sociological science

Supplement to:

Thomas, Kyla. 2018. "The Labor Market Value of Taste: An Experimental Study of Class Bias in U.S. Employment." Sociological Science 5: 562-595.

Methodological Appendix

This article draws on data from a five-month audit study of callback discrimination in four U.S. cities and a survey-experimental study of 1,428 U.S. hiring managers to investigate how cultural signals of class shape the labor market outcomes of U.S. workers. Below, I provide additional details regarding the cultural signals of class I selected for the study, the design of the audit study and survey experiment, and the design of the résumés.

Culture Signals of Class

To confirm that the combinations of highbrow versus lowbrow cultural tastes I selected for this study not only objectively reflect differences in social class but are subjectively perceived to do so, I conducted a series of online surveys on Amazon.com's crowdsourcing platform, Mechanical Turk. First, I distributed three exploratory surveys to measure the extent to which status-based differences in cultural taste are perceived to reflect differences in social class. Approximately 200 U.S. residents were sampled for each survey. One survey asked respondents a series of questions about people with differing musical preferences; the second asked respondents about people with differing sports preferences; the third asked respondents about people with differing food preferences.

To measure the perceived status of the music, sports, and food preferences, I used a well-established measure of status developed by Fiske et al. (2002) and applied widely in social psychology. Respondents were asked to rate, on a five-point Likert scale (1=not at all…5=extremely), how economically successful, well educated, and prestigious in occupation people with differing preferences were likely to be. To measure perceptions of class, respondents were asked to rate the social class (1=poor…5=upper class) of the same people. From these questions, I derived average rankings of the perceived status and perceived social class of individuals with various music, sports, and food preferences.

The Mechanical Turk sample was approximately 70% white, 55% female, 60% under the age of 35, and 40% college educated. Mechanical Turk samples more closely match the U.S. population than traditional university samples (Paolacci et al. 2010) but they are not nationally representative. To account for the distinctive demographic composition of Mechanical Turk users, I analyzed the effect of key sociodemographic factors on the rankings of taste reported in Figure 1. I do not find that the relative ranking of these preferences varies substantially by respondent gender, age, or educational background, so there is generally a consensus across demographic groups.

Mean results from the surveys are shown in Figure 1. From the surveys, it is clear that perceptions of status and class are strongly correlated. Statistical analyses confirm that, for music, sports, and food, the correlation between perceptions of status and perceptions of class is greater than 0.8. On average, individuals with a taste for stereotypically high status musical genres like classical music, opera, and jazz are systematically perceived as higher class than individuals with a taste for stereotypically low status musical genres ($p<0.001$). This extends to sports and food as well. For instance, individuals who participate in sports like sailing, golf, and tennis or who like traditionally European dishes like coq au vin and filet mignon are systematically perceived as higher status and higher class than individuals who participate in sports such as boxing, basketball, and bowling and prefer familiar dishes like cheeseburgers and tuna noodle casserole ($p<0.001$). I conclude that Americans systematically perceive highbrow

