



City Research Online

City St George's, University of London

Citation: Sherwani, D. & Stumpf, S. (2014). Toward helping users in assessing the trustworthiness of user-generated reviews. Paper presented at the 28th International BCS Human Computer Interaction Conference (HCI 2014), 9-12 Sep 2014, Southport, UK. doi: 10.14236/ewic/hci2014.13

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/15373/>

Link to published version: <https://doi.org/10.14236/ewic/hci2014.13>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Toward Helping Users in Assessing the Trustworthiness of User-Generated Reviews

Dara Sherwani and Simone Stumpf
Centre for HCI Design
City University London
London EC1V 0HB UK
{dara.sherwani.1, simone.stumpf.1@city.ac.uk}

Abstract. With the growing number of systems that provide user-generated reviews come new forms of trust relationships. User trust in vendors can be mediated by trust in either reviews or reviewers, or possibly both. These new forms of trust relationships might be affected by barriers to trust such as uncertainty about offered services and anonymity of users, as well as biased reviews. While current work is dedicated towards investigating the influential factors on trust, this study undertakes a different approach. First, it investigates the influence of interface design on users' assessment of trustworthiness. Second, it explores the way that interface design can affect users' perception of trustworthy and untrustworthy reviews and reviewers, by signalling the influential factors on trust. Third, this study investigates the effect of users' prior beliefs in the form of dispositional trust in the assessment of trustworthiness. To do so, an exploratory study gathering quantitative and qualitative data was conducted with 16 participants who interacted with a high-fidelity prototype. Our results show that users' assessment of trustworthiness is influenced by interface elements that relate not only to the review, but also to the reviewer, implying that trust in reviews is mediated by trust in reviewers. Furthermore, users' dispositional trust appears to affect the perceived trustworthiness of reviewers, especially because of elements relate to the reviewers' background, which transfers onto the reviews. Our results have implications for researchers and designers to help users' assessment of the trustworthiness of reviews and reviewers.

Keywords: User-generated reviews, trust signals, dispositional trust, user study

1. INTRODUCTION

The growing number of systems that provide user-generated reviews has dramatically changed the way users make purchase decisions in an eCommerce context (Liu et al., 2008). These systems, such as TripAdvisor, Yelp, etc. are extensions of online communities which are consumption-related. Users interact with these systems to read reviews and to assess vendors' services prior to purchasing them. Hence, users' trust in vendors can be mediated by the trust in the reviews or reviewers, or possibly both (Figure 1). This leaves users to look for signals of trustworthiness when deciding to trust reviews and reviewers (Sasse & Kirlappos, 2011).




Figure 1: Trust signals and mediated trust relationship

The importance of trust signals on systems that provide user-generated reviews arises because previous work has suggested that user trust in reviews and reviewers can be affected by a number

of barriers (Ku et al., 2012). First, trust could be lowered because of uncertainty about the reviewed services in which users cannot assess different aspects of the reviewed services, such as room size and cleanliness. Second, trust is difficult to establish in anonymous reviewers where users of these systems are not connected directly and do not engage in face-to-face interaction. Third, the reviews might be biased.

Our work aims to explore the way interface design can help users to overcome the barriers to trust, by signalling the trustworthiness of reviews as well as reviewers. This approach is inspired by Riegelsberger et al. (2005), and it explores the way interface elements can signal trust factors. This, in turn, can help users in distinguishing between trustworthy (truthful) and untrustworthy (untruthful) reviews. We also take users' backgrounds into consideration in the assessment of trustworthiness, particularly dispositional trust, which is suggested by previous studies (e.g. Hsu, 2008) as an influential factor in user trust. Therefore, our work aims to address the following research questions:

RQ-1: What aspects of interface design affect user assessment of trustworthiness?

RQ-2: How does interface design signal trust factors and affect the assessment of trustworthiness?

RQ-3: How does a user's dispositional trust affect the assessment of trustworthiness?

In the remainder of this paper, we first introduce related work to this study. Second, the study set-up and data collection methods are explained in detail. Results are then reported and the theoretical and practical implications of the study are discussed. Lastly, the conclusion summarises the results and points the way to future work.

2. RELATED WORK

2.1 Designing for user trust

User trust has been an increasing concern in the Human Computer Interaction (HCI) field, since users' behaviour toward and usage of systems rely heavily on trust. According to Riegelsberger et al. (2005), trust is defined as the "willingness to be vulnerable based on positive expectations about the actions of others". Most of the work conducted in the HCI field about user trust is within the context of eCommerce (e.g. Egger, 2001). These studies are dedicated to developing guidelines for increasing the perceived trustworthiness of online vendors by focusing on interface elements that can be interpreted as signifiers of trustworthiness. However, these guidelines can be problematic if implemented by untrustworthy vendors. Therefore, research in the area of user trust has to be widened and has to investigate ways of enabling users in assessing the trustworthiness of vendors, rather than focusing mainly on increasing user trust (Riegelsberger et al., 2005). This approach forms the basis of our study; however, our study focuses on helping users in assessing the trustworthiness of reviews and reviewers, since users run the risk of making wrong decisions by believing untrustworthy (untruthful) reviews (Ku et al., 2012).

