Foreignness, Ethnicity, and Housing Discrimination: Evidence from Singapore

Abstract

Discrimination at the inquiry stage shapes access to rental housing. Existing correspondence studies focus on ethnic-minority status and typically treat foreignness as a secondary attribute, often proxied by signals that also convey ethnicity. This paper isolates foreignness as a distinct source of differential treatment. We conduct a correspondence experiment in Singapore, where the ethnic majority includes both citizens and foreign nationals, and send 3,300 inquiries to 1,500 letting agents. Foreign applicants face a substantial reduction in the probability of receiving a response, with an odds ratio of 0.71. The effect operates in two layers. Foreignness compounds disadvantage for ethnic-minority applicants, and it also generates differential treatment within the ethnic majority, with British Chinese applicants receiving fewer responses than Singaporean Chinese applicants. A multinomial specification shows that the effect operates through non-response rather than explicit refusal. The results indicate that foreignness is a distinct and quantitatively important dimension of housing discrimination.

Keywords: Housing discrimination, foreignness, nationality, correspondence experiment, rental markets, Singapore
JEL Codes: R21, R31, J15, J71

1 Introduction

The first point of contact between a prospective tenant and a landlord is an important moment. Inquiries that go unanswered foreclose viewings, viewings that never occur foreclose offers, and the exclusion compounds across stages. A large literature establishes that landlords and letting agents respond differently to applicants whose names signal ethnic-minority status (Carpusor and Loges, 2006; Hanson and Hawley, 2011; Ahmed, Andersson, and Hammarstedt, 2010; Auspurg, Hinz, and Schmid, 2017; Flage, 2018; Liu and Rommel, 2023). A related strand documents disadvantages faced by foreign applicants in rental markets. However, in most settings, foreignness is proxied by names or identity markers that also signal ethnic-minority status. The estimated foreign penalty therefore reflects a composite of foreignness and minority status rather than a distinct mechanism. Evidence from Japan illustrates this limitation: foreign applicants receive fewer responses, but foreign status is conveyed through non-native names that also signal ethnic difference (Sugasawa and Harano, 2026).

This paper isolates foreignness as a distinct source of differential treatment. To identify this component, we require a setting in which foreign status and ethnic-majority status do not move together. Singapore provides such a setting. The Chinese ethnic majority comprises both citizens and several foreign-national subgroups whose immigration histories and cultural associations differ. The same names and visible markers can therefore signal majority ethnicity together with British, Mainland Chinese, Hong Kong or American nationality. Comparing responses across these combinations isolates the foreign component of differential treatment from the ethnic component.

The analysis uses a field correspondence experiment on the property portal 99.co. Twenty-two fictitious applicants send 3,300 inquiries to 1,500 letting agents between March and April 2024. Applicants vary along three dimensions, namely ethnicity, gender and nationality. Outcomes are classified into positive responses, non-committal inquiries, explicit refusals and no response. The design permits the estimation of foreign effects both within ethnic-minority groups and within the ethnic majority.

The empirical results show a robust foreign penalty in the probability of receiving any response. The estimated odds ratio is 0.71, which corresponds to a reduction of 29 per cent in the odds of receiving a response for foreign applicants. The effect operates in two layers. First, foreignness compounds disadvantage for ethnic-minority applicants. Second, foreignness generates differential treatment within the ethnic majority. British Chinese applicants face an odds ratio of 0.70 relative to Singaporean Chinese applicants, which shows that majority ethnicity does not insulate applicants from foreign-based exclusion. A multinomial specification shows that the effect operates through non-response rather than through explicit refusal.

The contribution of the paper is to identify and quantify foreignness as a distinct dimension of housing discrimination. Existing correspondence studies typically pool foreign applicants into a single category or conflate foreignness with ethnic-minority status. The Singapore setting allows the foreign component to be identified separately and shows that it is quantitatively important. The findings imply that estimates of minority discrimination in rental markets may combine multiple mechanisms, and that national-origin signals can generate exclusion even when ethnic-majority appearance is held constant.

The remainder of the paper proceeds as follows. Section 2 sets out the analytical framework. Section 3 describes the experimental design. Section 4 presents the empirical results. Section 5 concludes.

