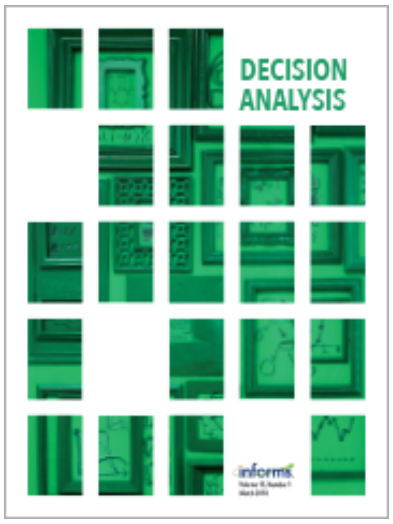

Decision Analysis

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Assessing and Forecasting Cybersecurity Impacts

Aitor Couce-Vieira, David Rios Insua, Alex Kosgodagan

To cite this article:
Aitor Couce-Vieira, David Rios Insua, Alex Kosgodagan (2020) Assessing and Forecasting Cybersecurity Impacts. Decision Analysis 17(4):356-374. https://doi.org/10.1287/deca.2020.0418

http://pubsonline.informs.org/journal/deca

DECISION ANALYSIS
Vol. 17, No. 4, December 2020, pp. 356–374
ISSN 1545-8490 (print), ISSN 1545-8504 (online)

Assessing and Forecasting Cybersecurity Impacts

Aitor Couce-Vieira,[a] David Rios Insua,[a,b] Alex Kosgodagan[a]

[a] Instituto de Ciencias Matemáticas, Consejo Superior de Investigaciones Científicas, Madrid 28049, Spain; [b] School of Management, University of Shanghai for Science and Technology, Shanghai 200093, P.R. China

Contact: coucevieira@outlook.com (AC-V); david.rios@icmat.es (DRI); alex.kosgoda@gmail.com,
https://orcid.org/0000-0003-2585-7695 (AK)

Received: September 6, 2019
Revised: March 17, 2020; June 1, 2020
Accepted: June 1, 2020
Published Online in Articles in Advance: October 26, 2020

https://doi.org/10.1287/deca.2020.0418

Abstract. Cyberattacks constitute a major threat to most organizations. Beyond financial consequences, they may entail multiple impacts that need to be taken into account when making risk management decisions to allocate the required cybersecurity resources. Experts have traditionally focused on a technical perspective of the problem by considering impacts in relation with the confidentiality, integrity, and availability of information. We adopt a more comprehensive approach identifying a broader set of generic cybersecurity objectives, the corresponding set of attributes, and relevant forecasting and assessment models. These are used as basic ingredients for decision support in cybersecurity risk management.

Funding: This work was supported by the European Union's Horizon 2020 project [Grant 740920 Supporting Cyber Insurance from a Behavioural Choice Perspective]. The work of D. Rios Insua is supported by the Spanish Ministry of Economy and Innovation Program [Grants MTM2017-86875-C3-1-R and RTC-2017-6593-7] and the AXA-ICMAT Chair on Adversarial Risk Analysis. This work was partially supported by a National Science Foundation grant [DMS-1638521] to the Statistical and Applied Mathematical Sciences Institute.

Keywords: cybersecurity • multiattribute utility • expert judgement • forecasting • risk management

1. Introduction

From private corporations to governmental facilities going through critical infrastructures, all kinds of organizations may be critically impacted by cyberthreats (Andress and Winterfeld 2013). Indeed, cybersecurity has become a major global problem as reflected in the Global Risks Report of the World Economic Forum (WEF 2020). Risk and decision analysis are fundamental methodologies to help managing such problems (Cooke and Bedford 2001). Through their tools, an organization can assess the risks affecting its assets and what security controls and insurance decisions should be implemented to reduce the likelihood and/or eventual impacts of cyber threats.

The medium-term aim of the research here described is the provision of a decision-analytic based support system that facilitates strategic cybersecurity resource allocations to an organization, which we call *defender*, within its annual cybersecurity planning process. Such a system should facilitate (i) forecasting and assessing the impacts of various cyber threats over the organizational assets, (ii) assessing the improvements induced by various feasible cybersecurity portfolios, and finally, (iii) choosing the best portfolio. A key element within the system consists of providing a preference model that supports such assessments, as we do here. Specifically, we provide the following:

- A generic tree of potential cybersecurity objectives for defenders, including the corresponding attributes. Its purpose is to support the identification of all potential impacts of cybersecurity threats in terms of relevant stakeholders' assets.
- A set of models relevant to forecast the outcomes in the involved attributes associated with the envisioned threats.
- A generic multiattribute utility function to translate the previous objectives in quantified assessments of stakeholders' preferences and risk profiles.

One could argue that different organizations pursue different cybersecurity objectives entailing, therefore, different cybersecurity decisions. For example, a standard small- or medium-sized enterprise (SME) might be interested only in having its online shop available; on the other hand, a large information technology (IT)-based company would also care about maintaining their services to third parties and protecting their personally identifiable information (PII) records and strategic business plans; finally, an energy company would like to avoid attacks to its cyber-controlled production plants to prevent major environmental impacts. We emphasize the generality of the proposed model in that it aims to serve as an initial catalogue for all kinds of organizations, regardless of their type and stakeholders. Relying on the proposed tree, an organization would select some of the objectives and eliminate the others. They could even include additional ones that are specifically relevant to such organization. In any case, observe that, from an implementation point of view, when presenting the tree to the organization we could highlight the objectives typical of its organization profile to better focus the discussion and analysis.

Our approach is inspired by earlier work in aviation safety risk management, counterterrorism, homeland security, and cybersecurity financial risk management. Rios Insua et al. (2019a) provide a value model for aviation safety at a state agency, including models to forecast and assess the impacts in aviation safety events; our proposal addresses cybersecurity events in a general organization. Keeney (2007b) identifies and structures preferences in antiterrorism analysis from the perspective of a government; we aim at achieving a similar purpose in the cybersecurity domain, covering all types of organizations. Keeney and von Winterfeldt (2011) provide a value model to assess homeland security decisions; we pursue a similar goal to support cybersecurity decisions, both in private and public administrations. Eling and Wirfs (2019) provide models to forecast some of the cybersecurity financial costs; we complement them with other financial and nonfinancial impacts.

A key point that we stress with respect to commonly used cybersecurity approaches is that we move beyond the traditional information security attributes of confidentiality, integrity, and availability (CIA) (Mowbray 2013) to a set of attributes that more globally conveys the impacts of cyber threats and are more amenable of interpretation from a business perspective. Interestingly enough, the complexity of hyper-connected environments has led to demands to go beyond the CIA triad (Vacca 2013) but has focused mainly on additional technical proposals in line with objectives referring to concepts such as authentication, authorization, or auditability procedures (Krutz and Vines 2004).

Indeed, recent cyberattacks help to motivate the need to broaden of the scope of cybersecurity impacts as seen through a few examples. For instance, interest in *impacts to other organizations* (Section 2.2.2) stems from attacks like the Target case in 2014, in which hackers attacked the famous retailer through its air conditioning supplier, stealing millions of PII records; this kind of attack has led to the new field of supply chain cyber risk management (Torres et al. 2020). Some of the subobjectives that we shall propose entail impacts that have been rarely considered in cybersecurity. For example, cyberattacks with *physical impact* (Section 2.2.3) are unusual, but the emergence of cyber-physical systems and smart infrastructures brings these risks to the fore; recall, for example, the Stuxnet attack. As another example, in 2018, a benevolent hacker used an antenna to spy on hundreds of aircraft taking control of onboard systems to carry out surveillance on all connected passenger devices (Brewster 2018); it then becomes obvious to foresee the potential physical damage that such an action could entail. Similarly, in relation with *environmental damage* (Section 2.2.4), Sayfayn and Madnick (2017) report how a hacker caused 800,000 liters of untreated sewage to flood the waterways in a certain area. At this point, it is also relevant to mention the impacts pursued by different attack types. For example, distributed denial of service (DDoS) attacks will typically affect just downtime and, consequently, reputation; ransomware threats will usually affect availability of information; and finally, exfiltration attacks will affect third parties and reputation.

The cybersecurity topic has been growing in popularity over the last years. However, there still lacks of material concerning modeling and decision-oriented research from a probabilistic angle, with some exceptions being Bagchi and Bandyopadhyay (2018), who examine the role of espionage in defending against

cyberattacks, and Pala and Zhuang (2019), who review the cybersecurity information-sharing literature. We contribute with our proposed cybersecurity multiattribute preference model. We first provide a generic objective tree for cybersecurity risk management. Ideally this would be shown to managers who would pick the relevant objectives for their problem at hand or, eventually, use it to complete their own objectives tree. For each of the objectives, potential attributes to assess objective achievement are provided. We also present forecasting models for the attributes, with a focus on nonmonetary impacts in which there is little incident data available, typically requiring structured expert judgement techniques (Dias et al. 2018); our interest is on the whole forecasting distribution, and not just point estimates, to compute expected utilities for risk management purposes. Once the objectives and their attributes are specified, we build a utility function to assess the impacts: we provide a generic model whose parameters are obtained through a series of questions that would need to be addressed by the cybersecurity risk manager. A numerical example serves us to illustrate the framework.