2.2 User trust in online reviews

Previous research has already suggested that three factors play an important role in user trust in reviews: review quality (Lee et al., 2008), source expertise (Kim et al., 2008) and source bias (Ku et al., 2012).

First, it has been argued that review quality plays a significant role. This factor has been broken down previously into four dimensions; understandability, sufficiency, relevance and reliability. Understandability refers to the easiness of understanding content. Sufficiency refers to the level of detail in the information. Relevance is defined as the extent of congruence between information needed by users and the actual stated information. And finally, reliability refers to 'dependability of information' (Lee et al., 2008).

Second, source expertise has been considered an influential factor in user trust in online reviews. Source expertise refers to the reviewer's perceived knowledge of a particular domain. It has been

argued that this factor is influential because users tend to trust expert sources in similar areas of interest. Source expertise can be reflected as summarised ratings of a reviewer's generated content by other users. This can help users to identify trustworthy reviews by reviewers without previous interaction (Kim et al., 2008).

Third, source bias has been suggested as an important factor in regard to user trust. Biased reviews, which are referred to as untruthful reviews, have features similar to real consumer reviews which makes them hard to detect by users. Biased reviews can be either positive, which is characterised by the reviewer praising a service without justification, or negative, which is indicated by the reviewer being too critical with lack of reasoning (Ku et al., 2012).

These factors capture facets of the review and reviewer, and should reduce the uncertainty about the offered services, both by providing helpful as well as accurate information (Liu et al., 2008).

Investigating the influential factors in user trust can offer significant implications for systems that provide user-generated reviews to increase user trust. However, this approach can be problematic because increasing user trust in reviews may lead to users trusting untrustworthy (untruthful) reviews. Thus, we suggest research in this area has to be broadened and point the way toward helping users in assessing the trustworthiness. One way to achieve this is by exploring the way that interface elements can signal trust factors and helping users distinguish between trustworthy and untrustworthy reviews and reviewers.

2.3 Dispositional trust

Dispositional trust represents a type of user belief and is defined as "the extent to which one displays a consistent tendency to be willing to depend on others across a broad spectrum of situations and persons" (Hsu, 2008). Dispositional trust indicates user perception of the reliability of others, and it can be shaped by a user's cultural background and personality type. McKnight et al. (2002) suggested that dispositional trust involves four dimensions: integrity, competence, benevolence and trusting stance. Integrity refers to the trustee's honesty in dealing with the trustor. Competence refers to the trustee's ability to meet the trustor's needs. Benevolence refers to the trustee's caring to act in accordance with the trustor's interests. Finally, trusting stance means that "regardless of what one believes about people's attributes, better outcomes result from dealing with people".

Dispositional trust has been shown as a significant factor in user trust in online vendors (Hsu, 2008). Thus, it has been suggested that users with high dispositional trust tend to have high trust in

vendors, in contrast to low dispositional trust users. However, none of the previous studies investigated the potential role of dispositional trust in the assessment of reviews and reviewers' trustworthiness. Thus, our work takes dispositional trust into consideration and explores its potential influence.

3. METHODS

3.1 Study setup

Sixteen participants (7 female, 9 male, mean age of 30) were involved in an exploratory study in which we gathered quantitative and qualitative data. All participants experienced the same procedure, with each session lasting approximately 1 hour. First, the facilitator explained the purpose of the study, and participants signed informed consent forms. Then, participants filled out a background questionnaire which captured their demographics as well as their dispositional trust.

Afterward, participants interacted with a computer-based prototype (Figure 2) showing them different reviews to help them investigate hotels for a holiday based on reviewers' contributions. Throughout the study, participants verbalised their thoughts and also provided explicit feedback on each review by filling in a trust factor questionnaire. This questionnaire asked participants to rate their perception of review quality, source expertise, source bias, as well as the helpfulness and accuracy of reviews.

3.2 Prototype and reviews

The prototype was similar to a standard hotel page on TripAdvisor (Figure 2). To avoid the confounding effects that explicit ratings, price and brand could

have on participants' trust, the prototype excluded the explicit ratings and price information, and stated a fictitious name for the hotel.

