A Lemon by Any Other Label

Vaibhav Garg
Comcast Cable, Philadelphia, PA, U.S.A.

Keywords: IoT, Security, Labeling, Energy Star, Economics.

Abstract: Apparent under-investment in IoT security is often explained by the lack of consumer demand engendered by information asymmetries. One proposed solution is to create IoT security labels as a market signal of differentiation. Such labeling may be binary, graded, or descriptive. Each label type can be further differentiated based on distinct implementations. This paper surveys the existing efforts to create IoT security labels along with the inherent limitations of individual approaches. Overall, we find that there is limited research in this area, which makes it difficult to ascertain the components of an effective IoT security label. We recommend that label designs should limit complexity and leverage existing institutions, such as trade groups, for sustainability as well as adoption.

1 INTRODUCTION

Security is rated as the most important factor that drives consumer purchasing decisions in IoT (Ipsos Public Affairs, 2019). IoT security is not just a technical issue, but also one of economics (Anderson and Moore, 2006). If consumers are unable to distinguish between secure and insecure products, they will be unwilling to pay a premium for security and merely differentiate products on price. Manufacturers in response would have no incentive to invest in IoT security, thus creating a market of security lemons (Smith, 2019). One solution to address this information asymmetry is to create IoT security labels (Morgner et al., 2020; Knowledge, 2019; Communicating Upgradability and Improving Transparency Working Group, 2017). This has garnered new found advocacy with the emergence of highly publicized cyber attacks such as the distributed denial of service (DDoS) attack on Dyn, where attackers leveraged poorly designed internet-connected cameras.

Any effort to create IoT security labels incurs three distinct challenges. First, it necessitates defining a security assessment that any proposed label satisfies (Fagan et al., 2020; C2, 2019; DCMS, 2019). Second, it requires determining the kind of label that will usable for non-expert consumers (Harris Interactive, 2019; Knowledge, 2019; Johnson et al., 2020). Finally, to design an effective label it will be critical to determine how and what information is being communicated (Baldini et al., 2016; Emami-Naeini et al., 2020; Johnson et al., 2020; Morgner et al., 2020).

This paper surveys the extant research in these three distinct areas to provide a representative view. There are related questions that are outside the scope of this paper. For example, researchers have noted that consumers are willing to pay a 30% premium for security, and as such there is an incentive for manufacturers to invest in certification and labeling (Ipsos Public Affairs, 2019). These considerations are beyond the scope of this paper.

We begin in Section 2 by reviewing the current efforts to define 'security'. Section 3 presents an overview of labeling and associated efforts in other risk domains. In Section 4 we present examples of security risk communication, including current proposals for IoT security labels. Section 5 discusses the implications of this research for the design of IoT security labels. Finally, Section 6 concludes with a discussion of future work.

2 DEFINING SECURITY

Proponents of a security label for IoT reference the success of Energy Star labels for energy consumption (King and Gallagher, 2020). However unlike energy, there is no unit for security. Thus, any effort to create an IoT security label must begin with a definition of 'security'. The oft repeated assertion to follow 'best practices' can be less than promising, as many such efforts fail to address basic secu-

Garg, V.
A Lemon by Any Other Label.
DOI: 10.5220/0010295205580565
In *Proceedings of the 7th International Conference on Information Systems Security and Privacy (ICISSP 2021)*, pages 558-565
ISBN: 978-989-758-491-6

rity practices that will prevent brute forced attacks such as Mirai (Dingman et al., 2018). The average manufacturer is then left to differentiate between numerous distinct best practices, recommendations, and guidelines across industry, government, as well as academic sources (Bellman and van Oorschot, 2020).

The Department of Digital, Culture, Media, & Sport (DCMS) analyzed over 100 documents across 50 organizations and mapped them across 13 distinct guidelines, to create a Code of Practice for consumer IoT security (Copper Horse, 2018; DCMS, 2019). However, since their effort in 2018, many prominent additions have been made to the space of IoT security guidance (Copper Horse, 2018). In the United States, for example, the National Institutes of Standards and Technology, or NIST, in May 2020 published NISTIR 8259A (Fagan et al., 2020), which defines a set of six high level baseline capabilities that should be satisfied by a plurality IoT devices (or their controller/hubs). NIST notes that not all baseline requirement may be applicable to all IoT devices. Thus, in practice a device may be able to satisfy only a portion of the six capabilities and still adhere to the spirit of NIST's expectation. Furthermore, NIST's capabilities relate specifically to the device itself.

