Culture and Social Influence: Evidence From Online Reviews*

Cayruã Chaves†

This version: February, 2023

Click here for the latest version

Abstract

Evidence from surveys and lab experiments suggests that people's propensity to conform to the opinion of others is lower in more individualistic cultures. Do these findings hold in real-world settings? This paper quantifies the role of culture as a determinant of social influence in the context of online consumer reviews. Exploiting discontinuities in the way Tripadvisor displays average ratings, I estimate how reviewers from different countries respond to the average opinion of past consumers. A discontinuous increase of 0.5 stars in a restaurant's average rating leads reviewers from countries with the least individualistic cultures to report ratings that are 0.1 stars higher. The size of the effect reduces in individualism and becomes statistically insignificant for consumers from the most individualistic cultures. The negative relationship between individualism and reviewers' tendency to conform cannot be explained by country-level predictors of individualism, such as income or religion. Moreover, cross-regional variation within Italy reveals that the correlation between cultural values and social influence also holds across reviewers from different regions within the same country. These findings imply that average ratings converge faster to firms' real quality when reviewers are from more individualistic cultures.

JEL Codes: D83, D90, Z13.

Keywords: culture, individualism-collectivism, social influence, online reviews, information.

* I am grateful to Gerard Llobet for his guidance and support throughout the project. I also thank Pedro Mira, Diego Puga, Guillermo Caruana, Dmitry Arkhangelsky, Manuel Arellano, and participants of CEMFI's Firms and Markets Workshop for helpful comments and discussions. I gratefully acknowledge funding from Spain's Ministerio de Ciencia, Innovación y Universidades (PRE2018-084485) and Ministerio de Economía, Industria y Competitividad (María de Maeztu Programme for Units of Excellence in R&D, MDM-2016-0684).

† PhD Candidate, CEMFI. E-mail address: cayrua.chaves@cemfi.edu.es. Website: cayruachaves.github.io

1 Introduction

Social influence, the process by which a person's actions change due to what others do, is a critical force driving human behavior.[1] Moreover, the degree to which social influence affects individual choices depends on the cultural norms people experience growing up. Large-scale surveys reveal substantial cross-country heterogeneity in the importance people give to making their own independent choices versus following group norms (Schulz et al., 2019). Furthermore, lab experiments suggest that the propensity to conform to the majority opinion varies greatly across populations from different countries (Bond and Smith, 1996). Do these findings hold in real-world settings? Can we quantify the role of culture in shaping individuals' propensity to be influenced by others?

This paper studies the relationship between culture and social influence in a specific setting, online consumer reviews. Using data on reviews for restaurants listed on Tripadvisor, one of the world's leading travel and food services platforms, I address the following question. When submitting their own reviews, how are consumers from different cultures influenced by the average opinion of previous reviewers? I estimate the effect of restaurants' average ratings on the rating reported by its next reviewer by exploring discontinuities in how Tripadvisor displays information. Moreover, to evaluate the role of culture, I take advantage of Tripadvisor's international pool of reviewers and test the degree to which a country-level measure of individualism explains variation in social influence. To understand the psychological mechanism underlying social influence, I test for the effect of alternative cultural measures and explore variation across psychologically diverse regions within Italy.

Tripadvisor is an online travel company that manages a website and mobile application featuring over 1 billion user-generated reviews, mostly of restaurants and hotels. This paper focuses on reviews submitted to restaurants located across 105 countries worldwide. Tripadvisor provides an adequate setting to study cross-cultural variation in social influence for three reasons. First, the platform displays a restaurant's average rating over its previous reviews in a salient manner. New reviewers, who are likely to perceive these average ratings as the majority's opinion, may be influenced by them when submitting reviews of their own.

Second, Tripadvisor displays average ratings rounded to the nearest half-star.[2] This feature allows me to compare the rating behavior of consumers who experienced restaurants with similar levels of quality but with a 0.5 stars difference in average ratings. Third, the fact that Tripadvisor is the second most visited website worldwide in travel and tourism means that it has a culturally diverse user base.[3] Importantly, most user profiles contain information on the country where they live, which allows me to link reviewers to a given cultural background.

I quantify the level of social influence on Tripadvisor by estimating the effect of a restaurant's overall average rating (i.e., the opinion of the majority) on the value of the next review it receives. To assess the role of culture, I allow this effect to depend on the reviewer's cultural background. My baseline analysis focuses on cross-country variation in individualism, a dimension of culture that may be particularly relevant for social influence. Individualism

[1] For a broad discussion on the power of social influence in driving human behavior, see Frank (2021).

[2] See Appendix A for a more detailed description of Tripadvisor's interface.

[3] Source: www.similarweb.com/top-websites/category/travel-and-tourism.

measures the extent to which people perceive themselves as independent agents as opposed to interdependent members of larger groups. In individualistic cultures, people give particular importance to personal freedom and achievement, while cultures with low levels of individualism emphasize social harmony and conformity. Specifically, I use the country-level measure of individualism constructed by Hofstede (2001), which places countries on a scale between 0 (least individualistic) and 1 (most individualistic).[4]

Estimating the effect of a restaurant's average rating on subsequent reviewers is not a trivial task. Average ratings are endogenous because they are positively correlated with restaurant quality and, thus, with the expected value of its future ratings. To deal with this issue, I exploit discontinuities in how Tripadvisor displays information. Tripadvisor displays the averages of past reviews for each restaurant, rounded to their nearest half-star. Thus, restaurants with similar underlying (continuous) averages may end up with headline star ratings that differ by 0.5 stars. Under the assumption that restaurant quality is a smooth function of the underlying average rating, this empirical strategy deals with the fact that high-quality restaurants mechanically get higher ratings.

A second challenge related to identifying the causal effect of average ratings on the review reported by the next user concerns reviewer selection. For instance, suppose consumers decide which restaurants to visit based on their idiosyncratic tastes/information and restaurants' star ratings displayed on Tripadvisor. In that case, reviewer unobservable characteristics could change right at the rounding thresholds. This kind of selection, although theoretically possible, does not seem to be a first-order issue in my setting. Reviewer-level observable characteristics that should correlate with one's propensity to use star ratings as an input into the consumption decision exhibit no discontinuities around rounding cutoffs.

Assuming that neither reviewer nor restaurant characteristics change discontinuously at Tripadvisor rounding cutoffs, my estimates can be interpreted as the causal effect of average ratings on the review reported by the next reviewer. That is my measure of social influence. Reviewers from the least individualistic cultures respond to a 0.5 stars increase in a restaurant's headline average rating by increasing their own reviews by 0.1 stars. This result means that for these reviewers, there is evidence of conformity, the adjustment of one's review in the direction of the opinion held by the majority. Moreover, reviewers' response to average ratings depends on individualism in a statistically significant way. The conformity effect described above decreases with individualism and shrinks to zero for reviewers from the most individualistic countries.

Next, in exploring the heterogeneity of these effects, I show that conformity increases in the number of previous reviews held by the restaurant being rated. Moreover, I do not find evidence that the magnitude of this type of heterogeneity changes depending on cultural background. The dependence of social influence on restaurants' number of previous reviews may be interpreted as evidence that the opinion of others affects one's own judgment of the restaurant experience. An alternative explanation is that people report a rating that differs from their private opinion to avoid the (potential) psychological costs of looking too different from

[4]Hofstede (2001) uses the term "collectivist" to refer to the low end of the individualism scale.

others.[5] Under both interpretations, conformity is expected to increase in the size of the group whose opinions individuals are exposed to.

I also explore the robustness of individualism as a predictor of cross-country variation in social influence. By interacting fixed effects for the reviewer's country with an indicator variable for whether the average rating she faces is rounded up, I estimate a separate social influence effect for each country. Next, I test the ability of different country-level variables to predict variation in estimated social influence. First, country differences in reviewer or restaurant characteristics predict very little of the variation in estimated social influence. Second, I show that the effect of individualism on reviewers' tendency to conform is robust to controlling for country-level predictors of this cultural trait, such as income or the share of the Protestants in the population.

There is something particular to the individualism dimension of culture that predicts reviewers' lower propensity to conform. Two types of evidence provide support for this claim. First, still using the previously mentioned country-level social influence effects, I show that the relationship between social influence and individualism is not due to how Hofstede (2001) classified countries. I show that Schwartz (1994) measures of affective and intellectual autonomy, which capture concepts akin to individualism, also predict lower country-level conformity. Second, none of the other three core dimensions of culture in the original model of Hofstede (2001) predict social influence.

Finally, to mitigate concerns that some underlying unobserved variable generates a spurious the correlation between country-level individualism and conformity to the average on Tripadvisor, I explore a sample restricted to Italian reviewers. Using within-Italy cross-regional variation in cultural values from the European Social Survey (ESS), I show that culture also matters within a country. I use the same strategy based on the rounding of average ratings to estimate region-level social influence effects. These estimates correlate positively with region-level average answers to ESS questions that measure people's emphasis on the values of conformity and obedience. Given that the conformity index from the ESS measures values emphasized by cultures characterized by low levels of individualism, this is the kind of correlation we should expect.

I focus on a single channel of social influence, the extent to which reviewers are affected by the current average rating. However, social influence might also operate through other channels. For example, reviewers may be directly influenced by recent reviews, which are also easily visible on Tripadvisor. Alternatively, subsequent reviewers may also be influenced by the text content of other consumers' reviews. I chose to focus on the star rating because it is a critical summary measure commonly used by reviewers, and because of Tripadvisor's rounding policy, which provides quasi-random variation in the perceived majority opinion about restaurants of very similar quality.

My findings have implications for the speed of information flow in online review platforms. They imply that when a significant fraction of users come from cultures with low levels of individualism, average ratings will take longer to converge to firms' real quality. More broadly,

[5]Some people may derive psychological gains from deviating from the majority opinion. I emphasize the potential costs of speaking one's mind because empirically, I observe more conforming than deviating behavior.

this paper suggests that cultural factors should be considered in other situations where social influence plays a relevant role. For example, releasing surveys with vote intentions during political campaigns may affect final election outcomes more strongly in societies with less individualistic values. A second example relates to public efforts to induce behavioral change, such as reducing smoking prevalence. My findings indicate that governmental efforts may have higher "multiplier" effects in less individualistic cultures, where each person who decides to stop smoking may significantly influence the probability that the next person will behave the same way.

Related Literature. Theoretical studies of how individual behavior is affected by what others do or think feature in the economics literature at least since Akerlof (1980). Other classic examples are Becker (1991), Banerjee (1992), and Bernheim (1994). Each of these authors focuses on a different reason why people respond to the choices made by others. Commonly discussed reasons include information, norms, and image-related concerns. None of these studies, however, features a direct role of culture as a factor that can shape people's propensity to follow (or deviate from) others.

Theories of why people's inclination to rely on the behavior of others varies across cultures are based on evolutionary arguments. This literature argues that differences in environmental and social factors in the distant past generated evolutionary pressures which were society-specific (Boyd and Richerson, 1988; Henrich, 2020), which in turn implies that important psychological traits have evolved differently across societies. One such trait is people's tendency to conform to the opinions and actions of others when making their own decisions.

Empirical evidence on the association between culture and social influence mainly comes from lab experiments and surveys. Data from the World Value Survey or the European Social Survey shows that the importance people assign to values such as independence and autonomy (or, in the opposite direction, tradition, conformity, and obedience) varies substantially across countries (Schulz et al., 2019). These answers strongly correlate with the Hofstede (2001) individualism measure, which also comes from cross-country surveys. Concerning lab experiments, the classic conformity study conducted by Asch (1956) was replicated in several countries.[6] In a meta-analysis, Bond and Smith (1996) found that conformity effects were lower for subjects from more individualistic cultures. I contribute to this discussion by showing that culture in general, and individualism in particular, is a crucial moderator of social influence effects in a real-world setting, online review platforms.

My paper also relates to the large literature on online reviews. Within economics, most papers have focused on the consequences of online review platforms for the units being rated. For example, both Anderson and Magruder (2012), Luca (2016), and Anenberg et al. (2022) use a similar regression discontinuity design to study different restaurant outcomes. The variation

[6] Asch (1956) original experiment attempted to measure the extent to which people are willing to change their judgment to agree with judgments made by others. Participants are asked to match one of three line segments to a target line based on length. They carry out the task under two conditions, alone or with others participants around them. Other participants are actors working for the researcher and giving the wrong answers purposefully. Conformity effects are significant. The task was designed to be easy, and when alone, people respond correctly 99% of the time, but when they are together with other participants who gave the wrong answer, mistakes occur 37% of the time.

I use is similar to the one explored in these papers, but I look at its effect on the evolution of ratings themselves.

Other authors have also looked at the process of how ratings are generated. Acemoglu et al. (2017) develop a theoretical model to study how platform design affects the speed of learning from online reviews. My results indicate that the degree of applicability of their model, which does not feature social influence, depends on the reviewers' cultural background. Concerning empirical studies, social influence effects have been quantified in different ways. For example, Dai et al. (2018) estimated a rating choice model where one of the parameters could be interpreted as arising from social influence. However, our approaches differ in meaningful ways. Their focus is on optimal rating aggregation. Thus they include a large set of controls and use the full extent of variation in ratings. In contrast, my goal is to uncover the effect of a specific parameter (social influence). Thus, I estimate a more parsimonious model (fewer covariates) and use variation exclusively from Tripadvisor's rounding rule, which provides a plausibly exogenous source variation in average ratings for restaurants of similar quality.

One paper focusing exclusively on the question of social influence in online settings is Muchnik et al. (2013). They conduct a field experiment with an actual website and obtain credible estimates of social influence effects. Three main aspects distinguish our papers. First, their study is not concerned with cultural differences. Second, the kind of social influence we study differs. While I focus on a single channel, the effect of the current average rating on the next rating, they estimate social influence effects that accumulate over time in a sequence of user ratings.[7] Lastly, instead of a field experiment with truly randomized treatments, I use a natural feature of Tripadvisor, which generates variation in users' perception of the average opinion of others conditional on restaurants' underlying quality. Although their approach delivers clean causal estimates, my setup is easily transferable and could be applied to study other platforms which display rounded averages, such as Yelp or Amazon.

Finally, my paper also relates to the literature exploring the role of cultural factors in explaining consumer ratings. A study of particular relevance is Hong et al. (2016), which documents a positive correlation between a individualism in the consumer's country and the absolute distance between the review she reports and the current average rating held by a restaurant. The authors interpret this as evidence of cross-cultural variation in conformity. I complement their study by providing causal estimates of culture-specific social influence in this context.

Roadmap. The rest of the paper is organized as follows. Section 2 describes the data. In Section 3, I present my baseline empirical strategy and results on the impact of individualism on reviewers' tendency to conform to the average of past reviews. Section 4 investigates other dimensions of cross-country heterogeneity that could explain variation in social influence. Section 5 zooms into Italy and explores cross-regional variation in cultural values. I conclude in Section 6.

[7] The website they investigate is a news aggregator. Thus, user ratings refer to upvotes or downvotes to specific comments instead of ratings to firms.

2 Data

2.1 Tripadvisor

I extracted public-facing information directly from Tripadvisor's website. In order to obtain reviews submitted by consumers from several different cultures, I adopted the following procedure. First, I obtained a list of the 105 nations included in Hofstede's cultural dataset.[8] Second, for each country, I selected the city with the largest number of reviews on Tripadvisor.

My empirical strategy requires the reconstruction of average ratings that appeared on Tripadvisor at the time of each new review. Thus, I had to scrape restaurants' complete review histories. This process is time-consuming, so I restricted the analysis to a random sample of restaurants within each city. I obtained the history of reviews for a 6% random sample of the restaurants listed in each city. The data was collected using web scraping techniques and reflected Tripadvisor's content at the time each scraping was running. The process lasted between January and June of 2022.

All reviews include a numerical rating between 1 (terrible) and 5 (excellent), indicating consumers' assessment of their experience with the restaurant. Ratings constitute my primary outcome variable throughout the paper. Moreover, all reviews have a time stamp, allowing one to recover the average that appeared on Tripadvisor when each new consumer arrived to submit a rating of her own. Lastly, for most reviews (approximately 80%), I observe the country where users report to live in,[9] which I use to link them to an individualism score, taken from Hofstede (2001) and measured at the country level (more on this in the next section). The interaction of a reviewer's individualism score and the average rating she sees on Tripadvisor constitutes the primary explanatory variable in most of the analysis. Note that the high share of tourist reviews implies that the country where a reviewer lives often differs from the country where the restaurant is.

The raw data extracted from Tripadvisor has reviews from as early as 2004. To keep the level of individualism stable, I restrict the analysis to the period between 2015 and 2019. Before 2015, levels of individualism decreased over time because Tripadvisor was launched in the United States and quickly gained popularity in other relatively rich countries, which also scored high in individualism. By 2015, consumers from several countries were already using the platform, and levels of individualism stabilized. Regarding the years 2020 and 2021, due to the pandemic of COVID-19 and lockdown measures adopted worldwide, review activity dropped precipitously. Thus, I also exclude these years from the main analysis. Figure 11 in Appendix C shows the evolution of the total volume of reviews and average level of individualism in the raw sample.

The second choice I make in defining the final sample relates to the frequency with which reviewers use Tripadvisor. I focus on frequent reviewers, whom I define as those with more

[8] Hofstede's data is available at http://www.geert-hofstede.com/. The most modern version of the data includes six dimensions (i.e., variables) of culture, each capturing a different aspect of cross-country variation in cultural values. For 105 countries, there is at least one cultural dimension available. The number of countries with a valid measure of individualism is slightly lower, 97.

[9] When creating a profile on Tripadvisor, besides mandatory information such as an email address, users have the option to fill in a field called "add your current city". This field is not mandatory, and users can choose to leave it blank in their profile.

than 20 reviews submitted to Tripadvisor.[10] The reason to concentrate on frequent reviewers is to reduce the prevalence of under-reporting, consumers' tendency to post reviews exclusively in extraordinarily positive or negative experiences (Hu et al., 2009). My main goal is to test whether the average rating displayed on Tripadvisor influences reviewers' rating choices. Thus consumers who only use the platform to submit a specific (ex-ante chosen) rating are not my main population of interest.[11]. In Section 3.3.3, I provide further discussion regarding the differences between casual and frequent reviewers.

Lastly, in getting to my final sample, I implement a few cleaning procedures to guarantee the data's quality. First, for each restaurant, I only included reviews that are at least two days apart from each other. Two days is the approximate time between the moment a consumer submits a new review and when Tripadvisor effectively uploads it to its webpage.[12] For example, suppose there are two or more days between the submission dates of review n and $n+1$ (given to the same restaurant). In this case, I can be confident that the average seen by $n+1$ contained the rating submitted by n. However, had these two reviews been submitted in two consecutive days, I would not have known whether to include n in the average seen by $n+1$. I drop such cases to be on the safe side and avoid introducing measurement error in the explanatory variable.

Second, I only include ratings submitted when the restaurant had more than ten prior reviews. Before accumulating a certain number of reviews, a restaurant's average is too noisy and may not be a good proxy for quality. Moreover, the likelihood of a fake review, such as when someone submits a good rating to help a restaurant owned by a friend or family member, is higher when the restaurant is recently opened. Lastly, I drop reviews submitted by users who do not inform where they live and by those who report a location not covered by Hofstede (2001) data on culture (this step only removes 0.7% of observations). Unless otherwise specified, all analyses use this final sample.

2.2 Individualism

My second data source is Hofstede (2001), who created a framework to measure cross-country variation in different cultural dimensions. Mainly, my focus is on one of these dimensions, the collectivism-individualism dimension of culture. Hofstede's individualism score is a measure that places countries between 0 (most collectivist) and 1 (most individualistic).[13] Individualism captures the extent to which people perceive themselves as independent agents (high score) as opposed to interdependent members of larger groups (low score). People from countries with individualism scores place a particular value on freedom and status, while in countries with low scores, the local culture emphasizes social harmony and conformity.

[10]The number of reviews I use to classify reviewers as "frequent" does not refer to reviews in my sample. Instead, it refers to the total number of reviews a user has ever submitted to Tripadvisor. I have access to this variable because, for each review, I observe a summary of the Tripadvisor activity of the user who submitted it. This summary includes the total number of reviews posted on the platform.

