

The Effect of Security Education and Expertise on Security Assessments: the Case of Software Vulnerabilities

Luca Allodi, Marco Cremonini, Fabio Massacci, Woohyun Shim



Abstract—In spite of the growing importance of software security and the industry demand for more cyber security expertise in the workforce, the effect of security education and experience on the ability to assess complex software security problems has only been recently investigated. As proxy for the full range of software security skills, we considered the problem of assessing the severity of software vulnerabilities by means of a structured analysis methodology widely used in industry (i.e. the Common Vulnerability Scoring System (CVSS) v3), and designed a study to compare how accurately individuals with background in information technology but different professional experience and education in cyber security are able to assess the severity of software vulnerabilities. Our results provide some structural insights into the complex relationship between education or experience of assessors and the quality of their assessments. In particular we find that individual characteristics matter more than professional experience or formal education; apparently it is the *combination* of skills that one owns (including the actual knowledge of the system under study), rather than the specialization or the years of experience, to influence more the assessment quality. Similarly, we find that the overall advantage given by professional expertise significantly depends on the composition of the individual security skills as well as on the available information.

1 INTRODUCTION

Given the raising importance of cyber security, several Universities have recently introduced security courses and degree programs [40], [49]. Some governments already identified key knowledge areas for their workforce (e.g. [32]) whilst professional organizations also proposed curricula for Security IT management [22] and cyber security at large [33]. However, the set of core skills is widely debated with opinions often split between teaching adversarial thinking or principles and abstractions [27], [12], [38]. Widely debated is also the role of governments through public initiatives mandating academic training as a formal requirement for recognized professionals [8], [36].

In this scenario, it is surprising that only few studies considered measuring the effectiveness of security education and professional experience with respect to relevant security tasks. Experiments to quantitatively evaluate the benefits

of education and experience with respect to the ability to tackle a cybersecurity problem have been mostly performed by considering specific technical skills, often the ability to write a secure code or to recognize programming errors. In one of those studies [14], where the problem of finding vulnerabilities in code was considered, no correlation between education and code analysis effectiveness was found whilst a negative correlation emerged between the number of vulnerabilities found and analysts' years of experience in security. An hypothesis to explain these results is that complex relationships between professional knowledge and problem solving are at work [17], [13].

Goal and methodology. Our long term objective is to understand how different security education levels, professional experience, and personal skills affect the outcome of software security activities. The latter involve a combination of technical skills, system awareness, and management-level perspectives, forming the set of attitudes (or "mindset") [20] that is often part of the job of information security professionals [46] or managers [11]. This is also reflected in best practices for security education at large (e.g. the DHS table of minimal content [32]). Part of the challenge is to identify a task that is i) simple and structured enough to be amenable to controlled experiments whilst being ii) rich enough to provide a non obvious challenge to human participants to the experiments, and iii) decomposable so that one could probe different parts of the security skill sets.

We propose to use software vulnerability assessment as a *prima facie* proxy to evaluate the interplay between different aspects of a security task: understanding the impact of a software vulnerability on a system requires technical and operational skills (e.g. to evaluate impact vectors), as well as user-oriented and a management perspective (e.g. to estimate user requirements on the attack, or consequences of an attack). The choice of vulnerability assessment has another significant advantage: it comes with a pre-defined, world- standard evaluation framework that, for each of these aspects, provides an objective reference frame: the *Common Vulnerability Scoring System* (CVSS) [28]¹. It is of

- L. Allodi is at the Eindhoven Univ. of Technology, NL. (L.allodi@tue.nl).
- M. Cremonini is at the Univ. of Milan, IT. marco.cremonini@unimi.it.
- Fabio Massacci is at the Univ. of Trento, IT. fabio.massacci@unitn.it.
- W. Shim is with the Korea Institute of Public Administration, KR.

Presented at the Workshop of Economics of Information Security (WEIS'2018)

¹ CVSS is the most employed software vulnerability evaluation methodology in industry and government [35]. It is the result of the joint work of national standard bodies, software manufacturers, non-software companies, and Computer Emergency Response Teams (CERTs).

particular importance that its usage in corporate settings is not targeted to software security specialists. It is a tool that ‘general security practitioners’ with a CS background should be able to operate. For instance, the User Interaction (UI) metric only requires the assessor to comprehend how the vulnerable software is used and in particular whether or not a user’s action is required for the exploit to materialize; the Attack Vector (AV) metric, instead, requires evaluating aspects of the network stack of the vulnerable application. Obviously, such task is only a part of the skills of a software security experts. Security testing is another example of activities and related skills. The DHS Curriculum provides a good starting point for a list of such activities and skills set [32]. However, as one cannot test all them at once (too many confounding factors) we decided to focus on the assessment task for this study.

To investigate the relation between skills and quality of a software security task represented by a vulnerability assessment, we conducted an experiment with three groups of participants, asking them to score a set vulnerabilities following CVSS guidance. Our participants were both students and professionals, the former divided between those with or without specific security education, the latter with several years of experience in the security field but no security education at academic level.

Main contributions. A preliminary finding is that being competent in security, through education or experience, improves the ability to evaluate vulnerability severity. Albeit expected, the quantification of such improvement may represent the basis for cost/benefit analyses (Eg whether pursue a specialized but costly Master degree or a professional training).

Indeed, we found that experienced security professionals have no clear advantage over younger students with a security specialization. Other experiments in different fields have shown that expert performance is close to that of novices when solving problems in novel situations for which ‘precompiled’ rules are not helpful [41]. In an other case, *expertise reversal effect* has been detected when experienced learners decide to ignore novel but detailed instructional procedures in favor of ‘I-know-better’ mental schemes [23]. Hence, formal security education could produce a workforce able to compete with seasoned professionals when asked to perform well-formalized but relatively new security tasks.

Whilst this result is encouraging, we believe that more studies and tests are needed in order to establish evaluation methodologies and metrics for advanced education and professional training that often are considered as important investments.

2 RELATED WORK

2.1 Information Security Education

Education is a recurrent topic in information security, often related to the definition of guidelines for academic curricula needed by a well-trained workforce or considered in field studies comparing the performance of different groups of professionals when faced to a security problem.

Relevant initiatives for the definition of curricula and guidelines have been carried out by professional associa-

tions like ISACA [22] or ACM [27] as soon as cybersecurity has been recognized as an emergent IT profession. The most recent initiatives exhibit a remarkably improved maturity of proposals. In particular, a joint initiative between ACM, IEEE, AIS, and IFIP has produced a detailed framework for cybersecurity curriculum encompassing several knowledge areas [33]. The NSA, instead, together with the Department of Homeland Security, has been supported for years some Centers of Academic Excellence in Cyber Defense (CAE-CD) and, in connection with that initiative, the definition of a list of cybersecurity Knowledge Units (KUs) to be mapped on specific topics included into academic curricula [32]. With respect to the content of a cybersecurity curriculum, both the Joint Task Force and the CAE-CD propose a broad spectrum of competences, including technical and non-technical KUs (i.e., CAE-CD), or human, organisational, and societal security knowledge areas in addition to the more traditional system, data, software, component, and connection areas (i.e. the Joint Task Force). Other analyses have further contributed to the discussion [40], [26], [49], [38].

Relative to the professionalization of cybersecurity, some have discussed the difficulty that professionals have to agree on a set of well-established characteristics and the need for the cybersecurity profession to evolve [8], [36], [12]. These analyses are connected to our work because they reflect the lack of experimental studies on the effectiveness of security education and security experience, and on the emerging scenario composed by many security profiles and professions.

Several works obtained results from controlled experiments with students and professionals performing security tasks, although almost all of them focused on specific skills or even specific tools. In [47], a sample of developers recruited from GitHub was asked to complete simple security programming tasks. It turns out that the only statistically significant difference was determined by the years of programming experience, which indicates that familiarity with the task is the main driver, rather than professional or educational level. This is in line with results in [14], whereby security professionals do not outperform students in identifying security vulnerabilities. The usability of cryptographic libraries has been studied in [2]. Relevant for our work is the fact that different groups of individuals, with different levels of education or experience, have been recruited. They found that the participants programming skill, security background, and experience with the given library did not significantly predict the code’s overall functionality. Instead, participants with security skills produced more secure software, although neither as good as expected nor as self-evaluated by participants. We have found compatible results, although under very different settings and more nuanced. All these works differ from ours in studying the performance of individuals with respect to a specific technical security skill or tool rather than studying how education, experience, and the combination of skills correlate with more general security problem solving performances. The one closer to us is [3], where Android developers’ decision making performance are also analyzed with respect to education and experience. The experiment was based on observing how developers with different background perform when provided with different types of documentation

and it found a sensible difference between professionals and students. However, with respect to our work, the analysis based on the different backgrounds was limited and only a minor part of the study.

In other research fields, namely human-computer interaction and software engineering, there has been an ample discussion on the accuracy and precision of problem solving by different types of evaluation panels. Some studies have since long investigated the factors that distinguish experts from novices [44] or in evaluating software usability [18], [24]. In general, what studies in different fields have demonstrated is that, under specific circumstances, experts not necessarily perform better than novices in technical tasks, and that the relation between knowledge and problem solving ability is often unclear. These observations further motivate a specific investigation for security problems.

2.2 Vulnerability assessment

Vulnerability assessment represents a fundamental phase for security risk analysis and the prioritization of activities.

In this regards, it is both a technical task and a security management issue, which includes the adoption of a structured analysis methodology supporting a metric for assigning a qualitative or quantitative evaluation to vulnerabilities and the availability of assessors with adequate skills and training [20], [35], [9]. Vulnerability distributions have been studied

Related to our work, some research analyzed the distribution of vulnerability assessments performed by means of the CVSS [39], [25], [16]. Similarly, [5], [7], [6] investigated whether vulnerabilities with a high CVSS score corresponds to an equally high risk of exploitation in the wild.

