Vulnerabilities and Responsibilities: Dealing with Monsters in Computer Security

W. Pieters[1] and L. Consoli[2]
[1]Centre for Telematics and Information Technology, University of Twente, The Netherlands
[2]Institute for Science, Innovation, and Society, Radboud University Nijmegen, The Netherlands

Purpose of this paper	In this paper, information security assessment is analysed in terms of cultural categories and virtue ethics, in order to explain the cultural origin of certain types of security vulnerabilities, as well as to enable a proactive attitude towards preventing such vulnerabilities.
Design/methodology/approach	Vulnerabilities in information security are compared to the concept of "monster" introduced by Martijntje Smits in philosophy of technology. The applicability of different strategies for dealing with monsters to information security is discussed, and the strategies are linked to attitudes in virtue ethics.
Findings	It is concluded that the present approach can form the basis for dealing proactively with unknown future vulnerabilities in information security.
Research limitations/implications (if applicable)	The research presented here does not define a stepwise approach for implementation of the recommended strategy in practice. This is future work.
Practical implications (if applicable)	The results of this paper enable computer experts to rethink their attitude towards security threats, thereby reshaping their practices.
What is original/value of paper	The paper provides an alternative anthropological framework for descriptive and normative analysis of information security problems, which does not rely on the objectivity of risk.

Keywords: cultural categories, information security, monster theory, risk, virtue ethics, vulnerabilities

1 Introduction

Information security can be defined as the practice that deals with protecting information systems against attacks. It typically concerns technical means to avoid or eliminate security vulnerabilities in information system design. Such vulnerabilities form a special type of risk.

Rather than covering events in the natural environment, these risks are associated with intentional human action: an attacker may exploit a vulnerability in an information system to disrupt the confidentiality, integrity or availability of information.

As vulnerabilities can be classified as risks, the problems associated with risk assessment also apply to information security. How can we anticipate risks in the design of the technology, if we may never have seen similar risks before? How likely are the associated future events? Can risk be expressed objectively at all? Such philosophical questions about risk have been around for a while (see e.g. Shrader-Frechette, 1990; Cross, 1992; Jasanoff, 1998). These questions have also been discussed in the context of a critique of classic theories of technology assessment, where it is pointed out that either the future has too many variables to allow for accurate (or even approximate) predictions of the future, or that models fail to take into account the complexity of the social embedding (Adams, 1995; Bradbury, 1989; Pursell, 1979).

With the emergence of the new field of information security, the objectivity of risk and security is still often implicitly assumed (Pieters, 2006; see for example Evans and Paul, 2004; Nikander and Karvonen, 2000; Oostveen and Van den Besselaar, 2004; Riedl, 2004; Xenakis and Macintosh, 2005), but similar questions can be asked about the uncertainty involved in assessing these variables. The possibility of *intentional* behavior by an attacker to make the system fail - characteristic of security as opposed to safety - increases the uncertainty in the case of security, because predicting what humans may do is thought to be different from predicting what nature may do. There is a double contingency, and the vulnerabilities game is an act of communication between attackers and defenders (Luhmann, 1995, pp. 103-136).

Such discussions about the objectivity of risks have consequences for the possibility of developing an information security ethics. They make a consequentialist approach problematic, since the high level of uncertainty involved makes it easy to deny the relation between action and consequence, and thereby to avoid responsibility. Security assessment is nevertheless an inescapable necessity, and it is therefore necessary to put forward a framework which makes it possible to describe security practices and ethics without relying on the objectivity of risk assessment. At the same time, the framework needs to be proactive, meaning that the model should not only be useful to assess and explain past problems, but also suitable for preparing for future attacks.

In this paper, we will make the case for virtue ethics as suitable framework. We investigate the scientific endeavor of information security as a cultural phenomenon, and from this perspective, our main question is how participants in this community can and should deal with uncertain risks. We employ existing theories from the anthropological perspective to develop a descriptive as well as a normative framework to discuss the attitudes towards vulnerabilities in the information security community, and illustrate the approach with examples from the field.

Rather than describing which types of flaws exist, which can also be helpful in analyzing the domain of information security (see e.g. Landwehr et al., 1994), we focus on why humans were unable to prevent them. We will introduce the concept of “cultural category” as a useful basis for understanding the difficulties of security assessment. We will then show how certain types of vulnerabilities can be interpreted in terms of clashes of categories. This allows us to give an explanatory account of such vulnerabilities, and to address the question of how we should handle (moral) responsibility and accountability within information security from the point of view of virtue theory. To achieve this, we will describe how the emerging “monsters” can be dealt with, and how these dealing strategies can be linked to character traits in the context of virtue theory.

2 Vulnerabilities as monsters

Cultural categories are classifications that help us describe and understand the world. The concept has been used by Mary Douglas in her anthropological analysis of impurity and danger (Douglas, [1966] 1994). Impure and dangerous are those phenomena that do not fit into a

cultural scheme of classification of the beings in the world. In this sense, dirt is “matter out of place” in the system of cultural categories. Some African tribes considered twins as monsters, because, according to these categories, only animals produced more than one child, and twins thus had both human and animal traits. These “monsters” can come into being when cultural categories are inadequate to fit phenomena. Douglas held the view that this analysis applies to modern society as well as traditional societies.

