Self-Disclosure and Perceived Trustworthiness of Airbnb Host Profiles

Xiao Ma
Jacobs Institute, Cornell Tech
New York, NY, USA
xiao@jacobs.cornell.edu

Jeffrey T. Hancock
Stanford University
Stanford, CA, USA
jeff.hancock@stanford.edu

Kenneth Lim Mingjie
Cornell University
Ithaca, NY, USA
kl545@cornell.edu

Mor Naaman
Jacobs Institute, Cornell Tech
New York, NY, USA
mor@jacobs.cornell.edu

ABSTRACT
Online peer-to-peer platforms like Airbnb allow hosts to list a property (e.g. a house, or a room) for short-term rentals. In this work, we examine how hosts describe themselves on their Airbnb profile pages. We use a mixed-methods study to develop a categorization of the topics that hosts self-disclose in their profile descriptions, and show that these topics differ depending on the type of guest engagement expected. We also examine the perceived trustworthiness of profiles using topic-coded profiles from 1,200 hosts, showing that longer self-descriptions are perceived to be more trustworthy. Further, we show that there are common strategies (a mix of topics) hosts use in self-disclosure, and that these strategies cause differences in perceived trustworthiness scores. Finally, we show that the perceived trustworthiness score is a significant predictor of host choice—especially for shorter profiles that show more variation. The results are consistent with uncertainty reduction theory, reflect on the assertions of signaling theory, and have important design implications for sharing economy platforms, especially those facilitating online-to-offline social exchange.

Author Keywords
Airbnb; self-disclosure; trustworthiness; sharing economy; social exchange.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION
Airbnb is an online lodging marketplace for short-term peer-to-peer rentals, facilitating monetary and social exchange between individuals [29]. On Airbnb, *hosts* can list places (e.g. rooms, apartments, houses, or even boats and castles) for *guests* to rent. The guest is often a temporary visitor, and is not acquainted with the host beyond Airbnb. At time of writing, Airbnb reports two million listings, and 60 million guests on the platform [1].

The main utility of Airbnb—identifying potential lodging resources offered by unknown individuals—comes with risks that affect both guests and hosts who wish to participate in the exchange. A potential host may worry about guests damaging their property. A potential guest may fret about their physical safety, the truthfulness of the quality of the property being advertised, or whether the host would be kind enough to provide assistance in exigencies [17]. Establishing guest-host trust helps manage such uncertainties and risks—making trust a crucial factor for the success of such social exchange sites.

There are several ways that Airbnb designs for trust. Airbnb has an assurance policy and a reputation system in place, in addition to making information about the host and property readily available before booking. On Airbnb, each host has a profile page that includes photos, a text-based self-description, social media verification status, and reviews (if any) from other Airbnb users who have stayed with the host. These profiles contribute to a guest's decision making process [36], and help establish perceived trustworthiness [17]. In this work, we focus on host profiles, especially the text-based self-description and its role in establishing the perceived trustworthiness of hosts in the eyes of potential guests.

Emerging literature is examining how people assess trustworthiness through self-disclosures made in online profiles. The *Profile as Promise* [15] conceptual framework, for example, incorporates the risks and rewards associated with assessing signals in a profile for whether said profile's promises can be trusted. Researchers had examined how individuals produce and assess trustworthiness signals in online dating profiles [41] and in online résumés [22]. However, we still know very little about what people self-disclose, and how that information is evaluated for trustworthiness in the context of sharing economy platforms such as Airbnb.

CSCW '17, February 25-March 01, 2017, Portland, OR, USA

DOI: http://dx.doi.org/10.1145/2998181.2998269

Given the importance of profiles, in this work we aim to advance our understanding of the type of content Airbnb hosts self-disclose in their profiles, and to determine the impact of these disclosures on perceived trustworthiness and host choice. We build on the Profile as Promise framework [15], drawing on theories from economics and communication to predict what kinds of information hosts will disclose in their profiles, and what kinds of disclosure will enhance trust. In particular, we apply uncertainty reduction theory (URT) [5] to predict that both quantity and diversity of information increases the perception of trust. We also draw on signaling theory [14, 39] to predict that specific kinds of information can signal trustworthiness in a profile.

Specifically, we use a mixed-methods approach, with qualitative analysis, large-scale annotation, and an online experiment to examine the text-based self-descriptions of Airbnb host profiles. We qualitatively develop a categorization scheme that characterizes the primary self-disclosure topics in these profiles. We then quantitatively show how predictions from URT and signaling theory apply to this case, revealing that an increase in the quantity of content and the inclusion of specific topics can enhance perceptions of trustworthiness. Finally, we use an online experiment to show that the perceived trustworthiness of profiles is a significant predictor of host choice. Our results have practical design implications for platforms facilitating social exchange in the sharing economy.

BACKGROUND

User profiles are an important part of many online systems, and serve a variety of functions [42]. In social networking sites, profiles provide an identity for the user that persists over time and the myriad of interactions on the site [6]. In online dating sites, profiles provide self-disclosure that attracts the interest of other users, while limiting the risks associated with outright deception [20, 41].

In each of these contexts—especially in services that lead to offline interactions—a key function of the profile is to establish perceived trustworthiness [17, 23]. Here, we distinguish between *trustworthiness* and *trust*. Trustworthiness is an attribute of a trustee [24, 28], while trust is exhibited by a trustor (e.g. a decision to take risks in an economic game [4, 11]). As our focus is on the host, the trustee, we focus on perceived trustworthiness as an attribute of the host.

As noted above, one approach to understanding how profiles are used to establish trustworthiness is the Profile as Promise framework [15], which uses the notion of a promise to characterize the function of a profile. In this view, the profile is a psychological contract between the profile holder, and the viewer that future interactions (e.g. with a date, a car driver, or an Airbnb host) will take place with someone who does not differ fundamentally from the person represented in the profile. The notion of a promise has been successfully applied to various contexts that require good faith to operate, including online dating [15] and job hunting [37].

The Profile as Promise perspective argues that the content and characteristics of the disclosures, or promises, made in user profiles should be diagnostic of trustworthiness perceptions. Within this framework we draw on theories from communication, and from economics, to form specific expectations about disclosures and their perceptions. For example, communication scholars have used URT [5] to show that strangers are concerned with increasing the predictability about the behavior of both themselves and others in the interaction that occurs when they first meet. URT has been used in research on dating sites to explain how much information may be shared in profiles [20], using self-disclosure as the process of making the self known to others [13]. According to this approach, people should disclose as much information in their host profiles as they feel comfortable sharing—a directive that presents various challenges, including the risk of over-exposure if the profile is public [3, 20]. Nevertheless, the more information disclosed in a profile, the more likely it is perceived as trustworthy.

According to the Profile as Promise framework, one way to understand what kinds of information people disclose in the kind of static and asymmetric context of online profiles is signaling theory [15], which considers the relationship between explicitly stated signals and the underlying qualities they are likely to represent [14, 39]. Spence's original application of signaling theory [39] explored how potential employees in a labor market tried to convey attributes that cannot be observed directly, such as reliability or goodness-of-fit with a company's culture. Signaling theory was also concerned with how such signals can be assessed, for example, by the employer. Some signals, called conventional signals, are relatively easy to fake; such as asserting that one is a reliable and hard worker, when in fact one is not. Other signals, called assessment signals, are more difficult to fake; such as claiming that one has a degree from a prestigious institution. Gambetta used signaling theory to study how taxi drivers [19] and criminals [18] form impressions and assess the trustworthiness of other parties based on cues, especially in fleeting initial interactions where high uncertainty is involved. Here we examine how Airbnb hosts signal their trustworthiness in their profile through self-disclosure and how these signals are perceived.

Research Questions

In our examination of Airbnb profiles, our first objective, according to the Profile as Promise framework, is to determine what kinds of information hosts provide in their profiles to reduce uncertainty and signal trustworthiness.