Some studies have considered the difference in vulnerability scoring produced by assessors with different profiles [19], [20]. Differently from these works, we aim at studying to what extent security education and practical experience influence the outcome of a vulnerability assessment.

3 STUDY DESIGN

3.1 Analysis goals and research questions

In this study we evaluate the effect of different subject characteristics on *technical*, *user and system-oriented*, and *managerial* aspects of a security task. Specifically, our study aims at the following two goals:

Goal 1: Effect of security knowledge

We first evaluate the effect of skills in security on the accuracy of an assessment. We distinguish between skills acquired through *formal security education*, meaning academic-level specialized security courses, and through *professional experience*, as years of work in the security field. We formulate two research questions:

RQ1.1: Does formal security education have a positive impact on the accuracy of a structured security assessment?

RQ1.2: Does professional expertise systematically improve assessment accuracy over formal security education?

Goal 2: Effect of specific security competences

The second aspect we want to investigate is how specific skills impact the accuracy of the security assessment. Our underlying hypothesis here is that the analysis should consider the specific mix of competences, derived from education and professional experience, in order to identify meaningful relations between education/experience and problem solving ability in security. The provision of a ‘standardized’ portfolio of security competences has been pushed by both academic curricula [27] and industry [45], yet it is unclear whether this fits well the requirements of a real world scenario. This research goal addresses the following research questions:

RQ2.1: Does the mix of competences affect the overall accuracy of an assessment?

RQ2.2: Which specific aspects of a security assessment are more influenced by which specific skill?

3.2 Task mapping and vulnerability selection

The CVSS v3 framework provides a natural mapping of different vulnerability metrics on aspects of the larger spectrum of security competencies we are considering: technical, user-oriented, and management-oriented. Using CVSS, the assessor performs an evaluation of the vulnerability based on available information.

Table 1 provides a summary of CVSS’s Base metrics used in this study, their possible values, and their relation with the three competency levels we identify.

To guarantee the vulnerabilities’ representativeness the wider set of vulnerability characteristics, we chose the vulnerabilities for our experiment by randomly sampling thirty vulnerabilities from the one hundred used by the CVSS *Special Interest Group* (SIG) to define the CVSS standard. This also assures that the sample is representative of the distribution of CVSS vulnerability measures in the population of vulnerabilities (which are not uniformly distributed, see for example [39], [6]). The vulnerability descriptions are taken from the National Vulnerability Database (NVD), the reference dataset of disclosed software vulnerabilities.

3.3 Participants and recruiting procedure

We follow [29] and performed a natural experiment recruiting three groups of individuals (total $n = 73$ participants): 35 major students with no training in security; 19 major students with three to four years of specific training in security; 19 security professionals with several years of experience. Some participants knew what CVSS is used for and its scores associated to CVE vulnerabilities, but none had experience with CVSSv3 vulnerability assessment or knew the specific metrics used to produce the score.

The experiment has been organized by replicating the procedure performed by the CVSS Special Interest Group (SIG) for its own operations by asking each participant to complete 30 vulnerability assessments in 90 minutes by using only the CVE vulnerability description as the only technical information available. The official assessments produced by the CVSS SIG were used as our benchmark to compare between-group performance.

Unfortunately, recruiting subjects with very different profiles makes it hard to control for possible confoundings;

TABLE 1
Summary of considered CVSS v3 Base metrics and mapping to competency levels.

CVSS	Metric	Metric desc.	Values	Skill set
AV	Attack Vector	Reflects how remote the attacker can be to deliver the attack against the vulnerable component. The more remote, the higher the score.	Physical, Local, Adjacent Net., Network.	The assessor understands the technical causes and vectors of attack related to a software vulnerability. This encompasses knowledge of vulnerable configurations, local and remote attack delivery, and aspects related to attack engineering.
AC	Attack Complexity	Reflects the existence of conditions that are beyond the attacker's control for the attack to be successful.	High, Low.	
PR	Privileges Required	Reflects the privileges the attacker need have on the vulnerable system to exploit the vulnerable component.	High, Low, None.	The assessor understands the interaction between the vulnerable system, the user, and the attack. For example, attacks against administrative users may require specific attack techniques (e.g. spear-phishing); similarly, user behaviour may affect the outcome of a security problem (e.g. ignoring alert dialogues).
UI	User Interaction	Reflects the need for user interaction to deliver a successful attack.	Required, None.	
C	Conf.	Measures the impact to the confidentiality of information on the impacted system.	None, Low, High.	The assessors can evaluate the repercussions of a security problem over business-level aspects such as data exfiltration and system performance.
I	Integrity	Measures the impact to the integrity of information stored on the impacted system.	None, Low, High.	
A	Availability	Measures the impact to the availability of the impacted component.	None, Low, High.	

for example, some professionals may have received an education equivalent to that of (a group of) student subjects, or some students may have changed masters during their student career. As these effects are impossible to reliably measure, we explicitly account for the (unmeasured) in-subject variability in the analysis methodology and report the corresponding estimates.

3.3.1 Students

Students participating in our study are MSc students of two European Universities, both requiring proficiency in English and a background in computer science. The first group, SEC, is enrolled in the Information Security MSc of the University of Milan that completes a BSc in Information Security held at the same university. The second group, CS group, is composed of students enrolled in a Computer Science MSc at the University of Trento, Italy. SEC subjects were recruited during the *Risk Analysis and Management* course at the first year of their MSc; CS students were recruited during the initial classes of the course *Security and Risk Management*, the first security-specific course available in their MSc curriculum.

Table 2 reports students skills as core knowledge units (taken from the U.S. Center for Academic Excellence (CAE) Core Knowledge Units in Cyber-security) of respective BSc programs. In particular, we see that they share some core Computer Science competences, whereas only the SEC group has been trained on core Information Security competences.

3.3.2 Professionals

Subjects in the PRO group are members of a professional security community lead by representatives of the Italian headquarters of a major US corporation in the IT sector. Participants in our study have been recruited through the advertisement in the Community's programme of a training course on CVSS v3. Participants in the PRO group have different seniority in security and all professional profiles focus on security-oriented problems, technologies, and regulations.

To characterize PRO experiences, we asked them to complete a questionnaire detailing job classification and years of experience, education level, experience in vulnerability assessment, and expertise level in system security/hardening, network security, cryptography, and attack techniques. Of the 19 components of the PRO group, 13 completed the questionnaire. The median subject in the PRO group has six years of expertise in the security field, and roles comprise Security Analysts, CERT members, Pentesters and IT auditors. A detailed characterization of PRO subjects over the other dimensions is given in Sec. 4.2.

3.4 Data collection

Ahead of the experiment, participants attended an introductory seminar to CVSS v3 held by one of the authors with several years of expertise on CVSS. Content and delivery of the seminar have been identical for the three groups. After that, participants were given a printed sheet in tabular

TABLE 2
Core Knowledge Units for CS and SEC students

	Basic Data Analysis	Basic Scripting	IT Sys. Components	Networking Concepts	Sys Admin.	DataB. Mngmt. Sys.	Net. Techn. and Prot.	Op. Systems Concepts	Prob. and Stats.	Programming	Fund. of Sec. Design	Fund. of Crypto.	Cyber Defense	Cyber Threats	Net Defense	Policy/Ethics Compl.	Fund. of Inf. Assur.
CS	●	●	●	●	●	●	●	●	●	●	○	○	○	○	○	○	○
SEC	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	○

TABLE 3

Example of assessment by one randomly selected participant for each CS, SEC, PRO groups compared to SIG’s evaluation for CVE 2010-3974

Excerpt of the CVE 2010-3974: *fxscovr.exe in the Fax Cover Page Editor in Microsoft Windows XP SP2 and SP3, Windows Server 2003 SP2, Windows Vista SP1 and SP2, Windows Server 2008 Gold, SP2, R2, and R2 SP1, and Windows 7 Gold and SP1 does not properly parse FAX cover pages, which allows remote attackers to execute arbitrary code via a crafted .cov file, aka “Fax Cover Page Editor Memory Corruption Vulnerability”.*

Participant in group:	CVSS assessment								Confident
	AV	AC	PR	UI	C	I	A		
CS	N	H	L	N	L	L	L	L	Yes
SEC	L	L	L	R	H	H	H	H	Unsure
PRO	L	L	N	N	H	H	H	H	Yes
SIG	L	L	N	R	H	H	H	H	-

form reporting vulnerability description, CVSS metrics, and requesting an evaluation on the confidence in the assessment. All participants had at hand a summary description of CVSS v3 metrics obtained from the First.org’s website for reference during the exercise. Time for completing the test was 90 minutes, chosen on the basis of previous pilot studies. All participants completed the assessment in the assigned time with the exception of seven students in the CS group. All experiment material is provided for consultation at <https://github.com/cvssexp/cvssmaterial>.

Table 3 reports an example of vulnerability assessment. The answers from one participant, randomly chosen, for each group, are shown together with reference evaluations produced by SIG (bottom row). In this particular case, the CS student had all answers wrong and, despite this, declared to be confident in his/her evaluation. Both the SEC student and the PRO professional, instead, made one mistake, but exhibited different degree of confidence in their evaluation.