The prototype presented 8 positive reviews. We excluded negative reviews to better investigate and control the influence of trust factors. We chose to present a total of 8 reviews because we manipulated each of the three primary trust factors into two categories – high and low review quality, high and low source expertise, and high and low source bias – and, therefore, followed a factorial experiment design ($2 \times 2 \times 2$).

Constructing the reviews was done in 3 stages. First, we chose a set of real reviews from TripAdvisor. Second, the reviews were revised to be similar in length and review date. All reviews were set at about 10 lines long to avoid the influence of review length on review quality. Furthermore, the date of each review was amended to be no more than 1 month before the study's start date to prevent the influence of the review date on participants' responses toward the perception of accuracy (Lee et al., 2008). Last, we classified these reviews into either high or low categories for each trust factor. The reviews were presented in a different sequence to participants to avoid a confounding learning effect.

The classification of reviews into high or low trust factor categories was based on approaches suggested by previous studies (Lee et al., 2008; Kim et al., 2008; Ku et al., 2012). First, review quality was classified based on Lee et al.'s (2008) approach which suggests that review quality can be assessed based on 4 dimensions: understandability, sufficiency, relevance and reliability. It has been suggested that understandability refers to the easiness of understanding the content and it is mainly affected by the content structure, in which sentences are




Figure 2: Example of a review in the study prototype. Interface elements are labelled in red and correspond to the coding scheme in Table 1. (these labels did not appear in the prototype)

structured incorrectly, or by the occurrence of spelling errors. We assessed understandability simply based on the occurrence of spelling errors, in which a review was highly understandable when there were fewer than 5 spelling errors. Sufficiency was assessed based on how many of the following 6 information items were mentioned: service, sleep quality, cleanliness, location, food and room. These information items capture different aspects of hotel services and are suggested to be included in reviews by websites like TripAdvisor. A review was assessed as highly sufficient when at least 3 of information categories were mentioned in the review. The dimensions of relevance and reliability are aspects that are related to a user's perception and therefore could not be explicitly controlled prior to the study. Therefore, we assessed them by proxy based on sufficiency. Thus, a sufficient review which included detailed information about a hotel was considered as more relevant because it provided the user with potentially more related information. Meanwhile, a sufficient review was also considered as more reliable because the detailed information better justified the reviewer's opinion. Accordingly, a high quality review was understandable and included relevant and reliable information with sufficient details.

Second, source expertise was classified into low or high through 2 dimensions: reviewer experience and reviewer helpfulness, as suggested by Kim et al. (2008). Reviewer experience can be indicated by the total number of reviews posted by a reviewer and reviewer helpfulness can be indicated by the number of helpful votes given to a reviewer. This means that a reviewer acquires high source expertise not only by generating many reviews but other users of the community must also perceive the reviews as helpful. Thus, we classified a reviewer as having high expertise not only based on posting many reviews (e.g. 30) but the helpful votes must also have been at least 50% of the total reviews.

Last, classification of source bias was based on Ku et al. (2012) who suggested that source bias is indicated by the reviewer being too positive. We therefore classified source bias into high and low based on the positive emotions indicated in a review. We chose a list of words; awesome, amaze, best, fantastic, impress, love, great and surprise, and derivatives of these words to classify source bias, and we categorised a review as highly biased when the number of these words was greater than 5.

After the study was completed, we checked whether our classification of reviews into high and low categories of trust factors was successful. This was achieved by analysing the effect of the classification on the perceived factor ratings from the trust factors questionnaire. We treated each

trust factor as 2 conditions (independent nominal variable) which then allowed us to compare participants' perceived ratings (dependent variable). For example, we compared the mean *perceived* review quality ratings in the 2 constructed high/low review quality categories through a paired t-test. The results of each of these 3 tests showed that there were statistically significant differences in participants' ratings for each pair of conditions: (a) high and low review quality ($t=27.33$, $p<.01$), (b) high and low source expertise ($t=20.84$, $p<.01$) and (c) high and low source bias ($t= 23.29$, $p<.01$). These results support that our way of classification was appropriate.

3.3 Quantitative data analysis

Quantitative data were captured using two questionnaires. First, a background questionnaire was filled out by participants prior to their main task. This questionnaire captured participants' demographics and their dispositional trust. Our approach for measuring dispositional trust was directly taken from previous work (McKnight et al. 2002), which suggests that dispositional trust involves 4 dimensions: integrity, competence, benevolence and trusting stance. We also used the items from McKnight et al. (2002) to measure each dimension. For example, we captured participants' trusting stance through agreement with "*I usually trust people until they give me a reason not to trust them*". Each of the 4 dimensions was measured using a 5- point Likert scale, ranging from strongly disagree (1) to strongly agree (5). Dispositional trust was then calculated as the average value of the 4 dimensions. We captured data on dispositional trust to explore the way it can affect participants' trust in reviews based on interface elements. Our analysis involved exploring how dispositional trust influenced the interface elements attended to by participants.

