TrustCoM

Deliverable

15

Report on Legal Issues Appendix C

Analysis of Legal Issues Related to Data Protection Law

WP9 Legal Issues

Tobias Mahler, Fredrik Vraalsen (eds.)

TrustCoM

A trust and Contract Management framework enabling secure collaborative business processing in on-demand created, self-managed, scalable, and highly dynamic Virtual Organisations

SIXTH FRAMEWORK PROGRAMME

PRIORITY IST-2002-2.3.1.9

Prepared for the European Commission under FP6 Contract No. 01945 as a deliverable form

Deliverable datasheet

Project acronym: TrustCoM

Project full title: *A trust and Contract Management framework enabling secure collaborative business processing in on-demand created, self-managed, scalable, and highly dynamic Virtual Organisations*

Action Line:	**6**
Activity:	**6.2**
Work Package:	**9**
Task:	**6.2.1**

Document title:	D 15 Appendix C
Version:	1.0
Document reference:	
Official delivery date:	31/07/2005
Actual publication date:	N/A
File name:	
Type of document:	Report
Nature:	Public

Authors:	Jon Bing[1], Andrew Jones[2], Mass Soldal Lund[3], Tobias Mahler[1], Thomas Olsen[1], Xavier Parent[2], Ketil Stølen[3], Fredrik Vraalsen[3]
Reviewers:	CCLRC, Heather Weaver and Judy Beck

Approved by:

[1] NRCCL

[2] KCL

[3] SINTEF

Table of Content

1 Introduction

This appendix presents an analysis of legal issues in Virtual Organisations (VOs) in a scenario based on the Virtual Community (VC) scenario developed in WP11.

The research covered in this Appendix was concluded prior to the work described in Appendices A and B. The experiences with the analysis of the VC scenario formed the methodological basis for the following studies, where legal risk analysis was pursued in an integrated way.

This Appendix includes three separate analyses with a focus on data protection law: The conventional legal analysis focuses on the data protection issues involved when a VO uses a reputation system. Based on this conventional legal analysis, we have investigated the application of formal-logical tools to the formalization of some specific data protection regulations. Finally we have demonstrated how a risk analysis of legal issues can use existing risk analysis methods and tools.

Sections 3, 4 and 5 contain an analysis of data protection issues in relation to VOs. The first objective of this analysis is to demonstrate the feasibility of the selected methods in a practical scenario. The second objective of this analysis was the analysis of legal issues in VOs with relevance to TrustCoM. The conventional legal analysis in section 3 focuses on the data protection issues involved when a VO uses a reputation system. Based on this conventional legal analysis, section 4 investigates the application of formal-logical tools to the formalization of some specific data protection regulations. Section 5 shows how a legal risk analysis can use existing risk analysis methods and tools.

2 Scenario for the Analysis

This section introduces the factual basis for the analysis of legal issues. We have selected and adapted a scenario provided by the TrustCoM WP 11. The description of this scenario can be found in section 2.1. The original scenario was slightly amended and made more specific, in order to facilitate a specific legal analysis.

This scenario raises a number of legal questions, which are analysed from a legal point of view in section 3, then formalized in section 4, followed by a (legal) risk analysis in section 5.

2.1 The Virtual Community Scenario (VC)

The "Virtual Community" (VC) has set up a `Credit Union', that is, a system whereby members of the community can pay money into and withdraw money from a centralized fund. Credit Unions are popular today with disadvantages groups of people, for example poor third world farmers or unemployed tenants in a local housing authority. These people are usually very disadvantaged and unable to get the credit that is available to normal salaried people. They often have to resort to loan sharks who charge inordinate amounts of interest, thereby disadvantaging the people even more.

The VC website is started by a charity which helps recovering drug addicts get back on their feet and re-establish themselves in society. They have formed a small but lively and cohesive virtual community based around the website’s online forums. The website’s founder, Maggie May (M), having had experience of such things before, has established the credit union facility for members of the group, allowing them to help each other financially to add to the moral support they already give each other. The credit union is an agreement between a trusted group of people whereby each member of the group pays money into a central fund, allowing other members to take out loans from that fund that they may not have been able to get elsewhere. The experience M has had before has been based in communities formed on a tight geographical basis, where members knew each other, and there was thus a strong incentive to abide by the norms of the group. However, in this case, the credit union is established in Oslo, but members are dispersed over all the Nordic countries.

The issues of trust are dealt with, as new members of the union must engage with the community and build up their reputations with the group’s membership, in order to be allowed to loan money from the union.

After joining, new members can begin paying into and eventually withdrawing from the central fund of the credit union. Deposits into the credit union can be done at any time by any member of the VO, either physically or electronically. Physical deposits would be achieved by the member taking cash to the bank (although this is seen to be a short term expedient measure until electronic cash handling becomes established). Electronic deposits would be made by sending a digitally signed deposit slip and digital cash to the credit union account via the VO's trust gateway.

In order for this to work, a trust-based framework (reputation system) will need to be in place. This trust framework serves three main purposes,

- to act as a buffer between the bank (account) and the members of the credit union, as the deposit and withdrawal of funds takes place through this framework, meaning that the bank can deal with the VO as a whole, rather than with a collection of individuals;
- the trust framework records the credits made by the VO members, and uses this in its trustworthiness calculations; and finally
- the trust framework monitors the VO members' interactions via the web site, and the users apportion some value of trustworthiness to each of the members of the community according to how much they contribute to the well-being of the web community as a whole.

This reputation system allows members of the VO to input into the system information indicating their view on the performance of other members of the VO. This will take the form of a system whereby a user can rate the performance of another user through the portal, with this rating of performance becoming a part of the overall trust measurement for the user being rated. This measure of trust for each user will be combined with further trust information gathered from indicators of the amount of money being contributed to the central fund by each member, and the consistency of these contributions. The total sum of these trust measures is used to automatically indicate how much money a given member can request to withdraw from the union at any given time and also the weighting given to each users vote within the system.

The e-voting portion of the system will be based on common e-voting mechanisms and will allow members to vote on whether or not to allow other members to withdraw funds from the credit union. Voting for fund withdrawal by a member will be guided to some extent by the 'trustworthiness' (as described previously) of the member asking for a withdrawal. Note that it is important that "permission" to withdraw funds is agreed by the VO community membership as whole, via e-voting, since the VO members themselves are ultimately the persons who will suffer if the withdrawing member defaults on repayments.

3 Legal Analysis of VC Scenario with Respect to Data Protection Law

The following section focuses on the data protection issues involved when a reputation service operates with personal data.

3.1 Promises of Reputation Systems

Reputation systems collect information about a person or other entity (hereinafter “reputation subject”) in order to evaluate the reputation subject’s conduct and make this evaluation accessible for other users’ decisions. An example is when Internet marketplaces like eBay* and Amazon.com* enable users to provide feedback on other users. In this case, feedback ratings are based on a user's past transactions and help other users learn about the transaction partner they are dealing with. Other examples include credit reporting services, which collect information about an entity’s economic behaviour. This information is communicated e.g. to banks when they decide about credit. The latter kind of reputation systems has existed for a long time, but recent developments with respect to Internet based transactions have led to an increased need for reputation systems.

Reputation systems may be of particular value when there is uncertainty about another person or entity involved in a planned transaction that involves risk. Transactions on the Internet involve a number of uncertainties with regard to the identity of the transaction partner, his or her ability and willingness to perform, and the availability of realistic means of enforcement. The lack of experiences, knowledge or information about the other person or entity may lead us to refrain from the interaction. Reputation systems can provide us with relevant experiences others have had with this person or entity. Research indicates that reputation systems can encourage market actors to participate in transactions[4]. Reputation systems have also been considered as a compensation or supplement for lacking realistic means of enforcement on the Internet[5, 6, 7]. Thus, it is possible to think of new application scenarios for reputation systems, e.g. within virtual communities[8].

* Trademarks or registered trademarks of eBay Inc. and Amazon.com Inc.

[4] Keser, C., Experimental games for the design of reputation management systems, IBM Systems Journal, Vol. 42, No. 3, 2003, pp. 498–506.

[5] Friedman, D., Contracts in Cyberspace, available at http://www.daviddfriedman.com/Academic/contracts_in_%20cyberspace/contracts_in_cyberspace.htm, last visited 23 April 2004.

[6] Gilette, C.P., Reputation and Intermediaries in Electronic Commerce, Louisiana Law Review, Summer 2002, pp. 1165–1197.

[7] Block-Lieb, S., E-Reputation: Building Trust in Electronic Commerce, Louisiana Law Review, Summer 2002, pp. 1199–1219.

The possibilities offered by reputation systems are promising, but one should also pay attention to possible threats.

3.2 Objectives and Methodology

The objectives of this legal analysis are to investigate privacy and data protection problems related to reputation services. Introducing a reputation system requires a rather extensive collection, evaluation and disclosure of personal data. When deciding whether or not to participate in a reputation system, a potential user's concern may be whether the system will meet reasonable expectations with respect to privacy. Users may fear that too much information about them is collected and disseminated. There may also be concerns with regard to the "judging function" of a reputation system, where a user's conduct is evaluated. Such evaluations may be significant, since they are meant to be the basis for future decisions concerning him or her. This may raise questions with regard to how the user can dispute an evaluation he or she disagrees with. The lack of transparency and comprehensibility may increase these concerns. All these privacy-related concerns and fears may weaken the acceptance of a reputation system by potential participants.

Privacy concerns are not the only factors that can slow down or impede the uptake of a reputation system. From the perspective of the entity that uses the reputation profiles for decisions, the relevance and accuracy of the reputation data is essential. This decision-maker is interested in optimized data quality and has a separate interest in the quality of the process that generates reputation profiles. If the quality is not satisfactory from this perspective, this may weaken the value and utilization of the reputation system.

Data protection law provides rules that secure a fair processing of personal data. Furthermore, data protection law aims at enhancing data quality and contributes to increased transparency with respect to how data is processed. Therefore, data protection law can contribute to improve the value, acceptance and uptake of reputation systems.

The aim of this analysis is to provide guidelines for the design of reputation systems from a data protection perspective. It identifies legal and technical issues that should be addressed in order to design lawful and legitimate reputation systems. We will not analyse a specific reputation system, but rather explore different possibilities when developing a reputation system. Where appropriate, we mention the specific consequences for the VC scenario. Technical and organisational design choices may have legal consequences, particularly with respect to data protection law.

[8] Examples include https://shareyourexperiences.com/home.php, where apparently anybody can register his or her experiences with any other person. This system offers an identity protection service for those who provide information and those who seek for information. However, apparently the person of reference (reputation subject) is left rather unprotected. This example provides an indication of the possible threats to privacy of an extensive use of reputation systems. Similarly problematic, the Norwegian newspaper Dagsavisen has recently reported about a service where customers of prostitutes can share their experiences about individual prostitutes, cf. http://www.dagsavisen.no/innenriks/article1200650.ece.

3.3 Introduction to Data Protection Law

In Europe, data protection is subject to a rather strict legislation both on the European and national level. In this respect, reference will be made to the EC Directive on Data Protection (hereinafter EC Directive[9]) and its implementations in relevant national acts on privacy and data protection. A reputation service dealing with personal data is bound to follow the applicable national data protection law. In the VC scenario, we have selected Norwegian data protection law to be applicable. The Norwegian Personal Data Act is well-suited as a basis for an analysis, since it is an implementation of the EC Directive[10]. Norwegian Data protection law also includes some provisions of special relevance for reputation services (see section 3.4.6 below).

3.3.1 Who is Who in Data Protection Law

This section will introduce the central actors and terms used in this legal analysis to analyse reputation systems in the light of data protection law.

Personal data: This term is defined in the EC Directive, Article 2, as “any information relating to an identified or identifiable natural person”. “Any information” is a rather wide wording, which includes everything that can be perceived, sensed or registered etc. about a person. There are reasonable arguments to hold that also opinions, even false ones, must be qualified as personal data[11]. An “identifiable person” is one who can be identified, “directly or indirectly”. Some of the data processed by reputation systems can be personal data. However, the data will only fall into this category, if the data subject is a “natural person”.

Data subject: In data protection law, the data subject is the natural person (individual) to whom the personal data refers. However, reputation systems may also hold data that refers to other entities than individuals. We will therefore introduce the term “reputation subject”.

Reputation subject: We will use the term reputation subject when referring to the entity to which the reputation data relates. A reputation system can in principle administrate the reputation of individuals, groups, organisations, and collective entities. The latter ones may be juristic persons, but especially the case of virtual organizations shows that a collective entity can operate without legal personality. It is also possible to think of objects in reputation systems, but this legal analysis does not deal with objects in reputation systems, because information about objects (without any relation to a person) does not raise data protection concerns. In principle, reputation systems must only comply with data protection law when processing data on individuals, while most of data protection law is inapplicable

[9] Directive 95/46/EC, Official Journal L281, 23/11/1995 pp. 31–50. In this report, articles without reference to other instruments refer to this directive.