In contrast, the Council to Secure the Digital Economy (CSDE), in coordination with twenty other trade groups, created another IoT security baseline (C2, 2019), which consists of 10 device capabilities as well as three product lifecycle management capabilities. Others like Consumer Reports go even further by requiring specific business practices, such as making software open source, as well as non-security related requirements, such as protection from harassment (Ditigal Standard, 2020). Choosing the right IoT guidance is critical to ensuring good security outcomes (Momenzadeh et al., 2020).

To facilitate this discussion, researchers investigated the meaning of 'best practice' in IoT security (Bellman and van Oorschot, 2020). They note that most of the guidance, more than 90%, does not relate to actual practices but is outcome-based. Simultaneously, much of the guidance focuses on early stages of the IoT lifecycle and therefore is targeted towards manufacturers. More importantly they note the absence of a common and consistent vocabulary.

From a labeling perspective, this further complicates the situation. For an Energy Star style label, manufacturers must define security as well as differentiate between security level 1 and security level 1+. Is C2 better than NIST simply because it has a longer list of recommendations? Is the Digital Standard even better merely because it is more onerous?

Choosing the wrong set of requirements to satisfy can fail to provide security (Dingman et al., 2018). This may exacerbated by vulnerabilities introduced in implementation. Furthermore, manufacturers must have some confidence in the lifetime of both requirements and labels. For example, manufacturers who spent resources investing in Mozilla's Trustable Technology Mark will be disappointed that this label no longer exists (Trustable Technology Mark, 2020). Thus, choosing the wrong 'security guideline' may impose costs on both the manufacturers and down stream customers, without any long term security guarantees.

3 LABEL DESIGN

Section 2 notes the difficulty in defining security. Assuming that the manufacturers, along with appropriate stakeholders, are able to coalesce around a specific definition, e.g. the C2 consensus, the next step will be determine which kind of label will be the most effective at informing consumer behavior. Prior research notes that there are traditionally three different kinds of labels (Blythe and Johnson, 2018): 1) binary, 2) graded, and 3) descriptive. These are listed in increasing order of design complexity. In addition, new technical solutions also allow for Smart Labels (Emami-Naeini et al., 2020). These combine a traditional label with a QR code that can be used to find additional information via an online resource. This section details the properties of these different label types as well as their benefits and inherent limitations. However, given that most consumers do not find QR codes to be usable we do not cover Smart Labels as a separate category in this paper (Harris Interactive, 2019).

3.1 Binary Labels

Binary labels act as *seals of approval* to show a product satisfies a baseline set of criteria. They are considered to be more usable, compared to alternatives, as they merely require consumer awareness of the label's existence (Knowledge, 2019). Binary labels are commonly used in the food and agricultural industries, e.g. USDA organic, Fair Trade, etc. In these domains binary labels may, however, bring a false sense of security, i.e. halo effects (Andrews et al., 2011). They may also wrongly portray that the product with a label is better than the one without, i.e. dichotomous thinking (Blythe and Johnson, 2018).

Previous research in security labels suggests that a binary label may not always create the impression that a IoT device was unhackable (Johnson et al., 2020). However, the researchers did not use an existing IoT security label for their study. Instead they employed Secure by Design which has been designed for physical security rather than cybersecurity (Police Crime Prevention Initiatives, 2020). Participants may have had prior awareness of that label, a limitation that may confound their findings. Prior research in Web Assurance Seals notes that consumer behavior is informed by familiarity and awareness (Odom et al., 2002). Despite that Johnson et al. noted that while binary labels were effective, they may be less so compared to graded and descriptive alternatives at informing purchasing behaviors.