3.5 Analysis methodology

We formalize a CVSS assessment by assuming there exists a function $a_i(v_j)$ representing the assessment produced by assessor $i \in \{CS, SEC, PRO\}$ of vulnerability v represented as the vector of CVSS metrics to be evaluated ($j \in \{AV, AC, UI, PR, C, I, A\}$). We further define a function $e(a_i(v_j))$ that detects the error on metric j by assessor i on vulnerability v by comparing the subject’s assessment $a_i \in \{CS, SEC, PRO\}(v_j)$

with the assessment provided by the SIG $a_{SIG}(v_j)$ on the same vulnerability.

We observe subjects in our study multiple times (once per vulnerability). As each observation is not independent and subjects may learn or understand each vulnerability differently, a formal analysis of our data requires to account for the variance in the observation caused by subject (e.g. rate of learning or pre-existent knowledge) and vulnerability characteristics (e.g. clarity of description). To evaluate the effect rigorously, we adopt a set of mixed-effect regression models that account for two sources of variation: the vulnerability; and the subject [4]. The general form of the models is

$$g(y_{iv}^j) = \mathbf{x}_{iv}\beta + \mathbf{z}_i\mathbf{u}_i + \mathbf{h}_v\mathbf{k}_v + \epsilon_{iv}, \quad (1)$$

where $g(\cdot)$ is the link function, and y_{iv}^j denotes the observation on CVSS metric j performed by subject i on vulnerability v . \mathbf{x}_{ivj} is the vector of fixed effects with coefficient β . The vectors \mathbf{u}_i and \mathbf{k}_v capture the shared variability at the subject and vulnerability levels that induces the association between responses (i.e. assessment error on CVSS metric j) within each observation level (i.e. subject i and vulnerability v). ϵ_{iv} is the leftover error. We report regression results alongside a *pseudo-R*² estimation of the explanatory power of the model for the fixed-effect part as well as for the full model as specified in [30]. We report odds ratio (exponentiated regression coefficients) and confidence intervals (via robust profile-likelihood estimations [43]) for a more immediate model interpretation. Odds lower than one (with $0 \leq C.I. < 1$) indicate a significant *decrease* in error rates. These are indicated in Table 6, 7 with a * next to the estimate. *Borderline* results are those whose C.I. only marginally crosses the unity up to 5% (i.e. $0 \leq C.I. \leq 1.05$).

4 EMPIRICAL RESULTS

Our data collection comprises 2190 assessments performed by 73 subjects over 30 vulnerabilities. We consider an assessment as valid if the assessment is a) *complete* (i.e., the whole CVSS vector is compiled), and b) *meaningful* (i.e. the assessment is made by assigning a valid value to each CVSS metrics). This leaves us with 1924 assessments for our analysis, or $\approx 88\%$ valid records.

TABLE 4
Confidence assessments for the groups

Group	Confident			tot.
	Yes	No	Unsure	
CS	228	552	82	862
SEC	275	203	57	535
PRO	254	167	106	527
tot.	757	922	245	1924

4.1 Effect of security knowledge

4.1.1 Assessment confidence

We start our analysis by evaluating the level of scoring confidence for the three groups for each vulnerability. Table 4 shows the results for the subjects' reported confidence in the assessments. Overall, subjects declared to have been confident in their assessment in 39% (757) of the cases, and non-confident in 48% (922). For the remaining 13%, subjects were unsure. Looking at the different groups, a significant majority of scorings in the CS group (64%) was rated as low confidence, while for SEC and PRO groups approximately 50% were confident assessments. Even by considering 'Unsure' assessments as low confidence, the figures for the SEC and PRO groups are statistically indistinguishable ($p = 1$ for a Fisher exact test²), whereas the difference is significant between CS and SEC+PRO confidence levels ($p = 0.017$).

4.1.2 Severity estimations

Whereas technical details may significantly vary between vulnerabilities, for simplicity we grouped the vulnerability assessed into six macro-categories whose definitions have been derived from the *Common Weakness Enumeration* (CWE)³:

- `input`: vulnerabilities caused by flawed or missing validation (e.g. code injection);
- `information`: vulnerabilities regarding system or process specific (e.g. info disclosure);
- `resource access`: vulnerabilities granting the attacker access to otherwise unauthorized resources (e.g. path traversal);
- `crypto`: vulnerabilities affecting cryptographic protocols or systems;
- `other`: vulnerabilities that do not belong to specific CWE classes (taken as is from NVD);
- `insufficient information`: vulnerabilities for which there is not enough information to provide a classification (taken as is from NVD).

Figure 1 reports how severity estimations of vulnerabilities vary, w.r.t. the reference score computed by the SIG, between the three groups of participants and for each vulnerability category. A positive difference indicates an *overestimation* (i.e. participants attributed a higher severity score); a negative value indicates an *underestimation*. We observe that `Cryptographic Issues` and `Insufficient`

information categories were perceived as more severe by all participant groups than by the SIG, whereas for other categories the results are mixed. Following NIST guidelines, and over- or under-estimation of two points may result in an important mis-categorization of the vulnerability, whereas an error of ± 0.5 points is within accepted tolerance levels [1]. Overall, we find that experiment subjects' estimations of vulnerability severity are only marginally off with respect to the SIG estimations.

4.1.3 Assessment errors

In Figure 2 we have a more detailed inspection of scoring errors for the three groups by considering the specific CVSS metrics rather than the total score as computed by the CVSS for a vulnerability. We first evaluate the *sign* and the *size* of errors. With regard to the sign of an error, for instance, the PR metric could have three values (High, Low, None; see Table 1). Assuming that the SIG attributed the value Low for a certain vulnerability, if a participant selects High the error is an overestimation (positive error, +1), if he or she selects None it is an underestimation (negative error, -1). Errors may also have different sizes, which depend on the specific metric and the specific SIG evaluation. In the previous example, the size of the error is at most 1. However, for a different vulnerability the SIG could have evaluated as High for the PR metric. In that case, if a participant selects Low it results in a negative error of size 1 (i.e., -1), if he/she selects None the error size is 2 (i.e., -2), with different consequences on the overall scoring error for the vulnerability.

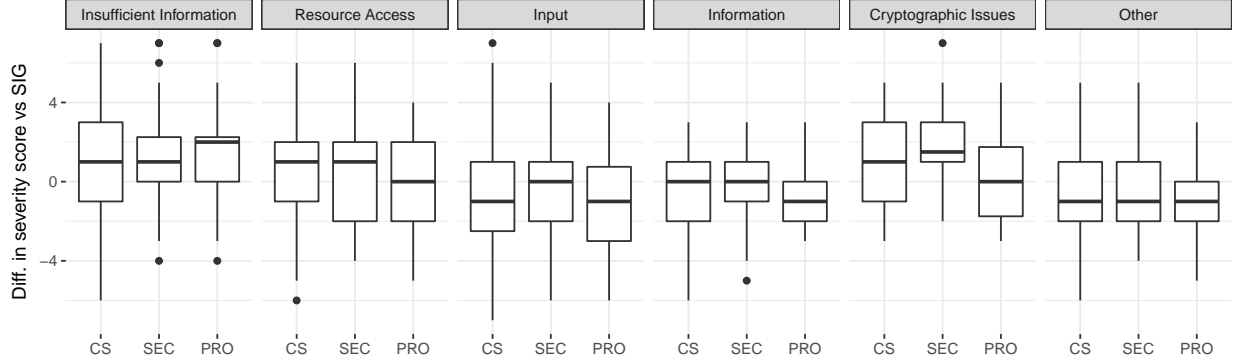
Given this computation of errors' sign and size, we observe that the frequency of large errors (defined as errors with size greater than 1), is small. This indicates that, in general, subjects did not 'reverse' the evaluation by completely missing the correct answer (e.g. assessing a High Confidentiality impact as a None), a situation that might have led to a severely mistaken vulnerability assessment. Whereas a detailed analysis of error margins is outside of the scope of this study we observe that, overall, most subjects in all groups showed a good grasp of the task at hand.

The large errors we observe on certain metrics (between 30% and 60% of tests, depending on the group of respondents and the metric, as discussed in the following) are mostly produced by errors of size 1. Error rates of this size are to be expected in similar experimental circumstances [34, finds error in the 30-40% rate over a binomial outcome], particularly considering that participants in our experiment have been explicitly selected with no previous experience in CVSS assessment, the limited amount of time, and the CVE description as the only technical documentation, this rate of small errors is unsurprising.

Overall, we observe that there is a clear difference in accuracy between the security unskilled CS and security skilled SEC+PRO for all metrics. This is particularly evident in the AV, AC and PR metrics, and all CIA impact metrics. This effect is also present in the UI metric, but here the CS and SEC students perform similarly, whereas professionals in the PRO group achieve higher accuracy. We observe an overall tendency in *over-estimating* PR and UI, and *under-estimating* AC, which may indicate that relevant information

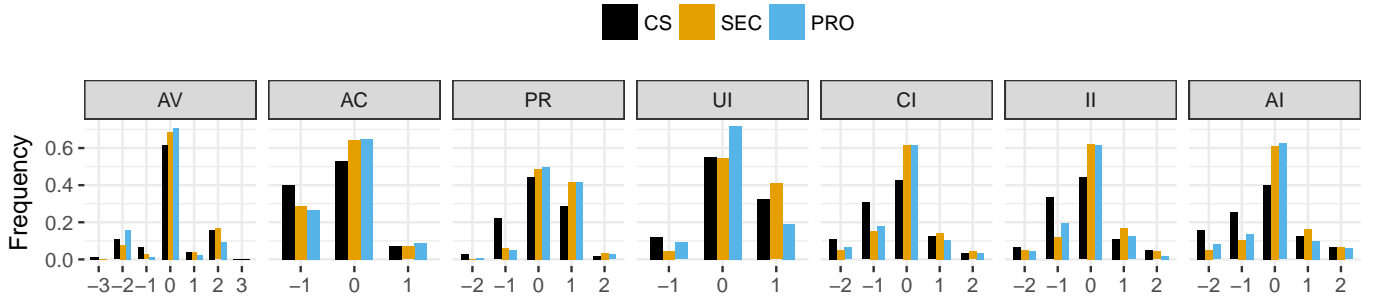
2. To avoid issues with dependent observations, we compute the rate of "Yes", "No", "Maybe" answers for each subject, and consider the highest rate to match the subject to a category.

