Reconciling Multiple Objectives – Politics or Markets?

Ross Anderson and Khaled Baqer

Computer Laboratory, University of Cambridge, UK
forename.lastname@cl.cam.ac.uk

Abstract. In this paper we argue that the evolution of protocols is one of the most important yet least understood aspects of the governance of information systems. At the deepest level, security protocols determine the power structure of a system: who can do what with whom. The development, adoption, spread and evolution of protocols, and competition between them, are both political and economic activities. They may reconcile multiple objectives or be the battlefield in which one interest defeats another. They occur at all levels in the online ecosystem, from individual and small-firm action, up through organisations and networks to whole ecosystems – and may eventually play a role in shaping culture, values and norms. They play a key role in innovation: early movers may use protocols to establish a strategic advantage and lock competitors out. How can we understand such complex behaviour? In this paper we sketch a possible framework inspired by research in institutional economics.

1 Introduction

Security protocols often support not just a community of users but a business community, and evolve with it. It should be obvious that if a protocol is designed by Alice, we might expect it to suit her strategic purposes. But this can make Bob unhappy, and the protocol eventually needs to be redesigned. A mature discussion of security needs to take account of how protocols evolve. However, we rapidly get out of the space of two-player games.

Previous workshops on security protocols have discussed several examples of protocols acquiring new features until eventually they broke, as a result of feature interaction or implementation complexity. The system designed by IBM and others to encrypt PINs in a bank's ATMs was extended by Visa and others to support worldwide networks of banks, ATMs and point-of-sale devices. That led to API attacks which arose from unanticipated feature interactions [9]. It was extended further to deal with EMV smartcards, leading to further vulnerabilities and attacks [14]. The same holds for SSL/TLS protocol suite, which is now more than 20 years old; it has accumulated options (such as export ciphersuites) and features (such as heartbeat) which have led to significant vulnerabilities. Because of the enormous number and diversity of users, fixes usually have to be applied at one end only; simultaneous upgrades to both clients and servers are hard [22].

This much is reasonably well known. At previous protocols workshops we have touched on economic models of protocols [5], asked whether in a democratic system there should be a 'loyal attacker', inspired by the 'loyal opposition' in parliament [7] and discussed crowdsourcing social trust [11]. In this paper we apply concepts from institutional economics to provide a framework for discussing the ways in which protocols evolve at different scales of organisation and time, and the effects on innovation.

The rest of the paper is organised as follows. Section 2 describes John Groenewegen's model of innovation in the electricity industry, and our proposed simplification of it for protocol analysis. Section 3 looks at protocol creep, bugfixes and tussles. Section 4 considers protocol complements and the scaling of innovation. Section 5 attempts to draw some conclusions – for the 'Internet of Things', for dispute resolution, and for the kinds of online conflict we have seen around recent elections on social media. Finally, we present our conclusions in section 6.

2 The Groenewegen model of innovation

The institutional economist John Groenewegen proposed a model of innovation in the electric power industry [16], which we copy in Figure 1.

The point of such models is that technology evolves in ways that depend on scale, on networks, and on history. Once you've invented the battery, you can perhaps electroplate spoons with silver. In theory you could make and sell torches, but in practice the electric lamp was not invented until the dynamo came along, and electricity got much cheaper. Given these technologies, any town that already has gas lighting should replace it with electric lighting, which is cheaper. Once several towns have electric power companies, it's worthwhile for them to build a grid and perhaps even merge to get economies of scale. The utilities then become natural monopolies, which brings in regulation.

At a workshop at Schloss Dagstuhl in November 2016 we discussed this with Johannes Bauer, Thomas Maillart, Barbara Kordy, Gabriele Lenzini and Sebastian Pape and simplified it to four levels. For present purposes, we simplify it still further as in Figure 2.

At the bottom, at level 4, are the actions of individuals or small firms; the person who invents a device, or writes a piece of software for their own use, and perhaps makes a few copies for friends. The time taken to do this can be very short, perhaps a matter of days to weeks.