Martijntje Smits (2002a; 2002b; 2006) argues that controversies surrounding the introduction of new technologies can often be explained in terms of such a clash between cultural categories. Here, again, danger and impurity are associated with a classification failure. For example, we may think of genetically manipulated food as an (unacceptable?) mixture of nature and culture. Smits argues that the concept of “monster” can be applied to controversies on new technologies as well, and proposes the term “monster theory” for a method of analysis of these issues. Smits argues (2002b, p. 143, our translation):

> From the monster theory it follows that waste and dangers are inevitable, because they are unintended by-products of cultural category classifications. On the borders of these classifications, ambiguities appear, that may, among other things, manifest themselves as monsters.

The latter happens when the ambiguity is experienced negatively, and cannot be resolved easily. The term “monster” can be understood by reference to monsters in stories and films, which often combine elements of different categories as well (Csicsery-Ronay, 2002).

We propose to generalize the concept of “cultural category” by using it not only in the broad context of the whole of society, as Smits does, but also in reference to practices pertaining to specific groups or “subcultures”. By “subculture” we understand in the context of this paper a group that can be identified as sharing some basic characteristics, like for example religious beliefs, moral standings, or professional occupation, in other words properties that determine the generalized “cultural categories” that they share. This allows us to apply the monster theory to the specific domain of information security.

As an aside, we should note that the way in which we use the concept of cultural category bears a passing resemblance to the “paradigm” idea developed by Kuhn (1964). However, while a paradigm shift is a rather dramatic event, in which the accumulation of anomalies leads to the abrupt replacement of a world view during a “revolutionary” period, the way in which categories react to monsters does not require or imply such a mechanism. This will become clear in the following sections.

In our case we focus on information security and propose that there is a strong analogy between the model proposed by Smits and the inherent fallibility of security assessment: security assessment can also be explained in terms of cultural categories, and phenomena that do not fit in these categories. Spectacular attacks on computer systems, having a major impact on the practice of security, often occur when the vulnerability does not fit into the existing *categories* of computer security. Therefore, we argue that at least some of the vulnerabilities emerging in computer security can be characterized as monsters. As much as society will always produce waste and dangers because of existing categories, computer security will always produce vulnerabilities because of existing security models.

On an even smaller scale, the cultural categories within a company may produce vulnerabilities in their information systems, when phenomena are not covered by the cultural categories of the company. To the outside world, the phenomenon appears as a failure, a mistake by the company in relation to the existing categories. However, from the company's perspective, the vulnerability may present a true challenge to their categories, and therefore appear as a monster. The company may not only have failed to address the vulnerability, *it may not even have been able to see it*, because it did not fit in their categories.

This analysis does not imply that the first attack of a certain type is always the most successful. It takes time to adapt when a challenge to the cultural categories is presented, and

therefore new attacks may be very similar to ones that were already known, but still be very successful (Kienzle and Elder, 2003). It is also important to stress that from this analysis it follows that "mistakes", understood as errors due to poor judgment or misapplication of knowledge, such as programming errors, are *not* to be categorized as monsters. These do not result from an inherent impossibility within the subculture to make sense of the phenomenon with their categories. Instead, they reveal a misapplication of existing categories. This also holds for errors due to time constraints, and ultimately money (Anderson, 2001). Something is only a "monster" if it is a classification failure. In the following, we will focus on such monstrous vulnerabilities.

3 Two examples of the clash of categories

A typical example of a clash of categories in computer security was the issue of viruses in Microsoft Word documents (Ford, 1996; Gordon and Ford, 1995). Up to a certain point in time, viruses were supposed to hide in executable files only, not in documents. Then, a virus was created that was capable of affecting Microsoft Word documents: the Concept virus. It was relatively harmless and meant to demonstrate the possibility of macro virus creation (Wallich, 1995). A more recent and infamous example of virus exploiting macros in Word is the Melissa virus (Garber, 1999).

It is interesting to note that both Garber and Wallich identify these viruses as barrier-breaking and radically new objects. In our framework, we can affirm that the viruses in Word documents were a clever example of the mixing of two cultural categories in computer science: those of programs and data. Virus scanners were supposed to focus on programs, not data. As a result, a vulnerability emerged that was not recognized as such: it was "matter out of place" in the category system. A monster had been created. The "success" of Melissa shows that monsters may continue to exist, even after they have been unveiled years ago.