[10] Act of 14. April 2000 No. 31 relating to the processing of personal data (Personal Data Act), an English translation is available at http://www.personvern.uio.no/pvpn/regler/index.html.

[11] Bygrave, L., Data Protection Law, Approaching Its Rationale, Logic and Limits, Kluwer Law International, The Hague, London, New York, 2002, p. 46.

with respect to collective entities. Collective entities' protection is usually limited to laws dealing with defamation, breach of confidentiality, unfair competition etc. This is in contrast to data protection law, which e.g. ensures data quality, i.e. that data are relevant, correct, complete and not misleading in relation to the purposes for which they are processed. Arguably, collective entities and individuals share some interests, particularly with respect to the quality of data[12]. Therefore, reputation system providers may want to choose to follow central data protection rules also when processing data on reputation subjects other than natural persons. In the VC scenario, all reputation subjects seem to be natural persons. Nevertheless, it is possible to think of a reputation system used in a virtual community where also collective entities can be members.

Data controller: In the EC Directive, the data controller is defined as anybody who determines the purposes and means of the processing of personal data. When deciding who is a data controller in a reputation system, one has to identify the person or organisation with decision making power. If the system is developed by one entity but independently used by another, the latter is the data controller, since this entity determines the purposes and means of the processing. In principle, it is not impossible to think of more than one data controller. Data controllers are responsible for the lawful processing and may be held liable.

3.3.2 Basic Principles in Data Protection Law

The most important rules in data protection law can be expressed in relation to a number of basic principles[13] to be found in most international and national data protection instruments and laws.

- Fair and lawful processing: Personal data must be processed fairly and lawfully.
- Purpose specification: Personal data must be collected for specified, explicit and legitimate purposes and not further processed for other purposes.
- "Minimality": The collection and storage of personal data should be limited to the amount necessary to achieve the purpose(s).
- Information quality: Personal data should be valid with respect to what they are intended to describe and relevant and complete with respect to the specified purpose(s).
- Data subject participation and control: Persons should be able to participate in the processing of data on them and they should have some measure of influence over the processing.
- Limitation of fully automated decisions: Fully automated assessments of a person's character should not form the sole basis of a decision that impinges upon the person's interest.

[12] Bygrave, supra note 11, chapter 12.

[13] Bygrave, supra note 11, pp. 57–68 and 2.

- Disclosure limitation: The data controllers' disclosure of personal data to third parties shall be restricted, it may only occur upon certain conditions.
- Information security: The data controller must ensure that personal data is not subject to unauthorised access, alteration, destruction or disclosure.
- Sensitivity: Processing certain categories of especially sensitive data is subject to a stricter control than other personal data.

3.4 Data Protection Law and Reputation Systems

In this section, we will correlate these principles of data protection law with some of the possible characteristics of reputation systems. When designing a reputation system, one is confronted with a number of technical and organisational choices. These choices have an impact on how the reputation system processes personal data.

3.4.1 Participation in Reputation Systems

The principle of fair and lawful processing generally requires data controllers to take into consideration the interests and reasonable expectations of data subjects. This also implies that data subjects should not be unduly pressured into participation in reputation schemes. The principle of fair and lawful processing is embodied in a number of requirements in data protection law. The data subject's consent is the most important criterion for making processing of personal information in reputation systems lawful. The EC Directive defines the data subject's consent as "any freely given specific and informed indication of his wishes by which the data subject signifies his agreement to personal data relating to him being processed" (Articles 2 (h) and 7). The Directive requires timely and comprehensible information to be provided to the data subject, and the consent should be expressed through a voluntary and positive action. One should note that the processing of personal data also may be lawful without consent if other criteria are fulfilled. This may be the case, e.g. if the data controller's interest overrides the privacy interest of the data subject, or if the processing is necessary in relation to a contract or a legal obligation (Article 7(b), (c) and (f)).

When implementing a reputation system, one may consider making it mandatory to achieve maximum participation and value of the system. Most auction sites have a mandatory reputation system where participation in the reputation system is a condition for using the services. Also in the VC scenario, the reputation system seems to be mandatory for all participants. A discretionary/optional reputation system might be an alternative to be considered, even though this may cause some practical disadvantages. In the VC scenario, it is possible to think of members who want to take part in the virtual community, but not in the credit union. These should have the opportunity to take part in the community without being registered in the reputation system.

3.4.2 Centralised and Distributed Reputation Systems

The current reputation systems that have seen some form of deployment are centralised in the meaning that there is one centralised reputation service provider. For example, in Amazon.com, information is centrally administrated. The VC scenario does not state whether it is based on a centralized reputation system, but this seems to be the case.

This can be compared to distributed reputation systems, where every entity runs a local instance of the reputation system. Also hybrid systems have been suggested, combining elements characteristic of both centralised and distributed reputation systems[14]. One advantage with a distributed reputation system from a data protection perspective could be that information is spread between all participants, thus hindering accumulation of information in one place. Even in a fully distributed system, the system designer should ensure that relevant data protection principles are respected, including the right to access own personal data and the possibility to rectify false sets of data. In some cases it may be difficult to identify the data controller(s) in distributed systems.

Obreiter has suggested the use of so-called “evidences” in distributed reputation systems[15]. These non-repudiable tokens describe the behaviour of a specific entity in a statement. Digital signatures are used to make sure that the statement can be passed on to others. For example one party in a transaction can pass an evidence token to the other party, declaring the receipt of the item they trade. This receipt can later be used in order to document the behaviour, i.e. that the item has been sent and was received. In a data protection perspective, the use of such tokens has the advantage that they are not controlled by a central instance, but by the data subject himself. However, if the statements are too detailed and the data subject is expected to transfer many such tokens in order to document trustworthiness, this could lead to an excessive dissemination of personal information.

3.4.3 Identity and Identification

Reputation subjects may participate in a reputation system disclosing their real life identity to the other participants, or they may act under a pseudonym. From a data protection point of view, this choice is one of the most fundamental issues. One has to consider the necessary functionality of the reputation system and should be aware of technical, organisational and legal means to protect the identities and the personal data of the users.

“Personal data” is defined in the EC Directive, Article 2, as “any information relating to an identified or identifiable natural person”. An “identifiable person” is one who can be identified, “directly or indirectly” within a reasonable time, considering the necessary effort taking account of all the means likely reasonably to be used.

[14] Fernandes, A. Kotsovinos, E., Östring, S. Dragovic, B., Pinocchio: Incentives for Honest Participation in Distributed Trust Management, in Trust Management, Second International Conference, iTrust 2004, LNCS 2295, pp. 63–77.

[15] Obreiter, P., A case for Evidence-Aware Distributed Reputation Systems, Trust Management, supra note 14, pp. 33–47, p. 39.

Existing reputation systems often use pseudonyms to hide the identities of the users. We are not aware of fully anonymous reputation systems. In eBay for example, users register their contact information and are provided with a pseudonym which is used for transactions on the marketplace. Since the person behind the pseudonym can be identified, the pseudonym itself and data related to the pseudonym are personal data in relation to the EC Directive. It is possible to think of “strong” or “weak” pseudonyms in relation to how difficult it is to reveal the real-world identity for other users of the reputation system[16].The disclosure limitation principle provides that strong pseudonyms should be preferred to weaker ones.

The VC scenario does not state whether members are identified with their full identity or by a pseudonym. From a data protection perspective, pseudonyms should be chosen.

When issuing a pseudonym, the reputation system has different possibilities to verify the identity of the person. If a strict verification procedure is implemented, this strengthens the possibilities of holding the user of a pseudonym liable for misconduct. Pseudonyms that are linked to a verified identity may be trusted more easily. It is also possible to think of reputation systems where parties could participate under different pseudonyms depending on the need for assurance and reliability[17]. If a reputation system allows the use of multiple pseudonyms, these should not be linked to one common reputation profile[18].

Reputation systems should be limited to a specific marketplace or environment. A general reputation service that covers all kinds of actions in different contexts may lead to an excessive disclosure of personal information. Therefore, one should be careful with linking profiles from different reputation systems. In the VC scenario, the use of the reputation system is rather extensive, since apparently any involvement with the community is covered. The reputation profiles are used *inter alia*[19] to calculate the voting rights of the member and to determine whether a member can withdraw money from the fund. These purposes would have to be specified, and the VC would have to make sure that members express their consent that their data is processed for these purposes.

3.4.4 Types of Data in Reputation Profiles

A reputation system can generate a reputation profile by combining elements of evaluation (“excellent eBay buyer”) with more factual elements regarding e.g. the timeliness of the transaction, its value or category. In this context, fact and

[16] Clarke, R., Identified, Anonymous and Pseudonymous Transactions: The Spectrum of Choice, http://www.anu.edu.au/people/Roger.Clarcke/DV/UIPP99EA.html, last visited 29 April 2004.

[17] See for example the RAPID-project, Roadmap for Advanced Research in Privacy and Identity Management, http://www.ra-pid.org/, last visited 29 April 2004.

[18] See also Seigneur, J.-M and Jensen, C. D., Trading Privacy for Trust, Trust Management, supra note 14, pp. 93–107.

[19] The Latin term "Inter alia" means "amongst other things," *see e.g.* http://www.clickdocs.co.uk/glossary/inter-alia.htm.

evaluation are not seen as two dichotomist categories. This is rather a question of degree. For example, the comment "timely delivery" may include elements of both facts (delivery date) and evaluation (relation of the delivery date to rules about delivery, e.g. in a contract).

Some reputation systems, e.g. within credit rating, are based fully or mainly on factual information. Facts can either be made available to the end-user as separate information in order to provide a more comprehensive picture of the reputation subject, or they can be combined with the evaluation. The reputation system can collect factual information from a party's declaration, or simply track some of the information that is processed in relation to a transaction. Any collection from the data subject must be done in a fair and lawful way. This may require an informed consent, i.e. the participant must fully understand what is being tracked and for what purposes the data will be used. Ideally this should be explained both in a detailed way and in a way that is understandable for the average participant. This must be done prior to the collection of information. The reputation system in the VC scenario seems to be based on a rather extensive collection of factual information. Prospective VC members should be asked for their consent to this collection when they apply for membership.

The other element in reputation systems consists of evaluations, normally provided by other participants. In a data protection context, this is classified as the collection of personal data from third parties. The reputation system must ensure that the data subject is informed about the fact that personal data is collected from others, in addition to providing all the information mentioned in the previous paragraph. Evaluations may be thought of as rather uncomforting by the reputation subject, since this can be perceived as a judgment about him or her. Two data protection principles can assist the reputation subject in such situations: The principle of data quality and the principle of the data subject's participation and control. Both principles are reflected in Art. 12 (b) of the EC Directive, according to which the data subject has a right to have incomplete or inaccurate data rectified, erased or blocked. Obviously, evaluations made by third parties are difficult to verify for reputation systems. To cope with this problem, some reputation systems allow participants to cross-comment evaluations. Interestingly, research has shown that this function in eBay's feedback system leads to an under-reporting of negative comments because of the fear for negative cross-comments[20]. However, while minor problems are under-reported, participants do report instances of fraud, which indicates that the system seems to work best when it is most needed[21]. The VC scenario does not specify what members can do if they disagree with an evaluation. This question should be decided, both in order to avoid possible conflicts, and in order to comply with the duty to rectify inaccurate data.

[20] Gilette, supra note 6, p. 1191.

[21] Block-Lieb, S., supra note 7.

3.4.5 Generation of Reputation Profiles

Reputation profiles can be generated by aggregating factual elements and evaluations. This can result in some kind of score, e.g. a number of stars (Amazon.com) to be communicated to other users. According to the EC Directive, Article 12 (a), the data subject has a right to access the "knowledge of the logic involved in any automatic processing of data concerning him at least in the case of the automated decisions". Automatic decisions based on reputation profiles will be discussed below. However, we can already state that this may require that reputation systems have to inform the data subjects about the algorithm that is used to generate the reputation profile. Additionally, the algorithm has to comply with the principle of information quality in the sense that it generates information that is relevant, adequate and not excessive in relation to the purpose of the processing. The VC scenario does not specify how a reputation profile is aggregated or how a reputation score is calculated.

3.4.6 Access to and Disclosure of Reputation Data

Any data subject has a general access right to data on himself or herself (Art. 12). In the following, we will explore the limitations with respect to the disclosure of data to third parties.

