



On Beta Models with Trust Chains

Tim Muller, Patrick Schweitzer

► **To cite this version:**

Tim Muller, Patrick Schweitzer. On Beta Models with Trust Chains. Carmen Fernández-Gago; Fabio Martinelli; Siani Pearson; Isaac Agudo. 7th Trust Management (TM), Jun 2013, Malaga, Spain. Springer, IFIP Advances in Information and Communication Technology, AICT-401, pp.49-65, 2013, Trust Management VII. <10.1007/978-3-642-38323-6_4>. <hal-01468207>

HAL Id: hal-01468207

<https://hal.inria.fr/hal-01468207>

Submitted on 15 Feb 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.



Distributed under a Creative Commons Attribution 4.0 International License

On Beta Models with Trust Chains

Tim Muller and Patrick Schweitzer
{tim.muller, patrick.schweitzer}@uni.lu

University of Luxembourg

Abstract. In a type of interactions over the Internet, a user (the subject) is dependent on another user (the target), but not vice versa. The subject should therefore form an opinion about the target, before possibly initiating an interaction. The scenario wherein a subject only relies on information obtained from past interactions with the target, is well-studied and understood. In this paper, we formally analyze the implication of allowing recommendations (statements of a third party) as source of information. We identify the family of valid models that admit recommendations. This allows us to verify particular existing models that admit recommendations.

1 Introduction

With the advent of the Internet, new types of interactions between different people arose. It is now possible, if not even common, to provide sensitive personal information to parties about which virtually nothing is known. For example, anyone can purchase goods from complete strangers on eBay. Contrary to purchasing goods in ordinary shops, buyers cannot inspect the commodities they acquire from an e-commerce website. Instead the shoppers have to wait and hope that everything will be delivered as ordered. In this paper, we focus on these kind of interactions, i.e. interactions where one party alone determines whether the outcome is beneficial or harmful to the other party. We call such interactions, interactions between a passive and an active party. The passive party attempts to avoid interactions with an active party that is likely to harm it. As a consequence, before potentially initiating an interaction, the passive party would like to estimate the likelihood with which the interaction outcomes are beneficial. We refer to such an estimate as a *trust opinion*. If a potentially passive party establishes a trust opinion about a potentially active party, the former is the *subject*, and the latter is the *target*¹.

In interactions over the Internet, the information which a subject has about (alleged) past behavior of a target is limited. Hence it might be beneficial to ask for the help of third parties. Third party statements about the target are called *recommendations*, hence we call these third parties recommenders. Trust opinions constructed with the help of recommendations are called *chained trust opinions*. In this paper, we formally study the implications of such recommendations.

¹ In the literature, the subject and the target are also referred to as trustee and trustor. This terminology may however lead to the incorrect conclusion that the trustee is being trusted and the trustor is trusting.

In the past, numerous formal models that derive trust opinions based on information about past behavior of active parties have been proposed. There exist simple models that allow a subject only to use his own past interactions with the target for information (see [12] for an effective method of gathering and filtering such information). For these approaches, a formal model, called the *Beta model* (or beta reputation system), has been derived [5,8].

To illustrate the Beta model, we introduce a simple running example.

Running Example. An economy teacher wants to teach her students about e-commerce with the help of a turn-based game. To set up the game, the teacher secretly distributes a random value $p_i \in [0, 1]$ to each student c_i for $1 \leq i \leq 30$. The value p_i represents the integrity of each student, and, similar to the integrity of users on an e-commerce website, it is unknown to the other players. On an e-commerce system this parameter models how likely the outcome of an interaction is to be successful. Each turn of the game follows the following pattern. First, in the turn of student c_i , the teacher assigns another student c_j to c_i . Then, c_i has the choice between trusting or not trusting c_j . In case c_i chooses to trust c_j , c_i gains two points with probability p_j , i.e. with the probability corresponding to the other student's integrity parameter. With the remaining probability of $1 - p_j$, c_i loses one point. If c_i chooses not to trust c_j , then he neither gains nor loses points. On an e-commerce platform winning points corresponds to a successful interaction (a success), losing points to a failed interaction (a failure). After every turn, the teacher updates the students' points, only revealing the outcome to c_i . Like in e-commerce, trusting someone with high integrity has a high probability to result in a successful interaction; trusting someone with a low integrity has a high probability to result in an unsuccessful interaction.

The classroom game can easily be analyzed within the Beta model. Assume that c_i previously had $s+f$ interactions with c_j . Of these $s+f$ interactions, s were successes and f were failures. With the help of the Beta model [4] we estimate the probability of a success when trusting c_j to be $\frac{s+1}{s+f+2}$, and the expected value of trusting c_j to be $2\frac{s+1}{s+f+2} - 1\frac{f+1}{s+f+2}$ points. When not trusting c_j , the points remain constant.

Suppose the next day, the teacher changes the rules of the game and allows c_i to query a classmate about his experience with c_j before having to choose whether or not to trust c_j . That expansion of the classroom game can no longer be expressed in the Beta model (as it does not admit recommendations), it requires an extension.

To overcome this challenge, many modern trust models use the Beta model as a foundation, and increase the model's expressivity and its (practical) applicability by including recommendations. We say that a model which uses the Beta model as a foundation is *in the Beta family*. If a model is in the Beta family and also supports trust chains, we say it is *in the Beta family with trust chains*. Many models in the Beta family with trust chains are ad-hoc. By ad-hoc models, we understand models in which the inventors define chained trust opinions according to their intuition. The existence of ad-hoc models is supported by the

fact that the research community has not yet settled on one trust model [7], not even under the assumption that the trust model is in the Beta family [6].

Rather than proposing a new model in the Beta family, we rigorously prove properties of trust chains valid in all models in the Beta family. We show the following properties. Chained trust opinions are modular (Proposition 4 and Theorem 3), meaning that complex trust opinions can be constructed from simpler ones. Every trust model makes implicit or explicit assumptions about how a recommender lies or about the number of interactions between users (Corollary 3). Chained trust opinions resulting have a different shape from the trust opinions in the Beta model (Theorem 4). Furthermore, Subjective Logic, an expressive ad-hoc extension of the Beta model, is not in the Beta family with trust chains (Corollary 5). The same conclusion can be derived for models similar to Subjective Logic, such as TRAVOS [13] and CertainTrust [11] (Corollary 4).