At the next level up we start to industrialise, with the emergence of a substantial firm or network. This might be a physical network, such as the electric wiring installed to convert a town's gas light to electricity, or a virtual network, as where software runs on a particular type of machine. Both can support innovation. A power company that recovers its startup costs from a street lighting contract can then start selling electricity to homes, and once that's available it starts to make sense for people to invent appliances. An example of a virtual network is where firms started selling home computers in the 1980s, which led

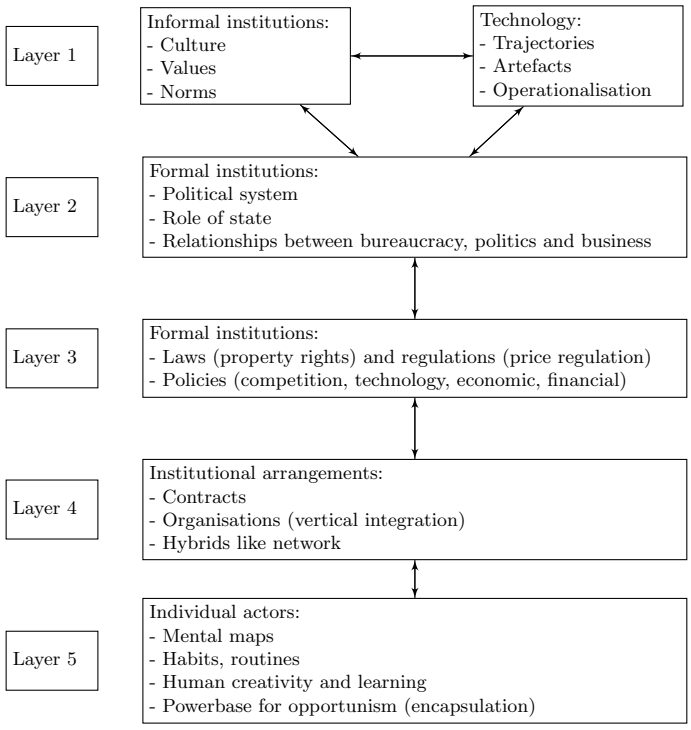

Fig. 1. Groenewegen's model of innovation, recreated from [16]

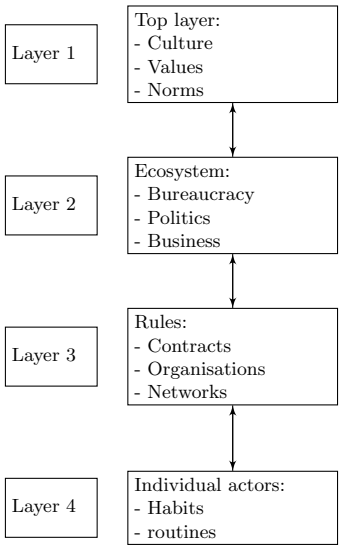

Fig. 2. Our innovation stack model

rapidly to the emergence of games companies that wrote software for them. The time constant for the emergence of a firm or network is typically 1–10 years.

The level above that is an environment or ecosystem, of which our starting example is a power grid plus its associated regulation, standards and markets. A grid may cover a section of a large country (as in the USA) or several smaller countries or regions (as in the EU, or the grid covering Ireland and Northern Ireland). Similarly, the Internet can be seen as an environment, while large service firms like Microsoft, Google and Facebook try hard to maintain their own ecosystems: their products work better together than with their competitors' offerings. The time constant to build an ecosystem is measured in decades rather than years.

Finally, the top level is one of culture, norms, religions, or (as Groenewegen puts it) social embeddedness. Culture is slow to change, but it does indeed change in response to changes in a civilisation's underlying technology package. Putnam documented, for example, how the arrival of television noticeably cut participation in club membership and group social activities, and no doubt the spread of smartphones [21], while Pinker ascribes the Enlightenment to the spread of the novel [20]. The rapid adoption of smartphones, messaging and social media over the past decade will no doubt have an impact too. In general, the time constant of cultural change is measured in centuries; however it does occasionally happen that a noticeable change occurs in the course of a human lifetime. One might argue that the Enlightenment and the late 20^{th} century rights revolution were such, and were at least in part facilitated by technological changes (printing and television respectively). Social conservatives may deplore one or even both of these developments; that such resistance can last for centuries merely underscores the slow pace of cultural change.