RQ 1. *What kinds of information do hosts self-disclose to signal their trustworthiness?*

The second question is concerned with how the information hosts disclose in their Airbnb profiles translates into perceived trustworthiness by guests. On Airbnb, like other peer-based sharing economy services, signals of trustworthiness are particularly important as trust is critical for social exchange [9, 29]. Trust on Airbnb and other online marketplaces is tied to ratings and reputation. However, ratings on Airbnb tend to not be informative as they are skewed high, and there is initial evidence that the number of reviews received is predictive of room sales even when controlling for scores [30, 43]. Research suggests that profile information matters on Airbnb: in one study, profile images were linked to the perceived trustworthiness of hosts and higher prices [17]. At the same time, the study showed that online review scores had no effect on the listing price, although profile text was not considered. We therefore ask:

RQ 2. *What is the effect of different types of self-disclosure on perceived trustworthiness?*

Note that it is not immediately clear that trustworthiness maps directly to the choice of host. Choice can clearly be influenced by other factors, such as assurance [8]. There is initial data-driven evidence that visual-based trustworthiness impacts choice [17], even when reputation scores are manipulated to increase their variance.

Given the difficulty in establishing the causal link between profile disclosures and guest decision-making, we conduct an experiment that isolates the trustworthiness of a profile's disclosures from other external factors—such as reputation indicators—and manipulates the effect of low versus high trustworthiness profiles on a decision-making task. In this experiment we focus on addressing the following research question:

RQ 3. *Do profile-based perceptions of trustworthiness predict choice of host on Airbnb?*

STUDY 1—HOW DO HOSTS SELF-DISCLOSE?

The primary goal of Study 1 is to uncover what Airbnb hosts self-disclose in the text-based self-descriptions in the profile. To the best of our knowledge, there is no established coding scheme for self-disclosure in this particular context. For this reason, Study 1 uses qualitative methods to develop, validate and apply a coding scheme for self-disclosure in Airbnb host profiles. We accomplished this task using a two-phase approach. In phase one, we developed and validated a coding scheme for topics of self-disclosure by qualitatively analyzing Airbnb profiles and using an inductive and iterative approach to identify categories. In phase two, we applied the coding scheme to a large set of host profiles, and examined patterns of self-disclosure on Airbnb.

Phase 1— Developing and Validating Coding Scheme

Step 1: Development

To create the self-disclosure coding scheme for Airbnb, we used an iterative, inductive analysis for content-topic categories of information in host profiles. As this study is exploratory in nature, we established some guidelines for developing the initial coding scheme. In particular, when creating an Airbnb profile page, the website prompts the host to share a few details about themselves, calling out three types of self-disclosure: "things you like", "style of traveling or hosting", and "life motto". We were cognizant of the Airbnb interface prompt and used it as a starting point, fitting codes to the prompt topics and refining them according to the content, but not restricting our coding to topics suggested by these prompts.

For this step, we constructed a *Development Dataset* consisting of 300 sentences randomly drawn from a weighted sample of 203 host profiles from 12 major U.S. cities. The profiles were extracted from an open-sourced Airbnb dataset collated by an independent organization, Inside Airbnb [26]. Non-English profiles were filtered out. We provide more details of the full dataset below.

Two authors independently coded the topics in each of the 300 sentences in the *Development Dataset*, using the qualitative data analysis and research software Atlas.ti. In addition to Airbnb prompts, the coders also considered topics used in previous self-disclosure studies [27, 31, 38]. After a full round of independent coding, the two coders compared their codes and deliberated, further merging the codes into concepts and topics. This analysis and coding process resulted in nine initial topic categories.

Step 2: Adjustment and Validation

In this step, we evaluated, adjusted, and validated the coding scheme for reliability and coverage, i.e. the percentage of sentences our codes applied to. In order to apply the coding scheme to a large set of host profiles on Airbnb, we designed a web interface to recruit annotators from Amazon Mechanical Turk (AMT). The web interface presented each individual sentence from the host profiles, together with the initial topics and descriptions (some with examples) that were developed through the coding process in the foregoing step. The annotator was instructed to tag all topics that appeared in the sentence (a sentence could mention multiple topics). If none of the topics applied, the annotator was instructed to choose "other".

We validated the reliability of the coding scheme using two metrics: the level of agreement among crowd workers, and the level of agreement between the crowd workers consensus and expert annotations, i.e. researchers from our team.

To compute the first metric, the level of agreement among crowd workers, we constructed the *Initial Validation Dataset*, consisting of 300 sentences drawn from a new sample of 203 profiles. Sentences in the *Initial Validation Dataset* did not overlap with those in the *Development Dataset*. We recruited crowd workers from AMT to annotate sentences in the *Initial Validation Dataset* using a web interface that we developed (paying $.02 per annotation), and computed Fleiss' kappa among the workers. There were four topics that had a Fleiss' kappa score lower than 0.5, indicating an unsatisfactory level of internal agreement. We iterated on the initially developed set of topics to address this issue, adjusting the name and description of two topics, and merged two closely related topics into one ("Hosting Attitude" and "Hosting Action" to "Hospitality"). After the edits to the coding scheme, eight topics remained, shown in the first column of Table 1.

To compute the second metric—the level of agreement between the consensus from crowd workers and expert annotations—we constructed the *Final Validation Dataset*, consisting of all 871 sentences from the text of a new batch of 203 profiles. The new profiles did not overlap with either those in the *Development Dataset* or those in the *Initial Validation Dataset*. Again, we asked three crowd workers to annotate each sentence, and used a majority voting rule to produce the final *vote* across three workers. A topic label for a sentence

Topic	Agreement			Description
	Vote-A1	Vote-A2	A1-A2	
Interests & Tastes	.77	.85	.78	Favorite books, music, hobbies, how I spend weekends and evenings, favorite ways of spending spare time.
Life Motto & Values	.51	.56	.52	Life motto, values, philosophies; e.g. "Live courageously, love passionately".
Work or Education	.83	.86	.79	Current or past job, school, major; e.g. "I'm an architect and designer".
Relationships	.69	.62	.60	Family, significant other, pet; e.g. "I have a beautiful 16 year old daughter, a little sweet terrier Nora, two fish & a frog."
Personality	.80	.70	.65	e.g. "I am extremely down to earth and I am a self-diagnosed work-a-holic".
Origin or Residence	.78	.69	.76	Where from, current residence, history of moving; e.g. "I lived in D.C. for 5 years and Philly for 2 years"; "We both really love how much Chicago has to offer."
Travel	.78	.72	.83	Love for travel; past travels; favorite travel destinations.
Hospitality	.73	.54	.66	Welcoming or greeting, reasons for hosting, demonstrating availability; e.g. "We're delighted to be your hosts and tour advisors during your stay here."

Table 1: Topics of self-disclosure in Airbnb host profiles.

was retained only if at least two out of three workers indicated that the sentence mentioned that topic.

In terms of coverage, in total, at least two voters agreed on one or more topics for 91.5% of the 871 sentences (if we consider the workers' using an optional "other" category, a majority vote was achieved for 97.4% of the sentences). The authors inspected the sentences where the workers did not reach an agreement, and verified that they did not contain significant missed themes.

Finally, two authors of the paper coded a 300-sentence sample from the *Final Validation Dataset*. We computed the agreement of these three different sources by calculating the pairwise Cohen's kappa scores for the worker majority vote (*vote* in Table 1), Author 1 (*A1*), and Author 2 (*A2*). The results of each pairwise agreement computation are shown in Table 1. The results suggest moderate to almost perfect agreement across all topics, and indicate that the coding scheme (the topic names, descriptions, and the set-up of majority votes from AMT) is reliable.

