

Taste and Trust

Audun Jøsang, Walter Quattrociocchi, Dino Karabeg

► **To cite this version:**

Audun Jøsang, Walter Quattrociocchi, Dino Karabeg. Taste and Trust. Ian Wakeman; Ehud Gudes; Christian Damsgaard Jensen; Jason Crampton. 5th International Conference on Trust Management (TM), Jun 2011, Copenhagen, Denmark. Springer, IFIP Advances in Information and Communication Technology, AICT-358, pp.312-322, 2011, Trust Management V. <10.1007/978-3-642-22200-9_25>. <hal-01568693>

HAL Id: hal-01568693

<https://hal.inria.fr/hal-01568693>

Submitted on 25 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Taste and Trust

Audun Jøsang¹, Walter Quattrociocchi², and Dino Karabeg³

¹ University of Oslo, Norway

josang@matnat.uio.no

² University of Siena, Italy

walter.quattrociocchi@istc.cnr.it

³ University of Oslo, Norway

Dino.Karabeg@ifi.uio.no

Abstract. Although taste and trust are concepts on clearly distinct ontological levels, they are strongly interrelated in several contexts. For instance, when assessing trust, e.g. through a trust network, it is important to understand the role that personal taste plays in order to correctly interpret potential value dependent trust recommendations and conclusions, in order to provide a sound basis for decision-making. This paper aims at exploring the relationship between taste and trust in the analysis of semantic trust networks.

Keywords: Trust, Reputation, Taste, Recommender, Semantic, Opinion, Knowledge, Agents

1 Introduction

Understanding the dynamics between social artifacts such as reputation and trust and their effects on both opinions formation, revision and decision making is an important challenge. Several studies have focused on modeling their interdependencies and on predicting their evolutions [3, 15, 19, 18], e.g. focusing on how the dynamics of decisions are affected by the society in which the individual is immersed. In fact, the dynamic nature of social interactions is considered to play a fundamental role [11] in the formation of taste and trust. Taste as an aesthetic, sociological, economic and anthropological concept refers to a cultural patterns of choice and preference. The term "taste" is also commonly interpreted as the bio-chemical assessment of food. In this study we primarily consider the sociological interpretation of taste, i.e. where it reflects the subjective judgment of things such as styles, manners, consumer goods and works of art. Social inquiry of taste is about the human ability to judge what is beautiful, good and proper.

The social artifacts such as reputation and trust are built upon the acceptance (or refusal) of a social standard or norm. Dealing with decision making, either reputation or trust act as selection criteria in the mind of the individuals and, in turn, are based upon cognitive attitudes.

Social psychology offers an extensive literature on attitude change models, as reviewed by [13]. Most influential in social psychology is the "The Social Impact Theory" [16], according to which the amount of influence depends on the distance, number,

and strength (i.e., persuasiveness) of influence sources. As stated in ([2]), an important variable, poorly controlled in current studies, is structure topology. Interactions are invariably assumed as either all-to-all or based on a spatial regular location (lattice), while more realistic scenarios are ignored.

Within this universe, in this paper we outline the role of taste relative to trust and reputation. We will start by describing two simple examples to illustrate the effect of taste in trust situations. For example, assume that a trusted (without being more specific) friend recommends watching a movie, but you find the movie so unpleasant to watch that you leave before the end. As a result your trust in your friend's movie recommendations will drop, but interestingly your friend might still insist that it was a good movie. It is perfectly plausible that your friend genuinely likes movies that you dislike, which is unproblematic, so your friendship is not challenged. Assume now a different situation where a trusted friend recommends a car mechanic for servicing your car, but you end up very disappointed because the mechanic left oil marks everywhere and charged an unreasonably high amount for a relatively simple job. In this situation it is not plausible that your friend claims that the mechanic did a good job, because oil marks and unreasonable prices are undisputedly negative, i.e. they are taste independent aspects. You would therefore not find it credible if your friend genuinely thinks that the mechanic did a great job, and if he did you might suspect that there are other motives behind. Not only would your trust in your friend's ability to recommend car mechanics drop, you would also find your friend irrational which could cause you to distrust him in general and even damage your friendship. The only rational response from your friend would be to agree that the mechanic did a horrible job, and to apologize for having recommended such a bad mechanic, in which case your mutual trust and friendship would survive. These two examples show that people can very well agree to disagree on taste dependent aspects, but can not agree to disagree on taste independent aspects. Determining what is taste dependent or independent thus has implications for analysing trust networks.

