

Argument Schemes for Reasoning about Trust

Simon Parsons^a, Katie Atkinson^b, Karen Haigh^c, Karl Levitt^d,
Peter McBurney^e, Jeff Rowe^d, Munindar P. Singh^f, and Elizabeth Sklar^a

^a *Brooklyn College and Graduate Center, City University of New York*

^b *University of Liverpool*

^c *Raytheon BBN Technologies*

^d *University of California, Davis*

^e *King's College London*

^f *North Carolina State University*

Abstract. Trust is a natural mechanism by which an autonomous party can deal with the inherent uncertainty regarding the behaviors of other parties and the uncertainty in the information it shares with those parties. Trust is thus crucial in any decentralized system. We build on recent efforts to use argumentation to reason about trust. Specifically, we provide a set of *schemes*, abstract patterns of reasoning that apply in multiple situations, geared toward trust. We describe, in the form of a set of critical questions, the situations in which the schemes may default.

Keywords. Argument schemes, critical questions, trust.

1. Introduction

Trust can be considered to be a mechanism for managing the uncertainty about autonomous entities and the information they deal with. As a result, trust can play an important role in any decentralized system. As computer systems have become increasingly distributed, and control in those systems has become more decentralized, trust has become steadily more important within Computer Science [7]. There have been studies, for example, on the development of trust in e-commerce [13,16,21], on mechanisms to determine which sources to trust when faced with multiple conflicting sources [5], and on mechanisms for identifying which individuals to trust based on their past activity [9,11]. Trust is especially important from the perspective of autonomous agents and multiagent systems.

In their influential model of trust, Castelfranchi and Falcone [4] maintain that trust should be reason-based, which suggests argumentation as a mechanism for constructing arguments (reasons) for and against adopting beliefs and pursuing actions, and explicitly recording the agents that need to be trusted. Several recent approaches examine such reasoning [14,15,17,18].

In order to develop a comprehensive model of argumentation for reasoning about trust, it is necessary to identify those patterns of argumentation that apply

to trust. We make a first attempt to do that in this paper. In particular, we follow Walton et al. [19] in identifying argument *schemes* (i.e., patterns for constructing arguments) and capturing, in the form of *critical questions*, the ways in which the schemes may fail (ways in which the arguments derived by the schemes may default).

2. Argument Schemes for Trusting

A quick perusal of the literature on trust suggests that there are a number of different ways in which one individual may infer trust in another. McKnight and Chervany [12] list a number of contexts in which trust may occur, and Jøsang et al. [10] distinguish between *functional* trust, the trust in an individual to carry out some task, and *referral* trust, the trust in an individual’s recommendation. If we are using argumentation to reason about trust, then we need to be able to capture these different approaches within an argument since identifying the steps in reaching a conclusion is fundamental to argumentation.

Here we collect a number of standard *schemes* or patterns of constructing arguments that relate to trust. In particular, we consider arguments about whether to trust an individual B . None of these patterns represent deductive reasoning—all may be wrong under some circumstances, and some may be wrong more often than they are right—but all represent forms of reasoning that are *plausible* under some circumstances. Because the schemes are only plausible, there are conditions under which they can fail. For each scheme we identify these failure conditions in the form of critical questions (see Section 3), in essence asking if the assumptions that underlie the schemes hold in the case we are interested in.

Note that in all the schemes, context is important. B is trusted only in some particular context for some behavior, and we assume that these schemes are only being applied in that context relating to that behavior. Thus, as we discuss later, there is a critical question relating to context that applies to every one of the schemes that we consider.

2.1. The Schemes

The schemes¹ are all stated in positive terms, as reasons to trust an individual:

Trust from Direct Experience (DE) If A has personal experience of B , and found B to be reliable, then A may decide that B is trustworthy. As a result, A accepts that B can be trusted. In this case, and in contrast to the cases below, A validates B for themself.

