Is the Sharing Economy about Sharing at All?
A Linguistic Analysis of Airbnb Reviews

G. Quattrone,[1] **S. Nicolazzo,**[2] **A. Nocera,**[2] **D. Quercia,**[3] **L. Capra**[4]

[1] Middlesex University, London, UK
[2] University Mediterranea of Reggio Calabria, Italy
[3] Nokia Bell Labs Cambridge, UK
[4] University College London, UK

Abstract

The sharing economy manifesto emphasises sharing under-used resources as a means to build stronger communities. This manifesto has however received strong critiques that claim these markets are all about access as opposed to sharing, and that consumers are after utilitarian, as opposed to social, value. Being able to assess whether an economy is about access or sharing has important implications for how companies operate, and compete, in this space. To help shed light onto this, we perform a linguistic analysis of the reviews that peers in a sharing economy platform leave to one another. We take Airbnb in U.S. as a use case, and identify the main themes that peers discuss in their reviews from 2012 to 2016. We find that, as one expects, utilitarian values (e.g., properties' facilities, convenience of location, business conduct) have been discussed much more frequently than social values (e.g., guest/host interactions), and, more interestingly, this gap has substantially increased over the years.

Introduction

The term "sharing economy" was originally coined to capture a class of economic arrangements whereby, rather than owning goods and services, individuals would share them with one another, under the premise that there exists substantial excess capacity in the system, and, with that, there exists an opportunity to optimize resources, and increase their *value*, through sharing. The sharing economy manifesto[1] emphasizes how *social* value is as important as *financial* value, with the practice of sharing seen as a means to build stronger communities.

This vision has however received strong critiques; for example, Eckhardt and Bardhi (Eckhardt and Bardhi 2015) observed that sharing is a form of social exchange that takes place among people who know each other, without any profit. The moment a company, like Airbnb or Uber, acts as a financial intermediary betweens strangers, transactions become economic exchanges. In these circumstances, peers are after utilitarian value, as opposed to social value.

Being able to assess whether an economy is about access or sharing has important implications for how companies operate, and compete, in this space. To gather in-formation about consumers' needs, market research often combines qualitative techniques (e.g., focus groups, in-depth interviews) with quantitative ones (e.g., customer surveys). These techniques require substantial financial and time investments. However, there is a continuous stream of ready-available secondary data that can be systematically analysed to extract such knowledge: reviews that peers leave to one another upon completion of a service exchange.

In this short paper, we propose a mixed-method technique to perform linguistic analysis of reviews in sharing economy markets, as a means to systematically gather knowledge of peers' values at a certain point in time and space. We take the case of AirBnB as an example, mainly for two reasons: *(i)* AirBnB has been widely adopted for several years; this will allow us to perform a comparative study at different times; and *(ii)* Airbnb is one of the most studied sharing economy platforms; this will allow us to leverage the existing literature to better understand our results. Overall we make the following main contributions:

1. We gathered, and subsequently analysed, a dataset of 49k AirBnB guests' reviews, and 22k hosts' reviews, produced in the U.S. over a period of 6 years of activity.

2. We systematically study the collected guests' and hosts' reviews in the U.S. per year of activity[2].

3. We find that utilitarian values (e.g., properties' facilities, convenience of location, business conduct) are discussed much more frequently than social values (e.g., guest/host interactions) and, more interestingly, this gap substantially increases over time.

By understanding how peers think about these novel sharing economy markets (e.g., convenience of location *vs.* fostering a guest/host social relationship), at different points in time, companies can tailor and market their services so to drive successful business models (e.g., highlighting the convenience and cost-effectiveness of access, as opposed to the financial obligations of ownership, and/or the emotional obligations of sharing).

[1]http://www.thepeoplewhoshare.com/

[2]Our lexicon is available on the project's site http://goodcitylife.org/airbnb/

Related Work

Sharing economy platforms like AirBnB have been studied by the academic community following two main lines of enquiries: understanding the impact of such platforms on related industries, and understanding motivation of both hosts and guests for participating.