Access to reputation data by third parties must be dealt with in light of the disclosure limitation principle. According to this principle, the data controller's disclosure of personal data to third parties shall be restricted, it may only occur upon certain conditions. The data subject may consent to such disclosure. Reputation data may be accessible all the time for any interested party, for example on a web site. Alternatively, access may only be given individually upon request.

When designing a reputation system, one should have in mind that some types of data are more sensitive than others. The EC Directive contains a catalogue of categories of especially sensitive data, which for example includes data concerning health or sex-life (Article 8.1.). This kind of data can only be processed under certain conditions. A reputation service that deals with data categories contained in this catalogue must restrict the disclosure to certain cases instead of allowing everybody to access the reputation data. Additionally, also other categories of personal data may have a strong impact or importance for a person, even though the category is not included in this catalogue. The sensitivity of personal data depends on its context. One example could be a person's credit history when applying for a credit. This data is of a major importance for the credit applicant, even though financial information is not included in the catalogue of sensitive data in the EC Directive.

Arguably, some of the data processed in the VC reputation system could be seen as sensitive. For example, we can not exclude that some health information will be processed, since the target group is recovering drug addicts, and health problems is probably among the most relevant issues for this group. If such data is processed, all the special rules about the processing of personal data must be observed. These rules can be found in Article 8 of the EC Directive and in Section 9 (1) of the applicable Norwegian implementation in the Norwegian Personal Data

Act[22]. According to Section 9 (2), there is a special exception for the processing of especially sensitive data by non-profit organisations: "Non-profit associations and foundations may process sensitive personal data in the course of their activities even if such processing does not satisfy one of the conditions laid down in the first paragraph, lit. a-h. Such processing may apply solely to data relating to members or to persons who, on account of the purposes of the association or foundation, voluntarily have regular contact with it, and solely to data which are collected through such contact. The personal data may not be disclosed without the consent of the data subject." Hence, although the VC reputation system processes data that falls into one of the special categories, one should be extremely careful with disclosing this data to all members of the community. In principle, a licence from the Data Inspectorate is required for the processing of sensitive personal data in Norway (Section 33 of the Norwegian Personal Data Act). The Data Inspectorate may set additional requirements when giving such a licence.

The financial information processed by the VC does not fall within one of the categories of especially sensitive data. As mentioned above, it may nevertheless be of central importance for the VC members. The importance of this type of data has led to special rules in some countries, even though financial information is not classified as sensitive. In this respect, reference should be made to the rules about credit reports under Norwegian[23] and Swedish law[24]. These rules regulate the disclosure of credit information by professional actors who specialise on trading such data. For example, disclosure of credit information is only allowed if the requestor has a legitimate interest in receiving the data. Additionally, every time the recorded information about an individual is disclosed, the credit information service has to contact this person (normally by letter). The data subject must be informed that data has been disclosed, who has requested it and what has been communicated. Here, also juristic persons are provided rights of access to information. This is one of the few examples where data protection law extends its scope to others than individuals.

The rules on credit reporting is the only set of rules that specifically regulates some reputation systems. However, it applies only to credit agencies. It is doubtful that the VC would be qualified as a credit agency according to section 4-2 of this regulation, since the information is only used within the community. Nevertheless, the main safeguards and procedures could be used analogously. Reputation systems like the VC should consider following some of these procedures in order to ensure the acceptance of their system.

[22] Act of 14. April 2000 No. 31 relating to the processing of personal data (Personopplysningsloven), English translation available at http://www.personvern.uio.no/pvpn/regler/index.html.

[23] Norwegian regulation on the processing of personal data, Forskrift om behandling av personopplysninger (personopplysningsforskriften) section 4.

[24] Sweden's Credit-Reporting Act, Kreditupplysningslag (SFS 1973:1173).

3.4.7 Decisions Based on Reputation Profiles

The major aim of reputation systems is to provide a basis for well-informed future decisions. In the cases of eBay and Amazon.com, the decision is whether or not to trust a certain pseudonym in the online market place. Decisions related to other reputation systems could include whether or not to participate in interaction with a subject in a virtual organisation, whether or not to allow a member of a virtual community access to a certain resource, whether or not to avail a credit to a person etc. There are basically two ways how these decisions can be made. Either the decision maker decides freely and uses the registered reputation as one of the premises for a decision. Alternatively, the decision can be made automatically on the basis of the calculated reputation score.

Automatic decisions are considered as problematic in a data protection perspective, and there are special rules for such decisions. Article 15 limits the use of certain automatic decisions based solely on automatic processing of data[25]. This applies only to decisions that are legally binding or which significantly affect the data subject. The data processed must be intended to evaluate certain personal aspects of the person who is targeted by the decision. These personal aspects include performance at work, creditworthiness, reliability, conduct etc, which all are aspects that could be evaluated in a reputation system. As mentioned above (p. 16), the data subject has a right to be informed about the logic involved in any automatic processing of data concerning him or her (Article 12 (a)). The data subject may also object to an automated decision and require a decision by a human (Article 15 (1), exceptions in (2)). Hence, when opening for automatic decisions based on reputation scores, one should be aware of these restrictions as implemented in national law.

In the VC scenario, the decision about the withdrawal of money from the common fund is made on the basis of an evaluation of a member's profile. However, the decision itself is not made automatically. Instead, VC members can decide about withdrawals through e-voting. Hence, this decision-making does not fall directly under Art. 15 of the EC directive.

3.5 Concluding Remarks with Respect to the Conventional Legal Analysis

Reputation systems should be carefully designed in order to comply with data protection law, if they (at least in part) deal with personal data. This will ensure a fair administration of information and users will more easily accept to participate in the reputation system. The above mentioned basic data protection principles can also be considered as a means to improve the data quality in a reputation system, which makes the reputation system more relevant as a basis for a decision and more attractive for the end-user. Below, we have tried to capture some relevant

[25] For more details about Art. 15 refer to Bygrave, L., Minding the Machine: Article 15 of the EC Data Protection Directive and Automated Profiling, Computer Law & Security Report, 2001, volume 17, pp. 17–24.

factors that should be considered to ensure that reputation systems respect data protection law.

- Participation in a reputation system should be limited to actors who have expressed their well-informed consent.
- The purpose(s) of the reputation system should be clearly defined.
- The collection, storage and dissemination of (personal) data should be limited to the amount necessary to achieve the purpose(s).
- The procedures regarding the collection and evaluation of personal data should be transparent and communicated in a comprehensible way.
- Reputation subjects should be allowed some participation and control with respect to the collection of data about them and with regard to the generation of their reputation profile.
- The quality of both the collected data and of the aggregated reputation profile should be valid with respect to what they are intended to describe and relevant and not incomplete with respect to the specified purpose(s).
- Fully automated decisions on the basis of reputation profiles should be avoided. If they are chosen, there should be full transparency regarding the algorithms used to calculate the reputation score and to make the decision. Additionally, the data subject should be able to claim a human decision.
- The security of (personal) data must be ensured.
- Reputation systems that deal with sensitive data should use a stricter policy to protect personal data.

These recommendations may assist in identifying legal problems, indicating that the reputation system developer and the data controller should seek legal advice to clarify how the law in the relevant jurisdiction solves these issues. The recommendations may also be used as a point of departure for future research on reputation systems with regard to data protection law.

4 Formal analysis of VC with Respect to Data Protection Law

This section is an investigation into the application of formal-logical tools to the formalization of some specific regulations relevant to the legal study of the Virtual Community (VC) scenario.[26] These include the EC Directive on Data Protection[27] and one of its implementations, the Norwegian Data Protection Act. It is natural to take the theory of normative-informational positions as one's starting point, because the building blocks of this theory include not only the usual deontic operators O_k and P_k (to be read as `it is obligatory for *k* that' and 'it is permitted for *k* that'), but also some information acquisition modal operator $I_j A$ with such reading as `agent *j* is informed that A'.[28] Such a notion can be related to the latter two regulations. There is no need to enter into the complexities of the formal theory.[29] Our goal here is just to try out the method in a rather simple case, by applying it to some of the basic notions involved in the latter two regulations, with a view to determining how far we can get.

This section can be viewed as an attempt to appreciate the extent to which formal conceptual analysis can contribute to the legal studies that will be conducted in TrustCoM. We will try to identify the kinds of nuances and distinctions that can be articulated in a logical framework such as the aforementioned one. We will also try to see how such nuances can arise in specific regulations like the EC Directive on data protection. This section is organized as follows. Subsection 7.1 focuses on the permission to disclose personal data. Subsection 7.2 introduces the obligation to use accurate data when disclosing personal data. Subsection 7.3 discusses ways of increasing the expressiveness of the framework so as to capture further aspects of relevance for the present analysis.

We choose here to proceed step by step. Simple structures convey very basic distinctions and insights, but might be gross oversimplifications. Complex structures may come closer to the contours of discourse, but can be extremely cumbersome to handle, with insights disappearing in a mass of details. It would not seem advisable to try to cover all complicating factors at once, but rather to get an initial appreciation of them a few at a time, only subsequently putting them together and investigating their interactions. Therefore, our policy in this study is to work with the simplest possible syntactic apparatus, reserving more complex machinery until the

[26] See supra section 5.

[27] See supra note 16.

[28] Agent *k* is the bearer of the obligation to send the information to *j.* The question of whether *k* is allowed to delegate such a task to another agent is left open. The person to whom the data refer (the data subject or, if you prefer, the reputation data – see section 5.2.3 below) is left implicit.

[29] For further detail on this logic, see Andrew J. I. Jones, "On Normative-Informational Positions", in A. Lomuscio & D. Nute, eds., *Proceedings of the Seventh International Workshop on Deontic Logic in Computer Science* (DEON04), LNCS/LNAI, Springer-Verlag, Berlin, Germany, May 2004, pp 182-190.

exact limits of the more Spartan one are clear – and only in so far as it is confirmed that its essential ideas are indeed "on the right track".

4.1 The Permission to Disclose Personal Data

We choose article 7 of the European Directive of 24 October 1995 as a running example, because the implementation of this article plays a central role in data protection laws in Europe. This article says:

> "Member states shall provide that personal data may be processed only if:
>
> - the data subject has unambiguously given his consent; or
> - processing is necessary for the performance of a contract to which the subject is party or in order to takes steps at the request of the data subject prior to entering into a contract; or
> - processing is necessary for compliance with a legal obligation to which the controller is subject; or
> - processing is necessary in order to protect the vital interests of the data subject; or
> - processing is necessary for the performance of a task carried out in the public interest or in the exercise of official authority vested in the controller or in a third party to whom the data are disclosed, except where such interests are overridden by the interests for fundamental rights and freedoms of the data subject which require protection under Article 1 (1)."

This article is implemented in Section 8 of the Norwegian Data Protection Act, which can be described with such semi-formal methods as Allen's Arrow diagrams.[30]

We start by simplifying the example in two ways. First, we substitute 'disclosed' for 'processed'. The processing of personal data is an umbrella category, which includes (but is not limited to) the disclosure of personal data. Second, we put aside the phrase "Member states shall provide that". This is a trivial alteration, since these words only mean that the Directive has to be implemented into national law by the member states.

The question is: how should the rest of Article 7 be analysed? This answer is complex; we will give it in layers.

First of all, it is worth mentioning that, from a legal perspective, the obligation not to disclose any personal data is in fact the general case. The argument is as follows. The right to freedom of expression, which is firmly embedded in, e.g., the European Convention on Human Rights and Fundamental Freedoms of 1950, seems prima

[30] This analysis is presented in appendix **Error! Reference source not found.** below. We believe that the use of such graphical tools can be misleading, if it is not supported by evidence from formalization. An illustration of this point is given in section 9.3 below.

facie to imply the right to disclose any personal data. But the right to privacy, which is also recognized as fundamental in the 1950 Convention, takes precedence over the right to freedom of expression. It is in this sense that the interdiction to disclose any personal data holds as a general rule. Article 7 lists a number of exceptions to this rule: the subject data has given his consent, etc. In this respect, the fact that Article 7 takes the form of a conditional statement is important. Formalizing the truth-functional (or propositional) structure of its antecedent is a straightforward matter; we omit the details here. We just stress that the locution ‘only if’ must be represented as a default-conditional $\Rightarrow$ as usually defined in non-monotonic logic. Indeed, Article 7 in turn allows for exceptions. The analysis of the atomic propositions appearing in the antecedent (“the subject has given his consent”, “the processing is necessary for the performance of a contract to which the subject is party”, etc.) goes beyond the resources of the present framework. For the purpose of the present discussion, these atomic propositions will be denoted by schematic letters p, q, r, etc, as usually done in elementary logic. The question of what the appropriate instances of these schematic letters would be is a topic for further investigation.