As an example, consider reasoning about whether a restaurant can be trusted to cook good food. When I have visited a certain restaurant in the past and found that the food is good, then I might trust the restaurant to produce good food in the future. My trust typically increases with the number of visits during which I

¹All the schemes are named by analogy to those in [19] which are called “Argument from X ” for various X .

have had a good meal. This argument is similar to “Argument from commitment” [19, page 335], where an individual is argued to be committed to some proposition on the evidence that what they said or did shows that they are so committed. Here, however, commitment is typically repeated in order to reinforce trust.

Trust from Indirect Experience (IE) If A does not have direct experience of B , but has observed evidence that leads it to believe that B has been reliable, then A may develop trust in B .

Thus I might develop a degree of trust in the quality of food served at a restaurant, not because I have visited it, but because I have walked passed it on a number of occasions and observed, for example, the fact that the tables are always full, and there are people waiting to get in.

This form of argument is distinguished from the previous case because of the fact that the experience is indirect. In the case of the restaurant example, I haven’t tried the food, I am just inferring from what I see that it is good. Unlike the direct experience case, my inference about the food could be wrong—the crowds might reflect the fact that it is a fashionable restaurant that serves mediocre food, that it has a famous DJ who regularly spins there, or that the owners are regularly providing free but indifferent meals for their friends in the hope that it will entice other people to come in.

Trust from Expert Opinion (EO) If B is an expert in some domain of competence, then A may trust B .

In this case, B is validated by some entity other than A —for example B is a doctor, and is validated by the appropriate medical board, or, to continue the restaurant example, that B is a graduate of the Culinary Institute of America. This validation is distinct from that of the next case in the sense that it is independent of the position that B holds, indeed, B holds the position *because* of the validation. The validation is also certificated—there is some concrete evidence of the validation. This is distinct from the previous case in that it is the expert who is providing evidence of B , not A .

In some cases—such as the case where B is a chef— A is able to validate B ’s trustworthiness for itself (by eating in the restaurant) and so augment an argument from expert opinion (a review by a critic who is expert on restaurants) with an argument from direct experience. In other cases—such as the case where B is a doctor—it isn’t clear that A can perform any direct validation, and instead A always infers part of its trust in B as a result of the certification of B ’s expertise.

This argument scheme is clearly related to the “Argument from expert opinion” from [19, page 15].

Trust from Authority (Au) If B is in a position of authority, then A may trust B .

We distinguish this case from the previous one in the sense that here the trustworthiness of B stems from the position that B holds rather than from external certification. A doctor is trusted (on medical matters) whether or not they

work for a hospital, as long as they have not been struck off (and possibly even then—a doctor who is struck off for financial irregularities presumably still knows how to diagnose). An employee of the MTA (the New York body that runs the public transport), for example, might well be trusted on MTA matters while they are on duty, but not if they no longer have any affiliation with the transportation authority and so might be presumed to no longer have access to up-to-date information.

Trust from Reputation (R) If B has a reputation for being trustworthy, then A may choose to trust B .

We consider that A may either have heard other individuals saying that B is trustworthy, or be aware of some aggregate measure of reputation that applies to B . To continue with the restaurant example, A may hear reports from friends, or may distill the reputation of B from a website like Yelp.

We distinguish the idea of an argument based on reputation from an argument based on expert opinion or on authority because the recommenders are neither experts (if A 's friends are food critics then, in our view, A would be using the expert opinion scheme), nor are they authorities. We also note that an argument from reputation is not the same as propagating trust. If A hears that B has a good reputation, this is a statement about B 's trustworthiness. It, like all the other schemes we consider here, is the establishment of a link between A and B . (Information derived from reputation may be used in propagating trust between individuals, but the use of the reputation scheme does not imply propagation.)

Trust from Moral Nature (MN) If A judges that B has a good character, then A may trust B .

Here, A is performing some inference about A that is grounded not in A 's knowledge of B 's past behavior (as in direct or indirect experience), nor is A 's view guided by B 's professional expertise, nor by B 's position in society. Rather, A is making some observation about other aspects of B and inferring trustworthiness from that. [19, page 141] classifies this kind of argument as the “Aristotelian ethotic argument”.