The observations that categories may be imperfect and that many people are aware of their limitations do not invalidate this argument. Even if people are aware of the limitations of the classes in the animal kingdom, the platypus may still be considered a "monster" when it is first discovered. In the same sense, computer scientists may be well aware of the limitations of the distinction between programs and data, but still develop their products from the perspective of this separation, and therefore be shocked when something shows up that challenges these classes. This is because even when people are aware of the limitations of the categories, they still need them as a basis for communicating in their subculture. And even if the developers of the macro-using software are aware that this feature constitutes a risk, the risk may not be perceived in the same way by the virus protection developers. Following the monster theory, any classification in computer science that affects or models security is therefore bound to create vulnerabilities as by-products that exploit the limits of the classification. In this sense, the conceptual separation of programs and data produced the text document viruses.

A second example is the separation of the hardware level and the software level in smart-cards. Normally, security models addressed either the software or the hardware, but not both. This enabled the power analysis attack, in which data (software) could be read by eavesdropping on the power consumption of the card (hardware) (Messerges et al., 2002). In one implementation, attacks are based on side-channel information gained by observing cache bits and misses in the current drawn by the smart card (Fournier and Turnstall, 2006). This, in turn, generated research on how to repair the vulnerability (Herbst et al., 2006). Tempest attacks, involving eavesdropping on information by capturing electromagnetic radiation emitted by electronic devices, are of a similar type. Although the possibility is well-known in military circles, it had not been thought of in the Dutch electronic voting regulation (Pieters and Van Haren, 2007). A British report on electronic voting did mention it (Fairweather and Rogerson, 2002).

An interesting question is who is responsible for such clashes. In case of the viruses in Word documents, was it Microsoft, which allowed macros to be executed in Word documents? Was it the virus writer, who exploited this feature in order to attack systems? Or was it the computer

science community, whose classifications were not suitable for all types of files? Different levels of responsibility can (and should) be identified. Of course, from a legal perspective, the virus writers are responsible for the problems involved. The Melissa virus writer has some moral responsibility as he willingly chose to write code for malignant purposes (Bissett, 2000), but this does not mean that we cannot discuss other responsibilities from an ethical perspective here. Also in the case of power analysis, we can pose the responsibility question in a way parallel to the Melissa example.

In both examples we can thus identify several levels of responsibility. We propose that, on a more fundamental level, the cultural categories *themselves* are responsible for providing the opportunities for attack. The community was inherently incapable of preparing itself for the new attacks, because their cultural categories *excluded* precisely the phenomenon that appeared, just as a categorization of animals may exclude the platypus. Such vulnerabilities can be understood as monsters. This observation allows us to put forward an approach for understanding the moral responsibilities and ethical attitudes of the security experts involved, in terms of their attitudes towards challenges to existing cultural categories.

4 Monsters and responsibilities

Based on the analysis and the examples of the previous sections, we aim at linking the descriptive analysis of practical reactions to and decisions on vulnerabilities, provided by application of the monster theory, with moral traits and moral responsibilities, yielding a normative view on the daily practice of computer security experts. We will first discuss the ethical approach taken.

The issue of moral responsibility of IT professionals is usually framed in the literature by discussing the issue whether computers create novel moral issues or, rather, reframe existing issues (Johnson, 2003; Johnson and Nissenbaum, 1994). Specifically referring to computer security, Nissenbaum and Felton (2002) have argued that new problems arise. Floridi (1999) focuses on the specific moral issues associated with information entities. Here, however, we are interested in the information security expert rather than the specific properties of the IT environment.

In information security, practitioners are dealing with situations with a high level of uncertainty. Moreover, the specific problems we focus on in this paper occur precisely when the existing cultural categories are insufficient to express the situation. In such settings, not all ethical frameworks are suitable. Approaches like deontology or utilitarianism take the agent as given and concentrate on other factors such as moral laws or consequences, which may be impossible to assess in such situations. By contrast, virtue ethics focuses precisely on character traits and dispositions of the agent, and can be used to define attitudes towards problems that involve a high degree of uncertainty.

For this reason, we express the link between reactions to monsters and ethical attitudes in terms of virtue theory (see Bowen (2000) for an example of an application in the context of computer science). In this approach, the focus lies not so much in a theoretical framework that can be used to analyze moral choices, but on attitudes of character (virtues and flaws), aiming at describing how to “live a good life” and be a “good human being”. The basic question we should be asking when dealing with ethical matters is for virtue ethics not: “What ought I do?”, but rather: “What kind of person ought I be?” One of us has argued elsewhere (Consoli, 2008) that professional conduct issues might greatly benefit from a virtue-theoretical approach, in that it helps eliminate the mismatch between externally imposed rules and the living morality of the community of practitioners.

Before proceeding, we must specify which model of virtue theory we have in mind. Virtue theory goes as far back as Aristotle (see Bowen, 2000), but has been revived in recent times by a number of authors (Anscombe, 1958; Foot, 2001; Hursthouse, 2006; Nussbaum et al., 1993) in very different contexts. We are concerned here with the formulation due to Alasdair MacIntyre (MacIntyre, [1981] 2007). There are two main reasons for choosing this approach in

our attempt to link strategies for dealing with monsters with moral attitudes of security experts.

