JOURNAL OF CULTURAL ECONOMY
https://doi.org/10.1080/17530350.2022.2098515

Routledge
Taylor & Francis Group

OPEN ACCESS Check for updates

# Characterising assurance: scepticism and mistrust in cyber security

Matt Spencer

Centre for Interdisciplinary Methodologies, University of Warwick, Coventry, UK

**ABSTRACT**

This paper presents an analysis of recent transformations in cyber security assurance, a field of evaluation that aims to establish whether technical products are secure. I work from a set of narratives about problems with assurance, drawn from interviews with practitioners based in the UK. I focus on characterisation: the stories practitioners tell, the cast of characters that populate them, and how such stories act to problematise the domain. Mistrust, it is argued, can be understood in terms of the capacities of sceptical narratives to efface the power of security certifications to be taken on 'face value.' A text-based view of mistrust is thus developed that can be differentiated from the conventional disposition-centred view. Examining mistrust, then, leads us to ask not how to change dispositions to make them 'more trusting,' but rather to critical questions about the palette of characters that feature in cyber security. I close the essay by offering a commentary on the way characterisation leads to the anticipation of experts in formulations of policy and on the possible 'counter-characterisation' that might be developed, for instance around 'caring' characters.

**ARTICLE HISTORY**
Received 22 October 2021
Accepted 20 June 2022

**KEYWORDS**
Cyber security; assurance; characterisation; mistrust; narrative

## Introduction

The desire for confidence that digital infrastructures are secure is acute, and increasingly so following high profile cyber-attacks (NotPetya, Wannacry, the Solarwinds breach, the US Colonial Pipeline incident, and the Hafnium Microsoft Exchange hack, to name a few). Technology assurance schemes aim to provide formal certification that digital products used in such infrastructures can be relied on to stand up against attempts to subvert them, but despite a 40-year history, the endea- vour of assurance remains today contested, imperfect and partial. Even among those who make and use formally assured technology, doubts about the worth of certifications abound.

When you ask cyber security practitioners what the problem is with technology assurance, they will often answer with a *story*: a story about what goes wrong, how assurance gets manipulated, or subverted, or misinterpreted. It may be a story told in abstraction, about generic characters and their relationships, or an account of a particular exemplary incident; in each case, narrative provides a form through which problems can be articulated, reiterated and addressed, rendering intelligible the complexity of securing.

It is hard to over-emphasise the depth and complexity of the socio-material context in which technology assurance takes place. It involves industries of product developers, designing and building (for example) servers, switches, firewalls, encryption devices, anti-virus, network monitoring,

**CONTACT** Matt Spencer m.spencer.1@warwick.ac.uk

operating systems (and so on). Some of which are generic, some specialised, and all with their own supply chains—including bespoke, commercial and open source software and firmware, produced through design procedures, test procedures, infrastructures for product development and management of versions and source code, all constructed and negotiated in relation to design constraints deriving from protocols and standards for interfaces and hardware components. Where formally assured, such products are assessed by evaluators using specialist and commodity testing technologies, held to standards for the evaluation of the security of products and for the audit of development processes, working under systems for the appraisal and recognition of evaluations, formats and channels through which certifications can be disseminated, not to mention the myriad sites of information technology infrastructure into which such assured components may be integrated, sites near and far of adversarial activity, and the wider ecology of institutions producing and circulating knowledge about threats and threat actors to which cyber security measures are expected to respond.

This context is, furthermore, highly dynamic. Technologies and threats are changing, and so is assurance (Spencer 2021). I draw my examples from a set of interviews conducted in a moment of upheaval in the UK's technology assurance policy landscape. Established schemes closed down, and a new 'principles based assurance' approach was piloted by the National Cyber Security Centre (NCSC). Practitioners knew that things needed to change and were changing, but were uncertain about the future. It is not surprising, given this background of change and complexity, that narrative plays a key role, populating a 'storyworld' with just those specific characters and (inter)actions relevant to the transformation.

This paper is about characterisation, and I use the term broadly, playing on the generic sense of 'characterising assurance' as the process of grasping this complex domain through narrative, as well as the characterisation that occurs within such narratives, which is often the crux of the issue: for instance, how unscrupulous parties may subvert the process of assuring technology, or how evaluators may be manipulated. Whether specific or generic, such characters rarely have much depth or nuance, appearing as simple 'straw men.' But addressing these characters *in* the text leads me to a third sense of characterisation, as what the text itself does, the performative capacity of texts that amplify the relevance of certain figures to the problems of the domain. Stories about problems with assurance, I suggest, enact mistrust; where mistrust is to be understood as a form of textual agency (Cooren 2004). Characterisation, then, does not just describe assurance, it intervenes: it makes assurance problematic (by complicating interpretations) and it shapes how reforms to assurance are imagined.

I begin by setting out the argument for adopting this rather text-centric view of mistrust, relating characterisation to classic work in science and technology studies. In the second section, I provide some historical background on cyber security assurance in the UK. I then move on to the main part of the paper: an analysis of characterisation and mistrust in a set of six excerpts from interviews with cyber security practitioners. I return to the theoretical discussion at the end, tying together the argument with broader approaches to trust, and examining the implications for a critical cyber security.

## Mistrust and characterisation

Character has been a central theme in scholarship on the production of objective knowledge. Steven Shapin, for instance, argues that in 17th century Europe, the integrity and character of the scientist featured prominently in epistemological discourse – but from the 18th century onward, this 'people-knowledge' was increasingly rendered invisible. Authority and testimony remained pivotal to the practice of science, but scientific knowledge was increasingly represented as if it were a kind of 'object-knowledge' free of reference to persons (Shapin 1994).

Characterisation, the process of creating character in a text (Margolin 1986, Culpeper 2014) is an effect of narrative that stands in tension with object-knowledge. More than indicating or naming a character, characterisation 'invests [them] with an attribute or set of attributes' (Garvey 1978, p. 63).

The certificates produced by cyber security assurance schemes declare an evaluated product secure, and their ability to do so as ’object-knowledge’ rests upon a format and style of text that minimises investment in the attributes of the character, for instance, of the evaluator herself, about whom minimal information is provided, aside from their formal role in the certification scheme.

But characterisation is not easily suppressed. The philosopher Mikhail Bakhtin wrote powerfully on the dialogical nature of life and discourse and their constitution in relation to the historical situation. For Bakhtin, texts that speak in a singular monological voice do so only through the suppression of dialogical, polyphonous tendencies – the multiplicity of voices that otherwise may be heard (1984).

> All words have the “taste” of a profession, a genre, a tendency, a party, a particular work, a particular person, a generation, an age group, the day and hour. Each word tastes of the context and contexts in which it has lived its socially charged life; all words and forms are populated by intentions (quoted in Cooren and Sandler 2014, p. 225)

What is particularly pertinent here is the potential for dialogical evocation that accompanies any factual assertion: ‘Who says?’ and ‘In whose interests is this fact?’