Trust from Social Standing (SS) If A judges that B would have too much to lose by being untrustworthy, then A may trust B .

Here, A performs a kind of expected utility calculation in terms of B 's position, asking what B could gain by not being trustworthy and what B could lose by being exposed, and deciding that the former is less than the latter. The possible loss is, roughly speaking, related to B 's position in society. In a real community, this is, we believe, a large part of what motivates trust between peers (as opposed to between folk who are externally certified), and is one of the reasons that trust is a challenge online, where it is so easy to hide behind anonymity.

Trust from Majority behavior (M) If A has found most people in the set from which B is drawn to be trustworthy, then A may choose to trust B .

This is an even less deductive form of trust derivation than most of the above, but it is still one we use. For example, when buying online, many of us are happy to trust our data to merchants we have no specific recommendation about because of our overall good experience with online merchants. In effect, we are generalizing from some experiences with some entities—individuals or organizations—to all entities in a group (possibly while being aware that not all entities in the group are trustworthy).

Trust from Prudence (Pru) *A* may decide to trust *B* because it is less risky than not trusting *B*.

The key to the prudence scheme is the assessment of risk involved in trusting *B*. As an example, consider the case where you are running late for an important meeting, but are now lost and almost certain to miss the meeting. This is a situation in which it makes sense to ask for directions from *B*, even though you don't know whether *B* is particularly trustworthy. There is a chance that the directions will be good, and you'll get there on time, and taking this risk is better than continuing to blunder around not knowing your way.

Trust because of Pragmatism (Pra) *A* may decide to trust *B* because it (currently) serves *A*'s interests to do so.

The pragmatism scheme only considers *A*'s current interests. If these are aligned with *B*'s, then *A* will trust *B* despite any other reasons to distrust *B*. An example of pragmatic trust is that where one coalition partner trusts another just because the two are part of a coalition, and so are currently working towards the same goals.

2.2. Discussion

The schemes break down into four categories. The first category is one in which *A* has collected evidence about the trustworthiness of a given individual. This category includes two schemes, direct and indirect experience, relating to whether the evidence that *A* has collected is direct evidence of *B*'s trustworthiness, or evidence of something from which trustworthiness can be derived.

The second category includes three schemes—expert opinion, authority and reputation—where *B* is considered to be trustworthy on some subject, or in some role, not because *A* has observed them doing this (directly or indirectly) but because there is some validation of *B* that can, in theory at least, be verified by *A*. If *B* is somehow validated as an expert or an authority, *A* can reasonably trust them. Of course, there are many cases of supposed experts or authority figures being shown to be frauds, but the fact that the argument scheme can lead to misplaced trust is not an argument against the scheme so much as an argument for careful posing of the critical questions. In the reputation scheme, *A* bases their decision on a distillation of reported evidence from other individuals. Since those individuals aren't typically validated as experts or authorities (someone reviewing a restaurant on Yelp might be a chef, but it is hard to establish whether this is the case), what *A* relies on in the reputation scheme is the “wisdom of the crowd”. In

other words that the average of a number of reported experiences is likely close to the experience that *A* will have. It is important to distinguish reputation from reasoning based on *referral trust*. Reputation is established from many individuals and refers to one specific recommendation. Referral trust is trust in an individual's ability to make trustworthy recommendations, and so is established about one individual, typically based on many recommendations.

The third category of argument schemes contains two schemes, moral nature and social standing, that are based on *A*'s observations of *B*, but observations that are not directly linked to trust. *A* might use the moral nature scheme to infer that *B*, whom *A* has observed to be very correct in their dealings with others, will act in a trustworthy way in a business deal. Similarly, *A* might use the social standing scheme to infer that *B*, who is a pillar of the local community, will not default on a loan. In neither case does *A* have any information about *B*'s trustworthiness, but is prepared to use aspects of *B* for which *A* does have information as a proxy for such evidence.

