Istituto Nazionale Previdenza Sociale

agosto 2024 – numero 82

WorkINPS *Papers*

The End of Tourist Traps: The Impact of Review Platforms on Quality Upgrading

Dante Donati

ISSN 2532 -8565

In copertina: uno storico "Punto cliente" a Tuscania
INPS, Direzione generale, Archivio storico

I WORKINPS PAPER

Le basi dati amministrative dell'*INPS* rappresentano una fonte statistica unica per studiare scientificamente temi cruciali per l'economia italiana, la società e la politica economica: non solo il mercato del lavoro e i sistemi di protezione sociale, ma anche i nodi strutturali che impediscono all'Italia di crescere in modo adeguato. All'interno dell'Istituto, questi temi vengono studiati sia dai funzionari impiegati in attività di ricerca, sia dai VisitInps Scholars, ricercatori italiani e stranieri selezionati in base al loro curriculum vitae e al progetto di ricerca presentato.

I **WORKINPS** hanno lo scopo di diffondere i risultati delle ricerche svolte all'interno dell'Istituto a un più ampio numero possibile di ricercatori, studenti e policy markers.

Questi saggi di ricerca rappresentano un prodotto di avanzamento intermedio rispetto alla pubblicazione scientifica finale, un processo che nelle scienze sociali può chiedere anche diversi anni. Il processo di pubblicazione scientifica finale sarà gestito dai singoli autori.

Tommaso Nannicini

The End of Tourist Traps:

The Impact of Review Platforms on Quality Upgrading

Dante Donati

(Columbia Business School)

INPS.0072.26/07/2024.0000763 E

The End of Tourist Traps: The Impact of Review Platforms on Quality Upgrading

La Fine delle Trappole per Turisti: L'Impatto delle Piattaforme di Recensioni sulla Qualità*

Dante Donati
Columbia Business School and CESifo

July 18, 2024

Abstract

Asymmetric information can distort market outcomes. I study how cheaper access to review platforms affects consumers' behavior and firms' incentives to upgrade product quality in markets where information is traditionally limited. I first build a model of consumer search with firms' endogenous quality decisions. In this model, lower search costs reallocate demand toward higher-quality producers, raising firms' incentives to upgrade quality, and more so for firms selling *ex-ante* lower-quality products. I then use the access to online reviews to proxy for reductions in consumers' search costs and estimate its impact on the restaurant industry in Rome, exploiting the abolition of mobile roaming charges in the EU in 2017 for identification. Based on a unique dataset combining monthly information from Tripadvisor with administrative social-security records, I find that revenues and total employment in restaurants with *ex-ante* higher ratings grew by 3-10% after the internet price drop. In turn, the probability for lower-rating restaurants to exit the market doubled, while surviving lower-rating establishments hired workers with higher wages and better curricula, eventually improving their online reputation. My findings have implications for the role of review platforms in industries with information asymmetries.
Keywords: Review platforms; Asymmetric information; Search costs; Service industry; Quality

L'informazione asimmetrica può distorcere i risultati di mercato. In questo articolo, studio come una riduzione dei costi di accesso alle piattaforme di recensioni influenzi il comportamento dei consumatori e gli incentivi delle aziende a migliorare la qualità dei prodotti in mercati dove l'informazione è tradizionalmente limitata. In primo luogo, costruisco un modello di ricerca dei consumatori con decisioni endogene di qualità delle aziende. In questo modello, costi di ricerca più bassi riallocano la domanda verso produttori di qualità superiore, aumentando gli incentivi delle aziende a migliorare la qualità, soprattutto per le aziende che vendono prodotti di qualità inferiore *ex-ante*. Successivamente, utilizzo l'accesso alle recensioni online come proxy per la riduzione dei costi di ricerca dei consumatori e stimo il suo impatto sull'industria della ristorazione a Roma, sfruttando l'abolizione delle tariffe di roaming nell'UE nel 2017 per l'identificazione. Basandomi su un dataset unico che combina informazioni mensili da Tripadvisor con dati amministrativi dei registri INPS, trovo che i ricavi e l'occupazione totale nei ristoranti con valutazioni più alte *ex-ante* sono cresciuti del 3-10% dopo la riduzione dei costi di roaming. Di conseguenza, la probabilità di uscita dal mercato per i ristoranti con valutazioni più basse è raddoppiata, mentre i ristoranti con valutazioni più basse sopravvissuti hanno assunto lavoratori con maggior esperienza e salari più alti, migliorando alla fine la loro reputazione online. Questi risultati hanno implicazioni per il ruolo delle piattaforme di recensioni in industrie con asimmetrie informative.
Parole chiave: Piattaforme di recensioni; Informazione asimmetrica; Costi di ricerca; Industria dei servizi; Qualità

*Dante Donati, Assistant Professor, Columbia Business School; dd3137@gsb.columbia.edu; Personal website.
La realizzazione del presente articolo è stata possibile grazie alle sponsorizzazioni e le erogazioni liberali a favore del programma 'VisitINPS Scholars'. I thank Maria Petrova, Ruben Durante, Gianmarco León-Ciliotta and Sandro Shelegia for their invaluable mentorship. I am also grateful to Francesco Amodio, Edoardo Di Porto, Ruben Enikolopov, Chiara Farronato, Rosa Ferrer, Renata Gaineddenova, Jared Gars, Nicola Gennaioli, Matthew Gentzkow, Fabrizio Germano, Albrecht Glitz, Jonathan Hall, Eliana La Ferrara, Massimo Motta, David Krisztian Nagy, Paolo Naticchioni, Victor Orozco, Christian Peukert, Giacomo Ponzetto, Davide Proserpio, Nandan Rao, Abel Schumann, Ananya Sen, Marta Reynal-Querol, Alessandro Tarozzi and various seminar attendees for their comments and useful insights. I gratefully acknowledge generous financial support from IPEG. I thank Claudia Trisna Amici and Saidevi Fioravanzo for invaluable research assistance. All remaining errors are my own. The findings and conclusions expressed are solely those of the author and do not represent the views of INPS.

1 Introduction

Asymmetric information can distort market outcomes in different ways. For example, when product quality is imperfectly observed, its equilibrium levels are too low, and both consumers and producers can incur significant welfare losses (Akerlof 1970; Leland 1979). By guiding purchase decisions toward better quality products and services, review platforms improve consumer welfare (Reimers and Waldfogel 2021; Fang 2022). But how do they affect firms' quality choices and industry composition? The question has both managerial and regulatory implications. For firms, intensified competition from online word-of-mouth requires managers to strategically allocate resources and invest in reputation building. For regulatory bodies, customer reviews could represent a cost-effective complement to time- and labor-intensive monitoring interventions, such as sanitary inspections. Theoretically, review platforms should create reputation mechanisms and enhance firms' incentives to upgrade product quality (Tadelis 2016; Goldfarb and Tucker 2019). Yet, empirical evidence is scarce[1] due to two main limitations: (1) online feedback is usually endogenous, exacerbating any causal assessment;[2] (2) data on product quality are difficult to obtain and, if available, they often come from online ratings, which are not informative about firm-specific investment decisions and industry composition.

In this paper, I exploit a natural experiment and unique data from Tripadvisor matched to confidential administrative employer-employee records to study how cheaper access to review platforms affects consumers' behavior and firms' incentives to upgrade product quality. I derive theoretical predictions from a consumer-search model and empirically investigate the demand- and supply-side effects of online reviews on the restaurant industry in the province of Rome. I assemble a novel dataset combining monthly information on the entire historical records of reviews collected from Tripadvisor with rich administrative establishment-level data. For identification, I take advantage of the 2017 policy that abruptly abolished mobile roaming charges in the EU, generating an arguably exogenous variation in the costs for travelers to access online reviews. I find that cheaper access to reviews increases average product quality in the service industry. Consumers reallocate their demand toward higher-quality restaurants, while production choices of lower-quality establishments change. Some of them exit the market, while the surviving ones invest in quality upgrading.

While existing studies on review platforms rely exclusively on online ratings to proxy for quality (e.g., Proserpio and Zervas 2017; Ananthakrishnan et al. 2023), I leverage detailed firm-level employment data and directly look at changes in production costs (salaries) and hiring decisions (workers' curricula) to measure investment into service quality. In this respect, this paper expands upon previous work on the gig economy (e.g., Shin et al. 2023; Alyakoob and Rahman 2022), and sheds new light on the labor market consequences of review platforms. Moreover, this unique context allows me to estimate the aggregate effects of online word-of-mouth on firm dynamics (entry/exit),

[1] Some exceptions include Farronato and Zervas (2022) who examine restaurants' incentives to comply with hygiene regulation and Liu et al. (2021) who compare efficiency of Uber and taxi drivers.

[2] This is a well-known challenge to identifying the causal impact of any type of reputation and information disclosure on demand and supply (e.g., Eliashberg and Shugan 1997).

by leveraging a policy that was not deliberately designed to affect the restaurant industry.[3]

To guide the empirical exercise, I first build a model in which consumers with heterogeneous search costs engage in costly sequential search to buy one unit of a vertically differentiated product while firms with heterogeneous abilities endogenously select into production and compete in quality. I show that, when consumers' search costs decrease: (i) the demand faced by firms that were *ex ante* selling high-quality goods increases; (ii) the overall quality level in the industry improves, and this is driven by both the exit of lower-quality providers and the investment in quality upgrading of all surviving firms, especially of those selling *ex-ante* lower-quality products.

To test these predictions empirically, I assemble a novel dataset, which combines information from Tripadvisor, the most popular travel guidance platform in the EU, with administrative employer-employee records maintained at the Italian National Social Security Institute (INPS). For about 5,500 matched restaurants in the province of Rome, the data contain time-invariant information on firms' characteristics, as well as time-varying information on the number of Tripadvisor reviews, rating, date of opening and closure, revenues, number of employees, wages and the full employment history of the workers, observed at monthly or annual intervals between 2007 and 2019.

I then use the cheaper access to online reviews to proxy for reductions in consumer search costs and estimate its impact on the restaurant industry. The identification strategy exploits the approval by the European Parliament of new EU roaming legislation,[4] which abolished all charges for temporary roaming within the European Economic Area (EEA) as of June 15, 2017. In practice, before that day, all EU residents traveling within the EEA were charged an additional price for data consumption on top of the home network rate. After the policy, the same home network rate is applied, resulting in virtually free internet access for tourists while traveling.

I show that, after the policy, Tripadvisor users from EU countries became 1.4 times more likely to post reviews on their mobile devices as opposed to PC. Moreover, the volume of reviews written by EU users substantially increased after the policy compared to those from locals. By contrast, reviews from extra-EU and Italian travelers did not change significantly. Importantly, these results are not driven by an increase in international tourist flows toward Italy. Thus, the policy provides an abrupt and arguably exogenous source of variation in Tripadvisor use by EU travelers, whose reviews constitute about 30% of the total volume in restaurants located in the most tourist areas of Rome.

I combine the temporal variation induced by the policy with the spatial variation in tourist demand. In particular, I take advantage of the granularity of my data and construct two measures of restaurants' exposure to tourist clientele that account for the intensity to which each restaurant is

[3]By contrast, Jin and Leslie (2003) and Klein et al. (2016) study very market-specific quality disclosure programs and feedback systems in online platforms.

[4]Specifically, the Regulation (EU) 2015/2120 of the European Parliament and of the Council of 25 November 2015 "aims to establish common rules to safeguard equal and non-discriminatory treatment of traffic in the provision of internet access services and related end-users' rights." The subsequent Commission Implementing Regulation (EU) 2016/2286 of 15 December 2016 states that "roaming providers should not levy any surcharge additional to the domestic retail price on roaming customers in any Member State [...]." `https://europa.eu/youreurope/citizens/consumers/internet-telecoms/mobile-roaming-costs`.

potentially affected by the lower information costs induced by the policy. The first measure reflects the probability of finding a restaurant given its location with respect to the closest tourist attraction and the road network around it. In practice, I use information from Google Maps API to define the partial road network that leads to the closest restaurants around each attraction and compute the probability of finding each of these restaurants while walking away from the attraction. I show that higher probability values are positively correlated with the share of reviews from foreigners, while they are negatively correlated with the restaurants' average rating.[5] The second measure exploits the variation in the number of attractions across ZIP codes as a proxy for potential exposure of all restaurants in a ZIP-code to tourist clientele and, therefore, to the change in information costs.

The identification strategy relies on a Difference-in-Differences specification, which compares the changes in the outcomes over a 5-year window across restaurants that are more and less exposed to tourist clientele. Particularly, in the baseline specification, I use the median value of the previously-described probability measure to create a binary treatment indicator for high *vs.* low tourist exposure. I conduct the analysis on the sample of restaurants with at least one review in the pre-policy period, as well as on three equally sized sub-samples defined using the tertiles of the restaurants' rating at the time of the policy — namely, low, medium and high rating —, and run the regressions on each group separately. This allows me to study heterogeneous effects across establishments. The identification assumption requires that, within each group, changes in the outcomes across restaurants with high and low exposure to tourists would have been the same in the absence of the policy.

I then test the theoretical predictions of the model. From the first one, I expect consumers to reallocate their demand toward restaurants with *ex-ante* higher Tripadvisor ratings, which therefore should expand their sales and employment (output). I obtain several empirical results supporting this prediction. After the policy, annual revenues increased by almost 7% in high-rating restaurants, by approximately 3% in mid-rating ones, and remained the same in the low-rating category, with this positive gradient being statistically significant. As a result, revenues increased by almost 5% overall, pointing out an average growth in sales by approximately 32.5 Thousand Euros a year.[6] Total monthly employment also expanded by approximately 4% in more tourist restaurants, compared to less tourist ones. With an increase by 10%, the mid-rating category is mostly responsible for the overall growth in firm size. Moreover, the impact is negative for low-rating restaurants, while positive for high-rating ones, yet, in both cases coefficients are smaller and insignificant, suggesting that high-rating establishments were already producing at full capacity. Finally, consistently with consumer learning from online maps, I also document demand reallocation over space, with restaurants located in "hidden" alleys growing more than those in front of tourist attractions.

The second set of theoretical predictions concerns the supply side. I expect to observe higher exit rates for lower-rated establishments, as well as investment into higher quality inputs (such as hiring

[5] As theory predicts (e.g., Chan and Leland 1982) firms that more frequently engage with uninformed consumers tend to under-provide quality.

[6] These results are quantitatively similar to those from papers studying the impact of review platforms on firms' sales (e.g., Chevalier and Mayzlin 2006; Anderson and Magruder 2012; Luca 2016; Hollenbeck 2018, among others).

more qualified workers) for all operating restaurants, and particularly for those with *ex-ante* lower ratings. I start with the analysis of firm exit. I find that, the probability for low-rating establishments to exit the market doubled after the policy, compared to the baseline period. By contrast, the policy did not significantly impact the probability that mid- and high-rating restaurants left the industry.[7] Moreover, by aggregating observations at the ZIP-code level to study exit and entry jointly, I find that the share of low-rating firms operating in the most touristy neighborhoods decreased by 2.5 percentage points (p.p.) after the policy, compared to non-touristy ZIP codes. These results suggest that lower search costs through online reviews — even when experienced by only a fraction of consumers — can alleviate adverse selection and make the industry more quality-oriented.

Then, I analyze the behavior of operating firms. I consider hiring decisions as a proxy for the restaurants' effort to upgrade the service quality through the recruitment of more experienced staff. I find that the probability of hiring a worker with previous experience in the restaurant industry increased overall by 10% with respect to the pre-policy mean. This effect is driven by low- and mid-rating establishments, where such probability went up by 9% and 16%, respectively. By contrast, the coefficient for high-rating restaurants is close to 0 and not statically significant. Additional evidence suggests that low-rating establishments accumulated human capital at the expenses of high-rating ones. As a result, daily salaries paid by low-rating restaurants increased by more than €1 (2% of their pre-policy mean), while they decreased by a similar amount in high-rating restaurants. Eventually, these opposite recruiting strategies had heterogeneous effects on the online reputation of these establishments, as measured by the average 5-month Tripadvisor rating. I find that restaurants in the low- and mid-rating groups received better reviews after the policy, as their 5-month Tripadvisor rating improved by 0.09 points (2.5%) and 0.08 points (1.9%), respectively.[8] By contrast, reputation remained unchanged in restaurants that were already at the high-end of the rating distribution. Importantly, all estimates are unaffected by the exclusion of restaurants that exited the market after the policy. Overall, these findings point out the role of review platforms in alleviating the moral hazard problem for experience goods.

Several placebo exercises validate the identifying assumption. For example, event-study estimates confirm the absence of diverging trends in the outcomes before the roaming regulation became effective. Moreover, a series of policy-permutation tests conducted in the pre-policy period provides further evidence on the exogeneity of the policy date with respect to other potential factors (such as seasonality) that might explain the observed results. Finally, I show that the main estimates are robust to the use of alternative measurements, samples and clustering units.

Review platforms have substantial industry-wide consequences. Back-of-the-envelope calculations suggest that abating the costs for all consumers to access online reviews leads to an overall increase in restaurant revenues, employment and exit rate by 1.6%, 1.5% and 0.5 p.p., respectively. The first two figures correspond to about 12% and 5% of the overall growth in revenue and employment experienced

[7]This result is in line with the empirical study by Hui et al. (2018) on eBay's reputation mechanism.
[8]These results are similar to those by Ananthakrishnan et al. (2023) and Proserpio and Zervas (2017).

by Italian restaurants between 2016 and 2019, respectively. While the last figure corresponds to almost 3% of the exit rate faced by the industry during the first year of the Covid-19 pandemic. Altogether, these results indicate that facilitating access to review platforms has consequential effects on the labor market, and on the performance and composition of firms operating in industries generally affected by asymmetric information.

This paper contributes to several strands of literature. First, recent studies find that online ratings are a significant driver of sales (Chevalier and Mayzlin 2006; Resnick et al. 2006; Cabral and Hortacsu 2010; Anderson and Magruder 2012; Luca 2016; Lewis and Zervas 2019) and they increase consumer welfare (Reimers and Waldfogel, 2021), especially for tourists (Fang, 2022). Others examine the interaction between reviews and firms' advertising decisions (Chen and Xie 2005, 2008; Hollenbeck et al. 2019), and between firms' use of management responses and their online reputation (Proserpio and Zervas 2017; Chevalier et al. 2018; Wang and Chaudhry 2018). A common conclusion is that review platforms have not only changed how consumers make decisions, but also how firms behave in the marketplace. However, empirical evidence on whether sellers react to reviews by boosting quality is scarce, with the exceptions of Ananthakrishnan et al. (2023), who exclusively rely on online ratings as a proxy for hotel quality, and Hunter (2022), who finds that platforms' rounding of ratings incentivizes firms near the rating thresholds to take strategic actions. In contrast, I am able to assess the impact of online reviews on quality upgrading, firms' employment decisions and industry composition, thanks to the richness of my data and the exogenous abolition of roaming fees.

Second, I contribute to the literature on information and product quality. Studies in this area focus on market-specific quality disclosure programs like certification badges, labels and scores (Jin and Leslie 2009; Elfenbein et al. 2015; Vatter 2021; Barahona et al. 2023; Hui et al. 2023), as well as the introduction of feedback systems in online marketplaces (Klein et al. 2016; Hui et al. 2018; Dai and Luca 2020). This research shows that such programs affect consumer choices, firm financial performance and market composition, as well as incentives to reformulate products and upgrade quality. For example, Jin and Leslie (2003) find that restaurants hygiene grade cards guide consumers decisions and incentivize restaurants to be clean. My paper innovates and expands upon this previous work in a twofold way. First, I show that widely-used informal sources of information on the consumer side — i.e., user-generated content available on review platforms — can affect an entire industry by providing sufficient incentives for firms to invest in quality upgrading. I do so by relying on an exogenous policy that, differently from the programs described above, had nothing to do with the industry under analysis. Second, I leverage a rich dataset to uncover how firms change their organizational structure when their output objectives change. While the above papers are silent about this aspect, others have looked at the relationship between firm organization and performance (e.g., Hansman et al. 2020). I focus on an overlooked channel: how firms engage in strategic hiring decisions over workers' curricula to reformulate their offering and upgrade service quality.

Finally, this paper also relates to the literature looking at information frictions as one source of demand constraints impeding firm growth (Jensen 2007; Allen 2014; Anderson et al. 2018; Jensen and

Miller 2018; Hjort et al. 2020; Bai 2021). The general lesson from this work is that more information in the market enhances growth through a reallocation of market share toward the most productive firms. While this literature has primarily focused on product markets, my paper shows that similar results hold in the service industry.

The paper is organized as follows. Section 2 presents the theory and Section 3 discusses the study setting and data; Section 4 describes the empirical strategy and Section 5 reports the main results; Sections 6 and 7 show placebo exercises and robustness checks; finally, Section 8 discusses the industry-wide consequences and Section 9 draws the conclusions.

2 Theoretical framework

2.1 A model of consumer search and quality upgrading

Restaurant meals, like other *experience goods*, can only be assessed through consumption (Nelson, 1970). When consumers cannot directly verify product quality prior to purchase, the market tends to under-provide it (Riordan, 1986). Review platforms offer readily available details on restaurant quality and other characteristics (such as price, location, and cuisine), empowering informed choices. As such, online search has made restaurant meals more akin to *search goods*.[9] Because this paper focuses on the reduction in costs to access information online, following Goldmanis et al. (2010), I model restaurant meals as a search good.

For conciseness, I describe the main concepts and predictions of the model in the text and provide the details in Online Appendix 1. The model extends the work of Goldmanis et al. (2010) to account for limited information about quality rather than prices. This allows to study how reductions in consumer search costs affect the equilibrium quality in the market through their effect on both (1) the behavior and (2) the type of producers. To do so, I borrow elements from two distinct theoretical literatures. The first is the set of models on sequential consumer search and endogenous producer choices (e.g., Anderson and Renault 1999; Bar-Isaac et al. 2012). The second is the set of industry equilibrium models with heterogeneous producers and endogenous selection into production (e.g., Hopenhayn 1992; Melitz 2003; Syverson 2004). The combination of these two types of frameworks allows to investigate the effects of search costs on quality provision through moral hazard and adverse selection, separately. This is a novel characteristic with respect to existing models on consumer search and product quality.

The model features consumers facing heterogeneous search costs s to learn about the quality q of a vertically differentiated product. They buy one unit of good and gain higher utility from better qualities. Consumers have full information of the price of the good, which is exogenous and independent of quality. Yet, they must visit stores one-by-one to know about its quality and, after

[9] According to Klein (1998), it is the relatively higher cost of search with respect to direct purchase that makes a good an experience good. Thus, when consumers can obtain important product information via new interactive media at decreasing costs prior to the purchase, the product can be considered a search good.

every visit, they compare the expected benefit and cost of continued search. Following McCall (1970), a consumer stops searching when the expected utility gain from another search is equal to the search cost s. Firms have heterogeneous abilities λ, which affect the efficiency (costs) of their production process. First, firms must pay a sunk cost to learn their abilities and decide whether to stay in the market or leave. Next, those that stay decide the quality level of their good and produce. Production requires a fixed cost of operation $C(q, \lambda)$, which depends positively on the chosen quality level q and negatively on the exogenous ability parameter of the firm λ. Hence, firms face a trade-off between revenues and costs, as both schedules are increasing in quality. In equilibrium, firms set the quality that maximizes profits given consumers' optimal search behavior (demand function) as well as their own and their rivals production costs. Equilibrium qualities, quantities and profits are increasing in the ability parameter λ. It follows that only firms with sufficiently high ability — i.e., above a cut-off level $\underline{\lambda}$ — will be able to earn positive profits and stay in the market.

2.2 Comparative statics

All equilibrium functions and values depend on the search costs that consumers face. My goal is to determine how a decrease in these costs will affect the equilibrium quality $q(\cdot)$ and cost function $C(\cdot)$ of operating firms, their demand as well as the operating cut-off level of ability $\underline{\lambda}$. For this purpose, I impose some further assumptions that allow to keep the algebra tractable and align the model with the empirical exercise explained in Section 3.

Assumption 1: *Search costs are uniformly distributed on* $[0, a]$ *for* $a > 0$*.*

This assumption allows to study changes in search costs that are heterogeneous across consumers. In particular, I will focus on a cost reduction that only affect consumers with *ex-ante* the highest costs (a).[10]

Assumption 2: *The firms' cost function takes the form*

$$C(q, \lambda) = \frac{q}{1-q}\,\frac{1}{\lambda},$$

which satisfies the requirements described in Online Appendix 1.3 for $q \in (0, 1)$ *and* $\lambda > 0$*.*

This assumption has two direct implications. First, the quality level chosen by any firm is bounded between $(0, 1)$. This is consistent with the way individuals value and rate the quality of a meal on review platforms, which exhibit a finite scale (e.g., 1 to 5).[11] The second implication is that costs are increasingly steeper in quality, and tend to infinity as $q \to 1$, for a given λ. This assumption reflects the idea that superb quality requires the owner to completely restructure the business model, which is unfeasible in the short-term.

The comparative statics results follow (derivations are in Online Appendix 1):

[10] In the empirical framework these consumers are identified with foreign tourists, which bear the highest costs to browse the Internet in the presence of roaming tariffs.

[11] In this respect, the upper bound equal to 1 becomes the natural limit of quality via the reputation mechanism, so that firms have no incentive to deliver a level of quality beyond that value.

Figure 1: The effect of lower search costs on quality, ability threshold and production costs

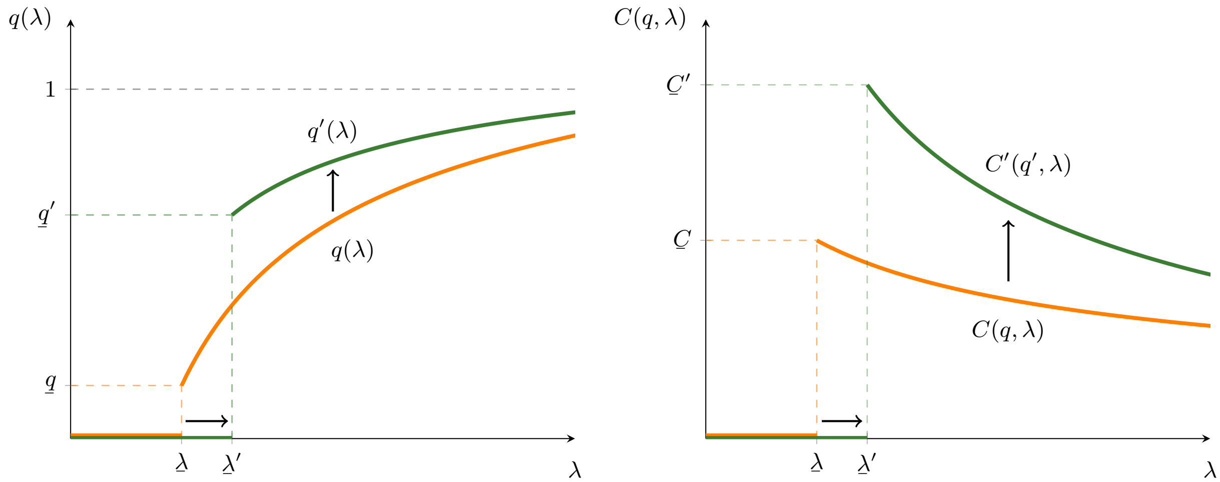

Notes: The orange (green) lines report the equilibrium schedules when search costs are higher (lower), that is $a : a \to a' < a$.

Proposition 1: When search costs decrease, the quality $q(\cdot)$ produced by a firm with ability λ increases $\forall \lambda \geq \underline{\lambda}$, and more so for firms with lower ability. That is, $q'_a(\cdot) < 0$ and $q''_{a\lambda}(\cdot) > 0$.

Proposition 2: When search costs decrease, the production costs $C(\cdot)$ of a firm with ability λ increase $\forall \lambda \geq \underline{\lambda}$, and more so for firms with lower ability. That is, $C'_a(\cdot) < 0$ and $C''_{a\lambda}(\cdot) > 0$.

Proposition 3: When search costs decrease, the cut-off ability value $\underline{\lambda}(\cdot)$ increases. That is, $\underline{\lambda}'_a(\cdot) < 0$.

Corollary 1: A decrease in search costs causes the demand $x(\lambda; a)$ faced by all firms with sufficiently high ability to increase.