2. Cybersecurity Risk Management Objectives

2.1. Context and Process

As frequently emphasized (WEF 2020), cyberattacks may entail very negative consequences in terms of costs, loss of reputation, or even casualties. We track and manage them through objectives or performance measures that we want to optimize. We suggest attributes to assess the objectives.

The problem is contextualized through an organization whose objective is to introduce a strategy to improve cybersecurity: we aim at supporting them in their annual cyber risk management planning that leads to choosing their portfolio of cybersecurity controls, possibly including cyber insurance. This entails forecasting and assessing the potential impacts of all relevant cyber threats and how the controls help in mitigating them to finally choose the optimal portfolio. A full framing of the problem may be seen in detail in Rios Insua et al. (2019b) and an initial decision support system (DSS) implementation in Couce et al. (2019). Our focus here is on providing the list of relevant cybersecurity objectives and attributes, as well as on parametrized models to be used as default options to forecast and assess such impacts. Through risk management, we aim at implementing cybersecurity controls to perform optimally with respect to such objectives.

As mentioned, the objectives will typically vary from state organizations to private ones and, among these, will differ from standard SMEs to IT-based ones, or large companies. They may also vary in different countries and domains (e.g., air traffic management, healthcare, manufacturing). We present a generic list of objectives arranged as a tree, from which an organization may choose when undertaking their cyber risk management process.

The process followed to build such tree essentially consisted of an in-depth review of numerous cybersecurity standards and frameworks and related literature to identify common known impacts. We also reviewed related security and safety literature and recent cybersecurity events to complete the initial list of objectives merging those that were similar. Adopting the terminology in Keeney (2007a), our focus was on determining the fundamental objectives for cybersecurity risk management, stemming from lists of (mostly fundamental) objectives identified in other standards. We did not explicitly use a means-ends objective network, but rather used a mind map to group the identified objectives and then displayed in a tree to construct our cybersecurity value model. As Brownlow and Watson (1987) point out, a tree structure helps an organization in mitigating the cognitive overload when solving large complex issues, as in cybersecurity risk management. Finally, we made the name and definition of these categories as generic as possible. This was aimed at reflecting standard requirements (Keeney and Gregory 2005) that must be met by an objectives tree to be useful for decision support: it should be *comprehensive*, *measurable*, *nonoverlapping*, *relevant*, *unambiguous*, and *understandable*. Our initial version was assessed in-depth and complemented by the technical partners of our sponsoring project and then validated in two workshops with cybersecurity staff in both large and IT-intensive SMEs.

The lowest nodes in the tree provide q dimensions used to (i) describe, forecast, and assess the cybersecurity status of an organization so as to ascertain whether a risk management program needs to be implemented; (ii) forecast and assess the implementation

of risk management alternatives, which in our case are cybersecurity and cyber insurance portfolios; and finally, (iii) support the decision of choosing the best cybersecurity portfolio. Each of these scales should be quantified with an *attribute*, allowing each consequence to be represented as a vector of attribute levels $c = (c_1, c_2, \ldots, c_q)$.

2.2. Proposed Cybersecurity Objectives

The mainstream approach for describing cybersecurity objectives is in terms of the popular CIA information security triad (Mowbray 2013). However, the increasing variety of attacks and attacker interests and the need to better integrate cybersecurity in operational terms renders such objectives as IT oriented (and somewhat obsolete). They are useful for expressing security from a perspective in which systems are described in terms of sets of pieces of information that are stored, processed, and/or transmitted. We can think of this as the *technical perspective*. Yet a *business perspective* would focus more on assets and activities relevant for the organization and its stakeholders. This is even more relevant if we reflect the principles introduced previously: objectives should cover the consequences over organizational assets and activities expressed in variables understandable in the language of the incumbent organization. This is a major emphasis in our work and should be of major interest to the decision analysis community.

Our initial search covered the main cybersecurity frameworks that provide catalogues of concepts analogue to our objectives. We included the European Telecommunications Standards Institute ETSI GS ISI 002 v1.2.1 (2015a), ISO 22317 (ISO 2015b),[1] the Open Web Application Security Project (OWASP) business impacts (2017), the Organisation for Economic Cooperation and Development (OECD) cyber losses types (OECD 2017), the European Union Agency for Cybersecurity (ENISA) Information Package for SMEs (ENISA 2007), the ENISA report on IT business continuity management for SMEs (ENISA 2010), the Sherwood Applied Business Security Architecture (SABSA 2009), and Magerit (2012). We also included the list of impacts identified by Hubbard and Selersen (2016). All in all, they depict several general categories of impacts (legal and regulatory, productivity, financial reputation, and loss of customers) with some examples of subcategories. However, they did not meet well the standard decision analytic requirements: most of them provide a list of recurrent or important business impacts rather than a comprehensive list encompassing less typical impacts (e.g., physical ones); similarly, they provide objectives or impacts that somehow overlap: most of them affect monetary objectives, and thus, some categorization among them is convenient. For instance, some costs affect specific assets (e.g., activity interruptions), whereas others affect less tangible assets (e.g., competitive advantage, reputation). Of course, creating a comprehensive and nonoverlapping set of objectives may have disadvantages, namely, the addition of concepts. One example in business terms is that income generation is a clear and major objective for companies making money through sales. However, organizations have alternative means to earn money, including grants, investments, or licenses. This is especially relevant, for example, for nongovernmental organizations (NGOs). A second example refers to the emergent and potential impacts of cyber risks involving physical and psychological aspects. As a consequence, our approach is to list objectives and impacts in cybersecurity and sort them in a hierarchy in a more comprehensive, measurable, nonoverlapping, relevant, unambiguous, and understandable manner. Comprehensiveness and no overlaps involve, mostly, careful addition of novel concepts. Relevance and understandability are more related with translating impacts from the traditional CIA realm to another one based on assets, activities, and stakeholders.

Besides existing lists of cybersecurity impacts, other important influences derive from asset management and law. The first discipline (see ISO 55000 (ISO 2014) on general asset management or ISO 19770 (ISO 2015a) for IT assets) helps us in conceptualizing the different status that an asset could attain, so that engineers may characterize how an asset affects a system or the organization in terms of reliability and predictability. The second one stems from the distinction between damage on property (economic or pecuniary) and persons (general or nonpecuniary). This facilitates the distinction between objectives that can be assessed in monetary terms (directly or through estimation) and others that are nonmonetary and thus need special consideration when it comes to their evaluation. It also helps with the distinction between the objectives' owners

Figure 1. (Color online) Cybersecurity Objectives

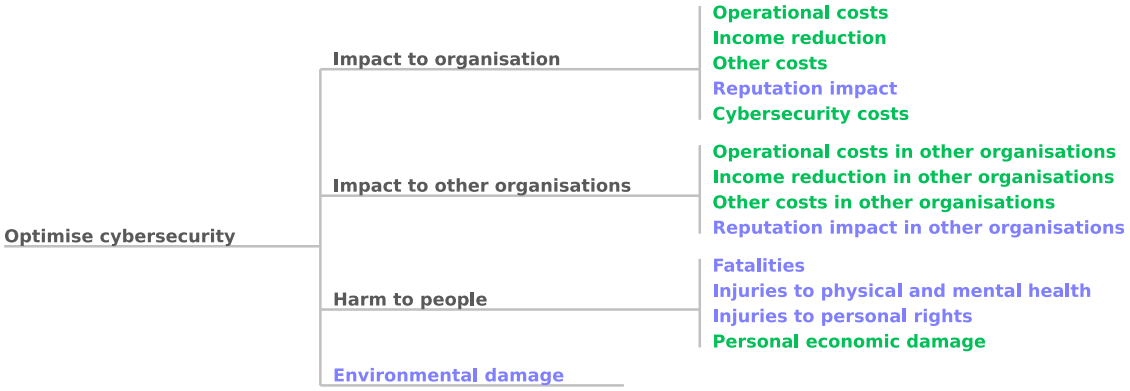

Note. Gray, assessed in monetary terms; light gray, not directly measurable in monetary terms; black, with both types of subobjectives.

(e.g., environmental damage suffered by third parties besides the monetary, or reputational consequences that such attacks could cause to the organization).

After an in-depth validation task, the previous process led us to develop the generic tree of cybersecurity objectives summarized in Figure 1, which complies with the properties mentioned previously. When it comes to comprehensiveness, we evaluated existing categories of impacts to, at least, have some that cover them. Solving eventual overlaps would entail creating more abstract concepts: we believe that this question should be addressed when performing specific cyber risk assessment exercises; should a threat involve impacts on several categories, it would be necessary to check that impacts included in one category are not included in a different one. For example, a big contract lost could be included either in loss of contracts or in loss of market share but not in both. Similarly, service unavailability may induce economic losses and reputational loss; once we forecast the downtime, we would need to reflect the entailed operational costs and reputation impact, as explained later. We also tried to bring more general terms for the objectives rather than more domain specific (e.g., organization instead of business). This may add a little more ambiguity and less understandability compared with domain-specific IT but provides undoubtedly a more comprehensive approach.

