Health policy 162 (2025) 105447

Contents lists available at ScienceDirect

Health policy

journal homepage: www.elsevier.com/locate/healthpol

Health hazard discrimination or prejudice? A correspondence experiment in Italy

Paolo Buonanno [a], Flavio Porta [a], Marcello Puca [b,*]

[a] University of Bergamo, Italy
[b] University of Salerno & CSEF, Italy

ARTICLE INFO

JEL codes:
C93
I18
J15

Keywords:
Discrimination
Prejudice
Social identity
COVID-19

ABSTRACT

We study how infectious-disease threats can spill over into discriminatory behavior. Using early COVID-19 in Italy as a case study, we ran an email correspondence experiment with 5356 tourism providers, randomly varying the sender's location and surname to signal origin from areas differentially hit by the first wave. Requests signaling origin from a highly affected area received about 5 percentage points fewer replies and more rejections than observationally equivalent requests; the penalty concentrated on North-sounding surnames and was absent for South-sounding surnames from the same city, pointing to prejudice rather than rational screening on contemporaneous infection risk. While our setting is tourism, the mechanism we uncover—disease-avoidance concerns activating social stereotypes—is general and consistent with theories of social stigma and the behavioral immune system. Such "health-hazard discrimination" can deter testing or travel, undermine equitable access to services, and amplify outbreaks when stigmatized groups avoid contact with providers. We discuss design and policy tools—bias-safe communication, temporary identity-blinding in first contacts, and platform-level fairness nudges—that can mitigate stigma-driven frictions during epidemics. Findings inform preparedness for future outbreaks beyond COVID-19.

1. Background

Infectious-disease threats routinely spill into the social sphere. When pathogen risk is salient, people rely on coarse signals of identity and place, which can produce stigma and discriminatory treatment even when actual infection risk is low. The COVID-19 pandemic, beyond its immediate health implications, has profoundly impacted social dynamics across various economic sectors. As countries grappled with the challenges of the pandemic, specific stigmas and biases emerged, particularly towards regions or groups perceived as more affected by the virus [1–4]. While this hidden cost of the pandemic might be overshadowed by the disease's primary toll, it holds significant implications for public health. Discrimination stemming from health crises can lead to unequal treatments and exacerbate societal inequities, especially for individuals migrating from an affected region or country. Furthermore, it may deter individuals from seeking testing to avoid association with the stigma, possibly exacerbating the hazard of spreading the virus.

The initial outbreak in Wuhan, China, inadvertently spurred an uptick in discriminatory behaviors and xenophobic sentiments towards Chinese individuals and those of East Asian descent worldwide. This discrimination was not merely rooted in the fear of potential virus carriers, but also targeted individuals of Asian origin based on their appearance. Essentially, Asian individuals encountered discrimination due to their ethnicity, regardless of their actual risk of carrying the virus. Such discrimination was not limited to Asian individuals; it also impacted those living in areas outside of China where the virus spread more rapidly during the early stages of the COVID pandemic. The observed stigmatization following outbreaks is not only a function of health-related fears but also reflects deeper mechanisms of group identity and exclusion. Drawing on Tajfel's social identity theory (1974), individuals may categorize others as part of an "out-group" based on regional origin, especially when that origin is associated with a perceived threat. This theoretical lens helps to explain why individuals from Bergamo—associated with the initial COVID-19 wave—may have been treated differently, irrespective of actual health risk.

Differentiating between "rational discrimination" associated with a genuine risk of contagion (i.e., statistical discrimination) and outright discrimination against specific populations regardless of the actual

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

* Corresponding author.
E-mail address: mpuca@unisa.it (M. Puca).

https://doi.org/10.1016/j.healthpol.2025.105447
Received 11 October 2024; Received in revised form 12 September 2025; Accepted 20 September 2025
Available online 3 October 2025

threat (i.e., prejudice) proves challenging. Data limitations often hide the manner in which individuals might respond when interacting with someone from a high-risk area. Our study overcomes these limitations exploiting a unique historical setting: Italy in the aftermath of the first COVID-19 wave in Spring 2020. Italy, especially the province of Bergamo in the Lombardy region, became a significant epicenter. Bergamo's overwhelming caseload led to harrowing scenes of overrun hospitals and military trucks transporting the deceased. As a result, Bergamo and its inhabitants faced inadvertent stigmatization, often being pinpointed as the primary source of Italy's COVID-19 spread. This association led to caution and prejudice against Bergamo's residents, fueled by fear and misinformation. Fig. 1 shows Italy's total excess mortality during the first COVID19 wave, a reliable epidemic severity indicator [5], by province. The areas highlighted in red distinctly reveal the uneven distribution of the disease, which is not attributable to the volume of testing as it is measured with excess mortality. Beyond documenting heterogeneity in excess mortality, maps like Fig. 1 can also seed durable place-based labels that later guide decisions under uncertainty—a mechanism relevant to any epidemic with uneven geography.