The concept of trust is relatively well studied in the trust management literature, where authors mostly agree that it is a rather overloaded concept with many different meanings. In order to foster a meaningful discussion it is therefore useful to define the exact type of trust that is intended when the term is used. In this paper we will distinguish between so-called *evaluation trust* and *decision trust* that can be defined as follows[9].

– *Evaluation Trust*

Evaluation trust is the subjective probability by which an entity, *A*, expects that another entity, *B*, fulfills a function on which *A*'s welfare depends.

– *Decision Trust*

Decision trust is the extent to which one party is willing to depend on something or somebody in a given situation with a feeling of relative security, even though negative consequences are possible.

Evaluation trust is conceptually much simpler than decision trust because evaluation trust only depends on assumed qualities of the trusted entity itself, whereas decision trust depends on additional parameters, such as utility and risk attitude of the trusting party, that in fact are external to the trusted entity.

Since taste and trust are both subjective, and they both express an entity's opinion about some concrete or abstract object there seems to be a close correspondence between the two concepts. A noteworthy aspect of trust is that it is typically assumed that relying party faces a (potential) risk exposure because of some value at stake with potential loss, and secondly that there is some uncertainty about whether the trusted object will fulfill its expected function. In case of taste, i.e. when the relying party likes or dislikes something, risk exposure or uncertainty are not normally assumed as important cognitive elements.

The term "trust" is both a verb and a noun, such as in "*Bob trusts the car mechanic*" and "*Bob has full trust in the car mechanic*", which both have the same semantical basis in the expectation of quality service. The term "taste" is also both a verb and a noun, such as in "*Bob tastes the food*" and "*Bob has a taste for food*", but which interestingly do not share the same semantic basis. To "taste the food" is the bio-chemical assessment of a particular food sample, whereas "taste for food" is the sociological ability to judge the quality of food in general. The verbs "to like" and "to dislike" correspond better to the sociological interpretation of taste as a noun, meaning that Bob's taste for food corresponds to which food he likes or dislikes. To say "*Bob likes the food*" would then express that the food has good quality according to Bob's sociological taste for food.

The concept "opinion" or more specifically the verb-like expression "to have an opinion about" is in many ways a generalisation of the verbs "to trust/distrust" and "to like/dislike". In that sense, the nouns "trust" and "taste" should be considered as subcategories of the more general noun "opinion". In this background we will consider "trust" and "taste" to be two variants of a more general concept which could be called "belief" or "opinion".

A *scope*⁴ expresses what the opinion is about, e.g. the specific type(s) of trust or taste assumed in a given situation. In other words, the object is assumed to fulfill certain quality functions such as being reliable/unreliable or good/bad, and the scope is what the subject assumes those functions to be.

A scope can be narrow or broad. For example, a relatively broad trust scope could be "to be competent in car mechanics", whereas a relatively narrow trust scope could be "to know how to change wheels".

Trust transitivity means, for example, that if Alice trusts Bob who has a specific opinion about something, then Alice will tend to adopt the same opinion. This assumes that Alice is actually aware of Bob's opinion. This could e.g. be achieved through a *referral* or *recommendation* from Bob to Alice.

Let us assume two separate scenarios. In scenario 1) Alice needs to get her car serviced, so she asks Bob to recommend a good car mechanic. In scenario 2) Alice considers going to a rock concert, but she does not know the band, so she asks Bob about his opinion. This situation is illustrated in Fig. 1.

Fig. 1. Scenarios of recommended taste and trust

⁴ The terms "trust context" [5], "trust purpose" [9] and "subject matter" [12] have been used in the literature with the same meaning.