The rest of this section describes the rationale behind such objectives. All of them refer to minimization. For example, when mentioning *impact to the organization*, we understand *minimizing impact to the organization*. Unless explicitly mentioned, the objectives will be expressed in monetary terms. Some of them refer to impacts that may last several years, even though our planning exercises refer to just one. The natural approach would be to aggregate them to equivalent impacts over the planning year; for instance, monetary impacts could be dealt with net present values (NPVs) (French and Ríos Insua 2000). The appendix maps two of the previously mentioned catalogues (SABSA and Magerit) within our proposed tree as an example.

2.2.1. Impact to the Organization. This objective consists of the following subobjectives:

2.2.1.1. Operational Costs. We cover here those costs related with the assets and activities involved in the organization's operations, the area responsible for producing goods or delivering services, and the cost of degradation, malfunction, abuse, unavailability, elimination, recovery, and unrecoverability of their associated assets and activities. We focus on assets such as software, IT devices, documents, and equipment and activities such as serving food, writing a report, or supporting administrative acts with citizens. These impacts can be represented with a monetary attribute. We include the following:

- Degradation, if the asset or activity performs its function in a less productive or more costly manner, for example, a text processor running slower than normal as an asset degradation, or slower document production as an activity degradation.

- Malfunction, if the asset or activity has disturbances or a hazardous behavior, for example, a text processor producing errors when writing several pages.
- Abuse, if the asset or activity is maliciously manipulated, for example, a malicious macro exfiltrating the document edited in the text processor.
- Unavailability of the asset or activity, for example, the employees cannot run the text processor.
- Recovery through the actions and resources to restore an asset or an activity back to normal. Some assets might be unrecoverable (e.g., a piece of art), and this might have an operational impact (e.g., uninstall a text processor with tailored macros that cannot be reprogrammed because the original programmer is unavailable).To assess these impacts, we just need to follow financial accounting practices, supported by analytic accounting with the aid of experts on the asset or activity. Basically, the value of the asset is its market value; changes on such asset will have an impact on future activities of the organization. For example, a degradation may affect the market value or reduce future income flows and similarly for malfunction and unavailability. Everything pivots around how less productive is the asset and whether this change in productivity affects the asset valuation and future income flows. Abuse is somewhat more difficult to assess; in this case, rather than calculating how less productive is the asset, we could look at the likelihood of a sabotage or incident and the costs that this incident may cause.

2.2.1.2. Income Reduction. Cyberthreats may impact the organization income, reducing it because of the loss of sales, contracts, market share, funding, or licenses. In a business context, they typically involve marketing and commercial aspects related to sales. However, we also take into account that some income does not necessarily have such origin, for example, in public and nonprofit organizations. All of them can be assessed in monetary terms. We include the following:

- Income reduction over sales flow, involving sales but also leads, quotes, postsale, and customer service.
- Loss of market share, expressed through the reduction over the sales flow. It could also be considered as an asset with an estimated economic value that can drop if market share is reduced.
- Loss of funding not directly related with sales flow, for example, through investments, grants, or public funding.
- Loss of licenses. It has a compliance origin, but their loss could reduce income.

Observe that when contracts are few but big, loss of contracts might be more of a practical indicator than sales and market share.

2.2.1.3. Other Costs. They include certain strategic, compliance, and financial assets or potential costs. Although their identification or estimation might be difficult, all of them may translate into income (e.g., technological advantages) or costs (e.g., less advertisement for a well-known brand). All of them can be represented through monetary attributes. We include the following:

- Loss of competitive advantage caused by leaked, spied, or publicly disclosed sensitive information, including intellectual property or commercial secrets. Although it could be correlated with income reduction or reputation impact, it is also considered an intangible, but defined, asset that can be estimated (Raggio and Leone 2019) or sold.
- Depreciation, abuse, unavailability, or elimination of financial assets. Examples are changes in stock value, financial blackmail, extortion or ransom, and theft of financial assets, including money or financial instruments.
- Costs from noncompliance with contracts, regulations, standards, or any other enforceable policy. Examples are fines and regulatory penalties, contractual and agreement penalties, and litigation costs.

2.2.1.4. Reputation Impact. We refer to impacts over reputation that affect the trustworthiness of the organization as an institution rather than those more directly measurable in monetary terms that impact brand value, reduce income or operations, or the activities toward recovering reputation. In principle, these impacts cannot be represented through monetary attributes.

2.2.1.5. Cybersecurity Costs. It is practical to separate the costs related with managing cybersecurity, because this is the activity we aim to support in our broader decision-making model (Rios Insua et al. 2019b). It covers the costs of preventive and reactive controls

and eventual (cyber) insurance. All of them are assessed in monetary terms.

2.2.2. Impact to Other Organizations. As mentioned, this is a relatively recent concern, being equivalent to third party liability in insurance. A cybersecurity incident in one organization might affect others, and thus the objectives should also involve minimizing damage to them. It replicates the objectives for our organization except that referring to minimization of cybersecurity costs, because we are not supporting their cybersecurity decision making. Therefore, it includes the following as subobjectives: *operational costs (to other organizations), income reduction (to other organizations), other costs (to other organizations), and reputation impact (to other organizations)*. These last ones are nonmonetary.

2.2.3. Harm to People. As exemplified, a cyber incident might also affect people such as employees, customers, or local communities. Therefore, the organization objectives could also involve minimizing harm to people, as described in the Introduction. We include the following: *fatalities* (nonmonetary), *physical and/or mental health injuries* (nonmonetary), *injuries to personal rights*, such as dignity or privacy (nonmonetary), and *personal economic damage*.

2.2.4. Environmental Damage. Similar to damage inflicted to people, the environment might be affected by cyberattacks against systems entailing physical operations. Indeed, in many industrial complexes, a network of sensors gathers and monitors data about equipment efficiency and materials flow. Information is sent to computer terminals and processed into commands for hardware elements like motors, pumps, and valves. This technology can control, for example, the flow of pipelines, the level of water in a reservoir, or the gates that hold in and control the release of sewage. Here we model the impact over the natural environment as such. As an example, the costs of cleaning pollution are an impact to organizations or people, but we include here, for example, the amount of land severely polluted as triggered by a cyber incident.

3. Attributes for Quantifying Nonmonetary Objectives

Several of the objectives that we identified are not measurable in monetary terms. We describe here how we may proceed for each of them. In general, we could start with the identification of the main scenarios that various cyber incidents could cause. These what-if scenarios should be comprehensive in terms of covering all feasible types of impacts related to the objective that the relevant stakeholders, assets, and activities of the organization may suffer if attacked. In our case, subject-matter experts within the project team helped in devising the scenarios, which were then validated in workshops with cybersecurity specialists. Once these scenarios are identified, they should be quantified for their use in the model. For this, we try to use a natural attribute, if available and simple enough to assess; when this is not possible, we could look for a constructed attribute or a proxy one (Keeney 1992).

3.1. Impact on Reputation

Hubbard and Selersen (2016) discuss how to assess reputation damage in cybersecurity. The authors demonstrate that evidence linking data breaches and stock prices of an attacked company is not strong enough. Rather, actual reputation losses may be more realistically modeled as a series of tangible costs and other internal and legal liabilities. In other words, organizations seem to engage in efforts to control damage to reputation instead of bearing what could otherwise be a much greater impact. The effect of these efforts on reducing the actual reputation loss tends to be enough so that the impact over sales or stock prices is hardly detected. As mentioned, this objective may impact brand value or reduce income. However, it also encompasses aspects related to trustworthiness, legitimacy, and image, potentially leading to reduced market share.

In the organizational theory literature, several authors apply an overall measure of reputation (Fombrun 2012), whereas others use an attribute-specific measure (Jensen et al. 2012) because organizations may have multiple types of reputation. To facilitate understanding, we can adopt the four categories in Carpenter and Krause (2011) with names adapted to our context: *moral reputation* (referring to how the organization treats stakeholders); *compliance* (related to how the organization follows legal and social norms); *performative* (concerning the capability of the organization for performing their job); and finally, *adaptability* (related to the capability of the organization to deal with complex environments different from a *business as usual* status).

Common ways of measuring or building attributes for concepts like reputation are interviews with stakeholder representatives or surveying a representative sample of such groups (van Riel and Fombrun 2007). This is meaningful when it is done for specific groups of stakeholders (Jensen et al. 2012), taking into account a competitor or a similar organization (Fombrun 2012) and past reputation performance (Jensen and Roy 2008).