Figure 1 graphically describes the consequences of a decrease in the parameter $a : a \to a' < a$ on the ability threshold and equilibrium qualities (left panel) as well as on the cost schedule (right panel). The first consequence of a decrease in consumer search costs is that some firms with the lowest abilities exit the industry, as $\underline{\lambda}$ shifts to the right. Operating firms ($\lambda > \underline{\lambda}'$) upgrade their quality when a decreases, and more so those that were initially producing lower-quality products (i.e., firms with lower abilities). As a result, production costs increase, indicating that the process of quality upgrading is overall costly. Particularly, firms that were initially producing lower qualities exhibit the largest cost increase. Hence, two simultaneous mechanisms make the industry more quality oriented when search costs fall. First, the upward shift in the equilibrium quality schedule for all operating firms is a consequence of a reduction in moral hazard: lower search costs make consumers more sensitive to the quality of the products, hence firms' incentives to upgrade it increase. Second, the rightward shift in the cut-off ability level is a consequence of a reduction in adverse selection: when search costs fall, demand to firms with the lowest abilities (and qualities) decrease, so do their profits, forcing them out of business.

The results of Propositions 1, 2 and 3 together with Corollary 1 constitute the main theoretical predictions of the paper. Declines in search costs in the restaurant industry driven by the advent and diffusion of online review platforms have heterogeneous effects across businesses. Low-type sellers are hurt, sometimes to the point of being forced to exit. Higher types, however, gain from the shift as their demand grows. Incentives to upgrade quality arise, resulting in higher quality levels, especially for surviving low-type firms.

2.3 Model limitations and extensions

To find a closed-form solution to the model, I impose prices to be exogenous and I abstract from horizontal product differentiation and preference heterogeneity. These assumptions are partly justified by the fact that prices and other product characteristics like the type of cuisine are more observable than qualities, especially in the restaurant industry.[12] Obviously, reviews could still affect prices by tightening competition, and they can help consumers find a better match value for their taste. In this respect, Fishman and Levy (2015) incorporate consumer taste heterogeneity in their search model, and study price and quality competition jointly. They show that lower search costs can have both positive and negative impacts on firms' incentives to invest in quality, because of their differential effects on prices.[13] However, their paper remains silent about industry composition effects, which represent an important feature of my framework.

Furthermore, I impose a sequential search rule for consumers, and a specific production cost function for firms. The results are qualitatively unchanged if these assumptions are relaxed. For instance, one could assume fixed sample size search behavior, where consumers sample a fixed number of stores and choose to buy the lowest priced/best quality alternative (e.g., Chan and Leland 1982; Cooper and Ross 1984). In this set of models, the probability for one uninformed consumer to face unfavorable conditions is proportional to the share of customers with imperfect knowledge about price/quality or to the cost of acquiring that information. Hence, even in this context, lower information costs would increase product quality. Yet, these frameworks do not directly allow to study composition effects.[14] Finally, my results are qualitatively unchanged when alternative convex cost functions are used (e.g., $C = 1/q\,\lambda$ or $C = q^2/\lambda$), provided that they satisfy the requirements described in Online Appendix 1.3 for quality $q \in (0, 1)$.

2.4 Hypotheses

To test the model predictions, it is necessary to identify the empirical counterparts of the theoretical quantities. Objective measures of quality are difficult to obtain, especially for the restaurant industry, where the quality of a meal reflects multiple dimensions (e.g., service and food) whose evaluation is

[12] This is in line with the idea that information on the price of a meal or the type of cuisine might be gathered before the purchase, for instance, by reading the restaurants' menu on the window.

[13] While they increase the market shares of high-quality firms, lower search costs also reduce their prices and profits more than those of low-quality firms, hence the effect on incentives to invest in quality is ambiguous.

[14] Another option would be to consider asymmetric information models with different types of producers, such as Akerlof (1970) and Leland (1979). However, these studies assume that information is unavailable at any price, thus leaving little scope for policies that reduce the cost of acquiring it.

to a large extent subjective. Following recent empirical literature (e.g., Ananthakrishnan et al. 2023; Chevalier et al. 2018; Proserpio and Zervas 2017), I use the Tripadvisor rating as a proxy for the dimension of quality that is mainly subjective and can be referred to as *reputation*. Importantly, this is the dimension that is revealed to consumers once they pay the search cost and visit the Tripadvisor profile of the restaurant. Hence, owners and managers seek to maximize it.

Moreover, in the absence of objective measures of output quality,[15] I rely on restaurant hiring decisions and production costs to proxy for it. In fact, high-quality output requires high-quality inputs (Kugler and Verhoogen, 2012) and, specifically, skilled workers (e.g., Verhoogen 2008; Brambilla et al. 2012). Similarly to Shin et al. (2023), I consider employee turnover, the curriculum of newly-hired employees and their wages to measure investment into service quality through the labor market. Finally, I use annual revenues and the total number of employees to proxy for output quantity and firm size, while information on the restaurant presence in the market comes from the official date of opening/closure of the business.

One potential caveat is that the empirical counterpart of the ability parameter λ remains unobserved. Nevertheless, the model suggests that there exists a one-to-one mapping between ability and quality. Hence, I rely on the Tripadvisor rating of the restaurant at baseline (i.e., before the reduction in internet tariffs) to proxy for the underlying ability parameter. Therefore, the theoretical predictions can be translated in the following testable empirical hypotheses. When search costs for consumers fall (through access to review platforms):

1. The demand faced by firms with *ex-ante* sufficiently high Tripadvisor rating increases: their revenues and number of employees grow.

2. The overall quality level in the industry improves:

   (a) Some of the firms with *ex-ante* the lowest ratings will exit the market: a reduction in adverse selection;

   (b) Surviving firms will invest in quality upgrading (e.g., hiring workers with better curricula) eventually improving their online ratings. These effects will be larger for firms with *ex-ante* lower ratings: a reduction in moral hazard.

3 Study setting and data

3.1 The EU roaming regulation

Following recent empirical literature on consumer search (e.g., Ershov 2020), I take advantage of online platforms to characterize reductions in search costs. In particular, I rely on an exogenous reduction in the costs of mobile internet caused by the abolition of roaming tariffs in the European Union, which promoted the use of review platforms like Tripadvisor.

[15] Previous work (e.g., Jin and Leslie 2003; Farronato and Zervas 2022) relied on health inspection scores to measure the hygiene dimension of quality. Unfortunately, these data are not available for Italy.

International mobile roaming regulations apply when customers use their mobile phones while occasionally travelling outside the country where they normally reside (specifically, outside the geographical coverage area of the home operator's network). This paper exploits the approval by the European Parliament of a new policy on the EU roaming regulation, which led to the abolition of all charges for temporary roaming within the European Economic Area (EEA) as of June 15, 2017.[16] In practice, if before that day all EU residents traveling within the EEA were charged at least € 0.05 per MB of data (on top of the home network rate), after the policy the same home network rate is applied with no additional charges.[17]

The EU roaming regulation consists of a series of policy packages that started in 2007 and regulate wholesale and retail international roaming tariffs. The policy was initially motivated by the large gap between the roaming prices charged to consumers and the actual cost of providing the wholesale service. Therefore, the underlying objectives of the reform were the intensification of the competition among providers and the promotion of market integration in the EU. The last decisive step took place on June 15, 2017, when wholesale and retail price caps for data were set to 0.[18]

For the purpose of this paper, the reform induced an exogenous shock to the costs of accessing online information for EU travelers while abroad. In particular, information contained in review platforms such as Tripadvisor became available to all EU travelers at virtually zero cost.[19] Hence, the search costs for certain tourists looking for restaurants while visiting a city drastically decreased compared to the pre-policy period.[20]

3.2 Data sources

To study the consequences of lower information costs on the restaurant industry, I focus on the whole Province of Rome. Looking at a large geographical area allows me to exploit spatial variation in the intensity of exposure to tourist clientele, an attractive feature for empirical identification. Specifically, I assemble a novel dataset combining three data sources.

The first source is Tripadvisor, the most popular travel guidance platform in Italy and Europe.[21] Listing an establishment on the platform is free and can be done either by the clients or by the owner/manager. I collect information on listed restaurants (name, address, price category, type of cuisine, etc.) as well as, for each restaurant, the entire historical record of reviews (date, rating, device, text, country and language of reviewer, etc.) from 2007 to 2019, for a total of approximately

[16] https://europa.eu/youreurope/citizens/consumers/internet-telecoms/mobile-roaming-costs

[17] Note that the policy also affected the prices of SMSs/phone calls, but only toward the country of origin. Calling/texting local restaurants while abroad remained equally expensive, hence this type of communication is unlikely to play any role in the observed results.

[18] Grzybowski and Muñoz (2020) show that the European Commission has succeeded to avoid unintended increases in domestic tariffs and induced operators to absorb the negative effects of the reform.

[19] For example, loading the Tripadvisor pages of 10 restaurants could cost at least 25 cents ($\sim$ 1MB) before the policy, while nothing after it.

[20] In this respect, Quinn et al. (2021) show that, after the policy, daily mobile data consumption (sum of uploads and downloads) for EU travelers while abroad grew by at least 54MB.

[21] SimilarWeb, unique users de-duplicated monthly, accessed in March 2022.

3 million reviews. Since the format of the data is unstructured, I combine them together to create a panel at the restaurant-month level. Importantly, I used the historical record of reviews to retrieve the average rating of the restaurants in any month between 2007-2019. The Tripadvisor sample contains information on 14,146 establishments with at least one review as of December 2019. Of them, 11,595 had at least one review in May 2017, i.e. the month before the roaming policy was effective.[22] Finally, from Tripadvisor I also gather information on location and attributes of the top-100 tourist attractions in the Province, according to their total volume of reviews.[23]

The second data source is provided by administrative Social Security records collected and maintained under restricted-use access at the Italian National Social Security Institute (INPS). For each establishment in the Province of Rome, the records contain information for the last 15 years on location (ZIP code), date of opening and closure, legal status of the firm, monthly number of employees, type of contracts and qualification of the workers, their wage bill and demographics, as well as their full employment history. According to this dataset, 10,391 restaurants operated in the Province and had an active profile at the Social Security Institute in at least one month between 2015 and 2019.[24]

The third source contains proprietary annual information on income statements and balance sheets originally collected and maintained at the Italian Business Registry (Chamber of Commerce) and accessed through the Cerved database.[25] This dataset provides information on revenues, costs, profits and other financial indicators, and it only covers firms with an LLC proprietorship status.[26] In particular, almost 5,000 restaurants in the Province were obliged to report their financial information to the Registry at any point in time between 2015 and 2018, which is the last available year.

3.3 Dataset construction

I matched the Tripadvisor sample with Social Security and financial records. Combining crowd-sourced data with administrative archives is challenging because of the different nature and confidentiality protocols of the two sources. The two main obstacles were (1) the anonymity of the administrative records and (2) the lack of official business identifiers in the Tripadvisor data. In fact,

[22]It is worth mentioning that having a Tripadvisor profile with a positive number of reviews at a particular point in time does not necessarily imply that a restaurant is an active business at that time. For instance, it might be the case that the restaurant has closed, but the Tripadvisor profile still exists.

[23]I consider as tourist attractions those activities that on Tripadvisor belong to the categories "Sites of interests" and "Monuments". With almost 128,000 as of 2019, the Colosseum is the most-reviewed attraction, while the National Roman Museum ranks 100th on the list, with almost 600 reviews.

[24]I use information on the primary activity of the firm (ATECO code) to identify restaurants. Particularly, I restrict the attention to the following ATECO codes: 56.10.11 (dine-in restaurants), 56.10.12 (agriturismi), 56.10.20 (take-away restaurants), 56.10.30 (bakeries).

[25]Particularly, I access the version of Cerved data that is available at the Social Security Institute in Rome, where the last available year is 2018.

[26]In the restaurant industry in the Province of Rome, LLC companies represent about 57% of the total. These businesses are owned by shareholders, who have limited personal liability for business related debts and are required by the law to report financial statement information at the Chamber of Commerce on an annual basis. By contrast, firms with no financial data are usually unlimited liability partnerships and sole proprietorship businesses, which are generally smaller and more likely to be family-owned restaurants.

for confidentiality purposes, Social Security records do not contain the names, addresses and unique business identifiers of the firms, i.e., the VAT codes.[27] On the other hand, Tripadvisor records do have the names and addresses of the restaurants, but not their unique business identifiers. Hence, the information was incomplete on both sides.

To overcome the limitation, I obtained additional data from the Italian Business Registry containing names, addresses and VAT codes of the restaurants. I then used the name and address to assign a VAT code to as many restaurants as possible in the Tripadvisor sample. For restaurants that could not be matched automatically using the name and address, I manually collected the VAT codes from their websites or from the pictures of receipts posted by the clients on Tripadvisor/Google. This procedure resulted in about 6,000 Tripadvisor restaurants with an associated VAT code, which was subsequently used to match restaurants with the Social Security records.[28]

The final matched Tripadvisor-Social Security sample comprises 5,472 firms that operated in the industry at any point in time between 2015 and 2019.[29] This sample represents almost 53% of the total number of active businesses in that period and it is employed in the market-level analysis to assess how entry/exit dynamics shape the composition of the restaurant industry. Moreover, among the matched restaurants, 4,628 had a Tripadvisor profile with at least one review (and, therefore, a rating) in May 2017. I employ this sample in the firm-level analysis.

3.4 Summary statistics

Table 1 shows descriptive statistics for the 4,628 matched restaurants with available Tripadvisor rating at the time of the roaming policy. For the main outcome variables, the table reports a series of statistics referring to the Jan 2015 - Dec 2019 period, as well as, for the sake of comparison, their mean in the pre-policy period (Jan 2015 - May 2017).

The average restaurant is a small business, with less than 6 employees and an annual revenue just below 700 Thousand Euros. Its employees work, on average, almost 16 days per month, and their adjusted full-time-equivalent gross salary is €67 per day. During the period of interest, the average establishment hires a worker with previous experience in the restaurant industry at a frequency of 8% of the total operating months and, when it does, the new employee has worked about 14.5 months in the sector. The 5-month rolling average rating that the typical restaurant obtains on

[27] Access to both INPS and Cerved databases was granted under a specific program (VisitINPS). In compliance with the program requirements, most of the empirical analysis presented in the paper was carried out at the data center in Rome, and no data has left the center except for the output tables and figures reported in the paper. Official information on the program is available here `https://www.inps.it/dati-ricerche-e-bilanci/attivita-di-ricerca/programma-visitinps-scholars`.

[28] Tripadvisor data had to be imported in the Social Security archives. To avoid that individual firms could be identified from the data, I selected only the most important Tripadvisor variables and grouped their values in categories. This simplified dataset was imported. The main data import took place in early 2019. For this reason, imported data on reviews, ratings and replies from Tripadvisor cover until December 2018.

[29] To minimize the risk of measurement error due to misreporting in the Social Security data, before conducting the analysis I trimmed observations with a number of employees above the 98th percentile. The final sample does not include these observations.

Table 1: Summary of firm-level outcomes

		Period: Jan 2015 - Dec 2019						Pre-policy
	Firms	Obs	Mean	SD	Min	Median	Max	Mean
N. of monthly employees	4628	219835	5.69	5.50	0.0	4.0	29.0	5.55
Annual revenues (Thousand, €)	2043	6677	692.18	1065.00	5.0	394.0	8752.0	646.60
Monthly working days	4628	219835	92.13	102.04	0.0	58.2	1922.9	90.96
Working days per worker	4517	197194	15.66	5.84	0.1	15.6	185.9	15.86
1 if firm exits (×100)	4628	219835	0.41	6.36	0.0	0.0	100.0	0.33
1 if firm hires worker w/ previous experience in restaurants	4628	219835	0.08	0.28	0.0	0.0	1.0	0.08
1 if firm hires worker w/o previous experience in restaurants	4628	219835	0.06	0.23	0.0	0.0	1.0	0.06
Months of experience in restaurants of newly-hired employees	3550	30133	14.46	22.96	0.0	3.8	157.0	13.01
Months of experience in restaurants of quitting/fired employees	3584	30911	27.21	29.83	0.0	17.0	152.0	25.12
Average daily salary (€)	4558	200402	66.60	19.12	24.6	61.1	156.8	64.88
Average 5-month Tripadvisor rating	4373	147274	3.96	0.65	1.0	4.0	5.0	3.98
N. of 5-month replies to reviews	4377	146713	2.46	11.47	0.0	0.0	313.0	2.56
N. of monthly Tripadvisor reviews	4572	178425	5.70	12.51	0.0	3.0	1110.0	6.18

Each observation is a restaurant-month-year, with the exception of revenues, which are observed at the restaurant-year level up to 2018. Data on Tripadvisor reviews, rating and replies refer to the period between Jan 2015 and Dec 2018. Daily salary is adjusted for part-time workers so to reflect the full-time equivalent salary.

Tripadvisor is almost 4, and the number of 5-month total replies to online reviews from the profile manager is 2.5.

To address concerns on the potential bias in the analysis introduced by the matching procedure, Appendix Table A1 compares the main descriptive statistics of the matched sample with those of the entire Tripadvisor and Social Security datasets, separately. While matched restaurants tend to be slightly larger in size, as well as closer and more exposed to tourist attractions than the average restaurant in the Province, they also appear to have similar Tripadvisor ratings and price tier distributions, recruit equally-experienced workers, and pay comparable salaries. Overall, this evidence seems to discard the possibility that results could be systematically driven by sample selection.

3.5 The roaming policy and the use of Tripadvisor

As the costs of mobile internet falls, its use is expected to increase. Consumers with free internet access have the possibility to search and compare products online before purchasing. Tripadvisor data allows to study the reviewers' behavior across types of device and nationality. Although Tripadvisor contributions reflect the supply of reviews and not necessarily their consumption, in the absence of better data on the demand side, I employ them here as a proxy for overall usage of the platform.[30]

Figure 2: The roaming policy and the use of Tripadvisor

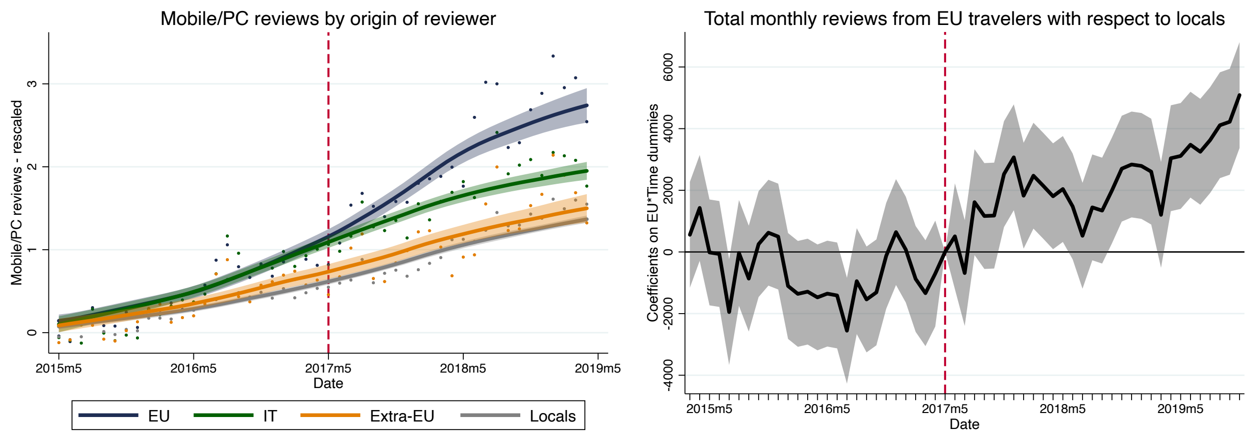

(a) (b)

Notes to Panel (a): Data on 14,146 restaurants with at least one review as of December 2019. Dots represent the monthly ratios, lines depict local polynomial fits with 95% confidence intervals. Values are re-scaled so that they are equal to 0 at the beginning of the period.
Notes to Panel (b): The graph reports estimated coefficients on the interactions between EU-dummy and time dummies, from a regression where each observation is a region of origin-month-year. The gray area reports 95% confidence intervals.

Panel (a) of Figure 2 shows the change in reviewers' behavior over time across device and origin of the reviewer, which is proxied by the language of the review.[31] The picture points out a shift from PC- to mobile-based contributions following the implementation of the roaming policy (the red dashed line). Importantly, this effect is remarkably visible only for EU tourists but not for Italians, extra-EU travelers and locals, who were not deliberately targeted by the new regulation. Regression analysis confirms the visual results. Appendix Table A2 compares the posting behavior of EU, extra-EU and Italian travelers with that of the locals. After the policy, users from EU countries became 1.4 times more likely to post reviews on their mobile devices as opposed to PC, while no significant change occurred for extra-EU and Italian users. The table also shows that the absolute number of monthly reviews posted on mobile devices by EU travelers increased by approximately 500 reviews per month after the policy, compared to the locals.

These patterns are not exclusively driven by PC-to-mobile substitution. Panel (b) of Figure 2 shows event-study estimates for the total (mobile+PC) number of reviews posted by EU users compared to the locals. While before the policy total contributions from both categories displayed similar trends, the volume of monthly reviews from EU travelers significantly and steadily increased after May 2017. Importantly, as Appendix Table A3 and Figure A1 show, these results are unlikely to be driven by a discontinuous increase in international tourist flows toward Italy around the time of

[30] Reviewers are likely to be a subset of the total users (visitors) of Tripadvisor, as posting reviews entails an additional effort that not all users are willing to bear. Hence, finding evidence on changes in reviewers' posting behavior likely implies that similar changes hold, more generally, for the broader set of users.

[31] For languages that imply multiple nationalities (e.g., English and Spanish), I use explicit information on the origin available on the Tripadvisor profile of the reviewer to identify non-EU users.

the policy. Overall, this evidence suggests that the new regulation provided an abrupt and arguably exogenous source of variation in the use of Tripadvisor services by EU travelers.

4 Empirical strategy

To estimate the demand and supply effects of increased access to information from online reviews, I combine the temporal variation induced by the policy with the spatial variation in restaurants' exposure to tourist clientele. Intuitively, and in line with Fang (2022), restaurants that more frequently cater to tourists are also more likely to be affected by the roaming regulation, as a larger share of their clientele experiences the decrease in mobile internet tariffs.[32]

4.1 Restaurants' exposure to tourists

I take advantage of the granularity of my data and construct two measures of exposure to tourist clientele that account for the intensity to which each restaurant is potentially affected by the policy. In both cases, I rely on the location of a restaurant with respect to tourist attractions to predict the composition of its clientele.

The first measure aims to capture the probability that a tourist finds a restaurant while walking away from an attraction site. To build this, I consider the top-100 attractions in the Province of Rome, according to their total volume of Tripadvisor reviews. For each attraction, I identify the closest Tripadvisor restaurants around it, and then use the Google Maps API to construct the shortest walking route from the attraction site to each of these restaurants. The procedure generates the partial road network around every attraction. I then assume that tourists follow a *random walk* while they move away from the attraction site, which allows me to assign equal conditional probabilities to every road at a same junction. Finally, I compute the joint probability to find the restaurant(s) located at any point along the network as a product of the conditional probabilities attached to all the consecutive roads leading to that point. The procedure, which is described in detail in Online Appendix 2.1, provides a continuous probability measure $P(i) \in [0, 1]$ that reflects the chances that restaurant i is visited by tourists while they move randomly toward the periphery of the road network, starting from the attraction site. Hence, by construction, this quantity only depends on (1) the location of the restaurant with respect to its closest tourist attraction, and (2) the shape and density of the road network around it.[33] Online Appendix 2.2 shows that such probability is positively associated with the average share of reviews from tourists, while it is negatively correlated

[32]In the 2.5 years preceding the regulation, Tripadvisor reviews from EU travelers accounted for about 30% of the total volume in restaurants located in the most touristy areas of the Province, while for less than 1.5% in the least touristy ones.

[33]The advantage of this approach is that it relies exclusively on location parameters. Alternatively, one could use other information contained on Tripadvisor, such as the origin of reviewers, to determine the level of exposure to tourists for each restaurant. However, such information is the result of past consumption patterns and reviewers' behavior that might interplay with the policy, influencing future consumers' decisions on a restaurant regardless of its actual level of exposure to tourists.

with the average Tripadvisor rating of restaurants.[34]

The probability measure described above varies across restaurants, which is an attractive feature to study firm-level response. Such a granular level of variation, however, does not allow to identify the aggregate (market-level) effects of the policy on the industry composition. For this purpose, I focus on the ZIP codes in the province of Rome, and construct an alternative measure of exposure to tourists at that level. Particularly, I focus on the number of top-100 tourist attractions in each ZIP code. This measure reflects the potential exposure of all restaurants in a ZIP code to tourist clientele and, therefore, to the change in internet tariffs induced by the policy. Among the 127 ZIP codes with at least one restaurant, about 25% have at least one tourist attraction, with the most touristy ZIP codes containing 25 sites.

4.2 Identification

The basic idea behind the identification strategy is to compare the evolution of firm-level and ZIP-level outcomes before and after the policy across firms/ZIP codes that are differentially exposed to tourists, and therefore to the change in the roaming tariffs.

4.2.1 Firm-level outcomes

The firm-level analysis employs the first measure of exposure to tourist clientele, i.e. the probability $P(i)$ previously defied. Online Appendix 2.3 shows that — because of the composition of their clientele — only restaurants with a sufficiently high probability are potentially affected by the policy, as both the pre-policy shares of EU reviews and their change across devices (from PC to mobile) after the policy are significantly higher for restaurants with probability values above the median. Thus, the baseline empirical specification of the paper relies on these facts to identify two equally sized groups of restaurants: the treatment group, composed of restaurants with a probability value above or equal to 0.17% (i.e., the median), and the control group with the remaining restaurants. Specifically, I consider observations within January 2015 and December 2019,[35] and estimate the following Difference-in-Differences model:

$$y_{i,t} = \beta \, Tourist_i \times Post_t + \alpha_i + \gamma_t + \phi \, \mathbf{x}_i \times Post_t + \varepsilon_{i,t} \tag{1}$$

where i is the restaurant, and t is time. Depending on the outcome, the analysis is conducted at the monthly or yearly level. $Tourist_i$ is a binary variable taking value 1 if the measure of tourist

[34] I investigate the robustness of my procedure to different data and assumptions. First, instead of focusing only on the shortest path, I also include all alternatives routes provided by the Google Maps API in the computation process. Second, rather than imposing equal conditional probabilities (random walk assumption), I assume tourists form educated guesses on which path to follow, based on importance (frequency) of each road. These procedures are explained in Online Appendix 2.1, and the resulting probability measures will be used in the analysis to conduct robustness checks.

[35] With the exception of financial and rating data, which are only available up to 2018.

exposure $P(i) \geq 0.17\%$.[36] $Post_t$ takes value 1 after the policy, that is for t after May 2017 when outcomes are monthly, while for t after 2016 when outcomes are annual.[37]

α_i and γ_t represent restaurant and time fixed effects, respectively. Their inclusion allows controlling for both time-invariant firm-level characteristics and aggregate trends (such as seasonality) that might affect the outcome $y_{i,t}$ while being simultaneously correlated with the main independent variable, $Tourist_i \times Post_t$. Nevertheless, there is still the possibility that some demand- and supply-side factors might influence the outcomes over time, while being simultaneously correlated with the main independent variable. To account for such potential endogeneity issue, vector $\mathbf{x_i}$ includes a series of time-invariant and predetermined restaurant-specific characteristics, which — once interacted with $Post_t$ — are allowed to have a differential impact on the outcomes over time. In particular, in all regressions I control for the distance (in km) of the restaurant to its closest attraction to account for factors, other than the presence of tourists, that correlate with proximity and could affect restaurant and consumer decisions.[38] It is worth mentioning that, once I control for the proximity to the attraction, the probability measure $P(i)$ mainly captures the "visibility" of a restaurant, i.e. whether the place is easy or difficult to be discovered by a tourist due to the shape and density of the road network. Other controls include restaurant price categories, a dummy indicating whether its cuisine is Italian or not, indicators for the concentration of restaurants within a 400-meter radius (reflecting the level of competition), indicators for the volume of reviews to the closest attraction (capturing the popularity of the whole area, potential congestion and rental costs), the classification of the main economic activity of the restaurant and its legal status (e.g., LLC vs sole proprietorship).[39] Finally, in all regressions, I also include ZIP-code linear time trends, as well as indicators for the distance to Rome city-center, to account for potentially diverging patterns across areas with different exposure to tourist demand and municipal regulations.[40] I cluster the standard errors at municipality level (86 clusters) to account for serial and spatial dependence in the errors.[41]

Theory described in Section 2 posits a differential impact of a reduction in search costs on consumers' demand and production choices across restaurants selling *ex-ante* different qualities. To empirically analyze the heterogeneous effects along the quality gradient, I estimate model (1) on different samples. First, I study the overall impact on all restaurants with available Tripadvisor rating at the time of the policy (N=4,628). Second, I use the tertiles of such average rating to split

[36]This is the median in the sample of 4,628 matched restaurants used in the firm-level analysis. For the sub-sample of restaurants with available revenue data (N=2,043), I use the respective median of the probability measure in the sub-sample, which is 0.35%.