In TNA-SL (Trust Network Analysis with Subjective Logic), there is a separation between functional and referral trust [8]. Alice’s trust in Bob is considered to be referral, because Bob will not service the car or play the music, Bob will just refer to something else. In contrast, Bob’s trust in the mechanic and Bob’s taste for the band is functional, because they will service the car and play the music respectively.

A recommendation is equivalent to a referral, and is precisely what allows trust to become transitive. At the same time, a referral always assumes the existence of a functional scope at the end of the transitive path, which in this example is about being a good car mechanic, or a good rock band.

A transitive trust path stops with the first functional taste/trust edge encountered when there are no remaining outgoing referral edges. It is, of course, possible for a principal to have both functional and referral trust in another principal, but that should be expressed as two separate trust edges.

In practice the last functional edge in a trust chain can very well be either described as ”likes/dislikes” or as ”trusts/distrusts”, or in more general terms as ”*has opinion about object Y.*” The referral edges are clearly meaningful when interpreted as ”trusts/distrusts”, but their interpretation can also be generalised as ”*has opinion about advice from entity X.*” This seems reasonable because the main quality of a referral is to correctly recommend somebody’s opinion, be it trust or taste. There is thus little room for taste dependence in the recommendations. However, the referral trust edge assumes that the relying party shares the same taste dependent scope as the functional edge.

The ”referral” variant of a scope can be considered to be recursive, so that any transitive opinion chain, with arbitrary length, can be expressed using only one scope with two variants. This principle is captured by the following criterion.

Definition 1 (Functional opinion derivation). *Derivation of functional opinion (e.g. taste or trust) through transitive referrals requires that the last edge represents a functional opinion (e.g. about taste or trust), and that all previous edges represent referral opinions.*

When generalising ”taste” and ”trust” as two specific types of a subject entity’s opinion about the quality of an object, it is possible to let any semantic opinion scope be part of a transitive trust path. This opinion scope can express trust, taste or other beliefs. Let the relying party be denoted by A , let the intermediate recommender agents be denoted by X_i and the target object by Y . Then a transitive semantic path can be expressed as:

$$A \xrightarrow{\text{referral opinion}} X_1 \xrightarrow{\text{referral opinion}} \dots \xrightarrow{\text{referral opinion}} X_n \xrightarrow{\text{functional opinion}} Y \quad (1)$$

In practical situations, a scope can be characterised by being general or specific. For example, knowing how to change wheels on a car is more specific than to be a good car mechanic, where the former scope is a subset of the latter. Whenever a given scope is part of all the referral and functional scopes in a path, a transitive path can be formed based on that scope. This can be expressed with the following consistency criterion.

Definition 2 (Opinion scope consistency). *A valid transitive opinion path requires that there exists a scope which is a common subset of all opinion scopes in the path. The scope of the derived opinion is then the largest common subset of all scopes along the path.*

Trivially, every edge in a path can carry the same scope. Transitive opinion propagation is thus possible with two variants (i.e. functional and referral) of the same opinion scope.

The examples above assume binary beliefs about the recommending agents along the transitive path, and about the target object. In reality an opinion is never absolute, and researchers have proposed to express trust as discrete verbal statements, as probabilities or other continuous measures. For example, in TNA-SL (Trust Network Analysis with Subjective Logic) [8] arguments are expressed as subjective opinions which can contain degrees of belief and uncertainty.

2 Trust and reputation with different preference cliques

It is common that different subjects express different scores and/or rankings of the same objects as a function of some quality criterion. When two subjects give a different assessment of the same object or observation there can be at least two explanations. One explanation is that they interpreted the observation differently so that the two subjects in fact perceived two different things, e.g. due to different observation conditions or cognitive capabilities. We will exclude this explanation in the analysis below, and assume that the observations are objective. The other explanation is then that the subjects have a different preferences for the same objects, e.g. due to a different internalised value set. The effect of different preferences is extensively used in collaborative filtering and recommender systems, where a group of users who express similar preferences is called a *clique* or *neighbourhood*. The existence of preference cliques is important for trust models and for the analysis of semantic trust networks.