3. Details at: <http://cwe.mitre.org>, last visited April 2018.



A positive difference indicates that subjects in that group overestimated the severity of the vulnerability w.r.t. the SIG's score. A negative difference indicates underestimation. Groups consistently overestimate vulnerabilities with *Insufficient information*, and *Cryptographic issues*.

Fig. 1. Distribution of difference in severity estimation by vulnerability type and subject group.



$error = 0$ indicates accordance. $error < 0$ indicates that subjects under-estimated that metric's assessment. $error > 0$ indicates that subjects over-estimated it. SEC and PRO subjects are consistently more precise than CS in assessing vulnerability impact.

Fig. 2. Distribution of assessment errors over the CVSS metrics.

for the assessment of these metrics are missing, a sensible problem already noted in the industrial sector as well (see for example the recent 'call for action' from NIST [31]). Conversely, the difference between SEC students and PRO professionals seems less pronounced, if present at all. The tendency of the error does not appear to meaningfully differ between groups, indicating no specific bias toward over or underestimation.

Table 5 reports quantitatively the comparison depicted in Fig. 2. As error sign does not present obvious between-group differences, for the sake of conciseness here we only consider differences in *absolute* error rates. At a first approximation, we evaluate significant differences between the groups by employing an unpaired Wilcoxon rank sum test,⁴ and test the alternative hypothesis that error size follows the intuition that $e(a_{i \in \text{PRO}}(v_j)) < e(a_{i \in \text{SEC}}(v_j)) < e(a_{i \in \text{CS}}(v_j))$, $\forall j \in \{\text{AV}, \text{AC}, \text{UI}, \text{PR}, \text{C}, \text{I}, \text{A}\}$. The statistical significance level is calculated using a Holm-Bonferroni correction for multiple comparisons. From the results it is apparent that subjects with security knowledge (i.e., SEC+PRO) are significantly more accurate than subjects with no security knowledge (i.e., CS) on all metrics. Interestingly,

4. This assumes independency between observations. This is violated in our data as the same subject is observed multiple times. Whereas this is formally incorrect, here it only serves the illustrative purpose of quantifying results in Fig. 2. A formal analysis follows in this Section.

TABLE 5
Preliminary comparison of group accuracy by CVSS metric.

	(1) CS		(2) SEC		(3) PRO		Signif.	
	err	sd	err	sd	err	sd	1v2	2v3
AV	0.67	0.91	0.56	0.86	0.54	0.87	*	
AC	0.47	0.50	0.35	0.48	0.35	0.48	**	
UI	0.45	0.50	0.45	0.50	0.28	0.45		**
PR	0.60	0.58	0.55	0.57	0.54	0.56		
C	0.72	0.70	0.48	0.66	0.48	0.67	**	
I	0.68	0.67	0.47	0.66	0.45	0.61	**	
A	0.82	0.77	0.51	0.70	0.51	0.73	**	

Significance is Holm-Bonferroni corrected and is indicated as: ** for $p < 0.01$ and * for $p < 0.05$ for a Wilcoxon rank sum test. A more formal evaluation accounting for within-subject dependency of observation is provided in Table 6.

we find that over all metrics (with the exception of UI as discussed above), the PRO group does not appear to perform significantly better than the SEC group. Whereas we can not say that the two groups perform *equally*, it is suggestive to notice that the computed error rates and standard deviations for the two groups indicate a substantial overlap between the two distributions. As discussed in the following, this result should not be considered completely surprising or counterintuitive, rather an interesting con-

vergence with similar experiments performed in different contexts and settings seems to arise, pointing to possible common characteristics of a class of problems.

As each metric has a different set of possible values, to simplify the interpretation of results, we here consider the binary response of *presence* or *absence* of error in the assessment. We define a set of regression equations for each CVSS metric j of the form:

$$g(e_{vi}^j) = c + \beta_1 \text{CONF}_{vi} + \beta_2 \text{GROUP}_i + \beta_3 \text{VULNTYPE}_v + \dots \quad (2)$$

where $g(\cdot)$ is the logit link function, e_{vi}^j is the binary response on presence or absence of error on metric j for subject i and vulnerability v , and $\beta_2 \text{GROUP}_i$ and $\beta_3 \text{VULNTYPE}_v$ represent respectively the vector of subject groups (CS, SEC, PRO), and vulnerability categories.⁵

Table 6 reports the regressions' results. We conservatively consider assessments with an 'Unsure' level of confidence (*ref.* Tab. 4) as non-confident. Effects for the group variables SEC and PRO are with respect to the baseline category CS. We report the estimated change in odds of error and confidence intervals of the estimation.

From the results it appears that *subjects with security knowledge, i.e. SEC+PRO, are significantly more accurate at the assessment than subjects with no security knowledge, i.e. CS, on all metrics.* Overall, SEC+PRO is between 30% to 60% less likely than CS in making an error. Interestingly, we find that *the PRO group does not appear to perform significantly better than the SEC group* (with the exception of UI metric, as reported in Fig. 2, for which PRO is approximately 60% less likely to err than SEC subjects). A borderline result is found for the AV and PR metrics, where the C.I. for SEC is only marginally crossing 1. This may indicate that the professional expertise that characterizes the PRO group does not necessarily improve the accuracy of the assessment over subjects with security knowledge but limited or no professional expertise. The effect of confidence on the assessment is relevant for the impact metrics CIA, indicating that a significant source of uncertainty may emerge from the effect of the vulnerability on the system. Interestingly, we find that some vulnerability types (Information and Resource access) are likely to induce error on the A metric, suggesting that specific knowledge or expertise may be needed to discern, for example, between *information* and *service* availability.

Variance by subject ($\text{Var}(c|ID)$) and by vulnerability ($\text{Var}(c|CVE)$) indicate that the intercept of the model may vary significantly for each observation (i.e. both different subjects and vulnerabilities have different 'baseline' error rates). This is interesting in itself as it indicates that neither the subject variables (GROUP_i) nor the vulnerability variables (VULNTYPE_v), whereas significant in explaining part of the observed error, may fully characterize the effect. For example, specific user characteristics or the thoroughness of the vulnerability description may play a substantial role in determining assessment accuracy. On

this same line, it is interesting to observe that the overall explicative power of the model is relatively small for all the considered metrics. This can be expected for random processes in natural experiments where the environment can not be fully controlled by the experimenter [4] (as exemplified by the variance explained by the full model as opposed to that of the fixed effects); still, the small R^2 values for the fixed effect parameters suggest that the sole presence of security knowledge, even when confounded by assessment confidence and vulnerability type, does not satisfactorily characterize the observation. This further supports that other subject-specific characteristics may drive the occurrence of an error. We investigate this in the following.

4.2 Effect of subject characteristics

To analyze results in finer detail, we use the answers from the questionnaire that characterizes PRO subjects as described in Sec. 3.5. This allows us to avoid possible bias in self-reporting by students and to focus on the target group of the professionals that eventually perform the analysis in the real world [37].

The median subject in the PRO group has six years of professional expertise in the security field, in a range between one and fifteen years ($\mu = 5.79$, $\sigma = 3.83$). Figure 3 reports the distribution of the levels for each measured variable. All factors are reported on an ordinal scale (with the exception of CVSS *experience* for which we have a nominal scale), codified in levels 1 \rightarrow 3, where *Education*: 1=High School; 2=BSc degree; 3=MSc degree. *Previous CVSS experience*: 1=None; 2=Yes; 3=NON-CVSS metric. *System security* \rightarrow *Attacks*: 1=Novice; 2=Knowledgeable; 3=Expert (a fourth level, 'None', is not reported as no participant rated him or herself less than novice on any of these dimensions). Most subjects obtained at least a BSc degree. From discussion during the initial CVSS training it emerged that none of the participants in the PRO group had a formal specialization in security at the University level. The group is evenly split between participants that have previous experience in vulnerability measurement (earlier versions of the CVSS or other methods); most participants rated themselves as 'Competent' or 'Expert' in Network Security, and are equally split between the levels 'Novice' and 'Competent or Expert' for all other variables.

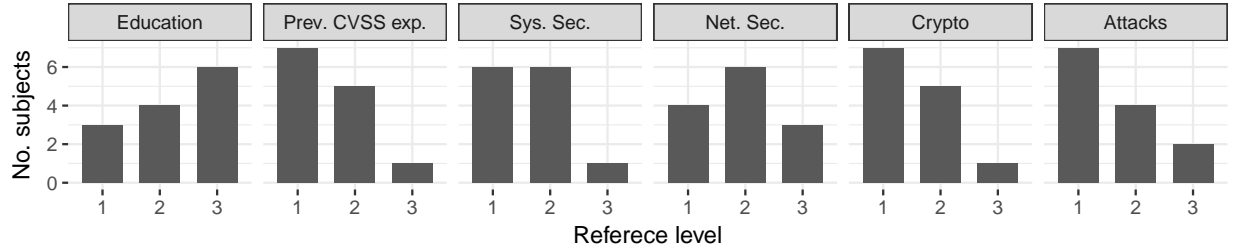
To evaluate the effect of subject characteristics on odds of error we first make two considerations: first, the subject distribution seem to be skewed toward presence or absence of expertise or education rather than being meaningfully distributed across all levels. For example, most subjects attended University with only a handful interrupting their studies after high school; similarly, few subjects rated themselves as 'experts' in any dimension, with most subjects being either 'novices' or 'competent' on the subject matter. We therefore collapse the levels to 'novice' or 'not novice' to represent this distinction. Second, some subject characteristics may show high levels of correlation: for example, subjects competent in system security may be likely competent on network security as well. Similarly, highly educated professionals may be (negatively) correlated with years of experience (as more time would be spent on one's studies than on the profession). We check for multicollinearity prob-

5. We did consider interaction effects between explanatory variables in the preliminary phases of this analysis, and found qualitatively equivalent results. To avoid complicating the notation and the result interpretation, we do not report those here.