Firstly, MacIntyre introduces the notion of *practice* (Macintyre, [1981] 2007, p.187):

> By a practice I am going to mean any coherent and complex form of socially established co-operative human activity through which goods internal to that form of activity are realised in the course of trying to achieve those standards of excellence which are appropriate to, and partially definitive of, that form of activity, with the result that human powers to achieve excellence, and human conceptions of the ends and goods involved, are systematically extended.

This notion of practice provides a way to link the activities within a subculture (as a practice with both individual practitioners and their communal interests and standards) directly with its moral significance, which is our goal. In fact, the activities characteristic of the practice are explicitly oriented at achieving a moral goal, namely the "standards of excellence". Furthermore, both the individual and the community level are considered: a practice is a form of "socially established co-operative human activity" and only a common effort of the practitioners van lead to the extension of the "human powers".

Secondly, MacIntyre looks for a way to give a "unitary core concept" of virtues (MacIntyre, [1981] 2007, p. 186). On the one hand, he says, we must recognize the relativity of the single virtues (or even what we take as constituting a virtue) on the historical and social context; on the other hand, the different account of virtues share (form) a tradition which has a conceptual unity. This approach provides the notion of virtue with a conceptual robustness that makes the situation-boundness of our analysis (even more so in a fast-changing field as the one of information technologies) not a weakness, but a reliable starting point to investigate moral responsibilities. In other words, the link we make between strategies and virtues does not purport to be absolute, but does claim to be embedded in a tradition which makes it at the same time recognizable and usable (because of its recognizability).

An added value of choosing virtue theory as our analysis tools is that, besides describing the situation, it can also provide the proactivity we referred to earlier in the paper. In other words, by linking monster-dealing strategies with virtues, it becomes possible for the security expert to be prepared for the future, and to deal efficiently with challenges that must yet be fully conceptualized.

5 Strategies and virtues

Smits (2002a; 2002b; 2006) considers four different ways of dealing with monsters: embracing, expelling, adapting, and assimilating. In this section, we investigate the extent to which they can be used for dealing with vulnerabilities in computer security as well. We also address the question how these strategies relate to basic ethical attitudes of the experts as individuals and the community as a whole. This allows us to evaluate the desirability of the different combinations of strategies and attitudes.

Before proceeding, we wish to stress the explorative character of the following analysis, and the need for more specific further research, as to the knowledge of the authors there has been very little research carried out in the field of virtue theory applied to computer security. Furthermore, the examples do not always represent the attitude of the community as a whole, as subcultures exist also within the information security community. In particular, there is often a gap between the attitude of industry professionals and the scientific community, especially when a company's view is mediated by its economic interests. The examples are thus only meant to explain *possible* attitudes.

Embracing as respect

Embracing the monster (for example, as if it were a wonder) may be interpreted as a sign of respect. During the introduction of plastics, some people thought this new material would be some kind of salvation from the limitations of nature (Smits, 2002b, p. 161). Thus, there is the possibility of admiring phenomena incompatible with the existing cultural order, and granting them a kind of holy status.

In computer security, this means granting a vulnerability, an inconsistency, the status of a kind of ultimate proof of the rightness of the existing order. In a way, this can be seen in the reaction of a Dutch voting machine manufacturer (Nedap) to the easy replacement of chips in their system by a pressure group: “We noticed that it was proved that the machine works excellently. The voting machine does exactly what is commanded.” Thus, if the attackers make the system count incorrectly, this proves that the system is correct. This is the “it's not a bug, it's a feature” approach, often induced by commercial interests: it is not a monster, it's a normal animal. It is not a problem with our categories; it is a problem with *your* categorization of the phenomenon.

This attitude is a form of respect, in which something is accepted without questioning. Respect can be useful in many situations, since it is impractical to question everything one encounters. As such, respect helps to maintain the existing category system, which helps us to deal with our environment in a practical way. The question is if this has to be considered a virtue for a computer security expert. Seeing a vulnerability as a confirmation of the existing order does not solve the problem, at least not from the point of view of the people who see it as a failure. Because security is inherently occupied with future events and unknown phenomena, respect for the monster stands in the way of dealing with it in a way which is effective for the security practice.

We can conclude that embracing is not a “virtuous” attitude for the community of practitioners and should therefore be discarded as a way of conduct. It is interesting to notice how, in the context of our analysis, the notion of respect – which is almost universally considered a virtue – turns out not to be a favorable or even positive attitude. In that sense, we see how the tradition within which virtues are embedded is constitutive of the virtues themselves. What is highly commendable in a situation, must be steered clear of in another one.

In fact, respect as attitude may happen (and even be appropriate) among hackers. They see a vulnerability they discover as a confirmation of their own place in the order of things, which is a “monstrous” place: a place which constantly seeks the border of existing categories (Nissenbaum, 2004). In *this* context, respect is not only commendable, but it also contributes to the self-definition of hackers.

Expelling as denial

Some people regarded plastics not as salvation, but as a disaster. Precisely the failure to fit existing classifications of materials made them filthy and dangerous. Luddites think such new technologies should be expelled or even destroyed (Smits, 2002b, pp. 149-150).