The easiest situation to analyse is when all subjects belong to the same clique, i.e. they have the same value set when assessing objects. For example assume that a digital music file is corrupted so that it can not be played on any playback device. All persons would normally assess a corrupted mp3 file to be bad. This situation is illustrated in Fig.2 where agents *A* and *B* belong to the same preference clique.

Fig. 2. Homogeneous preference clique

Reputation systems work best in case of a homogeneous preference clique. When it can be assumed that observations are objective, and that assessments are based on the same value set, then global ranking of the quality of objects is meaningful. Working within a single clique also simplifies the the analysis of trust networks.

Situations of different cliques with different value sets are very common. Assume for example that two subjects *A* and *B* assess a piece of classical music and a piece of rock music, where *A* likes classical music but dislikes rock music, and *B* likes rock music but dislikes classical music. In this situation, *A* and *B* clearly belong to different cliques regarding their taste for music. This is illustrated in Fig.3.

Fig. 3. Heterogeneous preference cliques with mutual acceptance

In this situation it would not be meaningful to let ratings from *A* and *B* be used to rank pieces of classical and rock music together because they have a different value set for judging music. Still, both *A* and *B* would normally respect and accept that other people have a different taste, and would therefore not question each other's moral integrity. The meaningful analysis of trust relationships between entities in different preference cliques would require to explicitly capture and express the difference in taste in the evaluation of trust opinions, so that a person would naturally distrust/discount a recommender who belongs to a different preference clique with respect to a specific scope. Several authors have proposed models inspired by this principle, e.g. [1, 4, 6, 14, 17, 20].

Another situation is if *A* has had her car serviced, and the mechanic has left oil marks on the car seats, which *A* assesses to reflect horrible car service. Let us assume that *B* assesses oil marks on car seats to reflect great car service. In this situation it is likely that *A* will think there is something wrong with *B*'s judgment and might question his moral integrity, as illustrated in Fig.3.

Fig. 4. Heterogeneous preference cliques with non-acceptance

In this situation it is plausible that the irrational assessment by *B* could make *A* suspicious about *B*'s moral integrity for other trust scopes in addition to car service, because if *B* clearly makes an irrational assessment within one scope it is likely that *B* might also be irrational for other scopes. A simple reputation system applied to this situation where ratings from *A* and *B* carries the same weight would not be very meaningful. Additional functionality would be needed to make it more useful, e.g. enabling ratings from *B* to be excluded or discounted. However, this principle would still not be useful in all situations, because it is *a priori* not evident whose ratings should be considered as irrational. In the case of election or voting for example, people often belong to different preference categories, but their votes carry the same weight. Voting systems allow people belonging to different preference cliques to cast their vote do decide who will be elected to political seats.

Trust functions can exhibit variability and still be the same type of function. For example, there can be an infinity of different movies that can all be judged as "quality

movie”. Similarly, there can be an infinity of different music performances that can all be judged as quality performances, some being classical music and some being pop music.

In theory it might be possible to refine variable scope functions so that in the end there is no room for variability, in which case it becomes a specific taste-independent scope. However, refining taste dependent scopes so that they become taste independent seems impractical, and nearly impossible in most cases.

Taste dependent scope seems to be manageable in case of trust systems where subjective trust is to be derived. However, for reputation systems, taste dependent scopes would significantly reduce the quality of the derived reputation scores. In case of reputation systems, it seems meaningless to let classical music lower the rate of pop music, simply because they will not be able to judge the quality of the music. In that case it would be necessary that the reputation system discounts ratings from classical music lower than about pop music, because their judgment will have a weak basis. This is precisely the principle behind reputation systems that take trust relationships into account, as e.g. proposed in [1, 4, 6, 14, 17, 20].

3 Semantic networks

Online social networks are growing fast these days, and it is fascinating to try to understand their meaning and implications. All social networks have in common a set of three fundamental building blocks which are entities, pointers, and attributes. These are in essence the same fundamental building blocks required for the Semantic Web and for Semantic Networks which in general can represent relationships between concrete entities or abstract concepts, or both. An entity is something that exists, virtually or physically. It can be a person like you and me, an organisation, a Web site, a document, or a service. An abstract concept can be thought of as an object class, where an entity is an object instance.