In Section 3, we formalize the notion of recommendations and add it to the Beta model, effectively formalizing all models in the Beta family with trust chains. Then, in Section 4, we study the most basic trust chains in the Beta family with trust chains. In Section 5, we prove that all models in the Beta family with trust chains have the property that trust opinions can be constructed modularly from the most basic trust chains. Finally, in section 6, we characterize trust models in the Beta family with trust chains, and show that existing models based on the Beta model are not in the Beta family with trust chains.

2 The Beta model

In this section, we introduce the Beta model. The formulation of the Beta model relies on well-known techniques from probability theory (see e.g. [1,3]). There are two concepts in particular that are important for our analysis. The first is conditional independence:

Definition 1 (Conditional independence of variables [2]). *Let (Ω, \mathcal{F}, P) be a probability space and let X, Y, Z be random variables (from Ω) with values in the measurable spaces (E_i, \mathcal{E}_i) , $i \in \{X, Y, Z\}$. Two random variables X and Y are conditionally independent given the variable Z if*

$$P(X \in A, Y \in B | Z \in C) = P(X \in A | Z \in C)P(Y \in B | Z \in C).$$

for each $A \in \mathcal{E}_X, B \in \mathcal{E}_Y$ and $C \in \mathcal{E}_Z$.

As shorthand we write $(X \perp\!\!\!\perp Y) | Z$ or even $X \perp\!\!\!\perp Y | Z$. Note that the definition is equivalent to $P(X|Y, Z) = P(X|Z)$.

And the second is the concept of beta distributions:

Definition 2 (Beta distribution). *A beta distribution is a family of continuous probability distributions in the interval $[0, 1]$, parameterized by two positive parameters, $\alpha, \beta \geq 1$. The probability density function of a beta distribution with parameters α and β is*

$$\beta(x; \alpha, \beta) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{\int_0^1 y^{\alpha-1}(1-y)^{\beta-1} dy}.$$

The expression under the fractions is known as the beta function on α and β , and for positive integers α and β , the beta function fulfills $B(\alpha, \beta) = \frac{(\alpha-1)!(\beta-1)!}{(\alpha+\beta-1)!}$.

We define the Beta model in a similar fashion to [9]. We first define a series of random variables. Let \mathbf{A} denote a set of agents. For $A, C \in \mathbf{A}$ and a set of events Ω , we then define:

- $E_C: \Omega \rightarrow \{S, F\}$ is a discrete random variable modeling the outcome of the corresponding interaction with target C .
- $R_C: \Omega \rightarrow [0, 1]$ is a continuous random variable modeling the (hidden) *integrity parameter* of target C which defines the probability of success.
- $O_C^A: \Omega \rightarrow \mathbb{N} \times \mathbb{N}$ is a discrete random variable modeling the *interaction history* of subject A about target C , representing the past interactions (number of successes and failures) between A as passive party and C as active party.

Running Example. In the classroom game, E_C models the outcome of an interaction with student C . The variable R_C describes the secret parameter initially assigned by the teacher to C and O_C^A expresses how many times student A interacted successfully with student C , and how many times it was a failure.

A trust opinion is a distribution over the integrity parameter of a target, based on the interaction history about the involved active parties. Hence, if a subject A establishes a trust opinion about a target C , the probability density function is of the form $f_{R_C}(x|O_C^A, \varphi)$, where φ may express additional conditions.

Next, we provide the assumptions of the Beta model, in the shape of dependencies and independencies of random variables, as formulated in [9]. For a more concise formulation of the (in)dependencies, we introduce sets of random variables.

$$\begin{aligned} \mathbb{E} &:= \{E_C : C \in \mathbf{A}\}, \\ \mathbb{R} &:= \{R_C : C \in \mathbf{A}\}, \\ \mathbb{O} &:= \{O_C^A : A, C \in \mathbf{A}\}, \\ \mathbb{W}' &:= \mathbb{E} \cup \mathbb{R} \cup \mathbb{O}. \end{aligned}$$

The size of the interaction histories is unknown. We therefore model it with a distribution λ , called the *entanglement*. Let $c \in [0, 1]$, $x_s, x_f \in \mathbb{N}$ and $\lambda: \mathbb{N} \rightarrow [0, 1]$ be a probability distribution. For all agents $A, C \in \mathbf{A}$ we set up the following dependency and independency relations as our assumptions.

D1 R_C is the uniform distribution on $[0, 1]$.

If we know nothing about the integrity of C , we assert all values equally likely. For specific applications, statistical data about behaviors of agents may be used to construct an alternative distribution. A suitable distribution has a probability density function that is non-zero on $(0, 1)$.

D2 $P(E_C=S|R_C=c) = c$.

We assume that the probability of good behavior of A is determined by an integrity parameter a .

D3 $P(O_C^A=(x_s, x_f)|R_C=c) = \binom{x_s+x_f}{x_s} c^{x_s} (1-c)^{x_f} \lambda(x_s+x_f)$.

Assumes that the probability that A and C had an interaction history with size x_s+x_f is $\lambda(x_s+x_f)$, and that each past interaction had success probability c .

I1' For $W \in \mathbb{W}' \setminus \{O_C^A\}$, it holds that $O_C^A \perp\!\!\!\perp W | R_C$.

The interaction history is completely determined by its size, and the probability of a success in a single interaction (by Dependency D3).

I2' For $W \in \mathbb{W}' \setminus \{R_C\}$, it holds that $R_C \perp\!\!\!\perp W | E_C \cap \bigcap_{D \in \mathbf{A}} \{O_C^D\}$.

The only indicators of the integrity parameter of C , are interactions with it.

I3' For $W \in \mathbb{W}' \setminus \{E_C\}$, it holds that $E_C \perp\!\!\!\perp W | R_C$.