Outside of academia, there are multiple proposals for binary IoT security labels. British Standards Institute, for example, simply extended their KiteMark effort to IoT security trustmarks (BSI, 2018). IoT Security Foundation offers a targeted trustmark which uses a lock icon to indicate security (IoT Security Foundation, 2020). Additionally, multiple governments are creating voluntary labeling programs that employ a binary certification label. For example, Finland's voluntary IoT security indicator, Tietoturvamerkki, also employs a lock icon (National Cyber Security Center, 2019). Australian government has similarly proposed a trustmark based on IoT Alliance Trust Framework (IoT Alliance Australia, 2020).

However, there is a dearth of published studies examining the effectiveness of binary labels for IoT security. In fact, given that the effectiveness of binary labels is highly correlated with familiarity, such studies may be difficult to design until a specific label reaches sufficient market adoption.

3.2 Graded Labels

Unlike binary labels, a graded label asserts both that a product meets a specific standard and the degree to which that standard is satisfied. The most prominent example of such a label is the Energy Star label. These are intended to inform consumer purchase decisions to decrease the overall energy consumption of the population. Thus, Energy Star labels report the energy efficiency of a specific device. Critics note that consumers' focus on energy efficiency ignores energy consumption. This may result consumer purchases with higher overall energy consumption (Waechter et al., 2015). However, most experts conclude that Energy Star labels are intuitive and in general drive purchase behaviors in the right direction.

The success of Energy Star has made graded labels the preferred choice of many government interventions in IoT Security. For example, the Singapore government proposed the introduction of a graded label based on EN 303 645 for Wi-Fi routers and smart home hubs (Cyber Security Agency of Singapore, 2020). In United States, the Cyber Solarium Commission's report similarly advances the need for an Energy Star-like label for IoT security (King and Gallagher, 2020). However, developing an Energy Star style label for IoT security not only requires one to determine the definition of 'security' but also be able to differentiate between distinct levels of security. Section 2 notes the difficultly in determining the former; the latter can only be less tractable.

Despite that Jameison proposed the first Energy Star-styled label for IoT security (Jameison, 2016). His proposal differentiates security based on *Logical Security Posture* or LSP, which assigns negative points for each additional logical interface. This proposal is yet to be examined for both its technical assurance as well as usability. UL, however, has proposed a commercial offering that also uses a five dimensional security label based on precious metals (UL, 2020).

Johnson et al. studied the effectiveness of a UK Energy Star style label for IoT security to inform consumer behaviors (Johnson et al., 2020). Their study also examined the interplay between security and other competing consumer interests such as functionality. Their study covered four different products - cameras, TVs, wearable, and thermostat. They note that participants were more likely to select the highest level of security, Grade A, rather than the middle level, Grade G. Simultaneously, participants were more likely to choose an unlabeled device than one with the lowest security level, Grade D. More interestingly, participants expected to pay less for a device with Grade G label than one without any label.

None of these proposals, academic or otherwise, have demonstrated external validity, i.e. there is no way to ascertain that the highest level of security offered by one of these labeling schemes exposes the device to less risk, compared to similar guarantees from the lowest level label. In fact, establishing a determinant of external validity in itself is a non-trivial issue. For example, it may be possible to correlate the devices and their associated labels with the number of compromised instances of those devices, such as through Shodan. However, more popular device types are more likely to be targeted by bad faith actors. Thus, an effort to demonstrate external validity may have to address popularity as an additional variable. Similarly, deployment context is important. A device deployed behind multiple layered controls will be less likely to be infected compared to another device that is simply exposed to the Internet.

Thus, developing a graded schema for IoT security labels is constrained by all of the challenges of a binary label with the added requirement to differentiate between different levels of security. Furthermore, research indicates that some label types, i.e. towards the middle or the lower end of the security spectrum, may either be ignored or, worse, penalized by the consumer. Consequently, while Energy Star styled labels have found much traction in policy discussions, few real solutions exist, with the exception of UL.

3.3 Descriptive Labels

To avoid the dichotomous thinking imbued in binary labels as well as the difficulty of grading security levels, a third option is to develop descriptive labels. These labels simply list the security properties satisfied by an offering, without suggesting an ordinal benefit as with graded labels. Descriptive labels offer the most amount of information, compared to alternatives. As such they may be more helpful for an informed consumer but offer unnecessary cognitive overload for the less technically literate (Klopp and MacDonald, 1981). For example, prior studies have noted that many consumers are unable to understand the information presented by nutrition labels, possibly the most popular instance of a descriptive label to date (Rothman et al., 2006).