[11]For example, the high share of 5-star ratings (over 50%) suggests that some people only submit 5-star reviews, choosing not to submit at all when their level of satisfaction was anything below 5-stars

[12]Tripadvisor's support page informs that most reviews are posted within 24-48 hours. In exceptional cases, when there is a potential conflict with Tripadvisor's review guidelines, processing and publishing a review may take longer.

[13]Hofstede (2001) original work computes the individualism scores on a scale between 0 and 100. I use a 0 to 1 scale to facilitate the interpretation of regression coefficients.

Figure 1: World Map of Individualism Scores

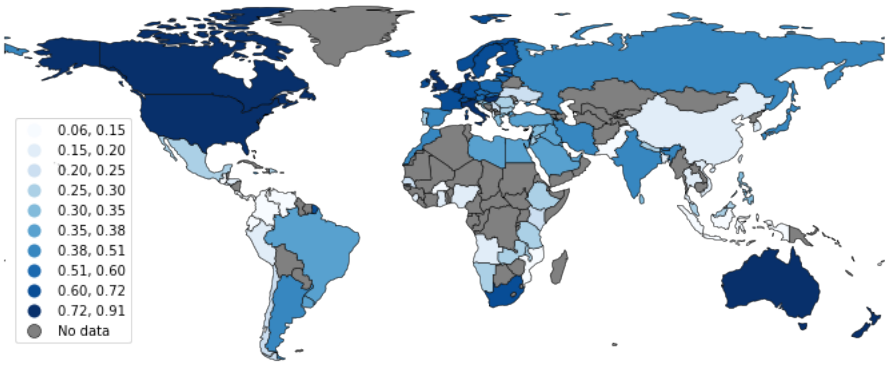

Notes: Individualism score for 97 countries. Source is Hofstede (2001)

Hofstede (2001) constructed his data from surveys of IBM employees that originally included respondents from 40 countries. Over time, he and his collaborators extended the data to almost 100 nations. He computed individualism scores by applying factor analysis to country differences in answers to several questions aimed at measuring values and work goals. The latent factor corresponding to individualism correlates positively with answers indicating a high value placed on personal achievement, individual freedom, and challenging/fulfilling work. On the other hand, it correlates negatively with answers to questions measuring the importance of harmony and cooperation between coworkers and of maintaining good relations with superiors.[14]

Figure 1 shows a world map of individualism. Darker shades of blue indicate a higher individualism score (grey indicates no data is available for that country). According to this measure, the Anglo-Saxon world stands out as particularly individualistic, with the United States, Australia, and the United Kingdom having the three highest scores. Western Europe, excluding Portugal and Spain, is also rather individualistic. On the other hand, the collectivist end of Hofstede (2001) index includes different countries in Asia, Africa, and Latin America, which hosts the three least individualistic in the world, Guatemala, Ecuador, and Panama.

2.3 Summary Statistics

Table 1 reports basic statistics of the main sample used throughout this paper. There is a total of 369,457 reviews, submitted by 265,394 different reviewers, to 10,282 unique restaurants. Most ratings are high, the average being 4.07 and the share of 5-star ratings being 40%. Regarding the spatial variation of the data, there are restaurants in 105 different countries and reviewers from 97 unique nations.[15] The distribution of reviews along consumers' country of origin is skewed towards countries with high levels of individualism. Across all reviews, the average individualism score is 0.64. Furthermore, the UK and the US, the third and the most individualistic countries, respectively, are the two countries with the largest number of reviews in the sample.

[14]Appendix B includes more details on the questionnaire used by Hofstede (2001).

[15]To avoid repetitive language, I will often refer to the reviewer's country as just country. When there is a need to discuss something related to the country where a restaurant is, which occurs less often, I will make a split reference to it using an expression like restaurant country.

Among the six most frequent countries, only one, Brazil, has an individualism score below 0.50, the midpoint of Hofstede (2001) scale.[16]

Table 1: Summary Statistics at the Review Level

Reported Ratings			Most Frequent Countries (of Reviewers)			
	Count	Share of Total		Count	Share of Total	IDV
Total	369,457	100%	United Kingdom	50,597	13.69%	0.89
1-star	11,496	3.11%	United States	37,211	10.07%	0.91
2-star	17,427	4.72%	Italy	35,915	9.72%	0.76
3-star	53,543	14.49%	France	27,371	7.41%	0.71
4-star	138,934	37.60%	Spain	21,922	5.93%	0.51
5-star	148,057	40.07%	Brazil	19,229	5.20%	0.38
			Australia	10,689	2.89%	0.90
	Mean	Standard Dev.	Russia	10,591	2.87%	0.39
Reported Rating$_n$	4.07	1.00	Germany	9,905	2.68%	0.67
Average Stars$_{n-1}$	4.08	0.42	Canada	8,884	2.40%	0.80
Individualism (IDV)	0.64	0.23	Argentina	8,072	2.18%	0.46
Num. Prior Reviews	308.15	486.48	Japan	7,809	2.11%	0.46
Tourist Dummy	0.54	0.50				
Unique reviewers	265,394					
Unique restaurants	10,282					
Unique countries (reviewers)	97					
Unique countries (restaurants)	105					

3 Baseline Analysis

3.1 Empirical Model

In this section, I test whether the average rating held by a restaurant affects the rating reported by its next reviewer. In particular, I estimate whether this effect varies systematically with the individualism score in the reviewer's country. I implement a regression discontinuity design that compares new ratings submitted to restaurants holding similar underlying (continuous) average ratings but with different headline stars prominently displayed on Tripadvisor.

The discontinuity in restaurants' star ratings arises because Tripadvisor displays averages rounded to their nearest half-star. For example, an underlying average rating of 4.24 appears as a 4.0-star headline average, while an average of 4.25 is displayed as a 4.5-star headline average. New reviewers easily observe headline stars but are unlikely to compute the underlying averages. In that sense, reviewers perceive headline stars as the average opinion of other consumers. Thus, this section's primary goal is to quantify the effect of the perceived majority opinion on the rating reported by a reviewer, conditional on the level of individualism in her culture.

A regression discontinuity model with additional interaction terms constitutes my baseline specification. The most important of these terms is an interaction between the discontinuity indicator and reviewers' individualism score. The second type of interaction term, between individualism and the running variable (continuous average), is only used in some specifications.

[16] Figure 12 in Appendix C shows the full distribution of individualism scores in the data.

It controls for the possibility that reviewer response to restaurant quality (proxied by continuous averages) may be culture-specific.[17] Formally, I estimate the following model:

$$\begin{aligned} Rat_{ijn} = \beta_0 + \left(\beta_1 + \beta_2\, Idv_{c(i)}\right) \times \mathbb{1}\{Avg_{jn-1} \geq k\} \;\; + \beta_3\, Idv_{c(i)} + \\ \beta_4 Avg_{jn-1} + \beta_5\left(Idv_{c(i)} \times Avg_{jn-1}\right) + X'_{ijn}\gamma + \epsilon_{ijn} \end{aligned} \tag{1}$$

Where k is a given rounding cutoff (e.g., 3.75, 4.25, 4.75). The dependent variable Rat_{ijn} is a rating written by reviewer i, to restaurant j receiving its n^{th} review. Avg_{jn-1} is the (continuous) average over previous ratings received by j, and $Idv_{c(i)}$ is the level of individualism in reviewer i's country. The subscript $c(i)$ explicitly indicates that the individualism score is a measure that varies with country c, where reviewer i lives.

The vector X_{ijn} includes a set of covariates that may affect ratings. These vary depending on the specification. Some examples are the number of prior reviews held by the restaurant, dummies for whether the consumer is a tourist or a local, and restaurant price indicators. The vector of controls also includes different sets of fixed effects. All specifications include time trends specific to the restaurant's city.[18] Additionally, some specifications include restaurant-fixed effects, which means I will use within-restaurant variation in ratings received on different sides of a rounding threshold.

The main coefficients of interest are β_1 and β_2. Together they determine the effect of a half-star increase in a restaurant's displayed average on the rating reported by consumers from a country with an individualism score of Idv_c. Under this specification, social influence is measured by $\beta_1 + \beta_2\, Idv_c$. Coefficient β_3, although not central to the analysis of social influence, is also informative about the role of culture in explaining consumers' rating behavior. It measures the correlation between individualism and reported ratings, controlling for restaurants' underlying average and whether Tripadvisor displays it rounded up or down.

Coefficient β_4 measures the linear dependence of new ratings on a restaurant's current underlying average. As a starting point, I estimated models where this linear effect does not change across sides of a rounding cutoff. There are two reasons for this choice. First, I tested the inclusion of an interaction term between Avg_{jn-1} and the rounding treatment (i.e., allowing the effect of the running variable to differ on each side of a cutoff), but this coefficient was never significant. Second, in some specifications I estimate slopes specific to each level of individualism (i.e., including the interaction term $Avg_{jn-1} \times Idv_c$). To keep the interpretation of the main effect of individualism reasonably simple, I prefer not to interact it with two different average ratings, below and above the cutoff. The results of this more flexible model, allowing for differential effects of the running variable for all combinations of cutoff side and reviewers' individualism score, are very similar to the ones presented here.

Before presenting the results of estimating Equation 1, I provide visual intuition on the kind of variation I am exploring. Figure 2 shows the distribution of average ratings, my running

[17] The motivation to test this type of specification arises from a simple rating choice model, where reviewers report a weighted average between their consumption utility (self-expression) and the star rating displayed on Tripadvisor (social influence). If one introduces culture-specific propensities to be influenced by others, this framework is consistent with culture-specific slopes of new ratings on the underlying average of past reviews.

[18] More specifically, I include fixed-effects for the interaction between period (month-year) dummies and dummies for the city where the restaurant is.

Figure 2: Distribution of Average Ratings (before each new review)

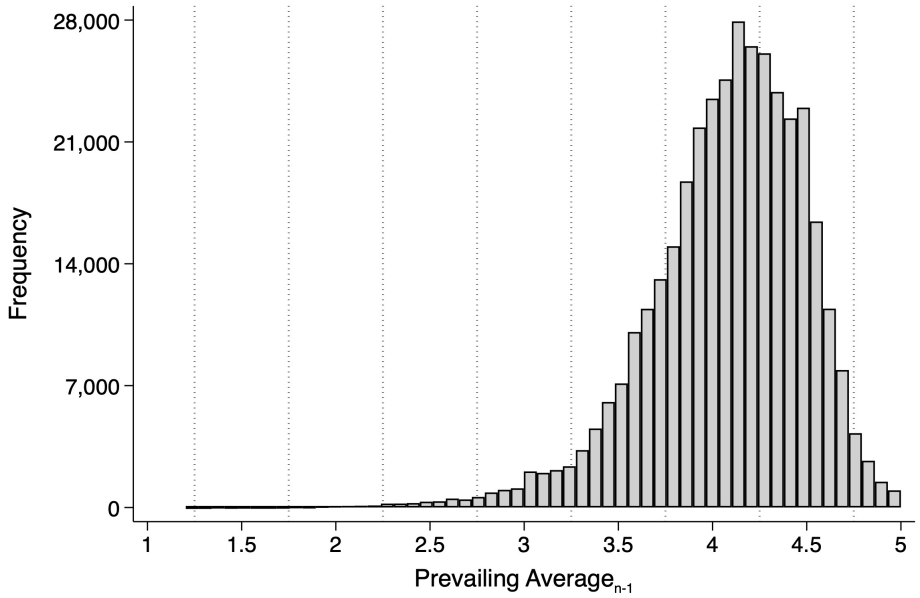

Notes: Vertical dotted lines indicate rounding cutoffs. Values in the x-axis with explicit labels (1, 1.5, 2...) are the headline averages displayed by Tripadvisor in between any given pair of rounding cutoffs.

variable. Specifically, for the n^{th} rating received by restaurant j, I compute Avg_{jn-1}, the average across all its $n-1$ preceding ratings. Avg_{jn-1} determines if the headline average shown to reviewer n is rounded up or down. Even though reviews are submitted to restaurants with averages across the entire range of possible values, most are written to restaurants with a 4.0 or 4.5-star headline average.

Figure 2 also displays Tripadvisor's rounding cutoffs, indicated with dotted vertical lines. In between these vertical lines, although underlying average ratings (and restaurant quality) vary, headline stars and, thus, reviewers' perception of the average opinion of others remain constant. In the estimation of Equation 1, I will only use observations within a small bandwidth of 0.1 stars around each cutoff. We can see in the picture that the neighborhood around the 4.25 cutoff is the one with more observations. Thus, this cutoff will receive special attention in the discussion of results.

Next, in Figure 3, I investigate the relationship between new ratings and the average of previous ratings appearing on Tripadvisor's page. I pool all rounding cutoffs together and focus on a small bandwidth around them. To make observations more comparable, I normalize ratings and averages by the distance to the closest rounding threshold.[19] To represent the role of reviewers' cultural background parsimoniously, I split observations into two groups, below and above the median individualism score in the sample. [20]

There are three takeaways from Figure 3. First, reviewers who consumed at restaurants with higher averages also tended to give higher ratings. This result is hardly surprising, as a

[19]For example, suppose a restaurant with an average of 4.20 receives a 5-star review, then the normalized review and average become 0.75 and -0.05, respectively (their distance to the 4.25 cutoff).

[20]Since my sample has a more significant number of reviews written by consumers from individualistic countries, the median IDV score is greater than 0.50. Across all reviews in the sample, the median IDV score is 0.71, which is the level of individualism found in France and Sweden.

restaurant's current average ought to be positively correlated with its quality and, thus, with the expected value of its future ratings. Second, Figure 3a indicates that reviewers from countries below the median individualism score seem to react to displayed averages. Ratings submitted by these reviewers increase discontinuously around Tripadvisor rounding cutoffs. Lastly, the behavior of reviewers in the group above the median individualism score differs. Ratings reported by these consumers still depend on the restaurant's underlying average but do not seem to be affected by whether this average falls to the left or the right of the cutoffs.

Figure 3: Effect of a Restaurant's Average on the Rating Given by its Next Reviewer

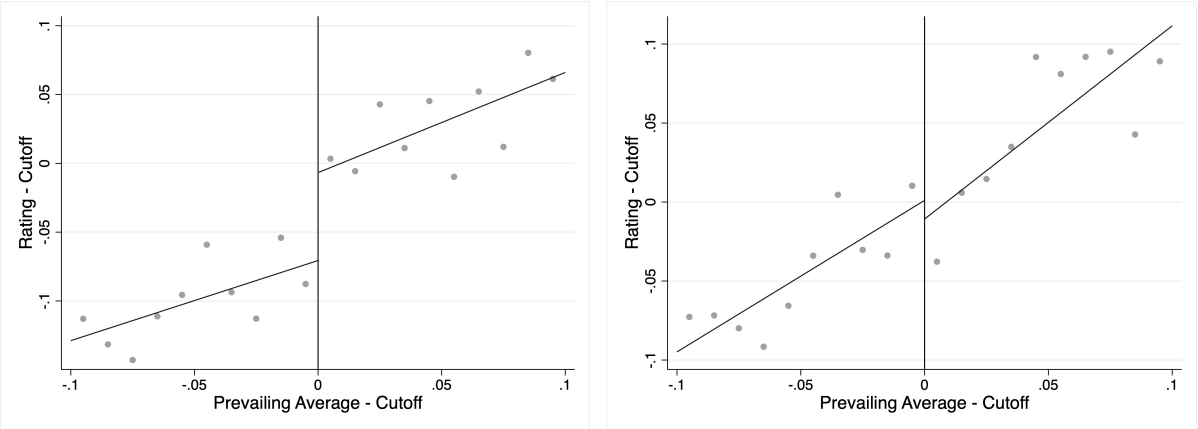

(a) Below Median Individualism

(b) Above Median Individualism

Notes: The median individualism score in my sample is 0.71. Dots represent sample averages within bins of size 0.01.

3.2 Results

3.2.1 Social Influence

Table 2 shows the results of estimating Equation 1. Regressions are estimated using observations within a bandwidth of 0.1 around the relevant rounding cutoffs.[21] Standard errors are clustered at the restaurant level.[22] I first show the results of estimating the model on a sample that pools all cutoffs together. Then, I focus on the most frequent cutoff, 4.25. In Appendix C Table 14, I discuss results for other cutoffs.

The coefficients of interest, β_1 and β_2, are shown in the first and second rows, respectively. Results confirm what's suggested in Figure 3. That is, reviewers' response to the discontinuity in averages displayed by Tripadvisor depends on their country of origin, and the role of individualism in moderating the effect is statistically significant. Specifically, reviewers from more collectivist countries tend to adjust their ratings toward a restaurant's displayed average. On the other hand, results suggest that consumers from individualistic countries do not respond to headline star ratings. Depending on the specification, reviewers from the most individualistic

[21] In Appendix C Table 18, I show that results are not sensitive to the choice of bandwidth.

[22] My main treatment variable, whether the average is displayed rounded up or down, is defined at the level of restaurants. Thus, I choose to cluster standard errors at the restaurant level. I also experimented with clusters at the level of the country of reviewers, which is the level at which individualism scores vary. Both approaches deliver very similar levels of statistical significance.

countries (i.e., Idv_c close to one) may even show a small (not statistically significant) tendency to deviate from the current average rating.

The first four columns present results for the sample that pools all cutoffs together. To make review instances more comparable, I include cutoff fixed effects in these specifications. The most important takeaway is that estimates of β_1 and β_2 from Equation 1 (first and second rows) are fairly stable across specifications. Adding controls (column 2), restaurant fixed-effects (column 3), or individualism-specific effects of the running variable (column 4) has a limited impact on the estimated effect of seeing rounded-up average ratings.

Columns 5 to 8 focus exclusively on the 4.25 cutoff. This threshold accounts for over half of the observations falling within a 0.1 bandwidth of some rounding cutoff. The implication is that results from the pooled sample are driven mainly by reviews submitted to restaurants around this threshold. Note that I only include restaurant fixed effects in columns 3 and 7, whereas they are not present in columns 4 and 8. I do this to compare these more demanding models (fixed effects and individualism-specific slopes) to the baseline specification shown in columns 2 and 6.

The effects estimated using the 4.25 cutoff sample present less stability than those obtained from the pooled sample. In particular, introducing individualism-specific effects of the underlying average Avg_{jn-1} (column 8) produces noisy estimates of this parameter, affecting the discontinuity's estimated effect. Still, in qualitative terms, the main message remains the same. The effect of headline displayed averages on the rating reported by the next reviewer is positive for consumers from collectivist countries and reduces with individualism. Point estimates for consumers from the most individualistic countries are below zero but not statistically significant.

Adding restaurant fixed effects has mixed consequences. On the one hand, the data's sequential nature poses challenges to using within restaurant variation. A positive change between the average consumer $n-1$ sees and the average consumer n observes must come from a high rating left by reviewer $n-1$, which creates a negative correlation between changes in average ratings and changes in individual ratings reported by new reviewers. That is why in columns 3 and 7, the coefficient on the average rating becomes negative.

Nonetheless, using restaurant fixed effects also has advantages. We effectively compare situations where a restaurant changes its underlying average without crossing a rounding cutoff with situations when the same level of change in the running variable leads the restaurant to cross a rounding threshold and, thus, to a change in its displayed average. In that sense, including restaurant fixed effects helps to mitigate concerns that a few restaurants, with lots of reviews and persistent averages, may show up on one side of the cutoff many times and drive most of the results. Given this trade-off, it is reassuring that both types of specifications, with and without restaurant fixed effects, deliver reasonably similar results.

In order to provide a more comprehensive picture of effect magnitudes, I plot linear predictions produced by the specification in column 7. Specifically, I plot predicted ratings given by consumers of cultures with differing levels of individualism under two scenarios, when the prevailing average is rounded down (4.0) and up (4.5). Figure 4 shows the results of this analysis.