The behavior of C is completely determined by its integrity parameter (by Dependency D2).

A trust opinion of A about C can now be seen as the probability density function given by $f_{R_C}(c|\varphi)$, where φ represents all knowledge of A about C , modulo the relations of the random variables. Typically, φ is equal to O_C^A , provided there are no recommendations. In this case, we call $f_{R_C}(c|\varphi)$ a *simple trust opinion*, to be able to distinguish it from trust opinions involving recommendations.

Theorem 1 (Axiomatization of the Beta model [9]). *The Beta model adheres to Dependencies D1–D3 and Independencies I1'–I3'. The simple trust opinion obtained from an interaction history with x_s successes and x_f failures is the beta distribution $\beta(c; x_s + 1, x_f + 1)$.*

Suppose there are two concurrently held trust opinions based on two different interactions with a single agent. It is desirable to combine these two trust opinions into a single trust opinion based on both interactions. We introduce a trust aggregation operator to accomplish that:

Definition 3 (Aggregation of trust opinions). *The aggregation of trust opinion $T = f(c)$ and $T' = g(c)$ is $T \oplus T' = \frac{f(c) \times g(c)}{\int_0^1 f(c) \times g(c) dc} \propto f(c) \times g(c)$.*

The trust aggregation operator correctly combines simple trust opinions:

Lemma 1. *Given trust opinions T and T' based on (x_s, x_f) and (y_s, y_f) , respectively, the aggregate trust opinion $T \oplus T'$ is based on $(x_s + y_s, x_f + y_f)$.*

Proof. $T \oplus T' \propto \beta(c; x_s + 1, x_f + 1) \times \beta(c; y_s + 1, y_f + 1) \propto \beta(c; x_s + y_s + 1, x_f + y_f + 1)$ \square

Our assumptions regarding simple trust opinions are in line with the Beta model. They are in fact sufficient to derive it (Theorem 1). Hence, those assumptions can be seen as valid for the numerous models that use the Beta model as a foundation [4,13,11].

3 Beta family with trust chains

According to the Beta model, a subject A constructs his trust opinion using only his own information, when planning to interact with a target C . Depending on the constructed trust opinion, A chooses to interact or not. Suppose that A wants to make a more informed decision. Then, the subject A may ask a third party, a recommender B , for advice. A recommender could provide an honest recommendation, or lie. Chained trust opinions are based on the notion that a trust opinion on the recommender B is a valid measure for the likelihood that B provides an honest recommendation about C . More formally:

Definition 4 (Chained trust opinions). *Every recommender (like every target) has an integrity parameter that determines the probability of a successful interaction. In case of a successful interaction, their recommendation is their trust opinion about the target. Chained trust opinions are trust opinions based on recommendations from recommenders.*

We add recommendations to the classroom game:

Running Example. After a number of turns, the students realize that the Beta model can be applied to construct a trust opinion about other students. This allows all students to make optimal choices. To keep the game interesting, as well as make it a more realistic emulation of e-commerce, the teacher adds recommendations in the following way: In the beginning of every turn, the teacher not only assigns a subject $c_i \in \{c_1, \dots, c_{30}\} =: S$ and a target $c_j \in S$, but also a set of recommenders $R \subseteq S \setminus \{c_i, c_j\}$ if c_i has never interacted with c_j . Every recommender $c_k \in R$ has to honestly provide their past interactions with c_j with probability p_k , or construct and provide a fake past history with c_j with probability $1 - p_k$. Again, students with a high integrity p_k are more likely to provide the past interactions rather than fake interactions. For a subject to construct the most accurate trust opinion, he needs to incorporate his opinion of c_k and the recommendation by c_k , for all $c_k \in R$.

To formally model recommendations in the Beta model, we introduce another random variable.

- $S_C^B: \Omega \rightarrow \mathbb{N} \times \mathbb{N}$ is a discrete random variable modeling recommendations of the recommender B about the target C , representing the alleged past interactions between B as passive party and C as active party.

We also introduce additional sets of random variables:

$$\begin{aligned} \mathbb{S} &:= \{S_C^B : B, C \in \mathbf{A}\}, \\ \mathbb{W} &:= \mathbb{W}' \cup \mathbb{S}. \end{aligned}$$

Let $a, b, x \in [0, 1]$, $n, k \in \mathbb{N}$ and $\lambda: \mathbb{N} \rightarrow [0, 1]$ as well as $\chi^B: [0, 1] \times \mathbb{N} \times \mathbb{N} \rightarrow (\mathbb{N} \times \mathbb{N} \rightarrow [0, 1])$, where $B \in \mathbf{A}$ be probability distributions. For all agents $A, B, C \in \mathbf{A}$ we set up the following additional dependency and independency relations as our assumptions. In fact, Independencies **I1'**–**I3'** from the initial Beta model only need to be generalized to encompass recommendations.

D4 $P(S_C^B=(w_s, w_f) | E_B=S, O_C^B=(w_s, w_n)) = 1$

Assumes that good behavior of B implies that the recommendation of B corresponds to his interaction history with C .

D5 $P(S_C^B=(y_s, y_f) | E_B=F, R_B=b, O_C^B=(w_s, w_f)) = \chi^B(b, w_s, w_f)(y_s, y_f)$

Defines the *lying strategy* χ^B of agent B . The lying strategy is a function, from a parameter and an interaction history $(k', n' - k')$ to a distribution of recommendations. A recommender (probabilistically) selects its fake recommendations.

I1 For $W \in \mathbb{W} \setminus \{O_C^A\}$, it holds that $O_C^A \perp\!\!\!\perp W | R_C$.

Similar to Independence **I1'**, except recommendations are also independent.