This 'innovation stack' seems an attractive way to describe and analyse innovation in fields other than electric power, as it reflects how technology entrepreneurs actually think. In the 1970s and 1980s, a small software company that had written a system to automate a sawmill (say) would then try to parametrise it and sell it to all sawmills; having done a project, you work to build it into a firm. From there, the next aspiration is to become a platform. The success of IBM, DEC and later Microsoft emphasised the value of being an ecosystem used by many industries rather than just one; during the dotcom boom, one startup after another would tell its investors it was going to be the Microsoft in some industry sector or another. Those who succeeded, such as Amazon, became household names, and this approach to business is discussed in detail by Shapiro and Varian's book "Information Rules" [23], which became the top business book in 1998. Now the largest firms such as Amazon, Apple, Google and Facebook attempt to embed themselves in culture, by not just building ecosystems around their core products but by shaping human behaviour around them and channeling it through them. At each step, ambitious firms try to move up to the next level in the stack; few succeed; but those who do reap tremendous rewards.

Can this framework give useful perspectives on how protocols evolve?

3 Protocol evolution

We first consider three forms of protocol evolution: growth, bugfixes and tussles.

3.1 What successful evolution looks like

Our first example is the world card payment network, which can be seen as emerging from the local innovation of Diners' Club in New York in 1950; both credit and debit cards were first offered by single banks to local customers, then networked together first nationally and then internationally. The protocol history really starts with the first cash machines fielded by Barclays in 1967. Early ATMs used ad-hoc mechanisms for PIN security but rapidly converged on a methodology pioneered by IBM. A single bank could give each customer a PIN generated by encrypting their account number with a secret PIN Derivation Key, available in all the bank's ATMs and also to its central mainframe.

This was then adapted to provide dual control, so that no member of bank staff would ever see any PIN other than their own, with a view to both minimising insider fraud and defending the bank against claims from customers who had been defrauded (or pretended to be). The mechanism was to embed the cryptographic keys in tamper-resistant hardware, both in the ATM and in a Hardware Security Module (HSM) at the bank. The cards and PINs sent to customers could be sent from different data centres, or from the same centre on different days; and likewise the pairs of keys used to initialise ATMs [6].

The natural next step was to extend the communications to support interbank networking, a project driven by Visa and other card brands in the early 1980s. Further features were added, such as allowing customers to choose their own PINs. This allowed convenient worldwide ATM networking, and in due course a serpent appeared in the garden: things had become so complex that attacks started to emerge based on careless design [4] and feature interaction [14]. No sooner had the worst of these attacks been mitigated but the banks rolled out EMV, adding a further layer of complexity; and the inevitable attacks were found in due course [24].

The card payment system evolved from being a single-bank system (level 3) to being an ecosystem (level 2), and assurance failed because of a combination of bloat, feature interaction and institutional economics. As noted in [18], the implementation of EMV in Britain involves thousands of pages of specifications, as the cost of building the ecosystem was accepting backwards compatibility with a lot of legacy systems. The combination of features that led, for example, to the No-PIN attack is scattered around this huge corpus. The bank brands building the ecosystem had insufficient incentive to insist on formal verification or even adversarial review. The member banks also have conflicts of interest internally [17]: an acquiring bank that uses a less secure PIN entry device increases the risk of fraud across all local card issuers, yet it's unlikely to face more than a few percent or so of this fraud itself. Ultimately, the limits to security are about governance.

3.2 Dealing with protocol failure

At Indocrypt 2011, Rescorla asked why the Internet runs on pre-historic cryptography [22]. His point was that while the basic TLS protocol may be secure in theory (it was verified by Paulson in the 1990s [19]), it has been broken repeatedly in practice because of features that were not verified (such as error handling), emergent properties of implementations (such as timing), forced parameter choices (such as export ciphersuites) and add-ons (such as session resumption – to which we would nowadays add heartbeats). Yet it has proved almost impossible in practice to fix these flaws properly. A robust solution would in most cases involve changes at both ends at once, which is impractical. So long as a few percent of online shoppers still use IE version 6, almost no merchant will be prepared to move entirely to modern ciphersuites that would abandon these customers and their business.