Sociological studies have long emphasised the contingency of semiotic conditions associated with monological discourse. In their study of the Salk Institute, Bruno Latour and Steve Woolgar described the work of the laboratory as oriented toward counteracting ‘pressures to submerge assertions in modalities such that they become artefacts,’ such as reference to the individuals involved in their construction (1986 [1979], p. 81; Fahnestock 1998). Susan Leigh Star and Karen Ruhleder documented an analogous phenomenon, running in a contrary direction: directives concerning the configuration of new infrastructures that seemed simple face-value instructions in their context of origination, picked up interpretive complications that transformed their performative capacities when they travelled across contexts, coming to signify, for instance, differential capacities of interpretation, differences in resources, and lack of cross-contextual understanding, thus hindering the process of infrastructuring (Star and Ruhleder 1996).

Where Latour and Woolger’s emphasis on the genesis of facts leads to an image of linear process in which modalities are shed, Star and Ruhleder point to how complications can be accrued through their example of crossing contexts, and in related work through ‘inversion’ (Bowker 1995). What interests me is the capacity of stories to open up complications that would otherwise have been black boxed (Latour 1996, p. 233). This requires recognising the *textual agency* of the stories people tell. François Cooren has made important contributions here, integrating the analysis of narrative with material semiotics. Drawing on Greimas and Derrida, as well as Latour and Callon, Cooren pushes speech act theory beyond Searle’s intentionalist philosophy (2000). Where Austin and Searle regarded the ‘act’ of speech acts as an act of the speaker, Cooren suggested we must be attentive instead to the potential of the utterance or text *itself* to act.

A very simple way in which stories about assurance act is in creating the conditions for their own re-iteration in future discourse, in other words their ‘entextualisation’ beyond the transitory moment of utterance (Cooren 2004, p. 389). Entextualisation can be achieved by writing things down (for instance when I transcribe an interview). But it is also broader than this, associated with ‘the means available … to render stretches of discourse discontinuous with their discursive surround’ (Bauman and Briggs 1990, p. 73–74). Some of the stories examined here, for instance, do this via ‘meta-narrative,’ presenting themselves, for instance, as reported and retold narratives (Bauman 1986).

The form of textual agency most pertinent here is the ability of narratives to enact mistrust. By this I mean that sceptical narratives intervene in the capacity of assurance’s ‘primary texts’ (such as the certificates given to products that passed an evaluation) to speak in a monological, authoritative voice about the security of technology. Stories that characterise assurance do not just describe it; they act within it, amplifying some voices over others. The problem of companies ‘gaming’ the system, for instance, is not just a matter of what those companies do; it is also a problem through the efficacy of knowledge about the issue, and how the propagation of stories informs the ways in which the objectivity of assurance is called into question.

Many contemporary scholars regard mistrust as a kind of personal disposition and it should be fairly clear that I am departing from this. Florian Mühlfried, for instance, writes of mistrust as an attitude of engagement or way of relating to the world (Mühlfried 2018, 2019; *see also* Breakwell 2020). Matthew Carey writes of mistrust as a disposition, though he also sees the stakes as importantly residing within communication (Carey 2017). Carey tells us that in the Moroccan High Atlas, 'certain classes of statement are systematically presented and understood as inherently unreliable, and the default listener position is thus one of mistrust' (p. 30). Character, Carey tells us, is not the topic of constant discussion that it is in the Western world. Indeed, he suggests that '[o]ne could not make durable claims about [specific people's] character' (p. 31).

While this 'meta' characterisation of persons as generically inscrutable leads Carey to the analysis of general 'classes of statement' and a socialised 'default listener position,' I would argue that the study of mistrust needs to be construed in broader terms, so that it is attentive to circumstances that have not settled into 'default positions,' in which mistrust is not necessarily socialised as a stable set of attitudes. If people are disposed to be sceptical of the claims of assurance schemes, that is not because this is how British people think, or how British cyber security professionals think; it is because they inhabit a socio-material environment populated, and framed (in the sense of Latour 1996), by texts, stories, rumours (and so on), and it is always within such an environment that interaction occurs. In essence, if there is an agent 'doing' the mistrust, I suggest, it is not a human interpreter imposing a schema, but rather the narratives themselves in their intertextual efficacy: the power of sceptical discourse to act upon what it is that other texts are capable of.

Mistrust understood in this way is not just about the conditions for the evaluation of a claim. It also 'acts' to define the contours of the problem and shape how future interventions are imagined. In his study of the domestication of Scallops in St. Brieuc Bay, Michel Callon emphasised the role of texts produced by researchers in constituting a movement of 'problematisation,' where this is understood as not the formulation of a question that needs an answer, but rather as a determination of who and what are the relevant actors and how their interests intersect (1986). The narratives we deal with here are more fragmentary, telling plural and partial stories, but likewise mobilise a cast of characters with distinctive roles and relations, a kind of 'moral universe', which sets the terms in which further policy interventions are formulated (Woolgar and Lezaun 2013, p. 331). Hence, we can see the development by the UK's National Cyber Security Centre (NCSC) of a new 'principles based assurance,' conceived explicitly in response to the kinds of problematic characters that assurance has involved in the past, and in anticipation of a better cast of characters for the future.

The moral universe articulated in narratives about assurance creates considerable friction where characters in these narratives intersect with personal identities. Characters are not just components of a narrative structure, but also have referential capacity (the literary theorist Alex Woloch suggests it is necessary to see character as a 'distributed field of attention' spanning semiotic structure and reference; Woloch 2004, p. 17). The problematisation of assurance can very easily get personal when people identify with problematic characters, or when they struggle to identify with the kind of character they seemingly ought to be. Here, the study of characterisation intersects with wider studies of personhood and subjectivity (Fiske 1992, Alexander 2011, Moor and Lury 2018, du Gay *et al.* 2019). Addressing this intersection does not, however, create a boundary where we have to shift to a psychological frame: as we will see, the politics of characterisation become the topic of further narratives about problems with reforms to assurance.

The six excerpts I work from are drawn from a series of 30 interviews conducted with a variety of practitioners (all based in the UK and working across risk management, cyber security policy, test labs, and product engineering) in the wake of recent changes to the UK's assurance landscape. Before we turn to these, however, I provide some brief background on cyber security assurance in the UK.

## Assurance in cyber security

The study of cyber security assurance schemes takes us to an interface between digital infrastructure and the intelligence services that has received little scholarly attention in the social sciences. Government Communications Head Quarters (GCHQ) has featured extensively in social scientific studies of data-driven surveillance, and of controversies such as that surrounding GCHQ's involvement in the NSA's PRISM programme (for instance, Aradau and Blanke 2015, Amoore and Piotukh 2015). But the emphasis has been on GCHQ's signals intelligence role. Assurance remains obscure, for it emerged from the sibling traditions of secure communications and secure computing.