Phase 2—Applying the Coding Scheme to Profiles

With the coding scheme validated, we could now annotate a large set of host profiles and examine the trends of self-disclosure. What might we expect for the disclosures? According to the Profile as Promise framework [15], hosts should disclose information they believe will signal to potential guests that they will be a trustworthy host. These disclosure goals should cause hosts to prioritize the disclosure of information that enhances trustworthiness. Signaling theory further suggests that perceptions of trustworthiness may be affected by the kind of signal [14, 39]. If this is the case, then hosts should disclose more assessment signals (i.e. disclosures that can be verified):

H1.1 *Hosts will disclose more about categories that have more assessment value, including Work or Education, and Origin or Residence, than about categories that have more conventional value, including Interests & Tastes, and Personality.*

The Profile as Promise framework also suggests that information should be disclosed about the most relevant underlying qualities that the host is promising to potential guests. In this context, the type of hosting situation, on-site versus remote, should lead to different disclosure patterns. The on-site hosts (who share their space with guests during their stay) need to signal what kind of person a guest might meet. The remote hosts (who are not present) need to signal that the guests will be taken care of in their absence. Previous work on Airbnb revealed that on-site versus remote hosting is an important part of sociability within the host-guest relationship [25, 29]. When hosting on-site, guests and hosts may have more substantial face-to-face interaction.

Given the increased likelihood of social interaction for on-site hosts, there is uncertainty about whether the guests and hosts will get along. We can draw on URT [5] to predict that on-site hosts will disclose more information than remote hosts in an effort to reduce the uncertainty for potential guests given that guests and hosts will socially interact. In particular, on-site hosts should disclose more information relevant to relationship development, such as one's preferences and personality.

H1.2 *On-site hosts will disclose more, especially for topics that can reduce uncertainty during the interaction of sharing spaces, such as Interests & Tastes, and Personality, than remote hosts.*

We first report on the dataset of Airbnb profiles we used for this analysis and throughout the rest of this paper. Then, we describe the process of applying the coding scheme to annotate a larger portion of the host profiles. Finally, we discuss the results of testing the hypotheses using the annotated data.

Airbnb Dataset

To apply the coding scheme to a larger portion of Airbnb profiles, we used the large-scale dataset collected by Inside Airbnb [26]. Inside Airbnb periodically scrapes the Airbnb website, making snapshots of Airbnb listings from 35 cities in 13 countries (at the time of writing) available for download. For each city, Inside Airbnb conducted a URL query through Airbnb search and scraped all public listings. We manually examined 10 samples from five cities in the dataset, visiting the Airbnb website for each entry to verify that the scraped data is consistent with the actual listing. Since each listing is always associated with a host, the free text portion of the host profile is available from the Inside Airbnb dataset. Some other metadata about hosts, in addition to the host self-description (the focus of the present paper) include: host ID (a unique identifier for a host across the Airbnb platform), first name, type of listing (Entire Home/Apt, Private Room, or Shared Room), and whether the host is a "superhost" on Airbnb.

We limited our scope of analysis to U.S. and English-language host profiles only. Host profiles from other countries may contain non-English phrases or characters, introducing sources of noise, and making it difficult for crowd workers to annotate. We performed source language detection using the Google Translate API [21] for each sentence in a host profile, and filtered out those containing non-English sentences.

The Inside Airbnb data included 93,361 listings across 15 U.S. cities. We first de-duplicated hosts from multiple listings by host ID. We used data from the 12 largest cities, excluding 3 cities with fewer than 1,000 unique hosts (Asheville, Oakland, and Santa Cruz County). We verified that the exclusion of these three cities did not affect the results from Study 1. For the remaining 12 cities, we deduplicated 89,965 listings to obtain 67,465 unique hosts. Out of these unique hosts, we further filtered out 20,710 (or 30.7% of the de-duplicated quantity) host profiles with empty self-descriptions, and 6,750 (10.0% of the de-duplicated quantity) that contained non-English phrases.

In the end, we had 40,005 non-empty, English-only unique host profiles from 12 U.S. cities, with the following breakdown: New York (data collected in Sep 2015; 14,513), Los Angeles (Jan 2016; 8,062), San Francisco (Nov 2015; 3,400), Austin (Nov 2015; 2,477), Chicago (Oct 2015; 2,149), Seattle (Jan 2016; 1,798), Washington D.C. (Oct 2015; 1,633), San Diego (Jun 2015; 1,522), Portland (Sep 2015; 1,415), New Orleans (Sep 2015; 1,173), Boston (Oct 2015; 922), and Nashville (Oct 2015; 941).

With our previously validated coding scheme, we annotated the topics of a larger portion of host profiles from the above-mentioned Airbnb dataset using AMT, following the exact same procedure as described for annotating the *Final Validation Dataset*. We constructed the *Annotation Dataset*, consisting of all 4,377 sentences from 1,031 profiles, randomly selected using a weighted sample according to the number of unique non-empty host profiles in each city. As the coding scheme was the same as that used for the *Final Validation Dataset*, we merged the results from the *Annotation Dataset* and the *Final Validation Dataset* , forming the *Experiment Dataset* to boost the amount of annotated data, resulting in 5,248 annotated sentences from 1,234 profiles.[1]

Self-Disclosure Trends

What do Airbnb hosts self-disclose in their profiles? We found that hosts were most likely to talk about *Origin or Residence* (68.8%), followed by *Work or Education* (60.29%) and *Interests & Tastes* (57.78%). There was substantial travel-related disclosure including writing about *Travel* (47.89%) and demonstrating *Hospitality* (52.76%). The topics that were less commonly mentioned were *Relationships* (27.88%), and *Personality* (26.58%). The topic that was least mentioned was *Life Motto & Values* (7.86%).

[1]The *Experiment Dataset* and other data used in this work are available from https://github.com/sTechLab/AirbnbHosts.

This pattern of results are partially supportive of H1.1 and the prediction from signaling theory that hosts would disclose more assessment signals than conventional ones. Consistent with the hypothesis was the frequent disclosure of assessment signals regarding *Origin or Residence* and *Work or Education*, and the low levels of disclosure regarding *Personality*. The frequent disclosure of *Interests & Tastes*, however, did not line up with the hypothesis. The analysis below on disclosures by host type provides some insight: the high rate of disclosure of *Interests & Tastes* was driven in part by on-site hosts, which may have been part of an effort to reduce uncertainty for guests who would be meeting their hosts.

Self-Disclosure Trends by Host Type

Addressing H1.2, we compared the self-disclosure of hosts based on the type of property they offered: on-site versus remote. As hypothesized, on-site hosts (M=66.12, SD=59.66) on average wrote longer profiles (measured by word count) compared to remote hosts (M=55.85, SD=52.09), $t(880)$=3.07, $p < .01$. Second, in terms of topics, we found that on-site hosts were more likely than remote hosts to write about topics that signal their personality and tastes. We calculated the percentage of profiles that mentioned each topic for different host types, shown in Figure 1a. For each topic of disclosure (identified by the first letter on the x-axis), we show the percentage of profiles of each host type that mentioned that topic (y-axis). As the figure reveals, on-site hosts were more likely to write about topics of *Interests & Tastes* (χ^2=18.57, df=1, $p < .001$) and *Personality* (χ^2=29.18, df=1, $p < .001$); and less likely to mention *Origin or Residence* (χ^2=4.17, df=1, $p < .05$). Note that the results for *Interests & Tastes* and *Personality* remain statistically significant with Bonferroni correction for multiple tests.

We also compared the *number of sentences* used for the different disclosure topics, which shows similar trends (Figure 1b). For *Interests & Tastes*, on-site hosts on average wrote more sentences (M=1.87, SD=.08) than remote hosts (M=1.59, SD=.05), $t(535)$=3.09, $p < .01$. For *Personality*, we also see that on-site hosts on average wrote more sentences (M=1.42, SD=.06) than remote hosts (M=1.17, SD=.03), $t(260)$=3.64, $p < .001$. Finally, on-site hosts on average wrote more sentences mentioning *Life Motto & Values* (M=1.56, SD=.19), $t(48)$=2.34, $p < .05$ than remote hosts (M=1.10, SD=.04). Again, the results for *Interests & Tastes* and *Personality* remain statistically significant with Bonferroni correction for multiple tests.