[37]In the annual analysis of revenues, year 2017 is assumed to be fully treated even if the policy was effective in June. If anything, this should reduce the size of estimated coefficients, thus providing a lower-bound for the effect of the policy.

[38]Examples include rent costs that tend to be higher closer to attraction sites, or congestion (in fact, restaurants in touristy areas can easily be overcrowded, thus leading to longer waiting times, more noise, and worse service).

[39]Appendix Table A4 reports the list of independent and control variables along with their descriptive statistics. For confidentiality concerns, continuous control variables were categorized before being imported in the INPS server. Being these controls, their simplification should not crucially affect the results.

[40]In Italy, sanitary and hygienic regulations of the restaurants as well as their structural standards (such as capacity and equipment) are generally established by the municipal councils.

[41]I provide robustness of the estimates to a different level of clustering, using ZIP codes.

the sample in three sub-samples of equal size and estimate the model on each group, separately. Appendix Figure A2 shows the overall rating distribution, and highlights the three subgroups of interest: low-rating restaurants, with rating < 3.85; mid-rating ones, with rating $\in [3.85, 4.25)$; and high-rating ones, with rating ≥ 4.25.[42] Note that Tripadvisor does not display the average rating of a restaurant, but rather its rounded value.[43] Therefore, these three groups contain restaurants whose displayed ratings are approximately below, around, and above 4, respectively.

By estimating (1) via OLS, the coefficient of interest β reflects the change in the outcomes before/after the policy across restaurants more and less exposed to tourists. In order for β to have a causal interpretation, the identification assumption requires that trends across the two groups would have been the same in the absence of the policy. I conduct a number of placebo exercises to provide plausible evidence in support of this assumption. These include (i) event study estimates, where the dummy variable $Tourist$ is interacted with semester dummies, which allow to both study the dynamic effects of the policy, and check for the presence of differential trends in the outcomes in the pre-policy period; (ii) a series of permutation tests, where the effect of several placebo policy-dates between 2012-2016 is assessed; (iii) specific placebo policy-dates coinciding with the months of May in years 2013-2016 to explicitly test whether seasonality could explain the observed results.

Finally, there is still the possibility that online information from mapping apps[44] helps tourists navigate the streets around attractions, allowing them to discover less visible restaurants — e.g., those around the corner or in hidden alleys of the city center —, which would have not been visited otherwise (as suggested by Ghose et al. 2013 and Dall'orso et al. 2016). I take advantage of the granularity of my probability measure to study the potential reallocation of consumption over space, from highly visible establishments to more hidden restaurants that are nevertheless easy to reach (i.e., within walking distance) for tourists. To do so, I allow for (1) to take a more flexible form, where I use deciles/quantiles of $P(i)$ instead of the dummy variable, and interact them with $Post_t$. I display these results in a series a figures.

4.2.2 ZIP-level outcomes

To study the entry/exit dynamics and the resulting effects of the policy on the distribution of equilibrium qualities, I group establishments at the ZIP-code level and exploit the variation in the number of tourist attractions to proxy for exposure to tourist clientele (as described in 4.1). In this setting, a Diff-in-Diff approach would compare changes over time across ZIP codes with a higher and lower number of attractions. Particularly, I focus on the matched sample of restaurants irrespective of their presence on Tripadvisor at the time of the policy (N=5,472).[45] As before, I consider the

[42] In the sub-sample with available revenue data (N=2,043), the rating tertiles are 3.80 and 4.20.

[43] In both Tripadvisor and Yelp the average rating is rounded at the nearest half integer. So for example, a 3.73 average rating would be rounded to 3.5. Some studies like Luca (2016) and Farronato and Zervas (2022) take explicit advantage of this feature in their identification strategy.

[44] These include Tripadvisor, which has a "find-near-me" option, but also other popular apps such as Google Maps.

[45] Note that this allows me to study the effects of the policy not only on the exit but also on the entry/type of new restaurants.

January 2015-December 2019 period and estimate the following equation:

$$y_{z,t} = \beta \, Attractions_z \times Post_t + \alpha_z + \gamma_t + \phi \, \mathbf{x}_z \times Post_t + \varepsilon_{z,t} \tag{2}$$

where z is the ZIP code, and t is time, measured in months. $Post_t$ takes value 1 after May 2017. $Attractions_z$ is a time-invariant variable containing the number of attractions located in z. α_z and γ_t are ZIP-code and month fixed-effects, respectively. To account for potentially diverging trends in the outcomes across different ZIP codes, I also include ZIP-level linear time trends. Moreover, in vector $\mathbf{x}_z$, I include categorical variables reflecting the average distance of restaurants in the ZIP code to Rome city-center, which I interact with $Post_t$ to control for factors, other than the presence of tourists, that correlate with proximity to the main city and might affect consumption and production choices over time. I cluster the standard errors at ZIP-code level (127 clusters) to account for serial correlation in the errors.

By estimating (2) via OLS, the coefficient of interest β reflects the change in the outcomes before and after the policy, across ZIP codes that are more and less exposed to tourist clientele. In order for β to have a causal interpretation, the identifying assumption requires that, in the absence of the policy, the outcomes of different ZIP codes would have changed similarly. To check the plausibility of this assumption, I perform a variety of event-study and placebo estimates similar to those described in the firm-level analysis of Section 4.2.1.

5 Results

5.1 Restaurant revenues and size

Theory presented in Section 2 predicts that higher-quality firms increase their output as a consequence of lower consumer search costs. To test this hypothesis in the absence of data on quantity, I rely on restaurant revenues and total employment as a proxy for output and size. The analysis considers the outcomes of restaurants while they are in operation, including those that exited the market after the policy. As discussed in Section 7, all results are robust to the exclusion of restaurants that went out of business after the policy.

5.1.1 Main results

First, I consider the sample of restaurants with available annual financial information and estimate equation (1). Column (1) of Table 2 shows that, after the policy, sales in more touristy restaurants increased by almost 5% compared to less touristy ones. The estimated coefficient is robust to the inclusion of additional controls (column 2), such as the price category and the type of cuisine, which might be correlated with both revenues and the level of exposure to tourists. The most conservative estimates imply an annual average increase in restaurant revenues of approximately 32.5 Thousand Euros, considering that mean revenue in the pre-policy period is around 650 Thousand Euros.

Table 2: Impact on restaurant revenues

	Y=log(annual revenues); years 2015-2018				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	0.047***	0.053***	-0.002	0.033**	0.069***
	(0.012)	(0.012)	(0.026)	(0.015)	(0.024)
Restaurant & Year FE	✓	✓	✓	✓	✓
ZIP-code*Year	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	6677	6652	2305	2299	2048
Restaurants	2043	2034	696	697	641
Clusters	57	56	39	40	41
Adj. R-squared	0.846	0.847	0.869	0.849	0.782
Mean Y pre-policy	646.6	648.8	977.4	558.0	360.7
DDD *p-value*				0.962	0.004

Post=1 if date is after 2016. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in$ $[1, 3.80), [3.80, 4.20), [4.20, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-year. The sample includes observations between 2015 and 2018. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating.* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Columns (3-5) of the table analyze the heterogeneous effects of the policy across restaurants with *ex-ante* different ratings, i.e., low, mid and high rating. Consistently with the model prediction, coefficients suggest that the overall increase in revenues is mainly driven by high-rating restaurants, whose sales expanded by almost 7%. Revenues in the mid-rating category also improved but to a smaller extent, by approximately 3%. By contrast, the policy had no impact on sales in the low-rating category, for which the estimated coefficient is virtually zero and not statistically significant. The last raw of columns (4-5) report the p-values of triple-difference coefficients, indicating a significant and positive revenue gradient along the rating dimension.

Data on revenues might be subject to measurement error, for instance due to firms misreporting their sales in the attempt to pay lower taxes. Therefore, I complement the analysis of restaurant output using monthly employment records.[46] To some extent, changes over time in the number of employees — i.e., firm size — reflect the variation in the restaurant's ability to attract clientele and fill-up the tables. However, the relationship between output and firm size is not necessarily linear (Basu and Fernald, 1997), especially when firms face capacity constraints, which likely imply decreasing returns to labor. In the case of restaurants, such constraints arise because of the narrow time-windows to serve a meal (launches and dinners) and limited physical space.[47] Figure 3 shows event-study estimates from separate regressions on the three sub-samples corresponding to the different rating categories previously defined. Table 3 reports the regression output.

[46] In this respect, labor information is generally more difficult to cover up and misreport to the authorities compared to financial data.

[47] In practice, an additional waiter would not be much productive when all the tables are already filled-up and clients have to wait in line outside of the restaurant for the next available seat.

Figure 3: Event-study estimates for restaurant employment

Notes: The graph reports estimated coefficients on the interactions of Tourist*Semester dummies from three separate regressions where each observation is a restaurant-month-year. All controls and fixed-effects from the main analysis are included. The omitted semester is 2017s1. Tourist restaurants are those with a measure of exposure above the median. The sample includes observations between Jan 2015 and Dec 2019. Shaded areas depict 95% confidence intervals.

Table 3: Impact on restaurant employment

	Y=log(monthly employees); Jan 2015 - Dec 2019				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	0.043**	0.042**	-0.024	0.103***	0.041
	(0.017)	(0.020)	(0.019)	(0.031)	(0.038)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	219835	217622	72133	76920	68569
Restaurants	4628	4576	1490	1571	1515
Clusters	86	86	59	71	71
Adj. R-squared	0.779	0.778	0.759	0.793	0.769
Mean Y pre-policy	5.5	5.6	6.9	5.7	4.0
DDD *p-value*				0.089	0.065

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating.* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

In line with findings on restaurant revenues, after the policy, total monthly employment expanded by approximately 4% in more touristy restaurants. The estimated coefficient does not change when additional controls are included in the regression (columns (1-2) of Table 3). Figure 3 reveals that the mid-rating category is mostly responsible for the overall increase. Here, on average, monthly employment grew by 10% after the policy, implying an average increase in the restaurant size by more than 0.5 workers, when compared to the pre-policy mean. The event-study estimates also suggest that such labor expansion did not take place immediately, but it rather happened around 6-to-12 months after the policy. This can be attributed to a delay in consumers' learning process and the restaurants' slow response to adjust their factors of production to meet increased demand.

Differently from revenues, total employment in the high-rating category did not change significantly. The estimated coefficient is around 4%, yet it is not statistically different from zero at any conventional confidence level. This results indicates the existence of diminishing returns to labor in the industry. Specifically, highly-rated restaurants catering to Italian tourists and locals were likely to operate at maximum capacity even prior to the roaming policy. The additional surge in demand from EU travelers leads to increased revenue but does not necessitate hiring more workers. Finally, and consistently with the revenue analysis, low-rating restaurants do not exhibit any significant change in their size. If anything, the sign of the estimated coefficient is negative (-2.5%) but not statistically different from zero.

5.1.2 Discussion

The above findings suggest that, because of the cheaper access to online word-of-mouth, restaurants with *ex-ante* a better reputation on Tripadvisor — namely a rating around or above 4 — attracted new clients and grew in size more than those with worse ratings (below 4). These findings are consistent with those from the previous literature (e.g., Chevalier and Mayzlin 2006; Anderson and Magruder 2012; Luca 2016; Lewis and Zervas 2019),[48] and add to this existing work providing new empirical evidence on the effects of online reviews on firm employment decisions.

The fact that gains at the top of the rating distribution are not compensated by losses at the bottom could be explained by several reasons. One is market expansion (e.g., columns (1-2) of Tables 2 and 3), which could occur, for example, if some consumers start substituting food from supermarkets/hotels with meals at the restaurants. However, such dynamic is unlikely to be the exclusive reason underlying the overall positive effects of the policy. Another possibility is demand substitution from restaurants with no Tripadvisor account (which are out of the sample and therefore not observed) to those with an active profile (in-sample). Nevertheless, even this type of substitution should not play a major role, as the great majority of the restaurants in the Province was on Tripadvisor around the time of the policy.

Alternative explanations bring into play supply-side dynamics. For example, firm exit (which is discussed in Section 5.2) might lead to a reduction in the total number of players in the market, leaving more clients — and therefore more revenues — to the surviving restaurants even if aggregate

[48]For example, using data on several online platforms such as Yelp, Tripadvisor, Expedia and Hotels.com, Luca (2016) and Lewis and Zervas (2019) have found that a one-star increase in rating leads to a 5-9 percent increase in restaurant/hotel revenues.

demand does not change. Moreover, upward price adjustments in high-rating restaurants could explain the overall larger revenues.[49] In fact, these restaurants could benefit from their online reputation to charge higher prices without loosing much of their clientele. Nevertheless, the increase in employment among mid-rating establishments clearly indicates that more than just a price adjustment is going on in this category, and some demand-side dynamics must be driving their expansion. Unfortunately, my data does not allow to disentangle the specific mechanism behind the overall positive effects of the policy, and all these hypotheses remain plausible.

5.1.3 Additional findings

So far, I relied on the extensive margins of employment — i.e., the total number of workers — to measure restaurant size and proxy for output. Table A6 in the Appendix presents the results using the (log) number of total working days per month, which also capture adjustments at the intensive margins. Coefficients are in line with the previous ones, with some additional evidence suggesting negative and significant effects (at the 10% level) in low-rating restaurants, where the number of working days shrank by almost 7% after the policy. Such a reduction might be the result of a negative demand shock that these establishments experienced. Moreover, to isolate changes at the intensive margins only, I consider the number of monthly working days per worker. Appendix Table A7 reveals that, after the policy, each employee in high-rating restaurants worked, on average, 0.4 days more per month, while those in mid-rating ones worked almost half-a-day less. The first finding is consistent with the presence of capacity constraints in high-rating restaurants: rather than hiring additional employees and expand in size, these firms demanded more days of work to their current personnel. The second result is consistent with the hypothesis (covered in detail in Section 5.3) that lower-quality restaurants may attempt at improving their service quality through strategic employment choices, for instance by hiring new dining room staff while guaranteeing them better working conditions (e.g., shorter shifts).

Moreover, previous work (e.g., Luca 2016; Lewis and Zervas 2019) found that online reputation is more important for independent restaurants, where asymmetric information is more severe compared to chains. My data lack information on resultants' affiliation, but they contain details on their price tier. In this respect, cheap (e.g., fast food) and fancy starred places are expected to gain less from online reviews than those in the middle segment, even when their ratings are high. For instance, low and mid-budget tourists (which represent the majority of visitors) are more willing to substitute a low-price restaurant with a medium-price one, once they are reassured about the good quality of the latter. Yet, at the same time, fine-dining restaurants would remain outside of their consideration set. Consistently with this hypothesis, Appendix Table A8 shows that high-rating mid-price restaurants expanded their total employment by approximately 10% after the policy, while the corresponding coefficients for low- and high-price restaurants are negative and not significant.

[49] This is in line with the evidence on profit margin reported in Appendix Table A5, which shows that profits increased after the policy only in the high-rating category.

By contrast, employment decreased in low-rating cheap and expensive restaurants after the policy.

Figure 4: The impact on restaurant employment across levels of exposure to tourists

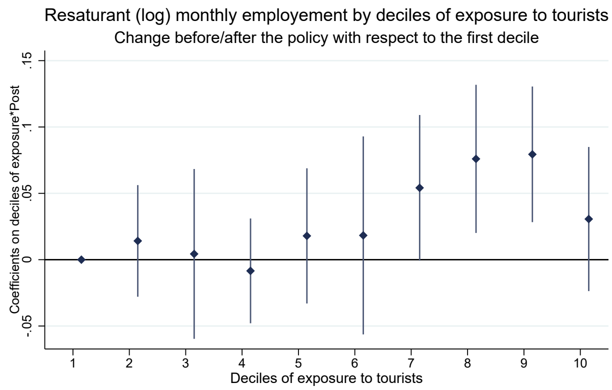

(a)

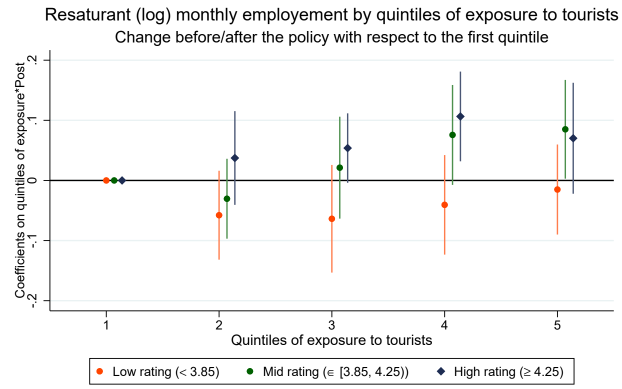

(b)

Notes to Panel (a): The graph reports estimates on the interactions of deciles of exposure*Post from a regression where each observation is a restaurant-month-year. The first decile is omitted. All controls and fixed-effects from the main analysis are included. Post takes value 1 after May 2017. The sample includes observations between Jan 2015 and Dec 2019. Heteroskedasticity-robust standard errors are clustered at municipality level. Bars depict 95% confidence intervals.
Notes to Panel (b): The graph reports estimates on the interactions of quintiles of exposure*Post from three separate regressions (low, mid, high-rating restaurants), where each observation is a restaurant-month-year. The first quintile is omitted. All controls and fixed-effects from the main analysis are included. Post takes value 1 after May 2017. The sample includes observations between Jan 2015 and Dec 2019. Heteroskedasticity-robust standard errors are clustered at municipality level. Bars depict 95% confidence intervals.

The baseline findings are robust to different specifications, measurements and samples. For instance, I examine whether my estimates are driven by the specific choice over the construction of the treatment variable $Tourist$. Particularly, rather than identifying just two groups of restaurants according to the median value of the probability measure, I consider its deciles and estimate a more flexible specification, interacting them with the dummy variable $Post$. Besides providing a robustness for the main results, this approach allows to study demand reallocation over space, e.g., from restaurants located in front of tourist attractions to those "hidden" in the surrounding alleys. To some extent, Panel (a) of Figure 4 suggests that such a reallocation is likely to take place. The impact of the roaming policy on employment is not statistically significant for restaurants at the 10th decile of exposure to tourists (the most visible ones), while the coefficients on the 7-8-9th deciles are driving the overall positive results. Moreover, coefficients on lower deciles are remarkably smaller in size, and always insignificant.[50] Section 7 conducts additional sensitivity analysis of the results.

5.2 Industry composition

The second set of hypotheses presented in Section 2 concerns the supply side, namely: (1) firms' decisions to stay in the market or not and, conditional on staying, (2) their level of investment

[50] In a similar fashion, Panel (b) of Figure 4 replicates the same exercise across both rating categories and quintiles of exposure to tourists. In this case, I use quintiles in order to have a sufficient and representative number of observations within each pair *(quintile, rating category)*. Point estimates displayed in the figure are qualitatively consistent with those from the main analysis and confirm that more touristy higher-rating restaurants drive the overall results. A similar conclusion holds for revenues, as discussed in Section 7.

into quality. This section covers the former, while the latter will be discussed in Section 5.3. Firm dynamics represents one potential mechanism through which the cheaper access to information from online reviews could affect the overall quality levels in the industry. Theory predicts that when consumers face lower search costs, those firms producing the lowest-quality products are more likely to be forced out of the market (i.e., a reduction in the adverse selection problem). To empirically test the effects of the roaming regulation on the industry composition and isolate the role of entry/exit dynamics (as opposed to quality upgrading), I track the presence of restaurants in the market over time by rating category. Particularly, I use the official date of registration and termination of the business and conduct the analysis both at the firm and ZIP-code levels.

5.2.1 Firm exit

Table 4: Impact on restaurant exit

	Y=1 if firm exits the market; Jan 2015 - Dec 2019				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	0.0011*	0.0016**	0.0031***	-0.0000	0.0015
	(0.0006)	(0.0006)	(0.0010)	(0.0009)	(0.0024)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	219835	217622	72133	76920	68569
Restaurants	4628	4576	1490	1571	1515
Clusters	86	86	59	71	71
Adj. R-squared	0.059	0.060	0.058	0.061	0.061
Mean Y pre-policy	0.003	0.003	0.003	0.003	0.004
DDD *p-value*				0.056	0.558

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in$ $[1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The firm-level framework described in equation (1) exploits the within-firm variation in the outcomes of interest over time, for the sample of firms with available Tripadvisor rating — and therefore operating in the market — at the time of the policy. As such, this setting only allows to study firm exit, and not entry.[51] Specifically, I construct a dummy variable that takes value 1 when a firm exits the market and 0 otherwise, and estimate (1) via OLS. Table 4 presents the estimation results of this linear probability model.

Columns (1-2) show that, after the policy, monthly exit rate among more touristy restaurants

[51]In equation (1), the coefficient of interest on $Tourist \times Post$ cannot be identified if firms' outcomes are not observed both before and after the policy. Moreover, if a firm enters the market later on, its Tripadvisor rating at the time of the policy would not be available. For these reasons, firm entry will be studied in the aggregate analysis presented later.

increased by 0.11-0.16 percentage points, with respect to less touristy ones. In comparison with the average pre-policy exit rate (0.3%), the frequency at which firms leave the market went up by approximately 35-55% during the 30 months after the new regulation. Column (3) shows that the overall increase in the exit rate is mainly driven by low-rating restaurants: their frequency to exit the market went up more than 0.30 percentage points, which corresponds to doubling the pre-policy average exit rate in this category. At the same time, exit rates in the mid- and high-rating groups were not significantly affected: the estimated coefficients in columns (4-5) are much smaller in size, and they are not statistically significant. Overall, this evidence supports the theoretical predictions of the model.[52]

5.2.2 Aggregate effects

Does the above result hold in the aggregate, when firm entry is taken into account? To answer this question, I rely on the ZIP-level framework described in equation (2), which provides a more suitable setting to study changes in the industry composition. For each ZIP code/month, I consider the (log) count of active restaurants — of any rating, as well as in the three rating categories previously identified — and regress it on the number of attractions, which is a proxy for exposure to tourist clientele. As in the above analysis, I use the Tripadvisor rating at the time of the policy to proxy for quality. To measure the quality of restaurants entering the market after the policy and assign them to one of the three rating categories, I use the most recent Tripadvisor rating.[53]

Table 5 presents the results. Column (1) indicates that, after the policy, the presence of one additional tourist attraction in the ZIP code is associated with a reduction in the overall number of active restaurants by 0.4%. Notably, columns (2-4) show that low-rating restaurants are the main drivers of such effect: their number decreased by 0.6% after the policy, for any additional attraction in the ZIP code. While the coefficient for the mid-rating category is also negative but not significant (-0.3%), the impact on high-rating restaurants was virtually null.[54] Columns (5-6) of Table 5 consider the percentage (expressed in 0-100 points) of active restaurants in each rating category, and show that one additional tourist attractions in the ZIP code made the proportion of low-rating restaurants shrink by more than 0.10 percentage points. Back-of-the-envelope calculations suggest that the share of low-rated restaurants operating in the most touristy neighborhoods (25 attractions) decreased by 2.5 percentage points after the policy, compared to non-touristy ZIP codes (0 attractions). Figure 5 plots the event-study estimates, which confirm the previous findings and provide visual evidence in favor of the parallel trends assumption.

Altogether, these findings bring empirical support for the hypothesis that lower search costs induced by cheaper access to online review platforms can make the industry more quality-oriented

[52] These results are in line with the findings of Hui et al. (2018) for online marketplaces (eBay).

[53] I focus on the most recent rating so to have a sufficient number of underlying reviews to compute it.

[54] This result challenges the view that cheaper access to information should reduce the number of high-demand firms, while increasing their market shares (Brynjolfsson et al., 2010). Most likely, this is due to the peculiarity of the restaurant industry, where capacity constraints limit the expansion of firm output above a certain threshold.

by forcing low-quality providers out of business (i.e., a reduction in the adverse selection problem). Sections 6 and 7 provide placebo exercises and robustness for the above results.

Table 5: The aggregate effects on industry composition

Y=	log(N. of active establishments)				% of active establishments	
	(1) All	(2) Low rating	(3) Mid rating	(4) High rating	(5) Low rating	(6) Mid rating
N. of attractions * Post	-0.004**	-0.006***	-0.003	0.001	-0.102**	0.010
	(0.002)	(0.002)	(0.002)	(0.002)	(0.049)	(0.096)
ZIP-code & Time FE	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓
Observations	7501	7501	7501	7501	7501	7501
ZIP-codes	127	127	127	127	127	127
Adj. R-squared	0.993	0.989	0.985	0.976	0.930	0.882
Mean Y pre-policy	29.90	10.43	10.31	9.16	32.76	34.65

Post=1 if date is after May 2017. Each observation is a ZIP-code-month-year. All regressions include the distance of the ZIP-code to Rome city center interacted with Post. The sample includes observations between Jan 2015 and Dec 2019. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. If the restaurant entered the market after the policy, the most recent rating is considered. Heteroskedasticity-robust standard errors clustered at ZIP-code level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure 5: Event-study estimates for industry composition

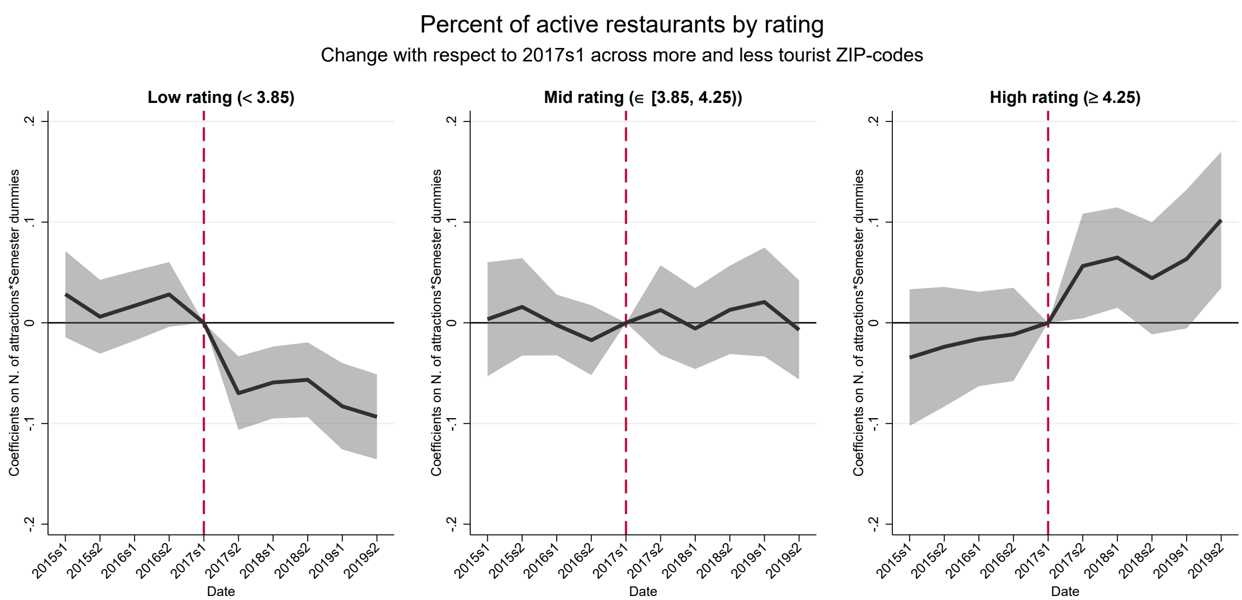

Notes: The graph reports estimated coefficients on the interactions of N. of attractions*Semester dummies from three separate regressions where each observation is a ZIP-code-month-year. Controls and fixed-effects from the ZIP-code-level analysis are included. The omitted semester is 2017s1. The sample includes observations between Jan 2015 and Dec 2019. Shaded areas depict 95% confidence intervals. Rating is computed at the time of the policy. If the restaurant entered the market after the policy, the most recent rating is considered.