To empirically assess the existence and the extent of these biases, we conducted an e-mail correspondence experiment, sending accommodation requests to various Italian touristic hosts. The primary variable was the perceived origin of the potential tourist, manipulated by altering the sender's address and surname. In a control condition, ad- dresses and surnames were neutral, with no association to any COVID-19 affected area. In treatment conditions, addresses and surnames indicative of provinces heavily impacted by the first wave, primarily Northern Italy, were used. To disentangle potential discrimination from regional stereotypes and not just health-related concerns, we also sent emails from the same location but with surnames indicative of specific Italian regions, either North-sounding or South-sounding. This strategy allowed us to measure pandemic-related biases related to the address of the sender and differentiate them from pre-existing regional biases based on the surname.

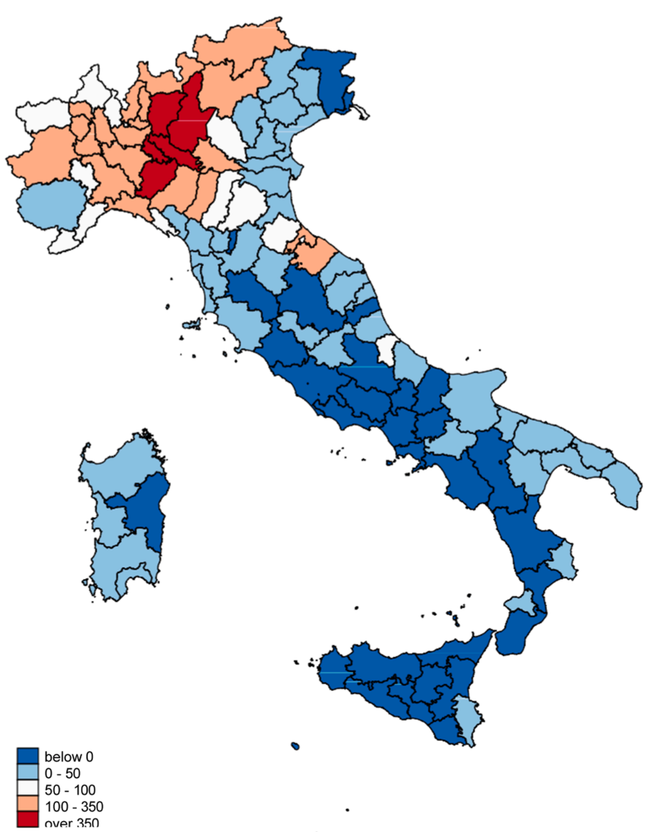

Fig. 1. Excess mortality during the first COVID19 wave, by Italian provinces.

Our findings indicate that senders coming from highly affected COVID-19 areas are approximately 5 % less likely to receive a response compared to neutral counterparts. This bias was more pronounced for senders with North-sounding surnames requesting accommodation in the South, with discrimination rate going up to 8 %. Senders from the same location but with a South-sounding name, instead, did not experience discrimination. This distinction suggests that regional stereotypes, rather than pandemic-related fears, might drive the observed heterogeneous rates of responses. This observation aligns with the notion that discrimination against Asian individuals was not linked to the actual risk of contagion but merely to their Asian identity. In our sample, individuals faced discrimination because their surnames reminded recipients of their origin from Bergamo, the Italian COVID epicenter. Additionally, our analysis revealed varied response patterns among accommodation types. Notably, B&Bs displayed a more pronounced bias than hotels, potentially due to the adaptable management protocols that smaller establishments like B&Bs can implement.

Furthermore, our data segmentation based on geographical markets in Italy reveals distinct discrimination patterns, with the Southern market exhibiting a bias against North-sounding names but not against South-sounding ones. This observation suggests potential regional biases where individuals might experience less discrimination in areas associated with their surnames. Another dimension of our analysis differentiates between hotels and B&Bs, noting that hotels frequently employ automated booking tools, unlike B&Bs. Discrimination patterns also differ between these accommodation types, with B&Bs showing a more pronounced bias, particularly against North-sounding names.

Our study contributes primarily to two branches of literature: the effects of discrimination on market outcomes and the societal implications of name-based biases. Labor market discrimination, especially based on names, has been a central theme in economic research. The seminal study by Bertrand and Mullainathan [6] highlighted that resumes with traditionally White-sounding names received 50 % more callbacks than those with African-American-sounding names, emphasizing the subtle biases influencing hiring decisions, even without explicit prejudice.

However, discrimination is not limited to the labor market. Studies by Guryan and Charles [7], Fryer Jr and Loury [8], and Fryer Jr and Levitt [9] delve into the societal repercussions of distinctively Black names and the intrinsic value of diversity. Bordalo et al. [10] further elaborates on the role of stereotypes in shaping perceptions and influencing behaviors. The digital realm also mirrors these biases, as evidenced by Ayres et al. [11] and Pope and Sydnor [12], who reveal racial disparities on platforms like eBay and Prosper.com. Bohren et al. [13] further explore the complexities of statistical discrimination, emphasizing the challenges in identifying its accuracy.