If we proceed with a constructed scale, we apply the principles presented in Keeney and Gregory (2005). We identify first the scenarios taking into account the previously mentioned components (e.g., type of reputation, type of stakeholders) identifying relevant cutpoints and thresholds. Once these scenarios are identified, they should be ordered from most to least preferred. Table 1 provides a simple example for a particular organization.

Alternatively, as proxy attributes, we could use the salience of cybersecurity incidents in news, media, and social networks or the cost of handling the reputation impact of the incident.

3.2. Harm to People: Fatalities and Injuries to Physical and Mental Health

Cyber incidents pose risks potentially triggering incidents that may affect people's health. Usually, cybersecurity physical risks would be a causing or facilitating event of an already existing safety risk that, most of the time, has been documented by the organization through industrial or occupational assessments. Indeed, from a physical point of view, cyber threats constitute a major concern when, for example, designing medical devices (Fu and Blum 2013), industrial control systems (Macaulay and Singer 2011), or autonomous driving systems (Taeihagh and Lim 2018). Moreover, mental health might be a relevant issue also, for instance, in relation to cyber bullying (Vandebosch and van Cleemput 2008).

Our first subobjective, minimizing fatalities, could be assessed with a natural attribute such as the number of fatalities. For the others, as an example, the World Health Organization (WHO) 2018 *International Classification of Diseases* (WHO 2018) provides a list with all types of injuries, diseases, and disorders together with several of their features. These classifications cover thousands of events or injuries. However, in a real case, our assessment will be more straightforward. Risk analysis typically distinguish between major and minor injuries (Rios Insua et al. 2019a). We could use them as the two natural attributes, also possibly discerning between physical and mental injuries.

They are also suitable for a constructed-attribute approach. Several methods may help us to create an ordinal scale (Hasler et al. 2012), such as the *injury severity score*, to assess the severity of injuries, the *global assessment of functioning* (GAF), or the *WHO disability assessment schedule* (Ustün et al. 2010) for physical or mental functioning. Table 2 provides an example with different levels of mental and physical impacts, based on some of the previous scoring systems, excluding those scores related to fatalities included in the corresponding subobjective.

Table 1. Example of Reputational Impact Scenarios Constructed Scale

Rank	Impact on reputation
1	No impact
2	Loss of moral or compliance reputation in up to 10% of employees, customers or the general public
3	Loss of performative or adaptability reputation in more than 25% of customers or general public
4	Loss of moral or compliance reputation in up to 50% of employees, customers or the general public
5	Loss of moral and compliance reputation in more than 50% of employees, customers or the general public

Table 2. Example of Physical and Mental Impact Scenarios Constructed Scale

Rank	Injuries to physical and mental health
1	No injury, emergency or functional impairment
2	Minor emergency that does not require medical intervention (NACA I); or minor injury (4 > ISS > 0); or absent or minimal psychological or physical symptoms, no more than everyday problems or concerns (GAF 81-90)
3	Slight to moderate non life-threatening emergency that requires medical intervention (NACA II and III); or moderate or serious injury (16 > ISS >= 4); or mild and moderate psychological or physical symptoms, causing slight to moderate impairment in social or occupational functioning (GAF 51-80)
4	Serious emergency that may be life-threatening and that requires medical care (NACA IV-VI); or severe to maximal (currently untreatable) injury (ISS >= 16); or serious psychological or physical symptoms or persistent danger causing serious to persistent inability in several areas of functioning including family, mood, relations, thinking or even danger of hurting self or others (GAF 1-50)

We combined in the same scale psychological and physical injuries (in the previous case based on the GAF score), but we could separate them.

Finally, the number of people entering into the hospital in relation to a cyber event could be used as a proxy attribute.

3.3. Harm to People: Injuries to Personal Rights

Cyberattacks may harm our dignity or privacy, accidentally or intentionally. Recall how large-scale activities of governments or companies on the Internet have become a major issue: the U.S. National Security Agency surveillance (Margulies 2013), the Great firewall of China (Lee and Liu 2012), or the scandal of Cambridge Analytica (Kurtz et al. 2018). In this context, governments and international institutions are pushing for a more secure and governable cyberspace. Namely, the United Nations (UN) Human Rights Council has stated that "the same rights that people have offline must also be protected online" (UN Human Rights Council 2015). See also the recent General Data Protection Regulation (GDPR) (European Parliament 2016) in Europe.

These rights could be identified from national jurisprudence, but the UN provides an international and overreaching framework. For our purposes, it is useful the classification system in the *Universal Human Rights Index Database* (UN Human Rights Council 2016), which covers human rights recognized by the UN under categories such as civil and political rights; economic, social, and cultural rights; or rights to specific persons or groups.

A constructed-attribute approach may be the best to make operational this subobjective. However, the nature of these rights, hardly commensurable, and their relatively large number makes this task demanding. One approach could be to create a hierarchy inspired in Maslow (1943) pyramid of needs. Most criticisms of this hierarchy focus on its last two categories, on what constitutes esteem and self-actualization or even whether the latter is more basic than the former. Table 3 provides an example of different impact levels over personal rights, using our modification of Maslow pyramid.

Alternatively, we could use as a proxy the number of legal actions against the organization from personal rights violations or the number of personal identifiable information records exposed in a cybersecurity event.

Table 3. Example of Personal Rights Impact Scenarios Constructed Scale

Rank	Injuries to personal rights
1	No personal rights violation
2	Violation of personal rights that may affect esteem and self-actualization needs
3	Violation of personal rights that may affect social belonging needs
4	Violation of personal rights that may affect safety needs
5	Violation of personal rights that may affect physiological needs, including safety needs that also affect physiological needs

3.4. Environmental Damage

As in Section 3.2, cyberattacks may trigger incidents with environmental impact (French and Ríos Insua 2000) for issues in environmental risk analysis. There are two relevant types of classifications for them: focusing either on the environmental impacts of normal operations or on those of incidents. For instance, the European eco-management and audit scheme (EMAS) (European Commission 2017) or the British environmental key performance indicators (UK Department for Environment, Food and Rural Affairs 2006) provide suggestions to assess the environmental impact of normal activities such as land use, energy efficiency, or emissions to air. These might be useful to identify impact scenarios in which the environmental performance of the organization is disrupted by a cyber incident. Additionally, frameworks like the Irish (Ireland Environmental Protection Agency 2010) and British Common Incident Classification Scheme (CICS) (UK Environment Agency 2006) facilitate the identification of environmental incidents such as the preservation of natural sites or habitats or contamination of water. They provide severity scores that might be helpful in deriving a constructed scale. However, they include impacts that we classify in other sections, such as human health or agricultural losses. Based on the British frameworks, we can suggest a constructed attribute for the environmental impacts. Table 4 provides a simple example.

Alternatively, the quantitative nature of environmental performance indicators might serve us to use them as proxy attributes. For example, we could use the variation in percentage of the most affected environmental indicator.

Table 4. Example of Environmental Damage Constructed Scale

Rank	Environmental damage
1	No environmental impact
2	Disturbance in the environmental performance indicators of the organization
3	Limited environmental damage, corresponding to CICS category 3 incidents
4	Significant environmental damage, corresponding to CICS category 2 incidents
5	Major environmental damage, corresponding to CICS category 1 incidents

3.5. Summary

Table 5 summarizes the cybersecurity risk management objectives and attributes that we could include in our cyber risk analyses. The advice would be to use the natural attribute whenever available; if not, use the constructed scale if it is not viewed as too ambiguous by the incumbent decision makers. Otherwise, the proxy attribute should be used (see Keeney and Gregory 2005 for further discussion). With this table, we facilitate this choice. Indeed, it highlights in bold the attributes that have actually been more useful to us in applications as validated in our workshops.

4. Models to Forecast Cybersecurity Impacts

When supporting risk management decisions, as in the cybersecurity resource allocation processes that we consider, besides assessing the consequences as described previously, we need models to forecast such consequences. The monetary attributes proposed may be dealt with relatively standard actuarial and financial models, as described in Eling and Wirfs (2019) and references therein.

We complement them by describing how to forecast the nonmonetary attributes of interest. We are interested mainly in assessing the entire forecasting distributions to be used later for risk management purposes. From the distribution we may eventually deduce the point estimates, should we need them for other purposes.

Throughout this section, we make several modeling assumptions. Our choices will be motivated by their mathematical tractability and physical or economical relevance and their flexibility, proving useful when expressing experts' opinions properly. In particular, the adopted models will depend on two parameters reasonably easily determined based on two expert judgements. This is especially important because we face a problem of lack of data, as, for reputational

Table 5. Summary of Objectives and Attributes.