However, descriptive labels are being explored as a potential solution for IoT security by many governments. The Department for Digital, Culture, Media, and Sport (DCMS) in the UK conducted a study that found that consumers preferred to have more information on their IoT security labels than less (Harris Interactive, 2019). Consequently, DCMS proposed a labeling scheme similar to a security lifetime label that requires the manufacturer to make declarations about three security requirements (DCMS, 2019):

- That the IoT device uses a unique password that cannot be reset to a non-unique factory default.
- That there is a public venue for third party researchers to report any security findings against the IoT device.
- That the length of time during which the manufacturer will provide security updates is specified.

Morgner et al. proposed the design of a security lifetime label that explicitly looks to avoid third party certification due to the associated costs (Morgner et al., 2020). Their proposal requires manufacturers to declare: 1) the duration for which automatic security updates will be made available to the customers and 2) the speed with which the manufacturer will provision an update to address a reported vulnerability. They found that the effectiveness of these labels to inform more secure purchases is highly correlated with the perceived risk of the IoT device.

Emami-Naeini et al. designed a descriptive IoT label with both security and privacy information (Emami-Naeini et al., 2020). Their label had two layers with a shorter label on the first layer and more details on the second layer. The security information on layer one consists of: 1) firmware version, 2) update date, 3) whether the security updates are automatic, 4) how long the security dates will be available, and 5) information about access control. They also found that while their label was helpful at informing purchase decisions for higher risk devices, it did not have an impact on lower risk devices.

Shen et al. also explore a descriptive label with five categories of ‘nutrition’: 1) system (security), 2) communication (security), 3) sensory (privacy), 4) data (privacy), and 5) connectivity (information) (Shen and Vervier, 2019). *System*, in this context, consists of information about certificates, secure boot, firmware/software update methods, passwords, and authentication. *Communication* constitutes information related to encryption, internet access, as well as the device’s ability/intention to communicate with other devices on the local network. Their proposed label style uses two additional design elements compared to prior efforts. The first style uses common knowledge color coding, i.e. red indicates severe. The second style complements colors with icons.

4 RISK COMMUNICATION

Section 3 notes the distinct types of labels as well as the associated limitations of each structure. Label design goes beyond just the structure and must address the content of the label as well, i.e. what information should be communicated to the end-user and how it may be framed. For binary labels, designers need to determine the icon used as the seal of approval. Similarly, the design of graded labels, for example, may choose between a three dimensional, four dimensional, or *n* dimensional scale for some *n*. The scale itself may be framed as a grade, a star rating, etc. As mentioned earlier, UL’s IoT rating, for example, is based on precious metals (UL, 2020). Finally, a descriptive label requires designers to determine the specific information that needs to be displayed as well as the associated format.

When designing a binary label, the choice of icon may leverage distinct mental models used to communicate security risks. Camp et al. notes that security experts employ one of five mental models: 1) physical, 2) medical, 3) criminal, 4) warfare, and 5) market (Camp, 2009). Physical and criminal mental models are employed through the use of a lock symbol to indicate the presence of https. However, this may have an unintended consequence: That non-expert consumers may assume the presence of security guarantees that are not, in fact, present. For example, the use of a lock icon in HTTPS has been known to confuse non-experts, who mistakenly conflate connection security with website security (Wu and Zappala, 2018). Similarly, for IoT security, a binary label like a lock icon may be misintepreted as a sign that the device is unhackable, as opposed to merely satisfying some baseline security requirements. Yet, both Finland's Tietoturvamerkki (National Cyber Security Center, 2019) and the IoT Security Foundation's (IoT Security Foundation, 2020) binary labels employ a lock sign to indicate security. Given the prior use of lock icons to indicate connection security, its use in IoT security might similarly confuse users who may assume that the IoT device merely provides the same and no greater security guarantees. Thus, lock-based IoT labels may potentially incur either underestimation or overestimation of risk by non-expert consumers.