2 Analytical Framework

Correspondence studies of housing markets have produced consistent evidence of differential treatment at the inquiry stage. Applicants whose names signal ethnic-minority status receive fewer responses than majority applicants across a wide range of settings. In many designs, foreignness is also present, but it is typically conveyed through the same signals that indicate ethnic-minority status. As a result, the estimated differences in response rates combine multiple dimensions of

applicant identity and do not isolate the contribution of foreignness as a separate mechanism.

To clarify this distinction, we treat foreignness and ethnic-minority status as analytically separable dimensions of differential treatment. Ethnic-minority status refers to group membership within a given society, while foreignness refers to national origin and the associated signals of cultural distance, language or institutional familiarity. The two dimensions often co-vary in practice, which makes them difficult to disentangle empirically. A research design that identifies one dimension while omitting the other will, in general, attribute to the included dimension any component that is correlated with the omitted one.

2.1 Economic mechanisms

Economic theory distinguishes three canonical mechanisms of discrimination. Taste-based discrimination arises when the discriminator has a preference against transacting with out-group counterparts and is willing to bear a utility cost to avoid such transactions (Becker, 1957). Statistical discrimination arises when individual-level attributes are imperfectly observed and the discriminator uses group-level signals to infer expected applicant characteristics (Phelps, 1972; Arrow, 1973). Identity-based discrimination arises when individuals derive utility from acting in accordance with the behavioural rules of the groups with which they identify, including prescriptions to favour in-group members in the allocation of scarce resources (Akerlof and Kranton, 2000).

The mechanisms are not mutually exclusive. A single screening decision may combine taste, inference and identity. A correspondence test at the inquiry stage observes the joint outcome of these mechanisms rather than the underlying motive. In housing markets, this distinction is particularly important because the applicant signal is noisy, the resource is excludable and rivalrous, and the screening decision belongs to the property holder or intermediary. Statistical discrimination is relevant when the agent believes that group signals predict contract-relevant attributes such as ability to pay rent, length of tenancy or maintenance behaviour. Taste-based discrimination is relevant when the agent attaches a direct cost to interacting with an out-group applicant. Identity-based discrimination is relevant when favouring an in-group applicant generates an identity payoff, or when allocating a scarce resource to an out-group applicant generates an identity cost.

Distinguishing these mechanisms empirically requires either a design that varies the information available to the discriminator, as in the conventional approach to statistical discrimination (Heckman, 1998; Altonji and Pierret, 2001), or a design that varies applicant signals in ways that expose different combinations of ethnicity, nationality and cultural distance. This paper follows the second route. It does not claim to separately estimate taste, statistical and identity-based discrimination. Instead, it identifies the dimension of applicant identity, foreignness, through which these mechanisms may operate.

2.2 Targets of housing discrimination

To organise the empirical setting, we use a tripartite classification of housing-discrimination targets. Colonial-type discrimination refers to differential treatment rooted in historical hierarchies asso-

ciated with empire and colonial-era population movements. Graylined-type discrimination refers to differential treatment rooted in symbolic association with historically contingent events or social narratives, such as event-triggered stigma. Foreign-type discrimination refers to differential treatment rooted in national-origin signals, language, accent or assumed cultural distance, in the absence of either a colonial linkage or a symbolic event that marks the group.

The categories map onto different mixtures of the economic mechanisms above. Colonial-type discrimination is theoretically connected to taste-based and statistical mechanisms inherited from historical hierarchies. Graylined-type discrimination is theoretically connected to identity-prescription and statistical mechanisms triggered by symbolic stigma. Foreign-type discrimination is theoretically connected to statistical and identity-based mechanisms tied to cultural distance. The classification is not intended to assign one motive to each empirical contrast. Rather, it clarifies why estimates that pool target categories may combine distinct economic channels. Online Appendix Table A8 summarises the classification and the associated economic mechanisms.

This paper focuses on foreign-type discrimination. The empirical challenge is to construct comparisons in which foreignness varies while ethnic signals are held constant, and vice versa. Most existing correspondence studies do not exploit this variation. Foreign applicants are frequently represented by names that differ from the ethnic majority, so that the observed foreign penalty reflects a joint effect. Evidence from Japan illustrates this issue: foreign applicants receive fewer responses, but foreign status is conveyed through non-native names that also signal ethnic difference (Sugasawa and Harano, 2026). The resulting estimates do not distinguish whether differential treatment arises from national origin, ethnic-minority status, or their interaction.