A study closely aligned with our research is by Ewens et al. [14], which differentiates between statistical discrimination and overt racial prejudice in the rental market. While our research employs a similar empirical strategy, it uniquely investigates discrimination within a public health context, especially relevant during the COVID-19 pandemic. Regional stereotypes and their implications have also been explored in the literature. Buonanno et al. [15] investigates biases associated with within-country stereo- types, suggesting that regional perceptions significantly influence interactions. This perspective is reinforced by Ge et al. [16] and Hudson et al. [17], which focus on racial biases in transportation network companies and the tourism sector, respectively.

This rest of the paper is structured as follows: Section 2 outlines the experimental design and methodology, Section 3 delves into the results and findings, and Section 4 discusses the mechanisms at play, as well as the limitation of our analysis. We conclude in Section 5.

2. Methods

2.1. Creating a database of email address

We designed our experiment to assess whether hosts discriminate against potential tourists perceived as SARS-CoV-2 carriers. Our initial step was to compile a list of potential hosts from popular Italian holiday destinations.

Our primary data source was the Italian National Institute of Statistics (ISTAT) database, which provides details on municipal touristic inflows. Out of 7914 Italian municipalities, the database identifies 2119 touristic municipalities across 11 regions (out of Italy's 20 regions). To ensure our sample represented popular destinations and to maximize the impact of our treatments, we selected municipalities based on two criteria: (i) those in the top 25*th* percentile for the ratio of tourists to population, and (ii) those in the top 25*th* percentile in terms of total tourist numbers. While criterion (i) targets popular destinations, criterion (ii) filters out smaller municipalities. This selection resulted in 331 municipalities. Using these municipalities, we sampled email addresses of touristic hosts from regional tourism portals, yielding 28,284 email addresses. From this list, we randomly selected 6000 emails. However, only 5356 were active, indicating potential outdated entries in our source. The email sample included bed & breakfasts (53.14 %), hotels (37.45 %), touristic farms (5.12 %), touristic resorts (1 %), and other accommodations (3 %). Armed with this list, we proceeded to send accommodation requests using three different sender profiles.

2.2. Creating tourist profiles

To credibly represent tourists seeking accommodation, we devised three distinct profiles without using actual identities. We established three Gmail accounts for this purpose. Two senders hailed from Bergamo, a mid-sized town in Northern Italy, and one from Rome, Italy's central capital. This distinction was designed to convey the sender's origin both explicitly and implicitly, using full street addresses and surnames. Indeed, each inquiry included multiple cues signaling the sender's origin: the name-surname combina tion, full mailing address, and email signature all explicitly stated the sender's city (e.g., Bergamo). While not every recipient may consciously recognize the geographic origin of a surname, many Italians associate specific names with regions. For instance, "Locatelli" is strongly associated with Lombardy, especially Bergamo, and "Capuozzo" with the South (e.g., Campania). We selected names that are both frequent and regionally distinctive, based on national telephone registry data. A map showing surname distribution (Fig. A1) illustrates the salience of the signal. As of July 2020, when our experiment took place, Italy had just emerged from its initial COVID-19 wave. This wave saw varied im- pacts across Italian regions, with Bergamo in Lombardy notably recognized as Europe's COVID-19 epicenter by international media.

However, the anticipated lower response rate for Lombardy tourists might also stem from factors unrelated to COVID-19. Italy's regions have historically exhibited cultural differences leading to varied social outcomes (e.g. [18]). Discrimination based on regional origin might be influenced by these disparities (e.g. [15]). To address this, we incorporated distinct name-surname combinations in the senders' signatures, aiming to discern cultural influences from the declared location in the emails. We derived these combinations from frequency data in the latest PagineBianche national telephone database managed by ItaliaOnline. After tabulating combinations by region, we selected three high-frequency combinations. Additionally, we incorporated a random zip code from each sender's municipality for authenticity. The resulting male sender profiles were: (i) "Antonio Mancini, Roma"; (ii) "Paolo Locatelli, Bergamo"; and (iii) "Salvatore Capuozzo, Bergamo". Fig. A1 depicts the distribution of these sur- names across Italy. While "Mancini" is widespread and serves as our baseline, "Locatelli" and "Capuozzo" resonate with Northern and Southern Italy, respectively. Thus, an email from Mr. Locatelli conveys both Northern Italy's cultural essence and Bergamo's location, whereas Mr. Capuozzo's email suggests a Southern Italian cultural background, despite his Bergamo residence.

2.3. Sending requests and measuring responses

With our three email profiles in place, we proceeded to send inquiries regarding accommodation availability. From July 1 to July 23, 2020, during a period when travel within Italy was unrestricted and Lombardy's infection rates had fallen substantially, each sender profile randomly sent emails to a third of the 6000 addresses in our host sample, totaling 2000 emails per profile. Importantly, the fieldwork took place in July 2020, a period when Italy had lifted internal travel restrictions and the incidence of new COVID-19 infections in Lombardy had dropped to very low levels. This timing is crucial in interpreting the results, as it indicates that the actual health risk was minimal, making the persistence of negative responses more likely attributable to prejudice than precaution. Each email, while bearing a distinct signature, contained an (almost) identical request: the availability and a quote of a double room from August 15 to August 22. We chose this timeframe, the peak of Italy's summer season, to maximize the likelihood of request congestion, thereby increasing the potential for discrimination.