Second, participants also filled out a trust factor questionnaire which captured their ratings of review quality, source expertise, source bias and perceived helpfulness and accuracy on a 5-point Likert scale, ranging from strongly disagree (1) to strongly agree (5). These scales were based on previous work on trust models (Kim et al., 2008; Lee et al., 2008; Ku et al., 2012; Liu et al., 2008). These data were used in establishing the relationship between the interface elements participants mentioned and trust factors.

3.4 Qualitative data analysis

Qualitative data were captured from participants' verbalisations as they thought aloud when reading the reviews. These data were analysed using a coding scheme to investigate the role of the interface elements in signalling the trustworthiness

of reviews and reviewers. The coding scheme was developed based on the design of our computer-based prototype (shown in Figure 2). Although the prototype presented 10 interface elements alongside the reviews, participants mentioned only 6 of the existing elements (which are labelled in Figure 2), and, thus, the coding scheme is restricted to codes that referred to these 6 elements. Also, participants' verbalisation included 2 new elements that did not exist in our prototype: *number of people who found the review unhelpful* and *photographic evidence* alongside to reviews, and we added these to the coding scheme. The resulting coding scheme is shown in Table 1.

For our analysis, we applied the codes to 'sentences'. A sentence is a continuous verbalisation by a participant until he or she stops. This way of breaking up the transcript into units of analysis resulted in 503 sentences across all participants. We applied the codes to a sentence when participants' verbalisations included a direct reference to interface elements, meaning that we followed an *in-vivo* coding approach (Lazar et al., 2010). If the same code occurred more than once in a sentence, only the first occurrence was coded.

Table 1: Coding Scheme

Label	Code	Example
1	City & country	The person is from the <i>same country</i> where the hotel is
2	Membership level	How important it's for me if he is a <i>top contributor</i>
3	Number of reviews	This guy has <i>41 reviews</i>
4	Number of cities	He travelled to <i>25 cities</i>
5	Number of helpful votes	He is a genuine reviewer because he has <i>18 helpful votes</i>
6	Number of people who found the review helpful	I'm more impressed by <i>5 people who found the review helpful</i>
New	Number of people who found the review unhelpful	If I just see <i>how many people vote dislike or maybe unhelpful</i>
New	Photographic evidence	If I'd be able to see <i>photos with this review</i> , I'd feel I can depend on it

4. RESULTS

4.1 Influence of interface elements on assessing trustworthiness

Interface elements can signal to users how much they should trust an online vendor (Egger, 2001);

however, there is little known about signals to trust a review or reviewer. A better understanding of this could lead to improved interface design which helps users, first, to assess the strength of trustworthiness and, second, whether to place trust either in the review or the reviewer, or possibly both. In this study, we investigated these two issues by analysing participants' think-aloud as to what interface elements were emphasised in relation to assessing trustworthiness.

4.1.1 Existing interface elements

Participants seemed to place more faith in some interface elements than in others to assess trustworthiness. Figure 3 shows how often interface elements were mentioned by participants as they discussed trust in each of the reviews. We note that four existing interface elements were mentioned the most by participants: *number of people who found the review helpful*, the *number of reviews* a reviewer has contributed, the *reviewer city & country* and the *number of helpful votes* a reviewer has gained. These elements seemed to give clear trust signals to the participants. However, some of the existing elements were not considered as frequently as trust signals. These elements were *membership level* and *number of cities* that have associated reviews (mentioned 7 and 6 times, respectively). A factor in this lack of focus seemed to be that participants struggled in understanding how these levels were earned and, thus, the signal was clouded: "*Senior contributor, I wonder what that means*" (P 12)




Figure 3: Interface elements mentioned by participants in order of their overall frequency: existing elements (dark) and new elements (light)

Some participants also paid more attention to other signals based on the reviewer's background instead of travelling experience. For example: "*I'm more inclined to think about the amount of reviews written by this person, I guess that makes it up for me to have an idea about him/her, I'm not sure if I care about how many places he's been in.*" (P 11)

Trust in a review could be eroded if a user does not trust the reviewer, or, more positively, high trust in a reviewer could be transferred to a review. Hence, we also investigated what types of interface elements play a role in users' ability to assess trustworthiness of the review and/or the reviewer. Overall, there seemed to be two broad types of interface elements. These types were: review-related and reviewer-related. The former type informed dimensions that related directly to the reviews, such as the *number of people who found the review helpful*. The second type revealed dimensions that related to the source of reviews, such as the background (i.e. city & country) and the reviewing behaviour (i.e. number of posted reviews and helpfulness of posted reviews).