In the 1960s, secure communications was, according to Richard Aldrich 'an unknown fourth British secret service' (alongside MI5, SIS and GCHQ; Aldrich 2013, p. 178). Having initially existed as the London Communications Security Agency, by the 1970s these activities were brought into GCHQ as CESG, the Communications-Electronics Security Group. While CESG has its roots in the security of voice and text communication, since the 1980s it gained responsibility for the UK's participation in, and design of, formal assurance schemes for secure computing. More recently, in 2016, CESG was one of the principal components from which the new National Cyber Security Centre (NCSC) was assembled – the part of GCHQ responsible for the UK's assurance policy today.

Assurance has been a core concern for computer security since its beginnings. Early discussions in the late 1960s and 1970s centred around the problem of access control in resource-sharing computing environments. James Anderson, in the pivotal 'Computer Security Technology Planning Study' of 1972, stipulated that the mechanism implementing access control logics 'must be small enough to be subject to analysis and tests, the completeness of which can be assured' (1972, p. 9). Debates over the nature of this analysis and testing culminated in the 1983 publication of the US's Trusted Computer System Evaluation Criteria (TCSEC), or 'Orange Book' as it is often known, which was the first formal assurance scheme (*see* Mackenzie 2004).

While the US established TCSEC, CESG was involved in the development of a rival scheme, ITSEC. Aspirations to create a common certification scheme that could be adopted internationally gathered pace during the 1990s. The result was the Common Criteria for Information Technology Security Evaluation (commonly 'Common Criteria' or 'CC'). Common Criteria became international standard ISO 15408 in 1999. Like ITSEC, CC evaluations are carried out by commercial evaluation facilities, private labs which in the UK operated under license from CESG.

Where TCSEC had defined a single set of expectations for the security of computer systems, CC was designed to respond to the expansion in form factors of computing that was gathering pace by the 1990s. It provided a flexible grammar for the articulation of expectations, enabling the production of many different 'Protection Profiles' specific to different kinds of products, systems or devices. A 'Security Target' is then defined for each evaluation, setting out how the product to be evaluated is to be understood in relation to the protection profile.

Although Common Criteria has seen success in terms of the breadth of its adoption (17 certificate authorising and 14 certificate consuming countries at present), doubts about its ability to prevent security problems from arising have been prevalent from the start. It was unclear, for instance, that certified products were in fact less vulnerable than uncertified ones (Lipner 2015, p. 29). Such doubts, as well as a perception that CC was overly expensive and time consuming, led CESG to develop the domestic Commercial Product Assurance (CPA) scheme, which launched in 2011 and certified its first product in 2012. CPA was designed to provide assurance of the security of various categories of commercial products, oriented primarily towards authorising their use in government networks, in a more rapid and lightweight manner than CC had done. Similar to CC, CPA evaluations are conducted by labs testing products against a set of security characteristics (the equivalent of CC's protection profiles).

CC and CPA are concerned with evaluating the security of commodity technology, understood in this context as technology generally available for purchase: routers, switches, firewalls, operating systems, VPNs, and so on. A parallel assurance tradition concerns the evaluation of special purpose ‘high assurance’ technology, such as high-grade cryptographic devices for use in secure government networks. This is the core legacy of the secure communications tradition, and these evaluations have remained within CESG/NCSC rather than being outsourced to private sector labs. Such evaluations also engage in a contrasting manner: where CC/CPA are designed for the evaluation of a finished product, high grade evaluations engage at an early stage in the design process, creating an assured product through co-design.

A third noteworthy tradition that sits alongside commodity and high-grade assurance of technology is the audit-based certification of organisational processes. In our domain, this primarily means gaining certification to standards such as ISO 27001 for information security management systems and ISO 9001 for quality management. These certifications apply to organisations, rather than to products, although there is a degree of overlap. While they are product-centred, CC involves evaluation of fault remediation processes, and CPA requires a ‘build standard’ audit of organisational processes for technical development. And although ISO 27001 is a certification for a security management system, the scope of the certification can be tailored to focus on a particular technical system or service (plus the security management wrapper placed around it).

A set of recent interventions into assurance policy sets the scene for this study. In 2019, the UK formally stepped back from Common Criteria, giving up its ‘certificate producer’ status, citing its ‘diminishing relevance’ in the contemporary world (National Cyber Security Centre 2019). In 2020, the decision was also made to phase out CPA, apart from for electricity and gas smart meters (for which CPA certification is required by law). And third, the NCSC began a pilot for a new style of ‘principles based’ assurance. The context, then, for these interviews was one of upheaval and uncertainty about the future.

## Assurance characterised

Everyone knows that assurance is essential to the security of digital infrastructures, but many – most – voice scepticism about the value of formal certification schemes. Many of my interlocutors urge caution about the meaningfulness of products having a ‘badge’ that says: ‘this is secure.’ Articulating the reasons for this, and thus the reasons government policy has been changing, featured prominently in the interviews I conducted for this study. These explanations relate the problems of doing assurance to a cast of characters, standing in various positions in relation to the production and consumption of technology, the production and consumption of knowledge about technology, and the architecture of assurance itself.

### *Hidden manipulators*

One of the prevalent ways in which concerns about assurance are articulated is through narratives about agents who are not visible to the consumer of a certificate, but that nevertheless may have intervened upstream of an evaluation to sway a result. If assurance schemes can be ‘gamed,’ any result should be treated with suspicion.

The following excerpt is drawn from an interview with a senior consultant who works with vendors of secure technical products. We were talking about the role that test labs play in the development process. He told me that customers who were interested in using his company’s product in their network would usually arrange for a test lab to run some penetration tests on the product as a standard activity during implementation. He then went on to contrast these bespoke engagements, specific to that customer’s project, with the ‘bad’ side of the labs, associated with formal certification schemes.

> *Interviewee*: One of the reasons why NCSC are moving away from Common Criteria, because you must have heard of that example they quote all the time, which is the Microsoft example with Windows NT. Which is

> here is a networked operating system, and here's a security target says it's perfectly secure if you don't connect it to a network.
>
> *Me*: –Yeah
>
> *Int.*: And a lab proved that.
>
> *Me*: Yeah?
>
> *Int.*: If you don't connect it to the network I can't break into it over the network. It's just like, oh well, well done guys!
>
> *Me*: (laughs) (0.4)
>
> *Int.*: So so so that that's that's kind of the the answer, so whilst I presented that tradeoff as a positive working with the labs, that that's our chosen way of engaging. That, that is the counter example as to why that's bad because Microsoft were able to (0.2) negotiate that to the point that it was (0.4) of no value to the customer.
>
> (Excerpt 1: Interview with a senior consultant for high assurance technology, 2021)

The account here is presented as a second-hand anecdote that is revealing about assurance more widely. It is meta-narrated as an account of the kind of story that the NCSC are telling 'all the time,' relating the problems of differential visibility between contexts of technology production, evaluation and use.