Our primary objective was to ascertain if a sender's profile influenced the likelihood of receiving a response and, if so, whether the response was affirmative. On August 15, 2020, we accessed the inbox and spam folders of each sender account, documenting the responses. We established two variables: (i) *Answer*, a binary variable set to one if a host responded to the sender's email and zero otherwise; (ii) *Yes*, another binary variable set to one if, contingent on *Answer* being one, the host's response was positive.

2.4. Weaknesses of the experimental design

While the power of experimental randomization makes our design suitable for drawing causal inferences, it does have limitations. First, our outcome measure is binary and thus cannot capture the full complexity of discriminatory behavior. In particular, we observe only whether an email received a response and whether that response was positive. While this provides a clean and comparable metric, it overlooks other dimensions, such as the tone, timing, or quality of the reply. Future research might expand this framework by incorporating content analysis of replies or implementing multi-stage correspondence designs to uncover subtler or longer-term biases. Second, another limitation arises from the indirect signaling of the sender's origin. Rather than explicitly stating the sender location's health hazard, we hint at it through a specific address and city name. This approach introduces multiple concerns. While city names are selected to emphasize origin, some hosts might overlook them or fail to associate them with the intended COVID-19 implications. Consequently, our results might understate the actual degree of discrimination.

A third challenge inherent in our study, as well as in traditional audit studies, is the exclusive reliance on email correspondence. As widely acknowledged, personal interactions and recommendations play a pivotal role in the hospitality industry, a dimension we cannot explore in this context. This oversight could skew our results if, for instance, tourists from certain regions predominantly rely on personal recommendations or if hosts who are less discriminatory lean more on direct interactions.

3. Results

3.1. Response and acceptance rates

In Table 1, we examine the response (*Answer*) and acceptance (*Yes*) rates based on the origin-surname combination of the sender and the host's location or typology. For the entire dataset, the neutral profile had

Table 1
Response rates by host area and sender's surname.

Hosts	Answer				Yes			
	Neutral sounding	*Covid*	*North sounding*	*South sounding*	*Neutral sounding*	*Covid*	*North sounding*	*South sounding*
All	55.32 %	52.71 %	52.24 %	53.15 %	77.14 %	75.65 %	75.00 %	76.26 %
Ratio		1.05	1.06	1.04		1.02	1.03	1.01
% difference		2.61 %	3.08 %	2.17 %		1.49 %	2.14 %	0.88 %
p-value		0.0353	0.0334	0.0956		0.1869	0.1373	0.3233
South	57.61 %	54.55 %	52.20 %	56.44 %	75.35 %	75.35 %	76.87 %	74.22 %
Ratio		1.06	1.10	1.02		1.00	0.98	1.02
% difference		3.06 %	5.41 %	1.17 %		0.01 %	−1.51 %	1.14 %
p-value		0.0596	0.0107	0.2990		0.4991	0.7097	0.3293
Center	49.65 %	46.55 %	47.74 %	42.66 %	67.61 %	66.43 %	66.67 %	65.57 %
Ratio		1.07	1.04	1.16		1.02	1.01	1.03
% difference		3.10 %	1.91 %	6.99 %		1.17 %	0.94 %	2.03 %
p-value		0.2520	0.3451	0.1185		0.4258	0.4421	0.4034
North	53.19 %	53.05 %	56.95 %	50.38 %	81.84 %	80.57 %	79.07 %	81.74 %
Ratio		1.00	0.93	1.06		1.02	1.04	1.00
% difference		0.15 %	−3.76 %	2.82 %		1.27 %	2.77 %	0.11 %
p-value		0.4758	0.8928	0.1516		0.3145	0.1952	0.4855
Hotel	58.40 %	56.23 %	58.01 %	54.61 %	86.56 %	86.41 %	86.06 %	86.75 %
Ratio		1.04	1.01	1.07		1.00	1.01	1.00
% difference		2.17 %	0.39 %	3.79 %		0.15 %	0.50 %	−0.19 %
p-value		0.1815	0.4439	0.0812		0.4729	0.4214	0.5312
B&B	53.35 %	50.54 %	49.07 %	51.85 %	71.16 %	68.14 %	66.98 %	69.11 %
Ratio		1.06	1.09	1.03		1.04	1.06	1.03
% difference		2.81 %	4.28 %	1.51 %		3.03 %	4.18 %	2.06 %
p-value		0.0754	0.0322	0.2510		0.1127	0.0811	0.2341

Notes: This table presents the response rates categorized by the host area and the sender's surname. The rates are further divided based on the nature of the email's content, whether it sounded neutral, related to Covid, or had a regional connotation (North or South sounding). The percentages represent the proportion of responses received, and the p-values indicate the statistical significance of the observed differences based on a *t*-test. The ratios provide a relative comparison between the different categories.

a 55.32 % probability of receiving a response. Upon receiving a response, there was a 77.14 % chance of it being positive. In contrast, requests originating from areas severely impacted by the COVID pandemic had a slightly lower (and statistically significant) chance of receiving a response at 52.71 %, with a 75.65 % probability of it being positive. When further segmented the COVID-senders into North-sounding and South-sounding surnames, these profiles had response probabilities of 52.24 % and 53.15 %, with positive response rates of 75.00 % and 76.26 %, respectively. Given the significant difference in response rates for the North-sounding profile ($p = 0.03$) compared to the non-significant difference for the South-sounding one ($p = 0.09$), we suspect that the observed discriminatory effect might be more related to regional stereotypes than to COVID-related concerns.