The interface element most frequently mentioned by participants to assess trustworthiness was review-related: *number of people who found the review helpful* was mentioned 46 times and considered by all participants. When using this element, participants relied extensively on the virtual communities' opinions to help them establish trustworthiness. This signal seemed to influence trust directly, leading to higher trust when there were more people who found the review helpful: *"I can depend on this...because 3 people have given helpful votes to this review"* (P6). Similarly, participants' trust seemed to decrease when reviews did not have any helpful votes by other users. For example: *"I can't trust this review...no one found it helpful"* (P6).

Participants also made extensive use of reviewer-related interface elements to assess trustworthiness of a review indirectly. In our study, participants frequently mentioned *number of reviews* a reviewer provided (34 times by 15 out of 16 participants), *reviewer city & country* (34 times by 13 participants) and *number of helpful votes* (24 times by 12 participants). In these cases, the level of trust placed seemed to transfer from reviewer to review: *"When I was reading this review I thought this is better than the one who has 2 reviews only, so this guy is alright and he is giving many reviews"* (P1) and *"Other people found him helpful because of the 18 helpful votes and that would make me trust it more"* (P3).

This transfer of trust also appeared negatively, when a reviewer's background information aroused suspicion. For example, participants often wondered about the accuracy of reviews when reviewers' *city & country* was the same as the hotel: *"I'm suspicious of this one...he or she is from Malaysia and it makes me wonder whether it's a real review or not...I'm not depending on it"* (P5). Intriguingly, this element which gave details of the reviewer also was used by participants in a different way, to assess their similarity to the

reviewer: *"It could be good for me the reviewer is from the UK like me...I can depend on it"* (P15).

Whereas most of the interface elements that are related to the review or reviewer give a direct, often numerical, trust signal which can be easily interpreted by users, this element is highly subjective and relies on the background of the user instead of contributions to the virtual community.

4.1.2 New interface elements

Participants also voiced the need for two new review-related signals embedded in the interface that could boost users' ability to assess trustworthiness. The first of these relates to *number of people who found the review helpful*. While this element already existed in the interface, 7 out of 16 participants also expressed the desire to see how many people found this review *unhelpful*, for 8 times. It seems that these participants intuitively understood that this signal captured only one aspect of a positive and negative dimension, and wanted a more complete and balanced perspective. Interestingly, participants did not expect this balance of perspective to be extended to the *number of helpful votes* a reviewer has gained, even though it also expresses positive and negative aspects, and it transmits a signal about helpfulness. Possibly, the overall helpfulness of a reviewer is more difficult to assess than the helpfulness of a review. Furthermore, participants' responses also revealed that photographic evidence alongside reviews may allow participants to assess the trustworthiness of a review: *"I can't judge accuracy of information completely...if I'd be able to see photos of the hotel with this review, I'd feel I can depend on this review"* (P13). Five of our participants' mentioned this 5 times, possibly because photos were seen as less subjective and also less prone to falsification. Similar to photos used for virtual re-embedding to add social cues (Riegelsberger & Sasse, 2003), photos in this instance could function to provide objective accuracy signals.

4.1.3. Summary

Taken together, these results indicate that not all interface elements are used equally. First, our study showed that interface elements varied in their strength, and the most used element by participants to assess trustworthiness related directly to the review. While reviewer-related signals also influenced the trustworthiness of a review, these signals appeared to be weaker, as participants used individual interface elements related to reviewers less frequently, even though there are more of them in the interface than review-related signals. This may be the root of participants asking for new review-related signals, such as *number of people who found the review unhelpful* and *photographic evidence* alongside reviews, that could give them further signals about

trustworthiness of a review and would allow them to judge related trust factors.

4.2 Relationship between Interface elements and trust factors

Previous work has identified a variety of factors that matter in users placing trust in reviews (Lee et al., 2008; Kim et al., 2008; Ku et al., 2012; Liu et al., 2008). Factors such as perceived review quality, source expertise and source bias as well as perceived helpfulness and accuracy are held to be important in existing trust models. We wondered how interface elements shaped the participants' perception of these trust factors, especially given that trust between vendor and user is mediated through reviews and reviewers. To better understand this relationship between interface elements and trust factors, we returned to the verbal think-aloud of participants. This time, however, we investigated the interface elements participants mentioned related to their perceived ratings from the trust factors questionnaire. To do so, we calculated the median of the perceived ratings and then divided the trust factor responses into high (above the median) or low (below the median). We then investigated the relationship between the interface elements that were mentioned and these trust factor groupings.