TABLE 6
Effect of security education on odds of error

Regression on odds of error accounting for presence or absence of security knowledge and professional security expertise. SEC+PRO are significantly more accurate than CS in the assessment. SEC does not perform significantly better than CS in the UI metric, whereas PRO does. Marginal results are obtained for AV and PR.

error	AV	AC	UI	PR	C	I	A
c	0.34 [0.11; 1.01]	1.11 [0.57; 2.14]	3.26 [0.91; 11.75]	3.16* [1.06; 9.52]	1.01 [0.38; 2.68]	1.48 [0.60; 3.66]	0.61 [0.22; 1.72]
SEC	0.70 [0.47; 1.04]	0.58* [0.38; 0.87]	1.05 [0.72; 1.53]	0.75 [0.55; 1.04]	0.41* [0.26; 0.64]	0.46* [0.32; 0.67]	0.36* [0.25; 0.52]
PRO	0.58* [0.39; 0.87]	0.59* [0.39; 0.89]	0.36* [0.25; 0.53]	0.72* [0.52; 0.99]	0.39* [0.25; 0.61]	0.47* [0.32; 0.68]	0.34* [0.23; 0.49]
Conf.	0.86 [0.65; 1.11]	1.00 [0.78; 1.27]	0.84 [0.64; 1.10]	1.01 [0.79; 1.28]	0.71* [0.55; 0.92]	0.64* [0.50; 0.82]	0.79 [0.61; 1.01]
<i>Vulnerability variables</i>							
Cryptographic Issues	0.43 [0.06; 2.90]	1.48 [0.51; 4.32]	0.36 [0.04; 3.19]	0.17 [0.03; 1.09]	1.38 [0.27; 7.08]	1.20 [0.27; 5.43]	3.74 [0.64; 21.83]
Information	2.15 [0.42; 11.21]	1.20 [0.47; 3.09]	0.19 [0.03; 1.29]	0.46 [0.09; 2.47]	2.82 [0.66; 12.00]	1.53 [0.40; 5.78]	4.58 [0.97; 21.87]
Input	2.69 [0.79; 9.24]	0.67 [0.33; 1.35]	0.23* [0.05; 0.94]	0.50 [0.15; 1.72]	1.59 [0.54; 4.66]	0.88 [0.32; 2.36]	2.91 [0.92; 9.38]
Resource Access	0.76 [0.16; 3.55]	0.88 [0.37; 2.11]	0.18 [0.03; 1.04]	0.19* [0.04; 0.87]	2.22 [0.58; 8.54]	1.21 [0.35; 4.06]	4.87* [1.15; 20.76]
Other	2.21 [0.42; 11.82]	0.48 [0.19; 1.20]	0.08* [0.01; 0.53]	0.51 [0.10; 2.65]	1.68 [0.39; 7.13]	1.12 [0.29; 4.25]	4.35 [0.92; 20.97]
$Var(c ID)$	0.25	0.33	0.19	0.93	0.38	0.22	0.20
$Var(c CVE)$	1.04	0.30	1.42	0.64	0.79	0.66	0.92
$PseudoR^2$ (fixed eff.)	0.09	0.04	0.12	0.13	0.07	0.06	0.11
$PseudoR^2$ (full mod.)	0.34	0.19	0.41	0.41	0.31	0.26	0.34
N	1924	1924	1924	1924	1924	1924	1924



All factors but CVSS exp. (nominal) are on an ordinal scale. Reference levels: *Education*: 1=High School; 2=BSc degree; 3=MSc degree. *Previous CVSS experience*: 1=None; 2=Yes; 3=NON-CVSS metric. *System security* → *Attacks*: 1=Novice; 2=Knowledgeable; 3=Expert.

Fig. 3. Education and expertise profile of professionals in the PRO group.

lems by calculating the Variance Inflation Factor of the categorical variables defined above, and drop the variables that show evidence of correlation; we keep: *years*, *attacks*, *system security*. We then define the following regression equation:

$$g(e_{vi}^j) = c + \beta_1 Years_i + \beta_2 Attacks_i + \beta_3 SysSec (3) + \beta VULNTYPE_v + ..$$

Table 7 reports the results. In general, we observe that not all expertise dimensions are relevant for all evaluation metrics. This is to be expected as, for example, knowledge of in attack techniques may have an impact on evaluating complexity of attack, but may make little difference on other more system-oriented aspects like requirements on user interaction. More in detail, we find that *Attack expertise* dramatically decreases error over the AV and AC metrics by almost 60%. *Years of experience* increases accuracy over the UI and PR metrics (by roughly 20% per year), explaining the

mismatch on UI between SEC and PRO subjects identified in Fig. 2. *System security* knowledge appears to have a positive impact on the accuracy of assessments on the C and A metrics, but we do not consider this effect to be highly significant. Results for vulnerability type are qualitatively equivalent to those reported for the evaluation by group in Tab. 6. Interestingly, the overall explanatory power of the model (accounting for both fixed and random effects) remains satisfactory, and the *subject characteristics are clearly effective in explaining the variance for most metrics*. The only low (< 10%) R^2 fixed-effect values is for AV and can be explained by the low incidence of error in this metric, which may be then simply be driven by random fluctuations. This is in contrast with the effect for, for example, the AC metric that is characterized by a high variability in error (*ref.* Fig. 2), and for which more than 20% of the variance is explained by the measured fixed effects. This is in sharp contrast with results in Tab. 6 where most of the variance was absorbed by the random effects.

TABLE 7
Effect of subject characteristics on odds of error in the PRO group

Regression on odds of error by subject characteristics and vulnerability category. Education, CVSSExp, NetSec, Crypto have been dropped because highly correlated with other factors in the regression; this is to avoid multicollinearity problems. Overall we find that different vulnerability aspects are covered by different subject characteristics.

error	AV	AC	UI	PR	C	I	A
c	0.79	2.70	3.42	59.80*	2.22	3.81	0.37
Years	[0.14; 4.36]	[0.56; 13.54]	[0.83; 14.79]	[1.83; 3027.71]	[0.35; 14.13]	[0.43; 34.31]	[0.04; 2.87]
Attacks	0.96	0.95	0.86*	0.80*	0.86	0.86	0.90
SystemSec	[0.84; 1.09]	[0.81; 1.11]	[0.76; 0.97]	[0.64; 0.99]	[0.71; 1.05]	[0.68; 1.10]	[0.74; 1.11]
	0.49*	0.42*	1.32	0.66	0.45	0.41	0.61
	[0.26; 0.89]	[0.19; 0.85]	[0.77; 2.25]	[0.24; 1.77]	[0.18; 1.10]	[0.13; 1.23]	[0.23; 1.53]
	1.14	0.74	0.77	0.74	0.48	0.43	0.42
	[0.62; 2.14]	[0.35; 1.54]	[0.44; 1.33]	[0.27; 2.03]	[0.19; 1.20]	[0.13; 1.35]	[0.15; 1.07]
<i>Vulnerability variables</i>							
Cryp. Issues	0.24	4.67	0.24	0.01	1.20	1.21	6.29
Information	[0.01; 3.53]	[0.65; 42.22]	[0.03; 1.79]	[0.00; 1.81]	[0.16; 9.43]	[0.15; 10.07]	[0.59; 76.88]
Input	2.13	1.03	0.14*	0.02	1.77	3.05	13.13*
	[0.23; 20.43]	[0.19; 5.63]	[0.02; 0.83]	[0.00; 2.33]	[0.30; 10.92]	[0.49; 21.10]	[1.64; 124.53]
Resource Access	0.81	0.21*	0.17*	0.20	1.08	0.59	4.07
	[0.15; 4.32]	[0.05; 0.72]	[0.04; 0.62]	[0.00; 8.24]	[0.28; 4.26]	[0.15; 2.40]	[0.82; 23.81]
Other	0.42	0.62	0.26	0.02	2.25	1.01	15.37*
	[0.05; 3.51]	[0.13; 3.00]	[0.05; 1.34]	[0.00; 1.91]	[0.42; 12.57]	[0.18; 5.76]	[2.20; 131.91]
	1.19	0.12*	0.12*	0.23	1.14	0.62	6.13
	[0.11; 12.64]	[0.02; 0.74]	[0.02; 0.73]	[0.00; 37.62]	[0.19; 7.10]	[0.09; 4.04]	[0.73; 57.79]
$Var(c ID)$	0.02	0.14	0.00	0.51	0.32	0.60	0.34
$Var(c CVE)$	1.59	0.74	0.83	1.58	0.83	0.90	1.19
$PseudoR^2$ (fixed eff.)	0.07	0.22	0.12	0.14	0.10	0.14	0.16
$PseudoR^2$ (full model)	0.38	0.39	0.30	0.47	0.33	0.41	0.43
N	357	357	357	357	357	357	357

5 DISCUSSION

Economics of technical education and of professional experience. The value of professional experience and of technical education could be measured by the remunerations offered to IT professionals and through the cost of achieving a degree in security-oriented courses or the cost for an advanced professional training program. Therefore, to know that information security knowledge significantly affects the accuracy of a security assessment brings no surprise. However, what is almost never measured, and might have important implications in investments for building a competent workforce, is the width of the gain. In other words, the economics of technical education and professional experience is seldom analytically investigated, more often is left to anecdotes or to political discourses.