The final category contains three schemes, majority, prudence and pragmatism. We can think of the majority scheme as a way of extrapolating any of the previous schemes. If we can show that any of those schemes for deriving trustworthiness apply to a suitably high proportion of a given population for which we do have evidence, we may be happy to infer that some member of the population for which we do not have evidence is also trustworthy. The prudence and pragmatism schemes are grouped with majority because all of them are about deriving trust in an individual without any specific knowledge of that individual. The majority scheme does this on the basis of the class the individual falls into. Prudence does it by assessing the comparative risk of trusting versus not trusting (where the reasoning is about the situation not the possibly trustworthy individual). Pragmatism derives trust from a consideration of *A*'s goals and *B*'s alignment with them, not on any information about whether *B* is trustworthy in their own right.

Two other schemes we considered but rejected are "Precedent", where *A* has trusted *B* before and decides *B* can be trusted again because there were no bad consequences from the previous time(s) *B* was trusted, and "Default", where *A* decides that *B* can be trusted despite having no evidence, perhaps with something like "tit-for-tat" [1] in mind². We do not list these schemes because we think they are already incorporated in schemes we have listed rather than new schemes. Precedent is a form of direct experience, and as we shall see, one of the critical questions that applies to direct experience addresses *B*'s past behavior. We consider Default to be a form of the majority scheme. After all, if it was not the case that an agent believed that the majority of individuals it was going to interact with would be trustworthy, then the default scheme would not be rational (in other words the default scheme is only rational for populations for which the majority scheme suggests trusting individuals).

Figure 1 shows an alternative characterization of our argumentation schemes that is based on a conceptual model of trust showing a trustor and a trustee interacting in an environment and in a society. *Pragmatism*, based on self interest, is an attribute of the trustor and *moral nature* is an attribute of the trustee.

²Where it is a winning strategy to start out trusting others to be cooperative.




Figure 1. Categorizing argumentation schemes for trust.

Expert opinion links the trustee to the environment and *prudence* (viewed as suitability of decision-making) links the trustor to the environment. The trustor has *direct experience* of the trustee and *indirect experience* via the environment. A trustee has *social standing* in the society and may be in a position of *authority* over the trustor. The society includes agents with resemblance to the trustor and trustee, respectively, and support distinct trust relationships: *reputation* via the former and *majority* via the latter.

The above formulation establishes a form of *conceptual completeness* for our schemes. Given the conceptual model of trust, there are no schemes that capture different relationships between the elements of the model (though there may be other schemes that capture aspects of the same relationships).

3. Critical Questions for Trusting

According to Walton et al. [19], the role of critical questions is to capture the defeasibility inherent in argumentation schemes. Critical questions are:

... questions that can be asked (or assumptions that are held) by which a non-deductive argument based on a scheme might be judged to be (or presented as being) good or fallacious. [19, page 15]

Here we list critical questions for each of the schemes presented above, each phrased so that if all questions for a given scheme are answered “yes”, then the scheme can be used to create a plausible argument. There are also two general questions, which can be posed to any argument:

CQ1 Is this context one in which we have established that *B* is trustworthy?

CQ2 Are the negative consequences of misplaced trust sufficiently small that we can discount the possibility that *B* is not trustworthy?

Most of the critical questions relate to *belief*—is it likely that *B* is or is not trustworthy? In many cases, the utility (positive if *B* is trustworthy, negative if *B* is untrustworthy) needs to be taken into account as well. CQ2 is one way to begin to capture that.

Having listed these general questions, we turn to argument-specific questions.

Direct experience There are two critical questions for arguments based on direct experience.

- DE1** Is *B*, the person in whom *A* is placing trust, really the person *A* has experience of, or is *B* an imposter?
- DE2** Is *A* sure that *B* has not been acting in a trustworthy way in order to obtain *A*'s trust and then deceive them?

The first of these questions aims to identify the situation in which we have a long series of interactions with *B* and learn to trust them. Then we are presented with a new interaction with someone who claims to be *B* but turns out to be *C*, someone we have never interacted with before. The second question identifies a “bait and switch” scenario in which *B* has acted as if they were trustworthy in order to later betray our trust.