Table 2: Effect of Displayed Average and IDV in User's Country on Reported Ratings

	All Cutoffs Pooled				4.25 Cutoff			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Main Effects								
Above Cutoff	0.072***	0.073***	0.086***	0.073***	0.075***	0.074***	0.105***	0.119***
	(0.018)	(0.018)	(0.020)	(0.018)	(0.025)	(0.024)	(0.028)	(0.042)
Above Cutoff × IDV	-0.095***	-0.095***	-0.079***	-0.095***	-0.105***	-0.103***	-0.105***	-0.172***
	(0.024)	(0.024)	(0.027)	(0.024)	(0.032)	(0.031)	(0.035)	(0.064)
Main Controls								
Individualism (IDV)	0.129***	0.080***	0.076***	-0.602***	0.158***	0.125***	0.134***	-2.743
	(0.019)	(0.019)	(0.021)	(0.152)	(0.025)	(0.025)	(0.027)	(2.276)
Average Rating	0.664***	0.652***	-1.121***	0.547***	0.710***	0.675***	-1.447***	0.234
	(0.096)	(0.095)	(0.118)	(0.098)	(0.133)	(0.133)	(0.168)	(0.359)
Average Rating × IDV				0.167***				0.683
				(0.036)				(0.542)
Additional Controls								
Tourist		0.129***	0.147***	0.127***		0.102***	0.115***	0.102***
		(0.008)	(0.009)	(0.008)		(0.010)	(0.012)	(0.010)
Ln Num. Reviews on Trip. (user)		-0.035***	-0.034***	-0.035***		-0.052***	-0.054***	-0.052***
		(0.003)	(0.003)	(0.003)		(0.004)	(0.004)	(0.004)
Smartphone		-0.030***	-0.025***	-0.030***		-0.017**	-0.012*	-0.017**
		(0.006)	(0.006)	(0.006)		(0.007)	(0.007)	(0.007)
Price: High		0.094***		0.094***		0.068***		0.068***
		(0.014)		(0.014)		(0.017)		(0.017)
Ln Prior Reviews (restaurant)		0.014***	-0.029**	0.015***		0.021***	-0.009	0.021***
		(0.003)	(0.014)	(0.003)		(0.004)	(0.025)	(0.004)
Cutoff FE	Yes	Yes	Yes	Yes	No	No	No	No
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant FE	No	No	Yes	No	No	No	Yes	No
Observations	142,206	142,206	140,963	142,206	75,598	75,598	74,914	75,598

Standard errors are clustered at the restaurant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: Dependent variable are individual ratings submitted when the prevailing average was close to a rounding cutoff. All regressions use a bandwidth of 0.1 around rounding cutoffs. Tripadvisor classifies restaurants into low, medium, or high price. Dummy for mid-range price (not shown) is never significant.

Figure 4: Effect of Displayed Average Rating, by Reviewer Individualism Score

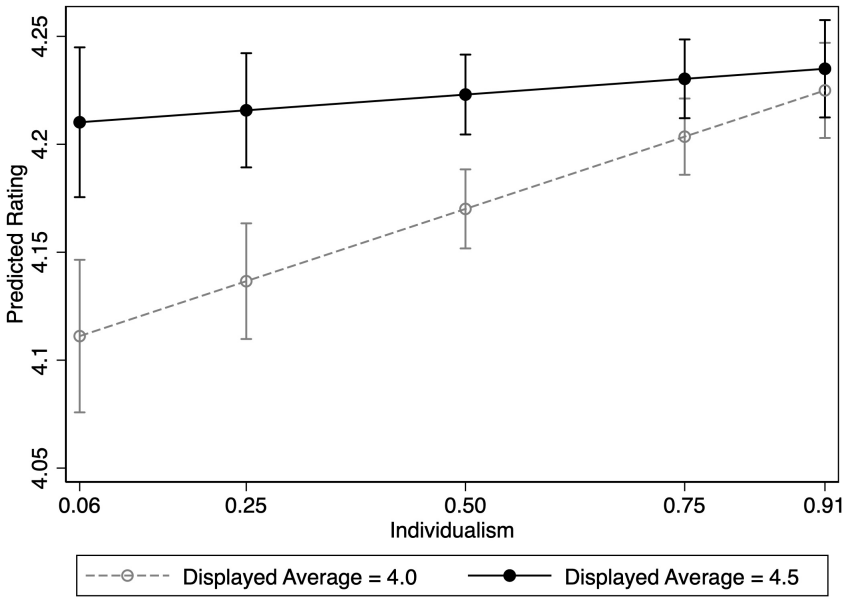

Notes: Linear predictions of the model in column 7 of Table 2

For a given individualism score, the vertical difference between the solid and the dotted lines measures the estimated effect of being shown a 4.5 rather than a 4.0-star average rating. This difference is highest for consumers from the least individualistic countries. They are the ones who, in response to a discontinuous increase in displayed averages, tend to increase their baseline ratings the most. As we move along the individualism score, the difference continuously shrinks towards zero, indicating that reviewers' rating choices become less and less influenced by the average they see on Tripadvisor's page.

Results suggest that the effect for reviewers from the most collectivist country, Guatemala, with a score of 0.06, is about 0.10 stars. One way to interpret this effect is as a share of the displayed average that reviewers incorporate into their rating decision when submitting reviews of their own. Thinking along these lines, the effect for consumers from the most collectivist country is 20%. That is, 20% of the change in headline averages gets reflected in changes in reported ratings. Specifically, a 0.10 stars increase in reported ratings when displayed averages increase by 0.50 stars.

An alternative way to interpret the effects is in terms of how much it changes the probability that reviewers will report specific (discrete) ratings. To quantify these effects, I estimate a specification similar to the one shown in column 6 of Table 2 but using an ordinary logit model rather than a linear regression.[23] Results show that most of the action happens along the 4-star versus 5-star rating choices. Reviewers from countries with low and high levels of individualism substitute between 4 and 5-star ratings in opposite ways. Specifically, the probability that reviewers from collectivist cultures report a 5-star rating increases by three percentage points (from 40% to 43%) when they see a 4.5 rather than a 4.0 headline average.

[23] I choose specification in column 6 because the number of fixed effects in specification 7 complicates the estimation of an ordinary logit model.

The opposite happens to reviewers from the most individualistic cultures, whose probability of reporting a 5-star rating decreases by two percentage points (from 46% to 44%) when headline averages increase from 4.0 to 4.5. Changes in the opposite direction happen to reported 4-star ratings, which, when averages change from 4.0 to 4.5, become less prevalent among collectivists and more prevalent among individualists. Table 17 in Appendix C shows the results of the estimation of this discrete choice model.

These effects may interact with other forces at play in these settings. Suppose, for example, that rounding up a restaurant's average causes one consumer, who was at the margin between submitting 4 or 5 stars, to finally write a 5-star review. There are two ways this single "inflated" rating can meaningfully change the trajectory of ratings received by this hypothetical restaurant. First, it will help the restaurant to maintain its position to the right of the rounding cutoff, keeping its displayed average rounded up, thus influencing future reviewers.

Second, and most importantly, previous research has identified that online settings are prone to significant social influence effects of a sequential nature (Muchnik et al., 2013). That is, the rating reported by consumer n might directly influence the rating reported by consumer $n+1$, holding fixed other types of social influence. Thus, a single reviewer who changes her rating due to the type of social influence effect I identify (i.e., the influence of majority opinion) may have sizable long-run effects on the restaurant's trajectory of ratings.[24]

3.2.2 Covariates

The coefficient on the running variable, a restaurant's underlying average, is generally positive and significant. Once more, that is not surprising. Better restaurants get higher ratings in general, which creates this positive correlation. The exceptions are specifications that include restaurant fixed effects, which, for reasons previously explained, have negative coefficients on Avg_{jn-1}.

A reviewer's individualism score is generally positively correlated with ratings, even when we look at variation in ratings within the same restaurant and include other covariates that correlate with individualism (columns 3 and 7). One potential interpretation of this positive coefficient is that reviewers from more individualistic countries are generally more lenient (less stringent) in their ratings. This suggests that culture may affect reviewers' rating behavior more extensively than just by impacting their propensity to conform to average ratings.

Finally, I briefly discuss the effect of the extra controls included in the analysis. In general, they reduce the baseline correlation between individualism and ratings. However, they do not affect the estimated impact of the discontinuity in displayed average ratings. This indicates that these variables, although correlated with individualism, do not change discontinuously near the rounding thresholds.

Tourists submit higher ratings than locals, which could be related to the pleasant feeling people experience when on vacation. Tripadvisor users with higher activity on the platform tend to report lower ratings, which could be connected with the previously discussed phenomenon

[24] Muchnik et al. (2013) run a field experiment with an online news aggregator and find that a single manipulated (fake) positive vote in the creation of new comments leads to substantial social influence effects. It increases by 32% the probability that the next vote is positive and by 25% its overall rating after five months (overall rating in this context is the number of positive minus negative votes).

of under-reporting. Restaurants of the highest price category received better ratings, which indicates that price correlates with quality.[25]

Lastly, the effect of a restaurant's number of prior reviews depends on whether a restaurant fixed effect is used. When we pool restaurants together, more previous reviews mean higher ratings because popular restaurants also tend to be better. However, when we look within a restaurant, this coefficient becomes negative. This suggests the existence of a negative trend in ratings over the life of a restaurant (Dai et al. (2018) find a similar pattern for restaurants listed on Yelp).

3.3 Evaluating the Identification Assumption

My empirical strategy assumes that expected ratings do not change discontinuously at Tripadvisor rounding cutoffs. This requires the determinants of ratings to be continuous functions of the underlying true average rating held by the restaurant being reviewed. Since the main determinants of ratings are reviewer and restaurant characteristics, I present this section in separate parts. I start by arguing that restaurant characteristics are smooth around rounding thresholds. Next, I discuss the issue of reviewer selection from two angles, selection into consumption and selection into reviewing. In both cases, I present evidence showing that they do not seem to be first order problems in my setting.

3.3.1 Restaurant Characteristics

There is no apparent reason why the characteristics of a restaurant should change discontinuously at rounding cutoffs. For example, restaurant quality likely increases in the value of the underlying average rating, but it is unlikely that it suddenly changes in the neighborhood of any threshold. Still, I test whether there is empirical evidence to support this reasoning. I estimate versions of the baseline specification in Equation 1, only switching the dependent variable from reported ratings to alternative restaurant characteristics. Table 3 presents the results and shows that restaurant attributes show no sign of discontinuities around Tripadvisor rounding cutoffs. Once we control for the underlying average rating, being above a rounding cutoff does not predict a restaurant's age, whether it actively manages its listing, its price category (low, medium, or high), or its number of prior reviews. I do not use restaurant fixed effects in this analysis because two of the four outcomes I study (a restaurant's price class and whether it manages its listing) are time-invariant for a given restaurant. Adding restaurant fixed effects to columns 1 and 4, which look at restaurant time-varying characteristics, does not change the results. The effect of rounding remains statistically non-significant.

Another potential concern with the identification strategy is that restaurants may try to manipulate their online ratings. Previous research has found that star ratings (i.e., rounded average), because they drive more demand to restaurants with better ratings, have a causal impact on restaurant performance, increasing their revenue (Anderson and Magruder, 2012;

[25]The correlation between ratings and restaurant price, although positive, is relatively weak. One explanation for this is that although prices correlate with restaurant quality, they also negatively impact consumer utility. Since ratings measure experienced utility rather than restaurant (absolute) quality directly, it is unclear what we should expect from the correlation between ratings and prices.

Table 3: Effect of Displayed Stars on Restaurant Characteristics

	Rest. Age	Manages Page	Price Class	Prior Reviews
Panel A: Pooled Cutoffs				
Above Cutoff	1.0586	-0.0055	-0.0136	-0.0173
	(1.6159)	(0.0339)	(0.0219)	(0.0735)
Above Cutoff × IDV	0.3002	0.0246	0.0229	-0.0947
	(2.3103)	(0.0425)	(0.0288)	(0.0926)
Cutoff FE	Yes	Yes	Yes	Yes
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes
Observations	142,206	142,206	142,206	142,206
Panel B: 4.25 Cutoff				
Above Cutoff	0.3143	0.0041	0.0011	-0.0251
	(2.5302)	(0.0607)	(0.0305)	(0.1083)
Above Cutoff × IDV	-0.0004	0.0363	-0.0233	-0.1094
	(3.5921)	(0.0739)	(0.0402)	(0.1370)
Rest. City x Month-Year FE	No	Yes	Yes	Yes
Observations	75,598	75,598	75,598	75,598

Standard error are clustered at the restaurant level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: All regressions use a bandwidth of 0.1. Column names refer to the dependent variable used in each model. Only coefficients of the discontinuity effects are shown. All specifications include Avg_{jn-1}, $Idv_{c(i)}$, and baseline covariates, except the dependent variable being studied, if it was one of the original covariates. *Restaurant age* is measured in years, *page claimed* is a dummy indicating whether the restaurant manages its own listing, *price class* is categorical variable with three levels, and *prior reviews* equals $n-1$ for the n^{th} review of restaurant j.

Luca, 2016), and reducing their likelihood of exit (Anenberg et al., 2022). Thus, restaurant owners have an incentive to manipulate their underlying average ratings so that they fall to the right of a rounding cutoff.

Although I cannot entirely rule out the possibility that there is some degree of manipulation, three factors suggest that it is not a first-order concern. First, as Anderson and Magruder (2012) show, restaurants' incentives to manipulate ratings do not jump discontinuously at a cutoff. They have incentives to submit fake positive ratings both when they are slightly below and slightly above a given cutoff. In the first case, they want to move over to the right side of the cutoff, while in the second, they want to remain there.

Second, it is hard for a restaurant to manipulate its average rating on a sustained basis. On Tripadvisor, a reviewer can only submit new ratings to the same restaurant once every three months. Thus, restaurants would have to constantly create user profiles (with an email account linked to them) if they wanted to submit fake positive reviews to themselves. Third, a McCrary (2008) type test finds no evidence of discontinuities in the density of observations around the 4.25 rounding threshold, suggesting that the volume of review activity seems to be a continuous function of restaurants' underlying average. Results of the density test are in Appendix C Table 15.

3.3.2 Selection Into Consumption

Systematic differences in reviewer characteristics may also introduce bias into estimates of social influence. A key concern is that reviewers may rely on Tripadvisor star ratings when deciding which restaurant to visit. For instance, if consumers' restaurant choices are strongly influenced by whether a restaurant has a 4.0 or a 4.5-star average, then reviewers' unobserved characteristics may change abruptly near these rating thresholds. It is important to note that my estimation strategy does not require reviewers to be completely homogeneous on both sides of the cutoffs. It is sufficient that their characteristics vary smoothly as a function of the underlying average rating.

To make this matter more concrete, we can describe consumers' choice of which restaurant to go to as a function of two components. Before searching for the best alternative, consumers may hold an (ex-ante) distribution of preferences, with their tastes for each restaurant. These ex-ante tastes relate to the extent to which they like different restaurant characteristics known to them before consumption, such as location or type of food. Second, consumers may use Tripadvisor's headline star ratings as extra information about the quality of different restaurants. By combining these two components, ex-ante tastes and headline star ratings, consumers decide which restaurant to visit.

As shown in Acemoglu et al. (2017), in as much as headline average ratings enter consumers' decision of which restaurant to eat, a selection effect may arise. Specifically, consumers who go to restaurants with low star ratings will tend to have high idiosyncratic preferences for these restaurants. Otherwise, they would not have chosen a restaurant for which the public signal (i.e., star rating) is low. On the other hand, when a restaurant's star rating is high, even reviewers with a lower idiosyncratic preference for it may choose to try it out.

To see how this selection would bias my estimates, we need to evaluate the systematic differences between reviewers who go to restaurants on different sides of the rounding thresholds. Suppose reviewer selection into consumption is present in the data. In that case, the average consumer who goes to a restaurant with a 4.0-star average will have a higher idiosyncratic taste for it than the average consumer who chooses to eat at a restaurant with a 4.5-star average. Moreover, given that these individual-level preference terms are a significant component of ex-post ratings, one would expect ratings to be lower above the cutoff if they were a strict reflection of experienced utility and social influence played no role. Thus, this paper's estimates of social influence effects likely represent a lower bound for the actual level of conformity to star ratings.

Having described the potential threats that reviewer selection into restaurants based on star ratings may bring to the identification strategy, I discuss why this does not seem to be a critical problem in my setting. First, reviewers potentially choose where to eat based on many factors besides Tripadvisor headline star ratings. For example, people choose which restaurants to patronize based on friends' suggestions, distance from home and work, or using ratings on other platforms, such as Google. Second, if reviewers on different sides of rounding cutoffs were different along unobserved characteristics, one would expect them to differ along observed ones as well. I can directly test this prediction in the data. Table 4 shows the results of estimating regression models analogous to my baseline specification but with reviewer characteristics as the dependent variable. Whether the sample pools all cutoffs together or looks separately at the 4.25 threshold, I fail to find evidence that reviewer attributes change discontinuously at Tripadvisor rounding thresholds.

Looking at each outcome separately, we see that the individualism score in the reviewer's country is not predicted by whether the restaurant's average rating falls to the left or the right of a cutoff. This indicates that reviewer selection issues, even if they exist, are unlikely to depend on individualism, providing additional support for interpreting the baseline estimates as originating from cultural differences in rating behavior rather than selection. Next, I show that neither the probability of being a tourist nor the total number of reviews a user has submitted to Tripadvisor display discontinuities around the cutoffs. These reviewer-level characteristics likely correlate with one's propensity to use Tripadvisor star ratings to decide which restaurant to eat at. The above cutoff indicator not having a significant effect on these variables provides additional evidence that selection is not quantitatively important.

A piece of additional evidence that reviewer selection into consumption is not a big problem comes from the extent to which new reviews changes with restaurants' underlying (continuous) average ratings. Suppose those who self-select to consume at restaurants with a rounded-down star rating are systematically different from those who select to eat at places with rounded-up averages. Then, the reviewers' response to changes in restaurant quality should differ over different sides of rounding thresholds. Empirically, a test of this prediction can be approximated by the idea that a strong selection effect suggests that the slope of new ratings on the underlying continuous average rating (i.e., a proxy for restaurant quality) should differ across sides of the cutoffs. However, going back to Figure 3, where I plotted the relationship between new ratings and a restaurant's previous average, we see that the slopes on both sides of the cutoff are very

similar. Indeed, when I estimate a model similar to Equation 1 but that allows the slope to be different on each side of the cutoff, I fail to find evidence that slopes are statistically different.

Table 4: Effect of Displayed Stars on Reviewer Characteristics

	Individualism	Tourist	Ln User Reviews	Smartphone
Panel A: Pooled Cutoffs				
Above Cutoff	-0.0028	-0.0106	-0.0060	-0.0035
	(0.0025)	(0.0200)	(0.0208)	(0.0093)
Above Cutoff × IDV		0.0172	0.0203	0.0009
		(0.0285)	(0.0271)	(0.0122)
Cutoff FE	Yes	Yes	Yes	Yes
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes
Observations	142,206	142,206	142,206	142,206
Panel B: 4.25 Cutoff only				
Above Cutoff	-0.0004	-0.0202	-0.0086	-0.0117
	(0.0036)	(0.0323)	(0.0285)	(0.0125)
Above Cutoff × IDV		0.0269	0.0045	-0.0032
		(0.0472)	(0.0371)	(0.0166)
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes
Observations	75,598	75,598	75,598	75,598

Standard error are clustered at the restaurant level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: All regressions use a bandwidth of 0.1. Column names refer to the dependent variable used in each model. Only coefficients of the discontinuity effects are shown. All specifications include Avg_{jn-1}, $Idv_{c(i)}$, and baseline covariates, except the dependent variable being studied, if it was one of the original covariates.

3.3.3 Selection Into Reviewing

Prior literature has found evidence that consumers' decision to submit a review depends on the utility derived from their consumption experience. This pattern, known as under-reporting, is such that consumers' propensity to post reviews is higher for extreme levels of consumption utility. In my data, approximately 50% of the observations are 5-star ratings, suggesting that some consumers only submit reviews when they are delighted with the restaurant and want to praise it.[26] Under-reporting by itself should not bias my estimates. It is defined by consumers being more likely to report excellent (5-star) or terrible (1-star) ratings, but it does not depend on restaurants' displayed averages. Thus, one would expect to get similar levels of under-reporting behavior across both sides of rounding cutoffs.[27]

[26]Due to the large fractions of both 1-star and 5-star ratings, researchers often use the term "J-shaped" to describe the distribution of ratings. However, this phenomenon seems to apply only to 5-star ratings in my data. The frequency of 1-star ratings is 5.4%, very similar to the share of 2-star ratings, 4.8%.

[27]Note that under-reporting differs from strategic rating, defined as a situation when consumers exaggerate their review in an attempt to change the headline average in the direction of what they think the "correct" rating is.