The setting used in this paper permits the required separation. In Singapore, the ethnic majority comprises both citizens and foreign nationals. Applicants can therefore share an ethnic-majority signal while differing in nationality. Comparisons within this group identify the effect of foreignness while holding ethnicity constant. Comparisons across ethnic groups, conditional on nationality, identify the effect of ethnic-minority status. Estimating both sets of contrasts within a unified framework allows the contribution of foreignness to be isolated from that of ethnicity.

This separation yields two testable implications. First, if foreignness is a distinct source of differential treatment, foreign applicants should receive fewer responses than otherwise comparable local applicants, even within the same ethnic group. Second, if foreignness operates independently of minority status, differential treatment should also be observed within the ethnic majority when nationality varies. The empirical analysis tests these implications directly.

3 Experimental Design and Implementation

The empirical objective is to identify the effect of foreignness separately from ethnic-minority status. The design exploits variation in nationality within the ethnic majority and variation in ethnicity within nationality, so that the two dimensions can be estimated jointly.

3.1 Identification strategy

The key comparison holds ethnicity constant while varying nationality. In Singapore, Chinese applicants constitute the ethnic majority and include both citizens and foreign nationals. We therefore compare Singaporean Chinese applicants with Chinese applicants of British, Mainland Chinese, Hong Kong and American nationality. This contrast identifies the effect of foreignness within the ethnic majority. A second set of comparisons holds nationality constant while varying ethnicity, which provides the corresponding contrasts for ethnic-minority status. Estimating both sets of comparisons within the same regression permits the foreign component to be separated from the ethnic component.

The identification relies on the fact that nationality is made visible without changing the ethnic signal. Names signal ethnicity and gender, while telephone country codes and a relocation sentence signal national origin. The British Chinese, Mainland Chinese, Hong Kong Chinese and American Chinese profiles therefore allow the foreign component to be observed among applicants who share the ethnic-majority signal. Conversely, Singaporean Chinese, Indian, Malay and White applicants allow ethnic differences to be observed while nationality is held fixed. This structure is central to the paper, because most correspondence studies observe foreignness and minority status through the same signal.

3.2 Applicant profiles and allocation

The analysis is based on a correspondence experiment conducted on the property portal 99.co. Twenty-two fictitious applicants sent 3,300 inquiries to 1,500 letting agents between March and April 2024. Each profile sent 150 inquiries. The full profile schedule is reported in Online Appendix Table A2, and the sampling structure is reported in Online Appendix Table A1.

Applicants vary along three dimensions: ethnicity, nationality and gender. The ethnicity categories are Chinese, Indian, Malay and White. The nationality categories are Singaporean, British, Mainland Chinese, Hong Kong Chinese and American. Names were selected to signal ethnicity and gender and were validated through a survey of Singapore-based respondents. Each applicant was assigned a country-specific telephone number, so that nationality was signalled through the country code displayed in the inquiry message. Foreign applicants included a short statement indicating relocation from the country of origin.

The Singaporean and British applicants were assigned to four groups of four, defined by gender and nationality. Within each group, the four profiles corresponded to the four ethnic categories. Each member sent inquiries to 150 randomly assigned agents. This part of the design generated 2,400 messages to 600 agents, with each agent receiving four inquiries within the relevant assignment group. The Mainland Chinese, Hong Kong Chinese and American Chinese profiles were assigned separately, with six profiles sending 900 messages to 900 additional agents. A four-day interval between inquiries to the same agent was maintained to minimise detection.

3.3 Messages, outcomes and property data

All inquiries used standardised message templates with randomised wording to avoid systematic differences in tone or content across applicant types. The message text was constructed from randomised opening, body and closing components. This reduced the probability that agents would recognise repeated wording across applicant profiles. The message templates are described in Online Appendix Section A.

Outcomes were recorded within seven days of each inquiry and classified into four categories: positive responses, non-committal inquiries, explicit refusals and no response. Positive responses include invitations to view, offers to discuss viewing, or offers of alternative listings. Non-committal inquiries include requests for additional information that do not yet imply movement towards a viewing. Explicit refusals include statements that the property is unavailable or that the applicant is not eligible. The primary outcome used in the analysis is an indicator for receiving any response.

Property listings were drawn from the active rental market without restriction on type or tenure. The sample includes 1,500 listings with a mean monthly rent of SGD 4,336, and it includes both condominium and HDB-flat listings. Property and landlord characteristics are included as controls in all specifications. Full property characteristics are reported in Online Appendix Table A3.

