

Uncertainty, Subjectivity, Trust and Risk: How It All Fits Together

Bjørnar Solhaug¹ and Ketil Stølen^{1,2}

¹ SINTEF ICT

² Dep. of Informatics, University of Oslo
{Bjornar.Solhaug,Ketil.Stolen}@sintef.no

Abstract. Trust management involves the identification and analysis of trust relations. However, adequately managing trust requires all the relevant aspects of trust to be addressed. Moreover, which aspects to address depend on the perspective of the trust management. In this position paper we relate the notion of trust to the notions of uncertainty, subjectivity and risk, and we explain how these aspects should be addressed and reasoned about from three different perspectives.

Keywords: Aleatory uncertainty, epistemic uncertainty, subjective, objective, trust, risk, trust management.

1 Aleatory Uncertainty vs. Epistemic Uncertainty

The future is uncertain and may be difficult to foresee. Nevertheless, we often make predictions about the future, be it about the weather for tomorrow, the outcome of a gambling, the stock prices, or the dependability of ICT systems. One way of making such predictions is to determine the uncertainty about outcomes of the future by means of statistical methods and probability theory. When we make predictions it is always with respect to some entity or part of reality. In the following we use the term *system* to refer to the object for prediction.

Uncertainty is often classified into two kinds [4]. On the one hand we may be uncertain about the future due to ignorance and lack of evidence. On the other hand uncertainty may be due to the inherent randomness of systems. The latter kind of uncertainty is commonly referred to as *aleatory uncertainty* and pertains to chance. Typical examples are the outcomes of the tossing of a coin, or the hands players of a game of poker receives. Aleatory uncertainty is the inherent randomness that cannot be removed from systems (without redesigning the systems). The former kind of uncertainty is commonly referred to as *epistemic uncertainty* and pertains to our knowledge about the system at hand. When making predictions about future behavior, the epistemic uncertainty is something we actively seek to reduce by gathering more information and evidence.

In determining future behavior, we may identify the possible outcomes of a situation and assign a probability $p \in [0, 1]$ to each outcome. In cases of perfect knowledge and where p is close to 0 or 1, the outcome is almost certain; there is no epistemic uncertainty and close to no aleatory uncertainty. However, if p gets

closer to 0.5, the outcome is increasingly uncertain, as for example the outcome of tossing a coin. Should knowledge be imperfect, on the other hand, we have a degree of epistemic uncertainty. This kind of uncertainty can, for example, be documented by using probability intervals $P \subseteq [0, 1]$ instead of exact values. The correct probability is then assumed to be a value $p \in P$. Using intervals instead of exact values is a form of underspecification that reflects epistemic uncertainty. While new knowledge is gathered, the underspecification can be reduced by narrowing the interval and thereby making a more precise prediction.

2 Objective vs. Subjective

The term *objective* commonly pertains to the existence of an object outside the consciousness and independent of the subject's perception of the object. Also the properties and qualities of the object are independent of the subject; objectivity and the objective are therefore typically associated with the true and the factual reality. The term *subjective*, on the other hand, pertains to the subject and how the subject perceives an object. The properties and qualities assigned to an object depend on the subject's perception, and may hence differ from one individual subject to the other; subjectivity and the subjective are therefore typically associated with the false and the possibility of wrong perceptions.

3 Trust vs. Risk

Trust is a relation between a trustor (an actor) and a trustee (an entity) and concerns the expectations of the trustor about the future behavior of the trustee. The same trustor can trust the same trustee for different purposes and in varying degrees depending on the purpose, for which reason there can be several trust relations between one pair of trustor and trustee.

The level of trust reflects the uncertainty about the future behavior of the trustee. When placing trust, the trust level is the prediction of the trustor and can be specified by a probability $p \in [0, 1]$ ranging from complete distrust to complete trust. This uncertainty can be both aleatory and epistemic. The aleatory uncertainty is associated with the inherent possibility of the trustee both to prove trustworthy and to deceive. The epistemic uncertainty reflects the degree to which the trustor has access to evidence about the trustee.

A further aspect of trust is that different actors may trust the same entity for the same purpose to different degrees, even if the histories and contexts are the same. This is because trust is about how the trustor perceives the trustee. In other words, trust is subjective. This is captured by Gambetta [2] who defines trust as the subjective probability by which the trustor expects that the trustee performs a given action on which the welfare of the trustor depends.

The notion of *trustworthiness*, on the other hand, pertains to the objective. The trustworthiness of an entity is an inherent quality of the trustee, and can be defined as the objective probability by which the trustee performs a given action on which the welfare of the trustor depends. When the subjective estimate,

i.e. the trust level, is too high or too low, we say that trust is ill-founded. In the former case the trust is higher than the trustworthiness, which means that the probability of deception is higher than what the trustor believes. In the latter case the trust is lower than the trustworthiness, which means that the probability of deception is lower than what the trustor believes.

Trust is inherently related to *risk*, because in a situation of trust the welfare of the trustor is at stake, and there is always a possibility of deception [1]. At the same time trust is related to *opportunity*, because if the trustor performs as trusted to it may have a positive outcome for the trustor [6]. Risk is commonly defined as the probability of the occurrence of a harmful event [3]. The risk level is given as a function from the consequence (loss) of the event and the probability of its occurrence. We define opportunity as the probability of the occurrence of a beneficial event [5]. The opportunity level is given as a function from the consequence (gain) of the event and the probability of its occurrence.

When interacting based on trust, the trustor typically seeks to maximize opportunities while minimizing risks. Importantly, the opportunity and risk estimates of the trustor are subjective beliefs of the trustor that the trustor acts upon; the objective opportunity and risk can only be derived from the trustworthiness of the trustee.