However, under-reporting may still present a challenge as it obscures a portion of the effect I aim to capture. My primary objective is to estimate whether reviewers' choice of rating to report is influenced by the average rating displayed on Tripadvisor, which assumes a certain degree of independence between consumers' decision to review and their choice of rating. For instance, if some consumers only use Tripadvisor to submit 5-star reviews, we will not observe any behavioral difference based on the displayed average. Therefore, for this group of reviewers, it may be more valuable to study their decision of whether to leave a review rather than their rating choice. Although I have limited data to investigate the review decision, it is an important area for future research.

Suppose a large subset of reviewers under-report in this systematic manner. In that case, the effect of displayed averages on consumers' rating choices may be close to zero when one looks at the entire sample of reviewers. There is no direct way to know which ratings in my sample result from under-reporting behavior and which came from consumers who, at least in principle, would be willing to submit different ratings for different levels of consumption utility. However, one reviewer-level characteristic likely correlates with this type of behavior, the frequency of activity on Tripadvisor. Users who only submit reviews occasionally are likelier to follow the rule of "review only if loved (or hated) the experience." On the other hand, consumers who write reviews fairly frequently are more likely to report on experiences of different quality levels and a broader range of ratings.

I find evidence in line with this hypothesis. The rating distribution differs significantly depending on the level of activity a given user presents on Tripadvisor. Specifically, conditional on submitting a review, the probability that it is a 5-star rating is substantially lower for consumers who use Tripadvisor more frequently. The definition of what is a frequent and what is a casual reviewer will always be subjective. There is a trade-off between mitigating the problems brought by under-reporting and keeping a sample that is still fairly representative of a large fraction of reviews on Tripadvisor. In trying to balance these two goals, I defined frequent reviewers as all those who submitted more than 20 ratings to Tripadvisor. Among all observations for which reviewers' country has an individualism score available, about 64% are submitted by frequent reviewers.

To give a sense of the potential difference in the extent of under-reporting behavior across these two types of consumers, I plot the rating distribution by reviewer type. Figure 5 shows that among reviews submitted by casual reviewers, the share of 5-star ratings is above 60%. For frequent reviewers, however, the same measure is approximately 41%. All regressions in the main body of the paper only use reviews submitted by frequent reviewers. In Appendix C Table 16, I show how baseline results change for different types of reviewers in terms of their level of activity on Tripadvisor. When I estimate the baseline model with all reviewers, social influence effects and its relationship to individualism get closer to zero. As I increasingly restrict the sample to include only reviews submitted by users with a certain minimum number of reviews

Under-reporting only refers to a relationship between consumption utility and propensity to report a rating. In contrast, strategic rating depends on the utility obtained relative to a restaurant's headline average at that moment. This section only discusses under-reporting. Strategic rating, a potential mechanism of how headline averages may affect reviewer behavior, is discussed in Section 3.4.

Figure 5: Distribution of Ratings by Reviewer Type

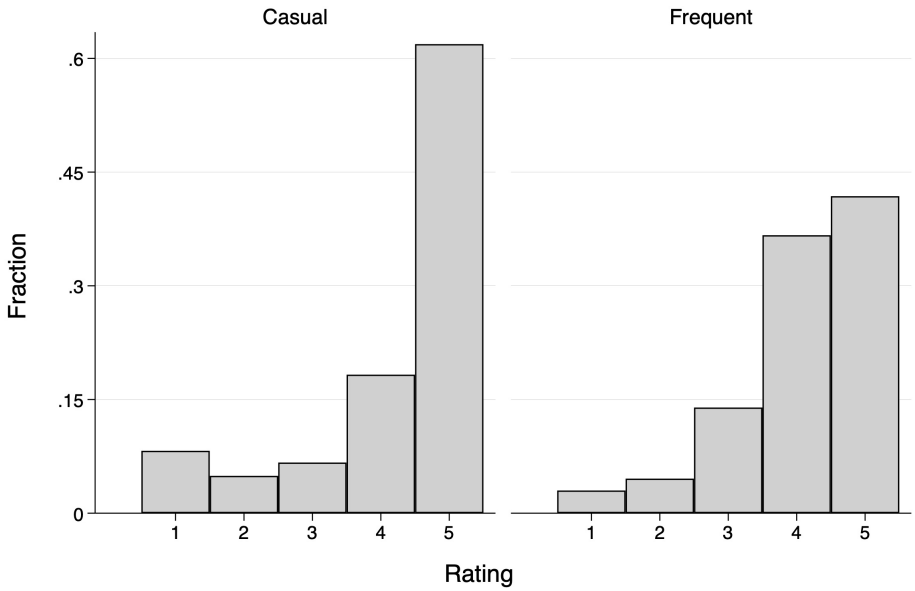

Notes: Reviewers are split according to the number of reviews they have ever posted to Tripadvisor. Casual reviewers are Tripadvisor users with 20 or less reviews. Frequent reviewers have contributed with more than 20 reviews.

on the platform, effects increase. Systematic differences in the propensity to under-report could be driving this result.

A related phenomenon is when people give good ratings to restaurants related to a person they know with the explicit intention of "helping" their business. For example, a person may submit a 5-star rating to a restaurant owned by a friend or where a neighbor works. Particularly problematic is the fact that people's willingness to post an excellent review to help a friend seems to vary systematically across countries. People in less individualistic countries report a higher degree of acceptance of such behaviors (Henrich, 2020).[28] This type of heterogeneity in rating behavior across countries would contaminate my analysis if prevalent in the data. I believe that excluding casual reviewers also helps to mitigate this concern.

3.4 Interpretation

There are two main results so far. First, some reviewers respond to discontinuous changes in a restaurant's average rating by adjusting their ratings in the same direction. In other words, there is evidence that some reviewers conform to average star ratings. Second, the extent to which this happens depends on culture. Specifically, reviewers from countries with higher individualism scores tend to display lower levels of conformity. This section discusses potential interpretations for these two types of effects.

[28] As discussed by Henrich (2020), individualism correlates with higher levels of impersonal fairness and stronger attachment to abstract notions of justice. On the other hand, collectivism correlates with the prevalence of contextual moral judgments. This difference may help explain why leaving a good review to help a friend is more socially acceptable in collectivist countries.

3.4.1 Conformity

I start by discussing why reviewers may react to the average of previous ratings independently of the role that culture may play. In order to do this, I proceed in steps. First, I consider, in general terms, the motivations why people conform to the behavior of others. Second, I discuss the reasons why consumers write online reviews in the first place. Third, I draw connections between these two analyses and present empirical tests on possible interpretations for patterns of social influence presented in Table 2.

Broadly speaking, the literature on social influence describes two types of conformity, *informational* and *normative* conformity (Deutsch and Gerard, 1955). The former is about social learning, which arises when there is uncertainty about the utility-maximizing choice. In these cases, the choices of others may convey valuable information about what the best action might be (Bikhchandani et al., 1998). The latter is inherently social and is rooted in humans' desire to maintain a good image in the eyes of others and cultivate a sense of belonging to a social group (Cialdini and Goldstein, 2004).

Next, concerning why consumers submit online reviews, the first thing to remember is that there are no material payoffs involved. Thus, consumers must act out of some alternative incentive. Previous literature has identified several psychological motivations underlying consumers' reviewing behavior. There are four key motivations. First, consumers might have an intrinsic desire to share their experiences. Second, people want to reward (punish) a business that served them well (poorly). Third, reviewers are concerned with helping other consumers to find (avoid) good (bad) firms. Lastly, people enjoy being part of the "online community" of reviewers and are motivated to maintain a good reputation in this environment (King et al., 2014).[29]

Given this discussion, some reasons why reviewers may react to the average rating given by past consumers become apparent. For example, the desire to earn a good reputation among reviewers may create incentives for individuals to submit ratings that echo the opinion of most consumers. The same motivation to earn a good reputation may also create incentives for individuals to deviate from the average. The direction of the incentive depends on what is rewarded in a given online community, conformity or assertiveness and uniqueness. Alternatively, reviewers' concern with being fair to the restaurant may imply that they will want, at least to some degree, to submit a rating that approximates true quality rather than just their personal consumption experience. In that case, they will incorporate information from the average rating in their own ratings.

In practice, it is hard to pin down and separate the relative importance of informational versus normative conformity. Even in controlled environments, things are not always as clear as they might seem. In the classic Asch (1956) conformity experiment,[30] given the low difficulty of the task, one may be tempted to interpret conformity behavior as resulting from normative reasons entirely. However, post-experiment interviews with participants revealed that some

[29]There are other motivations for writing online reviews and do not intend to provide an exhaustive list. See Yoo and Gretzel (2008); King et al. (2014); Lafky (2014) for a more detailed discussion on this topic.

[30]See Section 1, in the discussion of the related literature, for an explanation of the Asch (1956) experiment.

questioned the validity of their judgments and considered the possibility that the majority was accurate, suggesting that informational conformity may also have played a role, even if smaller.

With these caveats in mind, in Table 5, I show the results of some tests aimed at providing suggestive evidence on some of the mechanisms driving reviewers' conformity to the average rating. All regressions follow the structure of Equation 1, with the difference that I add interactions of some reviewer and restaurant characteristics with the indicator for whether the underlying average rating was above a rounding cutoff. The goal is to test whether any of these characteristics suggest shifts in overall levels of social influence independently of culture. Thus, I do not interact reviewer and restaurant characteristics with individualism scores. To have a more homogeneous sample, I focus on the 4.25 cutoff only.[31]

Table 5: The Role of Moderators of Reviewers' Response to Average Ratings

	Dependent Variable: Rating r_{ijn}					
	(1)	(2)	(3)	(4)	(5)	(6)
Above Cutoff	0.164***	0.148***	0.166***	0.151***	0.119**	0.134**
	(0.0463)	(0.0468)	(0.0471)	(0.0544)	(0.0546)	(0.0552)
Role of Culture						
Above Cutoff × IDV	-0.102***	-0.108***	-0.115***	-0.0986***	-0.109***	-0.114***
	(0.0312)	(0.0305)	(0.0315)	(0.0353)	(0.0354)	(0.0363)
Reviewer i Characteristics						
Above Cutoff × Ln User Reviews	-0.0210***	-0.0201***	-0.0208***	-0.0188***	-0.0184**	-0.0190***
	(0.00691)	(0.00690)	(0.00688)	(0.00720)	(0.00718)	(0.00717)
Restaurant j Characteristics						
Above Cutoff × Prior Reviews		0.0000849**	0.0000870**		0.000234***	0.000238***
		(0.0000423)	(0.0000421)		(0.0000531)	(0.0000530)
Above Cutoff × Rest. Reviews sq.		-2.89e-08**	-2.89e-08**		-7.57e-08***	-7.43e-08***
		(1.45e-08)	(1.40e-08)		(1.22e-08)	(1.22e-08)
Reviewer-Rest. ij Characteristics						
Above Cutoff × Tourist			0.00625			0.00193
			(0.0160)			(0.0189)
Above Cutoff × Smartphone			-0.0214			-0.0246*
			(0.0140)			(0.0143)
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant FE	No	No	No	Yes	Yes	Yes
Sample	4.25 cutoff	4.25 cutoff	4.25 cutoff	4.25 cutoff	4.25 cutoff	4.25 cutoff
Observations	75,598	75,598	75,598	74,914	74,914	74,914

Standard errors are clustered at the restaurant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: All regressions use a bandwidth of 0.1 and include the restaurant's underlying average ratings as well as a three-level categorical variable indicating restaurant price. Only the coefficients of discontinuity effects and its interactions are shown, but whenever a given interaction is included, the main effect of that variable (its effect without the interaction) is also part of the model.

In column 1, I test whether the total number of reviews a user has ever posted on Tripadvisor affects her response to average ratings. The number of reviews a user has contributed may serve as a rough proxy for her level of involvement with the Tripadvisor community, which includes discussion forums where consumers share tips and recommendations. The negative and significant coefficient of the interaction term in the third row indicates that users who have

[31] Focusing on a single cutoff helps to ensure that variables that might affect social influence have comparable values across observations.

submitted more reviews tend to conform less to the average rating. One interpretation of this result, through the lenses of social approval and image concerns, is that, on average, a user gains recognition on Tripadvisor by going against the majority and giving her own unique opinion. This type of argument is in line with results from Hong et al. (2016), who show that the number of "upvotes" a review receives increases in the absolute distance between its rating and the average of previous reviews.

Table 5 column 2 adds interaction terms related to the number of prior reviews held by the restaurant when it gets its n^{th} review. The number of past ratings is easily visible to new reviewers and appears next to a restaurant's average star rating.[32] Together, the coefficients of the linear and quadratic terms imply that conformity increases in the number of prior reviews.[33] In particular, it increases faster when the number of previous reviews is low and reduces the rate of increase as n gets larger. This result is consistent with the hypothesis that reviewers are somewhat concerned with being fair to the restaurant. Thus, when the average rating is very informative about true quality (large n), consumers put more weight on it when submitting ratings of their own.

In column 3, I test whether consumers reviewing outside their country (i.e., tourists) displayed differential levels of conformity. Being a tourist may correlate with having less information about the restaurant, which would suggest a higher incentive to conform to the previous average rating due to informational factors. On the other hand, relative to locals, tourists are less involved with a hypothetical group of reference (e.g., other reviewers of the same restaurant or potential readers), which suggests lower incentives to conform due to normative/psychological reasons. These factors may cancel each other, which may help explain why I do not find any significant effects of being a tourist on the magnitude of conformity.

Finally, in column 3, I evaluate if the propensity to conform to the average rating depends on whether the review was submitted using a smartphone or desktop computer. Average star ratings are more distinctively visible in the mobile application than on the website. Thus, if part of reviewers' response to displayed average ratings arises due to unconscious priming, we should expect conformity to be higher for reviews submitted from smartphones. I do not find strong support for this hypothesis in the data. Columns 4 through 6 include restaurant fixed effects and repeat the analyses. The results are very similar.

3.4.2 Cross-Cultural Variation

As for the second result, that consumers' reaction to average ratings varies systematically with culture, my preferred interpretation is that overall reviewer motivation is culture-dependent. That is, I speculate that, in choosing which rating to report, reviewers from different cultures assign different weights across the set of motivations described before (i.e., sharing personal experience, rewarding the restaurant, helping other consumers, and social benefits.). Findings from the psychology and economics literature documenting cross-cultural heterogeneity in motivations, beliefs, and norms support this argument.(Henrich et al., 2010).

[32] See Appendix A for a summary of Tripadvisor interface).

[33] Coefficients imply that conformity increases up to around $n = 1500$, which represent the 98^{th} percentile of observed values of n.

In particular, this literature points to the fact that people for more collective cultures have a stronger distaste for publicly expressing disagreement (Hofstede, 2001) and are more likely to conform to the majority opinion even when it contradicts their own private judgment (Bond and Smith, 1996). Furthermore, higher levels of individualism correlate with people's desire to feel they have unique tastes (Kim and Markus, 1999). These dimensions of psychological heterogeneity correlate with one another and have similar implications for what we should expect from reviewers' propensity to be influenced by the opinion of others. They all suggest that if it exists, conformity to the current average rating will be weaker for consumers from more individualistic countries.

As mentioned in Section 1, explanations for cross-country variation in people's propensity to conform are usually based on long-run cultural evolution arguments. This literature argues that differences in environmental and social factors in the distant past have set societies on different (cultural) evolutionary paths by modifying what kinds of social institutions and norms were more or less beneficial given the circumstances. For example, Fincher et al. (2008) highlights the role of exposure to diseases in increasing the relative benefits of adopting collective values and shows that pathogen prevalence predicts variation in collectivism-individualism. A second example is Schulz et al. (2019) and Henrich (2020), who argue that practices of the Catholic Church in the Middle Ages altering kinship structures in parts of Europe are behind a substantial fraction of today's variation in levels of individualism.

Although this literature points out specific factors that may have caused today's cross-country variation in people's tendency to conform, it does not intend to directly separate the roles of informational and normative types of conformity. From an evolutionary point of view, these two mechanisms are intertwined and evolve simultaneously. To the best of my knowledge, no studies empirically investigate whether cross-cultural differences in propensity to conform are more closely related to informational or normative (social) factors.

In Table 6, I augment the exercise from Table 5 by testing whether the effect of variables relevant to predict conformity changes with culture. I do this by adding triple interaction terms, which combine the discontinuity treatment, a given review-level variable, and the individualism score in the reviewer's country. The first column is identical in both tables, but in Table 6, I only show coefficients we are directly interested in. Column 2 shows no indication that the relationship between user activity level and propensity to conform depends on individualism. Columns 3 and 4 add the effect of restaurants' number of previous reviews. Once more, there is no evidence of a significant interaction between this moderator of conformity and culture. Finally, column 5 adds restaurant fixed-effect effects and presents the same qualitative patterns of the model in column 4.

I do not have enough statistical power to properly test the combined effect of individualism and reviewer-restaurant characteristics on their reaction to average ratings. However, point estimates indicate that individualism's negative association with conformity becomes stronger in magnitude when these dimensions of heterogeneity are taken into account. For example, the negative relationship between users' number of reviews and conformity is yet more negative when individualism is high. Moreover, the positive relationship between conformity and a restaurant's number of prior reviews is less pronounced for reviewers from countries with high

individualism scores. That is consistent with a story in which cross-cultural differences in an individual's propensity to conform to the group occur both when reputational/image-related concerns are dominant and when information plays the central role.

Table 6: Moderators of Cross-Cultural Variation in Response to Average Ratings

	Dependent Variable: Rating r_{ijn}				
	(1)	(2)	(3)	(4)	(5)
Reviewer i Characteristics					
Above Cutoff × Ln User Reviews	-0.0210***	-0.0168	-0.0156	-0.0147	-0.0165
	(0.00691)	(0.0140)	(0.0140)	(0.0140)	(0.0146)
Above Cutoff × Ln User Reviews × IDV		-0.00666	-0.00719	-0.00839	-0.00292
		(0.0202)	(0.0202)	(0.0201)	(0.0208)
Restaurant j Characteristics					
Above Cutoff × Prior Reviews			0.0000849**	0.000160**	0.000308***
			(0.0000423)	(0.0000641)	(0.0000793)
Above Cutoff × Rest. Reviews sq.			-2.89e-08**	-4.36e-08***	-9.03e-08***
			(1.45e-08)	(1.67e-08)	(1.53e-08)
Above Cutoff × Prior Reviews × IDV				-0.000111	-0.000110
				(0.0000912)	(0.0000914)
Above Cutoff × Prior Reviews sq. × IDV				2.17e-08	2.21e-08
				(1.60e-08)	(1.56e-08)
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes	Yes
Restaurant FE	No	No	No	No	Yes
Sample	4.25 cutoff	4.25 cutoff	4.25 cutoff	4.25 cutoff	4.25 cutoff
Observations	75,598	75,598	75,598	75,598	74,914

Standard errors are clustered at the restaurant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: All regressions use a bandwidth of 0.1 and include the restaurant's underlying average ratings as well as a three-level categorical variable indicating restaurant price. Only the coefficients of discontinuity effects and its interactions are shown, but whenever a given interaction is included, the main effect of that variable (its effect without the interaction) is also part of the model.

3.5 Additional Tests and Robustness Checks

I conducted a placebo test that further assures that the discontinuities in reported ratings result from Tripadvisor's rounding of average ratings. Instead of estimating reviewer response around the true cutoffs (i.e., 1.25, 1.75, ..., 4.25, 4.75), I check if they react to placebo non-rounding cutoffs at 1.1, 1.6, ..., 4.1, 4.6. When restaurants' underlying average ratings cross these cutoffs, headline star ratings displayed to new reviewers do not change. Thus I expect to find no discontinuities in their reporting behavior. Table 7 shows the results of this exercise and confirms that reviewers do not respond to these cutoffs. Similarly to the baseline analysis shown in Table 2, I first pool all cutoffs together and then look separately at the one with the most observations. None of the specifications deliver significant effects of the placebo cutoffs on the ratings reported by the next reviewer.