- I2 For $W \in \mathbb{W} \setminus \{R_C\}$, it holds that $R_C \perp\!\!\!\perp W | E_C \cap \bigcap_{D \in \mathbf{A}} \{O_C^D\}$.
 Similar to Independence I2', except recommendations are also independent.
- I3 For $W \in \mathbb{W} \setminus (\{E_B\} \cup \bigcup_{D \in \mathbf{A}} \{S_D^B\})$, it holds that $E_B \perp\!\!\!\perp W | R_B$.
 Similar to Independence I3', except recommendations not from B are also independent.
- I4 For $W \in \mathbb{W} \setminus \{S_C^B\}$, it holds that $S_C^B \perp\!\!\!\perp W | E_B = \mathbb{F} \cap R_B \cap O_C^B$.
 The choice of B for making fake recommendations about C is completely determined by $\chi^B(b, n, m)$ in Dependence D5.

Models in the Beta family with trust chains should adhere to Dependencies D1–D5 and Independencies I1–I4.

Definition 5 (Beta family with trust chains). *A model is said to be in the Beta family with trust chains, when it satisfies Dependencies D1–D5 and Independencies I1–I4.*

There are models that are inspired by the Beta model, and that include an operator \odot dealing with recommendations, but that are not models in the Beta family with trust chains. We argue that such models either are not Beta models or that \odot is not a trust chaining operator. If a model violates any of the Dependencies D1–D3 or Independencies I1'–I3', it is not a Beta model. We distinguish the possible violations of an assumption for each remaining assumption separately. If a model violates

- D4, then the model does not support trust chaining.
- D5, then another assumption must also be violated. This is due to the fact that under Dependencies D1–D4 and Independencies I1–I4 there exists a χ^B such that $\chi^B(b, w_s, w_f)(y_s, y_f) = P(S_C^B = (y_s, y_f) | O_C^B = (w_s, w_f), R_B = b, E_B = \mathbb{F})$.
- I1, then the model either violates Independence I1', or it assumes that some S_D^C are dependent with O_C^A given R_C . This is not in the spirit of the Beta model as the outcomes of the interactions between A and C should depend only on C .
- I2, then the model either violates Independence I2', or it assumes that some R_C are dependent with S_E^D given all observations of C . This is not in the spirit of the Beta model as the collection of all interactions with C should be an optimal estimator for R_C .
- I3, then the model either violates Independence I3', or it assumes that some E_C are dependent with S_E^D (for $D \neq C$) under all observations of C , which is not in the spirit of the Beta model as the probability of success of an interaction (given the integrity) should not be influenced by recommendations of others.
- I4, then in this model recommenders differentiate their strategy either on information they cannot know (e.g. interactions that the recommender did not participate in) or on information that is irrelevant for the recommendation (e.g. his opinion on yet another agent).

Not every model in the Beta family with trust chains is formalized our way. A model is already in the Beta family with trust chains when the assumptions can be reformulated to fit the assumptions up to isomorphisms.

4 Basic Trust Chains

The most basic scenario that involves trust chains, involves exactly one recommendation. This recommendation is given about a target with which the subject has no prior interactions. In other words, the recommendation is the only source of information that a subject has. This scenario is called *basic trust chaining*. It is studied in this section. In Section 5, we then prove that more complicated scenarios can be reduced to scenarios with basic trust chains.

Definition 6 (Basic trust chain, basic chained trust opinion). *A basic trust chain consists of three agents: the subject A , the recommender B , and the target C . The subject has an interaction history $x = (x_s, x_f)$ with the recommender. The recommender provides a recommendation $y = (y_s, y_f)$ about the target and, in reality, has an interaction history $w = (w_s, w_f)$ with the target. The trust opinion of subject A about target C with recommendations by recommender B is the basic chained trust opinion. It is depicted in Figure 1.*

Running Example. In the classroom game, basic trust chains appear when the teacher assigns only one recommender. Then, the subject is $c_i \in S$, the target is $c_j \in S \setminus \{c_i\}$ and the set of recommenders is $\{c_k\} \subset S \setminus \{c_i, c_j\}$.

We may now formulate the basic chained trust opinion of A about C with recommendations given by B as $f_{RC}(c|O_B^A=(x_s, x_f), S_C^B=(y_s, y_f))$. In other words, to formulate a trust opinion about the target, the subject uses its interaction history about the recommender as well as the (possibly fake) recommendation given by the recommender. If A has never directly interacted with B , the pair (x_s, x_f) equals $(0, 0)$.

Theorem 2 (Basic chained trust opinion). *Dependencies D1–D5 and Independencies I1–I4 are sufficient to derive the basic chained trust opinion of A about C with recommendations by B as: $f_{RC}(c|O_B^A=(x_s, x_f), S_C^B=(y_s, y_f)) =$*

$$\mathbf{eq}_1(y_s, y_f) \times \mathbf{eq}_2 + \sum_{w \in O_C^B} (\mathbf{eq}_1(w_s, w_f) \times \mathbf{eq}_3 \times (1 - \mathbf{eq}_2)), \quad (1)$$

where,

$$\begin{aligned} \mathbf{eq}_1(\varphi_s, \varphi_f) &= \beta(c; \varphi_s + 1, \varphi_f + 1), \\ \mathbf{eq}_2 &= \frac{\mathbf{eq}_4 \times (x_s + 1)}{\mathbf{eq}_4 \times (x_s + 1) + \sum_{w' \in O_C^B} \mathbf{eq}_5(w') \times (x_f + 1)}, \\ \mathbf{eq}_3 &= \frac{\mathbf{eq}_5(w_s, w_f)}{\sum_{w' \in O_C^B} \mathbf{eq}_5(w', w_f)}, \\ \mathbf{eq}_4 &= \lambda(y_s + y_f) \times \binom{y_s + y_f}{y_s} \times \frac{y_s! y_f!}{(y_s + y_f + 1)!} \\ \mathbf{eq}_5(\varphi_s, \varphi_f) &= \int_0^1 \chi^B(b, \varphi_s, \varphi_f)(y_s, y_f) \times \beta(b; x_s + 1, x_f + 2) db \\ &\quad \times \lambda(\varphi_s + \varphi_f) \times \binom{\varphi_s + \varphi_f}{\varphi_s} \times \frac{\varphi_s! \varphi_f!}{(\varphi_s + \varphi_f + 1)!} \end{aligned}$$




Fig. 1. Left: The view of subject A about target C , including the recommendation S_C^B from B about C . Right: The view of recommender B about target C .