It might be instructive to focus first on the obligation not to disclose any personal data. For clarity’s sake, consider the situation where the data are collected and stored on a computer, as is the case in the VC scenario. Let *f* be a one-place predicate; *f*(*l*) is read “the file or the record of agent *l* is this and this”. Here *l* denotes the reputation subject, i.e. the person to whom the reputation data relates. For immediate purposes, let *A* abbreviate *f*(*l*). The obligation for *k* not to disclose *l*’s record to agent *j* can be formalized as follows, where ¬ and ∧ denote negation and conjunction respectively:

$$\mathrm{O}_k (\neg \mathrm{I}_j A \wedge \neg \mathrm{I}_j \neg A)$$

Here, $\neg \mathrm{I}_j A$ denotes the fact that *j* is not informed that *A*, and $\neg \mathrm{I}_j \neg A$ the fact that *j* is not informed that not-*A*. The conjunction $\neg \mathrm{I}_j A \wedge \neg \mathrm{I}_j \neg A$ is an instance of what might be termed the *silence* position.[31] Finally, the permission for *k* to inform *j* that *A* can be rendered as:

$$\mathrm{P}_k (\mathrm{I}_j A \wedge \neg \mathrm{I}_j \neg A)$$

The formalization we have offered for the obligation of silence is adequate in many contexts. But there are cases where the proposed translation is too strong. In fact, the meaning of the legal rule is: *k* is obliged not to communicate the information himself (or through others), but he is not obliged to prevent *j* from obtaining the information elsewhere. The sentence $\mathrm{O}_k (\neg \mathrm{I}_j A \wedge \neg \mathrm{I}_j \neg A)$ suggests the opposite, and gives us to understand that *k* must secure the result that the silence position is maintained, no matter *j*’s source of information is. This nuance can be made explicit in the symbolism, by introducing some agency operator E_k to be read as ‘agent *k* brings it about that’. Thus, the following formalization is too strong:

[31] The fact that the above obligation to be silent allows for exceptions can be captured in the usual way. It suffices to treat this obligation as the consequent of a default-conditional having an arbitrarily chosen tautology as its antecedent.

$$\mathrm{O}_k\ \mathrm{E}_k\,(\neg\, \mathrm{I}_j\, A \wedge\ \neg\, \mathrm{I}_j \neg A) \qquad (1)$$

Intuitively, formula (1) captures formally the fact that *k* must secure the result that the silence position is maintained, regardless of *j*'s source of information. As we have just seen, this does not match the meaning of the legal rule we are trying to formalize. The legal rule just says that *k* is obliged not to communicate the information himself. Therefore, (1) is "too strong", in the following sense: It implies that *k* does not fulfil his obligation, even if *j* gets the information elsewhere (e.g. from another agent *l*). From a legal point of view, this is obviously wrong.

An alternative, weaker rendering is:

$$\mathrm{O}_k\ \neg\, \mathrm{E}_k \neg\,(\neg\, \mathrm{I}_j\, A \wedge\ \neg\, \mathrm{I}_j \neg A) \qquad (2)$$

When we say that the alternative rendering (2) is weaker than (1), we use the concept 'weaker than' as usually defined in logic: (1) implies (2), but not the other way around, according to the logical principles used for the modal operators.

The following sentences are each equivalent to (2):

$$\mathrm{O}_k\ \neg\, \mathrm{E}_k\ (\mathrm{I}_j\, A \vee \mathrm{I}_j \neg A)$$

$$\neg\, \mathrm{P}_k\ \mathrm{E}_k\ (\mathrm{I}_j\, A \vee \mathrm{I}_j \neg A)$$

$$\neg\, \mathrm{P}_k\ \mathrm{E}_k \neg\,(\neg\, \mathrm{I}_j\, A \wedge\ \neg\, \mathrm{I}_j \neg A)$$

The significance of (2) is best viewed in a 'possible worlds' semantic. The sentence within the scope of the obligation operator, then, says that *k*'s own actions are always compatible with the truth of $\neg\, \mathrm{I}_j\, A \wedge\ \neg\, \mathrm{I}_j \neg A$.

The above issue does not arise in the case of permissions, with which we will be mainly concerned in the rest of this section. Therefore, we can put aside the agency operator, which would complicate the presentation unduly.

4.2 The Obligation to Use Accurate Data When Disclosing Personal Data

In this section, we go one step further and investigate if, and to which extent, some other parts of the EC Directive should affect our initial understanding of (the consequent in) Article 7.

As just observed, there are circumstances in which the data controller is allowed to disclose personal data. Then we may ask whether there are any rules about the quality of the data that can be disclosed. According to EC Directive, Article 6 (1) lit. d personal data must be accurate. The term accuracy is not defined in the Directive, but "inaccurate" is defined in some national legislations as "incorrect or misleading as to any matter of fact".[32] The question may arise how this rule can be formalized. When we speak about the accuracy of data, it is important whether or not the content of the communicated information is true or not. Suppose it is *A* rather than ¬ *A* that, according to the evidence available to the data controller, is

[32] Section 70 (2) of the British Data Protection Act 1998, see also Carey, Peter: Data Protection. A practical Guide to UK and EU law, second edition, Oxford 2004.

true. The conjunction $I_j A \wedge \neg I_j \neg A$ is, then, an instance of what might be termed the *straight truth* position. The obligation to use accurate data when disclosing personal data can be rendered as follows:

$$O_k (I_j A \wedge \neg I_j \neg A).$$

One remark is to be made here. Besides the *straight truth* position and the *silence* position alluded to above, there are just two other scenarios that are logically possible: the *straight lie* position and the *conflicting information* position. The first one corresponds to the configuration $I_j \neg A \wedge \neg I_j A$, and the second one to the configuration $I_j A \wedge I_j \neg A$. It can readily be seen that these four conjunctions are mutually exclusive, and their disjunction is a tautology. Precisely one of these conjunctions must be true. Applied to legal contexts, the notion of the *conflicting information* position makes sense: the data controller can get conflicting information from different sources, and he or she may not be in a position to determine which information is true.[33] By contrast, the notion of the *straight lie* position seems strange, when applied to legal contexts. But the framework attempts to be perfectly general. It is not tied to any specific application area. The proposed tools might well be more general and more expressive than needed for the specific purposes of the legal analysis of a given scenario.

Keeping these remarks in mind, let us return to the example. We have identified $P_k (I_j A \wedge \neg I_j \neg A)$ as the appropriate 'normative-informational' position if the disclosure of personal information is allowed according to Article 7 of the EC Directive. But a central conjecture in the theory we are trying out here is that this might not be the only, not perhaps even an adequate, representation of what Art. 7 in combination with Art. 6 mean. As shown below, $P_k (I_j A \wedge \neg I_j \neg A)$ covers many cases. It is the systematic exploration of all such possible cases that motivates the theory of normative informational positions. There is, first, the problem of identifying the set of possible interpretations. This can be done by using formal methods. There is, then, the question of picking out one of these interpretations as the most likely one. This second question can be answered only by using further information (e.g., other parts of the regulation).

The method used for generating the set of possibilities that need to be considered involves many steps. The details are given in Andrew J. I. Jones, supra note 29, and will not be reproduced here. For simplicity's sake, we here opt for an alternative presentation, in terms of *maxi-conjunctions*.[34] This will help the reader appreciate better the basic idea of the construction, and convince him that the method used for

[33] Outside the legal domain, one may easily find other examples involving conflicting information from different sources. One can also find examples where the conflicting information comes from one and the same source. For instance, if you want some rival nation to be unsure about your own military plans, then a good strategy is to give this one contradictory information about them. Some other examples illustrating the usefulness of the notion of conflicting informational position can be found in Andrew J. I. Jones, supra note 29.

[34] A systematic study of this notion can be found in M. Sergot, "A computational theory of normative positions", *ACM Transactions on Computational Logic*, Vol. 2, N°4, 2001, p. 595-622. The author himself builds upon some observations made by D. Makinson in "On the formal representation of rights relation", *Journal of Philosophical Logic*, 1986, 15, 403-425.

generating the spectrum of possibilities is not as simple as the one used in propositional logic when generating a truth-table.

By a maxi-conjunction, we mean a maximal consistent conjunction. Consistent refers to some underlying logic, here the specified logic for O_k and I_j. “Maximal” means that addition of any new conjunctions (taken from a given set) yields an inconsistency. Let ± stand for the two possibilities of affirmation and negation. The basic idea can be summarized by saying that there are two main steps in the procedure. The first one involves computing the set of maxi-conjunctions of sentences of the form $\pm O_k \pm A$, where each A is itself a maxi-conjunction of sentences of the form $\pm I_j \pm A$.[35] It turns out that only 15 distinct conjuncts can be generated. The details turn out to be fiddly to state concisely. For simplicity’s sake, we will not give the full list, confining ourselves with enumerating below those that are relevant to the purposes of the present analysis. The second and final step consists in simplifying these 15 conjunctions, to remove redundant conjuncts (i.e., conjuncts that are themselves logically implied by some other conjunct in the same maxi-conjunction).

It is to be noted that, by construction, the maxi-conjunctions thus obtained are not only internally consistent, but also mutually exclusive and their disjunction is a tautology. In any given situation, precisely one of the conjunctions must be true. It is in this respect that the method can be said to provide a systematic exploration of all the logically possible situations.

We said earlier that $P_k (I_j A \wedge \neg I_j \neg A)$ covers many cases. By this, we mean that, among these 15 (non-equivalent) situations that are logically possible, 7 of them implies the truth of $P_k (I_j A \wedge \neg I_j \neg A)$. These are:

(N1) $O_k (I_jA \wedge \neg I_j\neg A)$

It is obligatory for *k* that *j* is told the *straight truth*.

(N4) $O_k (\neg I_jA \vee \neg I_j\neg A) \wedge P_k (\neg I_jA \wedge \neg I_j\neg A) \wedge P_k (\neg I_jA \wedge I_j\neg A) \wedge P_k (I_jA \wedge \neg I_j\neg A)$

The *conflicting information position* is forbidden for *k*, but the *silence position*, the *straight lie position* and the *straight truth position* are each permitted for *k*.

(N5) $O_k (I_jA \leftrightarrow \neg I_j\neg A) \wedge P_k (I_jA \wedge \neg I_j\neg A) \wedge P_k (I_j\neg A \wedge \neg I_jA)$

The *conflicting information position* and the *silence position* are both

[35] We draw here from an observation made by M. Sergot in “A computational theory of normative positions”, *ACM Transactions on Computational Logic*, Vol. 2, N°4, 2001, p. 595. The set of maxi-conjunction of sentences having the form $\pm I_j \pm A$ corresponds to the set of four informational positions alluded to above.

forbidden for *k*, but the *straight truth* and *straight lie positions* are both permitted for *k*.

(N6) $\neg P_k I_j\neg A \wedge P_k (\neg I_j\neg A \wedge I_jA) \wedge P_k (\neg I_j\neg A \wedge \neg I_jA)$

It is not permitted for *k* that *j* is told a lie, but the *straight truth* and *silence positions* are both permitted for *k*.

(N9) $O_k I_jA \wedge P_k (I_jA \wedge I_j\neg A) \wedge P_k (I_jA \wedge \neg I_j\neg A)$

It is obligatory for *k* that *j* is told the truth; the *straight truth position* is permitted for *k*, but so is the *conflicting information position*.

(N12) $\neg P_k (I_j\neg A \wedge \neg I_jA) \wedge P_k (I_jA \wedge \neg I_j\neg A) \wedge P_k (I_jA \wedge I_j\neg A) \wedge P_k (\neg I_jA \wedge \neg I_j\neg A)$

The *straight lie position* is forbidden for *k*, but the *straight truth position*, the *silence position* and the *conflicting information position* are each permitted for *k*.

(N15) $P_k (I_jA \wedge I_j\neg A) \wedge P_k (I_jA \wedge \neg I_j\neg A) \wedge P_k (I_j\neg A \wedge \neg I_jA) \wedge P_k (\neg I_jA \wedge \neg I_j\neg A)$

The *conflicting information position*, the *straight truth position*, the *straight lie position* and the *silence position* are each permitted for *k*.

So, we may ask, which of these 7 properly represents `it is permitted for *k* to inform *j* about *A*'? In most cases, conflicting information would not fulfil the requirement of accuracy.[36] Therefore, we suggest eliminating the last three conjunctions, (N9), (N12) and (N15), which each imply that the *conflicting information position* (vis-à-vis *j*) is permitted for *k*. For the same reason, we suggest eliminating (N4) and (N5), each of which implying that the *straight lie position* is permitted for *k*. Then there remain the two positions (N1) and (N6). The peculiar thing about (N1) is that it places agent *k* under an obligation to inform *j* about *A*. Such an obligation makes sense if (as we will see in a moment) agent *j* is the data subject, and has requested to be informed about *A*. However, with respect to the communication of personal data to other persons, there is no such obligation. Therefore, (N6) is the appropriate choice in most cases.[37]

[36] As already observed, it is possible to think of situations where conflicting information may be accurate: If the data controller has collected information from different parties who do not agree about factual information and the data controller has no possibility of verifying the information, then it may be accurate to communicate the conflicting information.