To rule out the possibility that these differences are due to a host's level of experience (e.g. hosts may modify their profiles to write about specific topics more as they host more guests), we conducted a similar analysis comparing average hosts with superhosts, a qualification type assigned by Airbnb [2] for hosts that meet several criteria, including frequent hosting, high response rate, and high review scores. We omit the details of this analysis for brevity, but note that despite the fact that superhosts wrote significantly longer profiles (a mean of 72.13 words compared to 57.74 words for non-superhosts, $t(220)$=3.01, $p < .001$), there was no significant difference between the groups in the likelihood of mentioning *Interests*

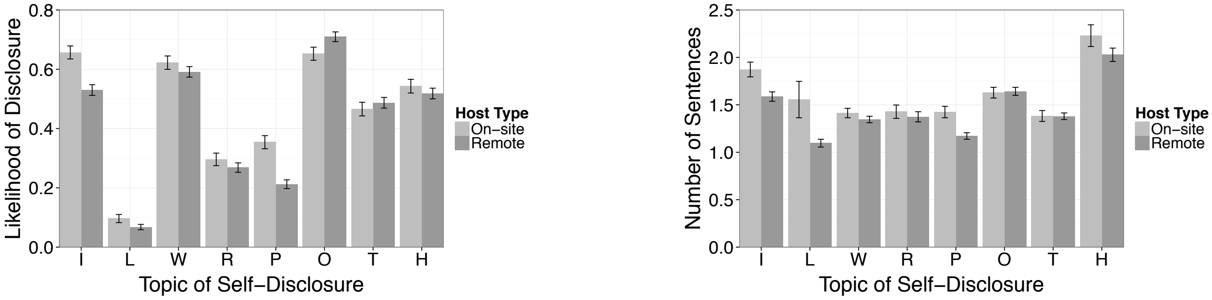

(a) Probability of disclosure per topic

(b) Average number of sentences per topic

Figure 1: Self-disclosure trends by topic and host type. The error bars represent one standard error.

& Tastes or *Personality*. This analysis suggests that the difference between on-site and remote hosts is not due to any difference in experience of hosting, but rather due to the expected differences in the type of interaction that is going to take place. Taken together, the results support URT's central contention that people seek to reduce uncertainty in the face of new relationships.

STUDY 2—SELF-DISCLOSURE AND PERCEIVED TRUSTWORTHINESS

Study 1 revealed that we can classify host disclosures into eight topics, and that these topics can be reliably assessed by independent coders. An important next question is whether the hosts were disclosing information that enhanced perceptions of their trustworthiness. That is, do the topics that hosts disclosed the most in Study 1 lead to higher levels of perceived trustworthiness? To examine this question we asked online participants to rate how trustworthy they found each profile.

One way to operationalize the concept of trustworthiness is by using three key dimensions: ability, benevolence and integrity [33]. These three dimensions are closely related but may have different effects on trust depending on context [10]. Further, these dimension are all likely to be relevant for Airbnb profiles. In the context of Airbnb, ability refers to domain-specific skills or competencies that the host has. Benevolence refers to the extent to which the host is believed to want to do good to the guest beyond profit-driven motives. Finally, integrity refers to the host adhering to a set of moral principles and rules.

How might the disclosures in Airbnb profiles influence these dimensions of trustworthiness? The Profile as Promise conceptualization of the profile as a psychological contract implies that information provided in the profile is an obligation by the host to a guest, namely that the information disclosed in the profile is trustworthy and will not misrepresent the host or the host's home [15]. This notion of the psychological contract suggests that hosts should be sensitive to how their promises will be evaluated for trustworthiness by potential guests. If this is the case, then hosts should produce promises that signal trustworthiness. We can draw on the same theoretical perspectives we used to characterize the production of disclosures to specify predictions about how the profile disclosures on Airbnb affect evaluations of trustworthiness. First, URT [5] predicts that the more information hosts disclose, the more the profile will reduce uncertainty for profile viewers, which should enhance how trustworthy they will be perceived. Note that more diverse information should lead to more uncertainty reduction. That is, profiles that disclose more kinds of information will be perceived as more trustworthy than profiles than simply say a lot about fewer things. We therefore predict that:

H2.1 *Longer and more diverse self-disclosures are perceived as more trustworthy.*

Secondly, Study 1 demonstrates that hosts communicate a variety of topics in their profiles. Signaling theory predicts that hosts use these disclosures to signal underlying qualities or attributes that should enhance the perceptions of their trustworthiness as a host. If the hosts have optimized their signaling behavior for trustworthiness, then the categories they disclose most often should be the categories of disclosure that are perceived as most trustworthy. Thus, profiles with disclosures that were observed frequently in Study 1, including *Origin or Residence*, *Work or Education*, *Interests & Tastes*, and *Hospitality* should be perceived as more trustworthy than profiles that do not contain these topics.

H2.2 *Self-disclosure topics used most frequently by hosts will be associated with increased perceived trustworthiness compared to less frequent topics.*

Methods

As mentioned above, we are interested in the perceived trustworthiness of host profiles. To measure trustworthiness, we developed a six-item perceived trustworthiness scale on three dimensions: ability, benevolence, and integrity [33]. Based on items in the scale developed by Mayer et al. for an organization context [32], we developed new items that measure trustworthiness in the context of hosting. These items are shown in Table 2. Items A1–A2 measure ability; items B1–B2 measure benevolence; and items I1–I2 measure integrity. When asking for profile ratings, these items were shown in a random order.

Procedure

To assess the perceived trustworthiness of host profiles, we recruited crowd workers from AMT to rate host profiles in the *Experiment Dataset* using the perceived trustworthiness

A1.	This person is capable of paying his/her own rent or mortgage.
A2.	This person maintains a clean, safe, and comfortable household.
B1.	This person will be concerned about satisfying my needs during the stay.
B2.	This person will go out of his/her way to help me in case of an emergency during my stay.
I1.	This person will stick to his/her word, and be there when I arrive instead of standing me up.
I2.	This person will not intentionally harm, overcharge, or scam me.

Table 2: Six-item perceived trustworthiness scale.

scale. We split the *Experiment Dataset* profiles into batches of 20, and had five different workers annotate each batch. Recall that these profiles were already labeled with the topics. We used 1,200 of the 1,234 profiles in the *Experiment Dataset* for this study. For each profile, workers were instructed to rate their level of confidence regarding each of the statements, on a scale from 0 to 100, with steps of 10. The task required that each worker only rate one batch of the profiles to prevent any single worker's perception from being over-represented in the results. Workers were paid $1.00 for each task.

At the beginning of the each task, we used a paraphrase question borrowed from [12, 35] to check the linguistic attentiveness of each worker. We re-issued the task if we received an incorrect response to this question. To create the perceived trustworthiness score, we calculated the perceived trustworthiness as the mean of responses for all six items by the five workers that rated the same profile. For some analyses, we also used three trustworthiness dimensions separately, with each score calculated as the average of the two relevant items.

Results

We investigated the effects of profile length and diversity (H2.1), and topic (H2.2) on perceived trustworthiness. Generally, the mean ability score of the 1,200 profiles was 68.82, *SD*=13.84; the mean benevolence score was 63.94, *SD*=13.97; the mean integrity score was 66.79, *SD*=13.37. Note that perceived trustworthiness scores across the three dimensions were highly correlated [pairwise Pearson's R correlation: A–B (initials): 0.86; A–I: 0.88; B–I: 0.92; $p < .001$].

Length, Diversity and Perceived Trustworthiness

To examine the effect of length (word count) on perceived trustworthiness, we plot the relationship between length (x-axis, log scale) and perceived trustworthiness (y-axis) on each of the three trust dimensions in Figure 2.