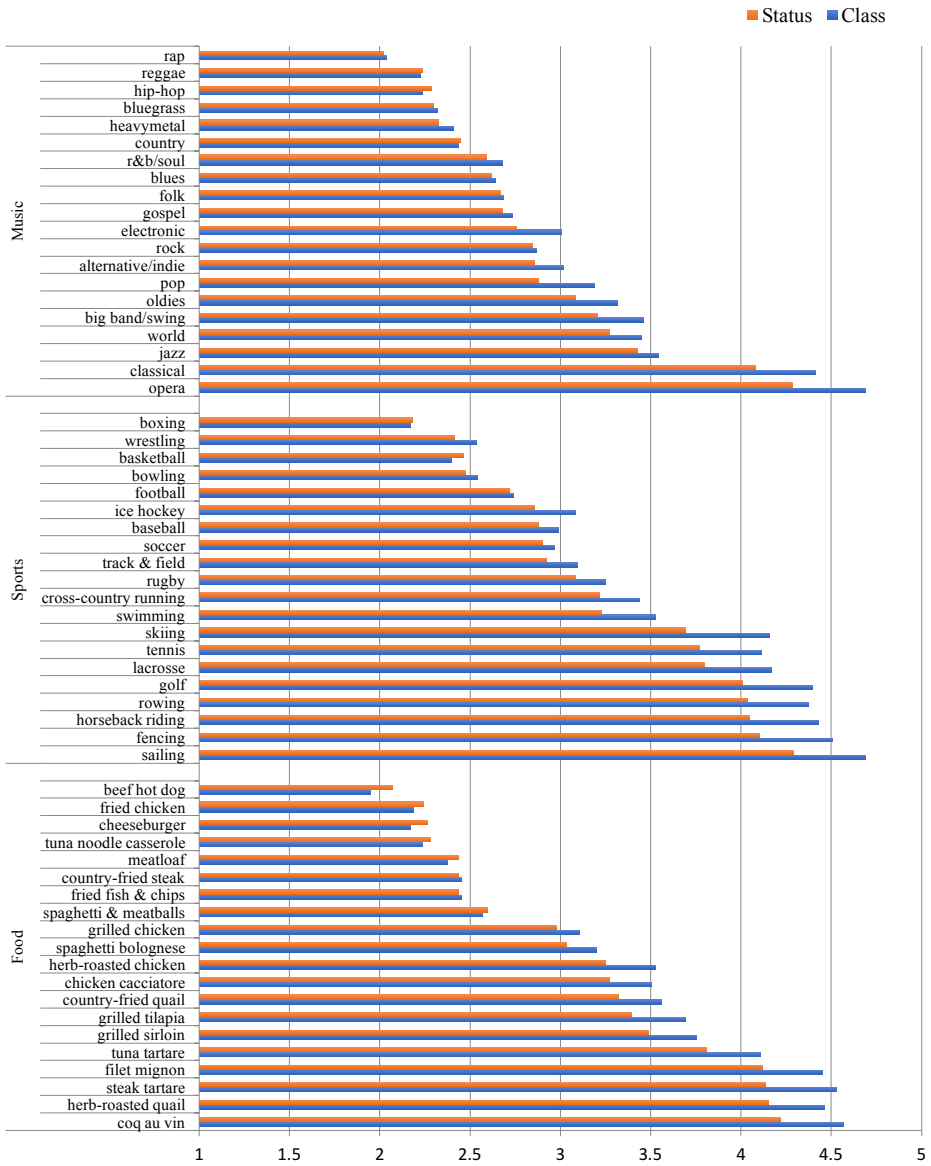

Figure 1: Ratings of Music, Sports, and Food Tastes.

Notes: This figure displays average ratings of status and social class for people with varying music, sports, and food preferences. To measure status, respondents were asked to rate on a five-point scale (1=not at all…5=extremely) how economically successful, well-educated, and prestigious in occupation they perceived the people to be. To measure class, respondents were asked to rate the social class (1=poor…5=upper class) of the same people.

and lowbrow cultural preferences as signals of differing class backgrounds and that there is substantial consensus across demographic groups.

To what extent do these perceived differences in social class reflect objective differences in structural circumstance? In Figure 2, I plot the perceived social class of various music consumers, taken from Figure 1, alongside their educational attainment (% bachelor's degree) and household income (% income of $75,000 and over), taken from the 2012 Survey of Public Participation in the Arts. Notably, the class differences perceived between people with high and low status tastes are greater than the educational and economic differences that exist between them. For example, a person who expresses a taste for opera music is 57% more likely to have a bachelor's degree and 25% more likely to make over $75,000 than a person who expresses a taste for rap. Meanwhile, he/she is attributed a social class position that is, on average, almost four times that of a person who likes rap music.

