CONSUMERS INTERNATIONAL
COMING TOGETHER
FOR CHANGE

HOW CAN CONSUMERS INTERNATIONAL CREATE POSITIVE CHANGE FOR CONSUMERS IN THE DIGITAL WORLD?

TECHNICAL OUTREACH (IDENTITY & PRIVACY), INTERNET SOCIETY ROBIN WILTON

The Internet Society advances the Internet as a global technical infrastructure, a resource to enrich people's lives, and a force for good in society.

Since 2001, **Robin Wilton** has specialised in digital identity, privacy and public policy, building a reputation as a thought-leader, communicator and translator between different stakeholder groups.

Before joining the Internet Society, Robin spent two years as a research analyst in Gartner's Identity and Privacy Strategies team, where in addition to his privacy work he specialised in public key infrastructure, electronic signature, single sign-on and federated identity.

Robin's experience includes: 12 years with IBM in systems engineering, technical support and consulting roles, at the UK and EMEA level; three years as Principal Consultant at JCP Trustbase Ltd, a start-up specialising in Java cryptography and PKI-enabling middleware; eight years with Sun Microsystems in technical pre-sales and the CTO team; 18 months establishing Future Identity Ltd. as an independent consultancy on privacy and digital identity. During his time at Future Identity he was also Director of Privacy and Public Policy for the Kantara Initiative.

To create positive change for consumers we need to look at two things; ethics and habits, as well as how we think about digital privacy.

WHAT'S THE DIGITAL PRIVACY PROBLEM?

A few weeks ago, preparing for a conference panel, I was asking people what they thought the problem was with digital privacy, from the consumer perspective. One person, after a little thought, replied that their experience of the privacy problem had gone through four phases:

- "I wasn't aware there's a problem."
- "OK, I see there's a problem, but why should I care?"
- "I care, but I don't know what I can do about it."
- "I tried to do something about my privacy, and now my browser/email/app doesn't work."

There is a model for this kind of experience: how people come to make decisions, and how those decisions turn into habits (or not). It is rather grandly called the 'transtheoretical model of behaviour change' – but it's a useful thing, despite the name, and we'll come back to it later.

The fourth phase is the point at which our attempts to improve privacy outcomes frequently break down. At that point, as a technologist, I'm often asked why there isn't some technical widget – an app, a browser plug-in, a black box – that can take care of a user's privacy on their behalf. The question is usually tinged with a degree of frustration.

I can sympathise with both the question and the frustration. After all, we can look back on some 30 years of data protection law, much of it based on the OECD's Guidelines governing the protection of privacy and transborder flows of personal data, which were adopted in September 1980. Those Guidelines, like the Council of Europe's Convention 108 of 1981, and the EU's Data Protection Directive of 1995, are long-standing enough to have matured and gone through at least one cycle of substantial review and renewal.

And yet, when we look at individuals' general experience of privacy and data protection, the outcomes don't appear to reflect either that maturity, or the effectiveness of the revised and updated guidance. Here are some of the common symptoms:

- Unexpected or excessive collection of personal data
- Insufficient care taken with its storage/use, leading to data breaches and inappropriate access
- Unexpected or unwelcome use
- Unexpected sharing

Individuals' expectations concerning their personal data are at odds with what actually happens - but why is there this misalignment? One answer I've been given is that "technology changes too fast; people just can't keep up, so their expectations lag behind reality". I'm not sure I buy that. After all, people seem to be adjusting fairly readily to the use of new technology; for instance, I've seen toddlers perfectly at ease with the user experience presented to them by a tablet computer, as is my 90-year-old mother.

Is the answer, perhaps, that the user experience gets a lot more design attention than the privacy experience? Possibly – and that's certainly the thinking behind the concept of 'privacy-by-design': to try and ensure that the privacy-related aspects of a product or service get as much attention as the rest of its design, and from as early as possible in the development process. But if that's the case, why is privacy is so slow to gain traction as a competitive differentiator? Technology products that make privacy a unique selling point such as Silent Circle's "Blackphone" handsets or Purism's laptops still seem to find favour with only a niche segment of the market, and that segment is often derided by other consumers as the 'tin-foil hat brigade'.[1,2] I know. I'm a card-carrying member of it, and even some of my colleagues can't understand why I get so concerned about potential privacy risks. So, the problem may be one of product design, but privacy-enhancing technologies still need to find the key to making usable privacy a 'must have' feature, rather than a 'why bother?'.