Objective	Natural attribute	Constructed attribute	Proxy attribute
Min. operational costs	Monetary Units		
Min. income reduction			
Min. other costs			
Min. operational costs in other orgs.			
Min. income reduction in other orgs.			
Min. other costs in other orgs.			
Min. personal economic damage			
Min. reputation impact		Yes	Media salience
Min. reputation impact in other orgs.			Public relations cost
Min. fatalities	Number of fatalities		
Min. injuries to physical and mental health	No. injured people	Yes	No. people in hospital
Min. injuries to personal rights		Yes	No. legal actions against organization; No. personal identifiable information records exposed
Min. environmental damage		Yes	Percentage of variation in environmental indicator

reasons, companies are reluctant to reveal details when they are cyberattacked. Thus, we typically need to rely on structured expert judgement techniques (Dias et al. 2018). In particular, when available, we could use the information from several experts possibly aggregated through a weighted additive combination with weights depending on the experts' performance, for example, based on the classic model of Cooke (1991).

4.1. Service Unavailability

We start with forecasting service unavailability given a certain type of attack and the eventual security configuration of the organization. This is a relevant ingredient when forecasting income reduction and operational costs and for the same objectives in relation with third parties (impact to other organizations; Section 2.2.2), because of its relevance in supply chain cyber risk management.

Given its flexibility in modeling various shapes over positive distributions, we use a gamma distribution $f(i_s|\mu_1, \mu_2) \sim Ga(\mu_1, \mu_2)$ for the duration i_s of the downtime. When the distribution is deemed to be multimodal, for example, because of having downtimes with very different origins (types of threats), we could use a mixture of gammas (Wiper et al. 2001) and proceed in a similar fashion as we now describe. The objective would be to obtain good estimates of μ_1 and μ_2. Given the lack of data, we obtain them through expert judgement. For this, because we have to assess two parameters, we may ask just two judgements of our experts. For example, we may ask the experts for the first (p_{25}) and third (p_{75}) quartiles of the i_s distribution and infer the parameters by solving

$$\min_{\mu_1,\mu_2} \left\{ \left(p_{25} - \mathrm{cdf}(.25,\, \mu_1,\, \mu_2)\right)^2 + \left(p_{75} - \mathrm{cdf}(.75,\, \mu_1,\, \mu_2)\right)^2 \right\}, \tag{1}$$

where $\mathrm{cdf}(\cdot, \mu_1, \mu_2)$ designates the cumulative distribution function (cdf) for the quantile of interest of the gamma distribution with parameters μ_1 and μ_2. Morris et al. (2014) provide web-based software to perform these computations. We then undertake consistency checks based on other quantiles. As service unavailability data become available, we could incorporate such information with the data updating the distribution in a Bayesian fashion (French and Ríos Insua 2000). Finally, when necessary, we could approximate the downtime, for example, through the expected value of the distribution.

For multiattribute aggregation purposes, we would typically multiply the downtime duration by the estimated expected cost of each unavailable unit. Recall that this would cover only unavailability operational costs, and we could include other impacts similar to those in relation with reputation. For example, if the unavailability period is too long, the organizational reputation could be severely damaged.

4.2. Reputation

As discussed in Section 3.1, there is no natural attribute that allows us to assess this impact. Our focus will therefore be on its consequence: the loss in market share induced by an attack over the organization. We designate it by l_s, the proportion of customers abandoning to a competitor.

The following general considerations can be made. If the organization is dominant in such service, the loss in market share would be negligible, and we assume that $l_s = 0$. On the contrary, if there are alternative service suppliers, the market loss could be nonnegligible, essentially depending on the reputation loss: the bigger it is, the bigger l_s will be. Given its flexibility in modelling distributions with support [0,1], and having only two parameters, we assume that l_s follows a beta distribution with parameters α and β. We would perform an assessment for each segmentation level, as those in Table 1. We would proceed in a similar fashion to Section 4.1 by asking two quartiles to experts and, subsequently, approximating the parameters, based on a least-squares cdf approximation, after appropriate consistency checks. Again, we could introduce schemes to learn about α and β through Bayesian updating, as data become available. Finally, when necessary, we could summarize the predicted proportion of lost customers through, for example, its mean.

All in all, for aggregation purposes, the loss could be quantified as $l_s \times k \times n \times c$, where k is the current market share, n is the market size, and c is the income produced per customer in the relevant risk management period, which would be available from in house accounting experts.

A similar approach could be adopted for the reputational impact to third parties.

4.3. Fatalities

The approach adopted considers that a cyberattack triggers an incident that causes the failure of a cyber-physical system, possibly leading to fatalities. As a consequence, in a given application, we need to forecast the number of fatalities associated with an accident in the corresponding system (besides the probability that the cyber event triggers the accident).

We could use an expert judgement-based approach using, for example, a Poisson distribution by asking several quantiles as we did in Section 4.1. As an alternative, we could adapt the fatality forecasting model for aviation safety accidents in Rios Insua et al. (2019a). To predict the number n_F of fatalities in an accident triggered by a cyberattack, we use a model $n_F = p_F \cdot q \cdot M$,

where p_F designates the proportion of fatalities; M, the maximum occupancy of the site that is the maximum number of persons allowed in it; and finally, q, the site occupancy degree, that is, the percentage over the maximum occupancy available at attack time. The parameters would depend on the type of site. Acknowledging that some attacks designed to kill people may actually kill none, some, or all of the occupants of the installation, we use a mixture model for the proportion p_F, $\tau_1 I_0 + \tau_2 \mathcal{B}e(a,b) + \tau_3 I_1$, where τ_1, τ_2, and τ_3, respectively, designate the proportions of accidents triggered by cyber incidents with no fatalities, both fatalities and survivors, and finally, with no survivors, being $\tau_1 + \tau_2 + \tau_3 = 1$, $\tau_i \geq 0, i = 1,2,3$. I_0 is the degenerate distribution at 0 (no one dies); $\mathcal{B}e(a,b)$ models the distribution of the proportion of fatalities in accidents when there are fatalities and survivors; and finally, I_1 is the degenerate distribution at 1 (all die). A priori, $(\tau_1, \tau_2, \tau_3) \sim \mathcal{D}ir(a_1, a_2, a_3)$, $p_F \sim \mathcal{B}e(a,b)$, a standard assumption in mixture models (Wiper et al. 2001). For the occupancy proportion q, the prior distribution would be $q \sim \mathcal{B}e(c,d)$, with the same choice arguments as in Section 4.2.We would assess all these parameters with expert judgement as in Section 4.1. In the presence of data, we may make inferences about the weights τ_i with a Dirichlet-multinomial model, and about p_F, when $0 < p_F < 1$, and q with Beta-binomial models (French and Ríos Insua 2000).

Finally, we use the concept of statistical value of life (Viscusi and Aldy 2003) to aggregate the fatalities in the multiattribute model.

4.4. Injuries

Some cyberattacks might produce injuries. As in Section 3.2, we distinguish between minor and major injuries. We consider three proportions p_{h_i}, $i = 1,2,3$, for the three types of survivors ($i = 1$, minor injured; $i = 2$, major injured; $i = 3$, uninjured), following a model

$$p_H = (p_{h_1}, p_{h_2}, p_{h_3}) \sim \alpha \cdot I(0,0,1) + (1-\alpha) \cdot \mathcal{D}ir(h_1, h_2, h_3),$$

where α designates the occurrence proportion in which none is injured, and $I(0,0,1)$ is the degenerate distribution in which there are no injuries. We assume an $\alpha \sim \mathcal{B}e(a,b)$ prior assessed with expert judgement; as data become available, we would learn about α with a Beta-binomial model. Similarly, we would proceed about the proportions of injured occupants with a Dirichlet-multinomial model. Then, the number $n_H = (n_{h_1}, n_{h_2}, n_{h_3})$ of injuries for an occurrence is predicted with a model

$$n_H = p_H \cdot q \cdot (1 - p_F) \cdot M,$$

where p_H designates the proportions of injuries and p_F, q, and M are as in Section 4.3. We could use the concept of statistical value of an injured person to aggregate the injuries in our multiattribute model (European Organisation for the Safety of Air Navigation 2013).

4.5. Personal Rights

There are several models in the literature related to forecasting incidents that can impact personal rights, such as the chance of cybersecurity breaches (Liu et al. 2015, Sarabi et al. 2016) or a proposal regarding litigation forecasting (Brown et al. 2004).

We focus on forecasting the number n_p of personal identifiable information records exposed in an attack as a proxy. For that, we use a model $n_p = q_p N_p$, where q_p is the proportion of exposed records and N_p is the maximum number of records available. A priori, we could use a beta model for q_p and introduce a beta-binomial model to learn about it as data accumulate. Similarly to the minimization process (1), we would infer the distribution parameter for the quantity q_p through expert judgement. We could segment the model for q_p, for example, depending on the economic sector considered or other organizational features.

Our example in Section 5 describes a method to aggregate this impact by deriving an economic valuation of an exposed record.

4.6. Impact over the Environment

As with fatalities, we could view the cyberattack as a triggering event that would start an environmental accident and use a pertinent environmental risk assessment model. Forecasting environmental impacts depends on the specifics of the sector and territory under assessment (European Food Safety Authority 2017). A general method that help forecasting environmental impacts is life cycle assessment (LCA) (Hellweg and Milá i Canals 2014). Fjeld et al. (2007) provide an in-depth treatment of environmental impact risk assessment and management.