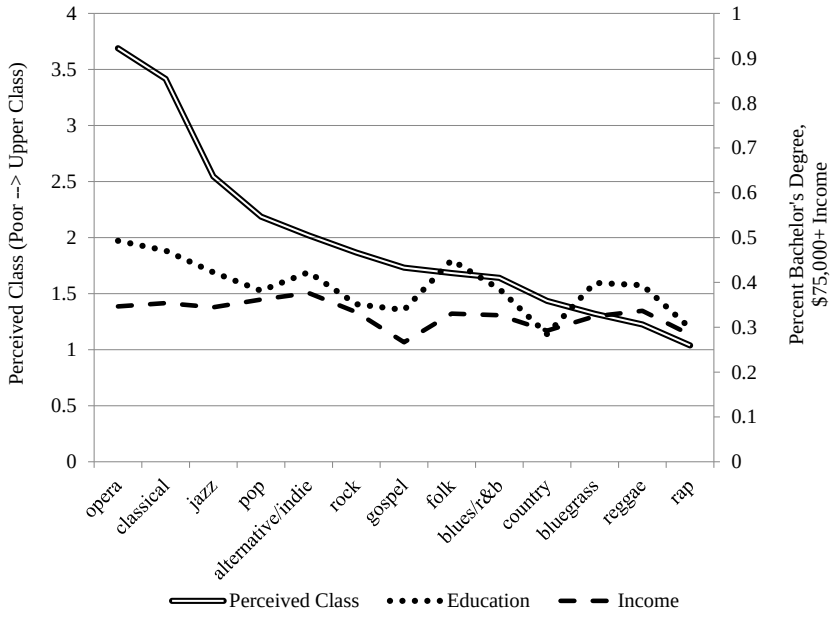

Figure 2: Musical Taste, by Perceived Class, Educational Attainment, and Household Income.
Notes: In this figure, data on the perceived class of music consumers come from a Mechanical Turk survey of 200 U.S. residents. Data on the educational attainment and household income of music consumers come from the 2012 Survey of Public Participation in the Arts.

This discrepancy between the real and perceived probability of an opera music fan belonging to a higher social class than a rap music fan reflects a classic reasoning error, what Kahneman and Tversky (1972) refer to as the representativeness heuristic. Americans likely overestimate the class difference between opera and rap music fans because they do not take into

account the relative size of the class categories—namely, that the upper class is relatively small in size and that a person's likelihood of belonging to the upper class is therefore quite low.

That said, the perceived social class of various music consumers is still strongly correlated with their educational background (correlation=0.74) and moderately correlated with their household income (correlation=0.52). The stronger correlation found between educational attainment and perceived social class is not surprising given that educational attainment is often found to be a better predictor of musical taste than income. Based on these findings, I conclude that Americans systematically perceive indicators of taste as indicators of social class. Furthermore, individuals' subjective, or perceived, differences in social class are associated with objective differences in structural circumstance. By both subjective and objective measures, people with highbrow cultural tastes are more likely to belong to the upper-middle class than people with traditionally lowbrow cultural tastes.

To confirm that the combinations of highbrow versus lowbrow cultural tastes I selected for the audit study and survey experiment also signal class differences when included in résumés, I surveyed an additional 401 U.S. residents on Mechanical Turk. Each respondent was provided a résumé with a male or female applicant name and signals of highbrow or lowbrow cultural taste. Average ratings of social class are presented in Figure 3 below.

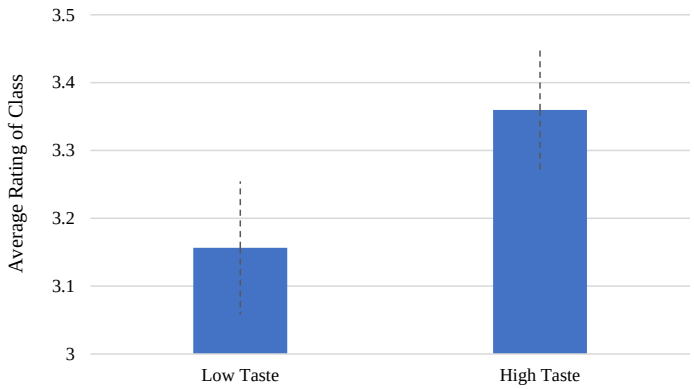

Figure 3: Effect of Signals of Cultural Taste on Perceived Class of Applicants.
Note: The ratings reported above are based on a scale where 1=Poor, 2=Working Class, 3=Lower-Middle Class, 4=Upper-Middle Class, and 5=Upper Class. Dashed lines represent confidence intervals.