Impact. Scholars from Business Studies and Economics have extensively analysed the relationship between sharing economy services and their regulated counterparts (e.g., Uber *vs.* taxis, AirBnB *vs.* hotels), to shed light onto the increasing accusations about the former being predatory and exploitative, causing severe externalities onto the latter. For example, (Zervas, Proserpio, and Byers 2016; 2015) focused their analysis on Texas and showed that indeed an increase in the number of AirBnB listings in this city leads to a decrease in the number of monthly hotel room accommodations. Other scholars have looked more broadly at the impact that sharing economy services have had on the economy. For example, (Fang, Ye, and Law 2016) found that such platforms have brought benefits to the broader tourism industry, since a reduction in accommodation prices has increased the number of tourists, and that has been driving the opening of many new job positions in related industries (e.g., restaurants). (Quattrone et al. 2016) conducted a temporal analysis on the growth of Airbnb in London, UK from 2012 to 2015. They found that Airbnb properties and hotels are not located in exactly the same urban areas: while the latter are mostly present in the city center and close to tourist attractions, the former are increasingly spreading out to areas where low and middle income people reside. These results collectively point to the fact that Airbnb and hotel industries are two distinct services, although they do no shed light onto what motivates customers to gravitate around one or the other.

Motivation. A parallel stream of works has looked into motivational factors behind sharing economy platforms' uptake. For example, (Varma et al. 2016) conducted an online survey of both Airbnb customers and traditional hotels' ones. They found that the two are quite distinct in terms of their needs and motivation. (Satama and others 2014) focused on sharing economy participants, and Airbnb hosts in particular, and delved further into their motivation to take part in these markets. By means of an online survey distributed via social media, they discovered that primary drivers of participations were convenience (e.g., speed and ease of finding accommodation) and hedonism (e.g., having fun); social influence (e.g., perception that others have of AirBnB users) was also mentioned, though to a lesser extent. (Ikkala and Lampinen 2015) focused on AirBnB hosts, and used in-depth interviews to elicit their motivation. Findings for the city of Helsinki revealed that money was often the initial driver; however, over time, social factors started to gain importance. A follow-up study (Lampinen and Cheshire 2016) conducted with 12 host interviewees now based in San Francisco, revealed that the financial benefits of hosting were indeed not a key factor, while social exchange and interpersonal interactions were. These findings suggest that, while hotels are business structures only, AirBnB is a platform where business and social factors come together. For platform owners to act upon these findings, they would need to gain confidence in their validity, and assess to what extent they hold across a larger sample of users and over time. To do so, a methodologically different approach is needed. In this paper, we propose one based on the linguistic analysis of the reviews that peers leave to one another.

Dataset

We selected Airbnb users at random for six months around the globe, and crawled their reviews and corresponding listings. In so doing, we obtained 282k Airbnb reviews: 203k guests' reviews, and 79k hosts'. We removed those that were too long (with more than 150 words) or too short (with less than 3 words). Since we were interested in a large scale geographic analysis of Airbnb, we needed to determine for which countries we had enough reviews. We decided to analyse one country only, the one producing the highest number of reviews. This country resulted to be the U.S. (49k reviews written by guests and 22k reviews written by hosts). We inferred the language of these reviews with the R package *textcat* and the Rosette Text Analytics API[3]; 99% of collected reviews resulted to be written in English.

Method & Results

Before starting our analysis, we built a dictionary of "topical" words (Task 1) that was then used to parse and analyse Airbnb reviews over six years (Task 2).

Task 1. Building a Dictionary

We built our dictionary in four steps: we developed a coding scheme by qualitatively analysing a random *sample* of Airbnb reviews (step 1), validated the coding scheme with a crowd-sourcing study on Crowdflower (step 2), labeled a larger set of reviews with another crowd-sourcing study on Crowdflower (step 3), and built a dictionary that consisted of the labels derived in the previous step (step 4).

Step 1. Developing a Coding Scheme. We performed a thematic analysis on 100 random Airbnb reviews. To ensure a representative sample of reviews across years, the 100 reviews were selected through stratified sampling. In a way similar to (Braun and Clarke 2006), two independent annotators coded these 100 reviews by performing three steps: *(i)* familiarising with the data, *(ii)* generating the initial codes and searching for themes among codes, and *(iii)* defining themes. After a full round of coding, the two coders compared their results, and agreed on which themes to remove, amend, or merge. As a result, they agreed on having five main themes. These were named 'Property', 'Location', 'Professional Conduct', 'Personality', and 'Social Interaction'.