I conduct a few additional robustness checks, shown in Appendix C. First, I show that results are robust to changing bandwidth sizes (Table 18). Second, I estimate models that allow the effect of the underlying (continuous) average rating to differ across sides of the rounding cutoffs or have a quadratic term (Table 19). These variations produce little change in estimated social

Table 7: Placebo Test for Discontinuities Around Non-Rounding Cutoffs

	Dependent Variable: Rating r_{ijn}							
	Pooled (Placebo) Cutoffs				4.6 Cutoff			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Above Cutoff	-0.001	-0.004	0.002	0.007	0.005	0.008	0.013	0.025
	(0.018)	(0.030)	(0.030)	(0.032)	(0.025)	(0.042)	(0.041)	(0.044)
Above Cutoff × Individualism (IDV)	0.003	0.007	-0.000	0.024	0.015	0.009	0.004	0.013
	(0.023)	(0.044)	(0.044)	(0.046)	(0.032)	(0.063)	(0.062)	(0.066)
Individualism (IDV)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Average Rating	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Average Rating × IDV	-	Yes	Yes	Yes	-	Yes	Yes	Yes
Average Rating × Above Cutoff	-	-	Yes	Yes	-	-	Yes	Yes
Average Rating × Above Cutoff × IDV	-	-	Yes	Yes	-	-	Yes	Yes
Rest. City x Mon-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cutoff FE	Yes	Yes	Yes	Yes	-	-	-	-
Restaurant FE	-	-	-	Yes	-	-	-	Yes
Observations	154,907	154,907	154,907	153,738	78,065	78,065	78,065	77,346

Standard errors clustered at the restaurant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: Dependent variable is a rating submitted when the restaurant's prevailing average was within a small bandwidth around placebo (non-rounding) cutoff. All models use a bandwidth of 0.1. Placebo cutoffs are 1.1, 1.6, 2.1, 2.6, 3.1, 3.6, 4.1, 4.6.

influence. Moreover, I also show that results hold in discrete outcome models, such as ordinary probit or logit (Table 17).

Finally, I evaluate the robustness of baseline results to the different sample choices described in Section 2. In summary, results are robust to the addition of all years of data and to the inclusion of reviews submitted when there were only a small number of previous reviews of the same restaurant.[34] Adding casual reviewers has a larger impact on estimated social influence, which becomes smaller across all cultures. In Section 2, I provide details on the sample construction. For a discussion on why casual reviewers are not suited for my analysis, see Section 3.3.3.

4 Other Determinants of Social Influence

In Section 3, I presented evidence that, when submitting ratings of their own, reviewers' response to discontinuous changes in the average displayed by Tripadvisor correlates with the level of individualism in their country's culture. Nevertheless, whether this association emerges due to a real relationship between individualistic cultural values and a person's propensity to react to what others think or do is still an open question.

The observed cross-country variation in response to displayed average ratings may have alternative explanations. First, individualism strongly correlates with other crucial dimensions of heterogeneity across countries, such as income and religious affiliation. Thus, individualism may be picking up the effect of another underlying non-culture-related variable. Second, culture is a multidimensional object, and countries vary in cultural values other than just individualism.

[34] Remember that the baseline sample is restricted to the years from 2015 to 2019 and does not contain the first ten reviews of each restaurant.

Some of these other cultural traits are correlated with individualism and may drive the results presented so far.

Does individualism remain a relevant predictor of reviewers' propensity to go along with headline stars displayed on Tripadvisor once we control for other dimensions of heterogeneity across countries? Is individualism the most relevant aspect of culture for explaining patterns of cross-country heterogeneity in reviewer behavior? This section argues that the answer to both of these questions is yes.

4.1 Empirical Model

I use a two-stage approach in the spirit of Combes et al. (2008), who explain variation in wages as a combination of worker and city characteristics. Here I adapt their idea to my setting, where I explain heterogeneity in ratings as a function of characteristics specific to a given review (consumer-restaurant) and characteristics shared by all reviewers from the same country, such as culture. In practice, the first stage is similar to Equation 1, with a critical difference. I substitute the individualism score variable with the reviewer's country fixed effects. Thus, in the first stage, I estimate country-level social influence effects. Then, the second-stage regression investigates which country-level characteristics, such as individualism or other dimensions of culture, better explain these fixed effects.

4.1.1 First Stage

In the first stage, I estimate the following regression:

$$Rat_{ijn} = \alpha Avg_{jn-1} + X'_{ijn}\gamma_1 + \theta_{c(i)} + \left(\beta_{c(i)} + X'_{ijn}\gamma_2\right) \times \mathbb{1}\{Avg_{jn-1} \geq k\} + \epsilon_{ijn} \tag{2}$$

Rat_{ijn} is still defined as a rating written by reviewer i to restaurant j, which is receiving its n^{th} review. Avg_{jn-1} also retains its meaning, j's underlying (continuous) average over ratings prior to n. X_{ijn} is a vector of reviewer and restaurant attributes. Specifically, it includes an indicator for whether i a tourist (i.e., $c(i) \neq c(j)$), the log of the total number of Tripadvisor reviews ever posted by i, an indicator for whether the review was submitted using a smartphone, the number of previous reviews accumulated by j (and its square). To account for the non-linear effect of the number of previous reviews accumulated by a restaurant, X_{ijn} also includes its square.[35]

To take into account average cross-country variation, Equation 2 includes two types of (reviewer) country fixed effects. The first, θ_c, captures cross-country variation in reported ratings across reviewers who are shown the rounded-down version of the restaurant's current average. The second type of country fixed effects is represented by β_c and captures the average social influence effect over consumers from country c. That is, across consumers from a given country, it measures the difference between the value of ratings submitted when Tripadvisor's headline average was rounded up versus when it was rounded down.

[35] In terms of the direct effect of n on reported ratings, the quadratic form is usually interpreted as arising from selection into consumption, where early reviewers usually have a higher taste for the restaurant than late ones (Dai et al., 2018). Regarding the effect of n on users' response to the average, it may affect the propensity to conform both due to informational and psychological reasons (see more in Section 3.4).

Unlike in Equation 1, here I allow the vector of reviewer and restaurant characteristics X_{ijn} to affect not only the baseline ratings but also reviewers' response to rounded-up headline averages. I follow Combes et al. (2008) and normalize reviewer and restaurant attributes to represent deviations from country means. For example, if x_{ijn} is a variable in X_{ijn}, the actual value used in the first-stage regression will be $\tilde{x}_{ijn} = x_{ijn} - \bar{x}_{c(i)}$, where $\bar{x}_{c(i)}$ is the mean of x over reviews (observations) submitted by reviewers from the same country as user i. This type of transformation allows the estimation of the effect of reviewer and restaurant characteristics (captured in the parameters γ_1 and γ_2) while also allowing average cross-country differences in ratings to be loaded in the fixed effects θ_c and β_c. This type of specification assumes that a unit increase in a given characteristic has the same impact on reviewers from different countries. For example, relative to being a local, being a tourist has the same impact on ratings, regardless of where the tourist comes from. An alternative is to include interactions between characteristics in X_{ijn} with country fixed effects. This approach generates too many parameters, which are estimated with low precision. Another possibility is interacting X_{ijn} with a country's individualism score. However, since individualism is used to explain country fixed effects in the second-stage regression, I prefer not to include it in the first stage to avoid the risks of creating a mechanical correlation between first-stage and second-stage estimates of the effect of individualism.

4.1.2 Second Stage

My goal is to understand what predicts the observed cross-country variation in reviewer reaction to displayed stars. Thus, the second-stage analysis focuses on the potential explainers of β_c. Formally, I run the following linear regression at the country level:

$$\beta_c = Z'_{trip,c}\,\delta_{trip} + Z'_{ext,c}\,\delta_{ext} + \nu_c \tag{3}$$

Terms $Z_{trip,c}$ and $Z_{ext,c}$ are vectors of country-level variables that could explain variation in average social influence effects. I explicitly indicate them with different subscripts to call attention to the fact that while some variables are computed from the Tripadvisor data directly, others, such as different cultural measures, are obtained from external data sources. In the next section, when discussing the results, I will be specific about the contents of $Z_{trip,c}$ and $Z_{ext,c}$.

The empirical model in Equation 2 is more demanding than the one in Equation 1 in terms of the number of parameters to be estimated. In particular, the estimation of country fixed effects requires enough ratings from reviewers from each country. Thus, I pool all cutoffs together to estimate social influence effects for as many countries as possible. To mitigate problems related to estimating fixed effects with too few observations, I restrict the sample to countries for which there are more than 100 reviews, which reduces the number of countries from 97 to 79. However, it only decreases the number of observations by less than 0.5%, from 142,434 down to 141,843. Bandwidth size remains at 0.1 stars.

Figure 6: Distribution of Estimated Country-Level Social Influence

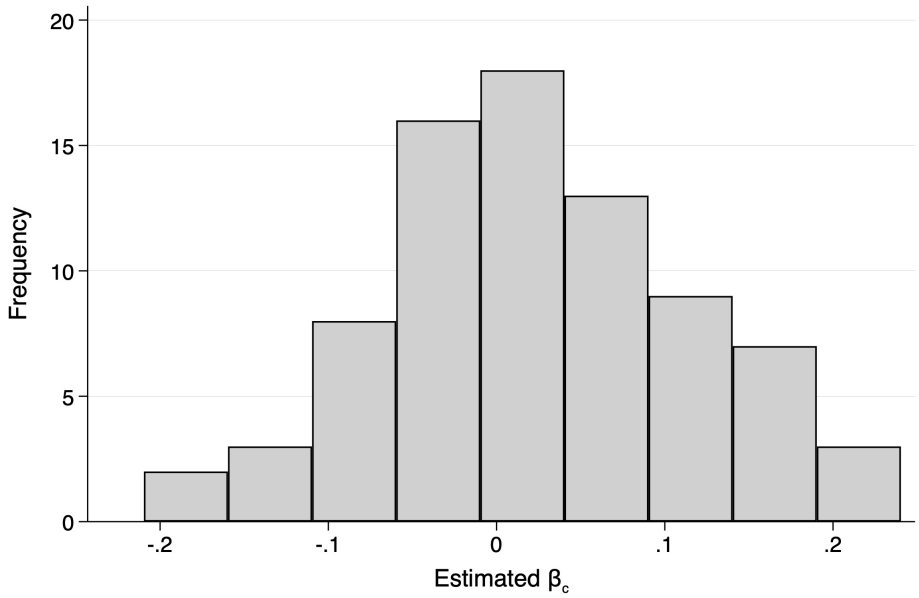

Notes: Distribution of the estimated values of β_c from Equation 2.

4.2 Results

4.2.1 Preliminaries

As previously mentioned, this section investigates alternative explanations for observed cross-country variation in social influence effects. Thus, in presenting the results, I focus exclusively on exploring the estimates of β_c. I start by showing the distribution of estimated β_c. Figure 6 shows that estimated social influence effects vary between -0.2 and 0.2. This range of estimates implies that, for some countries, reviewers' response to the discontinuous increases in the average rating displayed on Tripadvisor is relatively large, a 0.2-stars increase (or decrease) in reported ratings from a 0.5 stars increase in headline displayed averages.

For my purposes, examining the cross-country variation in estimates is more important than looking at the parameter values per se. I start by plotting the correlation between estimated β_c and individualism scores Idv_c. Figure 7 shows that estimated social influence negatively correlates with a country's level of individualism.[36] Naturally, this pattern is in line with the discussion in Section 3.2.1, whose main finding was that the effect of being shown rounded up average ratings decreased in the individualism score of the reviewer's country. However, the approach here is more flexible in allowing social influence effects to differ for countries with the same individualism score. For example, in Figure 7, we can see that Hungary (HUN), the Netherlands (NLD), and Canada (CAN), which all have an individualism score of 0.80, differ in their reviewers' average response to Tripadvisor headline stars.

Next, I test the robustness of the correlation shown in Figure 7 to the inclusion of three types of country-level covariates. First, I check whether the correlation between social influence and individualism can be attributed to cross-country differences in the Tripadvisor population (i.e.,

[36] Figure 7 only includes countries for which I have at least 400 reviews.

Figure 7: Estimated Country Level Social Influence and Individualism

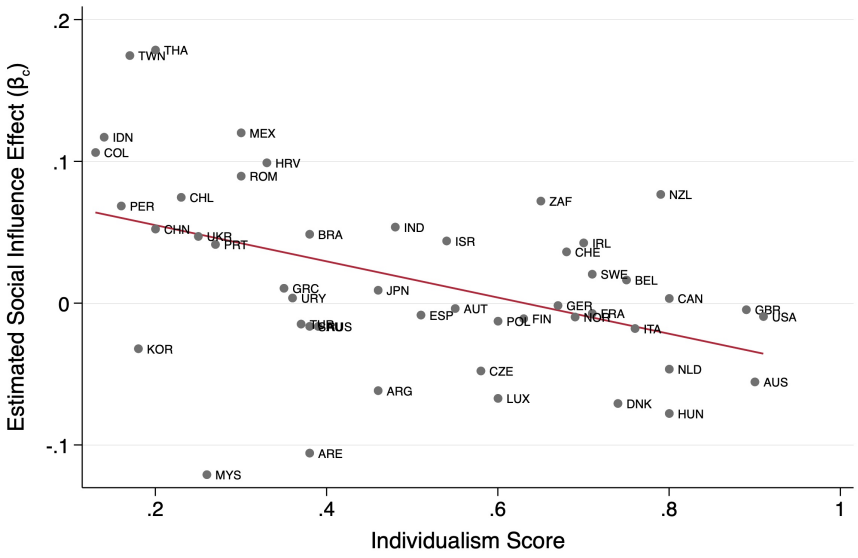

Notes: Scatter plot includes countries for which I have at least 400 reviews, $N = 48$

reviewer and restaurant characteristics). Then, I investigate the effect of variables that strongly correlate with a country's individualism score, namely income and share of Protestants in the population. Finally, I explore the role of culture more broadly, looking at other cultural values besides individualism.

4.2.2 Reviewer and Restaurant Characteristics

In further investigating the correlation displayed in Figure 7, the first test I carry out is whether basic differences in the Tripadvisor population can explain the negative correlation between country-level social influence and individualism score. By Tripadvisor population, I refer to reviewers' characteristics and attributes of the restaurants they face. For example, suppose the effect of average headline ratings on the rating choice of subsequent reviewers is stronger for smartphone users than desktop users (due to differences between the design of the website and the mobile application). If the within-country share of smartphone users correlates with individualism scores, this could drive part of the cross-country differences in estimated social influence effects depicted in Figure 7.

Table 8 shows the results of this analysis. Panel A contains estimates of simple OLS regressions. In Panel B, I weight observations by the inverse of standard errors from first-stage estimates of β_c. Some country-level social influence effects are estimated with higher precision than others, and this second approach is one way to give more weight to countries with more precisely estimated β_c.[37]

Country-level average Tripadvisor characteristics predict a large share of the variation in individualism scores. Generally, reviewers from countries with higher individualism scores face restaurants with more reviews, submit ratings more often from desktop computers, and rate more often as international tourists. The number of reviews submitted by a country's average

[37]Weighting by a country's number of observations in the first stage regression delivers similar results.

Table 8: Predictors of Social Influence Effect (β_c): IDV versus Tripadvisor Characteristics

	Dep. Var: IDV		Dep. Var: Social Influence (β_c)			
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Unweighted Regression						
Individualism (IDV)					-0.1011**	-0.1212**
					(0.0425)	(0.0494)
Mean Prior Reviews (j)	0.0027***	0.0023***	0.0003	0.0007**		0.0010***
	(0.0005)	(0.0005)	(0.0003)	(0.0003)		(0.0003)
Mean Prior Reviews Squared (j)	-0.0000*	-0.0000*	-0.0000*	-0.0000***		-0.0000***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)		(0.0000)
Mean Ln. User Reviews (i)		0.0471		0.0771**		0.0828**
		(0.0590)		(0.0350)		(0.0342)
Share via Smartphone		-0.1402		0.1260		0.1090
		(0.2302)		(0.1439)		(0.1426)
Share Tourist		0.2492*		-0.1547**		-0.1245*
		(0.1400)		(0.0625)		(0.0628)
Observations	79	79	79	79	79	79
R^2	0.357	0.398	0.062	0.164	0.066	0.221
Panel B: Weighted by inv. of SE(β_c)						
Individualism (IDV)					-0.0936***	-0.0958***
					(0.0235)	(0.0322)
Mean Prior Reviews (j)	0.0023**	0.0026**	-0.0001	0.0001		0.0004
	(0.0011)	(0.0010)	(0.0002)	(0.0003)		(0.0003)
Mean Prior Reviews Squared (j)	0.0000	-0.0000	-0.0000	-0.0000		-0.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)		(0.0000)
Mean Ln. User Reviews (i)		0.0279		0.0367		0.0394
		(0.0842)		(0.0333)		(0.0307)
Share via Smartphone		-0.4768**		0.0218		-0.0239
		(0.2058)		(0.1123)		(0.1139)
Share Tourist		0.4939**		-0.1004**		-0.0530
		(0.2082)		(0.0454)		(0.0445)
Observations	79	79	79	79	79	79
R^2	0.472	0.584	0.062	0.140	0.171	0.215

Robust standard in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Outcome variables is Idv_c in columns 1 and 2 and first stage estimates of β_c in columns 3 through 6. Independent variables are country level means (or shares) of different reviewer or restaurant attributes. For each of these variables, X_{ijn} in Equation 2 included deviations from country means. In Panel B, observations are weighted by the inverse of the standard errors of first stage estimates of β_c

user positively correlates with individualism, but this association is not statistically significant. Taken together, analyses in columns 1 and 2 show that some Tripadvisor characteristics are systematically different across reviewers' countries of origin in a way that also correlates with countries' individualism scores. This correlation suggests we should investigate whether such systematic differences are behind individualism's association with social influence effects.

These Tripadvisor characteristics also correlate with estimated social influence effects (β_c), particularly when individualism is not included as a separate regressor. However, the statistically significant effects of Tripadvisor characteristics in the unweighted regression (columns 3 and 4) are driven mainly by countries with noisy first-stage estimates. Most of these effects drop to zero in the weighted regressions, which gives more weight to precise first-stage estimates. The exception is the share of reviewers from a country who submit reviews as tourists, which has a negative and significant effect, even in the weighted regression. Rather than capturing an actual effect of being a tourist on reviewers' propensity to conform, this result arises due to a correlation between tourism activity and individualism scores. The specification in column 6 introduces individualism as a regressor, and the share of tourists is no longer significant in the weighted model, my preferred specification.

Overall, the main takeaway of this analysis is that individualism's negative association with reviewers' tendency to adjust their ratings toward the displayed average cannot be explained by country-level differences in the Tripadvisor population of reviewers and restaurants. Individualism is the only significant variable in the weighted model (Table 8, Panel B, column 6). Moreover, the fraction of the variation in β_c explained by individualism alone, 17%, is non-trivial (column 5). Finally, the point estimate is about 0.1, similar to what we obtained in Section 3.2.1. Next, I investigate the relationship between individualism and social influence by checking whether it is robust to other critical dimensions of cross-country heterogeneity, such as income or religion.

4.2.3 Correlates of Individualism

Individualism correlates with crucial dimensions of country-level heterogeneity. For example, Hofstede (2001) discusses the strong correlation between a country's individualism score and income level.[38] A second example is Protestantism, which since Weber (1905) has been described as one of the forces that helped to set western mentality on a path towards an increasing emphasis on values characteristic of individualistic cultures. Indeed, among countries in my sample, income per capita and the percentage of Protestants in the population are positively correlated with a country's individualism score. Together, these two variables predict approximately 50% of the cross-country variation in individualism scores.

Table 9 shows results of regressions that include income and the share of Protestants in the population as potential explainers of social influence effects. Columns 1 through 3 show that both variables are negatively correlated with first-stage estimates of β_c, especially income, which remains significant when both variables are simultaneously included in the regression.

[38] Hofstede (2001) attributed this correlation to a causal pathway from economic development to individualism. However, the current literature suggests a more complex relationship in which there is a feedback loop between individualism and economic development reinforcing each other (Henrich, 2020).

However, columns 4 to 7 show that this pattern can be explained by cross-country variation in individualism. Both income and prevalence of Protestantism become irrelevant to predict social influence once we include individualism as a regressor. Finally, column 8 shows that results are robust to the inclusion of continent fixed effect.

Altogether, these results support the argument that the correlation depicted in Figure 7 captures a real effect of individualism on people's tendency to conform. Next, I investigate the extent to which this is something particular about individualism or whether culture, more generally, can account for the observed cross-country variation in reviewer response to displayed average ratings.