Proof. The equations **eq1**–**eq5** represent the following probabilities:

$$\mathbf{eq1}(\varphi) = P(R_C=c|O_B^A=x, S_C^B=y, E_B=u, O_C^B=w),$$

$$\mathbf{eq2} = P(E_B=s|O_B^A=x, S_C^B=y),$$

$$\mathbf{eq3} = P(O_C^B=w|O_B^A=x, S_C^B=y, E_B=F),$$

$$\mathbf{eq4} = P(S_C^B=y|O_B^A=x, E_B=s),$$

$$\mathbf{eq5}(\varphi) = P(S_C^B=y, O_C^B=\varphi|O_B^A=x, E_B=F).$$

A proof of correctness of **eq1**–**eq5** can be found in [10]. The correctness from Formula (1) follows from the correctness of **eq1**–**eq5**, given that, for all $W \in \mathbb{W}$: $S_C^B \perp\!\!\!\perp W|E_B=s \cap O_C^B$ follows from Dependency **D4**. \square

Although Formula 1 may seem complicated, it can abstractly be viewed as a (infinite) weighted sum of beta distributions:

Proposition 1. *For every entanglement and lying strategy, a basic chained trust opinion is a weighted sum of beta distributions.*

Proof. If we collect factors that do not contain the variable c in the scalars k and k_{w_s, w_f} , Formula (1) simplifies to

$$k \cdot c^{y_s} (1-c)^{y_f} + \sum_{w_s, w_f \in \mathbb{N} \times \mathbb{N}} k_{w_s, w_f} c^{w_s} (1-c)^{w_f}. \quad (2) \quad \square$$

Furthermore, for some specific models in the Beta family with trust chains, the formula significantly simplifies. Particularly, for a lying strategy that consists of constructing truthful recommendations (see dash-dotted graph in Figure 2), the trust opinion is a beta distribution:

Proposition 2. *If $\chi^B(b, w_s, w_f)(y_s, y_f) = 1$ iff $(w_s, w_f) = (y_s, y_f)$, then the trust opinion from Formula (1) simplifies to $\beta(c; y_s + 1, y_f + 1)$.*

Taking an arbitrary entanglement λ and a lying strategy that consists of constructing completely informationless recommendations (see dashed graph in Figure 2), the trust opinion is a weighted sum of a beta distribution and the uniform distribution:

Proposition 3. *If $\chi^B(b, w_s, w_f)(y_s, y_f) = \frac{1}{y_s + y_f + 1}$ iff $w_s + w_f = y_s + y_f$, then the trust opinion from Formula (1) simplifies to $\frac{x_s + 1}{x_s + x_f + 2} \beta(c; y_s + 1, y_f + 1) + \frac{x_f + 1}{x_s + x_f + 2}$.*

An immediate consequence of Theorem 2 and Proposition 1 is that a model that supports basic chained trust opinions, makes assumptions about the entanglement and lying strategies.

Corollary 1. *It is not possible to compute basic chained trust opinions without knowledge of the entanglement λ and the lying strategy χ^B .*

Proof. Proposition 2 and 3 are not equal, hence the choice of χ^B matters. \square

Running Example. In terms of the classroom game, the corollary states that it is relevant how many turns have been played and how students lie. If a recommendation states “8 successes and 2 failures”, but each student has played 9 turns, the recommendation is clearly fake, whereas the same recommendation may be likely true when each student has had 100 turns. Suppose, a student c_k provides a recommendation to c_i that is likely to be fake. If c_k and c_i are good friends outside of the game, c_k might have a lying strategy of creating fake recommendations that strongly resemble the truth. Otherwise, c_k provides recommendations unrelated to the truth. Then, it is wise for c_i to rely on the recommendation of his friend, but not on recommendations of other arbitrary classmates.

Corollary 1 implies that without assumptions on λ and χ^B , no model can provide trust opinions. Therefore, any trust model in the Beta family with trust chains either implicitly or explicitly makes assumptions about numbers of interactions and about the lying strategy of recommenders. We believe that making implicit assumptions about lying strategies is critical, as it obfuscates the analysis of a model or hides undesirable consequences of a model. Hence, we suggest that new proposals for models in the Beta family with trust chains explicitly (and formally) provide the lying strategy of the recommenders.

Corollary 2. *For every entanglement λ and lying strategy χ^B , the subject can calculate the basic chained trust opinion.*

Proof. Apply Formula (1), with the relevant instantiations of λ and χ^B . \square

Thus, when the number of turns in the classroom game is known, and it is known what kind of lying strategy each student has, the subject can correctly compute the trust opinion, whenever the teacher assigns only one recommender.

A positive consequence of Corollary 2 is that defining the entanglement and the lying strategy is sufficient to explicitly define a model in the Beta family with trust chains. Not only is it mathematically possible, but we have developed a tool named Canephora² that can compute basic chained trust opinions, when χ^B and λ are provided. The tool is a proof of concept, that creating a model in the Beta family with trust chains is merely a matter of defining an entanglement and lying strategies. It is a prototype that allows the numerical comparison between different models (i.e. different choices of entanglements and lying strategies).

In Section 5, we see that defining the entanglements and the lying strategies is sufficient to explicitly define models in the Beta family with trust chains (not just models restricted to basic trust chains).

² <http://satoss.uni.lu/software/canephora>




Fig. 2. The same trust chain, $x = (6, 5)$ and $y = (8, 4)$, with different lying strategies. Solid: lies opposite of his true opinion. Dashed: lies independent of his true opinion. Dash-dotted: lies similar to his true opinion. Dotted: lies with a positive bias.

Determining the entanglement λ is usually simpler than finding the lying strategy. On many e-commerce systems, the number of interactions between users is known to the system. For example, eBay knows if a product is sold, even if it does not know whether the transaction was a success for the subject. Or in the classroom game, the teacher announces the number of turns, explicitly providing λ . Even if the entanglement is unknown, by restricting the choices of χ^B , the entanglement λ can be eliminated from Formula (1).