The combination of scale and network effects means that it is cheap to tweak the code, but extremely expensive to change the architecture. TLS 1.0 dates to the 1990s; versions 1.1 and 1.2 were minor tweaks; the more significant upgrade to 1.3 is still a work in progress. Similarly, the move from magnetic-strip-based ATMs and POS transactions to the smart-based EMV involved pilots in the early 1990s, standardisation work in the late 1990s and finally a roll-out in 2002 – which is still not complete. This roll-out has involved substantial local customisation, and incremental changes are still being made to support new services and patch up vulnerabilities that get exploited. Even there, it can be slow; it took UK banks over 5 years to block the No-PIN attack [18]. In biological terms, we can see protocols as a case of punctuated equilibrium rather than gradualism. Cladogenesis – the creation of new species – is not however down to a near-extinction natural event, but to richer forms of competition and conflict.

3.3 Dealing with conflict

Many of the interesting cases arise when an incumbent protocol is challenged. In 2012 we wrote of how the overlay payment service Sofortüberweisung (Sofort) challenged the incumbent giro payment system operated by banks in Germany and Austria [7]. Germany in particular has many small banks who did not operate very competitive payment service; Sofort built a system that would enable account holders to log into their bank accounts via its systems and make easy, rapid payments for online purchases. In effect, it was performing a Man-in-The-Middle (MiTM) attack on the German banking system, and providing a better, cheaper payment service that pleased both customers and merchants. The banks sued, accusing Sofort of inducing their customers to break the terms and conditions under which their accounts were operated. The competition authorities intervened, advising the court that the competition brought by Sofort was socially beneficial. The case was stayed and Sofort duly acquired a banking license.

In that case, the effect of a challenge to an established ecosystem was that the challenger established itself as a player and slightly improved or extended the ecosystem's services. The same pattern can be found in a number of other cases:

– During the dotcom boom, a number of new payment service firms sprung up to challenge the credit card industry, which is dominated by the Visa-Mastercard duopoly and thus somewhat slow-moving; the winner, PayPal, has in effect become a leading credit card acquirer, although it also has its own payment mechanism for personal account holders to use.
– The establishment by Android and Apple of phone payment mechanisms from 2011 was also intended as a challenge to the credit card duopoly[1] but did not really take off until after contactless credit cards became mainstream, leading retailers to upgrade their terminal fleets; it is now, like PayPal, a minor add-on to the credit card ecosystem.

There have also been cases where the absence of a viable incumbent enabled a protocol to establish a new ecosystem rather than just slightly extending an existing one:

– About 200 mobile-phone payment systems have been set up in various less developed countries, and about 20 of them have become mainstream[2]. In such protocols, the payer sends an encrypted SMS message or USSD command to a payment server ordering it pay a certain sum to a given phone number; the payee gets a message certifying receipt. The killer app was migrant remittances in an environment where most people had no access to conventional banking services but where mobile phones were spreading rapidly. The best-known is Kenya's M-Pesa whose operator Safaricom processes payments amounting to a substantial proportion of Kenya's GDP[3], becoming a larger payment service provider than any local bank, and has also become the monopoly Mobile Network Operator. In this case, regulation can be used after a monopoly establishes its presence to thwart attempts by newcomers to disturb the innovation stack and undermine the monopolist's control: the monopoly (Safaricom) uses regulation to accuse innovators of undermining its service, and thereby stalling the innovators' efforts in lengthy legal battles[4]. Moreover, attempts over the past ten years to set up similar phone payment systems in developed countries (such as Pingit and PayM in Britain) have generated little traction, as people in such countries already have plenty ways of paying both online and offline.
– The SSL/TLS protocol was different. It was the outcome of a challenge by Netscape in 1995 to a protocol suite, SET, that was getting too complex

[1] Full disclosure: the first author assisted with the design of Android Pay while a scientific visitor at Google.

[2] See GSMA's 'Mobile for Development' website (using the *Mobile Money* filter): `http://www.gsma.com/mobilefordevelopment/tracker`

[3] See mobile payment statistics at the Central Bank of Kenya: `https://www.centralbank.go.ke`

[4] For example, see Safaricom's battle with Equity bank over the bank's use of overlay SIMs to circumvent the restrictions on phones (`https://www.standardmedia.co.ke/business/article/2000135850/safaricom-loses-battle-to-block-equity-bank-s-thin-sim-card`), and Safaricom's battle with BitPesa (`https://www.bitpesa.co/blog/bitpesa-v-safaricom/`)

and difficult to implement. SET was a large collaboration between industry players (including Netscape and Microsoft) and the banking industry (Visa and Mastercard) to design a standard online credit card payment system. However, there were too many players, the specification got too complex, and implementing it would have required every bank to certify the public keys not just of its merchants but all its cardholders. It became clear that many banks would take years to even get online; large-scale certificate issuance was going to leave many potential customers stranded. SSL gave a quick and dirty way to get card transactions online, with only the merchants having to buy certificates – which they could do from third-party CAs.