What interests me here is not the truth of the claim, i.e. whether Microsoft did in fact manipulate the Common Criteria process in quite such a brazen way as my interlocutor suggests. Indeed, a similar claim, which may be the root of this story, concerns the conditions under which Windows NT obtained its TCSEC 'C2' certification (Mackenzie and Pottinger 1997, p. 56). What interests me is the relations between characters, and how the explication of these relations serves as a contestation of the authority of certifications.

In contrast to the certificate, which says 'this is secure,' the security target (the far less visible documentation of assumptions) adds 'if you don't connect it to a network.' In the story, Microsoft makes use of this differential visibility between the core texts of the scheme, to obtain a correct, but deceptive, result. A contrasting character here is the lab that 'proved' the result, and thus appears to be a dope, manipulated via these assumptions. The lab's kind of agency, or lack thereof, is evident in this exaggerated reference to 'proof.' CC evaluations are not proofs in a mathematical or logical sense, but calling evaluators' work 'proving' suggests a determining relationship between premises and the result that follows from them, and thus the limited freedom for those who carry out the evaluation to judge whether it is reasonable.

A second example of a similar kind of narrative is useful, because this kind of scepticism is not specific to CC (or, of course, to Microsoft). This time, I am discussing ISO 27001 with a risk manager from a telecommunications organisation.

> *Interviewee*: The problem is, as they say, they are snapshots and >you can game them< and (.) I'm, I mean, a telco who I will not name, because it's probably not fair on them, set up a laptop that they called, the laptop's name was <telco network.> They then had a 27001 certificate for that laptop, which is dead easy because it's a secure device, stands there, and they could put on their wall that they had (0.2) 27001 certification for <telco network.> Now nobody could > stop them doing that because <that's what that device was called, and that's the scope of that 27001 certificate
>
> *Me*: Did you see it?
>
> *Int.*: Yeah
>
> *Me*: Is this, is this a legend or were you there in the presence of said certificate?
>
> *Int.*: No, no, no, this is not this is not a legend. I was not there when it was done. It was at a telco, but it is a known thing that that was done (0.4) as an experiment

> (Excerpt 2: Interview with a risk management professional, 2020)

The absurdity of a confusion of names (is 'telco's network' the telco's network or anything whatsoever that has been given those words as a name?) vividly shows the potential for manipulation that is otherwise concealed in the certificate. The interviewee clarified that the telco in the story was not trying to dupe customers or regulators with this certified 'network;' they were doing it to prove a point, a test of the test. Just as we saw above, this little narrative constitutes mistrust in its intertextual relations with the certificates, insofar as it makes visible the fact that certificates, as intermediaries between contexts, can (be used to) conceal their own conditions of production/manipulation as well as to reveal security. And further, propagating itself by meta-narrating the very manner of its propagation, the telco not only demonstrates the absurdity by producing it, but in anticipation of its audiences, displays the certificate *on the wall*.

## Box tickers

The events related in these accounts may or may not be fairly described. The lab and auditor presumably could have protested about how these evaluations were being manipulated. But they are placed in similar positions. Although they do the work of certification, they are characterised as having very little agency. Whatever they do is 'really' produced through them by other agents: on the one hand a hidden manipulator and on the other a prescriptive set of assurance procedures they are bound to obey.

Assurance schemes are notoriously bureaucratic, and bureaucracy is associated with the kind of character we might call a 'procedural dope,' someone who follows procedure to a pathological extent (the analogy is with the species' of 'judgemental dopes' that Harold Garfinkel suggested were implicit in conventional sociological explanations; 1967, p. 68). If there is a stereotypical activity associated with such characters, it is 'box ticking,' an activity that serves as a source of perpetual critique when it comes to modern knowledge practices. As Michael Power observes in relation to risk management,

> for all the stridency of … ubiquitous criticisms, and their near unanimous acceptance by both regulator and regulated, there is a striking and puzzling fact to explain: "box-ticking" as a finely grained process in some broad sense persists, with at best only incremental diminution. (Power 2007, p. 153)

In cyber security assurance, it is not just the evaluator that is characterised in this way. A second vein of mistrust concerns the certificate's authority, or its powers to authorise decisions, insofar as this can be understood to render its consumer or 'reader' passive (Cooren 2000, p. 197). Of course, if the certificate were always right, there would be little cause for concern, but if the certified option is sometimes *not* the best choice, a 'box ticking mentality' that selects a certified product anyway becomes a source of problems. Such a box ticker will not see the context (which might for instance be a context of newer, better, but as yet uncertified product versions) in which the 'tick' ought not to be taken on face value.

Consider the three reasons offered by the NCSC for the closure of CPA:

> The length of time it takes to certify a product using the current assurance process may restrict the number of times a vendor submits a product for certification. Cost can also cause this effect.
>
> When building IT systems, there is a tendency to specify and use older, CPA-certified components to tick a risk management box rather than use newer uncertified components, that might be more functional and more secure because they have no badge.
>
> There is a clear imperative to operate a scheme that is capable of certifying a larger volume and range of products. (National Cyber Security Centre 2020)

CPA did not achieve comprehensive coverage of products in the market. The first rationale conveys the NCSC's explanation for why this was (it was too expensive and time consuming), while the second relates this issue to the unintended side-effect of incentivising worse security decisions via ‘box ticking.’ The third concerns the desire thus elicited to fix the problem with a new, reformed scheme.

In one interview, discussing the reasoning behind the CPA closure, a member of the NCSC’s technical team gave an account of assurance that related the gap between the ‘right answer’ and the best choice to the misaligned temporalities of assurance:

> *Interviewee*: If we pick on an example, let’s take erm -routers for example.
>
> *Me*: Yeah
>
> *Int.*: Giving names I, I’m not sure if they ever did, but Cisco for example, Cisco, a huge company, they are producing a router today which we could certify.
>
> *Me*: Yeah
>
> *Int.*: No problem. Job done. Six months’ time, they are going to produce version two of this router. Now -version two is almost certainly going to have a significant improvements to it, over the one we have certified, but if -if a government depar (.) body says you must use a certified device. They going to say I’ve got to use that one, the older version, which probably isn’t as good, but it’s >but it’s certified,< whereas what we want them to do is actually use the latest version, you know, like operating systems, there’s a reason why they are updated. They usually they usually fix more bugs than they introduce and you know. So we want people to be using the latest one. But because it has to be certified. Ooh you=know it’s a, we can’t do it. Can’t keep up.
>
> (Excerpt 3: Interview with NCSC Technical Team Lead, 2020)

The issue here is not with the validity of the evaluation, but rather with the manner in which certifications are interpreted. We hear the voice of the procedural dope: ‘Ooh.. we can’t do it.’ It ‘has to be certified.’ The overdetermining rules, along with a character inclined to follow them, turn a bad choice into the ‘correct’ one, achieving the opposite goal to the scheme’s overarching intention of improving security.