Note that the first question, as we have described it, covers the case where *C* steps in to deceive *A* by pretending to be someone, *B*, that *A* already trusts. The question also covers an alternative scenario where *C* has pretended to be *B* during a series of interactions with *A* so that *A* believes that the person that they trust and have potentially recommended to others, *C*, is actually *B*.

Indirect experience Again there are two critical questions:

- IE1** Can trust be inferred from the evidence?
- IE2** Is *B*, the person in whom trust is being inferred, really the person who should be trusted?

The first of these addresses the fact that evidence is uncertain and so, to reprise our restaurant example, the fact that there is a large number of people waiting for a table in the restaurant might not indicate that the food is good. The second question makes a finer distinction. If the food in the restaurant is good, is this an indication that we should trust food provision by the restaurant owner (so that we can infer that other restaurants owned by the same person are also good), or by the chef (so that if the chef moves to a different kitchen, we should infer that food from that source is good too)?

Expert opinion Here we have critical questions that are analogous, though distinct, from the identity question for direct and indirect experience—is *B* really an expert?

- EO1** Does *A* have proof that *B* is really an expert?
- EO2** Is *B*'s expertise relevant in this case?
- EO3** Is *B* suitably credible as an expert [19, page 15]?
- EO4** Is *B*'s opinion consistent with that of other experts?
- EO5** Does *A* know that *B* will not benefit as a result of the opinion they stated?

The difference between the first question here and DE1 is that with DE1, we are interested in *who* *B* is, namely whether *B* is the person we have learnt to trust. With expert opinion, we are vesting all our belief about trustworthiness in some certifying body. EO1 challenges this (and is really a multi-part challenge—is there a certificate stating that *B* is an expert, is it authentic, and does it really belong to *B*?). EO2 addresses the fact that *B* might be indisputably an expert, but just not the best expert (*B* may be a doctor, but in a very specific medical context,

although *B* might have some idea of the best course of action, might not be as good an expert as a doctor who specializes in that specific context). EO3 doesn't ask whether *B* is the right kind of doctor, but whether we think that *B* is a good doctor (maybe based on where they studied), or, as another example, not whether *B* is a lawyer, but whether *B* is an experienced lawyer. EO4 is a check that *B* doesn't hold a maverick opinion, and EO5 seeks to question whether *B*'s views have some benefit to *B* (we might not trust a restaurant critic's view of the food at a particular establishment if we knew he was being paid by the restaurant to write his review).

Authority As we said when we introduced the arguments, the argument from expert opinion and the argument from authority are similar. However, the critical questions show some of the differences between them.

Au1 Is *B* really in a position of authority?

Au2 Is *B*'s authority relevant in this case?

We are assuming here that expertise can be certified—that is the reason, after all, that doctors hang their medical certificates on their walls, mechanics hang their “authorized dealer” notifications, and restaurants in New York (have to) display the hygiene rating they were awarded by the city Department of Health. Authority, on the other hand, may sometimes be certified (by a uniform) but in other cases may be very hard to demonstrate.

Reputation Deriving trust from reputation requires that *B* has a good reputation and some assurance that reputation means something

R1 Does *B* have a good reputation?

R2 Are we sure that *B*'s reputation has not been manipulated to make it more positive?

Moral nature The questions for this scheme are derived from the critical questions that [19] gives for the ethotic argument.

MN1 Is *B* a person of good moral character?

MN2 Is character relevant in this case?

MN3 Is the degree of trustworthiness being inferred supported by the evidence?

Social standing Social standing is only a guarantee if *B* has significant social standing and there is a mechanism by which standing can be lost. The critical questions address this.

SS1 Does *B* have any social standing to lose?

SS2 Does *B* value social standing?

SS3 Would *B*'s standing be dented (or raised) by being exposed as untrustworthy?

SS4 If *B* is untrustworthy, can they be exposed in a meaningful way?