[37] Although we need to subject this point to further investigation, we believe that the context of contract formation also provide examples where (N6) is the appropriate choice. This issue is related to the principle of good faith, of which the obligation to provide information is a special case. An interesting discussion of the problems raised by these notions can be found in E. M. Weitzenbock, 'Good Faith and Fair Dealing in the

4.3 More Expressivity

We now indicate how to increase the expressiveness of the framework so as to capture further aspects of the EC Directive. We present below some structures that are relatively more complicated, and show how to analyse them within the existing framework. One of these more complex structures is mentioned in the Conceptual Framework for the Legal Risk Analysis.

7.3.1 Notifying

One interesting feature of the logic is that it allows us to iterate information acquisition operators. The logical machinery allows such an embedding, at least as a formal possibility. Does the EC directive on data protection law provide examples of such iterations? Yes. Consider the obligation for *k* to notify *i* that his personal data have been disclosed to *j* [this is Article 11]. The informational position within the scope of the deontic operator has the following more complex form:

$$\mathrm{I}_i\,(\,\mathrm{I}_j A \wedge \neg\,\mathrm{I}_j \neg A) \wedge \neg\,\mathrm{I}_i \neg\,(\,\mathrm{I}_j A \wedge \neg\,\mathrm{I}_j \neg A)$$

This sentence is similar in pattern to the *straight truth* position with which we have been working so far. The former can be obtained by replacing, in the latter, *j* with *i* and *A* with $\mathrm{I}_j A \wedge\neg\,\mathrm{I}_j \neg A$. Using a self-explanatory terminology, this more complex form of informational position can be said to be an informational position of second-degree. The methodological procedures outlined above are meant for the first-degree (or flat) case, but it should be possible to carry them over to the second-degree case. Then one would be able to compare the pattern exhibited by the maxi-conjunctions of the second-degree with those exhibited by the maxi-conjunctions of the flat type. This is a topic for further investigation.

Of course, the analysis of the obligation to notify that we have just outlined might well need further refinement. In particular, it is natural to ask if such a duty implies the following norm: it is obligatory for the data controller that (by putting in place some appropriate mechanisms) he *makes it possible* for the data subject to have access to the information given in the notification. Indeed, permitting the data subject to have access to the information in question seems to be one thing, making it possible for him to realize this permission seems to be another thing (this point has been emphasized by Kanger, when discussing what he calls the ‘realization’ of rights).[38] Does such a nuance effectively arise in the EC Directive? This remains to be appreciated better. For present purposes, suffice it to observe that the above distinction can be made explicit in the symbolism, if we allow ourselves the use of an operator Can (for “it is practically possible that”).[39] The following rendering suggests itself:

Context of Contract Formation by Electronic Agents’, in *Proceedings of the AISB 2002 Symposium on Intelligent Agents in Virtual Markets*, 2-5 April 2002, Imperial College London.

[38] Kanger, S., ‘On realization of human rights’, in G. Hölmstrom and A.J.I. Jones (eds), *Action, Logic and Social Theory*, Acta Philosophica Fennica, vol. 38.

[39] A treatment of this modality can be found in, e.g., A.J.I. Jones, “A Logical Framework”, in Pitt J. (ed.), *The Open Agent Society*, John Wiley & Sons, UK (forthcoming in 2004).The distinction between

$$O_k \, \mathrm{Can} \, (I_i (I_j A \wedge \neg I_j \neg A) \wedge \neg I_i \neg (I_j A \wedge \neg I_j \neg A))$$

This would be the simplest way to express the idea that *k* should make it possible for *i* to have access to the information given in the notification.

7.3.2 Quantifiers

There are at least two good reasons to introduce quantifiers.[40] First of all, it is in general necessary to specify that every individual in a certain category occupies such-and-such normative positions. Second, processing certain categories of especially sensitive data is subject to a stricter control than other personal data.[41]

7.3.3 Conditions

As we have seen, norms are conditional by their very nature. Deontic conditionals raise complex philosophical and technical issues, which any theory of norms (of any kind) must address. As initially conceived, the theory of normative-informational positions puts aside such issues, since as such they are not essential to the analysis of the procedures used for determining the logical space of normative positions. Formally adequate treatments of conditionals are nevertheless available in the literature, and can be inserted more or less straightforwardly into the existing framework.

There is, then, the further question of formalizing the atomic propositions in the antecedent of Article 7: the subject has given his consent; the processing is necessary for the performance of a contract to which the subject is party, etc. This is a topic for further investigation.

7.3.4 Right to be informed

It is natural to ask whether `permission to be informed' is the same as `right to be informed'. The answer is clearly 'no': rights are essentially relational – a principal insight in Hohfeld's work, from which the theory originates. This means that we need to supplement the picture with a representation of the normative informational positions of other agents. A matter of some complexity, but there are also grounds to believe that this complexity may be systematically analysed along the lines just suggested.

permission and practical possibility is obvious. An agent may be permitted to bring about such-and-such state of affairs (e.g. going abroad on holiday) without being able to do it. Such a distinction is missing in both the Allen account and the EPAL language.

[40] Those symbols enable us to speak about 'all' or 'some' things that fall into a given category.

[41] See Article 8 of the EC Data Protection Directive.

4.4 Concluding Remarks with Respect to the Formal Analysis

This section reported a first attempt to formalize some of the specific types of rights in the EC Directive on data protection. Such rights are relevant to the legal analysis of the VC scenario. The risk analysis methodology makes use of UML models, and is not expressive enough as regard the normative modalities. This is why we used here tools from modal logic. A natural next step would be to investigate how the UML (and, by the same token, the risk analysis methodology) can be adjusted to support some of the expressiveness provided by the present framework.

5 Legal Risk Analysis of VC with Respect to Data Protection Law

This section presents the results of a risk analysis of legal issues in the VC scenario and the methods used. The objective of the risk analysis is to assess the applicability of risk analysis methods for analysis of legal issues, and to come up with potential domain specific extensions or specialisations to facilitate legal risk analysis. To perform this assessment we utilised the CORAS risk analysis methods and tools as a basis.

Section 5.1 gives a quick overview of the CORAS risk analysis process, while Sections 5.2 through 5.5 presents the main analysis results. Finally, Section 5.6 presents some concluding remarks about the applicability of the methods and potential changes.

5.1 The CORAS Framework for Model-based Risk Assessment (MBRA)

The legal risk analysis documented in this report is based on the CORAS-framework developed in the CORAS-project[42]. The CORAS risk management process, based on the Australian Standard AS/NZS 4360 [43], provides a sequencing of the risk management process into the following five sub-processes:

- *Context Identification*: Identify the context of the analysis. Describe the target of evaluation and its environment, and identify usage scenarios, assets and risk evaluation criteria.
- *Risk Identification*: Identify the potential threats to assets and the vulnerabilities of the assets. Identify unwanted incidents.
- *Risk Analysis*: Evaluate consequence and frequency of each unwanted incident identified in sub-process 2.
- *Risk Evaluation*: Identify the level of risk associated with the unwanted incidents already identified and assessed in the previous sub-processes. Evaluate the identified risks with respect to the risk evaluation criteria, prioritise the risks, and categorise risks into risk themes.
- *Risk Treatment*: Address the treatment of the identified risks, and how to prevent the unacceptable risks.

[42] Vraalsen, F., den Braber, F., Hogganvik, I., Lund, M. S., Stølen, K., The CORAS tool-supported methodology for UML-based security analysis, SINTEF Technical report STF90 A04015, SINTEF ICT, 2004

[43] Australian Standard (1999): Risk Management. AS/NZS 4360:1999

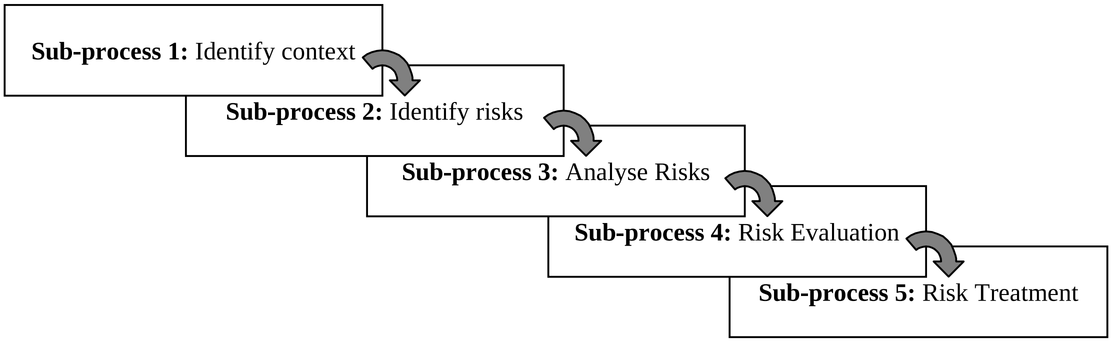
Sub-process 1: Identify context
Sub-process 2: Identify risks
Sub-process 3: Analyse Risks
Sub-process 4: Risk Evaluation
Sub-process 5: Risk Treatment

Figure 1: Overview of the CORAS risk management process

In the analysis of the VC scenario it was decided to focus most on risk identification (sub-process 2). In particular, we wished to assess the use of HazOp as a method for identification of legal risks. HazOp is one of the main methods used for identification of threats and unwanted incidents in CORAS, and is further described in section 5.3.1.

5.2 Context Identification

This sub-process consists of a number of activities. Some deal with "meta-information" about the risk analysis process itself, e.g. the participants involved and the schedule for the risk analysis. These activities have been left out in this case study.

Section 5.2.1 describes the target of evaluation. Section 5.2.2 presents the stakeholders, while section 5.2.3 presents the assets that were identified.

5.2.1 Target of Evaluation

The target of evaluation (ToE) is the subject of the risk analysis, in this case the Virtual Community credit union described in Section 2. The goal of this activity is to determine the scope of the analysis. To describe the target of evaluation, a number of different types of UML diagrams may be used, e.g.:

- use cases – high level description of usage scenarios and actors involved
- sequence and activity diagrams – more detailed description of use cases, showing processes and interactions between various actors
- class diagrams – a more static view of actors and other entities, their properties and relationships

An overview of the functionality of the VC credit union website is shown in the use case diagram in Figure 2. The usage scenarios are shown as ovals. A usage scenario may include other scenarios, as shown by the dashed arrows. The actors are shown as stick figures, with solid lines connecting them to the usage scenarios they are involved in.