## *Free thinking and rethinking assurance*

In Excerpt 3, both narrator and reader can appreciate that the use of the older version is not a good choice. We are presented with the reasons: the newer versions ‘usually fix more bugs than they introduce.’ Knowing this, we thus adopt a position from which box ticking is visible as a problem. From such a position one may voice scepticism, a sarcastic narrator’s ‘well done, guys’ (Excerpt 1), or justify an agenda for reform.

While the narratives thus far have concerned the production and use of certifications, the subject matter of assurance is not limited to this ‘ground level’ consideration. As we have seen, assurance has been perpetually reformed, and reform is itself a subject matter for further narrative problematisation. The following is drawn from a discussion with a veteran of technology assurance world:

> *Interviewee*: So before CPA, mainly what we did was was Common Criteria, which I guess you’ve come across that, ISO 15408 and that
>
> *Me*: Yeah
>
> *Int.*: So one of the common criticisms for that (.) that’s dogged it forever, is is that the evaluations >take too long and cost too much.< (.) So there came a point at which erm erm –some of the guys in CESG wanted to come up with something that still achieved that. (.) Basically, they felt you could focus in more on the <real security aspects and cut out a=lot=of=the bureaucracy (.) and that was what they tried then to do with CPA. (.) Now I –I’m going to give a little bit of a sort of meta comment. The number of times people have told me that we should do that, (.) as though that was some great discovery (.) is quite large and and (0.2) usually what happens is they have to bang their heads against it for a certain while to realize that (.) the reason things >cost too much and take too long is because it’s a <hard problem>, not because the rest of us are too stu:pid to have spotted those problems. But >every now and then you have to go< (.) that’s part of the cycle (.) somebody else comes up and says ^oh what you’ve missed is it’s taking too long and costing too much (.) (laughs)

> *Me*: (laughs)
>
> *Int.*: So I I say slightly jadedly that there was a bit of that element in CPA. (.) It was it was a oh well, we can solve those problems
>
> *Me*: Mmm
>
> *Int.*: And I think <that's relevant becau::se> one of the things that (.) >and I've been on a workshop only this morning where < precisely this was said (.) >the problem with CPA is it takes too long and it costs too much<
>
> (Excerpt 4: Interview with a Technical Manager of a UK Cyber Security Testing Evaluation Lab, 2020)

The politics of characterisation come to the surface here. The interviewee refers explicitly to how he is characterised (as a dope) in others' accounts of the problems: it is 'not because the rest of us are too stupid.' Claims that older schemes are too costly and slow, which have considerable driving force in reforming assurance, portray their architects as lacking insight. But in excerpt 4 it is the critic/reformer who is unable to see that they themselves are repeatedly announcing the same thing 'as though that were some great discovery' (that is, until they have 'bang[ed] their heads against it'). Parallelism is a common feature of storytelling (Bauman 1986), here acting as a device that portrays reformers as mindless: they think they have agency, and indeed they do sustain the constant churn of reform—but they are actually following a well-worn path that will lead them to the same place.

If the aspiration to efficiency (cost and time) is a source of perpetual challenge, some of the commitments associated with earlier schemes become decidedly dated. Ross Anderson remarks that 'the idea that a device should be secure because someone spent $100,000 getting an evaluation lab to test it five years ago would strike most people nowadays as quaint' (Anderson 2020, p. 1015; a 'quaint' idea being not simply wrong, but associated with a character stuck in the superseded thinking of the past). The point is not that assurance should be rapid and cheap, but rather that it must be dynamic and responsive to the pace of technological change, much as, in an earlier era, Common Criteria was borne out of the necessity to reformulate assurance for a world of increasingly diverse form factors of computing.

The latest attempts to reformulate cyber security assurance in the UK herald a new approach known as 'principles based assurance.' This was piloted by the NCSC in 2019–2020 in collaboration with a group of labs and technology vendors, producing principles-based evaluations of a set of 'cross domain solutions.' These included a 'browse-down' solution for accessing the internet from within a secure network, a system for implementing instant messaging/group chat across the boundary between networks, and an email gateway. While time and cost remain key considerations for policymakers, their major intervention that defines this approach lies elsewhere.

'Principles' in this context are placed in contrast to the prescriptive requirements according to which products are evaluated in schemes like CPA. The CPA's 'Security Characteristics' make explicit demands of assured technology. For the labs, these translate into a set of usually unambiguous tests, resulting in a series of pass/fail results, which for the customer of assured technology adds up to a simple binary outcome: the product in question either has the certificate or it doesn't. Principles, in contrast, set out design guidelines that need to be interpreted. Principles are conceived explicitly in opposition to the characterisation of prior schemes as overly determining for both evaluators and customers, and a form of assurance is sought that leaves room for context-sensitive judgement, both for the evaluator and for the consumer of the report.

This notion of interpretative judgement responds in part to the characterisation of procedural dopes, but it also draws on the tradition of the NCSC's high-grade evaluations. The pilot, then, tested out whether commercial labs were able to act like the NCSC's own high-grade assurance evaluators. Indeed, the principles used in the pilot as the basis of evaluation were developed from the NCSC's own evaluators' experience in conducting high-grade assurance activities for cross domain solutions, ported across into the setting of commercial labs. The following is from an excerpt of an interview with one of the architects of the pilot.

> *Interviewee*: When NCSC does assurance, we are quite erm free thinking >I think is the best way to describe it.< You know we're not driven by contracts or money or anything like that. Our job is to make things safe for the country. That's that's what we're here for. I think we're naturally inquisitive, so we'll we'll question things >partly because in some cases we have written the standards ourselves< so we know what the intent of them is. Therefore, we can question them.
>
> *Me*: Yeah,
>
> *Int.*: Whereas the labs are very much used to (0.4) here's the Common Criteria protection profile or here's the CPA security characteristic. Treat that as your bible thou shalt follow it. And I think that creates quite a (1.2) What's the word? Stovepiped way of thinking, rigid way of thinking. That's probably a better one, and it was trying to break that rigid thinking into a bit more freeform thinking. Erm the other thing to you know if you >if you look at the principles< it's, it says things like you know. You should have a hardware break in your import path or export path. (1.4) And then you've got to try and work out alright, what's the test I'm going to do to prove that's there? Whereas if you look at the (CPA) security characteristic, it will be >there will be this device in this place and you will test it by this.<
>
> *Me*: Yeah,
>
> *Int.*: So they had to think all of that through for themselves, which is not something, they've had to do to date.
>
> *Me*: Yeah yeah,
>
> *Int.*: So that that that took a lot of handholding by us.
>
> (Excerpt 5: Interview with a member of NCSC's product evaluation team, 2021)

The free thinking of NCSC's evaluators is related to their interests, to their lack of financial incentive, and their purpose – much less profane than profit and contracts, of 'making things safe for the country.' They can question the standards. The rigid thinking of the labs, in contrast, needs to be 'broken.' But it is not simply a rigid pattern of thought, like a deeply engrained habit, for rigid thinking is associated with a religious relationship with text, the language of commandments: 'thou shalt follow it.' In contrast to the absent agency of dopes who are manipulated into proving an absurd result, here rule following is a sign of submission to a transcendent text. Moving the labs away from that relationship with requirements, on the other hand, requires 'handholding,' recalling parent and child relations and representing free thinking as something needing nurture from those in a position of mastery.