We also segmented our data based on the geographical markets of Southern, Central, and Northern Italy and observed distinct patterns of discrimination. The Southern market displayed a pronounced bias against North-sounding names, while both the Central and Northern markets showed a bias against South-sounding names. Interestingly, in the Northern market, the North-sounding surname even outperformed the neutral profile, suggesting potential regional biases where individuals might face less discrimination in areas associated with their surnames. Finally, another dimension of our analysis differentiated between hotels and B&Bs. We observed that hotels often responded using automated booking tools, a trend not evident with B&Bs. Discrimination patterns also varied between these two accommodation types, with B&Bs exhibiting a more pronounced bias, especially against North-sounding names.

Result 1: Senders from COVID-impacted areas experienced noticeable discrimination in response rates.
Result 2: *Regional biases, combined with operational differences between hotels and B&Bs, played a significant role in shaping response patterns.*

3.2. Are "COVID-tourists" or Northern Italians less welcome?

Our analysis, as detailed in Table 2, reveals a pronounced disparity in response and confirmation rates based on the tourist's origin, discernible both in terms of the actual location and the surname of the applicant. Several underlying factors might be influencing these observed differences. The timing of the email or the nature of the respondent, be it a small B&B or a large hotel, could be potential contributors. To address these concerns and control for potential confounding variables, we conducted OLS regressions. The results, presented in Table 2, use the *Answer* and *Y es* dummies as dependent variables. The treatment variable, *Covid*, indicates if a request originates from Bergamo, recognized as the epicenter of the COVID-19 outbreak in Italy. Columns 1 and 5 of the table present the basic effect on response and positive response rates when an email is sent from Bergamo. In contrast, Columns 2 and 6 incorporate fixed effects for the day, time, and type of facility into the regression. The analysis is further refined in Columns 3, 4, 7, and 8 by introducing a dummy variable for facilities located in north- ern Italy, aiming to capture potential in-group biases [19]. A key observation from our findings is the consistently negative and statistically significant coefficient for *Covid* across all specifications for the answer rate. Notably, the interaction with the *North* dummy is not significant, suggesting that the facility's location does not influence response behavior.

The differential response rate observed between emails originating from Bergamo and those from a less affected location (i.e., Rome) prompts an essential inquiry: Is the observed effect a manifestation of rational statistical discrimination, or is it prejudice against individuals perceived to hail from a more infected region? To differentiate between these two scenarios, we introduce additional OLS regressions in Table 3. In this table, the treatment—specifically, an email from Bergamo—is split into two dummy variables: *Northern surname* and *Southern surname*. The *Northern surname* variable is assigned a value of one when the sender from Bergamo has the surname "Locatelli", emblematic of Bergamo, and zero otherwise. In contrast, the *Southern surname* variable is set to one when the email from Bergamo is signed by Mr. "Capuozzo", a distinctly Southern-sounding name.

Table 2
Baseline estimates.

	(1) *Answer*	(2)	(3)	(4)	(5) *Yes*	(6)	(7)	(8)
COVID	−0.026	−0.034	−0.040	−0.056	−0.015	−0.033	−0.011	−0.037
	(0.070)	(0.021)	(0.024)	(0.003)	(0.371)	(0.050)	(0.604)	(0.089)
North			−0.034	−0.089			0.074	0.014
			(0.161)	(0.002)			(0.006)	(0.677)
North x COVID			0.039	0.050			−0.002	0.015
			(0.198)	(0.102)			(0.961)	(0.652)
Constant	0.553	0.550	0.566	0.564	0.771	0.747	0.745	0.751
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
N. Obs.	5356	5356	5356	5356	2870	2870	2870	2870
Date FE	–	√	–	√	–	√	–	√
Hour FE	–	√	–	√	–	√	–	√
Type FE	–	√	–	√	–	√	–	√

Notes: This displays OLS estimates. In columns (1–4), the dependent variable *Answer* is a dummy variable set to one if the host responded, and zero otherwise. *Y es*, shown in columns (5–8), is another dummy variable, set to one for a positive host response, given an answer was received. The dummy variable (*COV ID*) takes a value of one for requests from Bergamo, Italy's most COVID-19 affected province, and zero otherwise. Columns (3–4) and (7–8) incorporate the interaction between the dummy variable *North*, indicating that an host is located in Northern Italy, and*COVID*. Control variables in columns (2), (4), (6), and (8) include: (i) email date, (ii) email hour, and (iii) host typology dummies. Standard errors are clustered at the firm level, with p-values denoted in parentheses.