Simultaneously, Johnson et al. note that participants recommended a more 'cyber' centric label as opposed to one that targeted physical security (Johnson et al., 2020). Their study uses a Secure by Design (Police Crime Prevention Initiatives, 2020) label that is currently used for physical security. Since this label is used by customers in the UK to differentiate products, it may benefit from existing consumer brand awareness. Previous research has shown that the successful adoption of binary labels is driven primarily by consumer awareness. Thus, BSI's approach to simply extend their well known KiteMark logo for residential IoT security may be a more effective approach to change consumer behavior (BSI, 2018).

The use of icons to indicate security is not limited to binary labels. For example, the UL's IoT security label employs a 'shield' icon in addition to their precious metals-based rating (UL, 2020). A shield as an indicator of IoT security has also been proposed by Public Knowledge, though they stopped short of proposing an actual icon (Knowledge, 2019). Shields, like locks, invoke a mental model of physical security. However, they also differ. While locks can be associated with a mental model of criminal behavior, shields are more aligned with a warfare mental model.

In addition to the icon, graded labels also have to pick the framing of their rating. As noted previously, the UL chose a rating based on precious metals. This employs a market-based mental model, i.e. as the professed security of the device goes up, so does its value, as indicated by the more expensive precious metals (UL, 2020). Even the least secure product then has some inherent value.

In contrast, Johnson et al. use a A-G grade based rating (Johnson et al., 2020). This is not grounded in any mental model or metaphor that the consumer can reference. However, the grades are combined with color codes from green to red, for A-G, respectively. Thus, green is the 'best' choice while red is the worst. This is again grounded in the notion of physical safety, where we often correlate green with safe and red with danger. Unsurprisingly, participants in the study expected a discount to purchase the red-labeled/G-grade IoT devices.

Descriptive labels, unlike graded labels, have to provide more details on the security of the IoT device. However, to balance awareness with usability, the label should only display the most relevant information that is relevant to a non-expert consumer. The label proposed by DCMS, for example, only conveys two pieces of information (DCMS, 2019): 1) whether important security features are included and 2) the duration of time for which the manufacturer will provide security updates. Johnson et al. extends this to two additional properties: 1) whether the device can connect directly to the Internet and 2) whether the information collected by the device is shared with third parties (Johnson et al., 2020). Emami-Naeini et al. in their study identify five items: 1) firmware version, 2) update date, 3) whether the security updates are automatic, 4) how long the security dates will be available, and 5) information about access control (Emami-Naeini et al., 2020). There is then little consensus on what are the must-have pieces of information that should be communicated to a non-expert to inform their device purchase based on security.

Even when researchers agree, e.g. duration of availability of security updates, this information can be presented using different language, such as duration of availability of security patches. A security update aligns more closely with the market mental model, whereas a security patch indicates a medical mental model. Furthermore, lifetime must be defined from either the time of production or the time of purchase. Additionally, once the device is past its update lifetime should it be discarded, will the manufacturer reduce functionality, or is the device expected to be only used in contexts with lower risk?

The individual risk items on any descriptive label can be combined with other risk indicators. For example, Shen et al. use colors to indicate the risk level of the specific information item (Shen and Vervier, 2019). The use of colors leverages affect heuristic (Garg and Camp, 2013); red is typically associated with 'bad' in a western context. These colors can then provide a shortcut for consumers to assess the risk of a device. But these shortcuts may not be appropriate for all information. Shen et al. label the ability of a device's Internet connection with a *yellow* color indicating 'caution' (Shen and Vervier, 2019). Yet, given that these labels are being generated for *Internet* connected devices, the presence of such connectivity can not reasonably be expected to be a cause for caution.

Aside from affect, other heuristics may also inform the effectiveness of a label design. For example, Shen et al. use a shield icon to indicate a category of device factors rather than the presence of appropriate security capabilities (Shen and Vervier, 2019). However, the *availability heuristic* may inform a different user understanding of this icon (Garg and Camp, 2013). Thus, users may conflate it with good security.