Privacy-by-design has to address a further design challenge, too. As an example here, think of a browser plug-in that alerts you every time a website tries to set a cookie. It probably wouldn't take long for most users to get bored and frustrated by constant warnings, and either ignore them, or disable the warning mechanism. Expose the function to the user in the wrong way, and no

THE DIGITAL PRIVACY PROBLEM IS HARD, NOT BECAUSE IT IS COMPLICATED, BUT BECAUSE IT IS SYSTEMIC

[1] Silent Circle's "Blackphone" handsets website, https://www.silentcircle.com/

[2] Purism's laptops website, https://puri.sm/

matter how worthy it is, they may reject it. However, at the other end of the scale, there's also a risk in shielding users too effectively from the complexities of what is being done on their behalf. Hide the function from the user completely, and they lose all awareness of what is happening – we don't want that outcome, either.

The ideal approach is to present these user-supportive functions in ways that increase comprehension and encourage adoption, rather than the reverse. This, too, relates to the behavioural model I referred to earlier, and to which I will return later.

And then there's the question of what happens to users' personal data when it is in the hands of third parties. So much of what happens on the Internet is driven by a powerful economic engine fuelled, in turn, by the monetisation of personal data. I have often heard monetization described as "the reason you can have free stuff, and cool innovation". The questions this raises for me are:

- Is that the only economic model available, or just the one with the greatest momentum?
- Is my privacy a fair price to pay for cool apps and free content?
- If I am paying for an app or service, or paying not to receive advertisements, does that guarantee that my data isn't being monetised?

In short, am I getting an honest bargain, and if not, how can I, as an individual, redress the balance between me and a multi-billion dollar corporation?

To recap, briefly: the problem of digital privacy involves elements of user awareness and choice; regulation and its effectiveness; technology design and adoption; data monetisation as an economic force… and that persistent mismatch between users' expectations and their experience.

ARE WE LOOKING AT THE PROBLEM IN THE RIGHT WAY?

As I noted above, individuals' reaction to the privacy problem is often accompanied by some frustration – and frankly, as a privacy advocate, so is mine. It often feels as though promising privacy-protecting efforts come to nothing, fizzle out without achieving critical mass, or fail to shift the behaviour of the market.

My theory is that this is not because the digital privacy problem is particularly complicated - or even particularly new, in some respects. After all, intermediaries have been collecting and monetising data about me since before the Internet. Rather, I think the digital privacy problem is hard because it's systemic. Multiple stakeholders are involved, many with differing motivations and sometimes conflicting interests; the influences that would change one stakeholder's behaviour won't work on some of the others, and the influences that work at one point in time may fail at another. The best way to change how we think about solving the problem is to change how we think about the problem. I packed a lot into this paragraph, so let's look at some specific examples, to make it less abstract.

First, what might motivate, say, the vendor of a connected object such as a smart light bulb? Probably, selling at a compelling price, achieving mass adoption, and maximising profit. Those motivations might lead to the following actions:

- ***Do as much as possible to minimise design/ manufacturing cost*** – If the cost of adding security or privacy functions doubles the price of your smart light bulb, relative to the competition, it might not sell.
- ***Sell on user functionality, not on vendor functionality*** – The user benefit is, say, the ability to control the lighting from your phone. The vendor functionality might include collection of data about usage patterns – but that isn't necessarily a compelling incentive for user adoption, so don't mention it.
- ***Increase your margins by monetising the data you collect about usage patterns, and the inferences you can draw from that data.***

I may be caricaturing slightly, here, but I think these are elements we can all see, to some extent, in the products and services offered to us in our connected lives. For a detailed examination of these issues, backed up by numerous case studies, I can recommend 'Networks of Control', by Wolfie Christl and Sarah Spiekermann of the Vienna University of Economics and Business.[3]