4.2.1 Review-related trust factors

Figure 4 shows the elements that participants mentioned when they were assessing reviews with high quality (above the median of 3.5). One interface element played a particularly important role in allowing participants to determine the perception of high review quality. We observed that participants frequently mentioned the element *number of people who found the review helpful*. 14 out of 16 participants referred extensively (35 times) to this element when they discussed high review quality, for example: "3 others like it, it has everything I care about" (P10)




Figure 4: Frequency of interface elements mentioned by participants in assessing high-quality reviews

In these instances, participants used the attitude of the virtual community to assess review quality, instead of dimensions such as understandability, sufficiency, relevance and reliability which have been previously suggested as the basis for users assessing review quality (Lee et al., 2008).

Participants supplemented their assessment of high review quality with interface elements related to reviewers. The next three most important elements that helped participants evaluate high quality reviews were *number of reviews* (10 participants, mentioned 21 times), *number of helpful votes* (10 participants, mentioned 16 times) and reviewer *city & country* (4 participants, mentioned 4 times). Here, in contrast to their overall frequency of use, the *number of help votes* a reviewer received outweighed *city & country of origin*. We found the same dependence on others' judgements when participants tried to assess low quality reviews. In this case, participants mentioned interface elements that involve votes about the review or reviewer 27 times in assessing low quality reviews: "It doesn't provide information and no one found it helpful" (P6). This provides further evidence that virtual communities' opinions are very important in trust and also that review quality is influenced by the trustworthiness of both the review and the reviewer.

There may be a need to provide further signals for low-quality reviews. Seven participants mentioned *number of people who found the review unhelpful* on 8 occasions in relation to low review quality. For example: "This one is bad, I wonder how many unhelpful votes it could get" (P12). None of our participants mentioned the need for this signal when they perceived high review quality, implying that this element could have a direct effect on assessing low-quality reviews.

Similar interface elements were mentioned by participants when assessing helpfulness (Figure 5). Participants referred to the elements *number of people who found the review helpful*, *number of reviews* and *number of helpful votes* for 18, 14 and 12 times, respectively, when they rated perceived helpfulness above the median of 3. This points to this trust factor being an important facet of reviews which are assessed in the same way as review quality through signals in the interface.




Figure 5: Frequency of interface elements mentioned by participants in assessing helpfulness of reviews

4.2.2 Reviewer-related trust factors

Source expertise and source bias are trust factors that are obviously related to the reviewer. Not surprisingly, we found that interface elements that exposed signals about the reviewer were used the most by participants to assess these factors. However, even within this range of interface

elements, some proved to be more helpful than others in identifying these two trust factors.

Source expertise relates to the reviewers' perceived knowledge in a particular domain and, as previously mentioned, is often assumed to be associated with summarised ratings of a reviewer's generated content by other users (Kim et al., 2008). Out of the four interface elements that could have given clues about source expertise, *number of reviews* and *number of helpful votes* for a reviewer seemed to particularly influence the perception of this trust factor. We noticed that when participants rated source expertise high in a reviewer (above the median of 3), they mentioned *number of reviews* 20 times and *number of helpful votes* 18 times (Figure 6). For example: "So I think this guy knows what he's talking about and he has 47 reviews and 29 helpful votes" (P2)




Figure 6: Frequency of interface elements mentioned for assessing high source expertise

Participants also seemed to consider the same signals for their perception of low source expertise but to a lesser extent. However, we observed that 5 participants mentioned the need for *photographic evidence* when they rated source expertise as low. This could be because participants perceived low-expertise reviewers as inexperienced and unhelpful, in comparison to high-expertise reviewers. This, in turn, could have led to uncertainty about reviewers and, therefore, a search for further signals to assess trustworthiness.

It has been proposed that users assess source bias based on reviews that praise a service without justification (Ku et al., 2012) and, therefore, on clues within the content of a review. Instead, we found that the interface element that participants most mentioned in relation to their high perception of source bias (above the median of 3) was *reviewer city & country*, which was mentioned 30 times by 12 out of 16 participants. A possible explanation for this is that there were no explicit signals about the level of source bias, and, hence, participants resorted to this interface element to give them clues about the reviewer's background. This is further supported by evidence that participants used the same interface element as a signal to assess review accuracy. We also observed that participants referred to *reviewer city & country*, which was mentioned 17 times when the rating of the review-related trust factor perceived accuracy was low (less than a median of 3).