Lemma 2. *For some lying strategies, the entanglement has no impact on the basic chained trust opinion.*

Proof. Consider the basic chained trust opinion given by Formula (1). For all $b \in \mathbb{R}$, and $w_s, w_f, y_s, y_f \in \mathbb{N}$ such that $w_s + w_f \neq y_s + y_f$, take $\chi^B(b, w_s, w_f)(y_s, y_f) = 0$. Then, $\lambda(\varphi_s + \varphi_f)$ cancels out of **eq5** unless $\varphi_s + \varphi_f = y_s + y_f$. In the reduced term, we can substitute $\lambda(\varphi_s + \varphi_f)$ for $\lambda(y_s + y_f)$. Then $\lambda(y_s + y_f)$ is a scalar that appears in every summand in the numerators and denominators of **eq2** and **eq3**. Thus λ cancels out of Formula (1). \square

Running Example. If a recommender makes a recommendation of which the size was impossible (or very unlikely), a student can identify the recommendation as a fake (or likely a fake). If all students take care never to fall into the pitfall of sizing fake recommendations according to a different distribution than the real interactions, sizing becomes irrelevant. Hence, the entanglement cancels out.

5 Modular Construction of Trust Opinions

In Section 3, the assumptions of the Beta model were formally extended to include trust chaining. We have formally derived a parameterized trust opinion in the case of basic trust chains. However, it is possible that a subject receives more than one recommendation, or that the subject also has a simple trust opinion of the target. Recall trust aggregation from Definition 3. We first prove that a basic chained trust opinion can be aggregated with a simple trust opinion.

Later, we prove that more complicated trust opinions can also be aggregated with basic trust opinions. The notion that aggregation of these trust opinions is possible, is called *modularity*.

Running Example. Imagine that the subject c_i constructs a trust opinion about the target c_j based on his past interactions (z_s, z_f) with c_j . However, the teacher also provides a recommender c_k , with which the subject has an interaction history of (x_s, x_f) . The student c_k himself, give the recommendation (y_s, y_f) about c_j . From the Beta model, the subject can construct his (simple) trust opinion based on (z_s, z_f) . From Section 4, the subject can construct his (basic chained) trust opinion based on (x_s, x_f) and (y_s, y_f) . The subject wants to construct a trust opinion based on (x_s, x_f) , (y_s, y_f) and (z_s, z_f) . We prove the subject merely needs to aggregate both trust opinions.

Many trust models in the Beta family with trust chains (such as Subjective Logic) assert modularity. A priori, it is not obvious that the assertion of modularity is justified.

Running Example. Consider a situation in the classroom game where a student first constructs a trust opinion T_d directly from all his information. Then he tries an alternative approach and constructs simple trust opinions based on only parts of his information. These simple trust opinions he then aggregates into a trust opinion T_i . Assume that the subject c_i has a strongly positive opinion T about the target c_j , and a mildly positive opinion T' about the only recommender c_k . Assume further that the lying strategy of c_k is probabilistic and unrelated to the actual interactions of c_k with c_j and that $\lambda(n)$ is irrelevant (Lemma 2). Moreover assume, the recommender c_k gives a mildly negative opinion R about the target c_j .

Constructing his trust opinion T_d directly, the subject c_i concludes that, even though he expected the recommender to give honest recommendations more often than fake ones, this particular recommendation is nearly certainly fake. The subject expects the recommendation to be fake because he is quite certain that c_j has a high integrity (due to his trust opinion T). In other words c_i does not think it likely that c_k has more failed than successful interactions with c_j (which honesty of R would entail). Therefore, in the resulting trust opinion T_d , the recommendation R does not have a large impact.

If the subject constructs his trust opinion T_i modularly, then he aggregates T with a basic chained trust opinion T_c based on T' and R , without applying his own experience with c_j . If the subject does that, he will accept (in T_c , thus in T_i) that it is likely that the recommender provided an honest opinion about the target.

In conclusion, we may expect that T_i is more influenced by R than T_d .

The naive intuition that a modularly constructed opinion (T_i) differs from a directly constructed opinion (T_d), is proven incorrect in Proposition 4 and Theorem 3. First, we prove modularity between a simple trust opinion and a basic chained trust opinion:

Proposition 4. For all models in the Beta family with trust chains, the chained trust opinion $f_{R_C}(c|O_B^A=(x_s, x_f), S_C^B=(y_s, y_f), O_C^A=(z_s, z_f))$ is the aggregate of the simple trust opinion $f_{R_C}(c|O_C^A=(z_s, z_f))$ and the basic chained trust opinion $f_{R_C}(c|O_B^A=(x_s, x_f), S_C^B=(y_s, y_f))$.

Proof. We require Independence **I1** and Dependence **D1**.

$$\begin{aligned}
& f_{R_C}(c|O_B^A=(x_s, x_f), S_C^B=(y_s, y_f), O_C^A=(z_s, z_f)) \\
&= \frac{P(O_B^A=(x_s, x_f), S_C^B=(y_s, y_f), O_C^A=(z_s, z_f)|R_C=c) \times f_{R_C}(c)}{P(O_B^A=(x_s, x_f), S_C^B=(y_s, y_f), O_C^A=(z_s, z_f))} \\
&\stackrel{\text{I1}}{=} \frac{P(O_B^A=(x_s, x_f), S_C^B=(y_s, y_f)|R_C=c) \times P(O_C^A=(z_s, z_f)|R_C=c) \times f_{R_C}(c)}{P(O_B^A=(x_s, x_f), S_C^B=(y_s, y_f), O_C^A=(z_s, z_f))} \\
&\stackrel{\text{D1}}{\propto} \frac{P(O_B^A=(x_s, x_f), S_C^B=(y_s, y_f)|R_C=c) \times f_{R_C}(c) \times P(O_C^A=(z_s, z_f)|R_C=c) \times f_{R_C}(c)}{P(O_B^A=(x_s, x_f), S_C^B=(y_s, y_f)) \times P(O_C^A=(z_s, z_f))} \\
&= f_{R_C}(c|O_C^A=(z_s, z_f)) \times f_{R_C}(c|S_C^B=(y_s, y_f), O_B^A=(x_s, x_f)) \quad \square
\end{aligned}$$