The questions address the following issues. If B has already been exposed as untrustworthy, B has nothing much to lose, and so standing is no guarantor of trustworthiness. If B doesn't care about social standing then nothing can be inferred from the potential loss of it. Further, social standing is only a deterrent if B 's peers will see exposure as reflecting badly on B . Finally, social standing is only a deterrent if it is possible to make the members of B 's social circle aware of the loss of trustworthiness. The assumption exposed by this final critical question means that SS isn't necessarily much help as an argument scheme in an online environment where tying reputation to an individual is complicated by the ease of acquiring a new identity and of maintaining anonymity.

Majority Since the majority scheme is a form of statistical argument, the need to consider Simpson's paradox³ forms the basis of a natural critical question M2 .

M1 Is B really in the class of individuals who are trusted?

M2 Is the class we are considering the most specific class that B is a member of?

As an example of M2, we might be prepared to trust online merchants in general on the basis of the majority scheme, but, on the basis of some bad experiences, might not be prepared to trust online sellers of electronics in particular.

Prudence Since the prudence scheme is about the risk of trusting B , the critical questions focus on this aspect:

Pru1 Is it riskier to not trust B than it is to trust B ?

Pru2 Is it possible to accurately estimate the risk in trusting and not trusting B ?

Pru3 Is there another individual we could trust where the risk would be lower than trusting B ?

Pragmatism The critical questions for the pragmatism scheme focus on the degree to which trusting B is in the best interests of the trusting agent:

Pra1 Does trusting B serve our best interests?

Pra2 Is there another individual we could trust such that trusting them would better serve our interests than trusting B ?

Note that all the critical questions discussed in this section are variations on what [19, page 93] calls the "trustworthiness question". In our context, where we are putting trust under the microscope, it makes sense to split the trustworthiness question into these more specific questions that are related to specific schemes.

³Formally, $\Pr(A|B) < \Pr(A|\neg B)$ but $\Pr(A|B, C) \geq \Pr(A|\neg B, C)$ and $\Pr(A|B, \neg C) \geq \Pr(A|\neg B, \neg C)$ [3].

4. Conclusion

This paper has taken a first step towards identifying argument schemes for reasoning about trust. The overall aim behind deriving these schemes is to provide a computational mechanism for establishing arguments about trustworthiness from a description of some scenario that does not itself include any explicit information about trustworthiness. We have identified ten general schemes that allow an individual to derive trust in another, and for each of these schemes, we have also identified a set of critical questions. The purpose of the critical questions is to identify cases in which the schemes might not apply—if the answer to any critical question is negative, then the scheme may not apply (or may be used to derive a lower level of trust).

A natural question to ask of such a set of schemes is whether it and the associated set of questions is exhaustive. This, of course, is hard to establish, and we believe that a full set of schemes will only emerge over time. This is certainly the case in the work of Walton, whose schemes for arguments on the basis of expert opinion have continued to develop, for example from those listed in [19] to those in [6]. In terms of the exhaustiveness of the critical questions, we can take a more pragmatic approach. One line of future work is to formalize these schemes, as in [2], as a precursor to being able to build them into a tool for reasoning about trust. As we progress with the formalization, we can ensure that questions are associated with every predicate and object in the set of formal schemes, making sure that, as desired, every aspect of a given scheme can be tested against the scenario in which it might be applied to check that it is valid.

A second line of future work is to extend the scope of the schemes. As they stand, the schemes are all concerned with generating arguments for trusting, with the critical questions capturing the cases where these schemes do not apply. They may also be the need to incorporate schemes that generate arguments for not trusting — for example *ad hominem* arguments [20] — and for schemes that generate distrust. There may even be a need for schemes that generate arguments against the arguments generated by other schemes for trusting.

Another important area of future work is to consider schemes for the propagation of trust. Propagation of trust captures the situation where, for example *A* has derived trust in *B* (perhaps on the basis of one of the schemes detailed here), *B* has derived trust in *C*, and then these two aspects are combined so that *A* has trust in *C*. As mentioned above, there are a number of standard patterns for propagating trust, for example those in [8]. Argument schemes are a natural method of capturing such patterns in a general way, along with the critical questions that identify the cases in which propagation is not permitted.