The challenge for the labs is to conduct evaluations in a rigorous and repeatable manner in conditions where the 'right answer' will depend upon where and how the product will be used and how it fits into a wider system or network architecture. Most principles allow flexibility in how they are addressed. The 'content-based attack prevention' principle for cross domain solutions, for instance, does involve a strict stipulation that a verification engine will be needed, but allows that this engine may be hardware-based, software-based, or hybrid in its architecture. What counts as a good choice will depend on the nature and form of the network it is going to be used in and the risks and threats facing that specific system. Rather than these considerations being decided in advance (during the formulation of evaluation criteria) they are to become part of the evaluation, and instead of producing a simple certificate, the evaluation will produce a report that sets out how the evaluators put the product into context, such that a reader might understand how appropriate this particular choice would be for *their* context. This is a very long way from the simple 'pass/fail' format of the past, and we move beyond the aspiration to a pure object-knowledge, towards an openly expert-constituted form of knowledge based on contextually situated judgement (Laurent and Thoreau 2019).

With this judgement at its core, principles-based assurance also makes new demands on the consumer of its outputs. If the boxes are taken away from the evaluators, the consumer of assurance is no longer working with simple 'ticks.' The intention is to drive them toward more engaged and considered thinking about what it would take to be secure in their particular context. However the capacity for engaged and considered thinking does not come for free. This was the problem articulated by an evaluator involved in the pilot:

> *Interviewee*: I >I fully support this idea the NCSC doesn't want to give a pass fail certificate.< I also envisage that everyone else >really, really liked< having a pass fail certificate that they did not have to think about. Not just the people buying it, but the the actual manufacturers as well, they want their certificate. They want to be able to put a stamp on their product that says <this is good> NCSC says so, and NCSC are the people who have >all of the power and the authority< so people are going to have to (0.2) go where they are leading, but I think that <there's got to be a piece of> education <for everyone that's being sort of dragged along this journey> a little bit, because I think that they are then going to get >instead of a certificate< a 30 page report that >they have to read and understand how it applies to them< like that's (.) that's a bi::g change, especially if you think about, you know, like who's buying this? Yeah, it might be big CNI, erm central government department. They have (.) tonnes of risk people to do that kind of >technical risk people to do that thing.< ^If you are talking about, I don't know a university with an IT team of three people, (0.2) like they barely are gonna have enough time to read through it, let alone really understand what it means for them,
>
> *Me*: Yeah
>
> *Int.*: So that's going to be really, it's it's the the, the smaller organizations who don't have that understanding of risk, are going to struggle the most with it,
>
> (Excerpt 6: Interview with a senior evaluator at a UK test house, 2020)

People prefer the certificates because they provide a simple answer 'that they did not have to think about.' In contrast with box tickers rendered problematic for their lack of contextual thought, this final excerpt shows that unthinking acceptance of assurance results can itself be put into a context, one of the resources (time and attention) it takes to understand.

This account is directly tied to the interviewee's experience *writing*: Writing the report for the pilot, something that raises the question of *for whom such a text is intelligible*, for instance for 'an IT team of three people.' If the older model of simple certificates was associated with the production of a determining text and an unthinking mentality, we start to see new concerns emerging about *who* will be able to recognise themselves as the kind of reader anticipated by the new kind of text.

As much as they are revisions to standards and styles of evaluation, then, reforms to assurance are also interventions in characterisation. The attempt to produce agents capable of genuinely taking responsibility and making optimal decisions about cyber security is deeply conditioned by narratives about how assurance happens and what goes wrong. The latter are not only descriptive, but constitute mistrust, performing assurance as problematic: Stories impinge upon the capacity of schemes to 'give' assurance in a face-value manner, and they give shape to a dynamic field of interventions that seek a better cast of characters, generating, in so doing, the potential for further problematisation.

We are thus left with the hint of a new kind of problem, one heralded by experiments with principles based assurance, in which the production of more complex 'primary texts' (reports, as opposed to simple certificates) raises new issues of interpretability and identification: where some of those characterised as decision makers struggle to identify themselves as 'that sort of person' capable of interpreting detailed context-sensitive information. This would bolster existing discourses in cyber security about *a lack of the right sorts of people*. Often known as the 'cyber skills gap,' this is the idea that the solution to our pressing problems with technological insecurities will need to involve the creation of an expanded contingent of highly trained technical experts: enough experts, and our problems would be kept in check.

Any such turn to expertise exacerbates challenges associated with the delegation of judgement and the consequential effects for the recognition of alternative voices. Where solutions to social problems with technology are delegated to experts, the potential for public debate can easily be closed down (Wynne 1988). 'Filling' the cyber skills gap might alleviate problems with technology, but it could easily exacerbate the challenge of engaging with societal issues, such as *whose* security is or should be at stake.

## Concluding remarks

This paper is a response to the stories practitioners have told me about assurance. As well as in interviews, similar stories are told in industry forums and working groups, as well as in informal moments. In a domain where discretion is key, details of technical and commercial sensitivity are routinely avoided, these rather generic kinds of stories are common currency. Such texts embody a distributed form of awareness of socioeconomic context, the market forces, fears of perverse incentives, overworked and beleaguered security teams, the power of government, and the responsibility to keep systems operational. They present ordered constructions of what is going on with assurance, picking out from a hugely complex field just those specifics that matter.

I want to avoid treating these accounts as 'containers' or 'conduits' of opinions, evidence of what people seem to think. Instead, I draw on the resources of narrative analysis, semiotics, and STS to interpret them in light of what they are able to *do*: their agency. I suggest that they 'do' mistrust, awkward though this turn of phrase may be. Such stories are not froth on the surface of social life, but play an active role in making the social-material assemblage of assurance. The security of digital infrastructures thus can be understood to emerge from the interplay of technical and interpretive construction.

Stories about problems impinge upon the capacity of certificates to make simple monological declarations. It is hard for people, unless they are very specifically isolated from this kind of talk, to treat certificates as simple reliable truths. The power of certificates to authorise the use of certain product choices is lessened, with implications for decisions and for policy. And in this latter respect, we see the future-making, prospective efficacy of characterisation: the manners in which the domain is understood as problematic defines the contours of the terrain on which new policy is formulated, eliciting a cascade of further narrative.