5.3 Restaurant quality upgrading

The last set of theoretical predictions state that reductions in consumer search costs affect restaurants' incentives to improve quality, and more so for lower-quality establishments (i.e., a reduction in the moral hazard problem). To test this hypothesis, I consider several proxies for both input and output quality, as discussed in Section 2.4. In particular, I focus on restaurants' hiring decisions over workers with different experiences, as well as their wages, as a proxy for changes in the service quality. Moreover, I use the online reputation of the restaurant as reflected in the average dynamic Tripadvisor rating to proxy for output quality. While hiring decisions capture a more objective dimension of quality, online ratings reflect the subjective experience of the consumers. Thus, to some extent, the two sources complement each other. The analysis considers the employment choices of restaurants while they are in operation, including those that exited the market after the policy. As discussed in Section 7, all results are robust to the exclusion of restaurants that went out of business after the policy.

5.3.1 Hiring decisions and salaries

I take advantage of the employer-employee matched data and investigate if, in their attempt at improving service quality, restaurants are more likely to hire workers with better curricula, as measured by their previous experience in the restaurant sector.[55] First, I consider the full employment history of every newly-hired employee in my sample of restaurants, and construct dummy variables indicating whether, by the time of their appointment, they had previously worked in the restaurant sector or not.[56] For each outcome, I estimate equation (1) on the entire sample of restaurants, as well as on the three sub-groups corresponding to the different rating categories.

In columns (1-4) of Table 6 the dependent variable takes value 1 in months in which the restaurant hires a new employee who had previously worked in other restaurants, and 0 otherwise.[57] Column (1) indicates that, after the regulation, restaurants more exposed to tourists became almost 1-percentage-point more likely to hire experienced employees, which is about 10% of the pre-policy value. Consistently with the theoretical prediction, columns (2-4) show that low- and mid-rating restaurants drive the effect. Their probability to hire workers with better curricula significantly went up by 0.9-1.1 percentage points, an increase of 9-16% with respect to the pre-policy mean. By contrast, the coefficient for high-rating restaurants is close to 0 and not statically significant.

These results could be due to increased employee turnover in lower-rated restaurants, rather than their recruiting strategy being intentionally targeted at more experienced workers. However, columns (5-6) of Table 6 indicate that this is unlikely to be the case. In fact, the probability of hiring workers with no experience in the industry decreased by 10% after the policy among more touristy low-rating restaurants, while it increased for the mid- and high-rating ones by almost 18 and 12%,

[55] I exclusively focus on the experience dimension of the worker's curriculum. Unfortunately, other factors (such as education) are not available in the data. Nevertheless, this should not pose a critical obstacle to my analysis, since in the restaurant sector previous experience is likely to be more informative than education to signal the skills of waiters and other room staff.

[56] The restaurant sector is defined by firms with ATECO codes 56.10.11 (dine-in restaurants), 56.10.12 (agriturismi), 56.10.20 (take-away restaurants), 56.10.30 (bakeries).

[57] To facilitate the interpretation of the coefficient estimates, I use a linear probability model (OLS) as the benchmark specification. Coefficients from a Logit model are qualitatively similar and are reported in Table A9 in the Appendix.

Table 6: Impact on restaurant hiring decisions (extensive margins)

Y=1 if firm hires worker	*with* previous experience in restaurants				*without* experience in restaurants		
	(1) Full sample	(2) Low rating	(3) Mid rating	(4) High rating	(5) Low rating	(6) Mid rating	(7) High rating
Tourist*Post	0.009***	0.009***	0.011***	0.002	-0.006**	0.011**	0.007**
	(0.002)	(0.003)	(0.003)	(0.004)	(0.002)	(0.004)	(0.003)
Restaurant & Time FE	✓	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
Observations	217622	72133	76920	68569	72133	76920	68569
Restaurants	4576	1490	1571	1515	1490	1571	1515
Clusters	86	59	71	71	59	71	71
Adj. R-squared	0.124	0.143	0.104	0.116	0.049	0.043	0.037
Mean Y pre-policy	0.08	0.10	0.07	0.08	0.06	0.06	0.06
DDD *p-value*			0.259	0.047		0.000	0.002

The sample includes observations between Jan 2015 and Dec 2019. Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (3-4) and (6-7) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

respectively, compared to their pre-policy values. Altogether, these estimates suggest that while low- and high-rating restaurants seem to focus their recruiting efforts on opposite types of workers (experienced vs. inexperienced, respectively), mid-rating ones hire from a more heterogeneous pool of candidates.

While the above analysis looks at the extensive margins of targeted recruiting strategies (i.e., whether or not restaurants hire experienced workers), consistent findings are obtained when considering the intensive margins of worker experience. In this case, I restrict the attention only to those months in which the restaurant hires/fires an employee, the contract terminates, or the employee voluntarily quits the job. I then measure the cumulative experience of such workers by counting the total number of months they have been employed in the restaurant sector in the past. This way, I can quantify the impact of the policy on the gain/loss in human capital that restaurants face.

Table 7 shows the results. Column (1) indicates that, after the policy, restaurants in more touristy areas hire workers with additional 1.5 months of previous experience in the industry, compared to less touristy ones. Columns (2-4) show that low-rating restaurants are the drivers of such a change: after the policy, they hire workers with 3 additional months of experience in the industry, which corresponds to 22% of the pre-policy value. The coefficients for the mid- and high-rating categories are much smaller and not statistically significant, suggesting that the accumulation of human capital mainly takes place in lower-rated restaurants. By contrast, high-rating establishments appear to loose human capital. Columns (5-7) consider the employment history of those workers who left the firm, either because they decided to quit, their contract expired or they got fired.[58]

[58]Note that this definition is intentionally broad, for instance it also includes workers who reached their retirement

Table 7: Impact on restaurant hiring decisions (intensive margins)

Y=Months of experience in restaurants of	newly-hired employees				quitting/fired employees		
	(1) Full sample	(2) Low rating	(3) Mid rating	(4) High rating	(5) Low rating	(6) Mid rating	(7) High rating
Tourist*Post	1.469*	2.977***	0.789	0.465	0.177	-0.440	2.375**
	(0.834)	(1.083)	(1.160)	(0.544)	(1.034)	(1.192)	(1.181)
Restaurant & Time FE	✓	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
Observations	30059	11318	10205	8536	12281	10395	8131
Restaurants	3531	1163	1220	1148	1197	1226	1136
Clusters	76	53	59	61	51	57	58
Adj. R-squared	0.117	0.109	0.117	0.127	0.190	0.170	0.183
Mean Y pre-policy	13.0	13.5	13.5	11.8	25.8	27.0	21.5
DDD *p-value*			0.034	0.000		0.598	0.460

The sample includes observations between Jan 2015 and Dec 2019. Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (3-4) and (6-7) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

While no effect is detectable for low- and mid-rating establishments, employees that are let go by high-rating restaurants are, on average, 2.4 months more experienced compared to those in less touristy restaurants.

Changes in the the labor quality should be reflected in the firm production costs. In a competitive labor market, firms must pay higher wages in order to attract employees with better skills. Figure 6 depicts the evolution of average gross daily salaries paid by the restaurants in the three rating categories.[59] Consistently with the evidence on workers' experience presented so far, these event-study estimates point out an increase in the salaries paid to employees of more touristy low-rating restaurants by more than €1 a day. By contrast, salaries in high-rating restaurants decreased by a similar amount, on average, while they did not change in the mid-rating category. At the same time, the figures provide evidence in support of the parallel trends assumption. Regression estimates reported in Table 8 confirm the graphical analysis. Salaries in low-rating (high-rating) establishments grew (shrank) by almost 2% (1.8%) with respect to their pre-policy values. By contrast, no significant change in average salaries is detected in the overall sample and in the mid-rating category.[60]

age. This is done to capture the overall loss in human capital that restaurants experience, by looking at any worker who left the firm, irrespective of the reason.

[59]To make the salaries of full- and part-time employees comparable, I compute the full-time equivalent salary for part-time employees, using the percentage of the part-time as reported in their contract.

[60]These results are robust to the use of a logarithmic scale of salaries, as shown in Appendix Table A10.

Figure 6: Event-study estimates for restaurant daily salaries (€)

Notes: The graph reports estimated coefficients on the interactions of Tourist*Semester dummies from three separate regressions where each observation is a restaurant-month-year. Full-time equivalent salary is computed for part-time employees, according to the percentage of the part-time as reported in their contract. All controls and fixed-effects from the main analysis are included. The omitted semester is 2017s1. Tourist restaurants are those with a measure of exposure above the median. The sample includes observations between Jan 2015 and Dec 2019. Shaded areas depict 95% confidence intervals.

Table 8: Impact on restaurant daily salaries

	Y=Average daily salary (€); Jan 2015 - Dec 2019				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	-0.010	0.038	1.312***	-0.120	-1.125**
	(0.245)	(0.257)	(0.374)	(0.440)	(0.448)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	200402	199026	67507	70593	60926
Restaurants	4558	4512	1492	1538	1482
Clusters	86	86	59	71	70
Adj. R-squared	0.467	0.469	0.485	0.465	0.451
Mean Y pre-policy	64.9	64.9	66.0	65.0	63.5
DDD *p-value*				0.206	0.015

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Altogether, these findings provide evidence that lowering the costs for consumers to access review platforms can affect firms' incentives to upgrade product quality through strategic employee turnover, especially for those producers with the highest margins of improvement. In particular, restaurants with initially lower qualities targeted their hiring efforts at more experienced workers (both at the intensive and extensive margins) and ended-up paying higher salaries. By contrast, restaurants that were already selling higher qualities took advantage of their established online reputation to divest in human capital and save in production costs.[61] Eventually, these opposite recruiting strategies might generate human capital flows from high- to low-rating establishments, contributing to the process of quality upgrading. As Appendix Table A11 shows, low- and mid-rating restaurants became more likely to hire workers previously employed in establishments with better Tripadvisor ratings.

5.3.2 Tripadvisor rating

I rely on Tripadvisor rating as a proxy for the online reputation of the restaurants. In particular, I compute the moving average of the monthly Tripadvisor rating over dynamic 5-month windows for all the restaurants in my sample. Employment choices are likely to impact ratings. Appendix Table A12 shows that recruiting workers with previous experience in the restaurant sector is associated with more positive Tripadvisor reviews in the subsequent months. The correlation is even larger when the new employee comes from a higher-rating establishment. By contrast, hiring non-experienced employees has no impact on subsequent rating. This evidence suggests that lower costs for consumers to access review platforms can affect the subsequent online reputation of restaurants through their employment decisions (as in Shin et al. 2023).

Figure 7 displays event-study estimates across the three rating categories, and shows that restaurants in the low- and mid-rating groups received better ratings after the policy, with peaks reaching almost 0.11 and 0.09 points, respectively. Coefficient estimates of model (1) reported in Table 9 confirm these patters. After the regulation, the 5-month Tripadvisor rating of more touristy restaurants improved by almost 0.05 points overall (i.e., a 1.3% increase with respect to the pre-policy mean), and by 0.09 points (2.5%) and 0.08 points (1.9%) in low- and mid-rating establishments, respectively. By contrast, the coefficient for the high-rating category is virtually zero and not statistically significant, indicating no change in the online reputation of those restaurants already at the top of the rating distribution.

These findings confirm the predictions of the model and are generally consistent with the evidence on the hiring decisions presented above. However, two facts might require further explanation. The first one is the rating improvement among mid-rating restaurants, whose employment strategies were not exclusively targeting experienced workers. One reason might be the use of management responses to consumer reviews as a way to obtain more positive feedback. In this respect, Appendix Table A13 shows a significant increase in replies among restaurants in the mid-rating category, pointing out a potential reason for their reputation upgrading.[62] Another possibility is the growth in personnel

[61] Appendix Table A5 indicates that these decisions eventually impact on the restaurant profitability, since the profit margin of high-rating establishments grew by 2.5 points after the policy. By contrast, profits in low- and mid-rating restaurants decreased.

[62] For instance, Proserpio and Zervas 2017 find that after responding to reviews on Tripadvisor, hotels' ratings increase by 0.12 points, an effect that is comparable to my estimates.

Figure 7: Event-study estimates for restaurant Tripadvisor rating

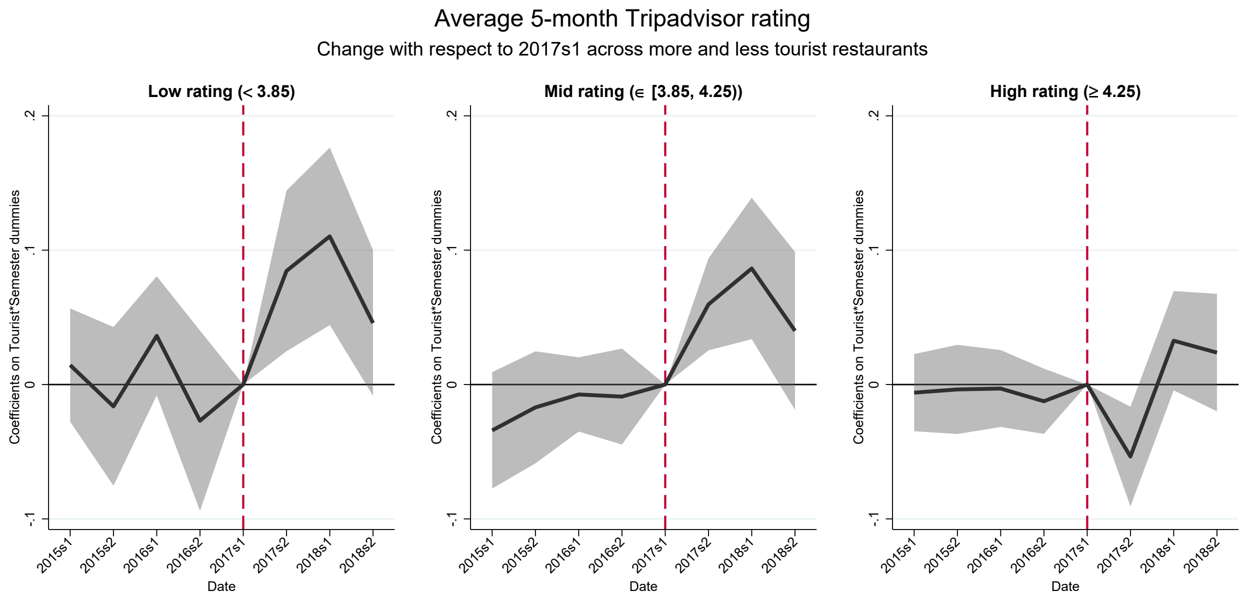

Notes: The graph reports estimated coefficients on the interactions of Tourist*Semester dummies from three separate regressions where each observation is a restaurant-month-year. All controls and fixed-effects from the main analysis are included. The omitted semester is 2017s1. Tourist restaurants are those with a measure of exposure above the median. The sample includes observations between Jan 2015 and Dec 2018. Shaded areas depict 95% confidence intervals.

Table 9: Impact on restaurant Tripadvisor rating

	Y=Average 5-month Tripadvisor rating; Jan 2015 - Dec 2018				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	0.040***	0.049***	0.087***	0.077***	-0.003
	(0.012)	(0.012)	(0.017)	(0.019)	(0.012)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	147274	146620	48577	53659	44384
Restaurants	4373	4330	1413	1499	1418
Clusters	86	86	59	70	70
Adj. R-squared	0.503	0.504	0.324	0.251	0.297
Mean Y pre-policy	3.98	3.98	3.51	4.05	4.43
DDD *p-value*				0.000	0.000

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

and improvements in working conditions discussed in Section 5.1, both of which could have positive effects on the service quality. Finally, mid-rating restaurants might have started using better raw materials in their kitchens, as their annual net purchases increased after the policy (Appendix Table A14), potentially indicating the use of higher-quality ingredients in preparing the recipes.[63]

The second empirical finding that might be puzzling is the absence of a decline in the online reputation for high-rating restaurants, despite the documented loss in human capital. For them, replies to reviews, expansions in their personnel and the use of better ingredients are not plausible explanations.[64] An alternative reason, which is in line with the theoretical model, is that these establishments are very capable (efficient) at managing their factors of productions. This efficiency advantage allows them to employ less skilled workers, save in costs and increase profits without compromising output quality.

Finally, two additional concerns on the use of online ratings to proxy for quality are worth mentioning. The first one is the presence of fake reviews (He et al. 2022), which could jeopardize the interpretation of the results. The second concern is the possibility that higher ratings reflect a better matching with the restaurants rather than an actual quality improvement. Regarding the former, Tripadvisor has significantly improved its effort and ability to detect and remove fraudulent reviews in a timely manner over the last years (2021 Review Transparency Report), which should attenuate this problem. Regarding both the former and the latter, the results on Tripadvsor ratings are observable and similar across reviewers of different nationalities.[65] This alleviates the concern that fake reviews drive the results, as they should be consistently present in multiple languages. Moreover, the fact that adjustments in rating occur also for Italians (who were not directly affected by the policy) speak in favor of an actual quality improvement rather than just a better matching.

6 Placebos

This section describes a series of placebo policy-permutation tests conducted in the period before the roaming regulation (Jan 2012 - Dec 2016) to assess its exogeneity with respect to other potential confounding factors or existing pre-trends in the outcomes that might explain the observed results (such as seasonality). In practice, I replicate the main analysis interacting the *Tourist* dummy variable with placebo policy-dates between May 2012 and May 2016 — for a total of 49 regressions for each outcome —, following the approach used to carry out randomization inference in experiments (Gerber and Green 2012).[66] Then, I plot the histograms of all estimated placebo coefficients for the whole sample of restaurants, as well as for the three different rating categories considered in the main

[63] Net purchases reflect any expenditure in inputs other than labor. As such, the variable might also include the purchase of specialized services from online advertising and customer management companies.

[64] Appendix Table A13 shows that, if anything, high-rating restaurants engaged less with reviewers, after the policy. Moreover, as discussed in Section 5.1, these restaurants did not significantly employ more workers. Finally, as shown in Appendix Table A14, annual purchases remained the same.

[65] This analysis is available upon request.

[66] In my case, rather than varying the composition of the control group, I modify the time dummy *Post*.

analysis. In this case, I use the Tripadvisor rating of the restaurant at the time of the respective placebo policy-date. Figure 8 shows an example of the output of this procedure for restaurant employment. The vertical dashed lines depict the respective policy coefficient estimated in Section 5.1. Red (black) lines indicate that the coefficient was significant (insignificant).

Figure 8: Permutation test for restaurant employment

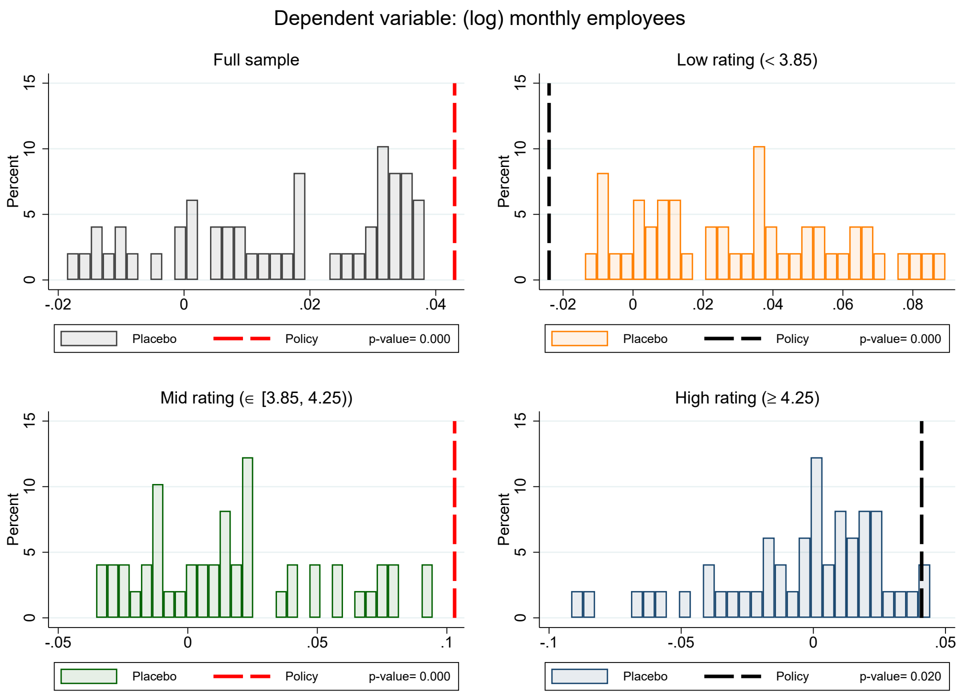

Notes: Each panel plots the distribution of coefficients on Tourist*Post-Month, where Month is between May 2012 and May 2016, estimated on a sample of observations between Jan 2012 and Dec 2016. Tripadvisor rating and the respective category are calculated in each month. The vertical dashed lines report the actual policy coefficients, as estimated in the main analysis. The line is red when the respective coefficient is significant at least at the 10% confidence level, and black otherwise.

The result of this exercise would speak in favor of the identifying assumption if the policy estimates lie at the extremes of the distributions of placebo coefficients. More formally, I compute and report the *p-value* of the permutation test by counting the number of times the placebo coefficients are more extreme than the policy estimate, and dividing it by the total number of coefficients. Low *p-values* imply that, most of the times, the policy estimates are larger than the placebos, alleviating concerns on potentially diverging pre-trends across more and less touristy restaurants. For example, in each subplot of Figure 8 the *p-value* is always below 2%. This means that the estimated 4.2% increase in overall restaurant employment and the 10% effect found in the mid-rating category can arguably be attributed to the roaming regulation.

I summarize the results for the other outcomes in a series of figures reported in the Online Appendix (Figures C1 to C6). When the policy estimates are significant, they are always at the extremes of the distribution of placebo coefficients. The associated *p-values* are consistently small and always below the conventional 5% level, implying that only in only a very few cases the placebos have magnitudes larger than the estimated policy coefficients.

Finally, additional placebo exercises are reported in the event-study estimates. These graphs provide visual evidence on the plausibility of the parallel-trend assumption underlying the empirical strategy of the paper. Figures 3, 5, 6, and 7 reported in the text suggest that this assumption is likely to hold for restaurant employment, industry composition, salaries and rating, respectively. Additional figures reported in the Online Appendix (C7 to C10) confirm that similar conclusions generally hold for restaurant revenues, exit and hiring decisions, although in some cases these graphs are less clear because of the nature of certain outcomes (e.g., binary variables for exit and hiring decisions). Altogether, the above placebo exercises corroborate the identifying assumptions and validate the use of the Difference-in-Differences strategy in this context.

7 Robustness

This section carries out a series of additional estimations to investigate the sensitivity of the main results to different measurements, samples and clustering units. First, a potential concern with the firm-level estimates is that they might be driven by both (1) the definition of the binary variable *Tourist* and (2) the construction of the underlying measure of exposure to tourist clientele $P(i)$, defined in Section 4. To address the first point, I replicate the baseline estimation using a more flexible specification, in which I consider dummies for quintiles of exposure to tourists — rather than the median value — interacted with the variable *Post*. This procedure generalizes the results to study the effects of the policy along a more continuous gradient of exposure and, at the same time, guarantees that estimates remain easy to interpret. I show the results using a series of figures, reporting the effects at each quintile (with respect to the first one, which is the omitted category) as well as across the three rating categories. Panel (b) of Figure 4 already discussed in Section 5.1 presents the results for restaurant employment. Consistently with the main analysis, better-rated restaurants at the higher quintiles of the tourist-exposure distribution are the drivers of employment growth. Online Appendix Figures C11 to C16 show that similar conclusions hold for restaurant revenues, exit, hiring decisions, salaries and rating, with point estimates being consistently larger for the highest quintiles. This evidence corroborates the benchmark specification and alleviates the concern that the main results are driven by the specific definition of the binary variable $Tourist$.

Moreover, the procedure I adopted to construct the probability measure $P(i)$ could also influence the firm-level outcomes. To address this concern, I study the sensitivity of the main analysis to the inclusion of all alternatives routes provided by the Google Maps API in the computation process. Tables D1 to D6 in the Online Appendix show that estimated coefficients are always qualitatively, and often quantitatively similar to those from the benchmark specification.

In addition, another potential concern is that sorting of restaurants into exit might bias the estimates. In fact, the baseline sample includes the approximately 560 firms that, at some point after the policy, ceased their operations and left the market. Hence, I replicate the estimations on the sample of restaurants that survived throughout the whole 30-month period after the roaming regulation. Online Appendix Tables D7 to D12 show that coefficients are very similar to those from the main analysis. Certain effects — such as the hiring of experienced workers and improvements in ratings for low-rating restaurants — are even stronger suggesting that, if anything, the presence of exiting firms might attenuate the results.

Finally, my estimates are generally robust to different clustering units. Online Appendix Tables D13 to D19 use the 127 ZIP codes (that are smaller than municipalities) to cluster the standard errors, and show that — with the exception of restaurant revenues, for which power is lower due to the annual frequency — the significance of the coefficients is not remarkably altered.

8 Industry-wide effects of reviews

What are the industry-wide effects of giving access to online reviews to all the customers? The Difference-in-Differences estimates provide reduced-form evidence on the impact of reducing internet fees for some tourists on restaurants in the Province of Rome. As such, the regression coefficients represent intent-to-treat effects of the provision of information, because in the treatment group (identified by the *Tourist* binary variable) only a fraction of the clientele — namely, the EU users — benefited from the cheaper internet costs. Moreover, the specificity of the context (Rome and its Province) poses a limit to the generalization of the results to the whole restaurant industry.

Recovering industry-wide consequences requires two further steps. First, I estimate treatment-on-the-treated effects, which can be obtained by re-scaling the above coefficients by the expected share of customers impacted by the roaming policy (EU users). Second, I generalize these effects to the entire Italian restaurant industry, re-weighting the estimates by the share of Italian restaurants that are likely to be exposed to tourists. The procedure requires three additional assumptions:

1. Among *Tourist* restaurants, take-up of the policy was 23%, which corresponds to the percent of Tripadvisor reviews from EU travelers in the post-policy period.

2. Non-tourist restaurants are not affected by the policy.

3. The share of tourist restaurants in Italy is 8%, which corresponds to the fraction of establishments located in ZIP-codes with at least one top-tourist attraction.

Assumption (1) relies on the share of Tripadvisor contributions from Europeans to proxy for their usage of the platform in the post-policy period.[67] Although a gap in the absolute volume of EU users who consume and produce reviews plausibly exists, the ratio EU/Total contributions should provide a reasonable approximation for the relative usage among Europeans compared to others. Assumption (2) is about the absence of spillovers and requires that outcomes in the control group (i.e., non-tourist restaurants) do not change after the policy. This assumption is likely to hold for at least two reasons. First, estimates of the policy by deciles/quintiles of exposure to tourists

[67] Tripadvisor usage statistics disaggregated by origin of reviewers are not publicly available.

show that effects are driven by restaurants at higher levels of exposure (e.g., Figures 4 and C11 to C16). Second, as shown in Online Appendix Figure C17, de-trended average employment in non-tourist restaurants remains stable after the roaming regulation. Finally, for assumption (3), I collect additional information from Tripadvisor on the top-100 tourist attractions in Italy (based on their total volume of reviews) and then compute the fraction of Italian restaurants that are located in the ZIP-codes with at least one of such attractions.

Table 10: Industry-wide effects of reviews

	Adjusted effect of reviews	2016-2019 growth rate	Percent of growth explained by reviews
Annual revenues	1.6%	13.2%	**12.1%**
Monthly employment	1.5%	29.7%	**5.1%**
	Adjusted effect of reviews	Exit rate in Covid year 2019-2020	Percent of exit rate explained by reviews
Annual exit rate	0.46 p.p.	15.7%	**2.9%**

Notes: Revenue growth rate refers to the 2016-2018 period.

I then consider the most conservative estimates of the policy on revenues, employment and exit as reported in columns (1-2) of Tables 2, 3 and 4. I divide these coefficients by 0.23 to recover treatment-on-the-treated effects, multiply them by 0.08 to re-weight for tourists areas, and then compare them with the aggregate trends in the industry. Results are reported in Table 10 and suggest that promoting access to review platforms has substantial consequences for the whole Italian restaurant industry. Back-of-the-envelop calculations point out that reducing the costs for all consumers to access online reviews leads to an overall increase in restaurant revenues, employment and exit rate by 1.6%, 1.5% and 0.5 p.p., respectively. The first two figures correspond to about 12% and 5% of the overall growth in revenue and employment experienced by restaurants between 2016 and 2019, respectively. While the last figure corresponds to almost 3% of the exit rate faced by the industry during the first year of the Covid-19 pandemic. Altogether, these results indicate that lower costs for consumers to access online word-of-mouth can have consequential effects on the labor market, and on performance and composition of firms operating in industries generally affected by asymmetric information.