Acknowledgments: This research was funded under Army Research Laboratory and Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

- [1] R. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.
- [2] T. Bench-Capon and K. Atkinson. Argumentation schemes: From informal logic to computational models. In C. Reed and C. Tindale, editors, *Dialectics, Dialogue and Argumentation: An Examination of Douglas Walton's Theories of Reasoning and Argument*, pages 103–114. College Publications, London, 2010.
- [3] C. R. Blyth. On Simpson's paradox and the sure-thing principle. *Journal of the American Statistical Association*, 67(338):364–366, June 1972.
- [4] C. Castelfranchi and R. Falcone. Trust is much more than subjective probability: Mental components and sources of trust. In *Proceedings of the 33rd Hawaii International Conference on System Science*, Maui, Hawai'i, January 2000. IEEE Computer Society.
- [5] X. L. Dong, L. Berti-Équille, and D. Srivastava. Integrating conflicting data: The role of source dependence. In *Proceedings of the 35th International Conference on Very Large Databases*, Lyon, France, August 2009.
- [6] T. F. Gordon, H. Prakken, and D. Walton. The Carneades model of argument and burden of proof. *Artificial Intelligence*, 171(10–11):875–896, 2007.
- [7] T. Grandison and M. Sloman. A survey of trust in internet applications. *IEEE Communications Surveys and Tutorials*, 4(4):2–16, 2000.
- [8] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *Proceedings of the 13th International Conference on the World Wide Web*, 2004.
- [9] C.-W. Hang, Y. Wang, and M. P. Singh. An adaptive probabilistic trust model and its evaluation. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, Estoril, Portugal, 2008.
- [10] A. Jøsang, E. Gray, and M. Kinateder. Simplification and analysis of transitive trust networks. *Web Intelligence and Agent Systems*, 4(2):139–161, 2006.
- [11] L. Li and Y. Wang. Subjective trust inference in composite services. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, Atlanta, GA., 2010.
- [12] D. H. McKnight and N. L. Chervany. The meanings of trust. Working Paper 96-04, Carlson School of Management, University of Minnesota, 1996.
- [13] L. Mui, M. Moteashemi, and A. Halberstadt. A computational model of trust and reputation. In *Proceedings of the 35th Hawaii International Conference on System Sciences*, 2002.
- [14] S. Parsons, P. McBurney, and E. Sklar. Reasoning about trust using argumentation: A position paper. In *Proceedings of the Workshop on Argumentation in Multiagent Systems*, Toronto, Canada, May 2010.
- [15] S. Parsons, E. Sklar, and P. McBurney. Using argumentation to reason with and about trust. In *Proceedings of the 8th International Workshop on Argumentation in Multiagent Systems*, Taipei, Taiwan, 2011.
- [16] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system. In M. R. Baye, editor, *The Economics of the Internet and E-Commerce*, pages 127–157. Elsevier Science, Amsterdam, 2002.
- [17] R. Stranders, M. de Weerd, and C. Witteveen. Fuzzy argumentation for trust. In F. Sadri and K. Satoh, editors, *Proceedings of the Eighth Workshop on Computational Logic in Multi-Agent Systems*, volume 5056 of *Lecture Notes in Computer Science*, pages 214–230. Springer Verlag, 2008.
- [18] S. Villata, G. Boella, D. M. Gabbay, and L. van der Torre. Arguing about the trustworthiness of the information sources. In *Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, Belfast, UK, 2011.
- [19] D. Walton, C. Reed, and F. Macagno. *Argumentation Schemes*. Cambridge University Press, Cambridge, UK, 2008.
- [20] D. N. Walton. *Ad Hominem Arguments*. University of Alabama Press, 1998.
- [21] B. Yu and M. Singh. Distributed reputation management for electronic commerce. *Computational Intelligence*, 18(4):535–349, 2002.