There are also elements, in my hypothetical example, of what economists call 'negative externalities'. That is: the vendor gets the benefits of data monetisation, while the

[3] Cracked Labs Website, http://crackedlabs.org/en/networksofcontrol

costs and risks associated with it fall on someone else (the consumer). For example, if the vendor suffers a data breach, and the personal data it has collected is abused, the resulting cost and harm fall on the consumer. Troy Hunt, a trainer and data breach consultant, has blogged recently about several worrying instances involving products aimed specifically at children.[4]

In an ideal world, some cost might return to the vendor in the form of legal penalties. However, the risk of that happening does not, currently, seem to influence vendor behaviour significantly in many jurisdictions - and particularly, when the hacking/abuse happens in a different jurisdiction from the vendor.

The topic of data breaches is one which the Internet Society examined in detail in its Global Internet Report for 2016, looking particularly at the economic factors and making five recommendations to build online trust.5 In summary, those are:

1. Put users at the centre of solutions; include externalities in the cost/benefit analysis.
2. Increase transparency through data breach notifications and disclosure.
3. Make data security must be a priority.
4. Make organisations should be accountable for their breaches.
5. Stimulate the market for independent security accreditation services.

Second, I want to return to the '*transtheoretical model*' I mentioned earlier. According to this model, people go through a number of stages in the course of making decisions. Good, clear explanations of the model can be found online, but this diagram gives a high-level summary:[6]

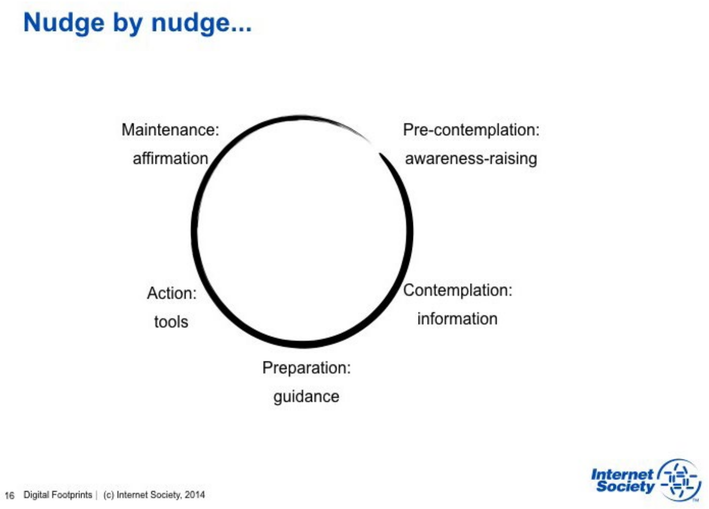
Nudge by nudge...

16 Digital Footprints | (c) Internet Society, 2014

USER ACTION ALONE IS UNLIKELY TO SUFFICE, BECAUSE OF THE POWERFUL NATURE OF THE ECONOMIC INFLUENCES THAT DRIVE SO MUCH INTERNET-RELATED COMMERCIAL ACTIVITY

According to the model, repeated iterations through this cycle can result in the formation of habits (whether good or bad); so, if we want to encourage consumers to form 'better' privacy habits, it should be useful to understand the formative process. What the model made clear to me was that, at each stage, the kind of intervention likely to succeed is different.

Think back to the first answer I got in the conversation I relayed at the beginning of this paper: "*I wasn't aware there's a problem*". That's the 'pre-contemplative' phase… I'm not even thinking about the problem, because I'm not aware of it. At that point, something has to make the individual aware of the problem.

Once they are aware of it, their next concern may be to find out if it's relevant to them: "OK, I see there's a problem, but why should I care?". At that point, they really want a binary answer: should I worry about this, yes or no?

At the next stage, though, a simple yes or no isn't enough: "*I care, but I don't know what I can do about it.*" Better outcomes depend on more information.