Table 9: Predictors of Country-Level Social Influence Effect (β_c)

	Dependent Variable: Estimated Social Influence (β_c)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Individualism				-0.087***	-0.077**	-0.094***	-0.086**	-0.080**
				(0.023)	(0.036)	(0.025)	(0.037)	(0.035)
Log Income pc	-0.023***		-0.022***		-0.004		-0.003	-0.001
	(0.008)		(0.008)		(0.012)		(0.012)	(0.012)
Share Protestant		-0.032*	-0.006			0.018	0.017	0.023
		(0.019)	(0.019)			(0.020)	(0.019)	(0.020)
Continent FE	No	No	No	No	No	No	No	Yes
Observations	77	77	77	77	77	77	77	77
R^2	0.105	0.017	0.106	0.153	0.154	0.157	0.158	0.183

Robust standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: All regressions are weighted by the inverse of the standard errors of first stage estimate of β_c. I only include countries for which all regressors are available.

4.2.4 Other Measures of Culture

Focusing primarily on individualism to measure cultural variation was a choice driven by the hypothesis that this dimension of culture was particularly relevant in determining people's tendency to conform to the majority opinion. This hypothesis, in turn, is based on the very definition of individualism and on how it is measured. People in countries with a high individualism score place a substantial value on personal freedom. In contrast, people in countries which score low on individualism put a stronger emphasis on harmony and conformity (Gorodnichenko and Roland, 2011). Thus, the argument is that online reviewers from countries with lower levels of individualism will have a higher tendency to adjust their reviews to more closely match the prevailing consensus, the average over all previous ratings.

However, individualism is just one of the dimensions along which cultures vary across countries. Other cultural values also play a crucial role in determining the nature of a nation's culture. Thus, it is plausible to ask whether some of these other cultural factors affect reviewer conformity behavior on Tripadvisor. The first issue is to define which additional cultural dimensions to consider. Given that the paper focuses on Hofstede (2001) individualism score, a natural approach is to look at other cultural dimensions defined in his model of national cultures. In his original work, based on factor analysis of tens of thousands of answers from

IBM employees to questions on attitudes towards life and work, Hofstede (2001) developed a model of a country's national culture consisting of four dimensions.

The first dimension is the previously discussed *individualism* index, measuring the extent to which people feel independent versus interdependent members of larger groups. Second, Hofstede (2001) defined the *power distance* dimension of culture, which measures the degree to which the less powerful members of a society accept that power is unequally distributed. A higher power distance index indicates more acceptance of inequality in the distribution of power. The third dimension of this model is the *femininity-masculinity* index, which measures the extent to which stereotypical gender roles are embraced in society. Higher scores indicate that gender roles are distinct, where men are supposed to be assertive and tough, while women are supposed to be more modest and tender. Lastly, the *uncertainty avoidance* dimension has to do with a society's tolerance for uncertainty. In societies with a high uncertainty avoidance index, people feel anxious in ambiguous situations and prefer having a stable set of habits and rituals.[39]

I start by quantifying the correlation between these cultural dimensions and individualism. Results are in the top panel of Table 10. The correlation between masculinity and individualism is minimal and not statistically significant. Both uncertainty avoidance and power distance are negatively correlated with individualism. The association is particularly strong for power distance, with a point estimate of almost -0.7 and R^2 of approximately 40%. Given the definition of these two cultural dimensions, this association should be expected. In countries where people emphasize individual independence, we should expect less acceptance of an unequal distribution of power.

Next, I study the correlation between country-level social influence and each cultural dimension. The bottom panel of Table 10 shows the results. In the first column, I repeat the analysis that uses individualism to explain cross-country variation in β_c. In the remaining columns, I show that none of the other dimensions of culture correlate with average country-level social influence effects. Given the negative and relatively tight correlation between individualism and power distance, one might have thought that the latter would positively correlate with estimated social influence effects. Point estimates are indeed positive, but they are small and not statistically significant. Overall, results suggest that, in terms of predicting reviewers' tendency to conform, there is something about individualism that is not captured by other cultural factors.

The set of countries for which each of Hofstede (2001) cultural dimensions is available varies. In particular, the individualism index is available for more countries, which explains the variation in the number of observations used in different models shown in Table 10. To guarantee that the set of countries used for estimation is the same across all cultural dimensions, in Table 11, I repeat a similar exercise which only includes the 68 countries for which all four dimensions are available. Moreover, instead of testing the effect of each cultural dimension separately, I run a "horse race" between individualism and other cultural variables.

[39]The most recent version of the data includes six dimensions. To the four dimensions previously mentioned, Hofstede (2001) added *short versus long-term orientation* and *indulgence versus restraint*. These two extra dimensions were added by incorporating data from the World Values Survey. In order to have a parsimonious treatment of culture, I restrict the analysis to Hofstede's original four dimensions.

Table 10: Relationship Between Estimated Social Influence (β_c) and Cultural Dimensions

	Independent variable			
	IDV	PDI	MAS	UAI
Dep. var: Individualism (IDV)				
Unweighted regression		-0.6872***	0.100	-0.352**
		(0.102)	(0.154)	(0.133)
Observations		73	73	73
R^2		0.395	0.007	0.093
Dep. var: Social Influence (β_c)				
Unweighted regression	-0.101**	0.030	0.009	-0.002
	(0.043)	(0.046)	(0.049)	(0.046)
Weighted by inv. of SE(β_c)	-0.094***	0.055	-0.002	0.006
	(0.024)	(0.035)	(0.028)	(0.023)
Observations	79	73	73	73
R^2	0.066	0.005	0.000	0.000
	0.171	0.038	0.000	0.001

Robust standard in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Dependent variable is the individualism score in the top panel and estimated social influence (β_c) in the bottom panel. For the latter, the table also includes results of a regression in which observations are weighted according to the inverse of the standard errors of estimated of β_c. All models are univariate regressions. Column names indicate the independent variable. IDV, PDI, MAS, and UAI, stand for individualism, power distance, masculinity, and uncertainty avoidance, respectively.

The main takeaway of Table 11 is the consistency in the effect of individualism, invariably negative, significant, and with a stable point estimate. Even when using the four cultural scores together (column 5), the effect of individualism remains statistically significant at the 1% level in the weighted model. The second lesson from Table 11 refers to the explanatory power of individualism. In the weighted model, this cultural dimension alone predicts 22% of the cross-country variation in the extent to which headline average ratings on Tripadvisor influence reviewers.

Finally, columns 4 and 5 show that, once we hold individualism fixed, uncertainty avoidance has a negative and significant effect on β_c. For countries with the same individualism score, conformity to the prevailing average is stronger for countries with more tolerance for ambiguity. I expected the opposite result, higher discomfort in the face of ambiguity to be associated with more conformity. Nevertheless, the effect's magnitude is half of the impact of individualism and only significant in the weighted model. Moreover, uncertainty avoidance only becomes significant when individualism is included as a regressor.

Before concluding this section, I present one last analysis providing evidence that individualism affects reviewers' propensity to react to the perceived majority opinion. Remember that Hofstede (2001) defines collectivism-individualism as two extremes of a single scale defining the extent to which agents perceive themselves as independent versus interdependent members of larger groups. If this particular aspect of culture indeed has important implications for people's propensity to conform to (or deviate from) others, attempts by other researchers to measure similar concepts should also correlate with my estimates of social influence effects on Tripadvisor.

In order to test this prediction, I rely on Schwartz's cultural database. Similarly to Hofstede (2001), Schwartz (1994) developed a framework to quantify cross-cultural variation in core values. He constructed his dataset from survey answers by school teachers and college students across 78 countries.[40] Based on respondents' stated importance of different values as guiding principles in their lives, Schwartz created different cultural measures as assigned country scores.

Notably, Schwartz (1994) model also includes a distinction between cultures in which people are more attached to groups and cultures in which people primarily think of themselves as independent individuals. To separate these two types of cultures, he uses the terminology *embeddedness* versus *autonomy*. In cultures with high levels of autonomy, individuals see themselves as independent bounded entities, similarly to Hofstede's concept of individualism. Moreover, Schwartz (1994) further splits the autonomy dimension of culture into two types. He argues that individuals in societies with high levels of autonomy are encouraged to find meaning by independently seeking their own ideas and intellectual orientations (intellectual autonomy) and engaging in positive experiences for themselves (affective autonomy). On the other hand, in cultures where autonomy is low (i.e., high embeddedness), meaning in life comes primarily from social relationships.

Table 12 shows the results of regressing my estimates of Tripadvisor country-level social influence on both types of scores for autonomy. All variables were normalized to have zero mean and unit standard deviation for the coefficients to be comparable. In the first two columns,

[40] For 58 of these 78 countries, I have estimates of social influence β_c.

Table 11: Relationship Between Estimated Social Influence (β) and Cultural Dimensions

	Dependent Variable: Social Influence (β_c)				
	(1)	(2)	(3)	(4)	(5)
Panel A: Unweighted Regression					
Individualism (IDV)	-0.134***	-0.143**	-0.135***	-0.140***	-0.155**
	(0.041)	(0.065)	(0.041)	(0.046)	(0.069)
Power Distance		-0.017			-0.024
		(0.070)			(0.078)
Femininity - Masculinity			0.026		0.033
			(0.039)		(0.042)
Uncertainty Avoidance				-0.026	-0.026
				(0.045)	(0.046)
Observations	68	68	68	68	68
R^2	0.123	0.124	0.127	0.127	0.132
Panel B: Weighted by inv. of SE(β_c)					
Individualism (IDV)	-0.103***	-0.116***	-0.107***	-0.131***	-0.134***
	(0.024)	(0.032)	(0.023)	(0.026)	(0.032)
Power Distance		-0.025			0.007
		(0.043)			(0.048)
Femininity - Masculinity			0.026		0.032
			(0.025)		(0.020)
Uncertainty Avoidance				-0.062**	-0.067**
				(0.026)	(0.030)
Observations	68	68	68	68	68
R^2	0.219	0.223	0.227	0.270	0.282

Robust standard in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Panel B weights observations (countries) by the inverse of the standard errors of estimates of β_c. I only countries for which all four cultural dimensions are available.

we can confirm that the negative correlation between individualism and social influence discussed throughout the paper does not come from anything particular to how Hofstede (2001) scores were constructed. Both autonomy scores from Schwartz's database have a negative and statistically significant association with reviewers' conformity to averages displayed on Tripadvisor.

Table 12: Effect of Individualism and Similar Cultural Measures on Social Influence

	Dependent Variable: Estimated Social Influence (β_c)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Intelectual Autonomy (std)	-0.014**		-0.011		-0.007**		-0.010**
	(0.006)		(0.008)		(0.004)		(0.004)
Affective Autonomy (std)		-0.016**	-0.010			0.004	0.010
		(0.007)	(0.009)			(0.006)	(0.007)
Hofstede Individualism (std)				-0.026***	-0.024***	-0.028***	-0.028***
				(0.005)	(0.005)	(0.006)	(0.005)
Observations	58	58	58	58	58	58	58
R^2	0.099	0.083	0.124	0.320	0.343	0.323	0.359

Robust standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$,$^{***}p < 0.01$

Notes: All regressions are weighted by the inverse of the standard errors of first stage estimate of β_c. I only include countries for which all regressors are available. All regressors are standardized to have zero mean and unit standard deviation.

Table 12 also highlights that, even though Schwartz (1994) measures of autonomy predict lower levels of conformity, they are not as powerful as Hofstede's omnibus individualism measure. Together, both measures of autonomy (column 3) predict only 12% of the variation in estimated β_c. On the other hand, the individualism score alone (column 4) predicts 32% of the variation in the same outcome variable. Another lesson from this analysis is that Hofstede's score seems more strongly correlated with affective rather than intellectual autonomy. The latter still predicts social influence even when I control for the effect of individualism (column 5). The former, however, becomes irrelevant, and its effect drops to zero when individualism is included as a regressor (column 6).

Finally, we can go back to one of the questions posed at this section's beginning. Is individualism the most relevant aspect of culture for explaining cross-country variation in reviewers' conformity to the prevailing average? Among the four core cultural dimensions in Hofstede (2001), the answer is yes. Individualism is the only one that predicts social influence on Tripadvisor. Moreover, looking at an alternative source of cultural data, variables that measure values similar to those that constitute Hofstede's individualism dimension also predict lower levels of conformity. These two pieces of evidence suggest that individualism is the most relevant aspect of culture to predict reviewers' propensity to conform to the average of previous reviews.

5 Within Country Evidence

In the previous section, I estimated cross-cultural variation in social influence on Tripadvisor from a cross-country perspective. Given the many dimensions along which people from different countries vary, cross-country comparisons may raise concerns related to omitted

variables. One particular concern is that consumers' perceptions of Tripadvisor's purpose may depend on where they are from. For example, among reviews left by American consumers, 72% were written by tourists, while among Italians, this number is 50%. [41] This illustrates the more general point that there could be selection on the type of reviewers using the platform in different countries.

5.1 Italian Regions

To mitigate concerns of this type, in this section, I focus on reviewers from Italy and explore variation in cultural/psychological traits across different regions of the country. There are a few reasons why I focus on Italy. First, it displays substantial variation in cultural norms across regions within its territory (see Putnam et al. (1992) for a historical account of this issue). Second, the European Social Survey (ESS), which reports participants' sub-national region of residence (NUTS 2 level), contains questions directly related to people's views of the importance of conforming to norms.[42] Lastly, Italians are among the most frequent reviewers in the sample, which implies that there are enough observations to conduct an analysis focused exclusively on them. Table 1 shows that only the UK and USA surpass Italy in terms of reviewer country of origin.

To measure variation in values across Italian regions, I use data constructed by Schulz et al. (2019).[43] They created a conformity-obedience index based on four questions from the European Social Survey (ESS). Respondents are asked to rate, on a six-point scale, the extent to which they think similarly to a person who says that: "(i) it is important to her/him always to behave properly. She/he wants to avoid doing anything people would say is wrong. (ii) She/he believes people should do what they are told. She/he thinks people should always follow the rules, even when nobody is watching. (iii) It is important to her/him to be humble and modest. She/he tries not to draw attention to herself/himself. (iv) Tradition is important to her/him. She/he tries to follow the customs handed down by her/his religion or her/his family.".

The final index from Schulz et al. (2019) is computed at the survey respondent level and is the average over these four questions, where higher scores indicate a higher motivation to conform.[44] To construct region-level scores, I take averages over ESS respondents from the same NUTS 2 region. Furthermore, to make units more easily comparable to previous analyses presented in the paper, I normalize the index such that its values fall between 0 and 1. Finally, to link this data with the Tripadvisor sample, I used the fact that most reviewers report the country and the city they live in. With information on the city of residence, I could assign a NUTS 2 region to approximately 96% of all reviews from Italian users.

[41] I define reviews as being submitted by tourists in all cases when the restaurant being reviewed is not in the country reported by the consumer as her place of residence.

[42] This allows me to compute averages over answers of respondents from the same region to obtain regional-level measures of people's self-reported propensity to conform.

[43] In Schulz et al. (2019), self-reported motivation to conform is one of many psychological measures used as dependent variables. Their main contribution is to quantify the effect of practices adopted by the Western Catholic Church during the Middle Ages on today's levels of individualism in the psychology of different populations.

[44] To avoid contaminating the conformity index with individual-level overall *closeness*, the authors subtracted from it the person's mean answer over all the 21 human value questions included in the survey.

5.2 Analysis and Results

I start with an analysis analogous to the one from Section 3. That is, I estimate a regression that includes the interaction between the discontinuity indicator and the conformity index from the ESS. Formally, I estimate the following equation:

$$\begin{aligned} Rat_{ijn} = \beta_0 + \left(\beta_1 + \beta_2\, Conf_{r(i)}\right) \times \mathbb{1}\{Avg_{jn-1} \geq k\} \; + \\ \beta_3\, Conf_{r(i)} + \beta_4 Avg_{jn-1} + X'_{ijn}\gamma + \epsilon_{ijn} \end{aligned} \tag{4}$$

In the equation above, most terms are the same as in Equation 1. The single difference is that, as a potential moderator of reviewers' response to Tripadvisor's headline average ratings, instead of Hofstede (2001) individualism score, I use the previously described conformity index, denoted by $Conf_{r(i)}$. The subscript $r(i)$ indicates that this variable is defined at the (Italian) region level. Given how $Conf_r$ is constructed, I expect it to positively correlate with reviewers' propensity to adjust their reviews toward a restaurant's average rating.

Table 13 shows the results of this analysis. Similarly to before, the bandwidth is 0.1 stars around Tripadvisor rounding cutoffs. Since I am focusing on a single country of origin, I pool all cutoffs together to increase the number of observations.[45] Column 1 shows the simplest model, where the only regressors are restaurants' underlying (continuous) average rating, reviewers' conformity index $Conf_{r(i)}$, an indicator for whether the underlying average is displayed rounded up, and its interaction with the conformity index. Discontinuity effects, shown in the first two rows, parallel the findings from Table 2, but with the opposite signs. This result is what we should expect, given that the conformity index $Conf_r$, in many respects, measures values emphasized by cultures with low levels of individualism.

The most critical coefficient is the interaction between the above cutoff indicator and the psychological measure of motivation to conform. It tells us that the difference in response to a 0.5-star increase in the average displayed by Tripadvisor varies by 0.29 stars between the regions with the lowest and highest values of $Conf_r$, a significant effect of culture on reviewer behavior. Regarding the overall social influence effect for consumers from a given region (i.e., the combination of coefficients in the first and second rows), my estimates suggest that reviewers from the less conforming regions of Italy tend to deviate from a restaurant's overall average when submitting their own ratings. On the other hand, on average, consumers from the areas where survey respondents reported caring more about conforming to social norms indeed display conformity behavior when rating Tripadvisor restaurants. Thus, my results help to validate the measure of conformity obtained from the ESS.

From columns 2 through 5, each specification becomes slightly more demanding by controlling for additional dimensions of heterogeneity. Column 2 allows the effect of the underlying average rating to differ depending on whether it falls above or below a rounding cutoff. Columns 3 and 4 add reviewer and restaurant-level covariates, respectively. Finally, in column 5, I add restaurant fixed effects. The estimated influence of a restaurant's headline average

[45] As in the cross-country analysis, the 4.25 cutoff is responsible for about half of the observations and, to a large extent, drives the pooled cutoff results.

Table 13: Prevailing Average Rating and Value of the Next Rating (Italian Reviewers)

	Dependent Variable: Rating r_{ijn}				
	(1)	(2)	(3)	(4)	(5)
Discontinuity Effects					
Above Cutoff	-0.179**	-0.181**	-0.175**	-0.172**	-0.228**
	(0.089)	(0.088)	(0.088)	(0.087)	(0.094)
Above Cutoff × Conform. Index	0.294**	0.296**	0.284**	0.279**	0.347**
	(0.141)	(0.141)	(0.141)	(0.140)	(0.149)
Main Controls					
ESS Conformity Index	-0.066	-0.068	-0.067	-0.068	-0.123
	(0.105)	(0.105)	(0.104)	(0.105)	(0.116)
Avg Rating	0.801**	1.072**	1.111**	1.130**	-0.090
	(0.350)	(0.495)	(0.497)	(0.497)	(0.679)
Avg Rating × Above Cutoff		-0.523	-0.594	-0.628	0.529
		(0.758)	(0.756)	(0.738)	(0.958)
Reviewer Controls	No	No	Yes	Yes	Yes
Restaurant Controls	No	No	No	Yes	Yes
Cutoff FE	Yes	Yes	Yes	Yes	Yes
Restaurant FE	No	No	No	No	Yes
Observations	8,958	8,958	8,958	8,958	7,811

Standard errors are clustered at the restaurant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: Sample only include Italian reviewers. The conformity index is based on answers to the European Social Survey (data from Schulz et al. (2019)) and computed at the regional level (there are 20 regions in Italy). All regressions use a bandwidth of 0.1. Reviewer Controls: tourist dummy, log of user number of reviews, smartphone dummy. Restaurant controls: 3-level categorical variable indicating restaurant price and the log of restaurant number of prior reviews.

rating on the value of its next rating is reasonably stable across specifications, the exception being the fixed effects specification, which delivers somewhat larger estimates.