---

[3]https://www.rosette.com/

Category	Description	Example
Property	Covering topics related to the property/house/room such as its cleanness, its decoration, size, or furniture.	"the room was bright and clean"
Location	Covering topics related to the neighbourhood where the property is situated such as nearby viewpoints, restaurants, metro stations, and so forth.	"the area was full of nice restaurants"
Professional Conduct	Covering topics related to the professional conduct of guest/host such as how good their communication has been, whether the guest has left the room clean and/or she has respected the house rules, or whether the host has been flexible with check-in/check-out times.	"she left the room clean and tidy"
Social Interaction	Covering topics related to the personality of guest/host or about guest/host spending social time together.	"I enjoyed chatting with her"

Table 1: Our four themes

Step 2. Validating a Coding Scheme. To ascertain the effectiveness of coding reviews with those five themes, we asked crowd-workers to annotate a new sample of 100 Airbnb reviews. To that end, we prepared a Crowdflower page[4] that consisted of three sections: *(i)* a list that showed our five themes; *(ii)* label examples of real Airbnb reviews (done by us); and *(iii)* the Airbnb reviews to be labelled. We paid 0.05$ per annotation, and each Airbnb review was independently annotated by at least four different workers. We computed the Fleiss' kappa agreement score for the five themes (Fleiss 1971), and two of them had a Fleiss' kappa score less than 0.5. We merged these two into one theme, resulting in four main themes: 'Property', 'Location', 'Professional Conduct' and 'Social Interaction' (Table 1). To ascertain the effectiveness of coding with those four themes, we again asked crowd-workers to annotate a new sample of 100 Airbnb reviews. All the four themes resulted into a Fleiss' kappa score higher than than 0.5, suggesting their validity.

Step 3. Labelling Reviews. We were then ready to label a larger set of Airbnb reviews using the newly found four themes. With Crowdflower, 600 reviews were annotated, each by at least four workers. We removed those on which annotators disagreed, and ended up with 550 labelled reviews. As one expects, the most popular theme was 'Professional Conduct', followed by 'Property' and 'Location'; and 'Social Interaction' was the least frequent theme.

Step 4. Building a Dictionary. Finally, with each of the four main themes, we needed to associate representative words (n-grams). We did so in a data-driven fashion in three steps. First, for each theme c, we split the 550 annotated reviews into two sets: Set_c and $Set_{\bar{c}}$. Set_c is the set of reviews

---

[4]Crowdflower is a crowd-sourced market of online workforce to clean, label and enrich data: https://www.crowdflower.com/.
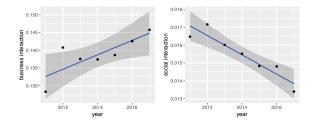

Figure 1: Fraction of mentions of business *vs.* social interactions in *guests*' reviews

labelled with theme c by at least three quarter of workers; while $Set_{\bar{c}}$ is the set of reviews labelled with theme c by no more than one quarter of workers. Second, we extracted all n-grams from Set_c and $Set_{\bar{c}}$, with $n = \{1, 2, 3\}$. For each n-gram t, we computed two measures $tf(t,c)$ and $tf(t,\bar{c})$, denoting the term frequency of t in Set_c and $Set_{\bar{c}}$, respectively. Finally, we computed $tf_{gain}(t,c) = \frac{tf(t,c)}{tf(t,\bar{c})}$, and, with each theme c, we associated all the n-gram t such that $tf(t,c) > th_{tf}$ and $tf_{gain}(t,c) > th_{gain}$, with $th_{tf} > 0$ and $th_{gain} > 1$. The use of the first threshold th_{tf} removed unpopular n-grams. The use of the second threshold th_{gain} included only the n-grams that were comparatively more popular in Set_c than in $Set_{\bar{c}}$. The values of th_{tf} and th_{gain} that produced a list with the least number of noisy and unpopular n-grams were $th_{tf} = 0.02$ and $th_{gain} = 3$. The resulting dictionary contained words arranged in the four themes.

Task 2. Analysing Airbnb Language

As a second and final step, we grouped Airbnb reviews by year and, for each group of reviews, we computed the fraction of n-grams in theme c:

$$adoption(c, text) = \frac{\sum_{t \in c} freq(t, text)}{\sum_{t \in v} freq(t, text)}$$

where $freq(t, text)$ is the number of occurrences of n-gram t in $text$, and v is the set of distinct n-grams in $text$. We identify two main types of interactions: the first is called 'business interaction' and groups our 'Property', 'Location', and 'Professional Conduct' themes together; and the second is called 'social interaction' and reflects our 'Social Interaction' theme.

For each of the two types of interactions, the fraction of n-grams (what we call adoption) in guests' reviews (Figure 1), and those in hosts' reviews (Figure 2) show strikingly consistent patterns. Reviews concerning business interactions (i.e., those mentioning property-related words and professional conduct) have been increasing over the years. Conversely, mentions of social interactions have been rarer and, interestingly, decreasing over time in both guests' and hosts' reviews.