Expelling the vulnerability-as-monster in computer security is in a practical way not feasible, because a threat to a computer system cannot be eliminated as easily as a new phenomenon in society, since the attacker is typically outside the control of the computer security community. There is an extra contingency here in the behavior of the enemy, who is committed to exploiting the monster. Where technological monsters in society are expelled by saying “we don't want this”, vulnerabilities-as-monsters are expelled by saying “there is no problem”. Diebold Election Systems used this strategy when they were accused of vulnerabilities in their voting machines (Gumbel, 2005, p. 260) as opposed to the Nedap reaction above. It might help to deny the problem and see if everything stays quiet, particularly because public perception of vulnerabilities makes it more likely that they are exploited.

Although the attitude of Diebold can not be considered typical for the behavior of security

experts (they are a commercial party with an economic interest), it makes nonetheless clear that such a take on the vulnerabilities problem cannot possibly be productive for a proactive attitude in the security community. From an ethical point of view, this disposition can be interpreted as one of denial towards the problem, as a way of not having to commit oneself to a choice and a course of action. Again, denial may be appropriate in situations where one has too many problems to deal with all of them, or even when the definition of a problem seems to be unreasonable. As in the case of respect, this attitude helps to maintain the existing categories, because problems with these are put aside. However, as dealing with future problems is constitutive of the information security practice, denial cannot be considered a virtue in this context. Context-dependence plays again a crucial role, exactly as one would expect from a virtue-ethical approach. In the case of computer security, leaving a monster arising from a category clash alone is often not only unfeasible, but counter-productive.

Adaptation as perseverance

Biodegradable plastics are an adaptation of the monstrous plastics to existing categories: they will no longer fail to rot when lying around (Smits, 2002b, p. 155; 2006, p. 501). In the strategy of adaptation, monsters are redefined such that they are compatible with existing categories.

Adapting vulnerabilities-as-monsters may be useful as well, for example by categorizing Word documents as executable files rather than data files, which was done in virus scanners. In such an approach, the threat becomes one of a known category: a virus in an executable file. In other words, the category stays fixed and the object is re-categorized.

From a virtue ethics point of view this reaction can be categorized as a form of perseverance, which denotes a) continuing a process and b) doing so despite difficulties. In this particular case, the difficulties are the monsters, and perseverance means trying not to compromise too much one's own world view by forcing the world to fit in it. Perseverance accepts that the existing categories may be imperfect, but refuses to change them based on new phenomena. Instead, it is said that “y is just another x”. Even in this case, perseverance is not inherently “bad”, because stable categories exist to create stable communities, and allow members to learn to avoid mistakes in their application. It could therefore be argued that such a strategy could be useful for information security experts in dealing with vulnerabilities. However, one of the priorities of information security is prevention and, in that sense, perseverance as a disposition does not seem to provide the proactive attitude which is required in order to prevent new security threats.

Moreover, the adaptation approach presupposes a unidirectional relation between categories and the phenomena they explain. Categories are given, and vulnerabilities have to fit into these categories to allow protection. This does not do justice to the complex and dynamical process in which vulnerabilities emerge in information systems. The approach cannot be generalized to deal with security problems, and we can see this in the virus vulnerability. The challenges of viruses in different file types, in the end, have not left the categories of virus protection unaffected, and we can now do “full scans”, “smart scans”, etcetera. Thus, the categories have been changed as well.

Up to now, we have tried to link three of the monster-dealing strategies with dispositions of character. However, respect, denial and perseverance are all ethical attitudes that are helpful in maintaining existing categories. Although these strategies may be useful to prevent “accidental” vulnerabilities introduced by mistakes or constraints in terms of time or budget, it became clear that these dispositions do not provide the security expert and her community with a useful attitude to deal with the monstrous vulnerability, because in this case it is the categories themselves that are being challenged. Since information security should, by its very nature, be prepared for unknown future threats, they may not even be virtuous within the practice at all. A different strategy is needed.

Integration as open-mindedness and courage

The last strategy Smits mentions is assimilating the monster, a pragmatic process in which both the monster and the cultural categories are being changed. Because the word "assimilation" seems to imply that something is made to fit rather than a mutual convergence process, we rather use "integration" instead. An example that Smits mentions is the shifting of the border between alive and dead due to the technology of organ donation. Here, "brain-dead" became a new criterion for deciding whether it would be allowed to remove usable organs from a body. Thus, a new category emerged for dealing with the new technology (Smits, 2002b, p. 159; 2006, p. 501).

Integration also happens in information security. In vulnerabilities-as-monsters, power analysis attacks on smart-cards now have their own field of research, and the power-analysis vulnerability has changed from a side-effect to something that can be prevented using appropriate tools. This means that both the categories and the technology have been changed, by integrating the monster of power analysis attacks.

From a virtue ethics point of view, the attitude of integrating the monster is related to the epistemic virtue of open-mindedness (Montmarquet, 2008). In order to change one's world view for new phenomena to fit in, such effort of flexibility of thought is required. Only when such effort is made, it becomes possible to change the categories of thinking themselves. This is not sufficient by itself, though, since a motivational attitude is also required to actually use the open-mindedness in a specific context, in this case information security.