Two features of the design are important for interpretation. First, the variation in nationality within the ethnic majority provides direct identification of foreignness, which is not available in most correspondence studies. Second, all contrasts are estimated within a unified regression framework, so that the effect of foreignness is identified while controlling for ethnicity and other applicant attributes.

4 Results

We report the effect of foreignness on the probability of receiving any response to a rental inquiry. Responses are recorded within seven days and include positive replies, non-committal inquiries and explicit refusals. The outcome variable equals one if any response is received and zero otherwise. Standard errors are clustered at the property level, and all specifications include controls for landlord characteristics and property attributes. Descriptive comparisons by ethnicity and nationality are reported in Online Appendix Figure A1 and Figure 1. Full preliminary tests are reported in Online Appendix Table A4.

Figure 1 illustrates the central descriptive pattern. Singaporean Chinese applicants receive any response in 46.67 per cent of inquiries, while British Chinese applicants receive any response in 38.00 per cent. Mainland Chinese and American Chinese applicants fall between these two values, and Hong Kong Chinese applicants receive a higher descriptive response rate. The figure therefore motivates the core empirical question: whether nationality generates differential treatment even when ethnic-majority status is held constant.

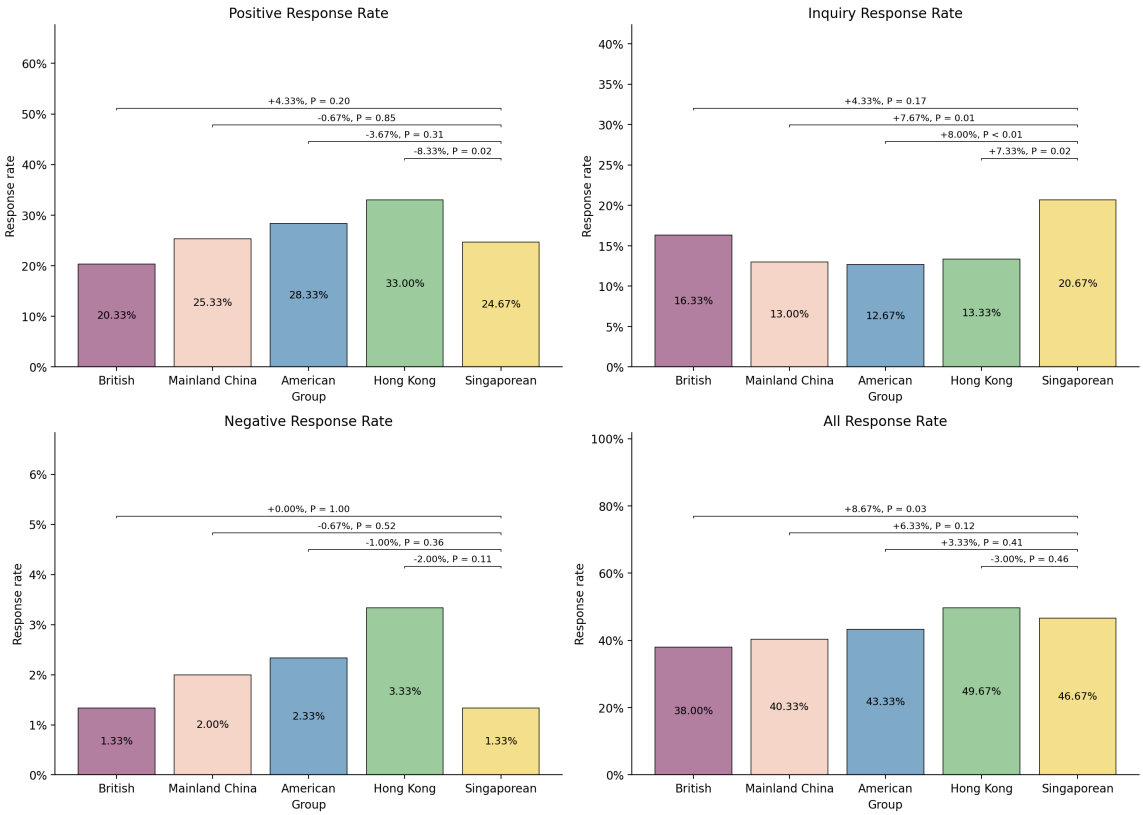

Figure 1: Response rates among Chinese applicants by nationality

Notes: The figure reports response rates among Chinese applicants by nationality and response category. Singaporean Chinese applicants provide the reference category for the within-majority comparison. The figure is descriptive; regression estimates with controls are reported in Table 1.