In the case of SSL/TLS, the cost of not authenticating both ends of the transaction was phishing; but that only got going until 2005, ten years after SSL was first launched. SSL/TLS does however requires certification of the public keys of servers at least. This brings us to the topic of how the evolution of a protocol can be affected by externalities, such as required complementary services.

4 Protocol complements and scaling

Public-key cryptography started off as a monopoly protected by the Diffie-Hellman and RSA patents; the operating company, RSA Data Security, issued certificates as a service to customers. This was spun off in 1995 as Verisign, and as the patents reached their end of life, this company became the next big money-spinner [3]. During the dotcom boom, it established a dominant position by purchasing rivals such as Thawte. However various nation states and others pressured Microsoft to put their root keys into the browser, so as to facilitate interception, leading to several hundred organisations becoming 'trusted'.

Iran was excluded, and hacked DigiNotar so as to be able to run MiTM attacks on dissidents. This led to the industry, led by Google, excluding Diginotar and thus terminating its business. The resulting disruption, which included multiple Dutch government departments not being able to offer their usual online services until they had acquired new certificates from elsewhere, was a shock to the industry. The result was a 'flight to quality' with most extended-validation certificates now issued by Verisign (now owned by Symantec), with Comodo the next largest. These CAs are considered "too big to fail" despite the fact that Verisign also suffered a data breach in 2010 and the Iranians also got Comodo certificates at one point [10].

This is not the first time the fight has been over complements to an ecosystem; IBM used its patents on punched cards as a key tool in its fight against Remington Rand and others for dominance of the tabulating machine market in the 1930s [2]. Shapiro and Varian discuss other cases of complements being used for leverage in battles to dominate markets with network externalities [23].

A key question however is scope. For example, the 'Internet of Things' is shaping up to be not one ecosystem but many, in each of which the objects sold by a vendor communicate with its cloud services and through them to a user's

phone. There is no real reason for the security protocols to be standard TLS, except perhaps in that the phone app may embed a lot of ad networks that need to rely on standards [8].

The lack of standards can cause real problems, such as in the failure to provide for the home area network (HAN) in the smart meter ecosystem promoted in the EU [1]. Different Member States adopted different standards for communication with appliances in the home, with none of them gaining enough traction to appeal to appliance vendors. As a result, it looks like the Google Nest standards may prevail. The EU regulators should have standardised local interconnectivity first if they wanted meters and appliances to communicate directly rather than via third-party service firms. In this case too, the scope of complements had a critical effect on scaling.

A further computer-science approach is to think in terms of centralisation versus distribution. In practical terms, the world of IT has spent 50 years centralising things and putting them online. But it does have limits, and these are often as much political as technical; data localisation laws, local telecoms regulation and bank supervision are three salient examples. Last year we talked about DigiTally[5], where an overlay network can be used to do an end-run round an incumbent payment system that's entrenched thanks to a mobile network monopoly; there too, however, there is a centralised issuer. In this particular case the multiple protocols (GSM for basic communications and then SMS or overlay for two competing or complementary payment systems) can compete in the normal commercial and political arenas. There we learned that one should always ask: where is the choke point? The incumbent's control of the agent network is one of the foundations of its monopoly, along with its control of the API and the dominant telco (that the government is a shareholder also doesn't harm them in this context). However, this is not always purely technological. For mobile payments in Kenya it is the 40,000 agents who turn banknotes into bits. The DigiTally case study reminds us that this is likely to come down to the sort of economics that competition policy people worry about, namely how institutions work in the real world, and the tussles between them.

5 New directions

The established literature on information economics already notes the importance of innovation in the success of a platform [23]. Windows (and before it Unix and the IBM ecosystem) succeeded by providing a platform on which millions of programmers could innovate. When people bought the platform, they gained access to the products of other innovators in the same ecosystem. The positive feedback thus created led to platform success. The smart meter HAN failed because people could not innovate; the default for an appliance maker is

[5] See our previous SPW paper [12], CCS 2016 keynote (`https://www.cl.cam.ac.uk/~kabhb2/DigiTally/docs/ccs-vienna-2016.pdf`), and the DigiTally project: (`https://www.cl.cam.ac.uk/~kabhb2/DigiTally/`)

to enable their equipment to be controlled remotely by the customer's phone via their own server.