Supporting H2.1, Figure 2 shows a clear relationship between increased profile length and perceived trustworthiness scores. This relationship is confirmed by linear regression with log transformation for profile length [$b = 7.89$, adjusted R^2=.38, $F(1, 1198)$=721.4, $p < .001$]. This means that when a profile doubles in length, the perceived trustworthiness score increases by approximately 5.47, suggesting a pattern of diminishing returns when hosts write longer profiles. To illustrate this pattern, we divide the profiles into deciles and calculate the average perceived trustworthiness score for each decile. Comparing profiles in the second decile (mean word count: 13) to those in the first (mean word count: 6), mean trustworthiness score increased 18.9%; whereas comparing profiles in the ninth decile (mean word count: 106) to those in the tenth (mean word count: 188), mean trustworthiness score only increased by 2.5%.

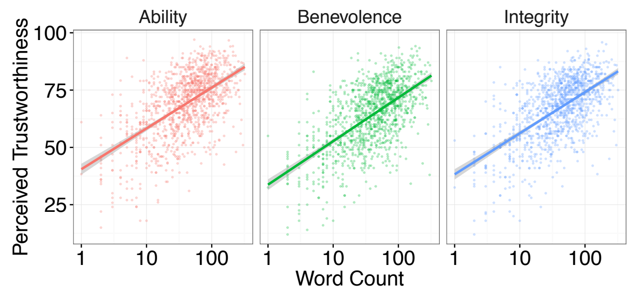

Figure 2: Perceived trustworthiness increases with profile length (x-axis on log scale).

H2.1 also predicts that, in addition to overall length, the number of topics will also have a positive impact on trustworthiness scores. We performed multiple linear regression analysis with the number of topics, length as control, and the interaction length × number of topics [adjusted $R^2 = 0.39$, $F(3, 1196) = 256.6$, $p < .001$]. The analysis showed that the number of topics contributes to perceived trustworthiness [$b = 4.47$, $t(1198) = 5.54$, $p < .001$] even when controlling for length [log scale, $b = 9.53$, $t(1198) = 15.03$, $p < .001$]. There was also an interaction effect between length and topic count [$b = -0.95$, $t(1198) = -5.04$, $p < .001$], indicating that for shorter profiles, the number of topics increased perceived trustworthiness even more.

Figure 3 visualizes the relationship between perceived trustworthiness and the number of topics mentioned in the profile. Each line represents the density distribution of perceived trustworthiness scores for profiles that mention a fixed number of topics. For example, the darkest lines represent the distribution of perceived ability, benevolence, and integrity of one-topic profiles. The figure shows that there is variation in trustworthiness score within each topic count bin, but as topic count increases, the trustworthiness scores also increase, and the variations become smaller. Note that here we are *not* showing the effect of profile length, which was illustrated in Figure 2.

Topic and Perceived Trustworthiness

We now analyze the effect of topic choice on trustworthiness scores. Recall that H2.2 predicted that the topics disclosed most frequently by hosts in Study 1, namely, *Origin or Residence*, *Work or Education*, *Interests & Tastes*, and *Hospitality*, would also be evaluated as most trustworthy.

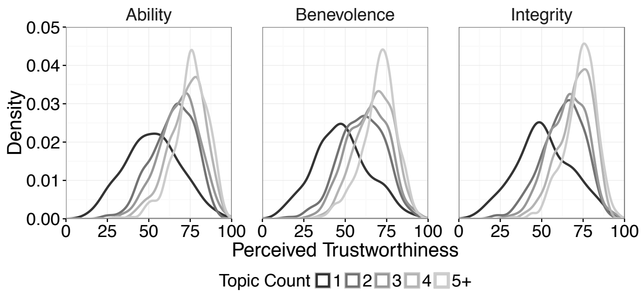

Figure 3: Perceived trustworthiness score distributions for profiles with different number of topics.

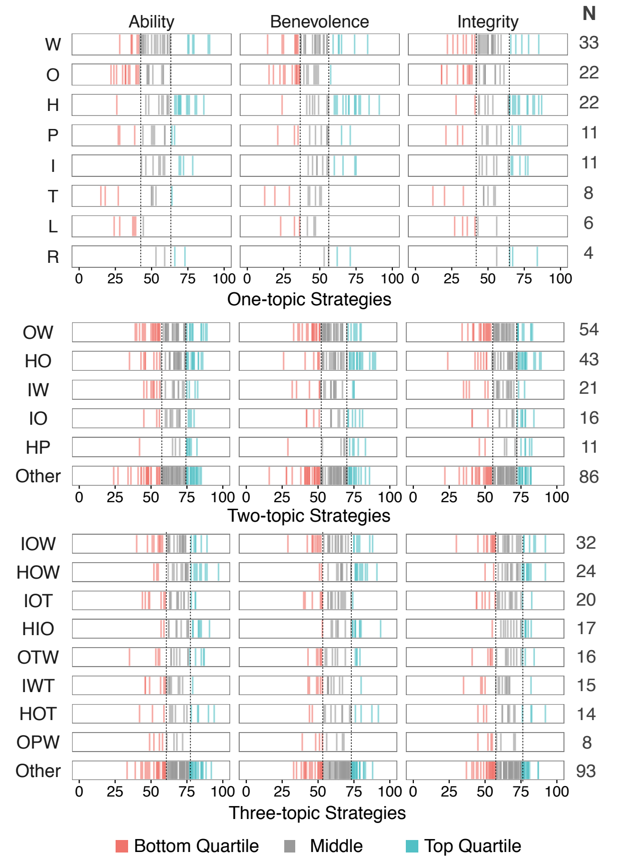

Figure 4: Comparison for different strategies, organized by the number of topics mentioned in the profile. The dotted lines indicate the bottom and top quartiles in each topic-count group.

In our dataset, there were eight profiles that did not mention any topics, 117 one-topic profiles, 231 two-topic profiles, 239 three-topic profiles, 269 four-topic profiles, and 336 profiles that mentioned five or more topics. We focus on profiles that are limited to one-topic, two-topic, and three-topic combinations. For example, looking at two-topic combinations, there are $8 \times 7/2 = 28$ different options, although, as we show below, there are some topic combinations that are more common than others. These 1-3 topic combinations have the most variation, but are also simpler to study, as understanding the impact of one single topic amid all combinations of different sizes is highly unlikely even with 1,200 profiles. We call these different combinations of topics "strategies", and compare the relative success of different strategies controlling for the number of topics.

We have shown that as the number of topics increase, the trustworthiness scores also increase. We computed one-way ANOVAs comparing the *relative* effectiveness of strategies within each of the one-topic, two-topic, and three-topic profile groups. For one-topic profiles, there was a significant effect of strategy on ability [$F(7,109) = 7.79$, $p < .001$], benevolence [$F(7,109) = 8.55$, $p < .001$], as well as integrity [$F(7,109) = 7.36$, $p < .001$]. For two-topic profiles, there was no significant effect of strategy on ability [$F(5,225) = 1.54$, $p = .18$], but a significant effect on benevolence [$F(5,225) = 3.95$, $p < .01$], as well as integrity [$F(5,225) = 2.74$, $p < .05$]. For three-topic profiles, there was a significant effect of strategy on ability [$F(8,230) = 2.83$, $p < .01$], benevolence [$F(8,230) = 4.21$, $p < .001$], as well as integrity [$F(8,230) = 2.95$, $p < .01$].

Figure 4 shows the raw data for this analysis, organized by the number of topics (panels), and three dimensions of trustworthiness scores (columns). Every row is marked on the left with the initials of the topics in the self-disclosure strategies (e.g. in the first row of the second panel, *OW* stands for the topic combination of *O*rigin or Residence, and *W*ork or Education). On the right, we show the number of profiles using this strategy (54 for *OW*, the most of all two-topic strategies). The vertical lines in each row represent profiles, positioned at the value of the perceived trustworthiness score on each dimension. The color indicates whether the profile falls within the bottom (red), or top (green) quartile of the profile group that used the same amount of topics (the dotted lines indicate the bottom and top quartile boundaries).