4.1 Main regression results

Table 1 reports the main logistic regressions. Column (1) includes the foreign indicator alone. Column (2) adds applicant ethnicity. Column (3) adds foreign-by-ethnicity interactions. Column (4) decomposes the foreign-Chinese interaction into country-specific Chinese categories.

The baseline specification in Column (2) gives an odds ratio of 0.709 for the foreign indicator. This corresponds to a reduction of approximately 29 per cent in the odds of receiving any response. The result is stable after applicant ethnicity and property controls are included, which indicates a substantial foreign penalty at the inquiry stage.

4.2 Foreignness and ethnic-minority status

Column (3) examines whether the foreign penalty reflects foreignness itself or its interaction with ethnic-minority status. The results show that foreignness compounds disadvantage for ethnic-minority applicants. The foreign–Indian, foreign–Malay and foreign–White interactions yield odds ratios of 0.611, 0.636 and 0.648, respectively. These estimates are statistically significant and economically large.

Table 1: Main logistic regressions of any-response receipt

	(1)	(2)	(3)	(4)
	Foreign only	+ Ethnicity	+ Eth Foreign	+ Country dummies
Male	0.980 (-0.19)	0.979 (-0.20)	0.979 (-0.20)	0.979 (-0.20)
Foreign	0.801** (-1.96)	0.709*** (-2.86)		
Indian (applicant)		0.721*** (-3.73)	0.909 (-0.84)	0.909 (-0.84)
Malay (applicant)		0.623*** (-5.47)	0.770** (-2.57)	0.770** (-2.57)
White (applicant)		0.731*** (-3.61)	0.897 (-0.94)	0.897 (-0.94)
Foreign Chinese			0.868 (-1.09)	
Foreign Indian			0.611*** (-2.89)	0.610*** (-2.89)
Foreign Malay			0.636*** (-2.60)	0.635*** (-2.61)
Foreign White			0.648** (-2.55)	0.647** (-2.55)
British Chinese				0.702** (-2.10)
Mainland Chinese				0.795 (-1.38)
Hong Kong Chinese				1.124 (0.71)
American Chinese				0.897 (-0.65)
Property and landlord controls	Yes	Yes	Yes	Yes
	3,300	3,300	3,300	3,300
Pseudo	0.010	0.016	0.017	0.019

Notes: Cells report odds ratios with z-statistics in parentheses. The dependent variable equals one if the applicant receives any response within seven days. Property and landlord controls comprise landlord gender, landlord ethnicity, property type, monthly rent, property size and number of bedrooms. Standard errors are clustered at the property level. *, ** and *** denote significance at the 10, 5 and 1 per cent levels.

The within-Singaporean ethnic-minority coefficients are smaller. The Indian and White coefficients are not statistically distinguishable from one once controls are included, while the Malay coefficient remains below one. The evidence therefore indicates that differential treatment for several minority profiles operates primarily through foreign status rather than through ethnic-minority status alone. This result supports the central claim that foreignness should not be subsumed within a broad minority category.

4.3 Foreignness within the ethnic majority

Column (4) examines whether foreignness produces differential treatment within the ethnic majority. It compares foreign-Chinese applicants of different national origins with Singaporean Chinese applicants. British Chinese applicants face an odds ratio of 0.702 relative to Singaporean Chinese applicants. This estimate is statistically significant in the unadjusted regression family and shows that majority ethnicity does not fully insulate applicants from foreign-based exclusion.

The estimates for Mainland Chinese, Hong Kong Chinese and American Chinese applicants are not statistically distinguishable from the Singaporean Chinese benchmark after adjustment for multiple testing. Online Appendix Table A5 reports the Holm-adjusted p-values. The appropriate interpretation is therefore not that there is a clean ranking of all foreign-Chinese subgroups, but that the design identifies a foreign penalty within the ethnic majority, with the British Chinese contrast providing the main empirical anchor.