Next, I estimate a model like the one in Equation 5 but substituting the conformity index $Conf_r$ by region fixed effects. That is, I estimate the model:

$$Rat_{ijn} = \theta_{r(i)} + \alpha Avg_{jn-1} + \beta_{r(i)} \times \mathbb{1}\{Avg_{jn-1} \geq k\} + X'_{ijn}\gamma + \epsilon_{ijn} \tag{5}$$

This type of specification still treats reviewers from the same region as equally prone to be influenced by the average star ratings. However, it is more flexible because the magnitudes of social influence effects across regions with similar levels of $Conf_r$ are allowed to be arbitrarily different. I am agnostic about the region-level characteristics that may affect reviewers' inclination to conform to (or deviate from) a restaurant's current average rating. This approach is the same from Section 4, where I include a more detailed discussion of this model.

Figure 8 plots the estimates of β_r against the survey-based measure of motivation to conform. Although there are few data points, we see a clear positive correlation. On average, Italian regions where ESS participants reported a stronger emphasis on conformity values are the ones for which estimates of social influence effects tend to be more positive (or less negative). The two regions with the largest (positive) estimates of social influence β_r, Liguria and Umbria, also happen to be the two regions with the largest values of the survey-based index $Conf_r$.

In general, there is a reasonable argument that people who grow up in cultures where drawing attention to oneself is considered a negative thing will probably be less willing to speak their minds when their opinion differs from the one held by the majority. In the context of online reviews, this translates as a higher propensity to conform to the average opinion of others. That is the positive correlation shown in Figure 8. To provide evidence that this correlation captures a true connection between the kind of cultural value measured in the survey and actual reviewer behavior, I switch attention to an alternative cultural value measured in the ESS.

I focus on generalized trust for two reasons.[46] First, trust, a fundamental component of many economic interactions, has been extensively studied in the literature and shown to correlate with important outcomes such as income per capita (Algan and Cahuc, 2010). Second, in the specific context of online consumer ratings, there are no apparent reasons to think that trust would affect reviewers' reactions to the average of past ratings. Figure 9 shows no association between estimated social influence and the ESS region-level trust index. This analysis provides evidence that the positive correlation between reviewers' propensity to conform and the ESS conformity index (Figure 8) is not accidental.

6 Conclusion

This paper provides evidence of cross-cultural variation in social influence. Focusing on the context of online reviews, I show that the influence of a restaurant's average rating on the next

[46]The ESS includes the following question: "Generally speaking, would you say that most people can be trusted or that you cannot be too careful in dealing with people?". Participants are asked to respond with a number between 0 and 10, where 10 means the highest level of trust. I compute region-level averages and normalize them to fall between 0 and 1.

Figure 8: Response to Headline Average Rating and ESS Conformity-Obedience Index

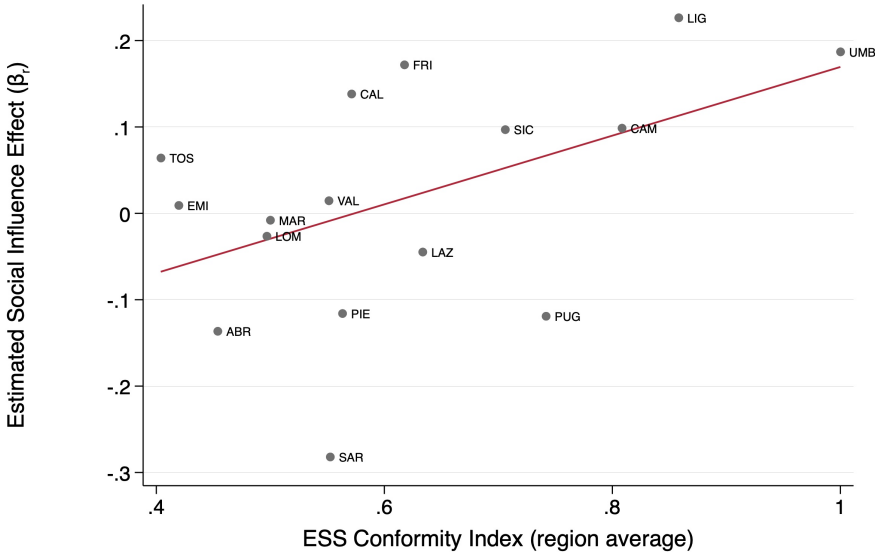

Notes: Each dot is one Italian region. Scatter plot includes regions for which I observe at least 100 reviews, ($N = 16$). In total, there are 20 regions in Italy. Conformity index computed using data from Schulz et al. (2019). It is based on answers to four questions in the ESS (European Social Survey). Higher scores indicate stronger emphasis on the values of conformity and obedience. Social influence effect estimated with specification analogous to Equation 2, but at the Italian regional level instead of country level.

Figure 9: Response to Headline Average Rating and ESS Trust Index

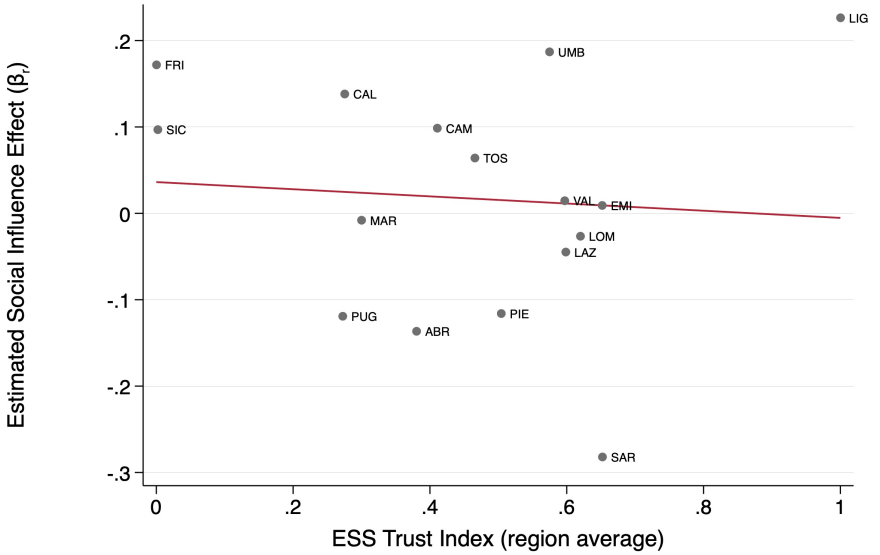

Notes: Each dot is one Italian region. Scatter plot includes regions for which I observe at least 100 reviews, ($N = 16$). In total, there are 20 regions in Italy. Trust index computed using data from Schulz et al. (2019). It is based on answers to the question: "Generally speaking would you say that most people can be trusted or that you can't be too careful in dealing with people?". Higher scores indicate more trust. Social influence effect estimated with specification analogous to Equation 2, but at the Italian regional level instead of country level.

review it receives depends on the reviewer's country of origin. To identify reviewers' response to the average of past reviews (social influence), I exploit that Tripadvisor displays average ratings rounded to the nearest half-star.

I provide three main results. First, I show that the individualism dimension of culture predicts cross-country variation in social influence. In particular, reviewers from countries with more individualistic cultures display a lower tendency to conform to the average of past reviews. Second, I examine the role of other dimensions of culture. My evidence suggests that individualism measures specific cultural and psychological traits connected to people's motivation to conform, which other cultural measures, such as generalized trust, do not capture. Third, using cross-regional variation within Italy, I show that the correlation between individualistic cultural values and social influence also holds across reviewers from the same country.

My findings contribute to the discussion on the consequences of our increasing reliance on aggregated crowd-sourced information to make decisions. Tripadvisor's evidence suggests that the extent of social influence in these environments is culture-specific, implying that the speed of learning from reviews also depends on culture. Specifically, the higher conformity displayed by reviewers from less individualistic cultures implies slower aggregate learning from online reviews. This finding indicates that the effects of the increased role played by online review platforms may differ depending on a society's cultural norms. For example, faster information flow is generally accompanied by increased firm turnover and higher levels of average quality provided by the market. My results suggest that online review platforms will more strongly affect these outcomes in societies with individualistic values.

More broadly, this paper also informs policy debates on issues where social influence plays an important role. For example, evidence exists that social influence is critical to understand phenomena as varied as the spread of fake news, voting choices, judicial decisions, or recreational habits such as smoking and drinking. Policies dealing with these kinds of issues must consider the cultural environment. This observation is particularly relevant in cases where international organizations, which work across different countries, have to design programs without enough knowledge of local norms.

My results do not imply a ranking of cultures. Although the connection between individualism and faster learning suggests positive welfare consequences, this feature is specific to this setting. In other situations, such as in collective action problems, too much individualism may be counterproductive. The main message is that culture shapes our underlying motivations, affecting behavior across various situations. This paper focuses on one of them. The literature on psychology and economics has successfully used experiments to document cross-cultural variation in values and beliefs. However, evidence using observational data in real-world settings is still sparse. More research in this direction is needed. The internet, with large amounts of data on the behavior of individuals worldwide, might prove helpful in this effort.

References

D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar. Fast and slow learning from reviews. Technical report, National Bureau of Economic Research, 2017.

G. A. Akerlof. A theory of social custom, of which unemployment may be one consequence. *The quarterly journal of economics*, 94(4):749–775, 1980.

Y. Algan and P. Cahuc. Inherited trust and growth. *American Economic Review*, 100(5):2060–92, 2010.

M. Anderson and J. Magruder. Learning from the crowd: Regression discontinuity estimates of the effects of an online review database. *The Economic Journal*, 122(563):957–989, 2012.

E. Anenberg, C. Kuang, and E. Kung. Social learning and local consumption amenities: Evidence from yelp. *The Journal of Industrial Economics*, 70(2):294–322, 2022.

S. E. Asch. Studies of independence and conformity: I. a minority of one against a unanimous majority. *Psychological monographs: General and applied*, 70(9):1, 1956.

A. V. Banerjee. A simple model of herd behavior. *The quarterly journal of economics*, 107(3): 797–817, 1992.

G. S. Becker. A note on restaurant pricing and other examples of social influences on price. *Journal of political economy*, 99(5):1109–1116, 1991.

B. D. Bernheim. A theory of conformity. *Journal of political Economy*, 102(5):841–877, 1994.

S. Bikhchandani, D. Hirshleifer, and I. Welch. Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of economic perspectives*, 12(3):151–170, 1998.

M. H. Bond and P. B. Smith. Cross-cultural social and organizational psychology. *Annual review of psychology*, 47(1):205–235, 1996.

R. Boyd and P. J. Richerson. *Culture and the evolutionary process*. University of Chicago press, 1988.

R. B. Cialdini and N. J. Goldstein. Social influence: Compliance and conformity. *Annu. Rev. Psychol.*, 55:591–621, 2004.

P.-P. Combes, G. Duranton, and L. Gobillon. Spatial wage disparities: Sorting matters! *Journal of urban economics*, 63(2):723–742, 2008.

W. D. Dai, G. Jin, J. Lee, and M. Luca. Aggregation of consumer ratings: an application to yelp. com. *Quantitative Marketing and Economics*, 16(3):289–339, 2018.

M. Deutsch and H. B. Gerard. A study of normative and informational social influences upon individual judgment. *The journal of abnormal and social psychology*, 51(3):629, 1955.

C. L. Fincher, R. Thornhill, D. R. Murray, and M. Schaller. Pathogen prevalence predicts human cross-cultural variability in individualism/collectivism. *Proceedings of the Royal Society B: Biological Sciences*, 275(1640):1279–1285, 2008.

R. H. Frank. Under the influence. In *Under the Influence*. Princeton University Press, 2021.

Y. Gorodnichenko and G. Roland. Which dimensions of culture matter for long-run growth? *American Economic Review*, 101(3):492–98, 2011.

J. Henrich. *The WEIRDest people in the world: How the West became psychologically peculiar and particularly prosperous*. Penguin UK, 2020.

J. Henrich, S. J. Heine, and A. Norenzayan. The weirdest people in the world? *Behavioral and brain sciences*, 33(2-3):61–83, 2010.

G. Hofstede. *Culture's consequences: Comparing values, behaviors, institutions and organizations across nations*. Sage publications, 2001.

Y. Hong, N. Huang, G. Burtch, and C. Li. Culture, conformity and emotional suppression in online reviews. *Journal of the Association for Information Systems, Forthcoming, Fox School of Business Research Paper*, (16-020), 2016.

N. Hu, J. Zhang, and P. A. Pavlou. Overcoming the j-shaped distribution of product reviews. *Communications of the ACM*, 52(10):144–147, 2009.

H. Kim and H. R. Markus. Deviance or uniqueness, harmony or conformity? a cultural analysis. *Journal of personality and social psychology*, 77(4):785, 1999.

R. A. King, P. Racherla, and V. D. Bush. What we know and don't know about online word-of-mouth: A review and synthesis of the literature. *Journal of interactive marketing*, 28(3):167–183, 2014.

J. Lafky. Why do people rate? theory and evidence on online ratings. *Games and Economic Behavior*, 87:554–570, 2014.

M. Luca. Reviews, reputation, and revenue: The case of yelp. com. *Com (March 15, 2016). Harvard Business School NOM Unit Working Paper*, (12-016), 2016.

J. McCrary. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of econometrics*, 142(2):698–714, 2008.

L. Muchnik, S. Aral, and S. J. Taylor. Social influence bias: A randomized experiment. *Science*, 341(6146):647–651, 2013.

R. D. Putnam, R. Leonardi, and R. Y. Nanetti. *Making democracy work: Civic traditions in modern Italy*. Princeton university press, 1992.

J. F. Schulz, D. Bahrami-Rad, J. P. Beauchamp, and J. Henrich. The church, intensive kinship, and global psychological variation. *Science*, 366(6466):eaau5141, 2019.

S. H. Schwartz. Are there universal aspects in the structure and contents of human values? *Journal of social issues*, 50(4):19–45, 1994.

K. H. Yoo and U. Gretzel. What motivates consumers to write online travel reviews? *Information Technology & Tourism*, 10(4):283–295, 2008.

Appendix

A Tripadvisor's Interface

Figure 10 shows an example of a restaurant listing on Tripadvisor and illustrates the process a reviewer likely goes through when rating a restaurant.[47] To submit a rating to a given restaurant, consumers need to find its listing page on Tripadvisor. They likely do that by searching for key terms, such as the restaurant's name and city. This step is represented in Figure 10a, where the search term "sucre madrid" is used. A list of options comes up, and the user chooses the relevant listing to click on. At this stage, the restaurant's current average rating appears highlighted and is proportional to the number of full and half bubbles filled in green. "Bubbles" are analogous to the concept of "stars", used by most other platforms. Given that in the previous literature and non-academic environments, most people refer to average ratings as star ratings, I stick to this terminology throughout the paper.

Once reviewers enter the restaurant listing, they see a summary with basic restaurant information. Some of these were already visible in the search scroll-down part of the process, such as the current number of reviews and average rating held by the restaurant, its rank within Tripadvisor, price category, and cuisine type. This stage of the process is shown in Figure 10b.

Next, users will likely scroll down the page to see the reviews section, shown in Figure 10c. Here, one can see a summarized version of previous reviews received by the restaurant. Notably, the (rounded) average rating is highlighted on top, both in numerical form and by the colored bubbles. The distribution of all previous reviews is also available. With this information, reviewers could, in principle, compute the restaurant's continuous average, which in this example is 4.5315. However, in practice, they probably stick to the rounded version, 4.5, as their assessment of what other people think of this restaurant.

Lastly, users willing to submit a rating of their own will click on the pencil icon on the bottom right of the page, which brings them to a page where they can choose how many bubbles (stars) to assign their restaurant experience. At this stage, shown in Figure 10d, the restaurant's current average does not appear anymore. Thus, social influence, if it exists, occurs when consumers are reading over the restaurant's listing and before clicking the pencil icon to submit their review. Finally, note that Tripadvisor explicitly asks "How would you rate your experience?", which suggests the platform's goal is for consumers to report their experience as it was, without conditioning their choice of rating on the current average held by a restaurant.

[47] Figure 10 shows the interface seen by a user of Tripadvisor's mobile application. The interface for users of desktop computers is similar.

Figure 10: Tripadvisor's Interface on its Mobile Application

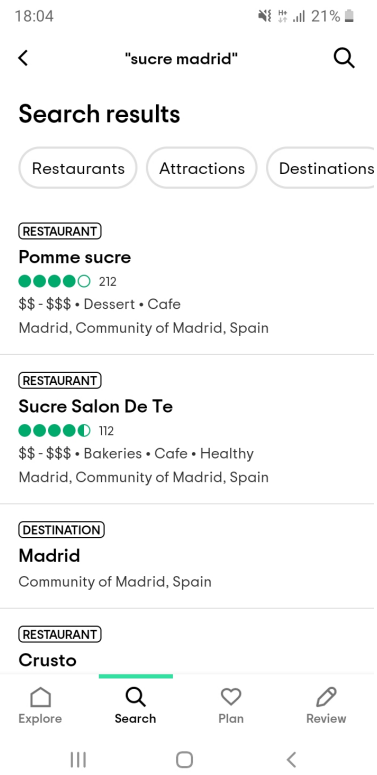

(a) Search Page

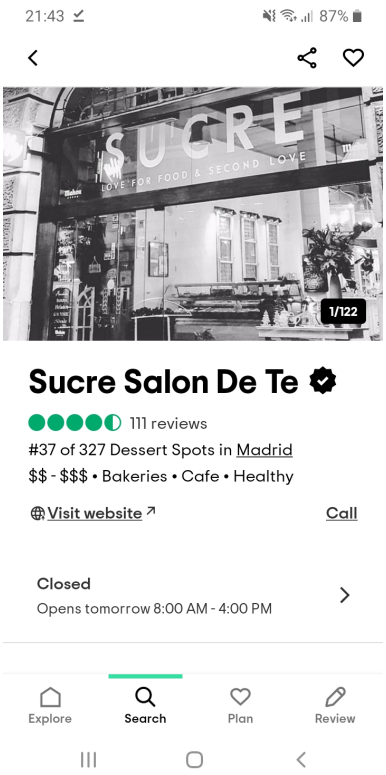

(b) Restaurant Information

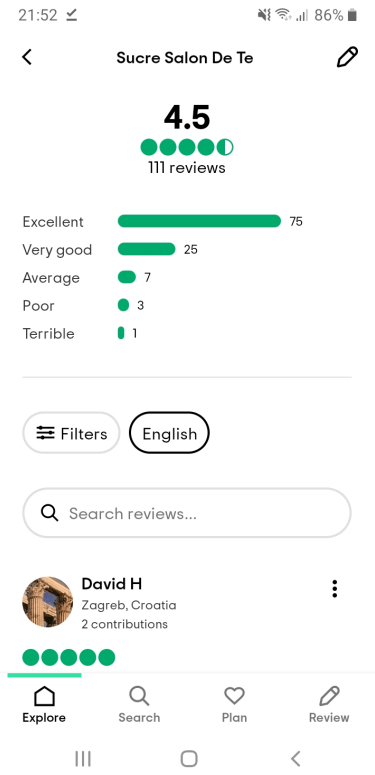

(c) Previous Reviews

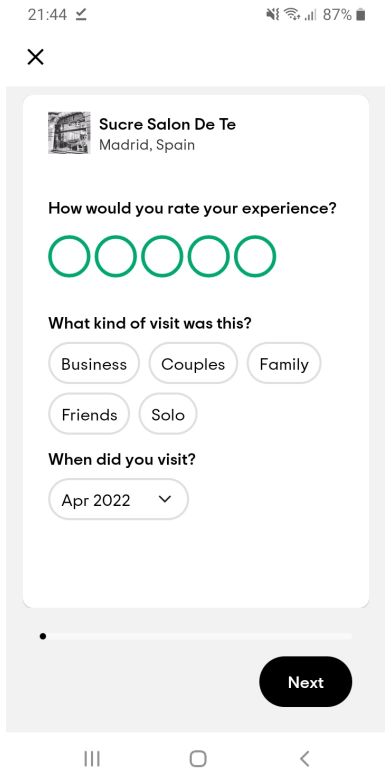

(d) Submit a Rating

B Hofstede's Individualism Score

Below I list the 14 core questions used by Hofstede (2001) to determine individualism scores. These questions were used in all survey waves. Final country scores were based on factor analysis applied to answers to these and some additional questions that were only used in specific waves. For the complete list of questions and more details on the factor analysis results, see Chapter 5 of Hofstede (2001).