I want to close the essay with the suggestion that this line of thinking takes us toward a conceptualisation of mistrust that does not, in any simple way, equate with an absence of trust as it has traditionally been understood, and I want to argue that this has implications for the way in which we imagine the impact or agency of scholarly analysis.

To say that the literatures on trust are expansive would be an understatement. But a very prevalent assumption is held across fields of study that trust is a fundamental element of sociality and, thus, when we study trust we are studying social foundations. This is true of the rational choice tradition, for which the game theoretical *prisoner's dilemma* is an archetype of trust (e.g. Bacharach and Gambetta 2001, Cook *et al.* 2005). It is also true for the phenomenologically informed tradition, for which Garfinkel's ethnomethodological *experiments in trust* play a similar role (1967). From a rational choice perspective, situations of trust resemble decision problems in which an agent must gamble (Coleman 1994, p. 99). A deficit of trust is anti-social: where gambling on the behaviour of others is not a rational strategy, the universe of rational action is severely curtailed, potentially leading to socio-economic stagnation (Fukuyama 1996). From the ethnomethodological perspective, the rational choice theorists have missed the point: in formulating the problem as an abstract problem of choice, trust's true importance is analytically unavailable, for trust is the constitution in the concrete happening of ordinary interaction of a situation as *that* kind of situation that it is (Anderson and Sharrock 2014, Watson 2014). A lack of trust is equally anti-social, but for different reasons: a breakdown in the coherence of situations being, for instance, associated with existentially threatening ontological insecurity (Giddens 1991).

In contrast with a lack of trust, the study of mistrust directs our attention to circumstances in which the explication of foundations is suspended. When Carey examines mistrust in the High Atlas, the point is not that people don't have intentions or character, but that making these things the overt topic of discourse strikes people as odd and unanchored. In the case of cyber security assurance, it is not that there is no 'real' quality of technology as secure or insecure, but that the knowledge of this quality becomes the topic of an avalanche of discourse on the character and intentions of those involved in its production and use. In both cases, what is studied in the analysis of mistrust is the manner of relating to the explication of 'real' intentions or 'real' quality, in other words the various ways in which *contingency* becomes relevant in social life.

So if we reflect on the production of scholarship on mistrust, we might observe that we work under the weight of expectations, seeded by history and narrative form, that scholarship produces authoritative accounts, capable of subsuming informants' stories. A text such as this carries the expectation of being a 'master' characterisation of assurance, one capable of taking these fragmentary narratives and wrapping them together to produce a definitive description. Yet, despite such an expectation, and my own implication in setting it up, there remains an important sense that this paper fails to be encompassing, and is better understood as *adjacent* to the stories others tell. Cyber security practitioners, after all, speak in familiar terms, appealing to the interests, incentives, institutional structure and political economy of assurance. Annelise Riles has written profoundly on the analytical use of this parallelism between the textual practices of modern organisational forms and academic social science (2000).

But what I want to pick up here is the implications of adjacency for how we might understand the efficacy of scholarly discourse. An authoritative account might promise to resolve the problem of assurance, to define what is really wrong and what might be done about it. An adjacent account's potential lies in its own capacity to amplify voices, and promises a critical cyber security insofar as those voices are otherwise unheard.

The examination of narratives about principles-based assurance led us to a concern about how the consumers of assurance reports are characterised. Where those reports are complex, and the interpretation potentially time consuming, we seem to be left with an anticipated expert reader, capable of exercising trained judgement in relation to this information and the context of implementation, or else with a non-expert, whose capability is, in contrast, lacking in relation to the expert they are not. This framing of the problem has potential consequences in disempowering those thus characterised, and in the characterisation of society more broadly as having a problematic deficit of expertise. As Brian Wynne pointed out years ago, delegation of socially complex technical issues to expert judgement can create challenges for public engagement (1988). One kind of answer to this would be to seek languages for characterising the consumer of assurance that do not fall into an expert/non-expert dichotomy, to seek to intervene through characterisation.

The language I have in mind here is the language of *care*. A mainstay of feminist scholarship, care falls outside of the parameters of most of the characters we met. We saw a glimpse, perhaps in the 'handholding' of excerpt 5. But care is not a kind of control, and it need not be patronising, for care is marked by an attentiveness to the needs of another. It is not free, but neither is a caring character a passive dope: care implies an existential commitment, a resolute ethical stance (Tronto 1995). In recent ethnographic work, Laura Kocksch draws on the notion of care to make sense of her observations of quotidian cyber security practice (2022). She sees care in contrast to the kind of mastery associated with expertise. 'To care' she writes, 'is engaging in efforts to do good without having a final answer' (Kocksch 2022, p. 238). If assurance is formulated as providing resources that support practitioners to take care of those digital infrastructures that fall under their stewardship, we would tell stories about a good relationship with technology, something that admits a plurality of forms and styles, and the possibility of engaging with the thorny questions of public interest, around what is worth securing and who is secured.

## Notes

I use the following transcription conventions: for emphasis, ^higher intonation, micro (.) pause, longer (0.4) pauses in seconds, >faster speech<, <slower speech>, dra::wn out words, words=run=together, and –stumbling over the start of a word.

## Acknowledgements

Many thanks to three anonymous reviewers and to the editors of the special issue, whose comments were very helpful in developing this paper.

## Disclosure statement

No potential conflict of interest was reported by the author(s).

## Funding

This work was supported by UK Research and Innovation [grant number MR/S037373/1].

## Data access statement

In order to protect participant confidentiality, supporting data cannot be made openly available. Further information about the data and conditions for access are available from WRAP at http://wrap.warwick.ac.uk/166796/.

## ORCID

*Matt Spencer* http://orcid.org/0000-0002-5146-6201

## Notes on contributor

***Matt Spencer*** is an Associate Professor at the University of Warwick's Centre for Interdisciplinary Methodologies. He is a UKRI 'Future Leaders Fellow' and his research concerns the social and cultural implications of computing. His current project develops a socio-technical analysis of cyber security. His previous research examined the computational transformation of the physical sciences.

## References

Aldrich, R.J, 2013. Whitehall wiring: the Communications-Electronics Security Group and the struggle for secure speech. *Public Policy and Administration*, 28 (2), 178–195.

Alexander, J. C, 2011. 'Market as narrative and character: for a cultural sociology of economic life'. *Journal of Cultural Economy*, 4 (4), 477–488.

Amoore, L., and Piotukh, V, 2015. Life beyond big data: governing with little analytics. *Economy and Society*, 44 (3), 341–366.

Anderson, B., and Sharrock, W, 2014. The inescapability of trust: complex interactive systems and normal appearances. In: R Harper, ed. *Trust, computing, and society*. Cambridge: Cambridge University Press, 144–171.