The motivation behind the open-mindedness can be related to the virtue of courage. Courage is a motivational virtue, as opposed to open-mindedness, which is an epistemic one (Montmarquet, 2008). In this context, we understand courage as a disposition of willingness to confront uncertainty or danger. As Mary Douglas showed, uncertainty and danger are closely related, precisely through the cultural categories we discussed earlier. Fixed categories provide humans with a sense of safety, and it therefore requires courage to subject those to challenges. In the context of information security, this also means that one may need to reduce one's own (epistemic) certainty in order to increase the security of the information systems one designs, and thereby the certainty of other people.

The virtue of courage may also be extended beyond the level of the individual. Even if individual open-mindedness and courage may allow one to identify challenges to the existing cultural categories, and thereby create the possibility to prevent vulnerabilities, lack of these virtues within the subculture may still prevent action. Thus, although our starting point is epistemic and individual, the fact that knowledge is embedded in the practice of information security requires not only open-mindedness at the individual level, but also the courage to change the practice itself.

Regarding assimilation this way, we are provided with the grounds we need for a proactive attitude towards security, which was missing in the adaptation strategy. Members of the computer security community are not only responsible for formalizing all aspects of existing categories, but rather for contributing to the evolution of the categories themselves, both individually and within the practice, so that they are better able to incorporate new phenomena, and thereby prevent new attacks. But this last step involves a definite amount of courage in the sense we defined above. Moreover, this courage can not be limited to one individual, but must be attained in a collective effort, through which the standards of excellence can be achieved. When acting courageously as a community, then, security experts form a practice in the full sense of the Macintyrean definition.

6 Conclusions and discussion

The treatment of deviant phenomena in a culture is a field of research with a long tradition. Based on the theories of Mary Douglas on impurity, danger and risk, Martijntje Smits analyzes

how our culture takes care of new technological phenomena. She calls this approach “monster theory”. We argue that this theory does not only make sense on a broad cultural level, but also within subcultures. These subcultures have their own sets of specific cultural categories.

We have proposed to adapt the monster theory to the practice of computer security. This enables us to frame the discussion about the meaning of new threats and the way to react to them in terms of strategies for dealing with monsters. The strategies that Smits distinguishes are embracing, expelling, adapting and assimilating. These can be used to describe reactions to vulnerabilities-as-monsters in information systems. Smits considers the assimilation strategy, which we prefer to call integration, the most promising one, since it does not consider the cultural categories as fixed and given. We argue that integration is also the best strategy in computer security as a subculture, for there is no final model of information security that incorporates all vulnerabilities, as there is no final set of cultural categories that fits all phenomena.

We have furthermore proposed to view the different strategies as linked to ethical attitudes in the context of virtue ethics *à la* Macintyre. We believe that this helps to re-conceptualize the discussion about responsibilities of (computer) scientists in a productive way. From a virtue ethics point of view, the strategy of integration requires both epistemic and motivational virtues. A certain amount of both open-mindedness and courage is desirable for members of the information security community to achieve the “standards of excellence ... appropriate to ... that form of activity”. Also within this perspective, integration turns out to be the most “virtuous” strategy. While the virtues associated with the other strategies may be helpful in specific situations, the relation of information security with the future and with attackers makes it necessary to be epistemically adaptive and have the willingness to confront the associated uncertainty. There is currently much research in a broad range of fields about the modern significance of courage as a virtue (Chun, 2005; Kateb, 2004; Schwartz, 2004). This paper aims to be a contribution to this discussion within the field of computer security.

Having made a case for open-mindedness and courage as the most virtuous dispositions security experts can deploy in order to deal with monsters, we are still faced with a crucial question: how can these virtues be translated into practice? How should open-mindedness and courage be operationalized in order to make it possible to go beyond established categories? This paper does not purport to give a stepwise approach for proactively identifying potential monsters, as its scope is to point at an alternative perspective on computer security and to give an exploratory account. Nevertheless, two suggestions can be made, which will have to be worked out in future work. Firstly, great emphasis on virtues should be put at the educational level. Teaching character may seem a strange endeavor, but virtue ethics shows that it is a feasible way (Hartman 2006). Virtuous attitudes can be taught in a biographical way, i.e. by using exemplary tales which are not only descriptive, but can help leverage our knowledge to deal with future situations (Consoli, in preparation).

Secondly, risk assessment methods should somehow reflect the attitudes of open-mindedness and courage. Information security experts already point out the need for “out-of-the-box” thinking in information risk assessment in personal communication. Also, on the lower level of organizational knowledge, many security incidents are related to “unknown unknowns”: phenomena that are not reflected in the organization’s categorization of its assets (Baker et al., 2008). This means that the need for open-mindedness and courage seems to be acknowledged within the community. “Out-of-the-box” may indeed be a good concept to denote precisely this combination of virtues: “out-of-the-box” is associated both with transcending existing conceptualizations and with the motivation that makes this possible. However, incorporating such proactive monster taming in risk assessment methods seems to be a *contradictio in terminis*, since monsters are by definition not included in existing categories, and can therefore not be part of a risk assessment process. How to resolve this paradox is an interesting topic for future research.