[4] 'Data from connected CloudPets teddy bears leaked and ransomed, exposing kids' voice messges', Troy Hunt's blog, 28/02/2017
[5] 'Internet Society's Global Internet Report for 2016', The Internet Society, 2016
[6] Study.com website; Transtheoretical model, http://study.com/academy/lesson/transtheoretical-model-definition-stages-of-change.html

So, even for these three 'informational' stages, we can see that different kinds of intervention are needed, if we are to respond to the individual's needs:

1. A compelling event that raises awareness of the problem
2. A quick, simple indication of its relevance to the individual and the need for action
3. Easy access to more information about what to do

The fourth answer I got was the point at which it all went wrong for the individual I was talking to. The model describes this as the 'Action' phase. The individual tried to fix the problem, only to find that the 'fix' broke their technology. This is the stage at which we are often inclined to expect some kind of technical widget to fix the problem, with greater or lesser success. We can also surmise that, if there is a working technical fix, but users aren't aware of the problem, or don't think it affects them, or don't know what to do about it, the technology is unlikely to see adoption. In other words, the 'action' phase can fail in numerous ways.

The final stage of the model is the 'Maintenance' phase. Here, the user's experience so far will influence whether or not they make the same choices next time they encounter the problem. Over time, whatever reinforcement they experience here (positive or negative) can lead to the formation of habit. So, for instance, if I make unhealthy eating choices but experience gratification in the 'maintenance' phase, and don't immediately keel over with a heart attack, I may well develop long-term bad eating habits, though in due course I may end up with furred arteries. Similarly, if I make poor privacy choices, experience gratification, and appear to suffer no ill consequences, I may continue with privacy-eroding habits until it's too late to repair the damage.

Part of the issue here is about the deferred consequences of poor privacy habits. The negative results of privacy-eroding behaviour are often remote, in time and place, from the action that caused them, so we tend not to 'learn the lesson'. By contrast, if I put my hand over a candle, I get negative feedback which I immediately associate with putting my hand in the flame, and I quickly form a habit of not doing so.

To recap: Habits develop as a result of an iterative decision-making cycle. To influence the formation of habits, we need to be able to intervene successfully in different ways, depending on the phase the individual has reached. Intervention at one phase may fail because it is poorly conceived (a technical 'fix' that breaks the user experience), or it may fail because previous phases have not been successfully addressed.

RECOMMENDATIONS FOR A NEW APPROACH

The digital privacy problem is hard, not because it is complicated, but because it is systemic. Different stakeholders have different incentives and will respond to different interventions. In the case of user motivations, there is a plausible model for how behavioural change takes place, and that model suggests that we should expect to intervene in different ways at different stages, if users are to develop awareness, motivation, capability, and privacy-enhancing habits.

However, user action alone is unlikely to suffice, because of the powerful nature of the economic influences that drive so much Internet-related commercial activity. Where market forces can be influenced, we should design the interventions that are likely to increase service providers' incentive to enhance privacy. Where market forces will predictably fail, there is a case to be made for regulatory intervention.

As noted earlier, this is a systemic problem - so the over-all approach should be prepared to apply different interventions to different stakeholders at different points in the process.

CHANGING STAKEHOLDERS' BEHAVIOUR

A key element of the Internet Society's proposed approach is to try to align the interests of consumers and service providers. We suggest that one way to do this is through the creation of a 'trust mark' that represents an organisation's commitment to ethical data-handling principles. Those principles, in turn, would reflect a set of policies and procedures that govern the organisation's collection and use of personal data. The organisation's entitlement to display the trust mark would be confirmed by an accreditation step and could then be monitored through external audit.

We believe this would give some service providers an incentive to distinguish themselves from the rest of the market, in much the same way as Fairtrade vendors do in the retail market. By analogy, the measure of success for a trust mark would not necessarily be 100% adoption by every vendor, but rather, the general shift in the market that results from consumers being aware of more ethical alternatives to existing products and services.