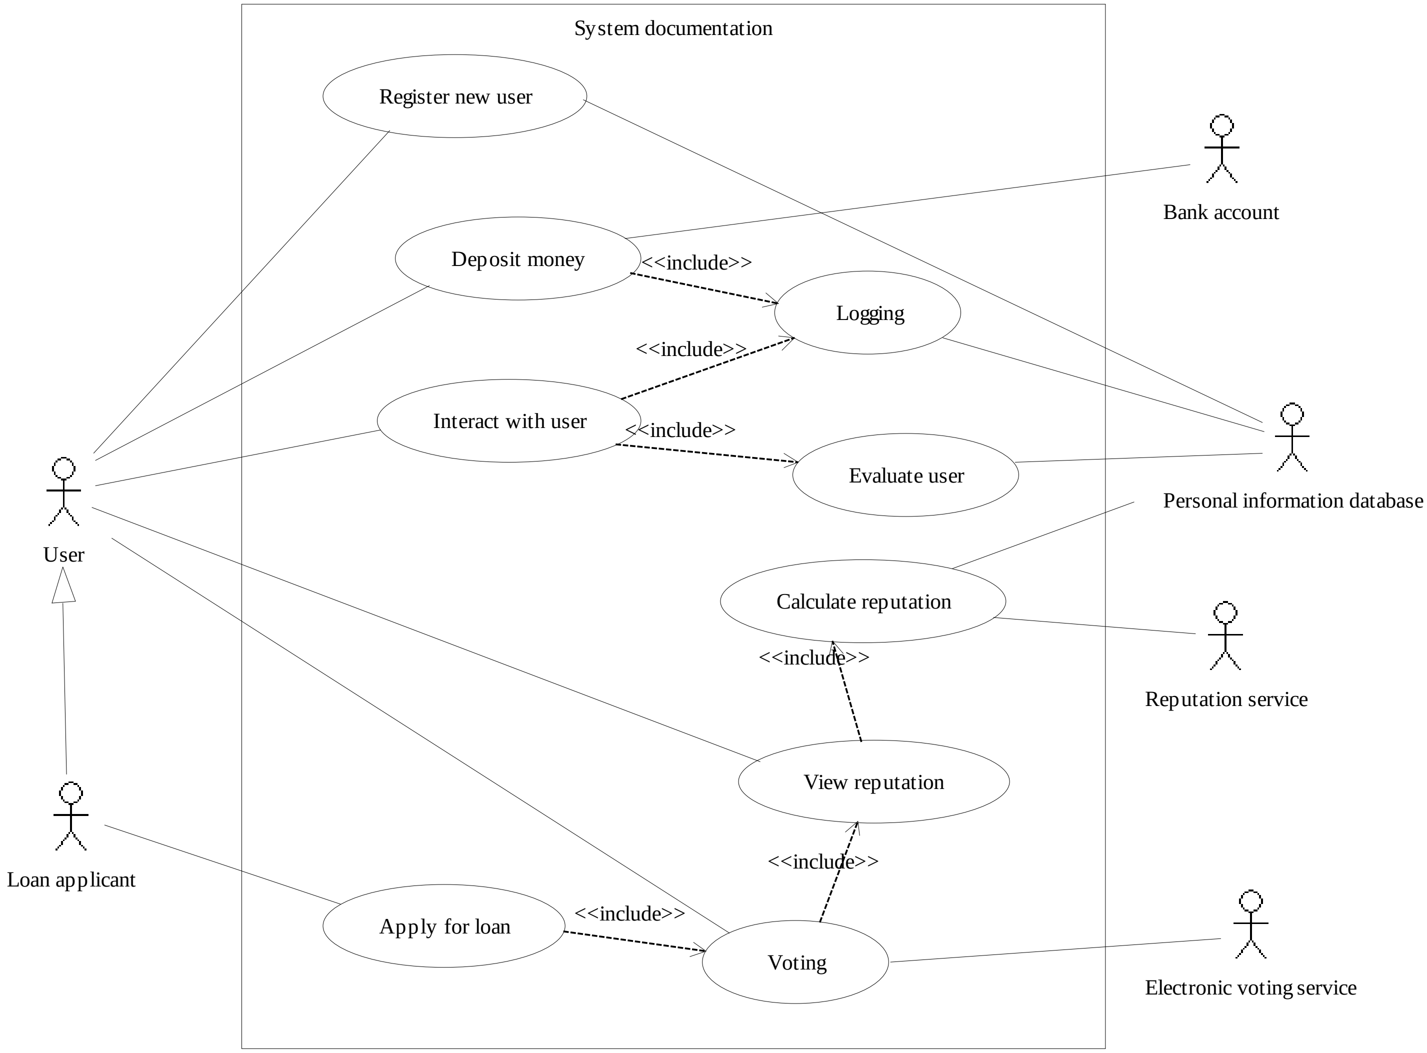

Figure 2: System description

5.2.2 Stakeholders

A stakeholder is defined as *a person or organisation that has interest in the target of evaluation*. The analysis was performed from the viewpoint of the VC credit union, which was defined as the sole stakeholder in this analysis. The information about stakeholders is recorded in a Stakeholder table, as shown in Table 1.

Stakeholder ID	**Stakeholder Role**	**Stakeholder Name**	**Description**
Credit union	System owner	VC credit union	The credit union owns and runs the VC systems and services

Table 1: Stakeholder table

5.2.3 Assets

An asset is defined as *a part or feature of the target of evaluation that has value for one of the stakeholders*. After describing the target of evaluation and identifying the stakeholder, the assets were identified and assigned values. As the analysis is based on a fictitious case, the risk analysts played the role of the stakeholder in this activity.

The Asset table is used to record information about assets, such as which stakeholder they are related to and the value assigned to the asset by that stakeholder. The asset table for the VC credit union is shown in Table 2. The assets are grouped into themes, such as information (all information in the system or that the system depends on).

Asset ID	**Stakeholder ID**	**Asset Theme**	**Asset**	**Description**	**Value**
Reputation	Credit union	Organisational	VC reputation	The reputation/brand name of the VC	Very high
VC fund	Credit union	Organisational	VC fund	The credit union fund	High
Charity fund	Credit union	Organisational	Charity fund	The funds of the charity organisation	High
Reputation service	Credit union	Software	User reputation service	The software which processes the user evaluations (collecting, distributing, calculating user reputation score)	High
Voting service	Credit union	Software	Electronic voting service	The software which processes the user votes in relation to loan applications	High
Legal record	Credit union	Law and regulation	Clean legal record	VC behaviour is in accordance with legal obligations	High
Evaluations	Credit union	Information	User evaluations	Evaluations of users' reputations by other users	High
Database	Credit union	Information	Personal information database	Database containing personal information about the users of the VC	High
Payment history	Credit union	Information	User payment history	Information about a user's payments into the credit union fund	Medium
Interaction history	Credit union	Information	Interaction history	Information about user's interactions with each other	Low

Table 2: Asset table

Asset diagrams are a special kind of UML class diagrams which can be used to show a graphical view of the assets. Asset diagrams are part of the CORAS UML profile for security risk analysis[44]. The asset diagram for the VC assets listed in Table 2 is shown in Figure 3. Assets can be subtypes of other assets, shown using arrows, similar to the subtyping of classes or concepts.

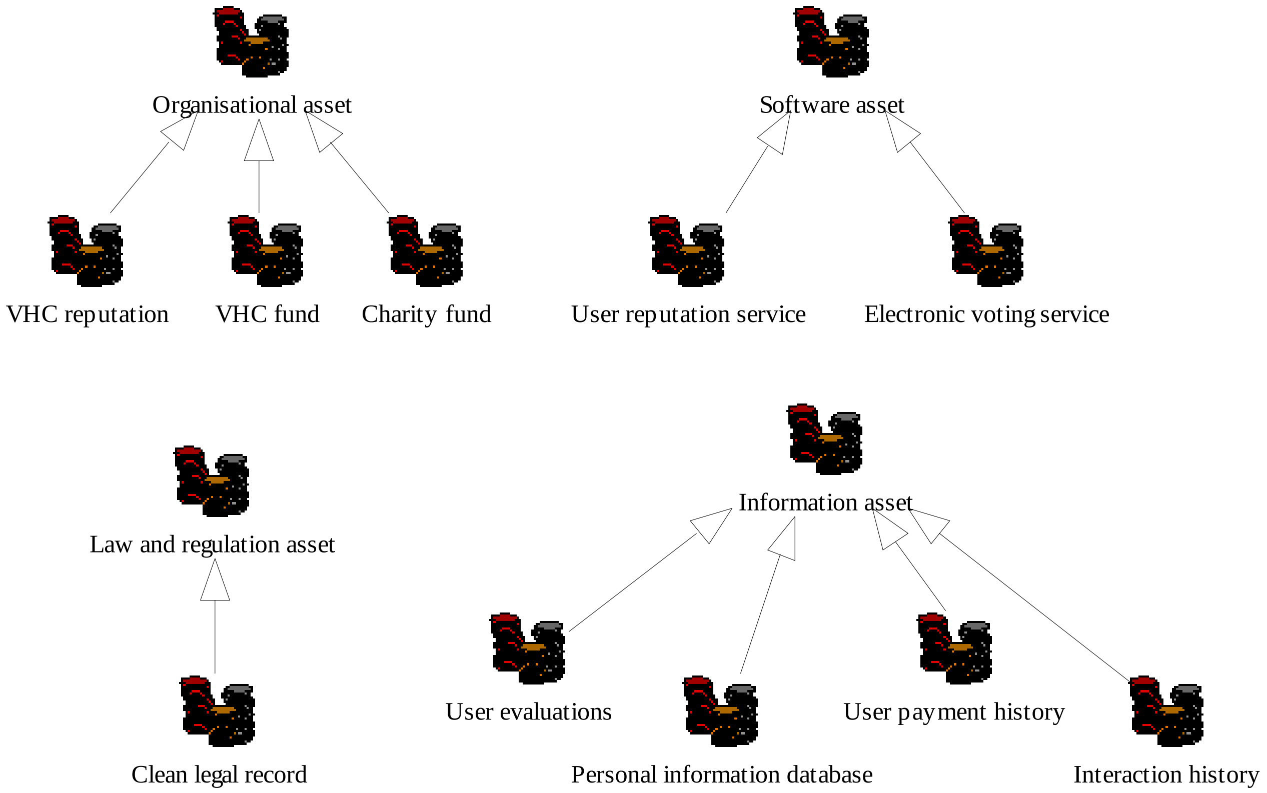

Figure 3: Asset diagram

5.3 Risk Identification

This section presents the risks identified during the analysis in the form of threat and unwanted incident diagrams specified using the CORAS UML profile. The method used for threat identification, HazOp, is introduced in section 5.3.1. In section 5.3.2, threat and unwanted incident diagrams are explained, and the results of the risk identification sub-process are presented in section 5.3.3.

[44] Lund, M. S., Hogganvik, I., Seehusen, F., Stølen, K., UML profile for security assessment. Technical report STF40 A03066, SINTEF, 2003

5.3.1 HazOp analysis

Hazard and Operability (HazOp) analysis[45] is a risk analysis method which may be used to identify threats and unwanted incidents. A HazOp analysis can be described as a “structured brainstorming”. The idea is to focus on the interesting parts of the target of evaluation, processing each individually and trying to identify threats connected to failure or incompleteness of these items. Usage scenarios may be described in use case diagrams and sequence or activity diagrams.

During a HazOp session, the knowledge and expertise of all the participants in the risk analysis session is exploited, in order to find as many relevant threats as possible. Guidewords and questionnaires may be used to guide the process and assist the participants in finding threats to the scenarios described in the diagrams. The session is led by the risk analysis leader while the risk analysis secretary is responsible for writing down the results of the analysis itself. The results are stored in a HazOp table. Table 3 shows an empty HazOp table.

Table 3 HazOp table template

Risk ID	*Asset ID*	*Item*	*Guideword/ Attribute*	*Security Scenario*	*Unwanted Incident*	*Consequence/Frequency*	*Treatment*

5.3.2 Threat and unwanted incident diagrams

When using the CORAS UML profile, the threats identified during the HazOp analysis are modelled in so called threat and unwanted incident diagrams. The threats are modelled using threat agents and threat scenarios, where the threat agents represent the active part of the threat and the threat scenarios model the behaviour of the threat agents. Each threat agent is related to one or more threat scenarios, which are again related to the assets they threat. However, even though a threat exists, this does not necessarily mean that something bad will happen, it merely represents the potential of this occurring.

The actual threat may also be the result of interplay between a number of threats. We model this by the use of unwanted incidents. Furthermore, an unwanted incident may lead to another unwanted incident, forming a chain of events. The unwanted incidents are also related to the assets they affect.

Threat and unwanted incident diagrams thus include four types of icons:

- **Threat agent:** Represent the active part of the threat, the initiator of the threat scenario. The profile has a number of predefined threat agents with

[45] Redmill, F., Chudleigh, M., & Catmur, J., Hazop and software Hazop. Wiley, 1999.

various graphical representations, depending on the type of threat agent, e.g. a stick figure is used to represent a human threat.

- **Threat scenario:** A description of the behaviour of a threat agent, i.e. how it may lead to an unwanted incident. The graphical representation is an oval with an ignited bomb.
- **Unwanted incident:** An undesired event that may reduce the value of an asset. The graphical notation is an oval with a warning sign.
- **Asset**: A part or feature of the target if evaluation that has value for one of the stakeholders. The graphical notation is a stack of coins.

The relationship between a threat scenario and an unwanted incident, expressed as a dashed arrow with the label <<include>> pointing from the unwanted incident towards the threat scenario, represents the fact that the threat scenario may be included in the description of how the unwanted incident may occur. Furthermore, the chain of events where an unwanted incident *P* may cause another unwanted incident *Q* is represented using a dashed arrow labelled <<initiate>> going from *P* to *Q.*

Figure 4 shows an example threat and unwanted incident diagram containing all the elements described above.