Innovation is not just a matter of software; people also innovate by setting up websites to support their hobbies and interests, or their social life. The effect of scale here is less studied. However, we might estimate that while millions of programmers benefited from the Windows platform as a space in which they could innovate, the world-wide web increased this by an order of magnitude. It enabled tens of millions of people to create websites by the mid-2000s; writing HTML is easier than writing programs.

Since then, social media has increased the scale by two further orders of magnitude. Facebook currently claims 1.9 billion active users, while networks in China and Russia have hundreds of millions more. All you need to do is upload a picture, click on a few friends and you're away. Part of Facebook's secret sauce was minimising the amount of expertise and effort required to join.

Once just about everyone's online, almost everything bubbles up, from cat pictures to hate speech. In fact, Grossman notes that much of the unpleasantness online is just a democratisation of the unpleasantness previously seen from tabloid newspapers and even state propaganda organisations: xenophobia, smears, doxxing and other personal attacks [15].

At this point, the protocol researcher might ask, 'What's missing?'

Our 'starter for ten' is recourse. Pinker describes how human conflict has dropped in stages from a high level in pre-state societies to the low levels found today in most societies, and that the largest single fall was associated with the emergence of the state.

The king provided justice, both as redress post facto, and as laws governing such matters as land and marriage to minimise the number of disputes arising in the first place. Gambetta's studies of organised crime show that where justice is not provided effectively by the state, there will be a temptation for private actors to do it – and we do indeed find various kinds of Mafia online, such as Dread Pirate Roberts and his Silk Road. Other options include industries able to provide recourse as a service, such as the credit card system.

But a purely technical solution cannot be enough. If the purpose of the state is to resolve disputes, it's not enough to have an army to monopolise military force in the jurisdiction and a college of judges to rule on individual cases. A state is more than that. States historically were tied up with culture and religion; they were as socially embedded as could be. From the elaborate negotiations used by pre-state tribes to end feuds through the gruesome public executions of early kings, justice evolved in society with splendid ritual: the English travelling judge in his wig and robes was the emissary of the King's Peace, and the jury men he swore to try cases embedded the verdicts in the local community (or at least in its elite). Law became democratic as society did; Blackstone describes its evolution as a 'long march from status to contract' [13]. This works at many levels; at the top level, of course, we elect our legislators, whose job it is to deal with gaps that emerge between laws and social norms. As a result, most of the people internalise the important rules most of the time.

The dispute resolution system must be one that people actually believe in. This belief ultimately has its roots in lived experience. The classic security engineering view that 'a trusted system is one that can break my security policy' [6] may be a good heuristic for reasoning about mechanism, but it is not enough to reason about how systems are embedded in society, especially when people feel anxious about globalisation, frustrated by call centres and infuriated when 'computer says no'. Perhaps it's not enough to assume we have an impartial "judge", who combines competence at cryptologic mathematics with effective competence in the real world. Perhaps we have to start asking "To whom am I really proving this?" Perhaps we have to start thinking about crowdsourced mediation or arbitration: in short, about juries.

6 Conclusions

Shapiro and Varian's classic "Information Rules" [23] is almost twenty years old but still has much to say about how security protocols succeed or fail. We have discussed a number of examples familiar to the protocols community, notably payment protocols and SSL/TLS. There are interesting issues around peer competition, and hacking everything from insulin pumps to tractors.

What we've tried to do here is organise the accumulated experience from 25 years of protocol workshops in an institutional economics framework.

Lessons learned include that if you want something to succeed you'd better think hard about how it will be open to innovation by others. Alternatively, as we've seen in mobile payment systems used in less developed countries, regulation can be used after a monopoly establishes its presence to thwart attempts by newcomers to disturb the innovation stack and undermine the monopolist's control.