We would expect trust-marked services to perceive a competitive advantage based on improved user trust - a concept which is explored in a set of over 50 case studies, assembled in 2016 by Gary Hasselbalch and Pernille Tranberg.[7] Depending on the regulatory environment, organisations able to show compliance with the trust-mark criteria might also perceive some form of regulatory benefit (a "safe harbor", in US terms).

In terms of user behaviour, the trust-mark approach fits well with the transtheoretical model. The trust mark itself serves as the simple, binary signal (or perhaps a three-value 'traffic light' model) that gives the consumer an instant indication of a service provider's privacy stance. The underlying principles would give further information in support of the consumer's decision, and ultimately, an organisation's accreditation and audit status could be open to inspection. Over all, it is conceivable that the increased transparency associated with trust-marked products would increase pressure on competitors to be more explicit about their own business models, or risk losing trust because of the implied inferiority of their offering.

ONE REASON IT CAN BE HARD TO SENSITIZE PEOPLE TO PRIVACY RISK IS THAT PRIVACY-ERODING BEHAVIOUR OFTEN APPEARS TO HAVE LITTLE OR NO ADVERSE EFFECT AT THE TIME

Designers and vendors would have an incentive to respond to any general shift in the market, generated by adoption of trust-marked alternatives, by improving the privacy design of their offerings. A similar trust mark model could, we believe, also be applied to apps (a privacy score, linked to more information about the permissions the app requests, the data it collects, and any back-end processing), and to connected objects (a score linked to more information about what data the object collects/generates, where it sends it, and what processing is done in the 'cloud'). The Internet Society's Global Internet Report (GIR) for 2016 discusses the role played by trust marks and similar 'credible signals' in establishing and reinforcing service providers' credibility. The GIR does this in the context of the economics of online security, but we believe there are direct parallels with the trust and privacy case.[5]

However, this does still leave at least one gap in the picture, concerning the transtheoretical model: what interventions are possible for the 'maintenance' phase of the cycle? What positive reinforcement can we achieve for users who are on track to develop positive privacy habits? Could it be turned into a game, for example, much as the "Cheevos" plug-in does for privacy features in Mozilla's Firefox browser?[8]

There is an alternative form of reinforcement, based on principles I heard about from the Design Thinking labs at Stanford University. In their experience, based particularly on projects to do with food labelling and healthy eating, was that the most effective approach is to influence the values that users apply to the decisions they make. Let's take donuts as an example. If you present the choice simply as 'have a donut or don't have a donut', the chooser has an 'instant gratification' incentive to take a donut, and (as remarked earlier) probably no instant heart attack to persuade them otherwise. However, if you present the choice as 'have a donut, or have an apple and live a longer and healthier life', you change the values the chooser applies to the decision.

[7] G Hasselbalch and P Tranberg, Data Ethics – The New Competitive Advantage 24/09/2016

[8] "Cheevos" Firefox add-on website, https://addons.mozilla.org/en-Us/firefox/addon/cheevos/

The Stanford labs found that this approach is more likely to result in sustained behavioural change. Fortunately, this too can be fitted into the transtheoretical model, at the informational phases of the decision-making cycle. One reason it can be hard to sensitize people to privacy risk is that privacy-eroding behaviour often appears to have little or no adverse effect at the time, and this can result in a dangerously low assessment of the risk of continuing. By analogy, it's not the first donut that fatally clogs the arteries, so we might persist with this potentially damaging behaviour until it produces serious physical symptoms, at which point much of the damage may already have been done.

In practical terms, this means that when we make those 'informational' interventions in the decision-making cycle, we need to do so in ways that are directed more towards influencing the values the individual applies to the decision, and less towards the possible consequences of that single act. So, for instance, we might frame the decision in terms of its long-term effect on the individual's credit rating, or the intimacy of the profile the service provider can build. It may be important to find ways of showing the disparity between, say, the trivial nature of the service and the intimate nature of the behavioural profile it creates. Making such disparities visible to the user may help them make their decisions less in terms of immediate convenience, and more in terms of their personal, long-term values.

WHERE MARKET FORCES WILL PREDICTABLY FAIL, THERE IS A CASE TO BE MADE FOR REGULATORY INTERVENTION.