Table 2: Binary logistic and multinomial logit specifications

Variable	Any response	Positive	Inquiries	Negative
Male	0.979 (-0.20)	0.897 (-1.25)	1.155 (1.41)	0.793 (-0.97)
Foreign	0.709*** (-2.86)	0.777*** (-2.65)	0.615*** (-4.40)	0.751 (-1.12)
Indian	0.721*** (-3.73)	0.629*** (-3.67)	0.857 (-1.09)	0.865 (-0.42)
Malay	0.623*** (-5.47)	0.545*** (-4.70)	0.737** (-2.11)	0.755 (-0.79)
White	0.731*** (-3.61)	0.751** (-2.36)	0.634*** (-2.97)	1.264 (0.75)
	3,300	3,300	3,300	3,300
Property and landlord controls	Yes	Yes	Yes	Yes

Notes: Cells report odds ratios with z-statistics in parentheses. The binary-logistic column estimates any response. The multinomial columns estimate a single model on the four-way response category, with no response as the base. *, ** and *** denote significance at the 10, 5 and 1 per cent levels.

4.4 Mechanism: silence rather than explicit refusal

Table 2 reports the multinomial specification. The model decomposes the any-response penalty into positive responses, non-committal inquiries and explicit negative responses, with no response as the base category.

The foreign coefficient is below one for positive responses and for follow-up inquiries. It is not statistically distinguishable from one for explicit negative responses. The result indicates that foreign applicants are not primarily refused in direct terms. Instead, they are less likely to receive the responses that would move the search process forward. The foreign penalty therefore operates through non-response.

This finding is important for interpretation. Explicit refusals are easier to observe and regulate than silence, but the empirical pattern shows that silence is the principal margin through which exclusion occurs. Studies that measure only explicit rejection may therefore understate inquiry-stage discrimination.

4.5 Robustness and heterogeneity

The remaining analyses support the main interpretation. Online Appendix Table A6 reports heterogeneity by property type and rent tier. The foreign penalty is larger in HDB-flat listings than in condominium listings and is concentrated in the middle rent tier. Online Appendix Table A7 reports the country-specific coefficients within each stratum. These patterns indicate that the foreign penalty is not uniform across the rental market, but the core result does not depend on these heterogeneity tests.

Gender does not generate a consistent pattern of differential treatment in this design. Online Appendix Figure A2 reports gender differences by ethnicity and response category. The male–female differences are small and do not alter the interpretation of the foreign penalty.

Taken together, the results identify a distinct foreign component of differential treatment. Foreignness reduces the probability of response on average, compounds disadvantage for ethnic-minority applicants, and generates measurable exclusion within the ethnic majority.

5 Discussion and Conclusion

This paper examines whether foreignness constitutes a distinct dimension of housing discrimination at the inquiry stage. The empirical results support this proposition. Foreign applicants face a lower probability of receiving a response, and the magnitude of the effect is economically large. The penalty operates in two ways. It compounds disadvantage for ethnic-minority applicants and, separately, produces differential treatment within the ethnic majority. The latter result shows that majority ethnicity does not insulate applicants from foreign-based exclusion.

The findings refine the interpretation of correspondence-test evidence in housing markets. Existing studies typically estimate differential response rates across groups defined by names or identity markers that jointly signal ethnicity and nationality. The resulting estimates therefore combine multiple mechanisms. The Singapore setting permits the foreign component to be isolated by comparing applicants who share an ethnic-majority signal but differ in national origin. The evidence indicates that foreignness accounts for a substantial share of observed differential treatment, and that it should not be subsumed within a broader ethnic-minority category.

The results also clarify the form through which exclusion occurs at the inquiry stage. The multinomial specification shows that the foreign penalty arises through a reduction in positive responses and follow-up inquiries, rather than through an increase in explicit refusals. Differential treatment therefore operates primarily through non-response. This pattern has implications for measurement, because correspondence tests that focus on explicit refusals may understate the extent of exclusion.

The findings are consistent with a combination of statistical and identity-based mechanisms. Nationality signals at the inquiry stage may be used as group-level information about unobserved tenancy attributes, including expected communication costs, tenancy duration or maintenance behaviour. At the same time, the presence of differential treatment within the ethnic majority suggests that cultural distance and group affiliation shape screening decisions even when observable ethnic signals are held constant. The design does not identify the relative contribution of statistical and identity-based channels, but it shows that foreignness is the applicant dimension through which these channels can operate separately from ethnic-minority status.