Similar to Proposition 4, we can even prove that modularity holds for all trust opinions. Let φ be a collection of basic trust chains and potentially the interaction history between the target and the subject. In other words, for some n , let φ be given by:

$$]O_C^A=(z_s, z_f),]O_{B_1}^A=(x_s^1, x_f^1), S_{C_1}^{B_1}=(y_s^1, y_f^1), \dots, O_{B_n}^A=(x_s^n, x_f^n), S_{C_n}^{B_n}=(y_s^n, y_f^n).$$

Theorem 3 (Modularity of trust opinions). For all models in the Beta family with trust chains, the trust opinion $f_{R_C}(c|O_B^A=(x_s, x_f), S_C^B=(y_s, y_f), \varphi)$ is the aggregate of the trust opinion $f_{R_C}(c|\varphi)$ and the basic chained trust opinion $f_{R_C}(c|O_B^A=(x_s, x_f), S_C^B=(y_s, y_f))$.

Proof. The only step of the proof in Proposition 4 that cannot be replicated (with φ substituted for $O_C^A=(z_s, z_f)$) is the application of Independence **I1**. Thus:

$$\begin{aligned}
& P(O_B^A=(x_s, x_f), S_C^B=(y_s, y_f), \varphi|R_C=c) \\
&\stackrel{?}{=} P(O_B^A=(x_s, x_f), S_C^B=(y_s, y_f)|R_C=c) \times P(\varphi|R_C=c)
\end{aligned}$$

The proof obligation can be reduced (with Independencies **I1** and **I4**) to $P(\varphi|R_C=c, E_C=u, O_C^B=(w_s, w_f), R_B=b) = P(\varphi|R_C=c)$, which follows from Independencies **I2** and **I3**. A more detailed proof can be found in [10]. \square

From Theorem 3, we can conclude that the subjects can compute a trust opinion based on their own history with the target, as well as on recommendations of an arbitrary number of other users, provided that the subject can compute basic chained trust opinions for all recommendations. More generally, Theorem 3 allows us to generate the following structures $S(\lambda, \theta) = (P, O, g: P \rightarrow O, c_{\lambda, \theta}: P \times P \rightarrow O, a: O \times O \rightarrow O)$, where P is the set of interaction histories, O is the set of opinions, g is the function that maps interaction histories to simple trust opinions, $c_{\lambda, \theta}$ is the function that generates basic chained trust opinions (for entanglement λ and assignment of lying strategies to users θ), and a represents aggregation of trust opinions. Depending on the choice of the entanglement and the assignment of lying strategies, the structures $S(\lambda, \theta)$ (generally) differ.

6 Analysis of the Models

The results from the last sections allow us to study the conditions that all trust opinions for in all models in the Beta family with trust chains must adhere to. If an existing trust model violates these conditions, it is therefore not in the Beta family with trust chains. Which, in turn, means that these trust models either break an assumption of the Beta model (on which they are based), or its operator dealing with recommendations does not actually model trust chains according to Definition 5.

First, we point out that the work in Sections 4 and 5 captures all models in the Beta family with trust chains up to isomorphism:

Corollary 3. *Every model in the Beta family with trust chains is isomorphic to a structure $S(\lambda, \theta)$ for an entanglement λ and an assignment of lying strategies θ .*

Proof. The corollary is a direct consequence of Corollary 2 and Theorem 3. \square

A consequence of the corollary is that if a model is in the Beta family with trust chains, there is a formulation of the model where the entanglement and the assignment of lying strategies are explicitly provided. This entails that if a formulation of a model does not explicitly mention the assignment of lying strategies, it is not an appropriate formulation as it obfuscates the lying strategies.

Furthermore, we prove a restriction on the shape of chained trust opinions:

Theorem 4 (Chained trust opinions are not beta distributions). *A basic chained trust opinion in any model in the Beta family with trust chains is in general not a beta distribution. An exceptional case arises when the lying strategy and the entanglement are set in such a way that the resulting opinion of a basic trust chain based on a recommendation (y_s, y_f) is always (y_s, y_f) .*

Proof. Expression (2) from Proposition 1 can only represent a beta distribution $\beta(c; S + 1, F + 1)$, if it can be simplified to $h \cdot c^S(1 - c)^F$ for some $S, F \in \mathbb{N}$ and $h \in \mathbb{R}^+$. Rearranging the coefficients specifying how the constants depend on \mathbf{eq}_2 (henceforth let $p = \mathbf{eq}_2$), we therefore have to prove the following:

$$p \sum_{\ell} a_{\ell} c^{\ell} + (1 - p) \sum_m b_m c^m = c^S (1 - c)^F$$

Comparing coefficients ($n < S$ or $n > S + F$), we obtain that $pa_n + (1 - p)b_n$ has to equal 0. This linear equation can only hold for a specific value of p unless $a_S = b_S = 0$. As we prove the general case, we may not restrict p to a specific value, seeing that \mathbf{eq}_2 depends on x and y . Hence the sum reduces to

$$p \sum_{\ell=S}^{S+F} a_{\ell} c^{\ell} + (1 - p) \sum_{m=S}^{S+F} b_m c^m = c^S (1 - c)^F$$

The sum $\sum_{\ell=S}^{S+F} a_{\ell} c^{\ell}$ is proportional to a beta distribution, namely $\mathbf{eq}_1(y)$. The only beta distribution with maximal exponent $S + F$, and minimal exponent S is $\beta(c; S + 1, F + 1)$. That means that the recommendation, $\beta(c; S + 1, F + 1)$, equals the resulting opinion, meaning that the equation can only hold if we are in the exceptional case. \square

Therefore, any model that represents all its chained trust opinions as beta distributions, is not in the Beta family with trust chains.