There are several takeaways from Figure 4. Consider first the one-topic strategies: clearly, *W*ork or Education, *O*rigin or Residence) and *H*ospitality are the most popular, representing 66% of the one-topic profiles. Visually, it is clear from Figure 4 that the most successful single-topic strategy is *H*, where the profiles trending more to the right and top-quartile profiles appearing more frequently. This visual examination is confirmed with post-hoc comparisons using Tukey HSD tests. Among one-topic profiles, *H*ospitality was the best-performing strategy, significantly trumping *L*, *O*, *P*, *T* for ability; *L*, *O*, *T*, *W* for benevolence and integrity ($p < .05$, same for all the post-hoc comparisons reported henceforth).

The second-best one-topic strategy was *I*nterests & Tastes, outperforming *L* and *O* for ability, and *O* and *T* for benevolence and integrity. Hosts were not very successful writing about *L*ife Motto & Values, although Airbnb explicitly prompted them to, which underperforms *H*, *I*, *R*, *W* for ability; *H* for benevolence, and *H* and *R* for integrity. Finally, as reflected in Figure 4, *W* outperforms *O* for ability; *O* outperforms *R* for benevolence, *R* outperforms *O* and *T* for integrity. None of the other pairwise comparisons were significant.

Moving on to the two-topic strategies, the dominant strategies are *OW* and *HO*, both combinations of the most popular single-topic strategies *W*, *O* and *H*, covering 42% of two-topic profiles. Interestingly, the *WH* strategy was not often used (for three-topic combinations, *HOW* is again popular). The next two popular strategies are *IW* and *IO*, indicating that hosts add on *I*nterests & Tastes as additional information. In terms of success for the two-topic strategies, post-hoc comparisons did not indicate significant differences among strategies for ability or integrity. However, for benevolence, *HO* outperforms *OW* and *Other* (all other two-topic combinations that are not explicitly listed in Figure 4).

Finally, we see in three-topic combinations that the most common strategies are *IOW* and *HOW*. In terms of relative success, post-hoc comparison indicated that *HOW* is clearly most successful, outperforming *IWT* for ability, *IOT*, *IOW*, *IWT*, *OPW*, *Other* for benevolence, and *IWT* for integrity. In addition, *HIO* outperforms *IOT* and *IWT* for benevolence. This may again be due to the high effectiveness of *H*ospitality as part of the disclosure strategy, when the host is making a direct promise to take care of the guests.

Overall, the pattern of results supports H2.2 and the prediction that profiles with topics most frequently disclosed by hosts are also those that are evaluated as most trustworthy. While hosts employed different combinations of topics as part of their self-disclosure strategies, it is clear that strategies that include the most frequently disclosed topics from Study 1 were the most successful in generating perceived trustworthiness: *Work or Education*, *Origin or Residence*, *Hospitality*, and *Interests & Tastes*. We now proceed to show that these trustworthiness scores are meaningful because they have a direct impact on host choice by potential guests.

STUDY 3—FROM PERCEPTION TO CHOICE

In this section, we examine how perception of trustworthiness leads to differences in host choice. As mentioned earlier, a number of factors may influence a potential guest's *decision* to stay with a host, such as availability, price, and characteristics of the property (e.g. location). Our primary question is whether the trustworthiness signaled by profile disclosures can influence a potential guest's decision-making outcome, all other things being equal. To address this question, we isolate disclosures in the profile by conducting an online experiment to examine the effect of perceived trustworthiness on host choice. In particular, we vary the level of perceived trustworthiness, and test the extent to which the perceived trustworthiness of profiles influences a potential guest's choice.

Understanding choice has important real world implications. In the face of a potential social exchange opportunity with multiple exchange candidates, those who portray themselves as untrustworthy can potentially be "punished". As shown above, the content of an Airbnb host profile affects perceived trustworthiness. We know that trustworthiness differences can affect choice [17] in other settings, and hypothesize that:

H3.1 *Higher perceived trustworthiness scores for text-based host profiles predict the likelihood of guest choice.*

Methods

To test whether perceived trustworthiness affects a potential guest's decision-making, we employed a pairwise experiment to elicit guest response. Since we obtained a perceived trustworthiness score for each profile in Study 2, we paired profiles with different scores to examine if the score predicts guest's preference between two hosts in a pair. If the value of the trustworthiness score perfectly predicts choice, the observed pairwise decisions we obtain from respondents should follow the Bradley-Terry model [7], which predicts the outcome of a comparison given associated values with each participant in the match.

To this end, we generated profile pairs that were comparable in length, but with one high and one low perceived trustworthiness score. We controlled for length for a number of reasons. Firstly, we showed above that the length is highly correlated with trustworthiness. Choosing high- and low-scoring profiles from a global sample is therefore likely to result in unbalanced short and long profile pairs. We therefore used an adaptive matching method that takes length, then score into account. First, we ranked 1,200 annotated host profiles based on word count. Then, from shortest to longest, we used a sliding window of roughly 240 profiles, with steps of size 120. All profiles within each window form a group. For each group, we calculated the bottom and top quartiles of mean trustworthiness score (the mean of the ability, benevolence, and integrity). We then iterated through every combination of two profiles, one from the bottom and one from the top quartile in that group, filtering out profile pairs where one profile is longer than the other by more than 20%. As the result of this process, we had 19,892 top-quartile-low-quartile (in the sliding window) profile pairs, representing 589 unique profiles.

The preference task for each profile pair was simple. First, each pair of profile descriptions was shown to a respondent. For the first five seconds, the profiles are shown but buttons were deactivated to encourage the respondent to read the profiles before making a decision. When the buttons become activated, the respondent click on one of the two profiles in response to the question, "Which of the two hosts do you feel more comfortable staying with?" We deliberately chose to ask about comfort, and not generally about host preference, e.g. "Which of the hosts would you choose to stay with?" Pilot studies we ran showed that, in making host preference decisions, people considered other personal and dyadic match factors, like their interest in staying in a location implied in the profile. While such considerations are generally interesting, the "comfort" phrasing was used to focus on the effect of the trustworthiness construct. A future study can address ecological validity by including other factors; we show that the trustworthiness construct *does* impact pairwise preference based on comfort, holding all else constant.

We recruited respondents from AMT for this task. Each respondent was asked to evaluate 50 profile pairs in each task. To improve the quality of the results, three pairs out of each batch of 50 were repeated in a random order, within each batch. If the respondent provided an inconsistent answer for more

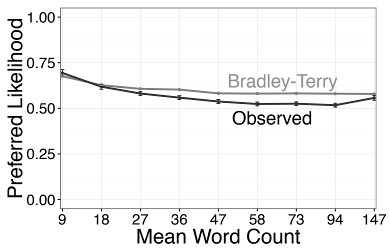

Figure 5: Observed likelihood that the host with high trustworthiness score is preferred and the probability predicted by Bradley-Terry model, by profile length. The error bars represent one standard error.

than one pair, we filtered their responses out of the analysis. Workers were paid $1.00 per task.

Experimental Results

We obtained choice results from 423 unique respondents. Consistency filtering, and removing four responses with missing values, left us with 355 responses consisting of 16,685 pairs of choices, representing the 589 unique profiles.

We used the Bradley-Terry model [7] to evaluate pairwise choice. The Bradley-Terry probability model predicts the outcome of a comparison given associated values with each participant in the match. If the perceived trustworthiness score accurately predicts choice, it should be a good fit to the theoretical Bradley-Terry model likelihood. Specifically, according to the Bradley-Terry model, the probability of profile i being picked compared to profile j is [7]:

$$P(i \text{ is preferred to } j) = \frac{\lambda_i}{\lambda_i + \lambda_j} \tag{1}$$

where λ is a positive-valued parameter associated with each individual option. In our case, λ is the perceived trustworthiness score. Intuitively, the larger the difference between λ_i and λ_j, the higher the probability of i being chosen (with a upper-bound of 1). In addition, if the difference between λ_i and λ_j is fixed, the probability of i being chosen decreases as the absolute values of λ_i and λ_j increase (with a lower-bound of 0.5).