As a final remark, the notion of monster can be seen as a tool for information security experts to focus on “what is in there” next to “what is out there”. If information security, and risk assessment in general, aim at better understanding of future phenomena, they cannot ignore

their own culture. For it may not be objective nature that they are investigating, but their own constructed version, if only for the fact that they are dealing with the future.

Acknowledgements

Part of the research described here was performed while the first author was employed by Radboud University Nijmegen and funded by a Pionier grant from NWO, the Netherlands Organisation for Scientific Research. The authors wish to thank Bart Jacobs, Wim Thijssen, Hub Zwart, the participants of the E-CAP 2006 conference and the anonymous referees for useful comments on ideas and earlier versions of this paper.

References

Adams, J. 1995, *Risk*, UCL Press.

Anderson, R. (2001), "Why Information Security is Hard - An Economic Perspective", in *ACSAC '01: Proc. of the 17th Annual Computer Security Applications Conference*, IEEE Computer Society, Washington, DC, p. 358.

Anscombe, G.E.M. (1958), "Modern Moral Philosophy", *Philosophy,* Vol. 33, p. 124.

Baker, W.H., Hylender, C.D. and Valentine, J.A. *2008 data breach investigation report*, Verizon Business, June 2008.

Bissett, A. (2000), "Some human dimensions of computer virus creation and infection", *International Journal of Human-Computer Studies*, Vol. 50, No. 5, pp. 899-913.

Bowen, J. (2000), "The ethics of safety-critical systems", *Communications of the ACM*, *Vol.* 43, No. 4, pp. 91-97.

Bradbury, J. (1989), "The policy Implications of Differing Concepts of Risk", *Science, Technology, and Human Values*, Vol. 14, No. 4, pp. 380-399.

Chun, R. (2005), "Ethical character and virtue of organizations: An empirical assessment and strategic implications", *Journal of Business Ethics*, Vol. 57, No. 3, pp. 269-284.

Consoli, L. (2008), "The intertwining of ethics and methodology in science and engineering: a virtue-ethical approach"*, Interdisciplinary Science Reviews,* Vol. 33, No. 3, pp. 233-242.

Consoli, L. (in preparation).

Cross, F.B. (1992), "The risk of reliance on perceived risk", *Risk*, Vol. 3, pp. 59-70.

Csicsery-Ronay, I. (2002), "On the grotesque in science fiction", *Science-fiction Studies*, Vol. 29, pp. 71-99.

Evans, D. and Paul, N. (2004), "Election security: perception and reality", *IEEE Security & Privacy*, Vol. 2, No. 1, pp. 24-31.

Douglas, M. ([1966] 1994), *Purity and danger: An Analysis of the Concepts of Pollution and Taboo*, Routledge, London.

Fairweather, B. and Rogerson, S. (2002), *Technical Options Report*, available at: http://www.communities.gov.uk/documents/localgovernment/pdf/155484.pdf (accessed 24 March 2009).

Floridi L. (1999), "Information ethics: on the philosophical foundation of computer ethics",

Ethics and Information Technology, Vol. 1, No. 1, pp. 37-56.

Foot P. (2001), *Natural Goodness,* Clarendon Press.

Ford, R. (1996), "Why viruses will always be a problem", *NCSA News*, pp. 5-7.

Fournier, J. and Turnstall, M. (2006), "Cache based power analysis attack on AES", in Batten, L.M. and Safavi-Naini (Eds.), *11th Australasian Conference on Information Security and Privacy — ACISP 2006*, Vol. 4058 of *Lecture Notes in Computer Science*, Springer, pp. 17-28.

Garber, L. (1999), "Melissa Virus Creates a New Type of Threat", *Computer*, Vol. 32, No. 6, pp. 16-19.

Gordon, S. and Ford, R. (1995), "Real-world anti-virus products reviews and evaluations", Proceedings *of Security on the I-WAY*, Crystal City, Virginia.

Gumbel, A. (2005), *Steal This Vote: Dirty Elections and the Rotten History of Democracy in America*, Nation Books, New York.

Hartman, E. M. (2006), "Can we teach character? An Aristotelian answer", *Academy of Management Learning and Education*, Vol. 5, No. 1, pp. 68-81.

Herbst, C., Oswald, C. E., and Mangard, S. (2006), "An AES smart card implementation resistant to power analysis attacks", in Zhou, J., Yung, M. and Bao, F. (Eds.), *Applied cryptography and Network security*, Vol. 3989 of *Lecture Notes in Computer Science*, Springer, pp. 239-252.

Hursthouse, R. (2006), "Are Virtues the Proper Starting Point for Ethical Theory?", in Dreier, J. (ed.), *Contemporary Debates in Moral Theory*, Blackwell, pp. 99-112.