9 Conclusions

The digital era has changed the way consumers get and share information. For instance, review platforms lower information frictions by helping consumers search and verify products before purchase. Yet, their consequences for markets with information asymmetries remain unclear. The implications are both managerial and regulatory. For firms, intensified competition from online word-of-mouth requires managers to strategically allocate resources and invest in reputation building. For regulatory bodies, customer reviews could represent a cost-effective complement to time- and labor-intensive monitoring interventions, such as sanitary inspections. While there is general optimism around the possibility for review platforms to create reputation and feedback mechanisms that incentivize firms to upgrade quality, empirical evidence is scarce. This paper shows that cheaper access to online reviews — induced by the abolition of internet fees — can change how firms operate and make the service industry more quality oriented.

First, I built a model in which consumers with heterogeneous search costs engage in sequential search to buy one unit of a vertically differentiated product, while firms with heterogeneous abilities endogenously select into production and compete in quality. The model predicts that lower search costs positively affect the equilibrium quality levels but have differing effects across businesses. Some of the lowest-quality firms exit the market while the surviving ones invest in quality-upgrading.

To test these hypotheses, I focused on the restaurant industry in the province of Rome and assembled a unique dataset which combines restaurants' information from Tripadvisor with rich administrative establishment-level data. I took advantage of a plausibly exogenous reduction in the cost of mobile internet — caused by a policy that abolished the roaming charges for tourists in the EU — to identify the effects of access to online reviews on consumers' behavior, restaurants' incentives to upgrade quality, and industry composition. Using a Difference-in-Differences strategy, I compared the variation (before/after policy) in the outcomes across restaurants that are more and less exposed to tourist clientele. I estimated the model on the whole sample and on three sub-samples of restaurants with different ratings at the time of the policy: namely, low, mid and high rating.

I showed that, after the policy, revenues increased in mid- and high-rating restaurants, while employment grew only in the mid-rating category, suggesting that high-rating establishments were already producing at full capacity. I then analyzed the supply side. First, I showed that for low-rating restaurants, the probability to exit the market doubled after the policy compared to the pre-policy period. Moreover, by aggregating observations at the ZIP-code level, I found that the share of low-rating firms operating in the most touristy neighborhoods decreased by 2.5 p.p. after the policy, compared to non-touristy ZIP codes. Then, I analyzed the behavior of surviving firms. I found that low-rating restaurants focused their recruiting efforts on workers with previous experience in the restaurant industry and ended-up paying higher salaries. Eventually, low- and mid-rating establishments improved their online reputation, as their dynamic Tripadvisor rating increased. The industry-wide consequences are substantial. Abating the costs for all consumers to access review platforms could lead to an overall increase in restaurant revenues, employment and exit rate by 1.6%, 1.5% and 0.5 p.p., respectively.

Altogether, my findings indicate that lower costs for consumers to access online reviews create the conditions that force some low-quality providers (i.e., the tourist traps) out of business and encourage others to produce higher-quality goods and engage in reputation building. These results highlight the possibility of quality-upgrading in the restaurant industry through policies that reduce the cost for consumers to browse the internet and promote the use of review platforms.

Funding and Competing Interests

The author has received partial research support from the Barcelona Institute for Political Economy and Governance. The author does not have any competing interests to this work.

References

Akerlof, G. A. (1970), 'The market for "lemons": Quality uncertainty and the market mechanism', *The Quarterly Journal of Economics* **84**(3), 488–500.
URL: *http://www.jstor.org/stable/1879431*

Allen, T. (2014), 'Information frictions in trade', *Econometrica* **82**(6), 2041–2083.

Alyakoob, M. and Rahman, M. S. (2022), 'Shared prosperity (or lack thereof) in the sharing economy', *Information Systems Research* **33**(2), 638–658.

Ananthakrishnan, U., Proserpio, D. and Sharma, S. (2023), 'I hear you: Does quality improve with customer voice?', *Marketing Science* .

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

Anderson, S. J., Chandy, R. and Zia, B. (2018), 'Pathways to profits: The impact of marketing vs. finance skills on business performance', *Management Science* **64**(12), 5559–5583.

Anderson, S. P. and Renault, R. (1999), 'Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model', *The RAND Journal of Economics* pp. 719–735.

Bai, J. (2021), 'Melons as lemons: Asymmetric information, consumer learning and seller reputation', *CID Working Paper Series* .

Bar-Isaac, H., Caruana, G. and Cuñat, V. (2012), 'Search, design, and market structure', *American Economic Review* **102**(2), 1140–60.

Barahona, N., Otero, C. and Otero, S. (2023), 'Equilibrium effects of food labeling policies', *Econometrica* **91**(3), 839–868.

Basu, S. and Fernald, J. G. (1997), 'Returns to scale in us production: Estimates and implications', *Journal of political economy* **105**(2), 249–283.

Brambilla, I., Lederman, D. and Porto, G. (2012), 'Exports, export destinations, and skills', *American Economic Review* **102**(7), 3406–3438.

Brynjolfsson, E., Hu, Y. and Smith, M. D. (2010), 'Research commentary—long tails vs. superstars: The effect of information technology on product variety and sales concentration patterns', *Information Systems Research* **21**(4), 736–747.

Cabral, L. and Hortacsu, A. (2010), 'The dynamics of seller reputation: Evidence from ebay', *The Journal of Industrial Economics* **58**(1), 54–78.

Chan, Y.-S. and Leland, H. (1982), 'Prices and qualities in markets with costly information', *The Review of Economic Studies* **49**(4), 499–516.

Chen, Y. and Xie, J. (2005), 'Third-party product review and firm marketing strategy', *Marketing science* **24**(2), 218–240.

Chen, Y. and Xie, J. (2008), 'Online consumer review: Word-of-mouth as a new element of marketing communication mix', *Management science* **54**(3), 477–491.

Chevalier, J. A., Dover, Y. and Mayzlin, D. (2018), 'Channels of impact: User reviews when quality is dynamic and managers respond', *Marketing Science* **37**(5), 688–709.

Chevalier, J. A. and Mayzlin, D. (2006), 'The effect of word of mouth on sales: Online book reviews', *Journal of marketing research* **43**(3), 345–354.

Cooper, R. and Ross, T. W. (1984), 'Prices, product qualities and asymmetric information: The competitive case', *The Review of Economic Studies* **51**(2), 197–207.

Dai, W. and Luca, M. (2020), 'Digitizing disclosure: The case of restaurant hygiene scores', *American Economic Journal: Microeconomics* **12**(2), 41–59.

Dall'orso, J., Gauriot, R. and Page, L. (2016), Disappointment looms around the corner: Visibility and local businesses' market power, Technical report, QUT Business School.

Elfenbein, D. W., Fisman, R. and McManus, B. (2015), 'Market structure, reputation, and the value of quality certification', *American Economic Journal: Microeconomics* **7**(4), 83–108.

Eliashberg, J. and Shugan, S. M. (1997), 'Film critics: Influencers or predictors?', *Journal of marketing* **61**(2), 68–78.

Ershov, D. (2020), 'Consumer product discovery costs, entry, quality and congestion in online markets', *Unpublished manuscript* .

Fang, L. (2022), 'The effects of online review platforms on restaurant revenue, consumer learning, and welfare', *Management Science* .

Farronato, C. and Zervas, G. (2022), Consumer reviews and regulation: evidence from nyc restaurants, Technical report, National Bureau of Economic Research.

Fishman, A. and Levy, N. (2015), 'Search costs and investment in quality', *The Journal of Industrial Economics* **63**(4), 625–641.

Gennaioli, N., La Porta, R., Lopez-de Silanes, F. and Shleifer, A. (2013), 'Human capital and regional development', *The Quarterly journal of economics* **128**(1), 105–164.

Gerber, A. S. and Green, D. P. (2012), *Field experiments: Design, analysis, and interpretation*, WW Norton.

Ghose, A., Goldfarb, A. and Han, S. P. (2013), 'How is the mobile internet different? search costs and local activities', *Information Systems Research* **24**(3), 613–631.

Goldfarb, A. and Tucker, C. (2019), 'Digital economics', *Journal of Economic Literature* **57**(1), 3–43.

Goldmanis, M., Hortaçsu, A., Syverson, C. and Emre, Ö. (2010), 'E-commerce and the market structure of retail industries', *The Economic Journal* **120**(545), 651–682.

Grzybowski, L. and Muñoz, A. (2020), 'Impact of roaming regulation on revenues and prices of mobile operators in the eu', *Unpublished manuscript* .

Hansman, C., Hjort, J., León-Ciliotta, G. and Teachout, M. (2020), 'Vertical integration, supplier behavior, and quality upgrading among exporters', *Journal of Political Economy* **128**(9), 3570–3625.

He, S., Hollenbeck, B. and Proserpio, D. (2022), 'The market for fake reviews', *Marketing Science* .

Hjort, J., Iyer, V. and De Rochambeau, G. (2020), Informational barriers to market access: Experimental evidence from liberian firms, Technical report, National Bureau of Economic Research.

Hollenbeck, B. (2018), 'Online reputation mechanisms and the decreasing value of chain affiliation', *Journal of Marketing Research* **55**(5), 636–654.

Hollenbeck, B., Moorthy, S. and Proserpio, D. (2019), 'Advertising strategy in the presence of reviews: An empirical analysis', *Marketing Science* **38**(5), 793–811.

Hopenhayn, H. A. (1992), 'Entry, exit, and firm dynamics in long run equilibrium', *Econometrica: Journal of the Econometric Society* pp. 1127–1150.

Hui, X., Saeedi, M., Spagnolo, G. and Tadelis, S. (2023), 'Raising the bar: Certification thresholds and market outcomes.', *American Economic Journal: Microeconomics* **15**(2), 599–626.

Hui, X., Saeedi, M. and Sundaresan, N. (2018), 'Adverse selection or moral hazard, an empirical study', *The Journal of Industrial Economics* **66**(3), 610–649.

Hunter, M. (2022), Chasing stars: Firms' strategic responses to online consumer ratings, *in* 'Proceedings of the 23rd ACM Conference on Economics and Computation', pp. 298–299.

Jensen, R. (2007), 'The digital provide: Information (technology), market performance, and welfare in the south indian fisheries sector', *The quarterly journal of economics* **122**(3), 879–924.

Jensen, R. and Miller, N. H. (2018), 'Market integration, demand, and the growth of firms: Evidence from a natural experiment in india', *American Economic Review* **108**(12), 3583–3625.

Jin, G. Z. and Leslie, P. (2003), 'The effect of information on product quality: Evidence from restaurant hygiene grade cards', *The Quarterly Journal of Economics* **118**(2), 409–451.

Jin, G. Z. and Leslie, P. (2009), 'Reputational incentives for restaurant hygiene', *American Economic Journal: Microeconomics* **1**(1), 237–67.

Klein, L. R. (1998), 'Evaluating the potential of interactive media through a new lens: Search versus experience goods', *Journal of business research* **41**(3), 195–203.

Klein, T. J., Lambertz, C. and Stahl, K. O. (2016), 'Market transparency, adverse selection, and moral hazard', *Journal of Political Economy* **124**(6), 1677–1713.

Kugler, M. and Verhoogen, E. (2012), 'Prices, plant size, and product quality', *The Review of Economic Studies* **79**(1), 307–339.

Leland, H. E. (1979), 'Quacks, lemons, and licensing: A theory of minimum quality standards', *Journal of political economy* **87**(6), 1328–1346.

Lewis, G. and Zervas, G. (2019), 'The supply and demand effects of review platforms', *Available at SSRN 3468278* .

Liu, M., Brynjolfsson, E. and Dowlatabadi, J. (2021), 'Do digital platforms reduce moral hazard? the case of uber and taxis', *Management Science* .

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

McCall, J. J. (1970), 'Economics of information and job search', *The Quarterly Journal of Economics* pp. 113–126.

Melitz, M. J. (2003), 'The impact of trade on intra-industry reallocations and aggregate industry productivity', *econometrica* **71**(6), 1695–1725.

Nelson, P. (1970), 'Information and consumer behavior', *Journal of political economy* **78**(2), 311–329.

Proserpio, D. and Zervas, G. (2017), 'Online reputation management: Estimating the impact of management responses on consumer reviews', *Marketing Science* **36**(5), 645–665.

Quinn, M., Godinho de Matos, M. and Peukert, C. (2021), 'The economic effects of mobile internet access – evidence from roam-like-at-home', *Unpublished manuscript* .

Reimers, I. and Waldfogel, J. (2021), 'Digitization and pre-purchase information: the causal and welfare impacts of reviews and crowd ratings', *American Economic Review* **111**(6), 1944–71.

Resnick, P., Zeckhauser, R., Swanson, J. and Lockwood, K. (2006), 'The value of reputation on ebay: A controlled experiment', *Experimental economics* **9**(2), 79–101.

Riordan, M. H. (1986), 'Monopolistic competition with experience goods', *The Quarterly Journal of Economics* **101**(2), 265–279.

Shin, M., Shin, J., Ghili, S. and Kim, J. (2023), 'The impact of the gig economy on product quality through the labor market: Evidence from ridesharing and restaurant quality', *Management Science* **69**(5), 2620–2638.

Syverson, C. (2004), 'Market structure and productivity: A concrete example', *Journal of Political Economy* **112**(6), 1181–1222.

Tadelis, S. (2016), 'Reputation and feedback systems in online platform markets', *Annual Review of Economics* **8**, 321–340.

Vatter, B. (2021), Quality disclosure and regulation: Scoring design in medicare advantage, Technical report, Working Paper.

Verhoogen, E. A. (2008), 'Trade, quality upgrading, and wage inequality in the mexican manufacturing sector', *The Quarterly Journal of Economics* **123**(2), 489–530.

Wang, Y. and Chaudhry, A. (2018), 'When and how managers' responses to online reviews affect subsequent reviews', *Journal of Marketing Research* **55**(2), 163–177.

Wolinsky, A. (1983), 'Prices as signals of product quality', *The review of economic studies* **50**(4), 647–658.

Appendix

Figures

Figure A1: International travelers to Italy (Δ from previous year, Thousand)

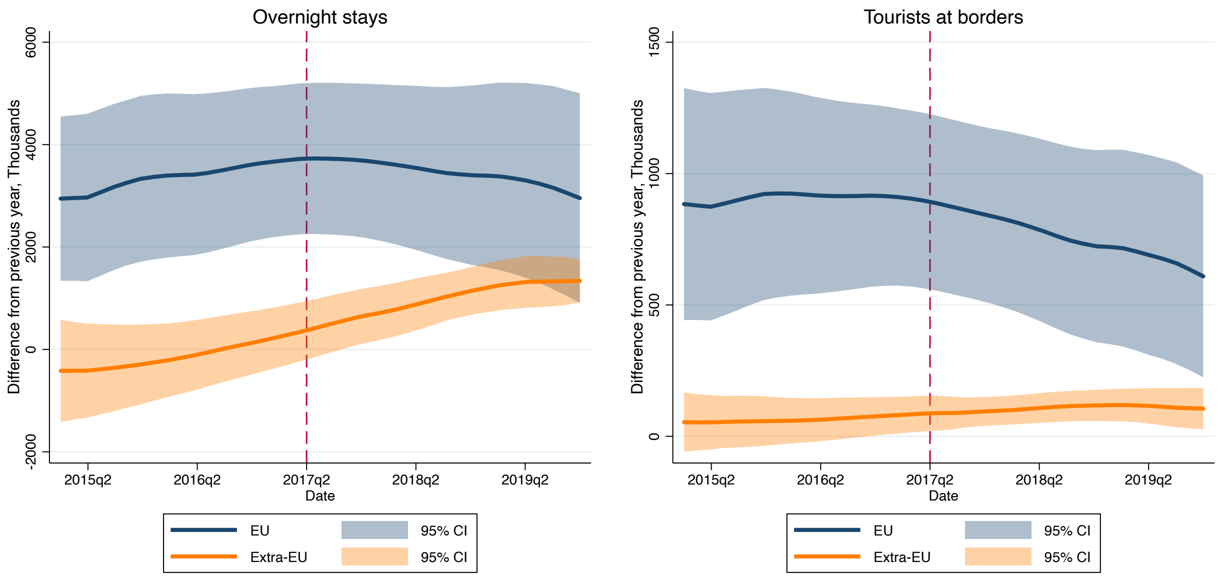

The data report international tourists traveling to Italy from different regions/countries of the world, in Thousand. The lines depict local polynomial fits of quarterly observations reporting the difference from the quarter of the previous year. EU refer to tourists from a EU country, and Extra-EU refer to all other tourists. Source: Bank of Italy.

Figure A2: Distribution of Tripadvisor average rating

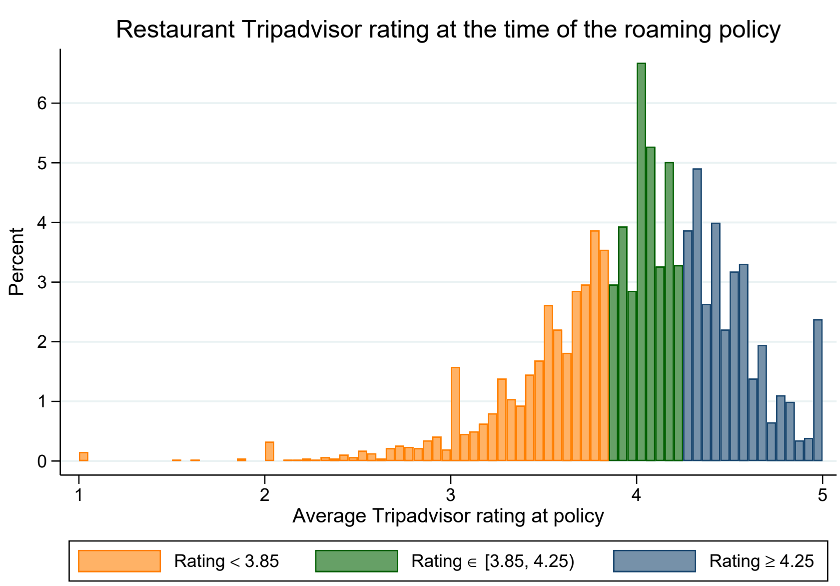

The figure shows the histogram of Tripadvisor average ratings of restaurants at the time of the policy, for the 4,628 matched restaurants with available information. Different colors split the overall sample in subgroups based on tertiles of ratings. In the sub-sample with available revenue data (N=2,043), the rating tertiles are 3.80 and 4.20.

Tables

Table A1: Comparison of main statistics across samples

	Tripadvisor sample	Matched sample	INPS sample
	Mean/(SD)	Mean/(SD)	Mean/(SD)
N. of tourist attractions in the ZIP code	3.01	3.10	2.59
	(6.65)	(6.59)	(6.13)
Probability of exposure to tourists	0.04	0.04	
	(0.12)	(0.10)	
Distance (km) from closest attraction	9.02	8.19	
	(11.14)	(10.71)	
Average 5-month Tripadvisor rating	3.97	4.00	
	(0.65)	(0.56)	
Average N. of 5-month replies to reviews	1.79	2.38	
	(8.98)	(9.58)	
Total Tripadvisor reviews	134.62	174.46	
	(320.06)	(297.80)	
Price €	0.25	0.25	
	(0.43)	(0.43)	
Price €€– €€€	0.71	0.71	
	(0.45)	(0.45)	
Price €€€€	0.04	0.04	
	(0.20)	(0.19)	
Average N. of monthly employees		5.24	4.31
		(5.02)	(4.60)
1 if firm exits market in Jan2015-Dec2019		0.21	0.29
		(0.41)	(0.45)
1 if firm exits after policy (Jun2017-Dec2019)		0.13	0.15
		(0.34)	(0.36)
1 if firm enters market in Jan2015-Dec2019		0.41	0.44
		(0.49)	(0.50)
1 if firm enters after policy (Jun2017-Dec2019)		0.15	0.17
		(0.36)	(0.38)
1 if firm hires workers w/ experience in restaurants at least once in Jan2015-Dec2019		0.76	0.69
		(0.43)	(0.46)
Average months of experience of newly-hired employees		13.64	12.72
		(14.55)	(14.45)
Average daily salaries (€)		65.98	65.88
		(10.48)	(11.41)
Observations	14146	5472	10391

Each observation is a restaurant. Data refer to the period between Jan 2015 - Dec 2019, unless otherwise specified. Data on Tripadvisor reviews, rating and replies refer to the period between Jan 2015 and Dec 2018. The matches sample is used in the market-level analysis.

Table A2: The roaming policy and the use of Tripadvisor by nationality of the reviewer

	Ratio Mobile/PC monthly reviews			Total monthly reviews from Mobile devices		
	(1)	(2)	(3)	(4)	(5)	(6)
EU	0.33***			-2301.21***		
	(0.08)			(111.86)		
EU*Post	0.47***			496.37***		
	(0.12)			(155.62)		
Extra-EU		0.18***			-2916.72***	
		(0.06)			(141.85)	
Extra-EU*Post		-0.12			129.76	
		(0.09)			(197.34)	
IT			0.99***			5617.90***
			(0.11)			(300.21)
IT*Post			0.02			-714.22*
			(0.15)			(417.65)
Month*Year FE	✓	✓	✓	✓	✓	✓
Observations	120	120	120	120	120	120
Adj. R-squared	0.840	0.824	0.764	0.923	0.908	0.880
Mean EU pre-policy	1.14			3162.97		
Mean Extra-EU pre-policy		1.00			2547.45	
Mean IT pre-policy			1.80			11082.07

Post takes value 1 after May 2017. Each observation is a region of origin-month-year. The regions of origin are EU, Extra-EU, IT and locals, which is the comparison (omitted) category in every column. The panel includes observations between 2015 and 2019. Standalone Post is absorbed by the Month*Year FE. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A3: The roaming policy and international travelers to Italy

	Δ from previous year, Thousdands					
	Overnight stays			Tourists at borders		
	(1)	(2)	(3)	(4)	(5)	(6)
EU	506.212***	530.153***		126.723***	133.083***	
	(83.752)	(118.615)		(17.343)	(24.561)	
Post	154.244*	172.200*	310.167	0.800	5.570	55.051
	(81.093)	(102.724)	(204.811)	(16.793)	(21.270)	(41.994)
EU*Post		-47.883	-47.883		-12.720	-12.720
		(167.747)	(161.642)		(34.734)	(33.143)
Origin FE			✓			✓
Year and quarter FE			✓			✓
Observations	320	320	320	320	320	320
Adj. R-squared	0.107	0.104	0.168	0.139	0.136	0.214

The data contain the number of international tourists traveling to Italy from different regions/countries of the world, in Thousdands. Each observation is a region of origin-year-quarter. The panel includes observations between 2015 and 2019. EU takes value 1 when the region of origin of the tourists is a EU country, and 0 otherwise. Post takes value 1 after 2017q2. Source: Bank of Italy. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A4: Summary statistics of independent and control variables

	Firms	Mean	SD	Min	Median	Max
Probability of exposure to tourists ($\times 100$)	4628	4.21	10.27	0.0	0.17	100.0
1 if tourist restaurant	4628	0.49	0.50	0.0	0.00	1.0
Tripadvisor rating at policy	4628	4.01	0.51	1.0	4.05	5.0
1 if low-rating ($\in [1, 3.85)$)	4628	0.33	0.47	0.0	0.00	1.0
1 if mid-rating ($\in [3.85, 4.25)$)	4628	0.34	0.47	0.0	0.00	1.0
1 if high-rating ($\in [4.25, 5]$)	4628	0.33	0.47	0.0	0.00	1.0
1 if restaurant is LLC	4628	0.56	0.50	0.0	1.00	1.0
1 if sole proprietorship	4628	0.23	0.42	0.0	0.00	1.0
1 if dine-in restaurant/bar	4628	0.86	0.34	0.0	1.00	1.0
1 if food truck	4628	0.04	0.19	0.0	0.00	1.0
1 if take-away only	4628	0.07	0.25	0.0	0.00	1.0
Distance (km) from closest attraction	4628	8.10	10.71	0.0	3.00	55.0
1 if distance to Rome city center <6 km	4628	0.49	0.50	0.0	0.00	1.0
1 if distance to Rome city center 6-15 km	4628	0.20	0.40	0.0	0.00	1.0
1 if distance to Rome city center >15 km	4628	0.32	0.47	0.0	0.00	1.0
1 if price is €	4576	0.26	0.44	0.0	0.00	1.0
1 if price is €€– €€€	4576	0.70	0.46	0.0	1.00	1.0
1 if price is €€€€	4576	0.04	0.19	0.0	0.00	1.0
1 if cuisine is Italian	4628	0.76	0.43	0.0	1.00	1.0
1 if no other restaurant in 400 m radius	4628	0.05	0.23	0.0	0.00	1.0
1 if 1-10 restaurants in 400 m radius	4628	0.25	0.43	0.0	0.00	1.0
1 if 11-30 restaurants in 400 m radius	4628	0.19	0.39	0.0	0.00	1.0
1 if more than 30 restaurants in 400 m radius	4628	0.51	0.50	0.0	1.00	1.0
1 if closest attraction has <1,000 reviews	4628	0.37	0.48	0.0	0.00	1.0
1 if closest attraction has 1,000-5,000 reviews	4628	0.36	0.48	0.0	0.00	1.0
1 if closest attraction has >5,000 reviews	4628	0.26	0.44	0.0	0.00	1.0

Each observation is a restaurant.