5 IMPLICATIONS FOR IoT SECURITY LABELS

IoT security labels have been proposed as one solution to the security information asymmetry between manufacturers and consumers, to thereby avoid a market of security lemons. Well designed labels may help make consumers choose products with a security exposure commensurate with their risk tolerance. Simultaneously, manufacturers may be able to differentiate their products on security and thereby charge a premium for greater security assurance. However, designing security labels that satisfy these criteria is non-trivial. Sections 2-4 detail three primary challenges. First, there has to be commonly accepted definition of 'security'. Second, stakeholders must determine the type of label that is appropriate for the context. Third, designers must identify the security information that should inform consumer behaviors as well as a presentation that makes it usable for non-expert consumers.

There is preponderance of security best practices, baselines, and standards in IoT (Copper Horse, 2018; Fagan et al., 2020; C2, 2019). This creates multiple challenges for manufacturers. They must 'define' security for their use case. This may include following one specific guidance or even a combination of multiple sources. This choice will be mediated by applicable regulatory mandates, which may vary or even conflict across jurisdictions. Furthermore, there will be a direct as well as an indirect opportunity cost to designing a product against a specific definition of security. If these definitions do not persist over an acceptable time frame, it will result in loss for the manufacturer. An example of this is the conclusion of Mozilla's Trustable Technology Mark. Manufacturers may have designed products to show compliance against TTM's specification (Trustable Technology Mark, 2020). However, they will no longer be able to use that to communicate security to their customers.

One solution may be for sector-specific trade groups to define baselines for individual products or even the sector itself. This may be particularly effective in mature sectors with long standing trade groups. They can bring to bear the combined expertise of the various constituent companies, while ensuring that the failure of any one company will not result in a defunct labeling scheme, assuring sustainability. Furthermore, trade groups will be able to assure consistent labeling within individual classes of IoT offerings. Different labels for the same device type for the same deployment context will likely make it challenging for consumers to compare them.

There are multiple extant examples of IoT security certifications from sector specific trade groups. For example, the trade association for the wireless communications industry in United States, i.e. CTIA, has an IoT certification program (CTIA, 2020). For now these certifications have not resulted in corresponding labels. In fact, of the multiple IoT security labels described in Section 3, none are based on a broader security specification. This argues the difficulty of translating technical security guidance aimed at engineers and experts into a public-facing label aimed at non-experts.

Ignoring the technical underpinnings, label design must at least be informed by user experience. Binary labels are more usable for consumers. Blythe et al. note that consumers would prefer security specific labels (Blythe and Johnson, 2018). This approach has, for example, been adopted by Tietoturvamerkki. However, given that the effectiveness of binary labels is driven by awareness, any new security specific labels may need to be complemented by public education campaigns. Additionally, there is need to understand whether consumer mental models align with the symbols currently being used to communicate security, i.e. lock and shield. These icons have not accurately communicated security exposure in other security contexts.

The alternative is to piggyback on existing indicators of trust, with security being appended to make explicit the property addressed by the label. This is the approach adopted by BSI. In this case security becomes one of the many other indicators of product quality. The latter may be easier to understand for consumers with low levels of technical literacy. Furthermore, this approach may be more effective where lack of security does not impact the customer themselves. Anecdotal evidence posits that reframing security as quality has led to software developers paying more attention to the former. Unfortunately, there is but one study that examines the impact of binary IoT security labels.

Beyond binary labels, graded and descriptive labels result in additional complexity. For graded labels, devices at the middle or bottom of the grade scale may be ignored or penalized by consumers. Furthermore, given the difficulty of defining security, it is more challenging to define multiple levels of security. Ideally, these levels should be externally validated in some manner, assuming that they are an indication of security risk exposure. Alternatively, they can be determined in terms of technical maturity. For example, it is reasonable to state that AES256 is more resilient than AES128. Regardless there should be clear distinctions between the different security levels, so that the average consumer can understand the differences. From a design perspective it may be better to have fewer levels to make the distinctions clearer.

For the descriptive labels discussed in this paper, effectiveness is correlated with the perceived risk of the device. As these label types are more information-rich there is a temptation to add more design elements, such as color and icons. However, a simpler label with fewer design elements may be better than one with more. Every additional element incurs a decision point that needs additional investigation to understand its impact on both consumer understanding and corresponding behaviors. While none of the studies covered in this paper address it, consumer understanding and behaviors may vary based on demographic factors, technical literacy, and prior risk experience. Thus, as labels increase in complexity, design becomes more challenging.