Figure 5 summarizes, for different buckets of profile length, the effectiveness of perceived trustworthiness in predicting choice compared to the prediction of the Bradley-Terry model. The x-axis shows the mean word count of the profile pairs in that length bucket, and the y-axis shows the likelihood of the profile with the higher trustworthiness score being chosen. Both theoretical (grey) and observed (black) likelihoods are plotted. The figure shows that for profiles in shorter length groups (the first two groups on the left), perceived trustworthiness predictions closely match the Bradley-Terry probability. For longer profiles, however, the observed likelihood is lower than what is theoretically expected, indicating less predictive power. Nevertheless, even for the longer profile, as Figure 5 shows, the likelihood of choosing the top-quartile profile was higher than chance (50%), as shown by Exact-Binomial tests for each length group ($p < .05$) except for the eighth (the bin with mean word count of 94, $p = .06$).

The divergence of the longer profiles from the model may be due to at least two factors. First, and more mundanely, we ran the task on AMT where workers may not be incentivized to spend time reading longer profile descriptions, skewing the results towards random chance. Second, the results may reflect the fact that higher trustworthiness scores for longer profiles are not as predictive for decisions. With longer profiles, other factors are more likely to be mentioned, such as interests and tastes, that may generate specific dyadic attractions, and play a more significant role in influencing choice. We discuss this possibility in greater detail below.

DISCUSSION

The three studies reported here make several contributions. First, we developed and validated a coding scheme for self-disclosure on the free-text portion of host profiles on Airbnb. The coding scheme describes eight topics that covers more than 90% of current discourse in host profiles. To the best of our knowledge, this is the first systematic coding scheme for analyzing self-disclosure in Airbnb profiles, or more generally, for profiles related to peer-to-peer sharing platforms.

The results of applying this coding scheme to the Airbnb profile dataset revealed that hosts most frequently write about *Origin or Residence*, *Work & Study* and *Interests & Tastes*. The least commonly disclosed topics were *Life Motto & Values*, *Relationships* and *Personality*. Study 1 also revealed that host type influenced the kinds of disclosures produced in host profiles, with on-site hosts revealing more information about their *Interests & Tastes* and *Personality* than remote hosts. These data are consistent with predictions from URT [5], which predicted that on-site hosts will disclose more information about their *Interests & Tastes* and *Personality* to reduce potential guest uncertainty about whether they would enjoy interacting with an on-site host.

Our studies also drew on signaling theory [39] to understand what topics hosts disclose, and how guests perceive those disclosures. To assess the implications for signaling theory, it is informative to consider the results across both Study 1, which focused on the production of disclosures in host profiles, and Study 2, which examined how those disclosures affected the perceptions of trustworthiness. Signaling theory predicts that hosts will signal their trustworthiness by disclosing more assessment signals (e.g. *Origin or Residence*, *Work or Study*), which are more difficult to fake than conventional signals (e.g. *Life Motto & Values*, *Personality*) [14]. Signaling theory also predicts that, if hosts are optimizing their disclosures for trustworthiness, then guests should evaluate profiles with the most frequently observed topics as most trustworthy. The data from Studies 1 and 2 largely confirmed both of these hypotheses: hosts disclosed more assessment signals than conventional ones, and guest perceived profiles with more assessment signals as more trustworthy.

There were, however, some important exceptions to the theoretical predictions. Certain strategies, such as demonstrating *Hospitality* or sharing one's *Interests & Tastes*, proved to be more successful than expected for conventional signals while other strategies, such as providing only one's *Origin & Residence*, proved less successful. Providing a welcome or greeting, or providing reasons for hosting, alone but preferably in combination with more assessment signal disclosures (e.g., *Hospitality* combined with *Origin & Residence*), was an important strategy that had a strong and positive effect on perceptions of trustworthiness. These findings suggest that signaling with conventional signals but that provide information about one's hospitality or interests can enhance trustworthiness in Airbnb profiles.

Finally, we demonstrated that perceived trustworthiness matters for decision-making in this context. The perception of trustworthiness from Study 2 predicted participants' decisions in a forced choice experimental task in Study 3, especially for profiles that are relatively short (less than 20 words). We also showed that when profiles are short, perceived trustworthiness almost perfectly predicts choice, whereas when the profile length increases, other factors appear to also influence choice. This may suggest a nuanced role of trust in decision making—there is a threshold of trust that is needed to pass muster, but other factors (e.g. homophily [34]) may weigh in once trustworthiness is no longer the issue.

Finally, this research suggests that the Profile as Promise framework [16] is a useful approach for understanding how hosts and guests produce and evaluate disclosures in Airbnb profiles. Hosts disclosed information about themselves that they perceived as relevant and of interest to potential guests, and their promises were evaluated based on their trustworthiness, as predicted by signaling theory and URT. This study suggests that the concept of a promise, or psychological contract, can be usefully applied beyond online dating profiles [15] and résumés [22] to peer-to-peer sharing platforms such as Airbnb.

Limitations and Future Work

There are some important limitations to this work. First and foremost, we opted to prioritize our theoretical understanding of trustworthiness in profiles, over developing an ecologically valid measure of the profile text's effect on host choice. As mentioned above, host choice on Airbnb can be impacted by many factors, including (most trivially) the price and characteristics of the rental property. Nevertheless, the experiment in Study 3 isolated and manipulated the perceived trustworthiness of the profile text, which allowed us to make causal claims regarding the profile text's impact on guest decision-making.

A related limitation of the work is the fact that it ignores dyadic and dynamic determinants of trustworthiness. A key mechanism of URT involves dyadic reciprocity and exchange [5]. In this work, we only examined a single-sided, one-time disclosure by hosts. It would be important to consider the effect of the dyadic properties of hosts and guest, and how they relate to trustworthiness and trust. Understanding how impressions of perceived trustworthiness form and evolve through conversations between hosts and guests would be another complex and interesting problem to tackle.

Our dataset only includes U.S. large cities. As a result, the findings may not generalize to hosts in smaller cities, though nothing in our findings would necessarily suggest that this would be the case. We did not consider gender and cultural differences in this work, either. In a preliminary investigation we inferred the gender of hosts from their first names, but did not find significant differences in disclosures between hosts of different gender. Future work can dive deeper into patterns of self-disclosure by individuals of different gender and cultures, potentially helping to combat discrimination or potential biases known to exist on sharing economy platforms [40].

Finally, while the paper uses a qualitative coding of profile text, another approach would have used other qualitative methods, such as interviews with hosts about their profile construction strategies. For example, how do Airbnb hosts present a trustworthy facade while balancing other important aspects (e.g. privacy)? Other research has qualitatively examined the experiences of hosts [29] but to date has not considered profile construction work.

Design Implications

Our findings have direct implications for improving the design of profile pages on sharing economy sites and services, with the view of encouraging trustworthiness and improving the rate of transactions. Our results suggest that hosts should be encouraged to disclose more information, and that this information should come from a diverse set of the eight categories identified in the coding scheme from Study 1. With knowledge of the profile features that may promote trust, interfaces for creating and editing profile text could encourage individuals to write more, and focus on the key categories exposed above. Automatic text analysis mechanisms could also be used to classify text into categories, and suggest other topics to improve breadth and ultimately perceived trustworthiness.

CONCLUSION

In the sharing economy, trust is a crucial social dynamic required to support the kinds of social exchange that the sharing economy was designed to unlock. Our study introduced a new coding scheme to categorize how hosts self-disclose in their profiles and how these disclosures affect perceived trustworthiness. We find that hosts use a variety of disclosure strategies, with some more successful than others, suggesting that platforms can support users to convey trustworthiness in their profiles and facilitate the trust dynamics required for the sharing economy.

ACKNOWLEDGMENTS

We sincerely thank the anonymous reviewers for their detailed and helpful feedback. This work is partially supported by AOL through the Connected Experiences Lab. We also thank Shahar Harel and Roi Reichart for the helpful suggestions during the early stage of this work, and colleagues at Cornell Tech who read and commented on the draft of this work.