Table A5: Impact on restaurant profit margin

	Y=Annual profit margin; years 2015-2018				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	-0.013***	-0.013***	-0.017***	-0.031***	0.025**
	(0.004)	(0.004)	(0.006)	(0.009)	(0.010)
Restaurant & Year FE	✓	✓	✓	✓	✓
ZIP-code*Year	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	6614	6591	2291	2283	2017
Restaurants	2026	2018	693	693	632
Clusters	57	56	39	40	41
Adj. R-squared	0.345	0.349	0.326	0.352	0.363
DDD *p-value*				0.061	0.096

Post=1 if date is after 2016. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.80), [3.80, 4.20), [4.20, 5]$ respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-year. The sample includes observations between 2015 and 2018. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A6: Impact on total working days

	Y=log(monthly working days); Jan 2015 - Dec 2019				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	0.087***	0.081**	-0.066*	0.188***	0.100
	(0.031)	(0.035)	(0.039)	(0.055)	(0.098)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	219835	217622	72133	76920	68569
Restaurants	4628	4576	1490	1571	1515
Clusters	86	86	59	71	71
Adj. R-squared	0.714	0.713	0.701	0.728	0.700
Mean Y pre-policy	91.0	91.7	113.8	95.4	63.0
DDD *p-value*				0.034	0.104

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A7: Impact on working days per worker

	Y=N. of working days per worker; Jan 2015 - Dec 2019				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	0.085	0.008	0.038	-0.466***	0.371**
	(0.094)	(0.092)	(0.133)	(0.142)	(0.159)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	196749	195348	64454	70021	60873
Restaurants	4517	4471	1454	1537	1480
Clusters	86	86	59	71	70
Adj. R-squared	0.727	0.727	0.741	0.742	0.695
Mean Y pre-policy	15.8	15.8	16.0	16.2	15.2
DDD *p-value*				0.263	0.294

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A8: Impact on restaurant employment by price category

	Y=log(monthly employees); Jan 2015 - Dec 2019								
	Low price (€)			Medium price (€€– €€€)			High price (€€€€)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Low rating	Mid rating	High rating	Low rating	Mid rating	High rating	Low rating	Mid rating	High rating
Tourist*Post	-0.057**	0.088**	-0.025	-0.008	0.117***	0.103**	-0.361***	0.113	-0.015
	(0.027)	(0.034)	(0.032)	(0.026)	(0.036)	(0.046)	(0.074)	(0.071)	(0.110)
Restaurant & Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	16919	17819	22560	54187	56170	42228	1027	2931	3781
Restaurants	337	365	493	1130	1149	935	23	57	87
Clusters	29	33	45	55	64	66	6	12	17
Adj. R-squared	0.762	0.781	0.769	0.754	0.792	0.756	0.801	0.806	0.791
Mean Y pre-policy	5.4	4.7	3.2	7.4	5.8	4.2	8.5	9.6	7.3

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A9: Logit estimates: impact on hiring decisions

Y=1 if firm hires worker	*with* previous experience in restaurants				*without* experience in restaurants		
	(1) Full sample	(2) Low rating	(3) Mid rating	(4) High rating	(5) Low rating	(6) Mid rating	(7) High rating
Tourist*Post	0.175***	0.150**	0.235***	0.032	-0.028	0.227**	0.160
	(0.048)	(0.074)	(0.087)	(0.095)	(0.084)	(0.093)	(0.100)
Restaurant FE	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
Observations	176898	61956	62097	52845	61361	63352	55925
Restaurants	3568	1229	1221	1118	1178	1221	1157
Pseudo R-squared	0.014	0.017	0.015	0.018	0.013	0.014	0.012
Mean Y pre-policy	0.12	0.14	0.10	0.11	0.14	0.08	0.07
DDD *p-value*			0.376	0.013		0.409	0.378

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (3-4) and (6-7) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A10: Impact on restaurant (log) daily salaries

	Y=log(average daily salary (€)); Jan 2015 - Dec 2019				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	-0.000	0.000	0.016**	-0.001	-0.014**
	(0.003)	(0.003)	(0.006)	(0.006)	(0.005)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	200402	199026	67507	70593	60926
Restaurants	4558	4512	1492	1538	1482
Clusters	86	86	59	71	70
Adj. R-squared	0.495	0.496	0.511	0.491	0.482
Mean Y pre-policy	64.9	64.9	66.0	65.0	63.5
DDD *p-value*				0.477	0.102

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A11: Impact on hiring from Tripadvisor restaurants

Y=1 if firm hires worker from	Tripadvisor restaurants with any rating				Tripadvisor restaurants with mid/high rating			
	(1) Full sample	(2) Low rating	(3) Mid rating	(4) High rating	(5) Full sample	(6) Low rating	(7) Mid rating	(8) High rating
Tourist*Post	0.004***	0.004	0.006**	-0.001	0.006***	0.008***	0.005**	0.001
	(0.002)	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)	(0.001)
Restaurant & Time FE	✓	✓	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓	✓
Observations	217622	72133	76920	68569	217622	72133	76920	68569
Restaurants	4576	1490	1571	1515	4576	1490	1571	1515
Clusters	86	59	71	71	86	59	71	71
Adj. R-squared	0.086	0.102	0.071	0.082	0.067	0.076	0.054	0.070
Mean Y pre-policy	0.04	0.05	0.04	0.04	0.03	0.03	0.03	0.03
DDD *p-value*			0.142	0.184			0.649	0.004

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (3-4) and (7-8) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A12: Correlation between Tripadvisor rating and restaurant hiring decisions

	Y=Average 5-month Tripadvisor rating; Jan 2012 - Dec 2018					
	(1)	(2)	(3)	(4)	(5)	(6)
Hire worker w/ experience in restaurants	0.0116***	0.0116***	0.0112***	0.0054*	0.0097***	
	(0.002)	(0.002)	(0.002)	(0.003)	(0.002)	
Hire worker w/o experience in restaurants		-0.0002	0.0001	0.0001	-0.0013	
		(0.003)	(0.003)	(0.003)	(0.003)	
Hire worker from higher-rating restaurant				0.0162***		
				(0.003)		
log(monthly employees)					0.0061**	
					(0.003)	
Years of experience in restaurants						0.0043***
						(0.001)
Restaurant & Time FE	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓
Controls			✓	✓	✓	✓
Observations	223494	223494	222405	222405	222405	30015
Restaurants	5147	5147	5089	5089	5089	3737
Clusters	89	89	89	89	89	76
Adj. R-squared	0.480	0.480	0.482	0.482	0.482	0.557
Mean Y	3.95	3.95	3.95	3.95	3.95	3.91

Each observation is a restaurant-month-year. The sample includes observations between Jan 2012 and Dec 2018. Heteroskedasticity-robust standard errors clustered at municipality level. Controls include distance (km) to closest attraction and indicators for distance to Rome city center, restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with time trends. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A13: Impact on restaurant replies to Tripadvisor reviews

	Y=N. of 5-month replies to reviews; Jan 2015 - Dec 2018				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	0.156*	0.308***	0.008	1.013**	-0.444*
	(0.083)	(0.094)	(0.247)	(0.442)	(0.258)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	146713	145085	48937	52172	43976
Restaurants	4377	4328	1412	1499	1417
Clusters	86	86	59	70	70
Adj. R-squared	0.704	0.704	0.663	0.647	0.752
Mean Y pre-policy	2.56	2.59	1.58	2.51	3.90
DDD *p-value*				0.006	0.000

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A14: Impact on restaurant net purchases

	Y=log(annual net purchases); years 2015-2018				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	0.056***	0.060***	-0.034	0.115***	0.019
	(0.012)	(0.010)	(0.029)	(0.041)	(0.026)
Restaurant & Year FE	✓	✓	✓	✓	✓
ZIP-code*Year	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	6677	6652	2305	2299	2048
Restaurants	2043	2034	696	697	641
Clusters	57	56	39	40	41
Adj. R-squared	0.858	0.859	0.882	0.865	0.801
Mean Y pre-policy	255.9	256.8	369.1	228.9	154.7
DDD *p-value*				0.839	0.027

Post=1 if date is after 2016. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.80), [3.80, 4.20), [4.20, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-year. The sample includes observations between 2015 and 2018. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating.* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The End of Tourist Traps:
The Impact of Review Platforms on Quality Upgrading

Dante Donati

Online Appendix

1 Model

1.1 Model setup

There is a continuum of firms, each of which sells one quality $q \in \mathbb{R}_+$ of a vertically differentiated good at a common and exogenous price p to a continuum of consumers whose total mass is fixed and normalized to one.[68] All consumers have perfectly inelastic unit demand but are heterogeneous in their search costs $s \in \mathbb{R}_+$, with $s \sim Z$ (and density z). A consumer that buys one unit of quality q at price p gets utility (net of any search costs) $u = q - p$. I normalize the price p to one, so that utility becomes $u = q - 1$. There is no outside good in the market. Firms are also heterogeneous, differing in their underlying abilities (types), which affect their cost of producing a good of a certain quality. The total mass of firms L is endogenously determined through a zero-profit condition.

The timing is the following. At the beginning of the period, potential firms consider entering the market. If a firm decides to enter, it pays the sunk cost of entry $\kappa \in \mathbb{R}_+$ and learns its own ability parameter $\lambda \in \mathbb{R}_+$, which is drawn i.i.d. from a publicly known probability distribution with cdf Γ and pdf γ. Next, firms decide whether to stay in the market or not. Those firms that choose to stay then decide the quality level of their good and produce. Production requires a fixed cost of operation $C(q, \lambda)$, which depends positively on the chosen quality level and negatively on the exogenous ability parameter of the firm (i.e., $C'_q > 0$ and $C'_\lambda < 0$). This cost can be avoided if the firm chooses to stay out of the market.

1.2 Consumers' problem

Consumers have full information on the price of the goods being sold. However, they only know the quality distribution, F (with density f) and must engage in costly search to learn the quality provided by any particular firm. This is in line with the idea that information on the price of a meal might be gathered before the purchase, for instance, by reading the restaurants' menu on the window. Consumers' search is undirected and sequential: they visit stores one-by-one to learn their quality and after every visit compare the expected benefit and cost of continued search. If the expected quality gain from visiting another store is lower than the marginal cost of search s, the consumer continues to search; otherwise, they buy the product with the highest quality in hand. Following McCall (1970), in this context, the optimal stopping rule is characterized

[68] Assuming exogenous prices has two advantages: (1) it keeps the algebra tractable and allows me to solve the model analytically; (2) it excludes the possibility that prices are used by firms to signal quality (as in Wolinsky 1983) and therefore to reduce the asymmetric information problem. At the same time, however, the assumption may sound implausible. To make the model more realistic, one could think of a segmented market (e.g., fast-food vs. starred restaurants) where firms compete in quality and charge the same price within each segment (but not across segments), and consumers search exclusively within a segment. This scenario would not qualitatively change the results of the model.

by a reservation quality level $\rho(s)$ (i.e., the minimum acceptable quality of a good) such that a consumer stops searching and buys only if they find a product with quality $q \geq \rho(s)$. Particularly, $\rho(s)$ is implicitly defined by

$$h(\rho, q) \equiv \int_{\rho(s)}^{\infty} [q - \rho(s)]\, f(q)\, dq - s = 0, \tag{3}$$

where the integral is the expected quality gain from another search, accounting for the option value of discarding lower quality draws. Using integration by parts, one can rewrite (3) as

$$h(\rho, q) \equiv \int_{\rho(s)}^{\infty} [1 - F(q)]\, dq - s = 0. \tag{4}$$

Applying the implicit function theorem to (4) yields $\rho'(s) = -1/[1 - F[\rho(s)]]$, that is, the reservation quality is strictly decreasing in the search cost (i.e., consumers with lower s are pickier). This also implies that $\rho(s)$ is invertible and its inverse is given by $\rho^{-1}(r) = \int_r^{\infty} 1 - F(q)\, dq$.

1.3 Producers' problem

Firms do not know the ability parameters and the qualities produced by their rivals in the market, but they do know their distributions (Γ and F). Moreover, firms only know the distribution of search costs Z, and not the search cost of any individual consumer. Each firm takes as given these distributions, and determines its optimal quality based on the demand it faces, characterized by the reservation quality rule $\rho(s)$ implied by (3).

I analyze the optimization problem of a firm with ability parameter λ that chooses to stay in the market. To determine the quantity as a function of the quality chosen by the firm, $x(q)$, one should start from the optimal search rule. Only consumers with reservation qualities $\rho(s)$ below q will buy from the firm. Consider a consumer with reservation quality $r < q$. Since the quality distribution in the market is F and the total mass of operating firms is L, the mass of firms producing quality q above r is $L[1 - F(r)]$. The consumer is equally likely to buy from any one of these firms. That is, the probability that they will buy from a particular firm producing quality q is $1/L[1 - F(r)]$. Integrating over all consumers with a reservation quality lower than q yields the following formula for quantity:

$$x(q) = \int_0^q \frac{g(r)}{L[1 - F(r)]}\, dr, \tag{5}$$

where $g(r)$ is the pdf of the reservation quality. This formula can be expressed in terms of the search cost and quality distributions. Using equation (4), the corresponding cdf can be expressed as:

$$G(r) = 1 - Z[\rho^{-1}(r)] = 1 - Z\left[\int_r^{\infty} [1 - F(q)]\, dq\right]$$

Taking the derivative of $G(r)$ with respect to r yields:

$$G'(r) = g(r) = -z\left[\int_r^{\infty} [1 - F(q)]\, dq\right] [F(r) - 1] = z[\rho^{-1}(r)]\, [1 - F(r)]$$

Finally, replacing $g(r)$ into equation (5) yields the following standard residual demand curve:

$$x(q) = \frac{1}{L} \int_0^q z[\rho^{-1}(r)]\, dr. \tag{6}$$

Equation (6) states that the demand faced by a firm is determined by its own quality as well as its competitors' qualities. Note that demand is increasing in quality, since $x'(q) = \frac{1}{L}\, z[\rho^{-1}(q)] > 0$. However, quality is costly. Higher-quality output requires higher-quality inputs, which come at a cost (e.g., searching for better suppliers and hiring workers with better curricula/experience). I assume that these costs do not depend on the quantity produced, yet they depend on the innate ability of the firm, which is governed by the parameter λ.[69] Hence, the cost function of a firm with ability λ is $C(q, \lambda)$, with $C'_q > 0$, $C''_{qq} > 0$, $C'_\lambda < 0$ and $C''_{q\lambda} < 0$. The last conditions imply that more capable firms (higher λ) are more efficient, so that their fixed cost to produce a given quality is lower or, alternatively, they produce a higher quality product spending the same cost.[70] Hence, the optimization problem of a firm with ability λ choosing to stay in the market is

$$max_q \quad \pi[q(\lambda), \lambda] = x[q(\lambda)] - C[q(\lambda), \lambda]. \tag{7}$$

The equilibrium quality function $q(\cdot)$ follows from the first-order condition for an optimum

$$x'_q[q(\lambda)] - C'_q[q(\lambda), \lambda] = 0, \tag{8}$$

with the second-order condition for a maximum requiring that

$$x''_{qq}[q(\lambda)] - C''_{qq}[q(\lambda), \lambda] < 0. \tag{9}$$

1.4 Equilibrium

Let $q(\cdot)$ and $x(\cdot)$ be the quality and residual demand function in equilibrium, respectively. Then, the following properties follow (proofs are in Appendix Section 1.6):

Property 1: The equilibrium quality function is increasing in the ability parameter: $q'_\lambda(\lambda) > 0$, $\forall \lambda$.

Property 2: The demand function is increasing in the ability parameter: $x'_\lambda[q(\lambda)] > 0$, $\forall \lambda$.

Property 3: The profit function is increasing in the ability parameter: $\pi'_\lambda[q(\lambda), \lambda] > 0$, $\forall \lambda$.

From Property 3 it follows that the decision rule for staying in the market or leaving is characterized by a cut-off value $\underline{\lambda}$ such that firms stay in the market if and only if $\lambda \geq \underline{\lambda}$, with $\underline{\lambda}$ satisfying

$$\pi(\underline{\lambda}) = x[q(\underline{\lambda})] - C[q(\underline{\lambda}), \underline{\lambda})] = 0. \tag{10}$$

In the initial stage, potential entrants have to decide whether or not to start producing. Assuming unlimited entry into the market, firms keep entering until the expected value of post-entry profits equals the sunk entry

[69] Assuming quantity-dependent costs does not qualitatively alter the results, yet it makes the algebra more cumbersome.

[70] The positive relationship between managerial skills (human capital) and firm productivity has been well established in the literature (e.g., Gennaioli et al. 2013). In my context, high-skill owners/managers/chefs make a more efficient use of their inputs, and therefore save in costs. For instance, managerial skills help owners self-train their room staff, or cooking abilities allow firms to avoid expensive products and nevertheless make fabulous dishes.

cost. That is, the entry condition requires that

$$\kappa = \int_{\underline{\lambda}}^{\infty} \pi(\lambda)\,\gamma(\lambda)\,d\lambda = \int_{\underline{\lambda}}^{\infty} \Big[x[q(\lambda)] - C[q(\lambda),\lambda]\Big]\gamma(\lambda)\,d\lambda. \tag{11}$$

Finally, it is possible to express the distribution of qualities F in terms of the distribution of abilities Γ. Property 1 implies that qualities will be distributed with support $[\underline{q},\bar{q}]$, where $\underline{q} = q(\underline{\lambda})$ and $\bar{q} = q(\infty)$. Thus, for $v \in [\underline{q},\bar{q}]$, the cdf will be given by

$$F(v) = Pr\{q(\lambda) \leq v \,|\, \pi(\lambda) \geq 0\} = \frac{Pr\{\lambda \leq q^{-1}(v) \;\&\; \lambda \geq \underline{\lambda}\}}{Pr\{\lambda \geq \underline{\lambda}\}} = \frac{\Gamma[q^{-1}(v)] - \Gamma(\underline{\lambda})}{[1-\Gamma(\underline{\lambda})]}. \tag{12}$$

Note that $F(v) = 0$ for $v \leq \underline{q}$ and $F(v) = 1$ for $v \geq \bar{q}$. I can now define the equilibrium in this market.

Definition 1: A search equilibrium is a set $\{\rho : \mathbb{R}_+ \to \mathbb{R}_+,\, q : \mathbb{R}_+ \to \mathbb{R}_+,\, x : \mathbb{R}_+ \to \mathbb{R}_+,\, F : \mathbb{R}_+ \to [0,1],\, \underline{\lambda} > 0\}$ satisfying (4), (6), (8), (9), (10), (11) and (12).

1.5 Comparative statics

Assumption 1: *The search cost distribution is uniform on $[0,a]$ for $a > 0$.*

Assumption 2: *The firms' cost function takes the form*

$$C(q,\lambda) = \frac{q}{1-q}\,\frac{1}{\lambda}, \tag{13}$$

which satisfies the requirements described in section 1.3 for $q \in (0,1)$ and $\lambda > 0$.

From Assumption 1, it follows that for $q \in (\underline{q},1)$ the demand function (6) becomes

$$x(q) = \frac{1}{L}\int_0^q \frac{1}{a}\,\mathbb{1}_{\{\rho^{-1}(r)\in[0,a]\}}\,dr = \frac{1}{aL}\int_0^q \mathbb{1}_{\{r\in[\rho(a),\rho(0)]\}}\,dr = \frac{1}{aL}\,[q-\rho(a)]. \tag{14}$$

Note that $x'(q) = 1/aL > 0$ and $x''(q) = 0$ so that, together with Assumption 2, the second-order condition (9) holds. Substituting (13) and (14) into (8), the first-order condition simplifies to

$$q(\lambda;a) = 1 - \sqrt{\frac{aL}{\lambda}}, \tag{15}$$

and the additional condition in order for the ability parameter to yield admissible quality levels follows:

$$q(\lambda) \in (0,1) \iff \lambda > aL. \tag{16}$$

In other words, firms need to have at least some ability in order to produce positive qualities.

Consistently with Property 1, the equilibrium quality schedule (15) is increasing in λ, that is $q'(\lambda) = \sqrt{\frac{aL}{\lambda}}\,\frac{1}{2\lambda} > 0$. Moreover, the function is concave, that is $q''(\lambda) = -\frac{3}{4\lambda^2}\sqrt{\frac{aL}{\lambda}} < 0$. This is a direct consequence of the functional form of the firm's production cost (13). As $q \to 1$, more capable firms will use their ability-advantage mostly to save in costs rather than to deliver a higher quality product. Hence, the quality decision becomes less sensitive to the ability parameter as λ gets larger.

From the equilibrium quality schedule, it follows that the demand, cost and profit functions reduce to

$$x(\lambda; a) = \frac{1}{a\,L}\left[1 - \sqrt{\frac{aL}{\lambda}} - \rho(a)\right], \tag{17}$$

$$C(\lambda; a) = \frac{1}{\lambda}\left[\sqrt{\frac{\lambda}{aL}} - 1\right] \text{ and} \tag{18}$$

$$\pi(\lambda; a) = \frac{1}{a\,L}\left[1 - \sqrt{\frac{aL}{\lambda}} - \rho(a)\right] - \frac{1}{\lambda}\left[\sqrt{\frac{\lambda}{aL}} - 1\right], \tag{19}$$

and the operating threshold value for ability, $\underaccent{\bar}{\lambda}$, follows from (10). That is,

$$\underaccent{\bar}{\lambda} := \pi(\underaccent{\bar}{\lambda}) = 0 \iff \underaccent{\bar}{\lambda}(a) = \frac{aL}{\left(1 - \sqrt{\rho(a)}\right)^2}, \tag{20}$$

such that only for firms with $\lambda \geq \underaccent{\bar}{\lambda}$ it is convenient to stay in the market and produce. Note that $\underaccent{\bar}{\lambda}$ satisfies condition (16). Finally, the lower and upper limits of the support of the equilibrium quality distribution are $\underaccent{\bar}{q} = q(\underaccent{\bar}{\lambda}) = \sqrt{\rho(a)}$ and $\bar{q} = q(\infty) = 1$.

To conclude the comparative statics exercise, it remains to demonstrate how a decrease in search costs for consumers with *ex-ante* the highest costs – i.e., a reduction in a – affects the above quantities, and how these changes depend on the ability parameter of the firm. For this purpose, it is convenient to formalize two preliminary observations that will be used to derive the subsequent results (all proofs are in Appendix Section 1.6). First, I define the quantity $\delta(a) \equiv a\,L(a)$, where I emphasize the dependence of L on a.

Lemma 1: $\delta(a)$ is increasing in search costs, that is $\delta'_a(a) > 0$.

Note that $\delta(a)$ can be interpreted as the inverse of the per-firm density of consumers with a given level of search costs. Hence, Lemma 1 states that such a density is decreasing in the search costs. The second observation is about the profit function of a firm with ability λ, described in equation (19). It is possible to show that, if an increase in search costs reduces the profits of any currently operating firm, it must also reduce those of all firms with higher abilities. Formally,

Lemma 2: If there exists $\lambda_0 \geq \underaccent{\bar}{\lambda}(a)$ such that $\pi'_a(\lambda_0; a) \leq 0$, then $\pi'_a(\lambda; a) \leq 0 \ \forall \lambda > \lambda_0$.

I can now state the following key results:

Proposition 1: When search costs decrease, the quality $q(\cdot)$ produced by a firm with ability λ increases $\forall \lambda \geq \underaccent{\bar}{\lambda}$, and more so for firms with lower ability. That is, $q'_a(\cdot) < 0$ and $q''_{a\lambda}(\cdot) > 0$.

Proposition 2: When search costs decrease, the production costs $C(\cdot)$ of a firm with ability λ increase $\forall \lambda \geq \underaccent{\bar}{\lambda}$, and more so for firms with lower ability. That is, $C'_a(\cdot) < 0$ and $C''_{a\lambda}(\cdot) > 0$.

Proposition 3: When search costs decrease, the cut-off ability value $\underaccent{\bar}{\lambda}(\cdot)$ increases. That is, $\underaccent{\bar}{\lambda}'_a(\cdot) < 0$.

Corollary 1: A decrease in search costs causes the demand $x(\lambda; a)$ faced by all firms with sufficiently high ability to increase: for each a, there exists $\hat{\lambda}(a) \geq \underaccent{\bar}{\lambda}(a)$ such that $x'_a(\lambda; a) < 0 \ \forall \lambda > \hat{\lambda}(a)$.

1.6 Proofs

Proof of Property 1: Applying the Implicit Function Theorem to the first-order condition (8) yields:

$$\begin{aligned}\frac{\partial q(\lambda)}{\partial \lambda} &= -\frac{x''_{qq}[q(\lambda)]\, q'_\lambda(\lambda) - C''_{q\lambda}[q(\lambda),\lambda] - C''_{qq}[q(\lambda),\lambda]\, q'_\lambda(\lambda)}{x''_{qq}[q(\lambda)] - C''_{qq}[q(\lambda),\lambda]} \\ &= -q'_\lambda(\lambda) + \frac{C''_{q\lambda}[q(\lambda),\lambda]}{x''_{qq}[q(\lambda)] - C''_{qq}[q(\lambda),\lambda]}\end{aligned}$$

$$\iff q'_\lambda(\lambda) = \frac{1}{2}\,\frac{C''_{q\lambda}[q(\lambda),\lambda]}{x''_{qq}[q(\lambda)] - C''_{qq}[q(\lambda),\lambda]} > 0$$

The latter inequality holds because the numerator is negative by assumption, while the denominator is negative by the second-order condition (9). ■

Proof of Property 2:

$$\frac{\partial x[q(\lambda)]}{\partial \lambda} = x'_q[q(\lambda)]\; q'_\lambda(\lambda) > 0$$

The latter inequality holds because $x'_q > 0$ as the demand function is upward sloping in quality, and $q'_\lambda > 0$ by Property 1. ■

Proof of Property 3: Applying the Envelope Theorem to the profit function (7) yields:

$$\frac{\partial \pi[q(\lambda),\lambda]}{\partial \lambda} = -C'_\lambda(q,\lambda) > 0$$

The latter inequality holds because of the assumption on the cost function. ■

Proof of Lemma 1: Replacing the equilibrium profit schedule (19) into the entry condition (11) yields the following identity

$$\Theta(\rho, L; a) \equiv \int_{\underline{\lambda}(a)}^{\infty} \left[\frac{1}{\lambda} - 2\sqrt{\frac{1}{aL\lambda}} + \frac{1}{aL}(1-\rho)\right] \gamma(\lambda)\, d\lambda - \kappa = 0, \tag{21}$$

where both ρ and L are functions of a, and $\underline{\lambda}(a) = aL(1-\rho)^{-2}$. Implicitly differentiating the identity with respect to a yields

$$\Theta'_a + \Theta'_\rho\, \frac{\partial \rho(a)}{\partial a} + \Theta'_L\, \frac{\partial L(a)}{\partial a} = 0. \tag{22}$$

These partial derivatives are

$$\Theta'_a = \int_{\underline{\lambda}(a)}^{\infty} \frac{1}{a^2}\left[\sqrt{\frac{a}{L\lambda}} - \frac{(1-\rho)}{L}\right]\gamma(\lambda)\,d\lambda < 0;$$
$$\Theta'_\rho = \int_{\underline{\lambda}(a)}^{\infty} -\frac{1}{aL}\,\gamma(\lambda)\,d\lambda < 0;$$
$$\Theta'_L = \int_{\underline{\lambda}(a)}^{\infty} \frac{1}{L^2}\left[\sqrt{\frac{L}{a\lambda}} - \frac{(1-\rho)}{a}\right]\gamma(\lambda)\,d\lambda = \frac{a}{L}\,\Theta'_a < 0.$$

The latter equality together with equation (22) yields

$$\frac{\partial\rho}{\partial a} = -\frac{1}{\Theta'_\rho}\left[\Theta'_a + \frac{\partial L}{\partial a}\Theta'_L\right] = -\frac{\Theta'_a}{\Theta'_\rho}\left[1 + \frac{a}{L}\frac{\partial L}{\partial a}\right]. \tag{23}$$

Since $\rho'_a < 0$ from equation (4), $\Theta'_a < 0$ and $\Theta'_\rho < 0$, then (23) implies that

$$1 + \frac{a}{L}\frac{\partial L}{\partial a} > 0 \iff a\frac{\partial L}{\partial a} + L > 0 \implies \delta'_a > 0. \blacksquare$$

Proof of Lemma 2: Taking the derivative of the equilibrium profit function (19) with respect to a yields

$$\pi'_a(\lambda; a) = \frac{1}{\delta(a)}\left[\frac{\delta'_a(a)}{\sqrt{\delta(a)}}\frac{1}{\sqrt{\lambda}} - \frac{(1-\rho)}{\delta(a)} - \frac{\partial\rho(a)}{\partial a}\right].$$

Thus, the sign of $\pi'_a(\lambda; a)$ depends on the sign of the term in brackets. Since $\delta'_a > 0$ by Lemma 1, this term is decreasing in λ. This implies that, if the term is negative for $\lambda_0 \geq \underline{\lambda}$, then it will also be negative $\forall\,\lambda > \lambda_0$. $\blacksquare$

Proof of Proposition 1: Lemma 1 implies that

$$q'_a(\lambda; a) = -\left[\frac{\delta'_a(a)}{2\sqrt{\delta(a)\lambda}}\right] < 0. \tag{24}$$

Moreover, taking the derivative of (24) with respect to λ, yields

$$q''_{a\lambda}(\lambda; a) = \frac{1}{4}\frac{\delta'_a(a)}{\sqrt{\delta(a)}}\lambda^{-3/2} > 0.$$

That is, the negative change in quality predicted by (24) becomes smaller (i.e., closer to zero) for larger values of λ. $\blacksquare$

Proof of Proposition 2: Lemma 1 implies that

$$C'_a(\lambda; a) = -\frac{1}{2\sqrt{\lambda}}\delta(a)^{-3/2}\delta'_a(a) < 0. \tag{25}$$

Moreover, taking the derivative of (25) with respect to λ, yields

$$C''_{a\lambda}(\lambda; a) = \frac{1}{4}[\lambda\delta(a)]^{-3/2}\delta'_a(a) > 0.$$

That is, the negative change in costs predicted by (25) becomes smaller (i.e., closer to zero) for larger values of λ. $\blacksquare$

Proof of Proposition 3: Consider the entry condition

$$\int_{\underline{\lambda}(a)}^{\infty} \pi(\lambda; a)\, \gamma(\lambda)\, d\lambda = \kappa.$$

Differentiating this with respect to a, yields

$$\int_{\underline{\lambda}(a)}^{\infty} \pi'_a(\lambda; a)\, \gamma(\lambda)\, d\lambda = 0. \tag{26}$$

Together with Lemma 2, this implies that $\pi'_a[\underline{\lambda}(a); a] > 0$, otherwise the integrand in (26) would be negative $\forall\, \lambda > \lambda$ (by Lemma 2), which would contradict (26). To see how the ability threshold $\underline{\lambda}(a)$ changes as search costs decreases, consider a shift in a: a_1 to $a_2 < a_1$. Then

$$\pi[\underline{\lambda}(a_2), a_2] = 0 = \pi[\underline{\lambda}(a_1), a_1] > \pi[\underline{\lambda}(a_1), a_2],$$

where the two equalities follow from the definition of $\underline{\lambda}$, and the inequality follows from $\pi'_a[\underline{\lambda}(a); a] > 0$. Since (by Property 3) $\pi'_\lambda > 0\; \forall\, \lambda$, it follows that $\underline{\lambda}(a_2) > \underline{\lambda}(a_1)$. $\blacksquare$

Proof of Corollary 1: Taking the derivative of $x(\lambda; a)$ with respect to a, yields

$$x'_a(\lambda; a) = \frac{1}{\delta(a)} \left[\frac{\delta'_a\, (\rho - 1)}{\delta(a)} + \frac{\delta'_a}{2\sqrt{\lambda \delta(a)}} - \rho'_a \right]. \tag{27}$$

The sign of (27) equals the sign of the expression in brackets. In particular, the expression is negative for sufficiently high values of λ, that is

$$x'_a(\lambda; a) < 0 \iff \lambda > \frac{\delta(a)}{4\Big[1 - \rho + \frac{\delta(a)\rho'_a}{\delta'_a}\Big]^2}.$$

Hence, there exists a $\hat{\lambda} > \underline{\lambda}$ such that $x'_a(\hat{\lambda}; a) < 0$. Since from (27) it is clear that x'_a is decreasing in λ (as $\delta'_a > 0$ by Lemma 1), this implies that $x'_a(\lambda; a) < 0\; \forall \lambda > \hat{\lambda}$. $\blacksquare$

2 Measurement

2.1 Defining exposure to tourists

I assume that search is sequential over space and bounded by the structure of the road network around a tourist site. In particular, the probability of coming across a restaurant is equal to the probability of ending up on the street where the restaurant is located, taking into account the previous path. Hence, two or more restaurants located on the same street have the same probability of being found, but this probability depends on the path to their location. Tourists start inspecting high-visible places around them - such as those in front of a tourist attraction - and then move to other less visible places until the marginal expected cost (time and fatigue of walking) becomes larger than the marginal expected benefit of finding a good deal.