With our study we have attempted a quantitative estimate of the benefit brought by knowledge. On average, it appears remarkable: those with security knowledge (SEC and PRO groups) show error rates reduced by approximately 20% (see Fig. 2). A second result appears by looking at the average confidence declared by participants: not just assessment accuracy improves with knowledge, but also confidence in assessments. In fact, the unskilled students CS are mostly not confident, while the skilled participants SEC+PRO declare higher confidence. Improving confidence in one's own work is as valuable as improving the accuracy, because it leads to a better self-evaluation, a better control of the task at hand, and optimizes time and efforts.

Implications for recruiting and training security professionals. What we have seen in our tests is that the combination of skills explains most of the subjects' variance. This is another observation often made anecdotally, but seldom

analytically tested in order to be translated into operational policies and tools useful in the definition of recruiting plans or training investments. Moreover, competences in different technical domains are correlated. For example, expertise in system security/hardening is highly correlated with expertise in network security; similarly, experience with previous vulnerability assessments and attack expertise go hand in hand. What emerges from our study is that professionals with same specialization, i.e., cybersecurity, exhibits different competence profiles defined by correlated skills, and different performance in assessing different security metrics are correlated with different profiles. This means that often the specification "security professional with x years of experience", or even a professional qualification like "Security Architect" and the like, could be a too coarse classification to be useful for finding a good match with a security task, especially when the task is mostly oriented to problem solving.

More specific and detailed studies about how specific technical skills correlate and influence problem solving abilities should be useful to inform recruiting and training investments. Recruiting should better know whether specific specialist skills are needed or technical problem solving abilities are required. This should be considered at fine-grained level, by analyzing skills correlation and complementarity. The same holds for professional training, today often unable to be tailored for a clear outcome when a problem solver profile is needed.

Studying socio-technical issues by testing students. One clear result from our study is that students with security training and security professionals could perform similarly for some class of problems and under specific conditions. Beside observations related to the economics

of technical education and of professional experience, the result is suggesting us another possibility: if we are able to control the conditions leading to similar performance between students and professionals, then for all practical scenarios in which those conditions are realistic, we could reliably study the socio-technical phenomena by testing students, typically easy to enroll in scientific studies, rather than professionals, notoriously very difficult to recruit for experiments.

In general, it seems that similar performance between students and professionals emerges when professionals are unable to address the problem by reusing past mental models or patterns of solution. It also seems to emerge for problems characterized by subjective interpretations and uncertainty about the consequences of an individual evaluation, whereas it disappears in situations where the amount of specialist technical skills or practice gained in years of experience is the only important factor.

We have already encountered such a situation in the software engineering literature. In [21], the authors explicitly tested whether CS students could be employed in place of professional programmers for assessing the project's lead time. Similar to our scenario, to perform the assessment, it is required to evaluate subjective data and the importance of contributing factors is uncertain. The study found no significant difference between students and professionals. In [37], again, the quality of programming between students and professionals has been tested. The results were that professionals perform significantly better when they could adopt a familiar development approach, but the equivalence between professionals and students returns when an unfamiliar development approach was tested. This effect of experience was also confirmed by cognitive studies [23].

Our work add a new scenario regarding information security and problem solving skills. Further experiments could shed more light on the specific conditions that may lead to a substantial equivalence in tests between students with certain knowledge and professionals with given profiles. This would contribute to develop a relevant body of knowledge about the preconditions for defining suitable natural experiments for specific technical aspects, like cybersecurity, without necessarily recruit the scarce and hardly available professionals.

Uncertainty in problem statement affects differently the different groups. Another result that consistently emerges from the analysis is that the assessment accuracy may vary significantly among vulnerabilities for all three groups of participants. In the Appendix, specific details and examples are provided. In general, this represents a confirmation of the bias possibly introduced by framing the security problem. Uncertainty, e.g. produced by an equivocal problem description, might introduce a random error, which may penalize either the experts or the non-experts. This suggests that follow-up studies may more formally consider the bias introduced by ambiguous vulnerability descriptions on the assessment accuracy and, in general, on problem solving when security threats are affected by uncertainty.

6 THREATS TO VALIDITY

We here identify and discuss the limitations of our study. We consider *Internal and External* threats to validity [48].

Internal. Subjects in all groups were given an introductory lecture on vulnerability assessment and scoring with CVSS prior to the exercise. The exercise was conducted in class by the same lecturer, using the same material. Another factor that may influence subjects' assessments is the *learning factor*: "early assessments" might be less precise than "late assessments". All subjects performed the assessment following a fixed vulnerability order. We address possible biases in the methodology by considering the within observation variance on the single vulnerabilities [4]. Further, the use of students as subjects of an experiment can be controversial, especially when the matter of the study is directly related to a course that the students are attending[10]. Following the guidelines given in [42], [10] we made clear to all students that the exercise is not part of their final evaluation, and that their assessments do not influence their grades or student career.

External. A software engineer investigating a vulnerability in a real scenario can account for additional information beside the vulnerability description when performing the analysis. This additional information was not provided in the context of our experiment. For this reason we consider our accuracy estimates as conservative (worst-case). Accounting for professionals directly addresses this concern and confirms previous studies on the representativeness of students in software engineering experiments [37]. On the other hand, the limited number of participants in the SEC and PRO groups, and the difficulty associated with recruiting large sets of professionals [37] calls for further studies on the subject.

7 CONCLUSIONS

In this study we evaluated the effect of security knowledge and expertise on security vulnerability assessment. Whereas the case study is specific to the application of CVSS, it nevertheless provides a useful framework that formalizes adversarial, system, and user perspectives in an assessment, and therefore well captures the overall skill set on which security practice is built. An important lesson learned with this work is that recurrent effects of expertise and education with respect to some classes of problems described in previous studies unrelated with security, emerge also with an important security task like the vulnerability assessment. Our hypothesis is that the task we have tested in our experiment exhibits some general characteristics common with similar tasks studied, for example, in software engineering contexts. This seems interesting to us since it suggests that possible fruitful approaches to some security problems could be derived from experience in different fields or applying methods and analyses not specifically produced within the security field.

Whereas it is not our goal to provide the ultimate guidelines for improving the shortage of well-trained security professionals, some useful lessons might be learned from our study. One is to avoid inflexible assumptions regarding the prevalence of experience over education or adversarial thinking over principles and abstractions. What we have

seen is that the ability to face to complex security problems cannot be explained by coarse-grained categories or single characteristics and skill. What matters is always a combination of skills and experience of the assessor and the specifics of the problem to be evaluated. Another lesson learned is that experience may replace formal education, but likewise specific training may provide abilities similar to those brought by experience. Therefore, for both academic education and professional training, the ability to face complex security problems mostly depends on a well-balanced skill set in breadth and depth of knowledge, combined with practical experience.

Future work will replicate this analysis with the final CVSS language [15], and include the *Scope* metric as well. We are particularly interested in cooperating with other researchers to replicate our study in different national and educational context as results might have important policy implication for university education in cybersecurity and eventually for cybersecurity in the field.

8 ACKNOWLEDGMENTS

This research has been partly supported from the European Unions Seventh Framework Programme for research, technological development and demonstration under grant agreement no 285223 (SECONOMICS), and from the NWO through the SpySpot project (no.628.001.004).