Table 3
Disentangling prejudice from discrimination.

	(1) *Answer*	(2)	(3)	(4)	(5) *Yes*	(6)	(7)	(8)
COVID	−0.031	−0.050	−0.084	0.035	−0.021	−0.063	−0.028	−0.060
	(0.067)	(0.007)	(0.002)	(0.269)	(0.275)	(0.003)	(0.349)	(0.065)
North	−0.022	−0.024	−0.014	−0.030	−0.009	−0.014	−0.023	−0.001
	(0.191)	(0.146)	(0.542)	(0.278)	(0.647)	(0.469)	(0.377)	(0.976)
North x COVID			0.039	0.050			−0.002	0.015
			(0.198)	(0.102)			(0.961)	(0.652)
Constant	0.553	0.556	0.582	0.496	0.771	0.758	0.759	0.895
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
N. Obs.	5356	5356	2816	1789	2870	2870	1566	950
Date FE	–	√	√	√	–	√	√	√
Hour FE	–	√	√	√	–	√	√	√
Type FE	–	√	√	√	–	√	√	√
Sample	All	All	South	North	All	All	South	North

Notes: This table presents the results of OLS estimates. The dependent variable *Answer* in columns (1–4) is a dummy variable equal to one if the host responded to the request, and zero otherwise. *Y es* in columns (5–8) is a dummy variable equal to one in case of a positive response from the host, conditional on receiving an answer. The primary explanatory variables are *Northern surname* (i.e., Locatelli) and *Southern surname* (i.e., Capuozzo). Both requests originated from Bergamo, Italy's most COVID-19 affected province. Columns (3–4) and (7–8) focus on southern and northern hosts, respectively. Control variables in columns (2), (4), (6), and (8) include: (i) date of request, (ii) hour of request, and (iii) host typology dummies. Standard errors are clustered at the firm level, with p-values presented in parentheses.

The coefficients associated with the *Northern surname* in Table 3 consistently lean negative, implying that emails with the "Locatelli" surname generally garnered fewer responses. Notably, this effect is statistically significant in models (2), (3), and (6).

This suggests that the email's origin, coupled with a Northern-sounding surname, significantly influenced the probability of securing a response or a positive acknowledgment. Conversely, the coefficients for the *Southern surname* are uniformly negative but lack statistical significance in most specifications. This indicates that while there might be a marginal decline in response rates for emails bearing the "Capuozzo" surname, the effect isn't as marked or consistent as its Northern counterpart.

Diving deeper into the data, focusing on hosts from the South (columns 3 and 7) and the North (columns 4 and 8), the coefficients underscore that the Northern surname faces more discrimination than the Southern one, even though both emails originate from Bergamo. This observation accentuates the pivotal role of regional biases in shaping response dynamics, overshadowing mere statistical discrimination. This pattern is consistent with predictions from social identity theory [19]: despite identical epidemiological risk, hosts may have perceived senders with Northern surnames as members of a stigmatized out-group ("locals from the hotspot"), while Southern surnames may have mitigated this perception by signaling "non-local" or culturally distant identities, thereby softening the stigma.

Result 3: *Only the Northern surname faced discernible discrimination, indicating that prejudice, rather than rational statistical discrimination, is the driving factor behind the differential response rates.*

To ensure the robustness of our findings, we re-estimated all models using logistic regressions and included host-level covariates such as accommodation type and region. The results, presented in the Appendix (Tables A3 and A4), show that point estimates and significance levels remain consistent across specifications. Since each host received only one inquiry, clustering at the host level is not required; however, we also re-estimated all OLS models from Tables 2 and 3 with standard errors clustered at the municipality level. These results, reported in the Appendix (Tables A5 and A6), are fully in line with the main findings. Overall, the consistency across specifications indicates that randomization successfully balanced the profiles and that the estimated treatment effects are robust to alternative modeling choices.

4. Mechanisms and interpretation

4.1. Interpreting answer rates

The observed discrepancies in response rates based on a tourist's origin, especially from Bergamo, the COVID-19 epicenter in Italy, hint at deeper underlying mechanisms at play. One plausible explanation is the pervasive influence of stereotypes and biases. The global narrative around COVID-19 has, at times, inadvertently associated certain regions with higher risks, leading to apprehensions about individuals from these areas. In the context of our study, tourists from Bergamo might be perceived as potential carriers of the virus, leading to hesitancy in responses. Additionally, the surname of the applicant, which can often indicate ethnic or regional origins, might trigger implicit biases in the respondent. Such biases, even if unintentional, can affect decision-making processes.

Drawing from Tajfel's social identity theory, these biases can be understood in terms of in-group and out-group dynamics Tajfel [19]. According to the theory, individuals categorize themselves and others into in-groups (groups they identify with) and out- groups (groups they do not identify with). This categorization often leads to in-group favoritism and out-group discrimination. Tourists from Bergamo, due to the association with COVID-19, might be perceived as part of an out-group, leading to differential treatment. Such perceptions are reinforced when the surname indicates a specific regional origin, further strengthening the in-group vs. out-group distinction.