All questions followed the format of "How important is it to you to ... ". For example, "How important is it to you to fully use your skills and abilities on the job?". Below are the 14 questions used in all survey waves:

1. Have challenging work to do—work from which you can get a personal sense of accomplishment.
2. Live in an area desirable to you and your family.
3. Have an opportunity for high earnings.
4. Work with people who cooperate well with one another.
5. Have training opportunities (to improve your skills or learn new).
6. Have good fringe benefits.
7. Get the recognition you deserve when you do a good job.
8. Have good physical working conditions (good ventilation and lighting, adequate work space, etc.).
9. Have considerable freedom to adapt your own approach to the job.
10. Have the security that you will be able to work for your company as long as you want to.
11. Have an opportunity for advancement to higher-level jobs.
12. Have a good working relationship with your manager.
13. Fully use your skills and abilities on the job.
14. Have a job which leaves you sufficient time for your personal or family life.

C Additional Analyses

Choosing the sample period. Figure 11 shows the evolution of the total number of reviews (right y-axis) and the average individualism score (left y-axis). It also highlights, within thin dotted lines, the period I include in the final sample on which social influence effects are estimated. Figure 11 illustrates the reasons for restricting the analysis to the period between 2015 and 2019. I drop years before 2015 because the average individualism score (IDV) decreases over time. I drop years 2020, 2021, and (part of) 2022 because, due to the pandemic, the number of reviews drops precipitously in this period.

Distribution of Individualism. Figure 12 plots the distribution of reviews by reviewers' individualism score. As discussed in Table 1 of Section 2, the distribution is far from uniform, with substantially more mass on the high end of the individualism score. Tripadvisor is more popular with consumers from countries that score high in individualism, especially in western Europe and the United States. There is a clear difference between these two regions, however.

Figure 11: Number of Reviews (right) and Average Individualism Score (left)

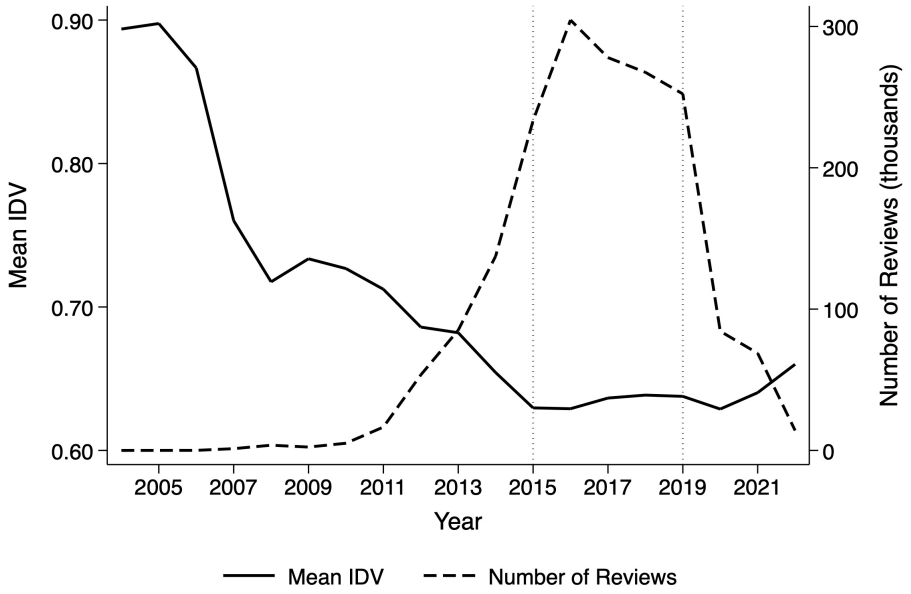

Notes: All numbers computed using the raw (full) sample scraped from Tripadvisor. Thin dotted lines demarcate the time period used in the final sample. "IDV" stands for Hofstede (2001) individualism score.

Europeans are as likely to submit a review to a restaurant in their home countries as they are for restaurants abroad. However, reviews from Americans in 75% of the time are submitted to restaurants outside the US.

Baseline Results by Cutoff. Table 14 shows the results of the baseline analysis for other cutoffs for which there is enough data. I focus on the model that includes individualism-specific slopes of ratings on the running variable (underlying average rating). Moreover, I show results with and without restaurant fixed effects. I include separate analyses for the 3.75 and 4.75 cutoffs, the second and third most frequent cutoffs regarding the number of observations. For the 3.75 threshold, results go in the same direction as for the 4.25. However, the coefficient on the role of individualism in shaping social influence is not significant. For the 4.75 cutoff, the rounding of average ratings does not affect reviewer rating behavior. In such situations, restaurants being reviewed have high quality and get a large fraction of 5-star reviews. Since ratings are censored at 5, there is little room for the rounding up of headline averages to exert an additional effect on reviewers' rating choices. Lastly, given the small fraction of observations around one of the cutoffs below 3.75, I do not look at them separately.

Test for the Continuity of the Density of Average Ratings. I carry out a density test based on McCrary (2008). Focusing on the cutoff with the largest number of observations, 4.25, I show that the density of observed average ratings does not present discontinuities around the cutoff. Figure 13 shows the distribution of average within a 0.1 bandwidth around the 4.25 cutoff. Some values are more frequent than others, but this is not related to whether it is above or below the rounding cutoff. Table 15 presents results of a regression version of the test. I split the range of values of average ratings used in the baseline analysis (Section 3.2.1) into 250 bins of 0.0008 stars in size. I then compute the fraction of all observations within each bin and run the regression using this fraction as the outcome. Results show that, even though the range of values of a given bin predicts its density (i.e., the negative coefficient on the second row), there is no evidence of a discontinuous jump in the relationship when we cross the rounding cutoff.

Table 14: Effect of Rounded Up Average and Individualism, by Cutoff

	All Cutoffs		3.75 Cutoff		4.25 Cutoff		4.75 Cutoff	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Above Cutoff	0.073***	0.088***	0.121*	0.160**	0.119***	0.154***	0.024	0.100
	(0.018)	(0.020)	(0.063)	(0.069)	(0.042)	(0.045)	(0.101)	(0.109)
Above Cutoff × IDV	-0.095***	-0.082***	-0.132	-0.146	-0.172***	-0.180***	0.019	-0.096
	(0.024)	(0.026)	(0.101)	(0.107)	(0.064)	(0.067)	(0.145)	(0.156)
Cutoff FE	Yes	Yes	No	No	No	No	No	No
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Restaurant FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	142,206	140,963	39,867	39,218	75,598	74,914	14,271	13,877

Standard errors are clustered at the restaurant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: All columns have culture specific slopes and basic covariates (same covariates as in Table 1).

Table 15: Testing the Continuity of the Density of Ratings

Dependent Variable: Fraction of Reviews in Bin			
	(1)	(2)	(3)
Above Cutoff	0.001	0.000	0.010
	(0.001)	(0.000)	(0.016)
Average Rating	-0.014***	-0.011***	-0.012***
	(0.005)	(0.004)	(0.004)
Average Rating × Above Cutoff	0.008	-0.004	-0.005
	(0.010)	(0.009)	(0.009)
Average IDV		-0.061***	-0.055***
		(0.011)	(0.014)
Above Cutoff × Average IDV			-0.015
			(0.024)
Observations	250	250	250

Robust standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: Number of observations equals number of 0.0008 bins between 4.15 and 4.35.

Figure 12: Distribution of Reviews by Individualism Score

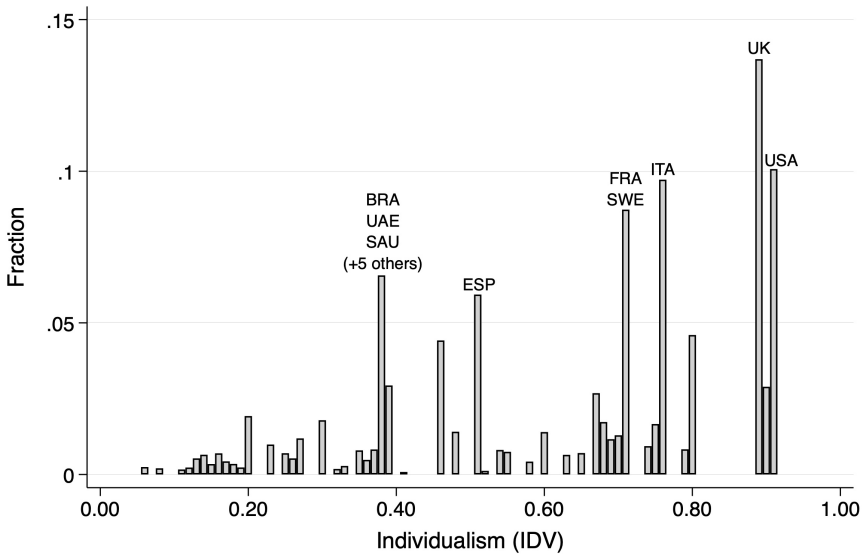

Notes: Histogram computed using the final sample. Each bar represents a distinct individualism score. Most frequent countries are indicated. Although not common, some countries have the exact same scores.

Baseline Results by Reviewer Level of Activity. I estimate the specification in column 2 of Table 2 for different subsamples, defined by the level of activity reviewers present on Tripadvisor. Specifically, I define subsamples by the number of reviews users had posted on the platform up to the point when the data was extracted. More casual reviewers, who have submitted only a few reviews to Tripadvisor, are more likely to be under-reporting. Thus, as explained in Section 3.3.3, I expect the effects of average ratings on consumers' rating choice to be closer to zero for these reviewers. That is what Table 16 shows, the social influence effect increases as I restrict the analysis to reviewers with higher activity levels.

Table 16: Baseline Results by User Number of Reviews

	User Number of Reviews				
	All	5+	10+	20+	40+
Above Cutoff	0.034**	0.038**	0.051***	0.073***	0.075***
	(0.016)	(0.016)	(0.017)	(0.018)	(0.020)
Above Cutoff × IDV	-0.034	-0.041*	-0.060***	-0.095***	-0.108***
	(0.021)	(0.021)	(0.022)	(0.024)	(0.026)
Cutoff FE	Yes	Yes	Yes	Yes	Yes
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes	Yes
Restaurant FE	No	No	No	No	No
Observations	209,280	183,030	166,658	142,206	109,847

Standard errors are clustered at the restaurant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: Dependent variable are individual ratings. Specification are the same as the one used on column 2 of Table 2, with a linear effect of the true underlying average, individualism, and baseline covariates.

Figure 13: Density Test Based on McCrary (2008)

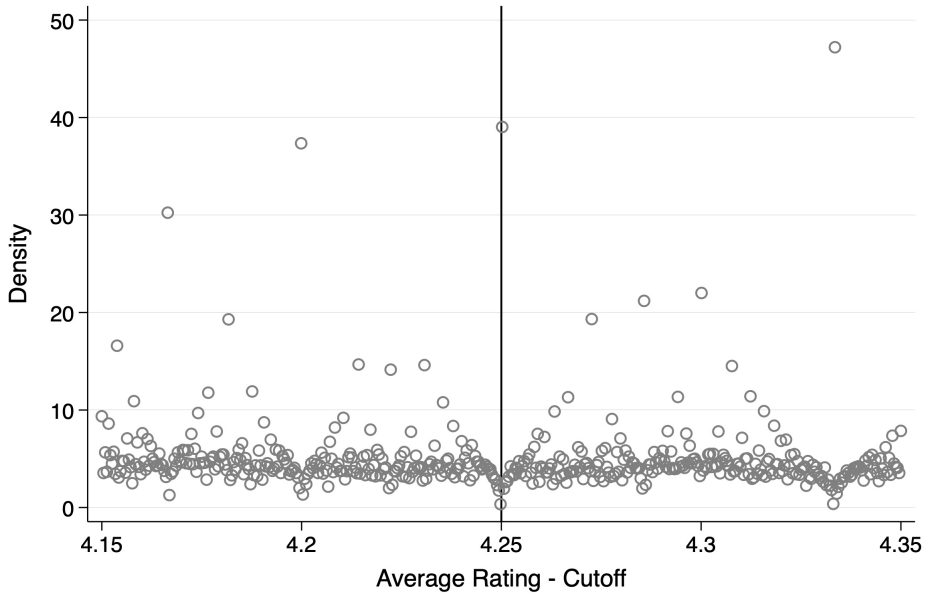

Discrete Outcome. Table 17 shows that results are robust to explicitly taking into account that ratings, the outcome variable, are discrete. I reproduce analyses analogous to what was presented in columns 2 and 5 of Table 1 and show that reviewers react to the headline average ratings they see displayed on Tripadvisor. In particular, when Tripadvisor displays a rounded-up average rating, reviewers from collectivist countries become more likely to report a higher rating. Similarly to the results from the linear regression shown in the main body of the paper, this effect decreases in the level of individualism in the reviewer's country, even becoming negative for the most individualistic countries. Table 17 also shows that there is no evidence that the effect of the running variable Avg_{ijn} changes depending on whether it falls to the left or the right of the rounding cutoffs.

Bandwidth size. Table 18 shows that the results of the baseline analysis are not sensitive to the choice of bandwidth size. I focus on the specification that allows the slope of the running variable to be culture-specific (i.e. interacted with individualism) and estimate the model for five different bandwidths. Moving from the smallest bandwidth (0.05) to the largest (0.15), the number of observations increases threefold, but estimated discontinuity effects remain stable. This stability in the estimates reassures us that results are not driven by bandwidth choice.

Variations in the Effect of the Running Variable. Table 19 shows that results are not sensitive to the choices of how underlying average ratings (Avg_{ijn}) affect the next rating received by a restaurant. This table brings different variations of the specification presented in column 2 of Table 2, with baseline results. Table 19 shows that the discontinuity effect is robust to allowing the linear effect of Avg_{ijn} to vary depending on whether it is below or above a rounding cutoff (column 2). Moreover, the table also shows that results are robust to using a quadratic rather than a linear effect of Avg_{ijn} (columns 3 and 4).

Effect of Different Sample Choices. Table 20 shows the effect of each sample choice made in getting from the raw data to the main sample used for estimation throughout the paper. There are three main dimensions over which the raw data differs from the main sample. First, the main sample is restricted to the years between 2015 and 2019. Second, it only includes ratings submitted when the restaurant had more than ten prior reviews. Third, it is restricted to

Table 17: Baseline Results Using Ordinary Regressions (discrete outcome)

	Dependent Variable: Rating r_{ijn}			
	Ordered Probit		Ordered Logit	
	(1)	(2)	(3)	(4)
Discontinuity Effects				
Above Cutoff	0.079***	0.077***	0.130***	0.126***
	(0.020)	(0.028)	(0.035)	(0.048)
Above Cutoff × IDV	-0.108***	-0.113***	-0.175***	-0.189***
	(0.027)	(0.037)	(0.047)	(0.062)
Selected Controls				
Avg Rating	0.805***	0.859***	1.440***	1.583***
	(0.154)	(0.208)	(0.261)	(0.353)
Average Rating × Above Cutoff	0.134	0.282	0.150	0.355
	(0.214)	(0.304)	(0.362)	(0.506)
Individualism (IDV)	0.106***	0.160***	0.200***	0.308***
	(0.023)	(0.030)	(0.038)	(0.051)
Cutoff FE	Yes	No	Yes	No
Restaurant City FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	142,434	76,100	142,434	76,100
Sample	All Cutoffs	4.25 Cutoff	All Cutoffs	4.25 Cutoff

Standard errors are clustered at the restaurant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: All regressions use a bandwidth of 0.1. In addition to the coefficients shown, all models include baseline covariates: tourist dummy, log of user number of reviews, smartphone dummy, a 3-level categorical variable indicating restaurant price, and log of restaurant number of prior reviews.

Table 18: Robustness to Different Bandwidth Sizes

	Bandwidth Size				
	0.050	0.075	0.010	0.125	0.150
Above Cutoff	0.061**	0.055**	0.072***	0.064***	0.059***
	(0.027)	(0.021)	(0.018)	(0.016)	(0.015)
Above Cutoff × IDV	-0.100***	-0.083***	-0.095***	-0.085***	-0.076***
	(0.035)	(0.028)	(0.024)	(0.021)	(0.019)
Individualism (IDV)	-0.916***	-0.799***	-0.583***	-0.587***	-0.606***
	(0.223)	(0.183)	(0.155)	(0.139)	(0.128)
Average Rating	0.649**	0.670***	0.555***	0.598***	0.597***
	(0.285)	(0.153)	(0.099)	(0.074)	(0.058)
Average Rating × IDV	0.258***	0.224***	0.174***	0.171***	0.174***
	(0.053)	(0.044)	(0.037)	(0.033)	(0.031)
Cutoff FE	Yes	Yes	Yes	Yes	Yes
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes	Yes
Restaurant FE	No	No	No	No	No
Observations	67,747	105,152	142,206	178,697	214,146

Standard errors are clustered at the restaurant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Dependent variable: individual ratings submitted when the prevailing average was close to a given rounding cutoff. Samples pools all cutoffs together. Specification allows for culture specific slope on true average ratings.

Table 19: Robustness to Different Ways to Model the Effect of the Running Variable

	Dependent Variable: Rating r_{ijn}			
	Linear Avg_{ijn}		Quadratic Avg_{ijn}	
	(1)	(2)	(3)	(4)
Main Effects				
Above Cutoff	0.0730***	0.0733***	0.0731***	0.0608***
	(0.0179)	(0.0179)	(0.0179)	(0.0221)
Above Cutoff × IDV	-0.0945***	-0.0944***	-0.0945***	-0.0944***
	(0.0237)	(0.0237)	(0.0237)	(0.0237)
Running Var. × Above Cutoff	No	Yes	No	Yes
Cutoff FE	Yes	Yes	Yes	Yes
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes
Observations	142,206	142,206	142,206	142,206
Sample	All Cutoffs	All Cutoffs	All Cutoffs	All Cutoffs

Standard errors are clustered at the restaurant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: All regressions use a bandwidth of 0.1. In addition to the running variable in the format specified in each column, all models include baseline covariates: tourist dummy, log of user number of reviews, smartphone dummy, 3 level categorical variable indicating restaurant price, log of restaurant number of prior reviews. In column 4, both the linear and quadratic effects of the running variable are interacted with the above cutoff dummy.

ratings submitted by frequent reviewers (with more than 20 reviews on Tripadvisor). Table 19 shows that the paper's main results are robust to the addition of all years and to the inclusion of reviews submitted when the restaurant had few previous reviews. Adding casual reviewers has a larger impact on results, which become smaller and lose statistical significance depending on the specification. See more details on sample construction in Section 2. For a discussion on why casual reviewers are not suited for my analysis, see Section 3.3.3.

Table 20: Effect of Different Sample Choices

	Main Sample	Main + All Years	Main + All n	Main + Casual Users
Panel A: Pooled Cutoffs				
Above Cutoff	0.073***	0.052***	0.059***	0.034**
	(0.018)	(0.016)	(0.017)	(0.016)
Above Cutoff × IDV	-0.095***	-0.072***	-0.074***	-0.034
	(0.024)	(0.021)	(0.023)	(0.021)
Cutoff FE	Yes	Yes	Yes	Yes
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes
Observations	142,206	190,037	154,112	209,280
Panel B: 4.25 Cutoff				
Above Cutoff	0.074***	0.052**	0.063***	0.030
	(0.024)	(0.022)	(0.023)	(0.023)
Above Cutoff × IDV	-0.103***	-0.073***	-0.093***	-0.038
	(0.031)	(0.028)	(0.030)	(0.028)
Rest. City x Month-Year FE	Yes	Yes	Yes	Yes
Observations	75,598	99,110	79,812	102,583

Standard errors are clustered at the restaurant level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Notes: Column names indicate the sample used. In Column 1, I work with the sample used throughout the paper. In each of the other columns, I add one of the dimensions over which the main sample differs from the raw data, year when the review was posted, number of prior reviews held by the restaurant, and casual reviewers, respectively. See Section 2 for details on the sample construction. All regressions use a bandwidth of 0.1. In addition to the coefficients shown, all models include the following covariates: the underlying continuous average rating, individualism score, tourist dummy, log of user number of reviews, smartphone dummy, a 3-level categorical variable indicating restaurant price, and log of restaurant number of prior reviews.