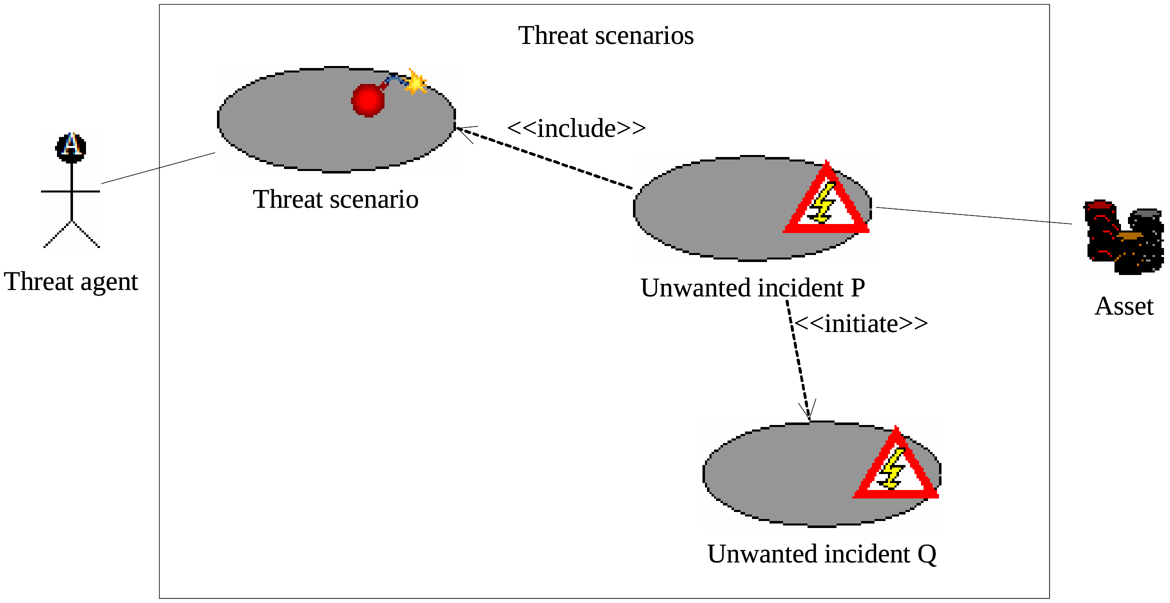

Figure 4: Example threat and unwanted incident diagram

5.3.3 Risk identification results

The threats and unwanted incidents identified during the HazOp analysis are presented in the threat and unwanted incident diagrams in Figure 5, Figure 6 and Figure 7. Figure 5 shows the threats of registering personal information without grounds and processing sensitive information without a license. Both of these threat scenarios are regarded as violations of the provisions of the Norwegian Data

Protection Law, and may therefore lead to a number of unwanted incidents, as shown in the diagram. A number of additional special threat scenarios which may lead directly to fines or imprisonment are shown in Figure 6. Finally, Figure 7 focuses on the notion of consent and threat scenarios related to this.

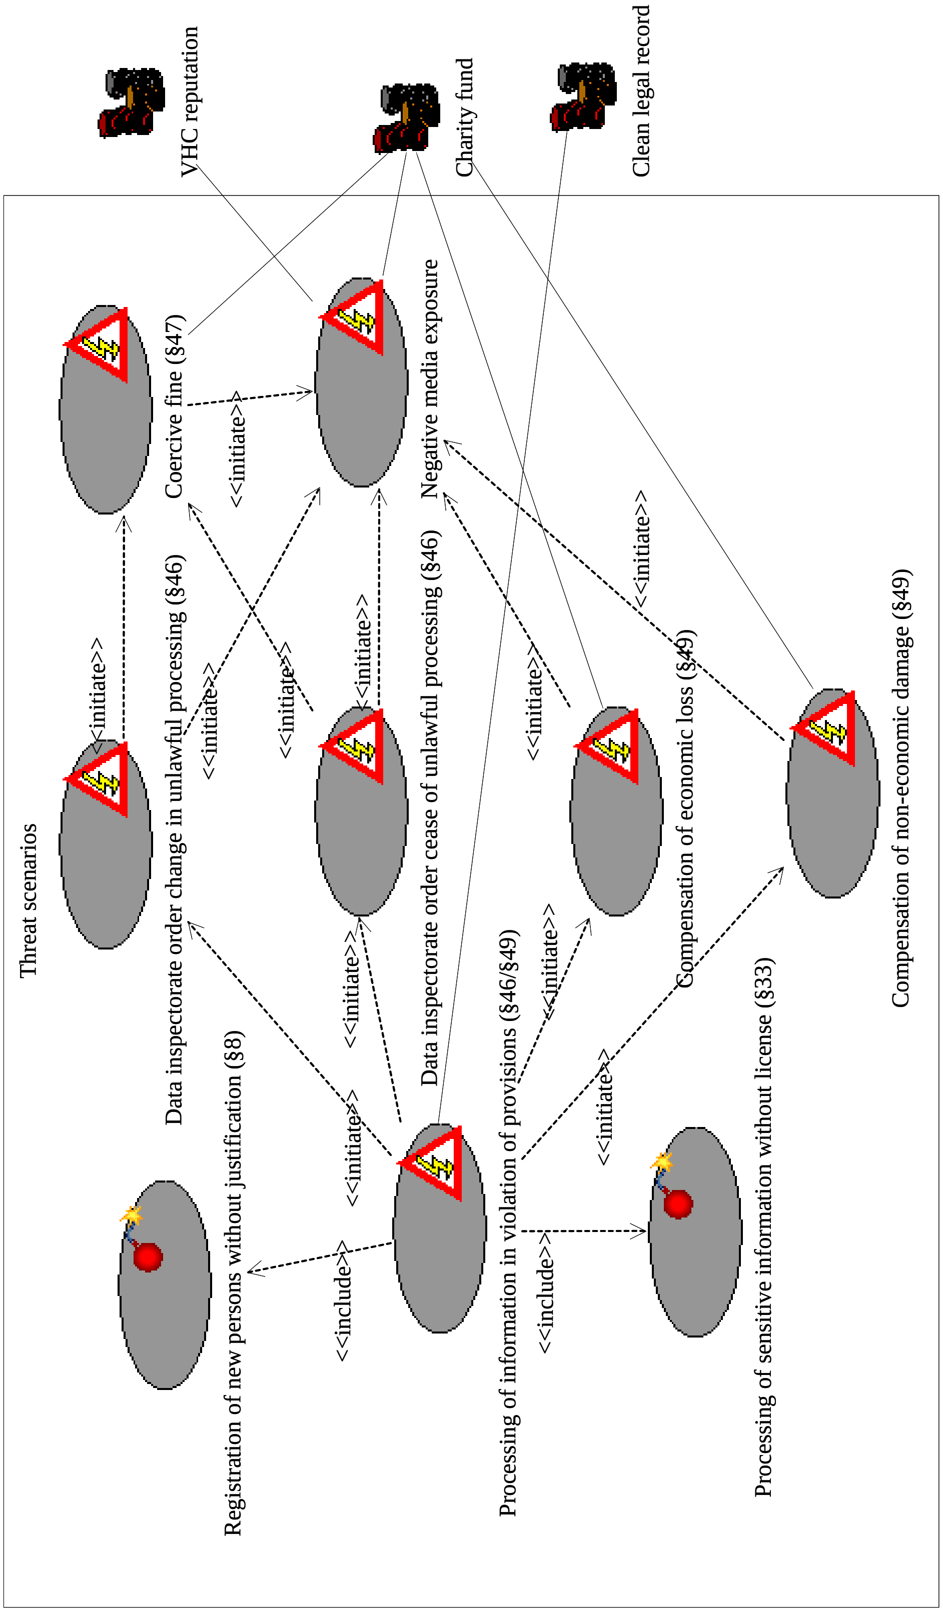
Threat scenarios
VHC reputation
Charity fund
Clean legal record
Coercive fine (§47)
Negative media exposure
Data inspectorate order change in unlawful processing (§46)
Data inspectorate order cease of unlawful processing (§46)
Compensation of economic loss (§49)
Compensation of non-economic damage (§49)
Registration of new persons without justification (§8)
Processing of information in violation of provisions (§46/§49)
Processing of sensitive information without license (§33)
<<initiate>>
<<include>>

Figure 5: Threat scenarios

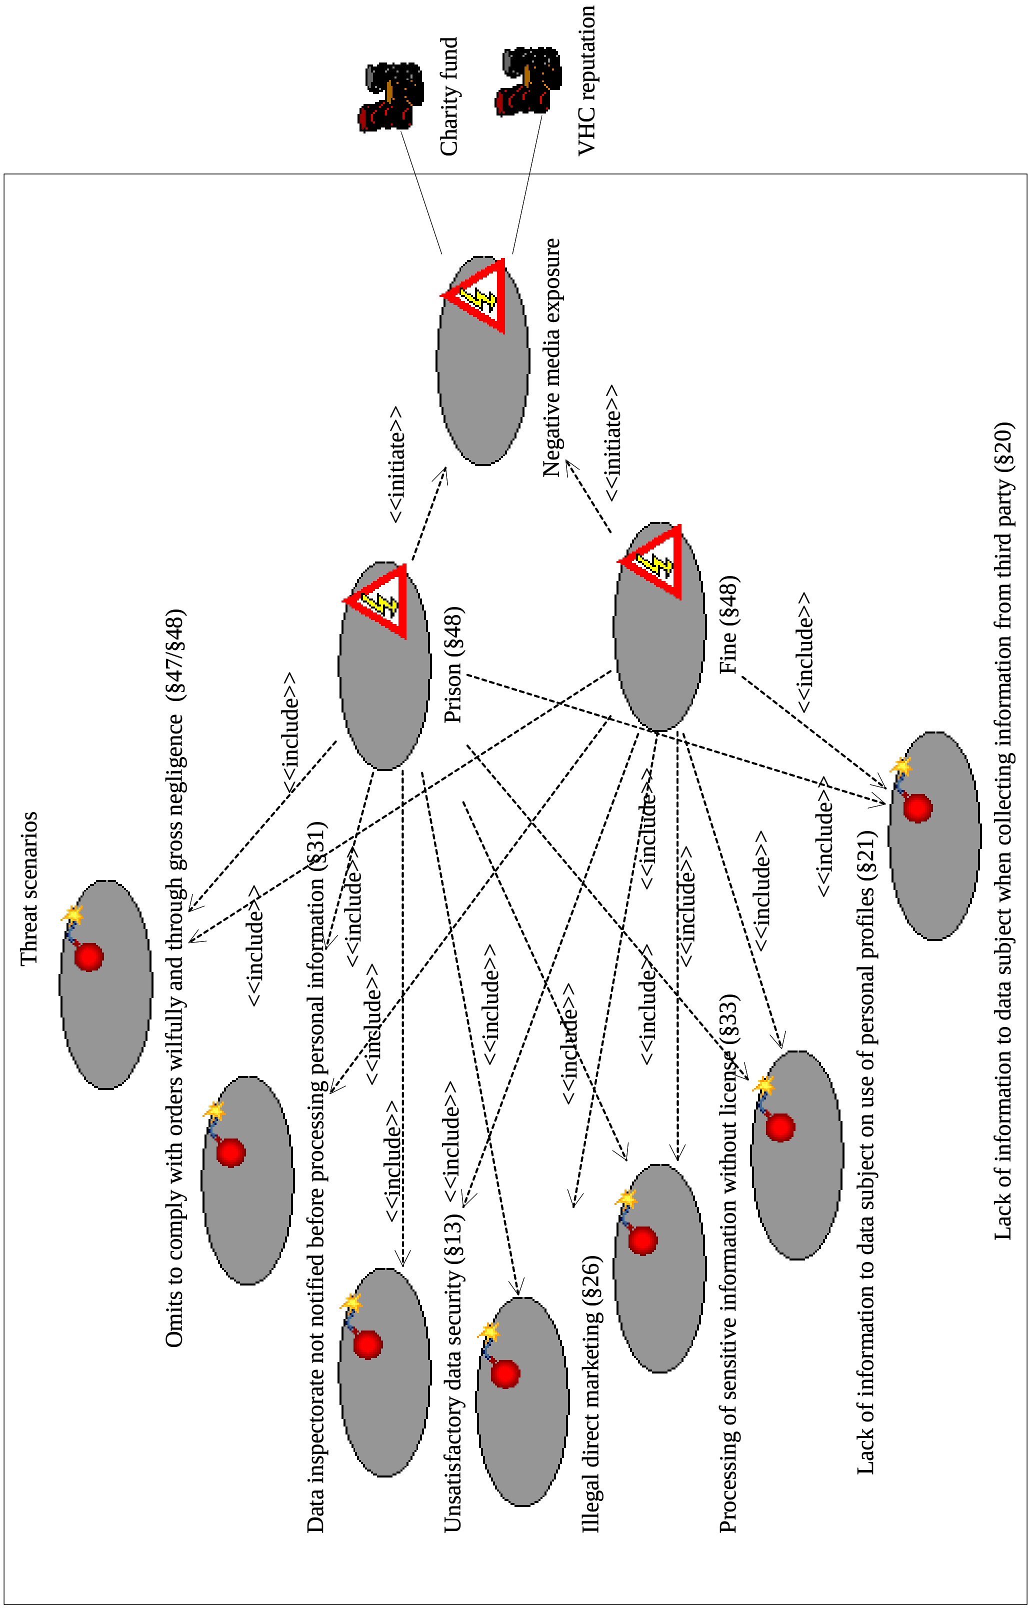
Threat scenarios
Omits to comply with orders wilfully and through gross negligence (§47/§48)
Data inspectorate not notified before processing personal information (§31)
Unsatisfactory data security (§13)
Illegal direct marketing (§26)
Processing of sensitive information without license (§33)
Lack of information to data subject on use of personal profiles (§21)
Lack of information to data subject when collecting information from third party (§20)
Prison (§48)
Fine (§48)
Negative media exposure
Charity fund
VHC reputation
<<include>>
<<initiate>>