Jasanoff, S. (1998), "The political science of risk perception", *Reliability Engineering and System Safety*, Vol. 59, pp.91-99.

Johnson, D. (2003), *Computer Ethics*, Pearson Education.

Johnson, D. and Nissenbaum, H. (1994), *Computers, Ethics, and Social Values*, Prentice Hall.

Kateb, G. (2004), "Courage as a virtue", *Social Research,* Vol. 71, No. 1, pp. 39-72.

Kienzle, D.M. and Elder, M.C. (2003), "Recent worms: a survey and trends", in: *WORM '03: Proc. of the 2003 ACM Workshop On Rapid Malcode*, ACM Press, New York, pp. 1-10.

Kuhn, T. (1964), *The structure of scientific revolutions*, University of Chicago Press.

Landwehr, C.E., Bull, A.R., McDermott, J.P. and Choi, W.S. (1994), "A taxonomy of computer program security flaws", *ACM Comput. Surv.*, Vol. 26, No. 3, pp. 211-254.

Luhmann, N. (1995), *Social Systems*, Stanford University Press, Stanford, CA.

MacIntyre, A.C. ([1981] 2007), *After Virtue*, University of Notre Dame Press.

Messerges, T.S., Dabbish, E. A., and Sloan, R. H. 2002, "Examining smart-card security under the threat of power analysis attack", *IEEE Transactions on Computers*, Vol. 51, No. 5, pp. 541-552.

Montmarquet, J. (2008), "The Voluntariness of Virtue – and Belief", *Philosophy*, Vol. 83, pp. 373-390.

Nikander, P. and Karvonen, K. (2001), "Users and trust in cyberspace", in Christianson, B.,

Crispo, B., Malcolm, J.A. and Roe, M., *Security Protocols: 8th International Workshop*, , *Cambridge, UK, April 3-5, 2000*. Revised Papers in Vol. 2133 of *Lecture Notes in Computer Science*, Springer, pp. 24-35.

Nissenbaum, H. (2004), "Hackers and the contested ontology of cyberspace", *New Media and Society*, Vol. 6, No. 2, pp. 195-217.

Nissenbaum, H. and Felton, E. (2002), "Computer Security: Competing Concepts", in *The 30th Research Conference on Communication, Information and Internet Policy*, September 28-30, Washington D.C.

Nussbaum, Martha C. and Amartya Sen (Eds.) (1993), *The Quality of Life.* Clarendon Press.

Oostveen, A. M. and Besselaar, P. van den (2004), "Security as belief: user's perceptions on the security of electronic voting systems", in Prosser, A. and Krimmer, R. (Eds.), *Electronic voting in Europe: Technology, Law, Politics and Society*, Lecture Notes in Informatics, Vol. P-47, Gesellschaft für Informatik, Bonn, pp. 73-82.

Pieters, W. (2006), "Acceptance of voting technology: between confidence and trust", in Stølen, K. et al. (Eds.), *Trust Management: 4th International Conference, iTrust 2006*, Vol. 3986 of *Lecture Notes in Computer Science*, Springer, pp. 283-297.

Pieters, W. and Van Haren, R. (2007), "Temptations of turnout and modernization: e-voting discourses in the UK and the Netherlands", *Journal of Information, Communication and Ethics in Society*, Vol. 5, No. 4, pp. 276 - 292

Pursell, C. (1979), "Belling the Cat: A Critique of Technology Assessment", *Lex en Scientia*, Vol. 10, pp. 130-142.

Riedl, R. (2004), "Rethinking trust and confidence in European e-government: linking the public sector with post-modern society". In *Proceedings of the Fourth IFIP Conference on e-Commerce, e-Business, and e-Government (I3E 2004).*

Schwartz, N. L. (2004), ""Dreaded" and "dared": Courage as a virtue", *Polity*, Vol. 36, No. 3, pp. 341-365.

Shrader-Frechette, K.S. (1990), "Perceived risks versus actual risks: Managing hazards through negotiation", *Risk*, Vol. 1, pp. 341-363.

Smits, M. (2002a), "Monster ethics: a pragmatist approach to risk controversies on new technology", paper presented at the Research in Ethics and Engineering Conference, Technical University of Delft, April 25-27.

Smits, M. (2002b), *Monsterbezwering: de culturele domesticatie van nieuwe technologie*, Boom, Amsterdam.

Smits, M. (2006), "Taming monsters: The cultural domestication of new technology", *Technology in Society*, Vol. 28. No. 4, pp. 489-504.

Wallich, P. (1995), "Meta-Virus: Breaking the Hardware Species Barrier", *Scientific American*, Vol. 273, No. 5, p. 34.

Xenakis, A. and Macintosh, A. 2005, "Procedural security and social acceptance in e-voting", in *Proceedings of the 38th Hawaii International Conference on System Sciences (HICSS'05*).

Corresponding author
Luca Consoli can be contacted at: l.consoli@science.ru.nl