Two limitations qualify the interpretation. First, the design observes only the initial stage of the rental search and does not capture discrimination that may arise at later stages, such as viewings or tenancy agreements. Second, the setting is specific to a market in which the ethnic majority includes both citizens and foreign nationals, which facilitates identification but may limit direct external validity. The empirical strategy, however, can be applied to other settings in which nationality varies within observable ethnic categories.

The main implication is that national origin should be treated as a separate dimension of analysis in studies of housing discrimination. Designs that pool foreign applicants into a single category or equate foreignness with ethnic-minority status risk conflating distinct sources of differential treatment. Isolating foreignness clarifies both the magnitude of the effect and the mechanism through which it operates.

Data and Code Availability

The data used in this study were generated through a correspondence experiment conducted on the 99.co rental platform between March and April 2024. The dataset records inquiry-level outcomes for 3,300 messages sent to 1,500 property listings.

The analysis code used to produce all results in the paper is available and fully reproducible. The data and code will be made publicly available in a recognised repository upon publication. To protect personal data and platform integrity, identifying information, including contact details and listing identifiers, will be removed or masked in the public version of the dataset. A replication file with anonymised data sufficient to reproduce all reported results will be provided.

Prior to publication, the data and code are available from the corresponding author upon reasonable request.

The study received approval from the relevant Research Ethics Committee. The experiment was not pre-registered. All analyses conducted are reported in the manuscript and the accompanying appendix.

References

Akerlof, G. A., and R. E. Kranton. 2000. "Economics and Identity." *Quarterly Journal of Economics* 115(3): 715–53.

Altonji, J. G., and C. R. Pierret. 2001. "Employer Learning and Statistical Discrimination." *Quarterly Journal of Economics* 116(1): 313–50.

Arrow, K. J. 1973. "The Theory of Discrimination." In *Discrimination in Labor Markets*, edited by O. Ashenfelter and A. Rees. Princeton University Press.

Acolin, A., R. Bostic, and G. Painter. 2016. "A Field Study of Rental Market Discrimination Across Origins in France." *Journal of Urban Economics* 95: 49–63.

Ahmed, A. M., L. Andersson, and M. Hammarstedt. 2010. "Can Discrimination in the Housing Market Be Reduced by Increasing the Information about the Applicants?" *Land Economics* 86(1): 79–90.

Auspurg, K., T. Hinz, and L. Schmid. 2017. "Contexts and Conditions of Ethnic Discrimination: Evidence from a Field Experiment in a German Housing Market." *Journal of Housing Economics* 35: 26–36.

Baldini, M., and M. Federici. 2011. "Ethnic Discrimination in the Italian Rental Housing Market." *Journal of Housing Economics* 20(1): 1–14.

Becker, G. S. 1957. *The Economics of Discrimination*. University of Chicago Press.

Carpusor, A. G., and W. E. Loges. 2006. "Rental Discrimination and Ethnicity in Names." *Journal of Applied Social Psychology* 36(4): 934–52.

Christensen, P. O., and C. Timmins. 2022. "Sorting or Steering: The Effects of Housing Discrimination on Neighborhood Choice." *Journal of Political Economy* 130(8): 2110–63.

Christensen, P. O., and C. Timmins. 2023. "The Damages and Distortions from Discrimination in the Rental Housing Market." *Quarterly Journal of Economics* 138(4): 2505–57.

Flage, A. 2018. "Ethnic and Gender Discrimination in the Rental Housing Market: Evidence from a Meta-Analysis of Correspondence Tests, 2006–2017." *Journal of Housing Economics* 41: 251–73.

Heckman, J. J. 1998. "Detecting Discrimination." *Journal of Economic Perspectives* 12(2): 101–16.

Hanson, A., and Z. Hawley. 2011. "Do Landlords Discriminate in the Rental Housing Market? Evidence from an Internet Field Experiment in US Cities." *Journal of Urban Economics* 70(2–3): 99–114.

Phelps, E. S. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review* 62(4): 659–61.

Liu, Z. M., and J. Rommel. 2023. "Unequal Treatment in the Rental Housing Markets of Four Chinese Megacities." *Journal of Housing Economics* 61.

Sugasawa, T., and K. Harano. 2026. "Disentangling Transaction-Stage and Management-Stage Discrimination in the Rental Housing Market: Evidence from a Correspondence Experiment in Tokyo." *Journal of Housing Economics* 72: 102134.