Figure B1: Example of partial road network

Segments represent roads, circles represent junctions

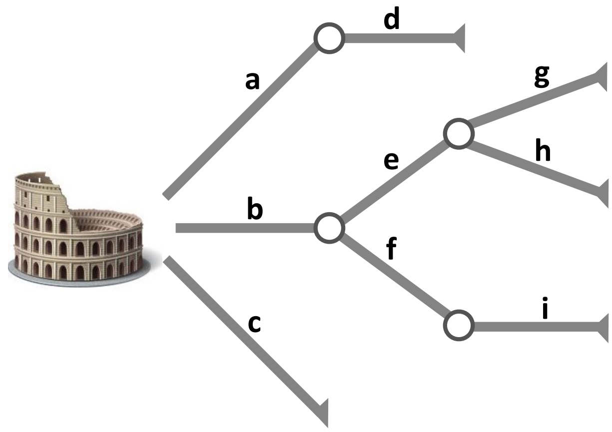

$p(a) = p(b) = p(c) = 1/3$
$p(e|b) = p(f|b) = p(g|e) = p(h|e) = 1/2 \Rightarrow p(e) = p(f) = 1/6$ and $p(g) = p(h) = 1/12$
$p(d|a) = p(i|f) = 1 \Rightarrow p(d) = 1/3$ and $p(i) = 1/6$

In practice, I use information from Google Maps and construct the partial road network that leads to the Tripadvisor restaurants around each attraction, and eventually compute the probabilities to find them while walking away from the attraction. The procedure works as follows. First, for each restaurant i, I consider its closest (shortest distance) tourist attraction t. Then, for each identified pair (t, i), I use the Google Maps API to find the directions (street names) of all the paths that lead from t to i on foot. In case more than one path is suggested by Google, I consider the shortest distance path to build the benchmark measure, while I also provide robustness to the use of all alternative paths. Then, I construct the partial road network around attraction t using the street names from Google Maps and, for each i, I compute the conditional probabilities of being in any of the roads that form the path from t to i (taking into account competing roads). Finally, I multiply them and compute the probability of finding i. In the process, I assume that all roads are equally weighted (random walk assumption). Figure B1 shows a simplified example of this calculation.

More formally, the probability that a tourist moving away from attraction t comes across restaurant i is equal to the joint probability of traveling the path defined by a vector of streets $(s_1, ..., s_{Ni})$ connecting t to i. Hence,

$$P(i) = P(s_1 \;\&\; s_2 \;\&\; ... \;\&\; s_{Ni}) = P(s_1|t) \prod_{j=2}^{N_i} P(s_j|s_{j-1}) \tag{28}$$

This probability measure reflects the chances that a restaurant is visited by a tourist and thereby the extent to which it is exposed to the policy. In particular, this measure not only reflects the "visibility" of restaurant i from attraction t, but also the effect of proximity of i to t. In fact, an increasing number of streets compete in the road network as the radius enlarges by moving farther away from the attraction, and this mechanically drives down the estimated probability. Hence, any observed differential impact of the policy along the probability measure could be in part explained by factors - other than the presence of tourists - that correlate with proximity and affect restaurants' decisions. For this reason, in the empirical analysis, I always control for the distance to the attraction and the distance to Rome city center. I use equation (28) to compute this probability for all restaurants in the Tripadvisor sample.[71] The empirical distribution of the probability measure is right-skewed. About 50% of the restaurants have roughly 0 or very low probability of being found by tourists. Specifically, in the Tripadvisor-INPS matched sample, the median value of the probability measure is 0.17%, and in the sub-sample with available revenue data is 0.35%.

2.2 Exposure to tourists, clientele and rating

Using the probability measure defined above, I examine how restaurants' type of clientele and ratings vary with the restaurant potential exposure to tourist demand. This exercise helps me to both validate the constructed measure as well as provide a description of the restaurant's industry. The left panel of Figure B2 exploits the origin of the reviewers to distinguish them between locals and tourists. I identify as foreign tourists all those reviewers writing in a language different than Italian. This means that the green line in the graph is probably under-reporting the share of total tourists, as Italian travelers are not accounted.[72] However, since the roaming policy did not affect Italians directly, excluding them from this group provides a more conservative picture of the potential effect of the policy across different levels of exposure to tourists. Particularly, for restaurants whose probability is below or equal to the median, the share of foreign clientele remains quite stable and below 10%. By contrast, this increases rapidly afterwards, and reaches almost 60% for restaurants at the top probability-decile.

Figure B2: Type of clientele and Tripadvisor rating by exposure to tourists

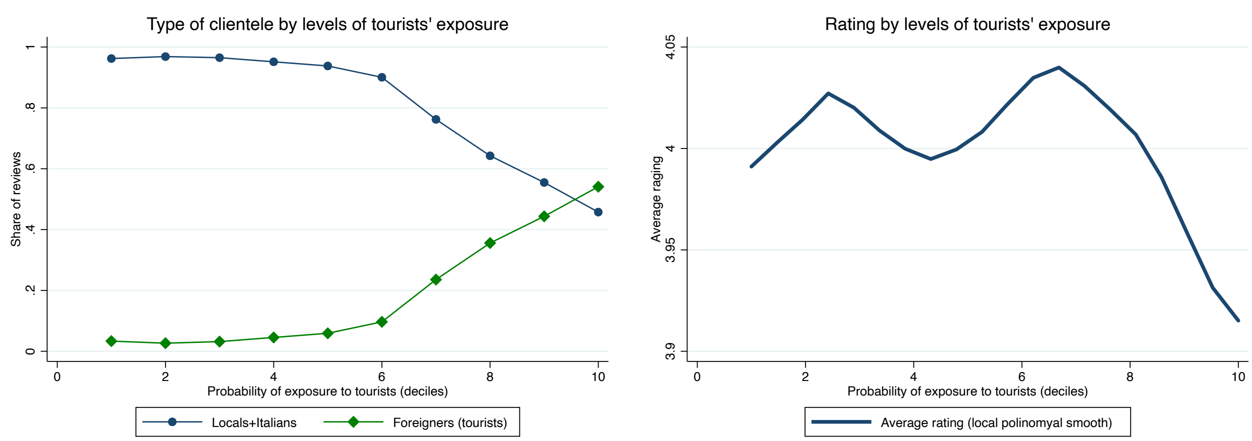

Data on 11,595 restaurants with at least one review as of May 2017. Reviews data refer to the pre-policy period Jan2015-May2017. Rating is computed at the time of the policy.

The right panel of Figure B2 sows how Tripadvisor rating varies across levels of exposure to tourists. For each decile of probability, it reports the mean of all restaurants' average rating at the time of the policy. Restaurants more exposed to tourist demand have, on average, poorer Tripadvisor ratings, in the order of 0.10-0.15 on a scale from 1 to 5. In line with the theory, restaurants that rely more on repeated and informed clientele sell higher quality meals.[73] The explanation is two-fold. First, locals are more likely to be

[71]Note that this measure could not be computed for the entire universe of restaurants in the Social Security database, as their exact location remains unknown to the researcher for confidentiality reasons.

[72]The main problem here is that, among Italians, I can identify the "locals" only for a subset of reviewers who explicitly indicate their town in their profile. The rest of Italians can either be tourists or locals, yet it is impossible to distinguish them.

[73]These results are in line with those of Dall'orso et al. (2016), who provide evidence on the existence of quality

informed. Restaurants located in areas where the share of informed consumers is higher (e.g. those with a lower probability measure), have incentive to provide a better quality product/service to stay in the market (Cooper and Ross, 1984). Second, locals exert control over quality through repeated purchases. Then, quality provision becomes a way to establish reputation in the market (Riordan, 1986).

2.3 The roaming policy and exposure to tourists

Figure B3 plots the share of Tripadvisor reviews from Europeans by different types of devices across deciles of probability, before and after the policy. The graph shows that, after the policy, the share of mobile (PC) reviews from Europeans increased (decreased). However, most of the change took place in restaurants with higher exposure to tourists, namely, those with a probability above the median, while virtually nothing happened for other restaurants. Because establishments with low probabilities have little chance of being affected by the policy, I consider them as a control group in the baseline specification.

Figure B3: Reviews from Europeans by types of devices and exposure to tourists

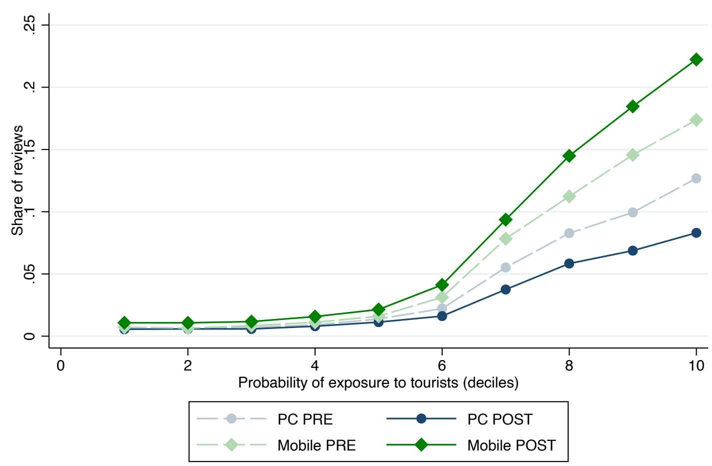

Data on 14,146 restaurants with at least one review as of December 2019. PRE refers to the period Jan 2015-May 2017. POST refers to the period June 2017-December 2019.

---

differential across more and less visible restaurants using data from Yelp on 10 large cities in Europe and North America. In particular, they show that restaurants with higher visibility from tourists - i.e. those located at street intersections - consistently exhibit a lower Yelp rating.

3 Online Appendix Figures

Figure C1: Permutation test for restaurant exit

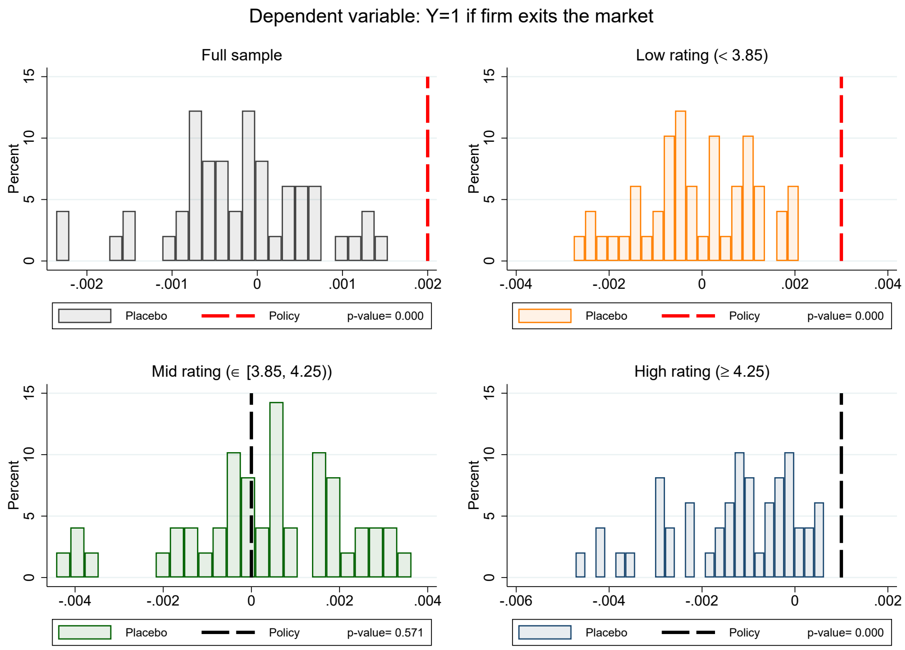

Notes: Each panel plots the distribution of coefficients on Tourist*Post-Month, where Month is between May 2012 and May 2016, estimated on a sample of observations between Jan 2012 and Dec 2016. Tripadvisor rating and the respective category are calculated in each month. The vertical dashed lines report the actual policy coefficients, as estimated in the main analysis. The line is red when the respective coefficient is significant at least at the 10% confidence level, and black otherwise.

Figure C2: Permutation test for industry composition

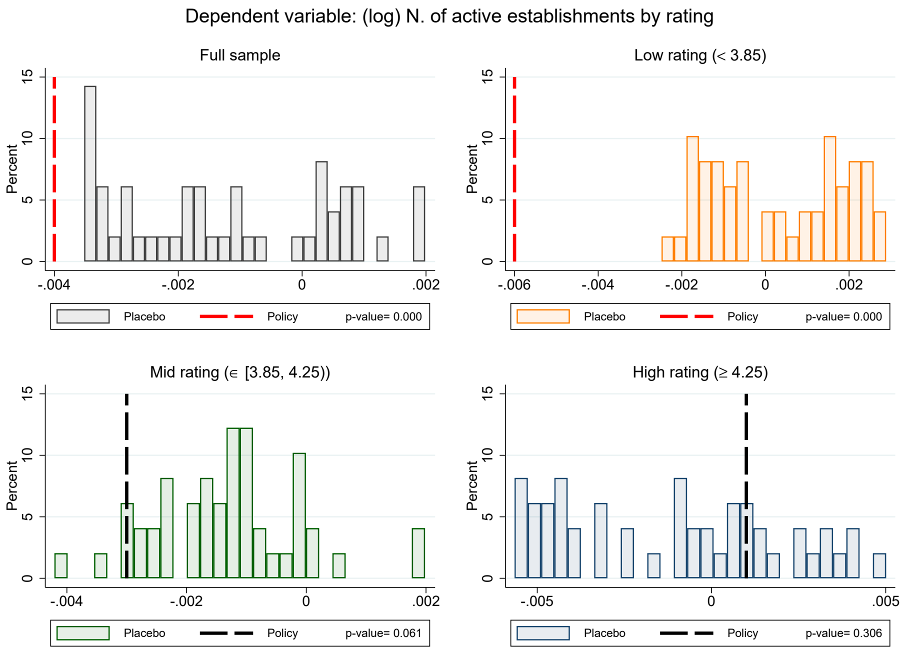

Notes: Each panel plots the distribution of coefficients on Tourist*Post-Month, where Month is between May 2012 and May 2016, estimated on a sample of observations between Jan 2012 and Dec 2016. Tripadvisor rating and the respective category are calculated in each month. The vertical dashed lines report the actual policy coefficients, as estimated in the main analysis. The line is red when the respective coefficient is significant at least at the 10% confidence level, and black otherwise.

Figure C3: Permutation test for restaurant hiring decisions (extensive margins)

Notes: Each panel plots the distribution of coefficients on Tourist*Post-Month, where Month is between May 2012 and May 2016, estimated on a sample of observations between Jan 2012 and Dec 2016. Tripadvisor rating and the respective category are calculated in each month. The vertical dashed lines report the actual policy coefficients, as estimated in the main analysis. The line is red when the respective coefficient is significant at least at the 10% confidence level, and black otherwise.

Figure C4: Permutation test for restaurant hiring decisions (intensive margins)

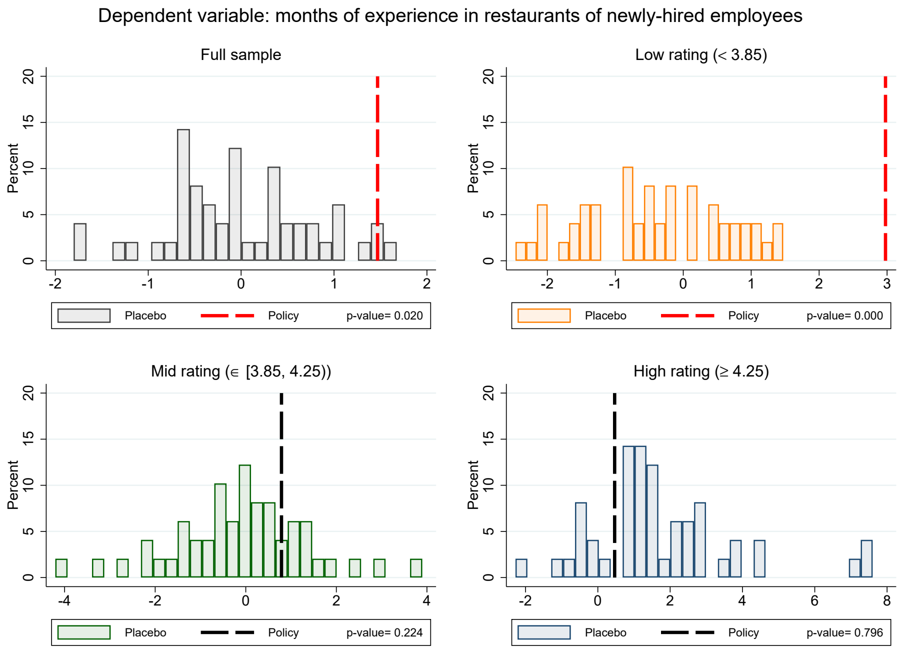

Notes: Each panel plots the distribution of coefficients on Tourist*Post-Month, where Month is between May 2012 and May 2016, estimated on a sample of observations between Jan 2012 and Dec 2016. Tripadvisor rating and the respective category are calculated in each month. The vertical dashed lines report the actual policy coefficients, as estimated in the main analysis. The line is red when the respective coefficient is significant at least at the 10% confidence level, and black otherwise.

Figure C5: Permutation test for restaurant daily salaries

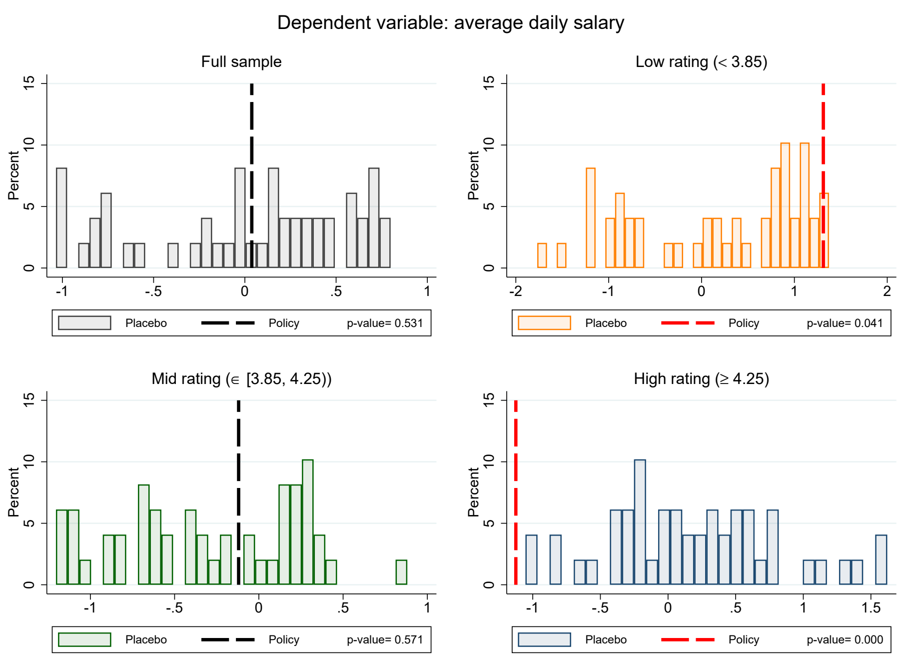

Notes: Each panel plots the distribution of coefficients on Tourist*Post-Month, where Month is between May 2012 and May 2016, estimated on a sample of observations between Jan 2012 and Dec 2016. Tripadvisor rating and the respective category are calculated in each month. The vertical dashed lines report the actual policy coefficients, as estimated in the main analysis. The line is red when the respective coefficient is significant at least at the 10% confidence level, and black otherwise.

Figure C6: Permutation test for restaurant Tripadvisor rating

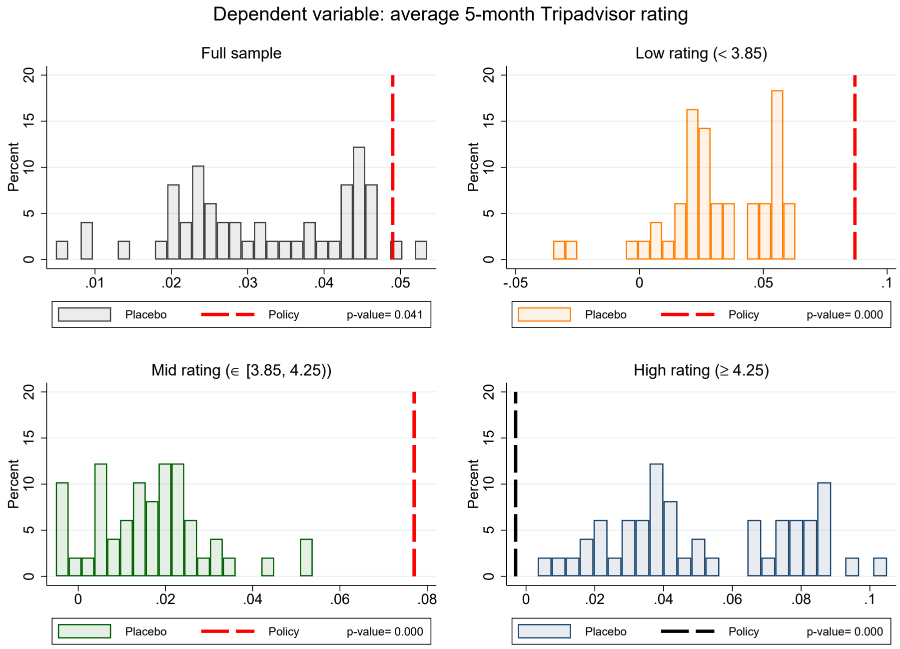

Notes: Each panel plots the distribution of coefficients on Tourist*Post-Month, where Month is between May 2012 and May 2016, estimated on a sample of observations between Jan 2012 and Dec 2016. Tripadvisor rating and the respective category are calculated in each month. The vertical dashed lines report the actual policy coefficients, as estimated in the main analysis. The line is red when the respective coefficient is significant at least at the 10% confidence level, and black otherwise.

Figure C7: Event-study estimates for restaurant revenues

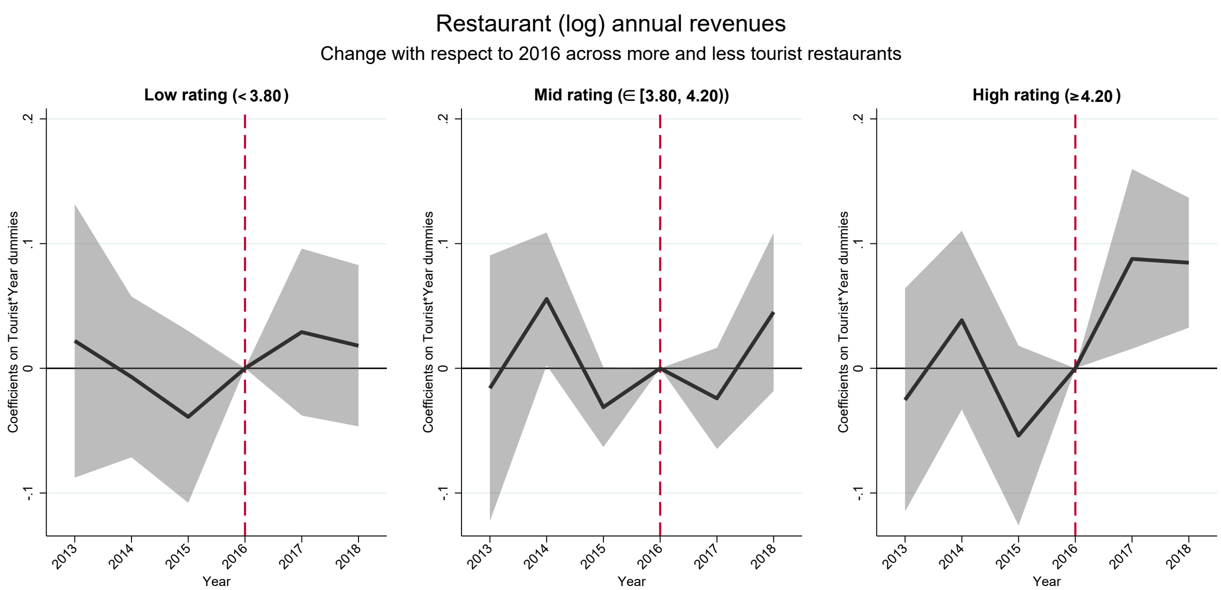

Notes: The graph reports estimated coefficients on the interactions of Tourist*Year dummies from three separate regressions where each observation is a restaurant-year. All controls and fixed-effects from the main analysis are included. The omitted year is 2016. Tourist restaurants are those with a measure of exposure above the median. The sample includes observations between 2013 and 2018. Shaded areas depict 95% confidence intervals.

Figure C8: Event-study estimates for restaurant exit

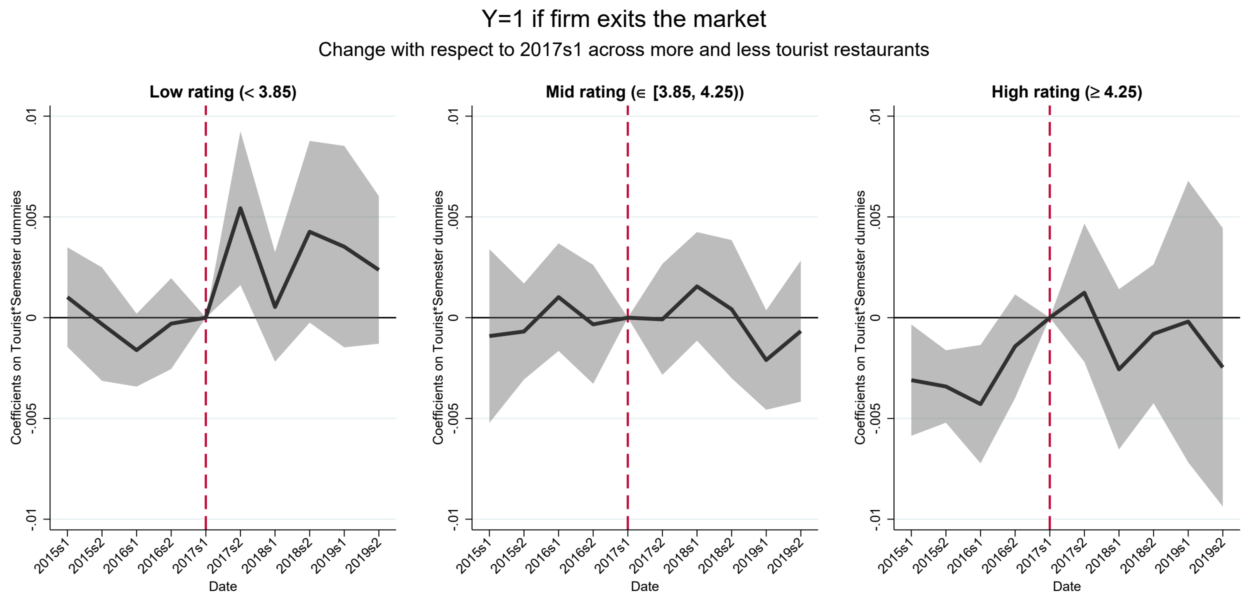

Notes: The graph reports estimated coefficients on the interactions of Tourist*Semester dummies from three separate regressions where each observation is a restaurant-month-year. All controls and fixed-effects from the main analysis are included. The omitted semester is 2017s1. Tourist restaurants are those with a measure of exposure above the median. The sample includes observations between Jan 2015 and Dec 2018. Shaded areas depict 95% confidence intervals.