Another factor could be the economic considerations of the hospitality industry. Businesses, especially in the tourism sector, have been grappling with the challenges posed by the pandemic. The perceived risk of hosting someone from a COVID-19 hotspot might be seen as detrimental to their reputation, leading to differential response rates.

4.2. Potential confounders

Our study, robust in its experimental design, acknowledges certain potential confounders that might influence the observed biases. One such confounder is the distinction in response patterns between hotels and B&Bs. While larger establishments like hotels might use automated systems for responses, B&Bs often involve more personalized interactions.

This difference in management style could introduce variability in our results. However, the ability of our experiment to capture these nuances across different accommodation types offers a comprehensive understanding of discrimination.

Another consideration is the reading patterns of email recipients. If some skimmed through the emails, neglecting parts where the sender's name is placed, it could influence their response. Yet, the consistent placement of the name towards the end in our experiment ensures uniformity and minimizes the impact of this potential confounder. Moreover, the potential disassociation between a name and its regional origin, perhaps due to factors like internal migration, could blur direct associations. But our focus on surnames that are strongly indicative of specific regions ensures that the majority of associations remain accurate.

A significant challenge is distinguishing between biases rooted in COVID-19 fears and those stemming from regional prejudices. Our experiment, however, by using both North- sounding and South-sounding surnames from COVID-19 affected areas, is uniquely positioned to tease apart these intertwined biases. Moreover, the consistent application of surnames across all emails ensures that any biases observed are more likely due to the experiment's variables rather than migration patterns.

While the power of experimental randomization enhances internal validity, we ac- knowledge that certain limitations remain. Unobservable host traits, such as personal risk tolerance, political beliefs, or regional sentiments, could influence responses. Moreover, some hosts may have overlooked the sender's surname or address altogether, attenuating treatment effects. If anything, this suggests that our estimated effects may be conservative. Another confounding factor could be the context of regional COVID experience. Hosts in areas that experienced high case counts themselves may be more empathetic, or conversely, more cautious. Future research might adopt multistage designs or content analysis to capture more nuanced behavioral signals, such as tone, delay, or conditional responses.

In conclusion, while potential confounders exist, our experiment's design and method- ology should be robust enough to provide significant insights into the nature of discrimination in the context of the COVID-19 pandemic.

4.3. External validity

While our experimental design focuses on a specific type of interaction, requests for short-term accommodation, there are reasons to believe that the observed discriminatory patterns may extend beyond this particular setting. The hospitality market, especially on digital platforms, offers a context in which hosts make rapid decisions based on limited information, often influenced by heuristics and perceived risk. These dynamics are not unique to tourism: similar mechanisms operate in other domains where individuals are evaluated remotely and under uncertainty, such as job recruitment, access to healthcare, or housing applications. In all these cases, perceived group identity and stigmatized associations—such as being from a former COVID-19 hotspot, may influence decision-making, consciously or not. Of course, the extent and form of discrimination may vary across sectors, and we acknowledge that our findings should be interpreted within the contextual specificity of our investigation. Nonetheless, the psychological processes underpinning our results, particularly the activation of social identity and out-group avoidance in the presence of a collective threat, are broadly relevant and help inform a more general understanding of pandemic-related stigma.

4.4. Broader implications for infectious-disease contexts

Our evidence contributes to a broader literature on stigma and discrimination around infectious disease, in which perceived contagion risk heightens out-group avoidance and activates pre-existing stereotypes. Classic accounts of stigma and social identity explain how labels linked to disease become socially discrediting [20,19]. Work on the behavioral immune system shows that pathogen-avoidance motives can generalize beyond true risk, producing excessive social distancing and prejudice [21,22]. Empirically, outbreaks such as SARS, H1N1, Ebola, MERS, and HIV/AIDS repeatedly triggered discriminatory responses toward salient groups—often decoupled from actual exposure (e.g., [23,24]). Our design—showing that a North-sounding surname from the same city faces a penalty while a South-sounding one does not—aligns with this pattern: when uncertainty is high, stereotype-consistent signals guide decisions even in low-risk periods (July 2020 in our case). Thus, we interpret Fig. 1's spatial heterogeneity (p. 3) not just as historical context, but as a trigger for durable disease-linked social labeling that can re-emerge in future epidemics.

5. Conclusion

In this study, we have shed light on the subtle biases that emerged during the COVID-19 pandemic, in the context of the tourism sector in Italy. Our experimental approach has allowed us to make a causal inference regarding the nature of these biases, distinguishing between health-related concerns and regional prejudices. Although our case study is tourism in Italy, the mechanism we uncover—disease-avoidance concerns amplifying pre-existing social labels—is general and likely to recur whenever infectious hazards are salient (e.g., during future respiratory outbreaks). Such "health-hazard discrimination" has consequences that extend well beyond market fairness: it can depress testing and disclosure, deter mobility when mobility is safe, and ultimately hinder epidemic control. It may also spill over to other mobile populations, including workers and economic migrants, narrowing their opportunities even in low-risk periods.