To cross validate our results, we considered another way of coding reviews that relies on the Linguistic Inquiry and Word Count (LIWC) dictionary. LIWC is a popular dictionary developed over the last few decades that includes 64
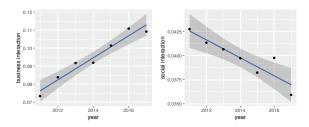
Figure 2: Fraction of mentions of business *vs.* social interactions in *hosts*' reviews
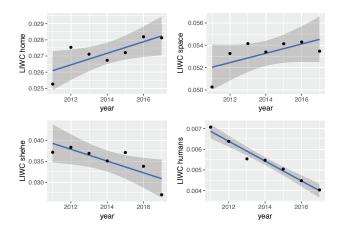

Figure 3: Fraction of mentions of *LIWC* categories in *hosts*' reviews (the same trend applies to guest reviews)

language categories, ranging from part-of-speech (e.g., articles or personal pronouns) to topical categories (e.g., family or money) (Tausczik and Pennebaker 2010). There are four categories that are related to our four themes the most and are:

- *Home*. This category includes words such as 'kitchen', 'bath', 'bed' and, as such, partly overlaps with our 'Property' category.

- *Space*. It includes words such as 'area', 'at', 'away', and partly overlaps with our 'Location' category.

- *She/he*. It includes declinations of the third-person singular pronouns 'she' or 'he'. In conversations, these have been found to indicate focus on people. On Airbnb reviews, they could indicate a social use of the platform.

- *Humans*. It includes words such as 'boy', 'girl', 'baby' and, on Airbnb reviews, might suggest mentions of interpersonal relationships and experiences.

LIWC analysis (Figure 3) supports our earlier findings: the use of words in the categories 'Home' and 'Space' increased over the years, while the use of words in the social categories of 'She/he' and 'Humans' consistently diminished.

Conclusion

As one might expect, for Airbnb guests and hosts, business considerations are more salient than social ones. But, as one might not expect, that social-business gap has not remained constant over the years: it has significantly increased instead. These results offer the first quantitative evidence on which extent the sharing economy is really about sharing.

However, some words of caution are in order. First, we do not know whether certain socio-demographic groups are more or less prone to write reviews than others and, therefore, whether they tend to be over- or under-represented. Second, we have only analysed Airbnb reviews produced in the United States of America. As a result, our findings may not hold in other countries. Despite those limitations, our method is readily applicable to other platforms such as Couchsurfing and TaskRabbit. Also, it offers a variety of practical implications, ranging from globally tracking the "sense of community", to ranking properties and listings based on hosts' sociability.

References

Braun, V., and Clarke, V. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3(2):77–101.

Eckhardt, G. M., and Bardhi, F. 2015. The sharing economy isnt about sharing at all. *Harvard business review* 28(01):2015.

Fang, B.; Ye, Q.; and Law, R. 2016. Effect of sharing economy on tourism industry employment. *Annals of Tourism Research* 57:264–267.

Fleiss, J. L. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin* 76(5):378.

Ikkala, T., and Lampinen, A. 2015. Monetizing network hospitality: Hospitality and sociability in the context of airbnb. In *Proc. of CSCW '15*, 1033–1044. ACM.

Lampinen, A., and Cheshire, C. 2016. Hosting via airbnb: Motivations and financial assurances in monetized network hospitality. In *Proc. of CHI '16*, 1669–1680. ACM.

Quattrone, G.; Proserpio, D.; Quercia, D.; Capra, L.; and Musolesi, M. 2016. Who benefits from the sharing economy of airbnb? In *Proc. of WWW '16*, 1385–1394.

Satama, S., et al. 2014. Consumer adoption of access-based consumption services-case airbnb.

Tausczik, Y. R., and Pennebaker, J. W. 2010. The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of language and social psychology* 29(1):24–54.

Varma, A.; Jukic, N.; Pestek, A.; Shultz, C. J.; and Nestorov, S. 2016. Airbnb: Exciting innovation or passing fad? *Tourism Management Perspectives* 20:228–237.

Zervas, G.; Proserpio, D.; and Byers, J. W. 2015. The impact of the sharing economy on the hotel industry: Evidence from airbnb's entry into the texas market. In *Proc. of EC '15*, 637–637. ACM.

Zervas, G.; Proserpio, D.; and Byers, J. 2016. The rise of the sharing economy: Estimating the impact of airbnb on the hotel industry. *Boston U. School of Management Research Paper* (2013-16).