Figure C9: Event-study estimates for hiring decisions (extensive margins)

Notes: The graph reports estimated coefficients on the interactions of Tourist*Semester dummies from three separate regressions where each observation is a restaurant-month-year. All controls and fixed-effects from the main analysis are included. The omitted semester is 2017s1. Tourist restaurants are those with a measure of exposure above the median. The sample includes observations between Jan 2015 and Dec 2018. Shaded areas depict 95% confidence intervals.

Figure C10: Event-study estimates for hiring decisions (intensive margins)

Notes: The graph reports estimated coefficients on the interactions of Tourist*Semester dummies from three separate regressions where each observation is a restaurant-month-year. All controls and fixed-effects from the main analysis are included. The omitted semester is 2017s1. Tourist restaurants are those with a measure of exposure above the median. The sample includes observations between Jan 2015 and Dec 2018. Shaded areas depict 95% confidence intervals.

Figure C11: Impact on restaurant revenues across quintiles of exposure

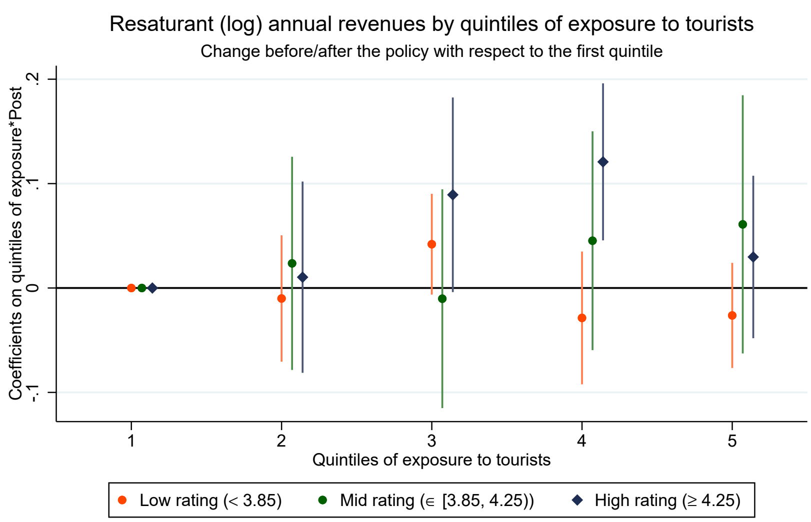

Notes: The graph reports estimates on the interactions of quintiles of exposure*Post from three separate regressions (low, mid, high-rating restaurants), where each observation is a restaurant-year. The first quintile is omitted. All controls and fixed-effects from the main analysis are included. Post takes value 1 after 2016. The sample includes observations between 2015 and 2018. Heteroskedasticity-robust standard errors are clustered at municipality level. Bars depict 95% confidence intervals.

Figure C12: Impact on restaurant exit across quintiles of exposure

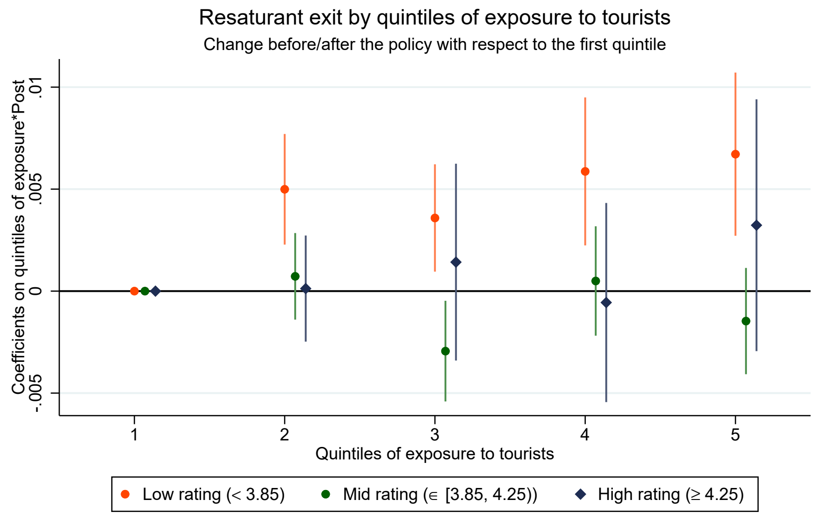

Notes: The graph reports estimates on the interactions of quintiles of exposure*Post from three separate regressions (low, mid, high-rating restaurants), where each observation is a restaurant-month-year. The first quintile is omitted. All controls and fixed-effects from the main analysis are included. Post takes value 1 after May 2017. The sample includes observations between Jan 2015 and Dec 2019. Heteroskedasticity-robust standard errors are clustered at municipality level. Bars depict 95% confidence intervals.

Figure C13: Impact on hiring decisions (extensive margins) across quintiles of exposure

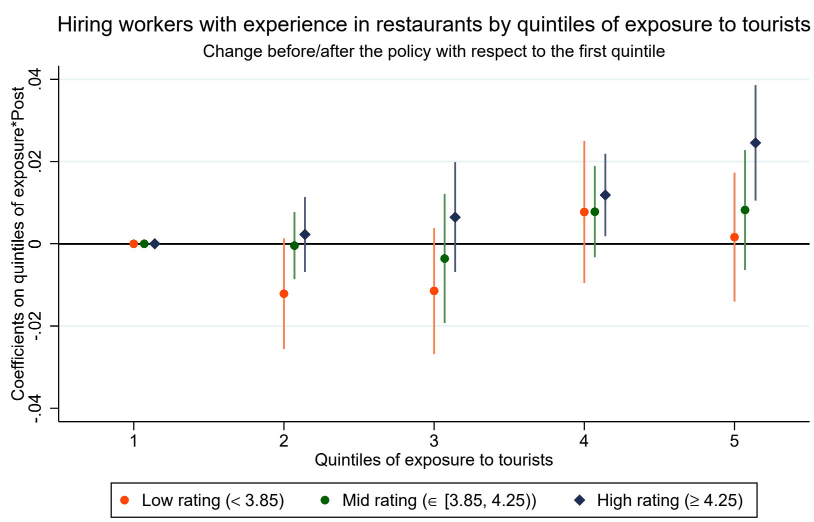

Notes: The graph reports estimates on the interactions of quintiles of exposure*Post from three separate regressions (low, mid, high-rating restaurants), where each observation is a restaurant-month-year. The first quintile is omitted. All controls and fixed-effects from the main analysis are included. Post takes value 1 after May 2017. The sample includes observations between Jan 2015 and Dec 2019. Heteroskedasticity-robust standard errors are clustered at municipality level. Bars depict 95% confidence intervals.

Figure C14: Impact on hiring decisions (intensive margins) across quintiles of exposure

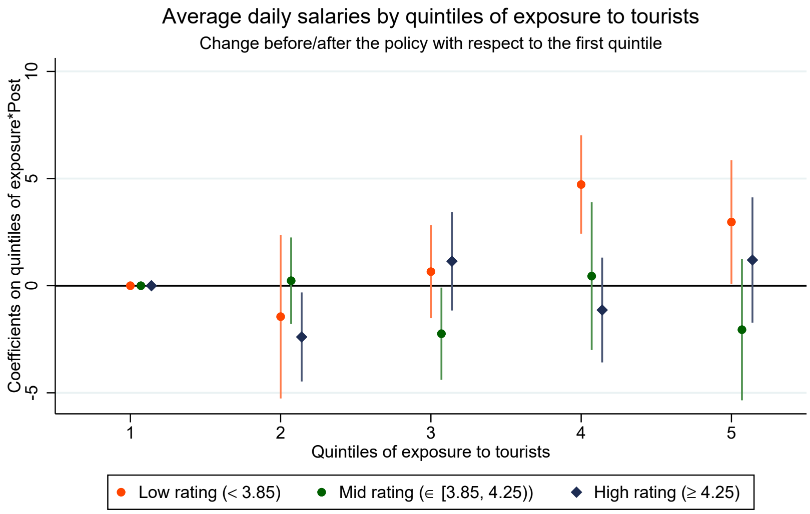

Notes: The graph reports estimates on the interactions of quintiles of exposure*Post from three separate regressions (low, mid, high-rating restaurants), where each observation is a restaurant-month-year. The first quintile is omitted. All controls and fixed-effects from the main analysis are included. Post takes value 1 after May 2017. The sample includes observations between Jan 2015 and Dec 2019. Heteroskedasticity-robust standard errors are clustered at municipality level. Bars depict 95% confidence intervals.

Figure C15: Impact on salaries across quintiles of exposure

Notes: The graph reports estimates on the interactions of quintiles of exposure*Post from three separate regressions (low, mid, high-rating restaurants), where each observation is a restaurant-month-year. The first quintile is omitted. All controls and fixed-effects from the main analysis are included. Post takes value 1 after May 2017. The sample includes observations between Jan 2015 and Dec 2019. Heteroskedasticity-robust standard errors are clustered at municipality level. Bars depict 95% confidence intervals.

Figure C16: Impact on Tripadvisor rating across quintiles of exposure

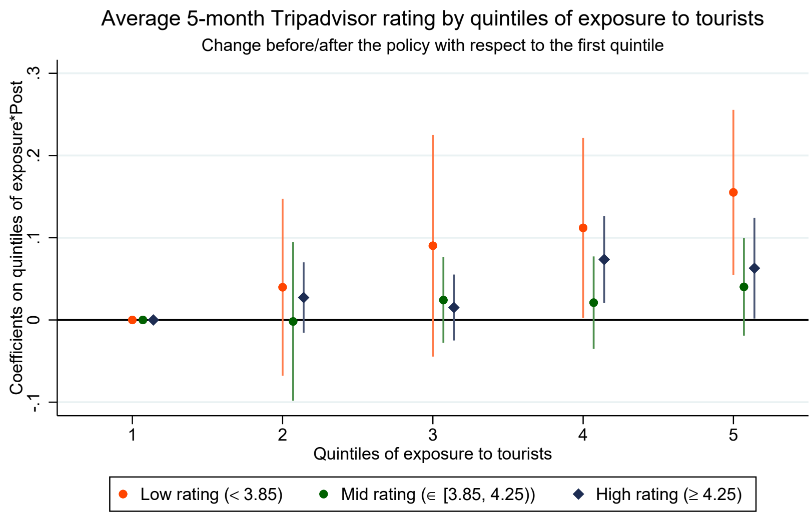

Notes: The graph reports estimates on the interactions of quintiles of exposure*Post from three separate regressions (low, mid, high-rating restaurants), where each observation is a restaurant-month-year. The first quintile is omitted. All controls and fixed-effects from the main analysis are included. Post takes value 1 after May 2017. The sample includes observations between Jan 2015 and Dec 2019. Heteroskedasticity-robust standard errors are clustered at municipality level. Bars depict 95% confidence intervals.

Figure C17: Impact of the roaming policy on de-trended monthly employment

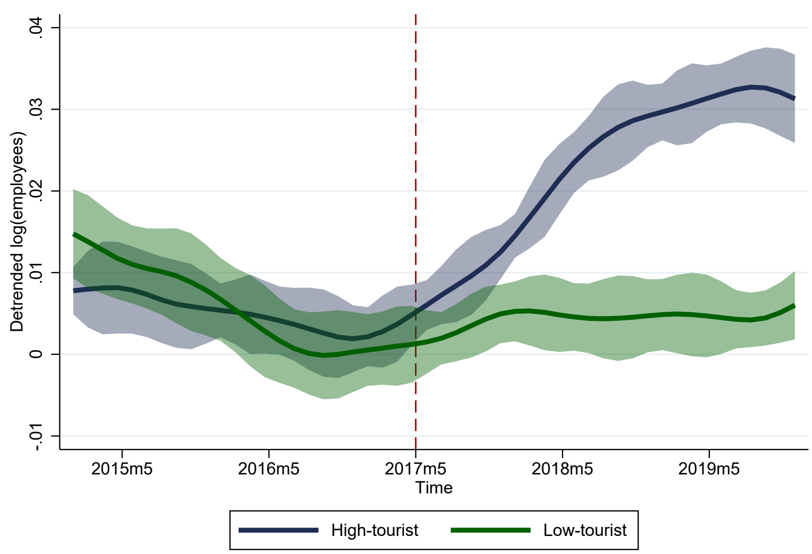

Notes: High-tourist restaurants are those for which the binary variable *Tourist*=1, while *Tourist*=0 for low-tourist restaurants.

4 Online Appendix Tables

4.1 Robustness of results: exposure to tourist including alternative routes

Table D1: Impact on restaurant revenues

	Y=log(annual revenues); years 2015-2018				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	0.048***	0.057***	0.010	0.001	0.063**
	(0.011)	(0.011)	(0.022)	(0.016)	(0.026)
Restaurant & Year FE	✓	✓	✓	✓	✓
ZIP-code*Year	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	6677	6652	2305	2299	2048
Restaurants	2043	2034	696	697	641
Clusters	57	56	39	40	41
Adj. R-squared	0.846	0.847	0.869	0.849	0.782
Mean Y pre-policy	646.6	648.8	977.4	558.0	360.7
DDD *p-value*				0.669	0.020

Post=1 if date is after 2016. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.80), [3.80, 4.20), [4.20, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-year. The sample includes observations between 2015 and 2018. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating.* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D2: Impact on restaurant employment

	Y=log(monthly employees); Jan 2015 - Dec 2019				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	0.041***	0.042***	-0.008	0.069***	0.052*
	(0.012)	(0.015)	(0.019)	(0.024)	(0.030)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	219835	217622	72133	76920	68569
Restaurants	4628	4576	1490	1571	1515
Clusters	86	86	59	71	71
Adj. R-squared	0.779	0.778	0.759	0.793	0.769
Mean Y pre-policy	5.5	5.6	6.9	5.7	4.0
DDD *p-value*				0.100	0.026

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating.* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D3: Impact on hiring decisions (extensive margins)

Y=1 if firm hires worker	*with* previous experience in restaurants				*without* experience in restaurants		
	(1) Full sample	(2) Low rating	(3) Mid rating	(4) High rating	(5) Low rating	(6) Mid rating	(7) High rating
Tourist*Post	0.007***	0.009***	0.004	0.005	-0.006**	0.006	0.006**
	(0.002)	(0.003)	(0.003)	(0.003)	(0.002)	(0.004)	(0.003)
Restaurant & Time FE	✓	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
Observations	217622	72133	76920	68569	72133	76920	68569
Restaurants	4576	1490	1571	1515	1490	1571	1515
Clusters	86	59	71	71	59	71	71
Adj. R-squared	0.124	0.143	0.104	0.116	0.049	0.043	0.037
Mean Y pre-policy	0.08	0.10	0.07	0.08	0.06	0.06	0.06
DDD *p-value*			0.985	0.113		0.006	0.013

The sample includes observations between Jan 2015 and Dec 2019. Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (3-4) and (6-7) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D4: Impact on hiring decisions (intensive margins)

Y=Months of experience in restaurants of	newly-hired employees				quitting/fired employees		
	(1) Full sample	(2) Low rating	(3) Mid rating	(4) High rating	(5) Low rating	(6) Mid rating	(7) High rating
Tourist*Post	0.678	2.478***	0.060	-0.780	-0.424	-0.051	1.412
	(0.851)	(0.802)	(1.264)	(0.656)	(1.389)	(1.113)	(1.855)
Restaurant & Time FE	✓	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
Observations	30059	11318	10205	8536	12281	10395	8131
Restaurants	3531	1163	1220	1148	1197	1226	1136
Clusters	76	53	59	61	51	57	58
Adj. R-squared	0.117	0.109	0.117	0.127	0.190	0.170	0.183
Mean Y pre-policy	13.0	13.5	13.5	11.8	25.8	27.0	21.5
DDD *p-value*			0.038	0.000		0.976	0.428

The sample includes observations between Jan 2015 and Dec 2019. Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (3-4) and (6-7) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D5: Impact on restaurant daily salaries

	Y=Average daily salary (€); Jan 2015 - Dec 2019				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	-0.121	-0.116	0.930***	-0.657	-0.680
	(0.232)	(0.259)	(0.343)	(0.680)	(0.431)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	200402	199026	67507	70593	60926
Restaurants	4558	4512	1492	1538	1482
Clusters	86	86	59	71	70
Adj. R-squared	0.467	0.469	0.485	0.465	0.451
Mean Y pre-policy	64.9	64.9	66.0	65.0	63.5
DDD *p-value*				0.181	0.095

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D6: Impact on restaurant Tripadvisor rating

	Y=Average 5-month Tripadvisor rating; Jan 2015 - Dec 2018				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	0.040***	0.040***	0.036*	0.076***	0.009
	(0.009)	(0.010)	(0.019)	(0.017)	(0.011)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	147274	146620	48577	53659	44384
Restaurants	4373	4330	1413	1499	1418
Clusters	86	86	59	70	70
Adj. R-squared	0.503	0.504	0.324	0.251	0.297
Mean Y pre-policy	3.98	3.98	3.51	4.05	4.43
DDD *p-value*				0.007	0.000

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.2 Robustness of results: excluding firms that exited the market after the policy

Table D7: Impact on restaurant revenues

	Y=log(annual revenues); years 2015-2018				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	0.068***	0.067***	0.033	0.043**	0.074**
	(0.016)	(0.017)	(0.030)	(0.016)	(0.030)
Restaurant & Year FE	✓	✓	✓	✓	✓
ZIP-code*Year	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	6176	6158	2148	2130	1880
Restaurants	1872	1865	642	637	586
Clusters	56	55	35	38	39
Adj. R-squared	0.852	0.853	0.871	0.858	0.790
Mean Y pre-policy	664.8	666.7	1011.7	563.7	366.1
DDD *p-value*				0.830	0.024

Post=1 if date is after 2016. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.80), [3.80, 4.20), [4.20, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-year. The sample includes observations between 2015 and 2018. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating.* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D8: Impact on restaurant employment

	Y=log(monthly employees); Jan 2015 - Dec 2019				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	0.050**	0.047*	-0.017	0.101***	0.047
	(0.022)	(0.026)	(0.020)	(0.034)	(0.051)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	199012	197346	65432	70772	61142
Restaurants	4071	4033	1315	1413	1305
Clusters	85	85	58	71	70
Adj. R-squared	0.784	0.783	0.764	0.797	0.777
Mean Y pre-policy	5.7	5.7	7.1	5.8	4.1
DDD *p-value*				0.099	0.063

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating.* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D9: Impact on hiring decisions (extensive margins)

Y=1 if firm hires worker	*with* previous experience in restaurants				*without* experience in restaurants		
	(1) Full sample	(2) Low rating	(3) Mid rating	(4) High rating	(5) Low rating	(6) Mid rating	(7) High rating
Tourist*Post	0.009***	0.010***	0.008**	0.002	-0.005**	0.010*	0.009***
	(0.002)	(0.003)	(0.004)	(0.003)	(0.002)	(0.005)	(0.003)
Restaurant & Time FE	✓	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
Observations	197346	65432	70772	61142	65432	70772	61142
Restaurants	4033	1315	1413	1305	1315	1413	1305
Clusters	85	58	71	70	58	71	70
Adj. R-squared	0.125	0.144	0.107	0.118	0.050	0.043	0.038
Mean Y pre-policy	0.08	0.10	0.07	0.07	0.06	0.06	0.06
DDD *p-value*			0.598	0.073		0.021	0.008

The sample includes observations between Jan 2015 and Dec 2019. Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (3-4) and (6-7) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D10: Impact on hiring decisions (intensive margins)

Y=Months of experience in restaurants of	newly-hired employees				quitting/fired employees		
	(1) Full sample	(2) Low rating	(3) Mid rating	(4) High rating	(5) Low rating	(6) Mid rating	(7) High rating
Tourist*Post	1.677*	3.294***	1.252	0.174	0.409	0.142	2.561**
	(0.911)	(1.156)	(1.270)	(0.612)	(1.034)	(1.195)	(1.103)
Restaurant & Time FE	✓	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
Observations	27448	10396	9414	7638	11228	9436	7171
Restaurants	3129	1035	1098	996	1056	1093	975
Clusters	74	53	57	59	51	56	55
Adj. R-squared	0.116	0.105	0.122	0.122	0.181	0.169	0.182
Mean Y pre-policy	13.1	13.7	13.4	11.8	26.0	27.4	21.8
DDD *p-value*			0.115	0.000		0.866	0.362

The sample includes observations between Jan 2015 and Dec 2019. Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (3-4) and (6-7) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D11: Impact on restaurant daily salaries

	Y=Average daily salary (€); Jan 2015 - Dec 2019				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	-0.007	-0.014	1.022**	-0.033	-1.025*
	(0.256)	(0.274)	(0.387)	(0.404)	(0.516)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	181903	180953	61649	64930	54374
Restaurants	4011	3979	1319	1384	1276
Clusters	85	85	58	71	69
Adj. R-squared	0.474	0.475	0.489	0.475	0.458
Mean Y pre-policy	64.8	64.9	65.9	65.0	63.5
DDD *p-value*				0.375	0.061

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D12: Impact on restaurant Tripadvisor rating

	Y=Average 5-month Tripadvisor rating; Jan 2015 - Dec 2018				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	0.053***	0.063***	0.112***	0.091***	0.011
	(0.012)	(0.012)	(0.018)	(0.022)	(0.013)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	131302	130941	43522	48710	38709
Restaurants	3849	3818	1248	1351	1219
Clusters	85	85	58	70	69
Adj. R-squared	0.503	0.504	0.324	0.252	0.300
Mean Y pre-policy	3.97	3.97	3.51	4.05	4.43
DDD *p-value*				0.000	0.000

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at municipality level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.3 Robustness of results: clustering standard errors at the ZIP-code level

Table D13: Impact on restaurant revenues

	Y=log(annual revenues); years 2015-2018				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	0.047	0.053	-0.002	0.033	0.069
	(0.030)	(0.033)	(0.075)	(0.052)	(0.071)
Restaurant & Year FE	✓	✓	✓	✓	✓
ZIP-code*Year	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	6677	6652	2305	2299	2048
Restaurants	2043	2034	696	697	641
Clusters	113	113	101	98	99
Adj. R-squared	0.846	0.847	0.869	0.849	0.782
Mean Y pre-policy	646.6	648.8	977.4	558.0	360.7
DDD *p-value*				0.965	0.419

Post=1 if date is after 2016. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.80), [3.80, 4.20), [4.20, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at ZIP-code level. Each observation is a restaurant-year. The sample includes observations between 2015 and 2018. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating.* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D14: Impact on restaurant employment

	Y=log(monthly employees); Jan 2015 - Dec 2019				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	0.043**	0.042*	-0.024	0.103***	0.041
	(0.021)	(0.022)	(0.043)	(0.036)	(0.033)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	219835	217622	72133	76920	68569
Restaurants	4628	4576	1490	1571	1515
Clusters	127	127	115	119	119
Adj. R-squared	0.779	0.778	0.759	0.793	0.769
Mean Y pre-policy	5.5	5.6	6.9	5.7	4.0
DDD *p-value*				0.078	0.051

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at ZIP-code level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating.* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D15: Impact on restaurant exit

	Y=1 if firm exits the market; Jan 2015 - Dec 2019				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	0.0011	0.0016*	0.0031**	-0.0000	0.0015
	(0.0008)	(0.0008)	(0.0015)	(0.0012)	(0.0018)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	219835	217622	72133	76920	68569
Restaurants	4628	4576	1490	1571	1515
Clusters	127	127	115	119	119
Adj. R-squared	0.059	0.060	0.058	0.061	0.061
Mean Y pre-policy	0.003	0.003	0.003	0.003	0.004
DDD *p-value*				0.365	0.621

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at ZIP-code level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D16: Impact on hiring decisions (extensive margins)

Y=1 if firm hires worker	*with* previous experience in restaurants				*without* experience in restaurants		
	(1) Full sample	(2) Low rating	(3) Mid rating	(4) High rating	(5) Low rating	(6) Mid rating	(7) High rating
Tourist*Post	0.009**	0.009	0.011**	0.002	-0.006	0.011*	0.007
	(0.004)	(0.007)	(0.005)	(0.007)	(0.006)	(0.006)	(0.006)
Restaurant & Time FE	✓	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
Observations	217622	72133	76920	68569	72133	76920	68569
Restaurants	4576	1490	1571	1515	1490	1571	1515
Clusters	127	115	119	119	115	119	119
Adj. R-squared	0.124	0.143	0.104	0.116	0.049	0.043	0.037
Mean Y pre-policy	0.08	0.10	0.07	0.08	0.06	0.06	0.06
DDD *p-value*			0.538	0.363		0.047	0.115

The sample includes observations between Jan 2015 and Dec 2019. Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at ZIP-code level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (3-4) and (6-7) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D17: Impact on hiring decisions (intensive margins)

Y=Months of experience in restaurants of	newly-hired employees				quitting/fired employees		
	(1) Full sample	(2) Low rating	(3) Mid rating	(4) High rating	(5) Low rating	(6) Mid rating	(7) High rating
Tourist*Post	1.469	2.977**	0.789	0.465	0.177	-0.440	2.375
	(0.924)	(1.149)	(1.584)	(1.680)	(1.548)	(1.928)	(2.135)
Restaurant & Time FE	✓	✓	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓
Observations	30059	11318	10205	8536	12281	10395	8131
Restaurants	3531	1163	1220	1148	1197	1226	1136
Clusters	121	108	111	114	107	110	113
Adj. R-squared	0.117	0.109	0.117	0.127	0.190	0.170	0.183
Mean Y pre-policy	13.0	13.5	13.5	11.8	25.8	27.0	21.5
DDD *p-value*			0.031	0.000		0.434	0.466

The sample includes observations between Jan 2015 and Dec 2019. Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at ZIP-code level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (3-4) and (6-7) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D18: Impact on restaurant daily salaries

	Y=Average daily salary (€); Jan 2015 - Dec 2019				
	(1) Full sample	(2) Full sample	(3) Low rating	(4) Mid rating	(5) High rating
Tourist*Post	-0.010	0.038	1.312**	-0.120	-1.125*
	(0.384)	(0.376)	(0.555)	(0.528)	(0.638)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	200402	199026	67507	70593	60926
Restaurants	4558	4512	1492	1538	1482
Clusters	125	125	114	118	118
Adj. R-squared	0.467	0.469	0.485	0.465	0.451
Mean Y pre-policy	64.9	64.9	66.0	65.0	63.5
DDD *p-value*				0.346	0.079

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at ZIP-code level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table D19: Impact on restaurant Tripadvisor rating

	Y=Average 5-month Tripadvisor rating; Jan 2015 - Dec 2018				
	(1)	(2)	(3)	(4)	(5)
	Full sample	Full sample	Low rating	Mid rating	High rating
Tourist*Post	0.040*	0.049**	0.087**	0.077**	-0.003
	(0.021)	(0.020)	(0.040)	(0.030)	(0.035)
Restaurant & Time FE	✓	✓	✓	✓	✓
ZIP-code*Time	✓	✓	✓	✓	✓
Controls		✓	✓	✓	✓
Observations	147274	146620	48577	53659	44384
Restaurants	4373	4330	1413	1499	1418
Clusters	127	127	115	119	119
Adj. R-squared	0.503	0.504	0.324	0.251	0.297
Mean Y pre-policy	3.98	3.98	3.51	4.05	4.43
DDD *p-value*				0.128	0.015

Post=1 if date is after May 2017. Tourist restaurants are those with a measure of exposure above the median. Restaurants in low, mid and high categories had their Tripadvisor rating at the time of the policy $\in [1, 3.85), [3.85, 4.25), [4.25, 5]$, respectively. Heteroskedasticity-robust standard errors clustered at ZIP-code level. Each observation is a restaurant-month-year. Distance (km) to closest attraction and indicators for distance to Rome city center interacted with Post are included in all regressions. Controls include indicators for restaurant price category, Italian cuisine, concentration of restaurants in 400m radius, n. of reviews to closest attraction, type of economic activity and legal status of the firm, all interacted with Post. Columns (4) and (5) report the p-values from a triple-difference estimation testing whether the DDD coefficients for mid and high ratings are equal to low rating. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$