Analyses of variance confirm that signals of cultural taste have a statistically significant effect on perceptions of social class. Both applicants are perceived as middle class, but respondents systematically perceive applicants with highbrow cultural tastes as closer to upper-middle class than applicants with lowbrow cultural tastes, who are generally perceived as lower-middle class ($F[1,399] = 8.73$; $p<0.01$). This positive effect of highbrow taste on perceptions of social class is greater for women than men. However, the interaction of cultural taste and gender is not statistically significant ($F[1,397] = 0.82$; not significant).

Audit Study Design

The audit study consisted of a 2 X 2 X 2 mixed factorial design, with two between-subjects factors and one within-subjects factor. The between-subjects factors were job type and applicant gender. Thus, job type and applicant gender varied *between* the résumé pairs submitted to each employer. The within-subjects factor was applicant cultural taste. Thus, applicant cultural taste varied *within* the résumé pairs submitted to employers—every employer received one résumé with signals of highbrow taste and one résumé with signals of lowbrow taste. These factors are outlined in Figure 4 below. Applicant cultural taste was signaled through the undergraduate clubs and personal interests listed in each résumé in a résumé pair. Applicant gender was signaled through stereotypically male and female names listed at the top of the résumés.

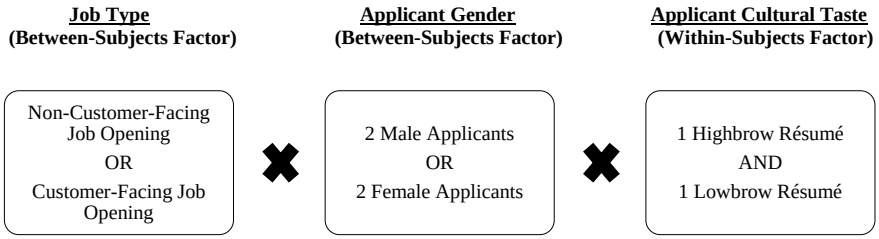

Figure 4: 2 X 2 X 2 Mixed Factorial Design of Audit Study.

The design of the audit study was as follows. I first identified job openings for customer-facing and non-customer-facing jobs on one of the largest employment websites in the U.S. I used a web crawler that collected information for jobs posted over varying 15- to 30-day periods. To search for customer-facing positions, I used the following key words: "customer service" and "sales." To search for non-customer-facing positions, I used the following key words: "administrative assistant" and "clerk." I then reviewed each job opening to confirm that my fictitious applicants had the appropriate educational credentials and years of work experience and to verify that the positions were correctly classified as "customer-facing" or "non-customer-facing."

In reviewing job opening collected by the web crawler, I sometimes came across positions with a title or set of primary responsibilities that spanned both customer-facing and non-customer-facing categories. In these cases, I classified the position based on my reading of the primary responsibilities included in the job description. Customer-facing responsibilities were given greater weight due to my hypothesis that employers would be more likely to discriminate on the basis of cultural signals of class for hiring positions with significant customer interaction.

For example, the job title "Sales Secretary" contains both customer-facing and non-customer facing keywords. However, the main responsibilities of a sales secretary tend to be administrative, so I typically coded sales secretary openings as "non-customer-facing." Meanwhile, "Receptionist" positions were typically coded as "customer-facing," despite their

administrative responsibilities, because one of the primary responsibilities of a receptionist is to serve as the "face" of an organization and greet and interact with clients.

"Sales Clerk" is another example of a common position with both customer-facing and non-customer-facing keywords. To classify each Sales Clerk position, I read through the job description to determine if the title referred to a retail cashier position or an administrative/data-entry position in a sales department. Administrative/data-entry positions in sales departments were included in the study and coded as "non-customer-facing." Retail cashier positions, which were typically listed as a "retail clerk" or "sales clerk" openings, were dropped from the study because the candidates were overqualified. I also eliminated all openings for jobs whose salaries were based only on commission.

