the journal of the American Medical Association* 313, 5: 459–460. https://doi.org/10.1001/jama.2014.14781.Conflict

144. Alan D Penman and William D Johnson. 2006. The Changing Shape of the Body Mass Index Distribution Curve in the Population: Implications for Public Health Policy to Reduce the Prevalence of Adult ObesityNo Title. *Preventing Chronic Disease* 3, 3. Retrieved from http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1636707/

145. RE Petty and JT Cacioppo. 1986. The elaboration likelihood model of persuasion. *Advances in experimental social psychology* 19. Retrieved August 29, 2011 from http://psychology.uchicago.edu/people/faculty/cacioppo/jtcreprints/pc86.part1.pdf

146. James O. Prochaska and Wayne F. Velicer. 1997. The Transtheoretical Model of Health Behaviour Change. *American Journal of Health Promotion 12*, 38–48. https://doi.org/10.4278/0890-1171-12.1.38

147. Stephen Purpura, Victoria Schwanda, Kaiton Williams, William Stubler, and Phoebe Sengers. 2011. Fit4life. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*, 423. https://doi.org/10.1145/1978942.1979003

148. Xiangang Qin, Chee-wee Tan, and Torkil Clemmensen. 2017. Context-Awareness and Mobile HCI: Implications, Challenges and Opportunities. In *HCI in Business, Government and Organizations. Interacting with Information Systems*, Fiona Fui-Hoon Nah and Chuan-Hoo Tan (eds.). Springer International Publishing, 112–127. https://doi.org/10.1007/978-3-319-58481-2

149. S. Reeves, M. Albert, A. Kuper, and B. D. Hodges. 2008. Why use theories in qualitative research? *BMJ* 337, aug07 3: a949–a949. https://doi.org/10.1136/bmj.a949

150. Joel R Reidenberg, Travis Breaux, Lorrie Faith Carnor, Brian French, Lorrie Faith Cranor, Amanda Grannis, James T Graves, Fei Liu, Aleecia Mcdonald, Thomas B Norton, Rohan Ramanath, N Cameron Russell, Norman Sadeh, and Florian Schaub. 2014. Disagreeable Privacy Policies: Mismatches Between Meaning and Users' Understanding. *Berkeley Technology Law Journal* 30, 1: 39–88. https://doi.org/10.15779/Z384K33

151. Wolfgang H. Reitberger, Wolfgang Spreicer, and Geraldine Fitzpatrick. 2014. Nutriflect: Reflecting Collective Shopping Behavior and Nutrition. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, 3309–3318. https://doi.org/10.1145/2556288.2557384

152. Yvonne Rogers. 2012. HCI Theory: Classical, Modern, and Contemporary. *Synthesis Lectures on Human-Centered Informatics* 5, 2: 1–129. https://doi.org/10.2200/S00418ED1V01Y201205HCI014

153. Milton Rokeach. 1973. *The Nature of Human Values*. New York, NY, USA.

154. Beate Rőssler. 2005. *The Value of Privacy*. Polity Press.

155. Ira S. Rubinstein. 2011. Regulating Privacy by Design. *Berkeley Technology Law Journal* 26: 1409.

156. Georg Ruckriem. 2009. Digital Technology and Mediation. In *Learning and Expanding with Activity Theory*, Harry Daniels and Kris D. Gutierrez (eds.). 88–111.

157. Scott J Savage and Donald M Waldman. 2013. The Value of Online Privacy. 1–48. https://doi.org/10.2139/ssrn.2341311

158. Donald a Schön. 1983. The reflective practitioner. *Pediatrics* 116, 6: 1546–52. https://doi.org/10.1542/peds.2005-0209

159. S. H. Schwartz, G. Melech, a. Lehmann, S. Burgess, M. Harris, and V. Owens. 2001. Extending the Cross-Cultural Validity of the Theory of Basic Human Values with a Different Method of Measurement. *Journal of Cross-Cultural Psychology* 32, 5: 519–542. https://doi.org/10.1177/0022022101032005001

160. Shalom H Schwartz. 1994. Are there universal aspects in the structure and contents of human values? *Journal of Social Issues* 50, 4: 19–45. https://doi.org/10.1111/j.1540-4560.1994.tb01196.x

161. Shalom H Schwartz. 2003. A Proposal for Measuring Value Orientations across Nations. *Questionnaire Package of the European Social Survey*: 259–319. https://doi.org/10.1016/B978-0-12-411466-1.00007-0

162. Evan Selinger and Kyle Whyte. 2011. Is There a Right Way to Nudge? The Practice and Ethics of Choice Architecture. *Sociology Compass* 5, 10: 923–935. https://doi.org/10.1111/j.1751-9020.2011.00413.x

163. Anselm Strauss and Juliet Corbin. 1994. Grounded theory methodology. *Handbook of qualitative research*, 273–285. https://doi.org/10.1007/BF00988593