Even when communicating limited elements, it is important to consider the usability as well as the whether the specific element is meaningful. One often-argued factor is the lifetime of software patches. However, it is unclear what happens if a company that advertised a long patch lifetime goes out of business. Alternatively, a company may release patches but choose not to address a specific vulnerability. Simultaneously, the same device factor may be made more or less usable for non-experts. For example, merely stating the firmware version on the label may be less usable than also declaring whether the firmware version was the latest version at the time of manufacturing.

Overall, there are many factors to be considered when designing an IoT security label. There are few studies that have studied these labels in detail and only one that we could find that compares multiple labels. Thus, the path to an effective label design remains unclear.

6 CONCLUSION

This paper surveys the current research in IoT security labels. Overall, there is only one study that compares different types of labels. This indicates the need for significant future investigation. Barring that, label designs should be kept simple with low levels of complexity. For binary labels it may be useful to persist with an existing label and extend it to security. For graded labels, designers should mitigate the possibility of low to medium tier devices being penalized. Finally, the descriptive labels may not be effective for low risk devices (or those that are perceived as such). Thus, the use of descriptive labels should be limited to high risk devices.

Manufacturers must start by differentiating between existing security best practices, baselines, and standards. Next these have to be translated into an appropriate label. Trade groups may be able to assist manufacturers in both areas, while ensuring labeling consistency and sustainability. As IoT devices differ significantly across risk - a car is not a camera, and as such different types of labels may be needed for distinct classes of devices.

Information asymmetry between IoT device manufacturers and consumers threatens to create a market of security lemons. One solution is IoT security labeling. However, poor designs will simply result in a lemons market of security labels. Once poorly designed labels have been deployed they may harm consumer confidence in security labeling for the long term. Thus, it is imperative that IoT security label design is addressed with the same diligence as technical aspects of IoT security.

ACKNOWLEDGEMENTS

I would like to thank Mark Walker of CableLabs and Prof. Eric Jardine from Virginia Tech for reviewing early drafts of the paper. The author, however, is solely responsible for any errors in this paper.

REFERENCES

Anderson, R. and Moore, T. (2006). The economics of information security. *Science*, 314(5799):610–613.

Andrews, J. C., Burton, S., and Kees, J. (2011). Is simpler always better? consumer evaluations of front-of-package nutrition symbols. *Journal of Public Policy & Marketing*, 30(2):175–190.

Baldini, G., Skarmeta, A., Fourneret, E., Neisse, R., Legeard, B., and Le Gall, F. (2016). Security certification and labelling in internet of things. In *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*, pages 627–632. IEEE.

Bellman, C. and van Oorschot, P. C. (2020). Best practices for IoT security: What does that even mean? *arXiv preprint arXiv:2004.12179*.

Blythe, J. and Johnson, S. (2018). Rapid evidence assessment on labelling schemes and implications for consumer IoT security. Technical report, DCMS: London.

BSI (2018). BSI launches kitemark for Internet of Things devices. Technical report, British Standards Institute.

C2 (2019). The C2 consensus on iot device security baseline capabilities. Technical report, CSDE.

Camp, L. J. (2009). Mental models of privacy and security. *IEEE Technology & Society magazine*, 28(3):37–46.

Communicating Upgradability and Improving Transparency Working Group (2017). Communicating iot device security update capabilityto improve transparency for consumers. Technical report, NTIA.

Copper Horse (2018). Mapping security & privacy in the Internet of Things. https://iotsecuritymapping.uk/.

CTIA (2020). Cybersecurity certification program for IoT devices. Technical report, CTIA.

Cyber Security Agency of Singapore (2020). Cybersecurity labelling scheme. https://www.csa.gov.sg/-/media/csa/documents/cos/2020/csa-cos-media-factsheet_cybersecurity-labelling-scheme.pdf.

DCMS (2019). Mandating security requirements for consumer 'IoT' products. Technical report, DCMS.