REFERENCES

1. Airbnb. 2016a. About Us. https://www.airbnb.com/about/about-us. (Retrieved May 2016).
2. Airbnb. 2016b. Airbnb Superhost. https://www.airbnb.com/superhost. (Retrieved May 2016).
3. Natalya N Bazarova and Yoon Hyung Choi. 2014. Self-Disclosure in social media: Extending the functional approach to disclosure motivations and characteristics on social network sites. *Journal of Communication* 64, 4 (2014), 635–657.
4. Joyce Berg, John Dickhaut, and Kevin McCabe. 1995. Trust, reciprocity, and social history. *Games and economic behavior* 10, 1 (1995), 122–142.
5. Charles R Berger and Richard J Calabrese. 1975. Some explorations in initial interaction and beyond: Toward a developmental theory of interpersonal communication. *Human communication research* 1, 2 (1975), 99–112.
6. danah boyd and Nicole Ellison. 2007. Social network sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication* 13, 1 (2007), 210–230.
7. Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika* 39, 3/4 (1952), 324–345.
8. Coye Cheshire. 2011. Online trust, trustworthiness, or assurance? *Daedalus* 140, 4 (2011), 49–58.
9. Coye Cheshire and Karen S Cook. 2004. The emergence of trust networks under uncertainty-Implications for Internet interactions. *Analyse und Kritik* 26, 1 (2004), 220.
10. Jason A Colquitt, Jeffery A LePine, Cindy P Zapata, and R Eric Wild. 2011. Trust in typical and high-reliability contexts: Building and reacting to trust among firefighters. *Academy of Management Journal* 54, 5 (2011), 999–1015.
11. Karen S Cook, Toshio Yamagishi, Coye Cheshire, Robin Cooper, Masafumi Matsuda, and Rie Mashima. 2005. Trust building via risk taking: A cross-societal experiment. *Social Psychology Quarterly* 68, 2 (2005), 121–142.
12. Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. A computational approach to politeness with application to social factors. In *Proceedings of ACL*.
13. Valerian J Derlega, Barbara A Winstead, Paul Wong, and Michael Greenspan. 1987. Self-disclosure and relationship development: An attributional analysis. *Interpersonal processes: New directions in communication research* (1987), 172–187.
14. Judith Donath. 2007. Signals in social supernets. *Journal of Computer-Mediated Communication* 13, 1 (2007), 231–251.
15. Nicole B Ellison and Jeffrey T Hancock. 2013. Profile as promise: Honest and deceptive signals in online dating. *IEEE Security and Privacy* 11, 5 (Sept. 2013), 84–88.
16. Nicole B Ellison, Jeffrey T Hancock, and Catalina L Toma. 2011. Profile as promise: A framework for conceptualizing veracity in online dating self-presentations. *New Media & Society* 14, 1 (2011), 45–62.
17. Eyal Ert, Aliza Fleischer, and Nathan Magen. 2016. Trust and reputation in the sharing economy: The role of personal photos in Airbnb. *Tourism Management* 55 (2016), 62 – 73.
18. Diego Gambetta. 2009. *Codes of the underworld: How criminals communicate*. Princeton University Press.
19. Diego Gambetta and Heather Hamill. 2005. *Streetwise: How taxi drivers establish customer's trustworthiness*. Russell Sage Foundation.
20. Jennifer L Gibbs, Nicole B Ellison, and Chih-Hui Lai. 2010. First comes love, then comes Google: An investigation of uncertainty reduction strategies and self-disclosure in online dating. *Communication Research* (2010), 0093650210377091.
21. Google. 2016. Google Translate API Documentation. https://cloud.google.com/translate/. (Retrieved May 2016).
22. Jamie Guillory and Jeffrey T Hancock. 2012. The effect of Linkedin on deception in resumes. *Cyberpsychology, Behavior, and Social Networking* 15, 3 (2012), 135–140.
23. Daniel Guttentag. 2015. Airbnb: Disruptive innovation and the rise of an informal tourism accommodation sector. *Current Issues in Tourism* 18, 12 (2015), 1192–1217.
24. Russell Hardin. 2002. *Trust and trustworthiness*. Russell Sage Foundation.
25. Tapio Ikkala and Airi Lampinen. 2015. Monetizing network hospitality: Hospitality and sociability in the context of Airbnb. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. ACM, 1033–1044. DOI: http://dx.doi.org/10.1145/2675133.2675274
26. Inside Airbnb. 2016. About Inside Airbnb. http://insideairbnb.com/about.html. (Retrieved May 2016).
27. Sidney M Jourard and Paul Lasakow. 1958. Some factors in self-disclosure. *The Journal of Abnormal and Social Psychology* 56, 1 (1958), 91.
28. Toko Kiyonari, Toshio Yamagishi, Karen S Cook, and Coye Cheshire. 2006. Does trust beget trustworthiness? Trust and trustworthiness in two games and two cultures: A research note. *Social Psychology Quarterly* 69, 3 (2006), 270–283.

29. Airi Lampinen and Coye Cheshire. 2016. Hosting via Airbnb: Motivations and financial assurances in monetized network hospitality. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, 1669–1680. DOI: http://dx.doi.org/10.1145/2858036.2858092

30. Donghun Lee, Woochang Hyun, Jeongwoo Ryu, Woo Jung Lee, Wonjong Rhee, and Bongwon Suh. 2015. An analysis of social features associated with room sales of Airbnb. In *Proceedings of the 18th ACM Conference Companion on Computer Supported Cooperative Work & Social Computing (CSCW '15 Companion)*. ACM, 219–222. DOI: http://dx.doi.org/10.1145/2685553.2699011

31. Xiao Ma, Jeffrey T Hancock, and Mor Naaman. 2016. Anonymity, intimacy and self-disclosure in social media. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, 3857–3869. DOI: http://dx.doi.org/10.1145/2858036.2858414

32. Roger C Mayer and James H Davis. 1999. The effect of the performance appraisal system on trust for management: A field quasi-experiment. *Journal of applied psychology* 84, 1 (1999), 123.

33. Roger C Mayer, James H Davis, and F David Schoorman. 1995. An integrative model of organizational trust. *Academy of management review* 20, 3 (1995), 709–734.

34. Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* (2001), 415–444.

35. Robert Munro, Steven Bethard, Victor Kuperman, Vicky Tzuyin Lai, Robin Melnick, Christopher Potts, Tyler Schnoebelen, and Harry Tily. 2010. Crowdsourcing and language studies: the new generation of linguistic data. In *Proceedings of the NAACL HLT 2010 workshop on creating speech and language data with Amazon's Mechanical Turk*. Association for Computational Linguistics, 122–130.

36. Riley Newman and Judd Antin. 2016. Building for trust: Insights from our efforts to distill the fuel for the sharing economy. http://nerds.airbnb.com/building-for-trust. (Retrieved May 2016).

37. Denise M Rousseau and Martin M Greller. 1994. Human resource practices: Administrative contract makers. *Human Resource Management* 33, 3 (1994), 385–401.

38. Zick Rubin. 1975. Disclosing oneself to a stranger: Reciprocity and its limits. *Journal of Experimental Social Psychology* 11, 3 (1975), 233–260.

39. Michael Spence. 2002. Signaling in retrospect and the informational structure of markets. *The American Economic Review* 92, 3 (2002), 434–459.

40. Jacob Thebault-Spieker, Loren G. Terveen, and Brent Hecht. 2015. Avoiding the south side and the suburbs: The geography of mobile crowdsourcing markets. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. ACM, 265–275. DOI: http://dx.doi.org/10.1145/2675133.2675278

41. Catalina L Toma, Jeffrey T Hancock, and Nicole B Ellison. 2008. Separating fact from fiction: An examination of deceptive self-presentation in online dating profiles. *Personality and Social Psychology Bulletin* 34, 8 (2008), 1023–1036.

42. Suvi Uski and Airi Lampinen. 2014. Social norms and self-presentation on social network sites: Profile work in action. *New Media & Society* (2014).

43. Georgios Zervas, Davide Proserpio, and John Byers. 2015. A first look at online reputation on Airbnb, where every stay is above average. *Social Science Research Network* (2015).