REFERENCES

- [1] NIST National Vulnerability Database (NVD), 2015. [online] <http://nvd.nist.gov>.
- [2] Y. Acar, M. Backes, S. Fahl, S. Garfinkel, D. Kim, M. L. Mazurek, and C. Stransky. Comparing the usability of cryptographic apis. In *Security and Privacy (SP), 2017 IEEE Symposium on*, pages 154–171. IEEE, 2017.
- [3] Y. Acar, M. Backes, S. Fahl, D. Kim, M. L. Mazurek, and C. Stransky. You get where you’re looking for: The impact of information sources on code security. In *Security and Privacy (SP), 2016 IEEE Symposium on*, pages 289–305. IEEE, 2016.
- [4] A. Agresti and M. Kateri. Categorical data analysis. In *International encyclopedia of statistical science*, pages 206–208. Springer, 2011.
- [5] L. Allodi and F. Massacci. A preliminary analysis of vulnerability scores for attacks in wild. In *ACM Proc. of CCS BADGERS’12*, 2012.
- [6] L. Allodi and F. Massacci. Comparing vulnerability severity and exploits using case-control studies. *ACM Transaction on Information and System Security (TISSEC)*, 17(1), August 2014.
- [7] L. Allodi, S. Woohyun, and F. Massacci. Quantitative assessment of risk reduction with cybercrime black market monitoring. In *In Proc. of IWCC’13*, 2013.
- [8] D. L. Burley, J. Eisenberg, and S. E. Goodman. Would cybersecurity professionalization help address the cybersecurity crisis? *Communications of the ACM*, 57(2):24–27, 2014.
- [9] A. Calder and S. Watkins. *IT governance: A manager’s guide to data security and ISO 27001/ISO 27002*. Kogan Page Ltd., 2008.
- [10] J. Carver, L. Jaccheri, S. Morasca, and F. Shull. Issues in using students in empirical studies in software engineering education. *IEEE Int. Symp. on Sof. Metrics*, 0:239, 2003.
- [11] P.-y. Chen, G. Kataria, and R. Krishnan. Correlated failures, diversification, and information security risk management. *MISQ*, 35(2):397–422, 2011.
- [12] W. A. Conklin, R. E. Cline, and T. Roosa. Re-engineering cybersecurity education in the US: an analysis of the critical factors. In *System Sciences (HICSS), 2014 47th Hawaii International Conference on*, pages 2006–2014. IEEE, 2014.
- [13] R. J. Dufresne, W. J. Gerace, P. T. Hardiman, and J. P. Mestre. Constraining novices to perform expertlike problem analyses: Effects on schema acquisition. *The Journal of the Learning Sciences*, 2(3):307–331, 1992.
- [14] A. Edmundson, B. Holtkamp, E. Rivera, M. Finifter, A. Mettler, and D. Wagner. An empirical study on the effectiveness of security code review. In *International Symposium on Engineering Secure Software and Systems*, pages 197–212. Springer, 2013.
- [15] First.org. Common vulnerability scoring system v3.0: Specification document. Technical report, FIRST, Available at <http://www.first.org/cvss>, 2015.
- [16] L. Gallon. Vulnerability discrimination using cvss framework. In *Proc. of NTMS’11*, pages 1–6, 2011.
- [17] P. T. Hardiman, R. Dufresne, and J. P. Mestre. The relation between problem categorization and problem solving among experts and novices. *Memory & Cognition*, 17(5):627–638, 1989.
- [18] M. Hertzum and N. E. Jacobsen. The evaluator effect: A chilling fact about usability evaluation methods. *International Journal of Human-Computer Interaction*, 13(4):421–443, 2001.
- [19] H. Hibshi, T. D. Breaux, and S. B. Broomell. Assessment of risk perception in security requirements composition. In *Requirements Engineering Conference*, Ottawa, Canada, 2015. IEEE.
- [20] H. Holm and K. K. Afridi. An expert-based investigation of the common vulnerability scoring system. *Computers & Security*, 2015.
- [21] M. Höst, B. Regnell, and C. Wohlin. Using students as subjects: a comparative study of students and professionals in lead-time impact assessment. *Empirical Software Engineering*, 5(3):201–214, 2000.
- [22] ISACA. Isaca model curriculum for information security management, 2nd edition, 2012.
- [23] S. Kalyuga, P. Ayres, P. Chandler, and J. Sweller. The expertise reversal effect. *Educational psychologist*, 38(1):23–31, 2003.
- [24] C. Ling and G. Salvendy. Effect of evaluators’ cognitive style on heuristic evaluation: Field dependent and field independent evaluators. *International journal of human-computer studies*, 67(4):382–393, 2009.
- [25] Q. Liu and Y. Zhang. Vrss: A new system for rating and scoring vulnerabilities. *Computer Communications*, In Press, Corrected Proof:–, 2010.
- [26] D. Manson and R. Pike. The case for depth in cybersecurity education. *ACM Inroads*, 5(1):47–52, Mar. 2014.
- [27] A. McGettrick. Toward effective cybersecurity education. *IEEE Security & Privacy*, 11(6):66–68, 2013.
- [28] P. Mell, K. Scarfone, and S. Romanosky. A complete guide to the common vulnerability scoring system version 2.0. Technical report, FIRST, Available at <http://www.first.org/cvss>, 2007.
- [29] B. D. Meyer. Natural and quasi-experiments in economics. *Journal of business & economic statistics*, 13(2):151–161, 1995.
- [30] S. Nakagawa and H. Schielzeth. A general and simple method for obtaining r^2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, 4(2):133–142, 2013.
- [31] NIST. Vulnerability description ontology (vdo): a framework for characterizing vulnerabilities. <https://csrc.nist.gov/publications/detail/nistir/8138/draft>.
- [32] C. of Academic Excellence in Cyber Defense (CAE-CD). 2019 knowledge units, 2018.
- [33] J. T. F. on Cybersecurity Education. Curriculum guidelines for post-secondary degree programs in cybersecurity (csec2017), dec 2017.
- [34] K. Onarlioglu, U. O. Yilmaz, E. Kirda, and D. Balzarotti. Insights into user behavior in dealing with internet attacks. In *NDSS*, 2012.
- [35] S. D. Quinn, K. A. Scarfone, M. Barrett, and C. S. Johnson. Sp 800-117. guide to adopting and using the security content automation protocol (scap) version 1.0. Technical report, NIST, 2010.
- [36] R. Reece and B. C. Stahl. The professionalisation of information security: Perspectives of UK practitioners. *Computers & Security*, 48:182–195, 2015.
- [37] I. Salman, A. T. Misirli, and N. Juristo. Are students representatives of professionals in software engineering experiments? In *Proc. of ICSE’15*, volume 1, pages 666–676, May 2015.
- [38] H. Santos, T. Pereira, and I. Mendes. Challenges and reflections in designing cyber security curriculum. In *World Engineering Education Conference (EDUNINE)*, IEEE, pages 47–51. IEEE, 2017.
- [39] K. Scarfone and P. Mell. An analysis of cvss version 2 vulnerability scoring. In *Proc. of ESEM’09*, pages 516–525, 2009.
- [40] F. B. Schneider. Cybersecurity education in universities. *IEEE Security & Privacy*, 11(4):3–4, 2013.
- [41] C. Singh. When physical intuition fails. *American Journal of Physics*, 70(11):1103–1109, 2002.
- [42] D. Sjøberg, B. Anda, E. Arisholm, T. Dybå, M. Jørgensen, A. Karahasanović, and M. Vokáč. Challenges and recommendations

when increasing the realism of controlled software engineering experiments. In *Emp. Methods and Studies in Soft. Eng.*, volume 2765 of *LNCS*, pages 24–38. Springer Berlin / Heidelberg, 2003.

- [43] H. Stryhn and J. Christensen. Confidence intervals by the profile likelihood method, with applications in veterinary epidemiology. In *Proceedings of the 10th International Symposium on Veterinary Epidemiology and Economics, Vina del Mar*, page 208, 2003.
- [44] J. Sweller. Cognitive load during problem solving: Effects on learning. *Cognitive science*, 12(2):257–285, 1988.
- [45] B. Von Solms. Information security governance: Cobit or iso 17799 or both? *Computers & Security*, 24(2):99–104, 2005.
- [46] R. Wash. Folk models of home computer security. In *Proc. of SOUPS'10*, 2010.
- [47] D. Wermke and M. Mazurek. Security developer studies with github users: Exploring a convenience sample. In *Symposium on Usable Privacy and Security (SOUPS)*, 2017.
- [48] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén. *Experimentation in software engineering*. Springer Science & Business Media, 2012.
- [49] C. Yue. Teaching computer science with cybersecurity education built-in. In *ASE@USENIX Security Symposium*, 2016.

APPENDIX

Here we discuss the assessment results for three vulnerabilities as examples of the way participants with different skills and experiences have interpreted uncertain information (see Figure 4 and the following discussion of the three examples). Finally, Figure 5 reports error rates for all vulnerabilities of the assessment. Figure 4 reports the assessment accuracy (expressed in terms of number of errors) for three vulnerabilities that represent typical outcomes: (i) the three groups perform similarly; (ii) SEC+PRO have a clear advantage over CS; (iii) we obtain mixed results over different metrics.

- *Similar accuracy over all metrics (CVE-2014-2005)*.
Sophos Disk Encryption (SDE) 5.x in Sophos Enterprise Console (SEC) 5.x before 5.2.2 does not enforce intended authentication requirements for a resume action from sleep mode, which allows physically proximate attackers to obtain desktop access by leveraging the absence of a login screen.

From this description, it is clear that the attacker needs to be *physically proximate* to the target system, which gives an obvious clue for AV; similarly, all groups showed low error rates over the CIA assessment, as it is clear that the attacker gets full (user) access to the system by impersonating the legitimate user. Whereas almost all SEC and PRO subjects understood that the attacker need not be logged in *ahead* of the attack and scored PR correctly, CS students were likely confused by the existence of an authentication mechanism for the attacker to bypass. This suggests that well-formalized security tasks may be accomplished comparably well by security experts and general IT experts.

- *Clear effect of security knowledge (CVE-2009-1136)*.
The Microsoft Office Web Components Spreadsheet ActiveX control (aka OWC10 or OWC11), as distributed in Office XP SP3 and Office 2003 SP3, Office XP Web Components SP3, Office 2003 Web Components SP3, Office 2003 Web Components SP1 for the 2007 Microsoft Office System, Internet Security and Acceleration (ISA) Server 2004 SP3 and 2006 Gold and SP1, and Office Small Business Accounting 2006, when used in Internet Explorer, allows remote attackers to execute arbitrary code via a crafted call to the msDataSourceObject method, as exploited in the wild in July and August

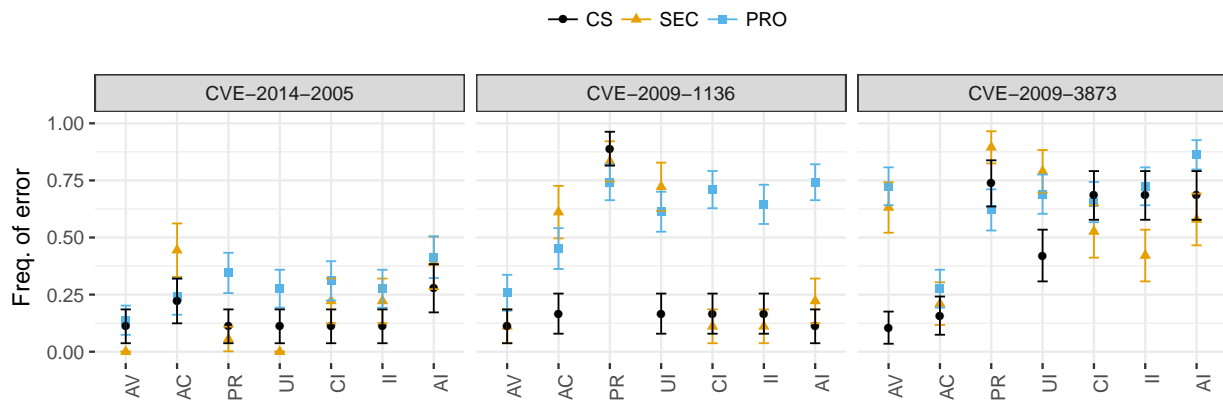
2009, aka “Office Web Components HTML Script Vulnerability.”