Dingman, A., Russo, G., Osterholt, G., Uffelman, T., and Camp, L. J. (2018). Good advice that just doesn't help. In *2018 IEEE/ACM Third International Conference on Internet-of-Things Design and Implementation (IoTDI)*, pages 289–291. IEEE.

Ditigal Standard (2020). Digital standard. https://www.thedigitalstandard.org/the-standard.

Emami-Naeini, P., Agarwal, Y., Cranor, L. F., and Hibshi, H. (2020). Ask the experts: What should be on an IoT privacy and security label? *arXiv preprint arXiv:2002.04631*.

Fagan, M., Megas, K. N., Scarfone, K., and Smith, M. (2020). IoT device cybersecurity capability core baseline. Technical report, NIST.

Garg, V. and Camp, J. (2013). Heuristics and biases: implications for security design. *IEEE Technology & Society Magazine*, 32(1):73–79.

Harris Interactive (2019). Consumer internet of things security labelling survey research findings. Technical report, Harris Interactive.

IoT Alliance Australia (2020). Internet of Things security guideline. Technical report, IoTAA.

IoT Security Foundation (2020). Best practice user mark FAQ and terms of use. Technical report, IoT Security Foundation.

Ipsos Public Affairs (2019). Product security: IoT and other internet enabled devices. Technical report, Centre for International Governance Innovation.

Jameison, A. (2016). IoT security: It's in the stars. https://www.slideshare.net/AndrewRJamieson/iot-security-its-in-the-stars-169-v201605241355.

Johnson, S. D., Blythe, J. M., Manning, M., and Wong, G. T. W. (2020). The impact of IoT security labelling on consumer product choice and willingness to pay. *PLOS ONE*, 15(1):1–21.

King, A. and Gallagher, R. M. (2020). Cyberspace solarium report. Technical report, CSC.

Klopp, P. and MacDonald, M. (1981). Nutrition labels: an exploratory study of consumer reasons for nonuse. *Journal of Consumer Affairs*, 15(2):301–316.

Knowledge, P. (2019). Security shield. Technical report, Public Knowledge.

Momenzadeh, B., Dougherty, H., Remmel, M., Myers, S., and Camp, L. J. (2020). Best practices would make things better in the IoT. *IEEE Annals of the History of Computing*, 18(04):38–47.

Morgner, P., Mai, C., Koschate-Fischer, N., Freiling, F., and Benenson, Z. (2020). Security update labels: Establishing economic incentives for security patching of iot consumer products. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 346–363.

National Cyber Security Center (2019). Tietoturvamerkki. https://tietoturvamerkki.fi/.

Odom, M. D., Kumar, A., and Saunders, L. (2002). Web assurance seals: How and why they influence consumers' decisions. *Journal of Information Systems*, 16(2):231–250.

Police Crime Prevention Initiatives (2020). Secure by design. https://www.securedbydesign.com/.

Rothman, R. L., Housam, R., Weiss, H., Davis, D., Gregory, R., Gebretsadik, T., Shintani, A., and Elasy, T. A. (2006). Patient understanding of food labels: the role of literacy and numeracy. *American journal of preventive medicine*, 31(5):391–398.

Shen, Y. and Vervier, P.-A. (2019). IoT security and privacy labels. In *Annual Privacy Forum*, pages 136–147. Springer.

Smith, M. W. (2019). Information asymmetry meets data security: The lemons market for smartphone apps. *Policy Perspectives*, pages 85–96.

Trustable Technology Mark (2020). The trustable technology mark is wrapping up. https://trustabletech.org/the-trustable-technology-mark-is-wrapping-up-\\trustable-technology-lives-on/.

UL (2020). IoT security rating level. https://ims.ul.com/iot-security-rating-levels.

Waechter, S., Sütterlin, B., and Siegrist, M. (2015). Desired and undesired effects of energy labels—an eye-tracking study. *PloS one*, 10(7):e0134132.

Wu, J. and Zappala, D. (2018). When is a tree really a truck? exploring mental models of encryption. In *Fourteenth Symposium on Usable Privacy and Security ({SOUPS} 2018)*, pages 395–409.