164. Haichun Sun. 2013. Impact of exergames on physical activity and motivation in elementary school students: A follow-up study. *Journal of Sport and Health Science* 2, 3: 138–145. https://doi.org/10.1016/j.jshs.2013.02.003

165. Alistair Sutcliffe and Abdallah Namoun. 2012. Predicting user attention in complex web pages. *Behaviour and Information Technology* 31, 7: 679–695. https://doi.org/10.1080/0144929X.2012.692101

166. Monika Taddicken. 2014. The "Privacy Paradox" in the Social Web: The Impact of Privacy Concerns, Individual Characteristics, and the Perceived Social Relevance on Different Forms of Self-Disclosure1. *Journal of Computer-Mediated Communication* 19, 2: 248–273. https://doi.org/10.1111/jcc4.12052

167. Richard H Thaler and Cass R Sunstein. 2008. *Nudge: Improving Decisions About Health, Wealth and Happiness*. Yale University Press, New Haven & London.

168. The European Parliament and the Council of the European Union. 2016. *Regulation (EU) 2016/679 General Data Protection Regulation*. Retrieved from http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679&from=en

169. Amos Tversky and Daniel Kahnemann. 1974. Judgment under uncertainty: Heuristics and Biases. *Science* 185, 4157: 1124–1131. https://doi.org/10.1126/science.185.4157.1124

170. University of Washington Value Sensitive Design Research Lab. 2011. Envisioning Cards. Retrieved August 10, 2017 from http://www.envisioningcards.com/

171. Blase Ur, PG Leon, and LF Cranor. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. *Proceedings of the Eighth …*. Retrieved November 26, 2012 from http://dl.acm.org/citation.cfm?id=2335362

172. W3C. 2012. Tracking Preference Expression (DNT). Retrieved from http://www.w3.org/2011/tracking-protection/drafts/tracking-dnt.html

173. SD Warren and LD Brandeis. 1890. The right to privacy. *Harvard law review* 4, 5: 193–220. Retrieved November 1, 2012 from http://www.jstor.org/stable/10.2307/1321160

174. Gordon Wells. 2011. Motive and Motivation in Learning to Teach.

175. Robert West. 2005. Time for a change: Putting the Transtheoretical (Stages of Change) Model to rest. *Addiction* 100, 8: 1036–1039. https://doi.org/10.1111/j.1360-0443.2005.01139.x

176. Alan F Westin and Robert R Belair. 2008. How Online Users Feel About Behavioral Marketing and How Adoption of Privacy and Security Policies Could Affect Their Feelings Results of a Harris Interactive / Westin Survey Sponsored by Privacy Consulting Group. *Group*: 1–10. Retrieved from https://www.ftc.gov/sites/default/files/documents/public_comments/privacy-roundtables-comment-project-no.p095416-544506-00052/544506-00052.pdf

177. Glen Whitman. 2006. Internalities and the Economics of Self-Control Against the New Paternalism. *Policy Analysis*, 563: 1–16. Retrieved from http://object.cato.org/sites/cato.org/files/pubs/pdf/pa563.pdf

178. T. D. Wilson. 2006. A re-examination of information seeking behaviour in the context of activity theory. *Information Research* 11, 4: 260. https://doi.org/260

179. Heng Xu, Robert E. Crossler, and France Bélanger. 2012. A value sensitive design investigation of privacy enhancing tools in web browsers. *Decision Support Systems* 54, 1: 424–433. https://doi.org/10.1016/j.dss.2012.06.003

180. Rayoung Yang, Eunice Shin, Mark W. Newman, and Mark S. Ackerman. 2015. When fitness trackers don't "fit": end-user difficulties in the assessment of personal tracking device accuracy. 623–634. https://doi.org/10.1145/2750858.2804269

181. Bieke Zaman and Nassim Jafari Naimi. 2015. A Value Sensitive Design Case Study : Why Values Do (not) Design. In *Charting the Next Decade for Value Sensitive Design Workshop*. Retrieved from https://lirias.kuleuven.be/bitstream/123456789/515711/1/Why+Values+do+(not)+design_Zaman_JafariNaimi.pdf

182. Zhao Zhao, S. Ali Etemad, and Ali Arya. 2016. Gamification of exercise and fitness using wearable activity trackers. In *Advances in Intelligent Systems and Computing*, 233–240. https://doi.org/10.1007/978-3-319-24560-7_30

183. KDD Cup '05 Challenge. Retrieved from www.sigkdd.org/kdd2005/kddcup.html#ldata

184. 2009. *DIRECTIVE 2009/136/EC OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL*. European Union.

185. 2015. The Scott Trust: values and history. *The Guardian*. Retrieved October 22, 2017 from https://www.theguardian.com/the-scott-trust/2015/jul/26/the-scott-trust