Who are the innovators – programmers, micro-entrepreneurs from bloggers to small firms, or everybody? Who will resolve disputes, and how? The framework we've proposed suggests that it's not enough to have a resolution mechanism that can be verified using a logic of belief. It must be a mechanism in which people actually believe. This is a topic that protocol researchers have barely touched; the closest attempts so far may be on human auditability of digital elections. Yet even those require people to trust mathematicians and other experts. The events of 2016 suggest that expert opinion is not always enough. If protocols are the tools of power in a digital age, how can they win genuine acceptance? The crowdsourced reputation systems operated by firms such as Tripadvisor and eBay may indicate a possible direction of travel. If people feel they don't have a voice any more, to the point that using customer feedback properly can give a travel agency or auction house a real advantage, then in what circumstances might it be possible to extend this to dispute resolution?

Acknowledgements

We are grateful to Johannes Bauer, Thomas Maillart, Barbara Kordy, Gabriele Lenzini and Sebastian Pape for the discussions at Schloss Dagstuhl referred to in section 2, and to workshop participants for the discussions there.

References

1. Directive 2009/28/EC of the European Parliament and of the Council of 23 April 2009 on the promotion of the use of energy from renewable sources and amending and subsequently repealing Directives 2001/77/EC and 2003/30/EC. *European Parliament*, 2009.
2. The IBM punched card. *IBM Inc*, 2017.
3. Verisign. *Wikipedia*, 2017.
4. R. Anderson. Why cryptosystems fail. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, pages 215–227. ACM, 1993.
5. R. Anderson. The initial costs and maintenance costs of protocols. In *International Workshop on Security Protocols*, pages 336–343. Springer, 2005.
6. R. Anderson. *Security engineering*. John Wiley & Sons, 2008.
7. R. Anderson. Protocol governance: the elite, or the mob? In *International Workshop on Security Protocols*, pages 145–145. Springer, 2012.
8. R. Anderson. Making security sustainable. *Communications of the ACM*, forthcoming.
9. R. Anderson, M. Bond, J. Clulow, and S. Skorobogatov. Cryptographic processors-a survey. *Proceedings of the IEEE*, 94(2):357–369, 2006.
10. A. Arnbak, H. Asghari, M. Van Eeten, and N. Van Eijk. Security collapse in the https market. *Communications of the ACM*, 57(10):47–55, 2014.
11. K. Baqer and R. Anderson. Do you believe in Tinker Bell? The social externalities of trust. In *International Workshop on Security Protocols*, pages 224–236. Springer, 2015.
12. K. Baqer, J. Bezuidenhoudt, R. Anderson, and M. Kuhn. SMAPs: Short Message Authentication Protocols. *International Workshop on Security Protocols*, 2016.
13. W. Blackstone. *Commentaries on the Laws of England*, volume 2. Collins & Hannay, 1830.
14. M. Bond and R. Anderson. API-level attacks on embedded systems. *Computer*, 34(10):67–75, 2001.
15. W. Grossman. Don't die interview. *www.nodontdie.com*, Oct 21 2016.
16. R. W. Künneke, J. Groenewegen, and J.-F. Auger. *The governance of network industries: institutions, technology and policy in reregulated infrastructures*. Edward Elgar Publishing, 2009.
17. S. J. Murdoch, M. Bond, and R. Anderson. How certification systems fail: Lessons from the ware report. *IEEE Security & Privacy*, 10(6):40–44, 2012.
18. S. J. Murdoch, S. Drimer, R. Anderson, and M. Bond. Chip and PIN is broken. In *Security and Privacy (SP), 2010 IEEE Symposium on*, pages 433–446. IEEE, 2010.
19. L. C. Paulson. Inductive analysis of the Internet protocol TLS. *ACM Transactions on Information and System Security (TISSEC)*, 2(3):332–351, 1999.
20. S. Pinker. *The better angels of our nature: The decline of violence in history and its causes*. Penguin UK, 2011.

21. R. D. Putnam. Bowling alone: America's declining social capital. *Journal of democracy*, 6(1):65–78, 1995.
22. E. Rescorla. Stone knives and bear skins: Why does the internet still run on pre-historic cryptography. *INDOCRYPT (Invited talk)*, 2006.
23. C. Shapiro and H. R. Varian. *Information rules: a strategic guide to the network economy*. Harvard Business Press, 1998.
24. P. Youn, B. Adida, M. Bond, J. Clulow, J. Herzog, A. Lin, R. L. Rivest, and R. Anderson. Robbing the bank with a theorem prover. Technical report, University of Cambridge, Computer Laboratory, 2005.