Participants' reliance on *reviewer city & country* to determine source bias and review accuracy could

be because there were no other signals that helped participants in this matter. In fact, 9 of our 16 participants specifically told us that they struggled when source bias was perceived as high: "Nothing there tells me these people are real customers" (P4)

4.2.3. Summary

Overall, our results suggest that users employ interface elements as signals that relate to trust factors; however, the way that these signals are used is not straightforward. It appears that virtual community judgements play a more important role than previously considered, especially in the assessment of review quality and helpfulness. Moreover, there is a disparity in the strength of these signals in assessing trust factors, which becomes especially important if there is a variety of signals available which may compete with one another. Last, there seems to be a need for further signals that can help users assess the trust factors, particularly in the case of direct signals for source bias and accuracy of reviews.

4.3 Effects of dispositional trust on using interface elements as trustworthiness signals

User belief can affect how trustworthy a person is considered from the outset by a user in real life, and this disposition to trust also applies online to the relationship between users and vendors (Hsu, 2008). Low dispositional trust could lead to users not trusting reviewers, regardless of signals that allow them to ensure that reviewers have adequate experience. Currently, there is a lack of insight into the role of dispositional trust in assessing the trustworthiness of reviews and reviewers, and this may hamper the design of interfaces that could assist users who need help the most. We therefore explored how interface elements were used to assess trustworthiness by participants in our study, given their dispositional trust measured in the background questionnaire.

We first of all divided participants into either having high or low dispositional trust. The participants in our study had a high spread of dispositional trust scores, ranging from 1.5 (minimum) to 4 (maximum), a mean of 2.56 and a standard deviation of 0.72 (Figure 7). We used the median value of 2.5 as the pivot point for dividing the participants into 2 groups: low dispositional, with a score of 2.5 and below were classed as low dispositional trust, and the rest as high dispositional trust. Ideally, the groups would have been divided based on a score of 3 since it represents the exact middle score, however, our participants' dispositional trust measures were not normally distributed so we used the median to divide participants into 2 groups. As a result, the groups were equal, each consisting of 8 participants.




Figure 7: Participants' dispositional trust measures, ranging from 1.5 (lowest) to 4 (highest), median is 2.5

We then investigated whether there was a pattern in how frequently interface elements were mentioned by the respective groups. Figure 8 shows the three interface elements which were used differently by participants with high or low dispositional trust; the remaining elements were used similarly by both groups, and we will not discuss them in detail.

Participants with low dispositional trust paid more attention to reviewer background details. They mentioned reviewer *city & country* on 22 occasions, almost twice as many times as participants with high dispositional trust. On almost all these occasions, participants with low dispositional trust used this interface element to judge the reviewer as less trustworthy. It appears that low dispositional trust led participants to focus on negative aspects of the reviewer, confirming their suspicions. Hence, interface designs may inadvertently provide signals which are taken by users with low dispositional trust to trust reviewers even less.




Figure 8: Interface elements that differed between participants with high dispositional trust (dark) and low dispositional trust (light)

Signals that directly relate to the reviewers' expertise were used extensively by participants with high dispositional trust. Participants with high dispositional trust mentioned *number of reviews* and *number of helpful votes* 21 and 19 times, respectively, whilst participants with low dispositional trust assessed the same elements only 13 and 5 times, respectively. It appears that these participants were more swayed by what other users thought of reviewers, which, by definition, having a disposition to trust others. We see again that user belief can reinforce how signals in the interface are interpreted: *"You can tell he/she has done reviews before and he/she has 47 reviews and 29 helpful votes ...so yes I can depend on it"* (P2 high dispositional trust participant) and *"It looks like the reviewer is a professional or experienced*

reviewer based on the 41 reviews and helpful votes but I don't know if just looking at these kind of push me to go for it" (P15 low dispositional trust participant)

Our results suggest that participants' dispositional trust affected the establishment of trust in the reviewer, and then this trust transferred onto the assessment of the review. Low dispositional trust seemed to be associated with a more critical stance toward the interpretation of trust signals, whereas trust appeared to be boosted by signals in the interface that worked only because these users already had a disposition to trust others.

5. DISCUSSION

Our study has three main implications for research. First, our study builds on previous work (Riegelsberger et al., 2005) that explored users' assessment of trustworthiness of online vendors through trust signals, and we extended this work to user-generated reviews. Systems that provide user-generated reviews can result in new forms of trust relationships, where reviews and reviewers mediate the trust between user and vendor. Therefore, there seems to be an obvious need to explore these new forms of trust relationships and better understand users' trustworthiness assessment.

Second, we contribute to the existing knowledge about trust factors in online reviews. Previous work has focussed extensively on exploring what factors influence trust in online reviews (Lee et al., 2008; Kim et al., 2008; Ku et al., 2012), and our study extended this by exploring the way interface elements are used as signals for these factors. Future research could investigate this relationship between signals and trust factors further, especially as they relate to review and reviewer.