Whereas all groups correctly understood that the attack can happen remotely (AV), the security knowledge of SEC and PRO has a clear effect on the CIA metrics. For this vulnerability, students in both the SEC and CS groups were likely confused by the long list of vulnerable systems, giving the impression that these are specific vulnerable software configurations (a criteria for AC:H [15]), as opposed to a mere list of vulnerable software. PRO subjects did not get confused by this. In this vulnerability the PRO advantage on the UI metric, discussed in the analysis, is apparent: PRO subjects are the only one that consistently understood that the attack process requires a user to load a webpage that will *then* load the vulnerable method. This may be easier for PRO subjects to grasp because of the typical attack dynamics of phishing or XSS attacks commonly received by organizations.

- *Mixed results (CVE-2009-3873)*.

The JPEG Image Writer in Sun Java SE in JDK and JRE 5.0 before Update 22, JDK and JRE 6 before Update 17, and SDK and JRE 1.4.x before 1.4.2_24 allows remote attackers to gain privileges via a crafted image file, related to a “quantization problem,” aka Bug Id 6862968.

The high error for SEC and CS students is likely caused by the misleading “*remote attackers*” reference in the description: the vulnerability requires the component to load an image file locally (irrespective of whether this is provided from remote), and qualifies for an AV:L assessment (see also [15, Sec. 3.3 of the User guide]). PRO subjects did not get tricked by the misleading wording. Again, PRO subjects outperformed both student groups in the UI metric, understanding that the file need be loaded by the user (e.g. through interaction in a web browser). Interestingly, all groups have a high degree of error in the CIA metrics, suggesting that they deemed “*gain privileges*” as a moderate impact, whereas in most environments Java’s JDK/JRE will be running with already high privileges, hence giving the attacker full access.



Fraction of erroneous assessments by group for three CVEs. Higher on the scale corresponds to higher error (lower is better). The vertical bars report the standard errors. For the first vulnerability, CVE-2014-2005, the three groups perform similarly over all metrics. In the second, CVE-2009-1136, security knowledge gives a clear advantage on assessment accuracy, particularly in the CIA impact metrics. Lastly, for CVE-2009-3873 we observed mixed results, where security expertise appear to help for PR, CIA, but not for AV and AC, with SEC performing worse than CS and PRO on the UI metric.

Fig. 4. Example of assessment error rates by group on three CVEs

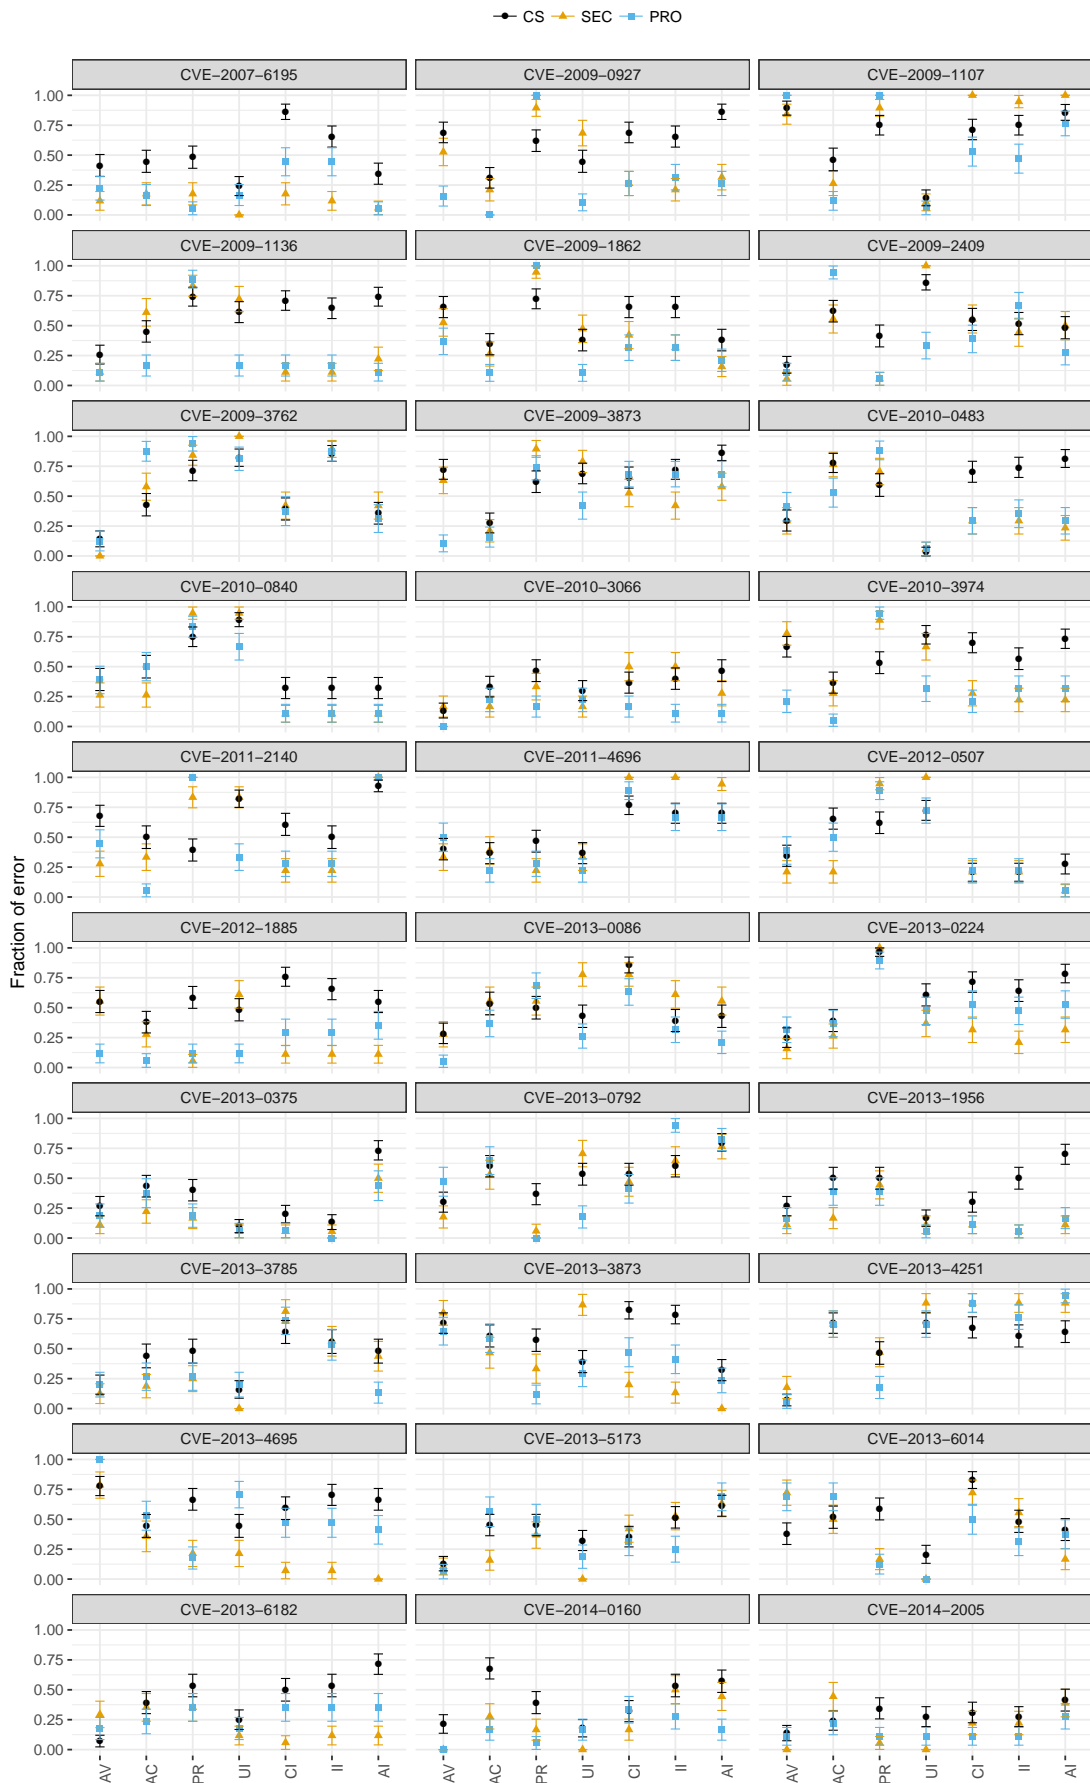


Fig. 5. Error rates for CS, SEC, and PRO by vulnerability and CVSS metrics