Anderson, J.P, 1972. *Computer Security Technology Planning study*. US Airforce Electronic Systems Division.

Anderson, R, 2020. *Security engineering: a guide to building dependable distributed systems*. 3rd ed. Hoboken: John Wiley & Sons.

Aradau, C., & Blanke, T, 2015. 'The (big) data-security assemblage: knowledge and critique'. *Big Data & Society*, 2 (2).

Aradau, C., and Blanke, T, 2017. Politics of prediction: security and the time/space of governmentality in the age of big data. *European Journal of Social Theory*, 20 (3), 373–391.

Bacharach, M., and Gambetta, D., 2001. Trust in signs. In: K. S. Cook, ed. *Trust in society*. New York: Russell Sage, 148–184.

Bakhtin, M, 1984. *Problems of Dostoevsky's Poetics*. Trans. C. Emerson. Minneapolis: University of Minnesota Press.

Bauman, R, 1986. *Story, performance, and event: contextual studies of oral narrative (Vol. 10)*. Cambridge: Cambridge University Press.

Bauman, R., and Briggs, C. L, 1990. Poetics and performances as critical perspectives on language and social life. *Annual Review of Anthropology*, 19 (1), 59–88.

Bowker, G, 1995. Information mythology—the world of/as information. In: L. Bud-Frierman, ed. *Information acumen: the understanding and use of knowledge in modern business*. London: International Thomson Publishing, 231–247.

Breakwell, G. M, 2020. Mistrust, uncertainty and health risks. *Contemporary Social Science*, 15 (5), 504–516.

Carey, M, 2017. *Mistrust: an ethnographic theory*. Chicago: HAU Books.

Callon, M, 1986. Some elements for a sociology of translation: domestication of the Scallops and the fishermen of St-Brieuc Bay. In: J Law, ed. *Power, action and belief: a new sociology of knowledge?* London: Routledge, 196–223.

Coleman, J. S, 1994. *Foundations of social theory*. Cambridge: Harvard University Press.

Cook, K. S., Hardin, R., and Levi, M, 2005. *Cooperation without trust?* New York: Russell Sage Foundation.

Cooren, F, 2004. Textual agency: how texts do things in organizational settings. *Organization*, 11 (3), 373–393.

Cooren, F., 2000. *The organizing property of communication*. Amsterdam: John Benjamins Publishing.

Cooren, F., and Sandler, S, 2014. Polyphony, ventriloquism, and constitution: in dialogue with Bakhtin. *Communication Theory*, 24 (3), 225–244.

Culpeper, J, 2014. *Language and characterisation: people in plays and other texts*. Milton Park: Routledge.

du Gay, P., *et al.*, 2019. Character and organization. *Journal of Cultural Economy*, 12 (1), 36–53.

Fahnestock, J, 1998. Accommodating science: the rhetorical life of scientific facts. *Written Communication*, 15 (3), 330–350.

Fiske, J, 1992. British cultural studies and television. In: R C Allen, ed. *Channels of discourse, reassembled: television and contemporary criticism*. Chapel Hill: University of North Carolina Press, 284–325.

Fukuyama, F, 1996. *Trust: the social virtues and the creation of prosperity*. New York: Simon and Schuster.

Garfinkel, H, 1967. *Studies in ethnomethodology*. Englewood Cliffs, NJ: Prentice-Hall.

Garvey, J, 1978. Characterization in narrative. *Poetics*, 7 (1), 63–78.

Giddens, A, 1991. *Modernity and self-identity: self and society in the late modern Age*. Stanford: Stanford university press.

Kocksch, L., 2022. *Fragile computing – theorising cybersecurity in practice*. Doctoral dissertation. Ruhr-Universität Bochum.

Latour, B, 1996. On interobjectivity. *Mind, Culture, and Activity*, 3 (4), 228–245.

Latour, B., and Woolgar, S, [1986] 1979. *Laboratory life: the construction of scientific facts*. Princeton: Princeton University Press.

Laurent, B., and Thoreau, F, 2019. Situated expert judgement: QSAR models and transparency in the European regulation of chemicals. *Science & Technology Studies*, 32 (4), 158–174.

Lipner, S.B, 2015. The birth and death of The Orange book. *IEEE Annals of the History of Computing*, 37 (2), 19–31.

MacKenzie, D.A, 2004. *Mechanizing proof: computing, risk, and trust*. Cambridge: MIT Press.

MacKenzie, D.A., and Pottinger, G, 1997. Mathematics, technology, and trust: formal verification, computer security, and the US military. *IEEE Annals of the History of Computing*, 19 (3), 41–59.

Margolin, U, 1986. The doer and the deed: action as a basis for characterization in narrative. *Poetics Today*, 7 (2), 205–225.

Moor, L., and Lury, C, 2018. Price and the person: markets, discrimination, and personhood. *Journal of Cultural Economy*, 11 (6), 501–513.

Mühlfried, F, 2018. Introduction: approximating mistrust. In: R. Mühlfried, ed. *Mistrust: ethnographic approximations.* Bielefeld: transcript, 7–22.

Mühlfried, F, 2019. *Mistrust: a global perspective.* London: Palgrave MacMillan.

National Cyber Security Centre., 2019. *Common criteria*. Available at: https://www.ncsc.gov.uk/information/common-criteria-0 [Accessed 10 Jul 2021].

National Cyber Security Centre., 2020. *Commercial product assurance*. Available at: https://www.ncsc.gov.uk/information/commercial-product-assurance-cpa [Accessed 10 Jul 2021].

Power, M, 2007. *Organized uncertainty: designing a world of risk management*. Oxford: Oxford University Press.

Riles, A, 2000. *The network inside out*. Ann Arbor: University of Michigan Press.

Shapin, S, 1994. *A social history of truth: civility and science in seventeenth-century england*. Chicago: University of Chicago Press.

Star, S. L., and Ruhleder, K, 1996. Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information systems research*, 7.1, 111–134.

Spencer, M, 2021. Creative malfunction: finding fault with rowhammer. *Computational Culture: A Journal of Software Studies*, 8.

Tronto, J. C, 1995. Women and caring: what can feminists learn about morality from caring? In: V. Held, ed. *Justice and care: essential readings in feminist ethics*. Milton Park: Routledge, 101–116.

Watson, R, 2014. Trust in interpersonal interaction and cloud computing. In: R Harper, ed. *Trust, computing, and society*. Cambridge: Cambridge University Press, 172–198.

Woloch, A, 2004. *The one vs. the many: minor characters and the space of the protagonist in the novel.* Princeton: Princeton University Press.

Woolgar, S., and Lezaun, J, 2013. The wrong Bin Bag: a turn to ontology in science and technology studies? *Social Studies of Science*, 43 (3), 321–340.

Wynne, B, 1988. Unruly technology: practical rules, impractical discourses and public understanding. *Social Studies of Science*, 18 (1), 147–167.