Third, our work also added that dispositional trust seems to influence the use of trust signals, in which low dispositional trust users seemed to be more critical toward establishing trust, more so than high dispositional trust users. Previous work (Hsu, 2008; McKnight et al., 2002) has paid scant attention to the role of dispositional trust in user-generated reviews, and this warrants further investigation.

Two specific questions arise from our study that should be addressed by future studies. First, previous work (Lee et al., 2008) suggested that negative reviews can be very influential in user trust, possibly more so than positive reviews. This in turn can affect the use of interface elements for assessing trustworthiness. Second, our study has provided qualitative data to deeper understand the relationship of interface elements and trust. Thus, a larger-scale investigation could provide quantitative evidence to validate the effect of trust signals.

Our results also have practical implications for designers of systems that provide user-generated reviews. First, signals that relate to the online community's opinions about reviews and reviewers should be transparent to help users assess trustworthiness, as our results revealed that participants relied on the community's opinions to assess the quality and helpfulness of reviews and the reviewers' expertise. Since this seems to matter the most in establishing trust, these kinds of interface elements play an important role in system design. However, the online community's positive opinions may not be enough to establish trust. Thus, we recommend capturing a more balanced perspective by signalling both positive and negative opinions.

Participants in our study told us that they would appreciate objective evidence to verify the information given. This evidence could come in the form of photographs, but it is less obvious how other service-oriented information items (e.g. service, sleep quality, etc.) that may be important in purchasing decisions could be represented in this way. Thus, more creative ways are needed to help users assess review quality and accuracy.

Last, our results showed that participants resorted to reviewers' background information, i.e. reviewer *city & country*, but struggled to determine source bias. Designers could help users by providing more direct signals for the assessment of reviews' untruthfulness, possibly by providing signals that prove the purchase transactions which could better help the assessment of the reviews' truthfulness.

6. CONCLUSION

We carried out an empirical study to explore the way interface elements are employed by users in the assessment of user-generated reviews' trustworthiness. We first found that participants used both review and reviewer-related signals, implying that trust in the review is also mediated by trust in the reviewer. Second, participants wanted signals of both positive and negative community's opinions as well as direct signals of source bias and accuracy of reviews. Third, dispositional trust appeared to influence the way users employed interface elements to assess trustworthiness, with a strong focus on reviewer-related information by participants with low dispositional trust.

Our work has identified important areas for future research and practical implications for the design of systems that incorporate user-generated reviews. This study presents the first steps toward understanding how interface signals are used in assessing trustworthiness — ultimately putting the power of decision-making back into the hands of the user.

7. REFERENCES

- Egger, F.N. (2001). Affective Design of E-Commerce User Interfaces: How to maximise perceived trustworthiness. CAHD: Conference on Affective Human Factors Design. ACM Press, New York, NY, pp. 317–324.
- Hsu, C.-J. (2008). Dominant Factors for Online Trust. International Conference on Cyberworlds, pp. 165–172. IEEE
- Kim, Y. et al. (2008). Building a web of trust without explicit trust ratings. 24th IEEE International Conference on Data Engineering Workshop, pp. 531–536. IEEE
- Ku, Y., Li, H., & Wu, F. (2012). Evaluation of Trustworthiness of Reviews in the Online Travel Community. IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications, pp. 1417–1422. IEEE
- Lazar, J. et al. (2010) Research Methods in Human-Computer Interaction. John Wiley & Sons Ltd. UK
- Lee, J., Park, D. & Han, I. (2008). The effect of negative online consumer reviews on product attitude: An information processing view. Electronic Commerce Research and Applications, 7, 3, 341–352.
- Liu, Y., Huang, X., An, A. & Yu, X. (2008). Modeling and Predicting the Helpfulness of Online Reviews. In: IEEE 8th International Conference on Data Mining, pp. 443–452. IEEE
- McKnight, D.H., Choudhury, V., & Kacmar, C. (2002). "Developing and Validating Trust Measures for E-Commerce: An Integrative Typology," Information Systems Research, Vol.13, No.3, pp. 334-359.
- Riegelsberger, J., Sasse, M. A., & McCarthy, J. D. (2005). The mechanics of trust: A framework for research and design. International Journal of Human-Computer Studies, 62(3), pp. 381–422.
- Riegelsberger, J. & Sasse, M.A. (2003). Shiny happy people building trust? Photos on e-Commerce websites and consumer trust. In: CHI (pp. 121-128)
- Sasse, M. A. & Kirlappos, I. (2011). Familiarity Breeds Con-victims: Why We Need More Effective Trust Signaling. In Wakeman, I., Gudes, E., Jensen, C. D., Crampton, J. (Eds.). *Trust Management V*. Springer-Verlag New York Inc