In managing trust we need to understand both aleatory and epistemic uncertainty, and to understand both the subjective and the objective aspects of trust. Moreover, all these aspects must be taken into account when assessing the involved risks and opportunities.

4 Trust Management from Three Perspectives

What trust management is about depends on behalf of whom the trust, risk and opportunity are to be managed. We use an example of online poker to explain the differences. The poker web application connects the players online and serves as the house. Two actors involved are Alice who is an online poker player and Bob who is the network provider. As a trustor in this setting, Alice needs to assess whether the deck of cards (i.e. the web application) is fair and whether Bob is just. As the network provider, Bob's concern is the trust of his customers, e.g. that the online players trust him not to act as one of the online players while observing the other players' hands. The two trust relations we consider are hence one between Alice and Bob and one between Alice and the deck. There are also other trust relations, but for simplicity we focus on these two only. Finally, we have Claire who is the owner of the online poker house and whose main concern is to make money.

Figure 1 illustrates various concerns related to trust management. In the upper part we have the relevant aspects of the factual reality, namely trust (the subjective perception of the reality) and trustworthiness (the objective qualities of the reality). In the lower part we have the target of investigation when uncovering the factual reality. In the case of trust, the target of investigation is the trustors, whereas in the case of trustworthiness, the target of investigation is the trustees.

<i>Factual reality</i>	Trust Subjective probability of Bob being just and deck being fair	Trustworthiness Objective probability of Bob being just and deck being fair
<i>Uncovering factual reality</i>	Target of investigation Alice	Target of investigation Bob and deck

Fig. 1. Uncovering the factual reality

4.1 Trust Management on Behalf of the Trustor

When the trust manager acts on behalf of the trustor, the task is to help establish the trustworthiness of the trustees involved. In other words, to help the trustor to reduce epistemic uncertainty and get the aleatory uncertainty right. With respect to our example, the task of the trust manager could be to help Alice to correctly assess the trustworthiness of the two trustees involved, namely Bob and the deck.

This kind of trust management might be thought of as a risk analysis with the trust relation as target of analysis and the stake of the trust relation as the asset to be protected. In order to correctly manage the involved risks, the trust must be well-founded, which is precisely the task of the trust manager to ensure.

4.2 Trust Management on Behalf of the Trustee

When the trust manager acts on behalf of the trustee, the task is to help the trustee to maintain or increase the trustee's reputation among a group of trustors; in principle, independent of whether this reputation is equal to, lower or higher than the trustee's trustworthiness. However, one may argue that an "honest" trust manager will not attempt to rise the reputation of the trustee beyond the trustees' trustworthiness.

With respect to our example, the trust manager might act on behalf of either the deck (i.e. the web application) or Bob. In the case of the former, the client (the one who pays for the trust management) might be the software manufacturer.

In trust management on behalf of the trustee, if the objective is to defend the trustees current reputation one could conduct a defensive risk analysis with the trust relation as target of analysis and the current reputation of the trustee as the asset. Note that the trustor's perception of the trustee is an important ingredient in such a risk analysis; it is the reputation of the trustee and not its trustworthiness that is to be defended. If the objective is to rise the trustees current reputation, one might employ the more offensive kind of risk analysis conducted by share traders where the level of risk tolerance is balanced against the level of opportunity.

4.3 Trust Management on Behalf of the System Owner

When the trust manager acts on behalf of the system owner, the task is to assess the impact of the trust relations within the system on the overall behavior or some quality of the system.

With respect to our example, the system in question includes Alice, Bob, the deck, the underlying software and infrastructure, and the two trust relations. The task of the trust manager might be on behalf of Claire to assess whether trust relations within the online poker games might be exploited to implement a scam that would ruin Claire by manipulating the players' trust in either Bob or the deck; to the extent that Alice and other players make trust-based decisions, the trust relations have direct impact on the business risks and opportunities for Claire.

Again, an important ingredient in this kind of trust management is risk analysis. Such a risk analysis would be rather conventional with the important exception that the description of the target of analysis would contain explicit trust relations. As in any risk analysis, a crucial ingredient is to make predictions and determine uncertainty. Because the target in this case includes trust relations, these relations have direct impact on the overall behavior of the system. In other words, the aleatory uncertainty of the future behavior of the system depends on both the subjective trust and the objective trustworthiness.

Acknowledgments. This work has been partially funded by the European Commission via the NESSoS (256980) network of excellence and by the Research Council of Norway via the DIGIT (180052/S10) project.

References

1. Castelfranci, C., Falcone, R.: Social trust: A cognitive approach. In: *Trust and Deception in Virtual Societies*, pp. 55–90. Kluwer Academic Publishers (2001)
2. Gambetta, D.: Can we trust trust? In: *Trust: Making and Breaking Cooperative Relations*, Electronic edn., ch. 13, pp. 213–237. Department of Sociology, University of Oxford (2000)
3. ISO/IEC: ISO/IEC 13335-1 Information technology – Security techniques – Management of information and communications technology security – Part 1: Concepts and models for information and communications technology security management (2004)
4. O'Hagan, T.: Dicing with the unknown. *Significance* 1(3), 132–133 (2004)
5. Refsdal, A., Solhaug, B., Stølen, K.: A UML-based method for the development of policies to support trust management. In: *Trust Management II – Proceedings of the 2nd Joint iTrust and PST Conference on Privacy, Trust Management and Security (IFIPTM 2008)*. IFIP, vol. 263, pp. 33–49. Springer, Heidelberg (2008)
6. Solhaug, B., Elgesem, D., Stølen, K.: Why trust is not proportional to risk. In: *Proceedings of the 2nd International Conference on Availability, Reliability and Security (ARES 2007)*, pp. 11–18. IEEE Computer Society (2007)