For health policy, the implication is that stigma mitigation is not only an equity goal but an efficiency goal. Preparedness plans should anticipate that place-based or group-based labels can persist after the peak of an epidemic and shape behavior in unrelated service markets. Digital platforms, which play a pivotal role in the tourism sector, can be designed to minimize potential bias triggers, ensuring a more equitable experience for all users. To translate this into practice, public messaging should avoid geographic or ethnic shorthand and pair maps with clear statements of current individual risk and screening rules, so origin is not used as a proxy for infection. Platforms that mediate first contacts can anonymize early exchanges (temporarily hiding surnames and precise locations), standardize request forms, and deploy light-touch fairness nudges when response gaps emerge. Authorities and platforms should monitor acceptance gaps by origin, while pairing any movement restrictions with anti-stigma campaigns and facilitation of safe travel and testing. Because digital platforms already structure these interactions, they are well placed to implement—and iteratively refine—this toolkit.

CRediT authorship contribution statement

Paolo Buonanno: Writing – review & editing, Validation, Methodology, Data curation, Conceptualization. **Flavio Porta:** Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Formal analysis, Data curation. **Marcello Puca:** Writing – review & editing, Writing – original draft, Validation, Supervision, Software, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declarations of competing interest

None.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.healthpol.2025.105447.

References

[1] Bagcchi S. Stigma during the COVID-19 pandemic. Lancet Infect Dis 2020;20(7):782.
[2] Chandrashekhar V. The burden of stigma. Science 2020;369(6510):1419–23.
[3] Fallucchi F, Faravelli M, Quercia S. Fair allocation of scarce medical resources in the time of COVID-19: what do people think? J Med Ethics 2021;47(1):3–6.
[4] Islam A, Pakrashi D, Vlassopoulos M, Wang LC. Stigma and mis- conceptions in the time of the COVID-19 pandemic: a field experiment in India. Soc Sci Med 2021;278:113966.
[5] Buonanno P, Galletta S, Puca M. Estimating the severity of COVID-19: evidence from the Italian epicenter. Plos One 2020;15(10):e0239569.
[6] Bertrand M, Mullainathan S. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. Am Econ Rev 2004;94(4):991–1013.
[7] Guryan J, Charles KK. Taste-based or statistical discrimination: the eco- nomics of discrimination returns to its roots. Econ J 2013;123(572):F417–32.
[8] Fryer Jr RG, Loury GC. Valuing diversity. J Political Econ 2013;121(4):747–74.
[9] Fryer Jr RG, Levitt SD. The causes and consequences of distinctively black names. Q J Econ 2004;119(3):767–805.
[10] Bordalo P, Coffman K, Gennaioli N, Shleifer A. Stereotypes. Q J Econ 2016;131(4):1753–94.
[11] Ayres I, Banaji M, Jolls C. Race effects on eBay. Rand J Econ 2015;46(4):891–917.
[12] Pope DG, Sydnor JR. What's in a picture? Evidence of discrimination from prosper. Com. J Hum Resour 2011;46(1):53–92.
[13] Bohren JA, Haggag K, Imas A, Pope DG. Inaccurate statistical discrimination: an identification problem. National Bureau of Eco- nomic Research; 2019. Technical report.
[14] Ewens M, Tomlin B, Wang LC. Statistical discrimination or prejudice? A large sample field experiment. Rev Econ Stat 2014;96(1):119–34.
[15] Buonanno, P., Puca, M., and Sbriglia, P. (2020b). Corruption and within-country stereo- types. *Available at SSRN 3617590*.
[16] Ge Y, Knittel CR, MacKenzie D, Zoepf S. Racial discrimination in transportation network companies. J Public Econ 2020;190:104205.
[17] Hudson S, So KKF, Meng F, Cárdenas D, Li J. Racial discrimination in tourism: the case of African-American travellers in South Carolina. Curr Issues Tour 2020;23(4):438–51.
[18] Putnam R. The prosperous community: social capital and public life. Am Prospect 1993;13(Spring). Vol. 4Available online: http://www.prospect.org/print/vol/13 (accessed 7 April 2003).
[19] Tajfel H. Social identity and intergroup behaviour. Soc Sci Inf 1974;13(2):65–93.
[20] Goffman E. Stigma: notes on the management of spoiled identity. Englewood Cliffs, NJ: Prentice Hall; 1963.
[21] Murray DR, Schaller M. The behavioral immune system: implications for social cognition, social interaction, and social influence. Advances in experimental social psychology, 53. Academic Press; 2016. p. 75–129.
[22] Schaller M, Park JH. The behavioral immune system (and why it matters). Curr Dir Psychol Sci 2011;20(2):99–103.
[23] Parker R, Aggleton P. HIV-and AIDS-related stigma and discrimination: a conceptual framework and implications for action. Soc Sci Med 2003;57(1):13–24.
[24] Person B, Sy F, Holton K, Govert B, Liang A, , S.A.R.S. Community Outreach Team, Garza B, et al. Fear and stigma: the epidemic within the SARS outbreak. Emerg Infect Dis 2004;10(2):358.