5. Utility Model

From the comprehensive list of cybersecurity objectives identified in Section 2, the incumbent organization would choose those objectives that are relevant in its cybersecurity risk management problem, possibly eliminating some of them and/or adding others.

Recall that our final aim is to develop a decision support system for cybersecurity risk management resource allocation; therefore, it is convenient to develop also a generic preference model over such impacts, as we do now through a multiattribute utility function. For this, we use the classic concepts of measurable multiattribute value function (Dyer and Sarin 1979) and relative risk aversion (Dyer and Sarin 1982). Essentially, we aggregate first the objectives with an additive multiattribute value function and then use an exponential utility function over the values if the manager is (constant absolute) risk averse or risk prone or a linear utility function if he/she is risk neutral. A detailed review with relevant underlying structures, independence conditions, and elicitation protocols for the weights, component value functions, and risk aversion coefficient may be seen in Ortega et al. (2018), with pointers that justify the appropriateness of such structure in Keeney (2007b) and Keeney and von Winterfeldt (2011). Rather than repeating their arguments, we provide an actual example from a case study.

5.1. Case Study

We integrated the previous developments within the CYBECO Toolbox described in detail in Couce et al. (2019). Such a decision support system aims at facilitating cybersecurity resource allocation processes (including cyber-insurance purchase) to SMEs that are nonintensive in IT use. We describe here the built-in preference model and sketch its use and relevance. As mentioned in the Introduction, the proposed objectives tree aims at being generic. We also mentioned how different organizations might hold different views on such tree, discarding and/or adding further objectives, as we illustrate now.

For system development, several relevant cybersecurity scenarios for SMEs (Musaraj et al. 2018) were identified and later validated in the previously mentioned workshops. In the scenarios, we covered and synthesized a number of cyber threats and impacts relevant for the SMEs of interest. We map them in Table 6 against our objectives from Section 2. All impacts are linked with monetary costs, except the loss of personal records, which is linked with damage to personal rights.

Assuming a risk-averse organization, if we apply the utility function sketched previously and detailed in Ortega et al. (2018), we use

$$u(m,r) = 1 - \exp\big(-\rho(v_m(m) + v_r(r))\big),$$

Table 6. CYBECO Impacts Mapped onto Our Objectives

Impacts in CYBECO Toolbox	Cybersecurity objective
Facilities: damages to physical properties	Min. operational costs
IT infrastructure: business downtime	Min. operational costs
Market share: percentage lost	Min. income reduction
Personal information: records lost	Min. injuries to personal rights
Personal information: privacy and security liability lost	Min. other costs
Customers: loss of customers due to brand reputation and damage	Min. income reduction
Production: interruption of provided services or products	Min. income reduction
Contractual and regulatory losses	Min. other costs
Recovery and other postincident expenses	Min. operational costs
	Min. cybersecurity costs

where m is the monetary impact, r is the impact on personal rights, and v_m and v_r are the component value functions.

The first objective, m, is measured through a natural attribute (monetary units) that we express in euros. This also includes the security costs of cybersecurity controls and cyber insurance, because they are related to the objective *Min. cybersecurity costs*. The second one, r, is measured with a proxy attribute (number of records exposed), associated with a parameter c_r, which monetizes them. Therefore, the utility function finally used is described as

$$u(m,r) = 1 - \exp\big(\rho(m + c_r r)\big).$$

To assess c_r, we should provide an economic value to privacy. The legal costs of injuries to personal rights are part of the monetary costs. However, there is no solid estimation for the *value of privacy* (Acquisti et al. 2013). Estimates based on British (Godel et al. 2017) and American (Hann et al. 2007) customers reveal that consumers' value of their personal information is up to £7.25 and $44.62, respectively. Assuming that they are risk neutral and they assign a probability of less than 1% to data exposure, then, taking the more conservative British figure (equivalent to € 8.25), their personal information should at least be valued at € 825; risk aversion would reduce this figure slightly. On the other hand, the American figures of a lower perception of the likelihood would increase it (e.g., more than € 4,000 with the American figures or € 1,650 if we assume a probability of breach of less than 0.5%). Thus, we use € 825 as a conservative estimate of the economic value of privacy per record.

Then, the utility function that we use adopts the form

$$u(m,r) = a \times \big(1 - \exp\big(\rho(m + 825r) + b\big),$$

with a and b chosen to scale it at [0,1]. To adjust it, we determine the worst reasonable cost that is $m_* + 825r_*$, where m_* is the sum of the maximum cost of the impacts and the cybersecurity expenditure, and r_* is the maximum number of records that can be exposed. For a specific organization, m_* was estimated at € 2,000,000 and r_* at 5,000, so that the worst cost is € 6,125,000. The best cost holds for $m^* = r^* = 0$. For such SME, we assessed one point in their utility function obtaining for c_1 (half of the worst cost) $u(c_1) = 0.8$, with the aid of the probability equivalent method (Farquhar 1984). Simple computations lead to the assessments $a = 1/15 = 0.066$, $\rho = 4.5267 * 10^{-7}$ and $b = 1$, and the utility function adopted is

$$u(m,r) = 0.066 * \big(1 - \exp\big(4.5267 * 10^{-7}(m + 825r)\big)\big) + 1.$$

We use it as the default utility in the CYBECO toolbox mentioned previously when we deal with non-sophisticated defenders. For sophisticated users, we allow them to assess their own utility function much as we have shown, based on determining m_*, r_*, and $u(c_1)$.

5.2. Uses

As mentioned, the previous generic preference model is implemented in our tool together with the forecasting models. An organization willing to use it introduces their defining parameters. This includes (i) the assets that they consider subject to cyber threats (e.g., their market share, their computer equipment, and their PII records); (ii) the threats that they envision as relevant to their cybersecurity (e.g., computer virus, flood, misconfiguration, DDoS attack by a competitor, DDoS or social engineering attack by a cybercriminal); (iii) the impacts of relevance from the previous list; and (iv) their budget, technical, and compliance constraints. With this information, we may proceed on to performing their cybersecurity risk management resource allocation, which would lead to budget allocation among controls and insurance products that best protect (in the sense of maximizing expected utility) the organization from all the threats that they deem relevant.

For this, given a cybersecurity portfolio p (typically composed of several control measures, such as installing a firewall, introducing an intrusion detection system and a sprinkler system, and acquiring insurance products, such as conventional product A and cyber insurance product B), the system makes probabilistic forecasts of the attacks to be received and their impacts, based on built-in parameterized forecasting models, summarized in a model $f(m,r|p)$ that describes the distribution over monetary losses m and lost records r that the organization might suffer, given that portfolio p has been implemented. Such impacts are assessed with the utility function $u(m,r)$ described previously, which is used to estimate the expected utility associated with the incumbent portfolio:

$$\psi(p) = \iint u(u,r) f\big(m,r|p\big)\, dm\, dr.$$

Based on this procedure to assess the portfolios' expected utilities, the system looks for the portfolio with maximum expected utility numerically solving the problem: $\max \psi(p)$. The system proposes the optimal portfolio p^* to the user, together with some sensitivity analysis information. Rios Insua et al. (2019b) and Couce et al. (2019) provide detailed examples.

6. Discussion

Cybersecurity risk management decisions require the definition of the cyberattack impacts relevant for an incumbent organization. This has been traditionally dealt with from a technical perspective focusing on the confidentiality, integrity, and availability triad. This is useful for assessing digital systems from an information systems perspective but still needs a translation to what really represents for an organization in business and operational terms. Other standards, such as Magerit (2012) or SABSA Institute (2009), which we analyze in the appendix, provide a more organizational perspective, although we consider their view of what is an organization rather than general, closer to the archetype of a big business corporation or a public agency.

We therefore provided a wider organizational-oriented perspective, delivering a generic tree of objectives for such purpose. From it, the incumbent organization may choose which are the relevant objectives for them or, alternatively, use it to complete their own objectives. We also provided the corresponding attributes that measure or estimate objective achievement, as well as generic models to forecast and assess them. These tools facilitate the formulation of the Defender's preference model by responding a few simple questions and have an orientation on the forecasting models to be implemented, facilitating his cyber risk management decision analysis. The importance of the simplification is brought by such concepts. Without them, the cybersecurity manager could find the task of preference elicitation too complicated cognitively. The proposed model is embedded in our CYBECO tool (Couce et al. 2019), which facilitates cybersecurity resource allocation processes to SMEs at the strategic and tactical levels.

In future work, we shall undertake a similar approach for attackers. This involves the generation of cyber-attacker objective trees, based on the assessment of different types of attackers (e.g., nation states, cybercriminals) and their motivations to undertake cyberattacks (e.g., financial, espionage). It also involves the use of random utility functions and random distributions when it comes to adding uncertainty about attacker preferences and information, in the spirit of adversarial risk analysis (Rios Insua et al. 2019b).

Acknowledgment

The authors thank the referees for excellent and constructive discussions.