Figure 6:Threat scenarios

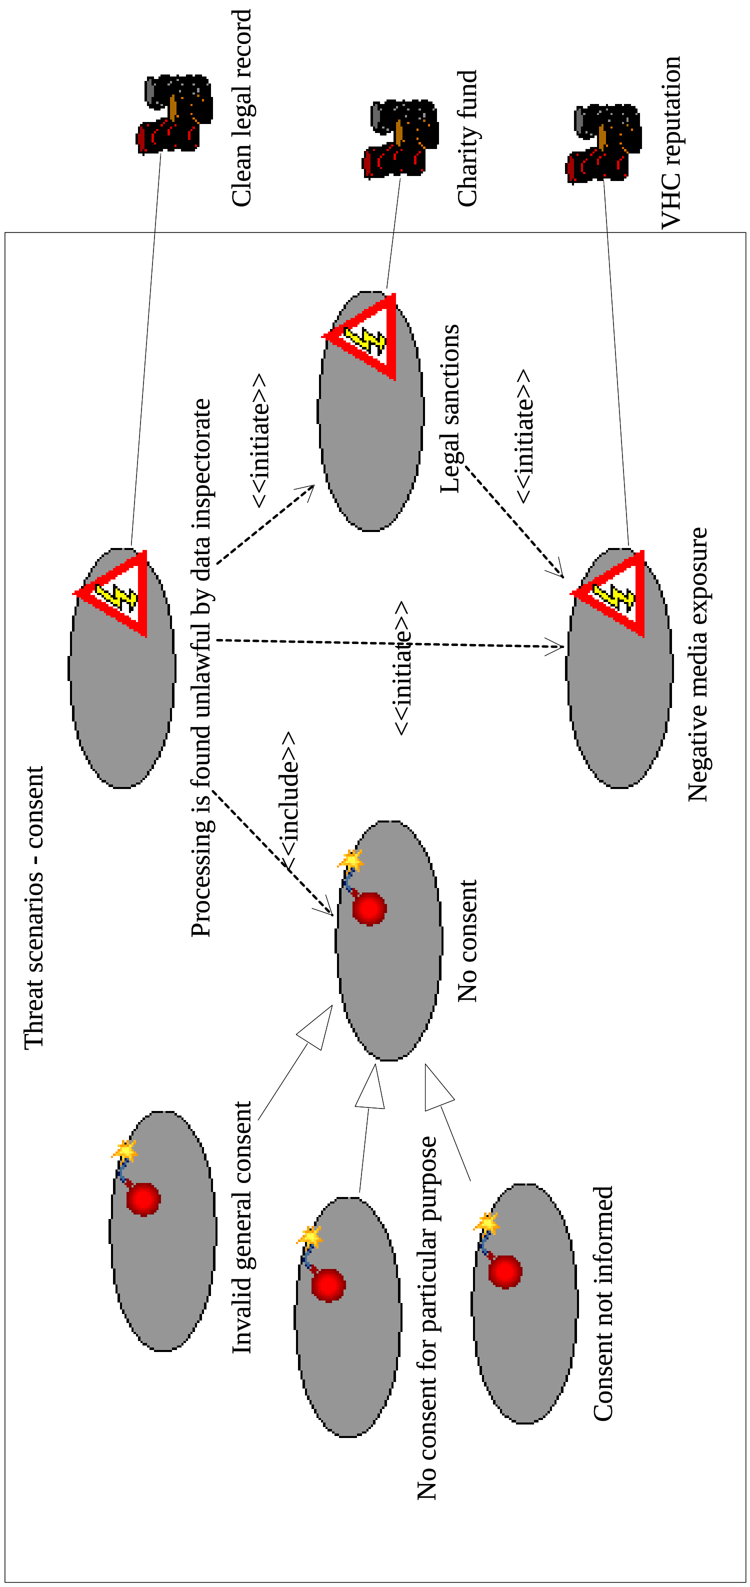

Figure 7: Threat scenarios

The structure of legal norms turned out to be difficult to represent in the threat and unwanted incident diagrams. The consequent was relatively easy to model, e.g. as an unwanted incident. However, the diagrams only allow binary relationships, making it hard to represent the relationships between the antecedent (requirements) and consequent. The antecedent may also be structured, e.g. norms combined into other norms, further exacerbating the problem. The antecedent, and thus a large part of the information contained in the legal norm, was left out of the diagrams. Other ways of incorporating this information in the model-based risk analysis need to be investigated..

5.4 Risk Analysis and Evaluation

The risk analysis and evaluation sub-processes further analyse the unwanted incidents in order to determine their risk value, forming the basis for determining which risks require treatment and which need to be monitored further. These sub-processes were not performed in this case study, but they are described briefly below.

5.4.1 Risk Analysis

After the unwanted incidents have been identified, they are further analysed to evaluate the consequences of the unwanted incidents occurring, e.g. in terms of monetary loss, and their frequency, i.e. the probability of how often they could occur. The consequence and frequency values can be either qualitative, e.g., low, medium or high, or quantitative, e.g. a concrete monetary amount. A number of methods exist for determining the consequence and frequency values, such as Fault Tree Analysis (FTA) and Markov analysis.

5.4.2 Risk Evaluation

The consequence and frequency values are combined to produce a risk value, an estimate for the level of risk Based on the risk evaluation criteria, a risk is either accepted or not, depending on its risk value. The risks that are not accepted are then prioritised, and may also be grouped into risk themes in order to make risk treatment more effective, e.g. based on similar treatment options.

5.5 Risk Treatment

The goal of the risk treatment sub-process is to determine treatment options for the risks that are not accepted during the risk evaluation sub-process. The treatment options are analysed with regard to their costs and benefits and then prioritised based on these results. Treatments can be categorised based on these general treatment strategies[46]:

- *Risk avoidance:* Avoid the risk by not performing the activity likely to generate risk
- *Reduction of frequency:* Reduce the probability of the unwanted incident occurring
- *Reduction of consequence:* Reduce the loss caused by the unwanted incident occurring
- *Risk transfer:* Transfer part or all of the risk to another party, e.g. through contracts, insurance or organisational structures
- *Risk retention:* The risk is retained, e.g. the risk is accepted but monitored

Treatments may be documented in treatment diagrams, which are similar to threat and unwanted incident diagrams. Treatment options are shown as ovals with a red cross in the corner, and are connected to the threats and unwanted incidents they treat with dashed arrows. The arrows are labelled with one of the treatment strategies listed above, e.g. <<ReduceLikelihood>>.

Figure 8 shows an example treatment diagram, containing some treatments to the threats and unwanted incidents related to consent in Figure 7. In this case there is a threat that personal information is processed for a particular purpose without the required consent. Two different treatments were proposed:

- Inform the users explicitly of how their personal information is processed and for what purposes
- Inform the users of all changes to the processing and purposes

These treatments both reduce the frequency with which the threat and related unwanted incidents occur, thus reducing the risk value.

[46] Australian Standard (1999): *Risk Management.* AS/NZS 4360:1999. Stathfield: Standards Australia.

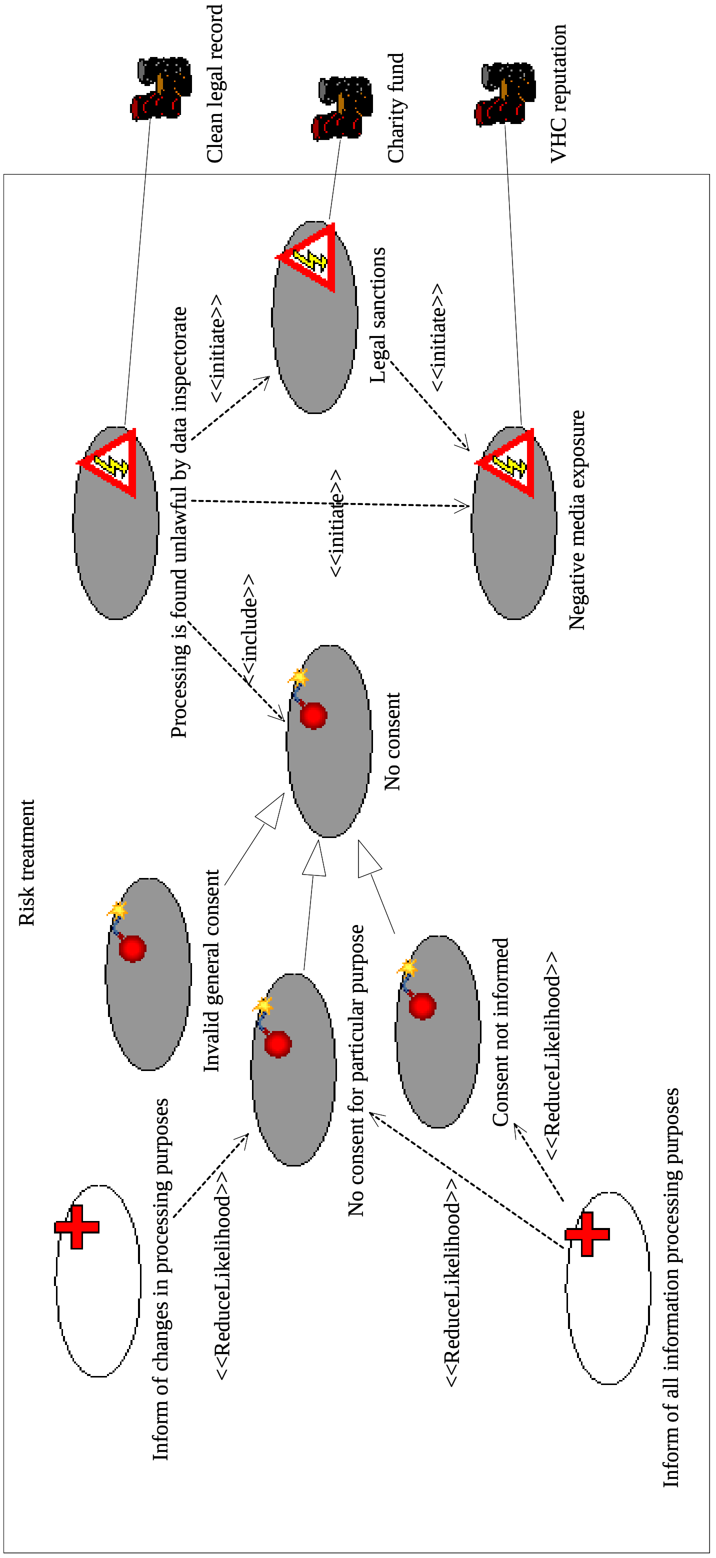

Figure 8: Treatment diagram

5.6 Concluding Remarks with Respect to Legal Risk Analysis

During this case study we have shown that methods from traditional risk analysis, in particular HazOp analysis, can also be applied successfully to identification and analysis of legal risks. The need for additional expressive power was identified. The efficiency of legal risk analysis as opposed to conventional legal analysis has not been studied. However, the idea is not to replace conventional legal analysis, but rather to provide a complementary method. The results of the risk analysis and the conventional legal analysis indicate the necessity of integrating these approaches. Hence, the Appendices A and B include integrated analyses, where conventional legal analysis and risk analysis are combined.

Legal risk analysis can be used to do an identification and prioritisation of legal risks, based on the assets determined by the stakeholders. Risk analysis may in this way help guide the application of conventional legal analysis to the areas which are of highest importance, thus increasing effectiveness.

Furthermore, legal risk analysis enables us to incorporate legal aspects into the overall risk management, producing a more complete picture of the risks to a system or organisation.

Legal risk analysis can also help provide additional insight to conventional legal analysis by including a greater range of consequences than purely legal ones. The other consequences, e.g. negative publicity, may be of higher importance to the stakeholders and are hence also major motivation factors for their actions. This may be valuable input to legal analysis, e.g. in determining the appropriate severity of legal sanctions such as fines.