An entity’s identity serves as the pointer for accessing or navigating to the entity. Entities can point to other entities, and to each other, to form a graph. An entity pointing to another entity represents a directed edge in the graph and implicitly suggests a semantic relationship between the entities. A set of attributes can be used to describe the nature of the relationship, e.g. as “belongs to”, “likes” or “trusts”. Without attributes it would only be possible to define pure mathematical graphs. Equipped with attributes, the graphs come alive and can represent semantic and social networks as illustrated in the somewhat ego-centric example below.

Fig. 5. Semantic Network

There is of course an infinity of entities and relationships out there, so it is impossible to formally represent them all. In the semantic network illustrated above, only Walter, Audun and Dino are person-entities, and they typically have some semantic relationship to each other. The Walter-Audun-Dino relationships are just a tiny part of the

whole semantic network, which shows that semantic networks are much more than just social networks.

The fact that Audun trusts Walter and Dino can be very useful in the sense that it can help Audun to learn something from them. This now becomes a semantic trust graph derived from the underlying semantic network. In fact, the missing - but potential - link between Audun and Mister X's TV station in the diagram, or the missing link between Audun and music Y, can be derived from the existing semantic trust graph.

An example of an implementation of this type of semantic network is the social website Rumble.com that allows people to express their taste for various things, and where they also can express levels of trust in each other. The computational trust engine inside Rumble.com uses Subjective Logic which represents trust by the three parameters Trusted - Distrusted - Don't Know. Some people might find it strange to include "Don't Know", but this is really essential for representing trust. In fact, there is a default "Don't Know" trust relationship between any pair of person-entities which in practice can take the value of the base rate trust between entities in the community. A link gets initialised with base rate trust, and then evolves dynamically with various degrees of trust.

4 Implications for knowledge management

We have a subjective view of the world philosophically seen. At the same time, we often strive towards a shared view, such as common knowledge. Wikipedia shows how technology may be combined with innovative social organization to establish a common body of knowledge from distributed subjective sources: Instead of only publishing distinct documents about a given subject, Wikipedia allows the global community to create and maintain a single document for each subject. While a single-document solution dramatically reduces effort and aggregates trust, Wikipedia-style unification is feasible only for undisputed facts, although various attempts have been made to add degrees of trust and/or reputation to Wikipedia [7] to reflect relative quality of contributions. Wikipedia therefore tends to exclude controversial content. Knowledge federation is a collaborative knowledge model aimed at exploiting possibilities between those two extremes. While aiming to identify and highlight that which is relevant and shared as community view, knowledge federation leaves room for dissent and difference of opinions. A federated organization of documents becomes in effect a representation of the collective state of mind, displaying both the points of agreement and the still contending individual positions.

While collaborative document federation *à la* Wikipedia is a proven way of creating shared documents, a specific value matrix object enables a complementary approach: by *crowdsourcing* or *federating* the information needed for corresponding judgment and decisions. A value matrix is an object – attached to a resource throughout its lifetime – that accumulates all data potentially useful for evaluating this resource [10]. The columns define the criteria and the rows define available ways of evaluating those criteria. Implemented as an object, the value matrix gathers and stores relevant information, and provides suitable functions for evaluating the resource. By separating data collection from decision making, a value matrix object accumulates every piece of data that

may be of value when making a decision, without judgment or bias, and affords free choice of criteria, including the ones reflecting taste and trust. In a socio-technical system where every resource (author, user, document, idea) has an associated value matrix object, creative systemic analysis and derivation of taste and trust become accessible. The development of a suitable theory, technical tools and practices are the subjects for further research.

5 Conclusion

Trust and taste are closely related, and it is necessary to take differences in taste into account when analysing trust networks. We argue that trust and taste are specific types of the more general concepts of belief or opinion with a given semantic scope, so that trust networks can be generalised as semantic opinion networks. The analysis of trust networks can therefore be generalised to the analysis of semantic opinion networks. In such networks it is possible, and necessary, to take differences in taste and preferences into account for modelling, analysis and decision making. This has important implications e.g. for the collaborative management of knowledge.