Corollary 4. *CertainTrust [11] and TRAVOS [13] are not in the Beta family with trust chains.*

TRAVOS is an interesting case, as the authors set out to do essentially the same as is done in this paper. Similar to this paper, they treat the Beta model formally (using random variables for the integrity, for the outcomes and the recommendations) and study the relation between honest recommendations and fake recommendations. However, TRAVOS asserts that the result of a trust chain (in their case called reputation) is a beta distribution. A similar argument hold for Subjective Logic:

Corollary 5. *Subjective Logic [4] is not in the Beta family with trust chains.*

Proof. Subjective Logic is isomorphic to a model where all trust opinions are beta distributions. \square

Hence, Subjective Logic breaks an assumption of the Beta model (on which it is based), or its operator dealing with recommendations (called trust transitivity or trust propagation) does not actually model trust chaining. Both can be argued, since in Subjective Logic the trust transitivity operator is based on fuzzy logic, rather than distributions over integrity parameters, yet trust opinions and trust aggregation (called fusion) are based on the Beta model (i.e. based on distributions). The latter, would entail that the fraction of Subjective Logic dealing with trust chaining is not useful; the former entails that usefulness of trust chaining does not follow from the theory surrounding the Beta model.

It is possible to alter Subjective Logic to incorporate a trust chaining operator such that it is isomorphic to a structure $S(\theta, \chi)$. However, the property of Subjective Logic that a trust opinion equates to a belief triple will no longer hold. Rather, a trust opinion will equate a weighted sum of belief triples, e.g. $\sum_i k_i(b_i, d_i, u_i)$. The fusion (trust aggregation) of two trust opinions $\sum_i k_i(b_i, d_i, u_i)$ and $\sum_j k'_j(b'_j, d'_j, u'_j)$ will then be $\sum_{i,j} k_i \times k'_j((b_i, d_i, u_i) \oplus (b'_j, d'_j, u'_j))$, where \oplus denotes unaltered fusion of belief triples from Subjective Logic. There are several valid variations for transitive trust operators (trust chains), and Proposition 3 shows that the operator need not be complicated.

7 Conclusion

We study a family of models based on the Beta distributions: the Beta family with trust chains. The models in that family are very similar to the Beta model, but more expressive. In particular, they can express trust chaining.

An important property, proven for all models in the Beta family with trust chains, is that trust chaining operations are modular (Proposition 4 and Theorem 3). So complicated trust opinions can be constructed by aggregating simpler trust opinions. Many existing trust models have asserted this property, which we now proved.

Another commonly asserted property in models inspired by the Beta model, is that all trust opinions can be represented as beta distributions. This property is disproven for models in the Beta family with trust chains (Theorem 4). This

result implies in particular that Subjective Logic, TRAVOS and CertainTrust are not in the Beta family with trust chains (Corollaries 5 and 4).

We have proven that, up to isomorphism, every trust model in the Beta family with trust chains implicitly or explicitly makes assumptions about lying strategies and (except in special cases) about the entanglement (Corollary 3). Conversely, we have shown that, up to isomorphism, all trust models in the Beta family with trust chains can be constructed by selecting lying strategies and an entanglement (Corollary 3). Moreover, we have created a tool (Canephora) that calculates chained trust opinions, when instantiations of an entanglement and lying strategies are provided.

In the future we want to study the effectiveness of lying strategies using game theory. That would enable us to calculate the optimal lying strategies of recommenders, providing powerful models. Furthermore, we want to formally extend the Beta family with trust chains with additional operators, such conjunction [9]; in particular it is interesting to discover whether a modularity result still holds.

References

1. Billingsley, P.: Probability and measure. Wiley, 3 edn. (1995)
2. Bouckaert, R.: Bayesian belief networks and conditional independencies. Tech. Rep. RUU-CS-92-36, Utrecht University, The Netherlands (1992)
3. Gut, A.: Probability: A Graduate Course (Springer Texts in Statistics). Springer (2007)
4. Jøsang, A.: Artificial reasoning with subjective logic. In: 2nd Australian Workshop on Commonsense Reasoning (1997)
5. Jøsang, A., Ismail, R.: The beta reputation system. In: Proceedings of the 15th Bled Electronic Commerce Conference. vol. 160, pp. 324–337 (2002)
6. Jøsang, A., Marsh, S., Pope, S.: Exploring different types of trust propagation. In: Stølen, K., Winsborough, W., Martinelli, F., Massacci, F. (eds.) Trust Management, LNCS, vol. 3986, pp. 179–192. Springer, Berlin / Heidelberg (2006)
7. Krukow, K., Nielsen, M., Sassone, V.: Trust models in ubiquitous computing. Royal Society of London Philosophical Transactions Series A 366, 3781–3793 (Oct 2008)
8. Mui, L., Mohtashemi, M.: A computational model of trust and reputation. In: Proceedings of the 35th HICSS (2002)
9. Muller, T., Schweitzer, P.: A Formal Derivation of Composite Trust. In: Proceedings of the 5th International Symposium on Foundations & Practice of Security (2012), to appear
10. Muller, T., Schweitzer, P.: On beta models with trust chaining (technical report). Tech. rep., University of Luxembourg (2013), available at <http://satoss.uni.lu/members/tim/papers/TrustChainingTechReport.pdf>
11. Ries, S.: Certaintrust: a trust model for users and agents. In: Proceedings of the 2007 ACM symposium on Applied computing. pp. 1599–1604. SAC '07, ACM, New York, NY, USA (2007)
12. Staab, E., Fussenig, V., Engel, T.: Towards Trust-Based Acquisition of Unverifiable Information. In: Cooperative Information Agents XII, LNCS, vol. 5180, pp. 41–54. Springer Berlin / Heidelberg (2008)
13. Teacy, W., Patel, J., Jennings, N., Luck, M.: TRAVOS: Trust and Reputation in the Context of Inaccurate Information Sources. Autonomous Agents and Multi-Agent Systems 12, 183–198 (2006)