Appendix

In this appendix, we compare our tree of cybersecurity objectives with the valuation criteria of two main standards, Magerit and SABSA.

A.1. Mapping Magerit Valuation Criteria to Our Cybersecurity Objectives Tree

Table A.1. Magerit Mapping to Cybersecurity Objectives Tree

Magerit (2012)	Cybersecurity objectives tree
Personal information	Injuries to personal rights (impacts to persons), Other costs (impacts to organization due to noncompliance regarding personal information) and Operational costs (information asset degradation).
Legal obligations	Other costs
Security	Cybersecurity costs
Commercial or economic interests	Income reduction or other costs (if strategic)
Service interruption	Operational costs
Public order	For most organizations, *Impact to other organizations*. For those organizations responsible for public order it might be necessary to create a new objective of nonmonetary nature to evaluate the potential states of public order: *Max. public order*.

Table A.1. (Continued)

Magerit (2012)	Cybersecurity objectives tree
Operations	Operational costs
Administration and management	Operational costs
Loss of confidence (reputation)	Reputation impact
Prosecution of crimes and law enforcement	For most organizations, *impact to other organizations*. For those organizations responsible for these tasks it is related with Operational costs
Service recovery time	Operational costs
Classified information	As a characteristic of information assets, Operational costs

Note. Some of the proposed objectives are missing within Magerit: fatalities, injuries to physical and mental health, personal economic damage, and environmental damage.

A.2. Mapping SABSA High-Level General Business Attributes to Our Cybersecurity Objectives Tree

Table A.2. SABSA Mapping to Cybersecurity Objectives Tree

SABSA Institute (2009)	Cybersecurity objectives tree
Financial—accounted	Other costs
Financial—AML compliant	Other costs
Financial—auditable	Other costs
Financial—benefit-evaluated	Income reduction
Financial—cash-flow forecast	Income reduction
Financial—credit controlled	Other costs
Financial—credit risk managed	Other costs
Financial—investment returnable	Other costs
Financial—liquidity risk managed	Other costs
Financial—market risk managed	Other costs (understood as financial market risks)
Financial—profitable	Income reduction
Financial—reporting compliant	Other costs
Physical (all attributes)	Operational costs. Note that some characteristics are related to security/risk characteristics of the assets (access controlled, damage protected, defended, secure, theft protected).
Human (all subattributes)	Characteristics related to human capital, which could be classified as an asset; therefore, the related objective is operational costs
Process (all subattributes)	Other costs
Strategic—administered	Other costs
Strategic—branded	Other costs
Strategic—communicated	Other costs
Strategic—competitive	Other costs
Strategic—compliant	Other costs
Strategic—financed	Other costs
Strategic—goal oriented	Other costs
Strategic—governed	Other costs
Strategic—logistically managed	Operational costs
Strategic—market penetrated	Income reduction
Strategic—market positioned	Income reduction
Strategic—reputable	Reputation impact.
Strategic—supply chain managed	Operational costs
System (all attributes)	Operational costs. Note that some characteristics are related to security/risk characteristics of the assets (access controlled, incident managed, risk managed).

Note. As with Magerit, some of the proposed objectives are missing in SABSA: fatalities, injuries to personal rights, injuries to physical and mental health, personal economic damage and environmental damage.

Endnote

[1]Standards in the ISO 22300 family are the continuation of BS 25999 (2007), one of the most popular standards in business continuity management.

References

Acquisti A, Leslie KJ, Loewenstein G (2013) What is privacy worth? *J. Legal Stud.* 42(2):249–274.

Andress J, Winterfeld S (2013) *Cyber Warfare: Techniques, Tactics and Tools for Security Practitioners* (Elsevier, New York).

Bagchi A, Bandyopadhyay T (2018) Role of intelligence inputs in defending against cyber warfare and cyberterrorism. *Decision Anal.* 15(3):174–193.

Brewster T (2018) This guy hacked hundreds of planes from the ground. *Forbes* (August 9), https://www.forbes.com/sites/thomasbrewster/2018/08/09/this-guy-hacked-hundreds-of-planes-from-the-ground/#21804f4946f2.

British Standards Institution (2007) *BS 25999-2:2007 Specification for Business Continuity Management* (British Standards Institution, London).

Brown S, Hillegeist SA, Lo K (2004) *Management Forecasts and Litigation Risk* (Elsevier, Amsterdam).

Brownlow S, Watson S (1987) Structuring multi-attribute value hierarchies. *J. Oper. Res. Soc.* 38(4):309–317.

Carpenter DP, Krause GA (2011) Reputation and public administration. *Public Admin. Rev.* 72(1):26–32.

Cooke RM (1991) *Experts in Uncertainty: Opinion and Subjective Probability in Science* (Oxford University Press, Oxford, UK).

Cooke RM, Bedford T (2001) *Probabilistic Risk Analysis: Foundations and Methods* (Cambridge University Press, Cambridge, UK).

Couce A, Rios Insua D, Koutalieris G, Chatgiannakis V (2019) Decision support systems for cybersecurity risk management and cyber insurance. Accessed August 15, 2019, https://www.cybeco.eu/images/items/CYBECO-D3.2_Improved%20Modelling%20framework%20for%20cyber%20risk%20management_v2.0.pdf.

Dias LC, Morton A, Quigley J (2018) *Elicitation: State of the Art and Science* (Springer, New York).

Dyer J, Sarin R (1979) Group preference aggregation rules based on strength of preference. *Management Sci.* 25(9):822–832.

Dyer J, Sarin R (1982) Relative risk aversion. *Management Sci.* 28(8): 875–886.

Eling M, Wirfs J (2019) What are the actual costs of cyber risk events. *Eur. J. Oper. Res.* 272:1109–1119.

ENISA (2007) Information package for SMEs with examples of risk assessment/risk management for two SMEs. Accessed August 15, 2019, https://www.enisa.europa.eu/publications/information-package-for-smes/at_download/fullReport.

ENISA (2010) IT business continuity management—An approach for small medium sized organisations. Accessed September 24, 2020, https://www.enisa.europa.eu/publications/business-continuity-for-smes/at_download/fullReport.

European Commission (2017) Commission Decision (EU) 2017/2285 of December 6, 2017 Amending the User's Guide Setting Out the Steps Needed to Participate in EMAS, under Regulation (EC) No 1221/2009 of the European Parliament and of the Council on the Voluntary Participation by Organizations in a Community Eco-Management and Audit Scheme (EMAS). Legislation, Publications Office of the European Union, Luxembourg.

European Food Safety Authority (2017) EFSA guidance document for predicting environmental concentrations of active substances of plant protection products and transformation products of these active substances in soil. *EFSA Journal* 13(4):4093.

European Organisation for the Safety of Air Navigation (2013) Annual report. Report, European Organisation for the Safety of Air Navigation (EUROCONTROL), Brussels, Belgium.

European Parliament (2016) Regulation (EU) 2016/679 of the European Parliament and the Council of April 27, 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation). Legislation, Publications Office of the European Union, Luxembourg.

Farquhar PH (1984) State of the art—Utility assessment methods. *Management Sci.* 30(11):1283–1300.

Fjeld R, Eisenberg N, Compton K (2007) *Quantitative Environmental Risk Analysis for Human Health* (Wiley, New York).

Fombrun CJ (2012) The building blocks of corporate reputation: Definitions, antecedents, consequences. Barnett ML, Pollock TG, eds. *The Oxford Handbook of Corporate Reputation* (Oxford University Press, Oxford, UK), 94–113.

French S, Ríos Insua D (2000) *Statistical Decision Theory* (Wiley, New York).

Fu K, Blum J (2013) Controlling for cybersecurity risks of medical device software. *Comm. ACM* 56(10):35–37.

Godel M, Landzaat W, Suter J (2017) Research and analysis to quantify the benefits arising from personal data rights under the GDPR. Report, Department for Culture, Media & Sport, London, UK.

Hann Ih, Kai-Lung H, Sang-Yong TL, Ivan PLP (2007) Overcoming information privacy concerns: An information processing theory approach. *J. Management Informs. Systems* 24:13–42.

Hasler RM, Kehl C, Exadaktylos AK, Albrecht R, Dubler S, Greif R, Urwyler N (2012) Accuracy of prehospital diagnosis and triage of a Swiss helicopter emergency medical service. *J. Trauma Acute Care Surgery* 73(3):709–715.

Hellweg S, Milá i Canals L (2014) Emerging approaches, challenges and opportunities in life cycle assessment. *Science* 344(6188): 1109–1113.

Hubbard DW, Selersen R (2016) *How to Measure Anything in Cybersecurity Risk* (Wiley, New York).

Industry Specification Group (2015) ETSI GS ISI 002 V1.1.1 information security indicators (ISI); event model, a security event classification model and taxonomy Annex B1.8—With what kind of impact. European Telecommunications Standards Institute, Sophia Antipolis, France.