References

1. Touhid Bhuiyan, Yue Xu, Audun Jøsang, Huizhi Liang, and Clive Cox. Developing trust networks based on user tagging information for recommendation making. In *The 11th International Conference on Web Information Systems Engineering (WISE 2010)*, Hong Kong, 12-14 December 2010.
2. Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591+, June 2009.
3. John Deighton and Kent Grayson. Marketing and seduction: Building exchange relationships by managing social consensus. *Journal of Consumer Research: An Interdisciplinary Quarterly*, 21(4):660–76, 1995.
4. Jennifer Golbeck. Generating predictive movie recommendations from trust in social networks. In *Proceedings of The Fourth International Conference on Trust Management*. Springer, 2006.
5. T. Grandison and M. Sloman. A Survey of Trust in Internet Applications. *IEEE Communications Surveys and Tutorials*, 3, 2000.
6. Claudia Hess and Christoph Schlieder. Trust-based recommendations for documents. *AI Communications - AICOM*, 21(2-3):145–153, 2008.
7. Christian Jensen. Supporting multi-agent reputation calculation in the wikipedia recommender system. *IET Information Security*, 4(4):273 – 282, 2010.
8. A. Jøsang, R. Hayward, and S. Pope. Trust Network Analysis with Subjective Logic. In *Proceedings of the 29th Australasian Computer Science Conference (ACSC2006), CRPIT Volume 48*, Hobart, Australia, January 2006.
9. A. Jøsang, R. Ismail, and C. Boyd. A Survey of Trust and Reputation Systems for Online Service Provision. *Decision Support Systems*, 43(2):618–644, 2007.
10. D. Karabeg and R. Lachica. Knowledge federation as a principle of social organization of knowledge creation and sharing. In *Proceedings of the First International Workshop on Knowledge Federation*. CEUR, Aachen, Dubrovnik, October 2008.

11. Jan Lorenz. Continuous opinion dynamics under bounded confidence: A survey. *Int. Journal of Modern Physics C*, 18(12):1819–1838, 2007.
12. G. Mahoney, W. Myrvold, and G.C. Shoja. Generic Reliability Trust Model. In A. Ghorbani and S. Marsh, editors, *Proceedings of the 3rd Annual Conference on Privacy, Security and Trust*, St.Andrews, New Brunswick, Canada, October 2005.
13. W.A. Mason, F.R. Conrey, and E.R. Smith. Situating social influence processes: Dynamic, multidirectional flows of influence within social networks. *Personality and Social Psychology Review*, 11(3):279, 2007.
14. Paolo Massa and Paolo Avesani. Trust-aware recommender systems. In *Proceedings of the 2007 ACM conference on Recommender systems (RECSYS 2007)*, RecSys '07, pages 17–24. ACM, 2007.
15. Maxwell McCombs. Setting the agenda: the mass media and public opinion. *Public Opinion Quarterly*, 2004.
16. A. Nowak, B. Latane, and J. Szamrej. From private attitude to public opinion: A dynamic theory of social impact. *Psychological Review*, 97(3):362–376, 1990.
17. John O'Donovan and Barry Smyth. Trust in recommender systems. In *Proceedings of the 10th international conference on Intelligent user interfaces*, pages 167–174. ACM, 2005.
18. W. Quattrociocchi, R. Conte, and E. Lodi. Simulating opinion dynamics in heterogeneous communication systems. *ECCS 2010 - Lisbon Portugal*, 2010.
19. W. Quattrociocchi and M. Paolucci. Reputation and uncertainty. a fairly optimistic society when cheating is total. In *ICORE 2009 - International Conference on Reputation*. Gargonza, Italy, 2009.
20. Patricia Victor, Chris Cornelis, Martine De Cock, and Paulo Pinheiro Da Silva. Gradual trust and distrust in recommender systems. *Fuzzy Sets and Systems - FSS*, 160(10):1367–1382, 2009.