Ireland Environmental Protection Agency (2010) Guidance to licensees/COA holders on the notification. Management and Communication of Environmental Incidents. Report, Office of Environmental Enforcement, Wexford, Ireland.

ISO (2014) ISO 55000:2014—Asset management—Overview. Principles and terminology. Report, International Organization for Standardization, Geneva, Switzerland.

ISO (2015a) ISO 19770-5:2015—IT asset management—Overview and vocabulary—Part 5. Report, International Organization for Standardization, Geneva, Switzerland.

ISO (2015b) Societal Security—Business Continuity Management Systems—Guidelines for Business Impact Analysis. Report, International Organization for Standardization, Geneva, Switzerland.

Jensen M, Roy A (2008) Staging exchange partner choices: When do status and reputation matter? *Acad. Management J.* 51(3): 495–516.

Jensen M, Kim H, Kim BK (2012) Meeting expectations: A role-theoretic perspective on reputation. Barnett ML, Pollock TG, eds. *The Oxford Handbook of Corporate Reputation* (Oxford University Press, Oxford, UK), 140–159.

Keeney R (1992) *Value Focused Thinking* (Harvard University Press, Cambridge, MA).

Keeney R (2007a) Developing objectives and attributes. Edwards W, Miles RF Jr, von Winterfeldt D, eds. *Advances in Decision Analysis: From Foundations to Applications* (Cambridge University Press, Cambridge, UK).

Keeney R (2007b) Modeling values for anti-terrorism analysis. *Risk Anal.* 27(3):585–596.

Keeney R, Gregory R (2005) Selecting attributes to measure the achievement of objectives. *Oper. Res.* 53:1–11.

Keeney R, von Winterfeldt D (2011) A value model for evaluation homeland security decisions. *Risk Anal.* 31(9):1470–1487.

Krutz R, Vines R (2004) *The CISP Prep Guide* (Wiley, New York).

Kurtz C, Semmann M, Schulz W (2018) Toward a framework for information privacy in complex service ecosystems. *39th Internat. Conf. Inform. Systems* (Association for Information Systems, San Francisco, CA).

Lee JA, Liu CU (2012) Forbidden city enclosed by the great firewall: The law and power of Internet filtering in china. *Minnesota J. Law Sci. Tech.* 13(1):125–151.

Liu Y, Sarabi A, Zhang J, Naghizadeh P, Karir M, Bailey M, Liu M (2015) Cloudy with a chance of breach: Forecasting cyber security incidents. *24th USENIX Security Sympos.* (The USENIX Association, Berkeley, CA), 1009–1024.

Macaulay T, Singer BL (2011) *Cybersecurity for Industrial Control Systems: SCADA, DCS, PLC, HMI, and SIS* (Auerbach Publications, Abingdon-on-Thames, UK).

Magerit (2012) *Metodología de Análisis y Gestión de Riesgos de los Sistemas de Información,* version 3 (Ministerio de Hacienda y Administraciones Públicas, Madrid).

Margulies P (2013) The NSA in global perspective: Surveillance, human rights, and international counterterrorism. *Fordham Law Rev.* 82(5):2137–2167.

Maslow AH (1943) A theory of human motivation. *Psych. Rev.* 50(4):370–396.

Morris D, Oakley J, Crowe J (2014) A web-based tool for eliciting probability distributions from experts. *Environ. Model. Software* 52:1–4.

Mowbray TJ (2013) *Cybersecurity: Managing Systems, Conducting Testing, and Investigating Intrusions* (Wiley, New York).

Musaraj Cousin, K M, Melvin V, Melvin C, Couce A, Rios Insua D, J Vila, et al. (2018) CYBECO Deliverable D4.1 Cyber insurance use cases and scenarios. Accessed March 30, 2018, https://www.cybeco.eu/.

OECD (2017) *Enhancing the Role of Insurance in Cyber Risk Management* (OECD Publishing, Paris).

Ortega J, Radovic V, Rios Insua D (2018) Utility elicitation. Dias LC, Morton A, Quigley J, eds. *Elicitation: The Science and Art of Structuring Judgement* (Springer International Publishing, New York), 241–264.

Pala A, Zhuang J (2019) Information sharing in cybersecurity: A review. *Decision Anal.* 16:157–237.

Raggio R, Leone R (2019) Drivers of brand value, estimation of brand value in practice and use of brand valuation: Introduction to the special issue. *J. Brand Management* 17(1):1–5.

Rios Insua D, Alfaro C, Gomez J, Hernandez-Coronado P, Bernal F (2019a) Forecasting and assessing consequences of aviation safety occurrences. *Safety Sci.* 111:243–252.

Rios Insua D, Couce-Vieira A, Rubio JA, Pieters W, Labunets K, Rasines D (2019b) An adversarial risk analysis framework for cybersecurity. *Risk Anal*, ePub ahead of print June 10, https://doi.org/10.1111/risa.13331.

SABSA Institute (2009) The SABSA White Paper (Sherwood Applied Business Security Architecture), Hove, UK.

Sarabi A, Naghizadeh P, Liu Y, Liu M (2016) Risky business: Fine-grained data breach prediction using business profiles. *J. Cybersecurity* 2(1):15–28.

Sayfayn N, Madnick S (2017) Cybersafety analysis of the Maroochy shire sewage spill. Working paper, MIT Sloan School of Management, Cambridge, MA.

Taeihagh A, Lim HSM (2018) Governing autonomous vehicles: Emerging responses for safety, liability, privacy, cybersecurity, and industry risks. *Transporation Rev.* 39(1):103–128.

The Open Web Application Security Project (2017) The OWASP risk rating methodology. Accessed September 23, 2020, https://owasp.org/www-community/OWASP_Risk_Rating_Methodology#.

Torres A, Redondo A, Rios Insua D, Domingo J, Ruggeri F (2020) Expert judgement methods in a supply chain cyber risk management. Hanea AM, Nane GF, Bedford T, French S, eds. *Expert Judgement in Risk and Decision Analysis* (Springer International Publishing, Cham, Switzerland).

UK Department for Environment, Food and Rural Affairs (2006) Environmental key performance indicators reporting guidelines for UK business. (DEFRA, London). Report, Department for Environment, Food and Rural Affairs, London.

UK Environment Agency (2006) Incidents and their classification: The Common Incident Classification Scheme (CICS), version 12. Report, Environment Agency, London.

UN Human Rights Council (2015) Resolution on the promotion, protection and enjoyment of human rights on the internet. Report, United Nations, Geneva, Switzerland.

UN Human Rights Council (2016) Universal Human Rights Index Database. Accessed October 2016, http://uhri.ohchr.org/search/guide.

Ustün T, Kostanjsek N, Chatterji S, Rehm J (2010) Measuring health and disability: Manual for WHO disability assessment schedule, WHODAS 2.0. Report, World Health Organization, Geneva, Switzerland.

Vacca J (2013) *Computer and Information Security Handbook*, 2nd ed. (Morgan Kaufmann, Burlington, MA).

van Riel CBM, Fombrun CJ (2007) *Essentials of Corporate Communication* (Routledge, Abingdon, UK).

Vandebosch H, van Cleemput K (2008) Defining cyberbullying: A qualitative research into the perceptions of youngsters. *Cyberpsych. Behav.* 11(4):499–503.

Viscusi K, Aldy J (2003) The value of a statistical life: A critical review of market estimates throughout the world. *J. Risk Uncertainty* 27:5–76.

WEF (2020) *Global Risks Report* (World Economic Forum, Cologny, Switzerland).

WHO (2018) *International Statistical Classification of Diseases and Related Health Problems*, 11th revision (World Health Organization, Geneva).

Wiper M, Rios Insua D, Ruggeri F (2001) Mixtures of gamma distributions with applications. *J. Comput. Graphic Statist.* 10:440–454.

Aitor Couce-Vieira has a BSc in economics and PhD and MSc in decision systems. He is a research assistant at the

Instituto de Ciencias Matemáticas (ICMAT), working in the H2020 project Supporting CYber Insurance from a BEhavioural ChOice Perspective (CYBECO) and participating in the development of decision support models for analyzing the risk of cybersecurity incidents. His research interests include risk analysis and decision support.

David Rios Insua is AXA-ICMAT Chair in adversarial risk analysis at ICMAT-Consejo Superior de Investigaciones Científicas (CSIC) and visiting professor at the School of Management, University of Shanghai for Science and Technology. He is a member of the Spanish Royal Academy of Sciences. His current research interests cover decision analysis, risk analysis, and adversarial risk analysis and their applications to aviation safety, cybersecurity, and adversarial machine learning.

Alex Kosgodagan is a postdoctoral fellow at ICMAT-CSIC who received a PhD in 2017 from the Institut Mines-Telecom (IMT)-Atlantique in probabilities, statistics, and operations research. His research interests lie at the boundaries of such domains as stochastic processes, high-dimensional statistical learning, and decision and information theory. He is involved in prediction-oriented methods, which he first put into practice in financial mathematics to subsequently move reliability and risk analysis.