Once the list of all eligible job openings had been finalized, I randomly assigned a male or female applicant gender to the pair of (highbrow and lowbrow) résumés being submitted to each job opening. I then randomly assigned applicant names and résumé templates to each résumé in a pair. To signal a male gender, each résumé in a pair was either assigned the name "Scott Ryan" or "Jack Miller." To signal a female gender, each résumé in a pair was either assigned the name "Hannah Ryan" or "Sarah Miller." I also randomly assigned the order in which the two résumés would be submitted. Résumés were submitted a day apart. Approximately 24 hours after submitting the first résumé, I submitted the second résumé for the same opening.

In sum, the order of résumé submission, along with the assignment of applicant gender, applicant name, and résumé template, were randomized and counterbalanced. For example, if, on the first day, I submitted a lowbrow résumé that had been randomly assigned a female applicant gender, the name "Hannah Ryan," and template #1, on the second day I submitted a highbrow résumé with a female applicant gender, the name "Sarah Miller," and template #2. I also included a short cover letter with each résumé, indicating each applicant's interest in the position and a summary of the occupational and educational histories already listed on the résumé.

Between May 2014 and September 2014, I conducted this experiment, which included three complete waves. In total, I submitted 2,096 résumés and cover letters in response to 1,048 job openings. The average rate of callback across all four cities was 11.5% (15% in Boston, 12.8% in Chicago, 7.6% in Los Angeles, and 10.1% in New York).

Survey Experiment Design

The survey experiment consisted of a 2 X 2 X 2 X 2 mixed factorial design, with three between-subjects factors and one within-subjects factor. As in the audit study, the between-subjects factors were job type and applicant gender, with the addition of firm status. Job type, applicant gender, and firm status were randomly assigned to respondents and varied between the résumé pairs assigned to respondents. The within-subjects factor was applicant cultural taste. Applicant cultural taste varied within the pair of résumés assigned to respondents—every respondent received a résumé with signals of highbrow taste and a résumé with signals of lowbrow taste. These factors are outlined in Figure 5.

The survey experiment was web based and distributed to a nationwide sample of 1,428 hiring managers drawn from a pre-existing panel of HR/Personnel Services managers maintained by the research company Research Now. Respondents qualified for the survey if they identified as (1) full-time or part-time employees, (2) who make decisions regarding whether or not to interview or hire job applicants, (3) were using a desktop computer, laptop, netbook, or tablet computer to take the survey, and (4) provided informed consent. I excluded mobile device users

to reduce the risk of receiving incomplete or low quality responses. Long surveys with matrix-type or open-ended questions do not always display well on a mobile device and are more difficult to complete on a mobile device.

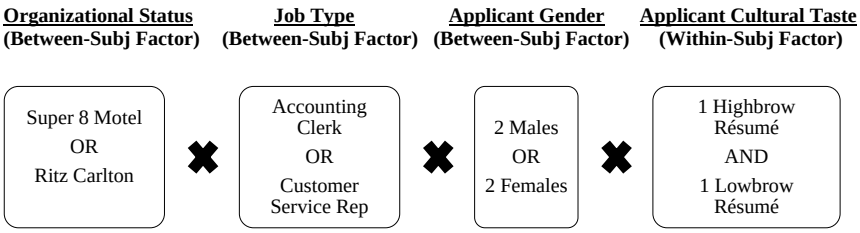

Figure 5: 2 X 2 X 2 X 2 Factorial Design of Survey Experiment.

In partnership with Research Now, I distributed the survey experiment in June 2014. 39,541 Research Now panel members were contacted via email and 4,292 clicked on the link to begin the survey. Of those 4,292 respondents, 2,241 were excluded from the survey because they did not meet at least one of the screening criteria detailed above. Of the remaining 2,051 respondents, 623 dropped out and did not complete the survey. The remaining 1,428 respondents completed the survey.

At the beginning of the survey, respondents were asked to evaluate two applicants seeking an entry-level position at a hotel. Each respondent was randomly assigned to one of eight possible hiring conditions: (1) the hiring organization was high in status (Ritz-Carlton Hotel) or low in status (Super 8 Motel), (2) the hiring position was either high in customer contact or low in customer contact, and (3) the gender of the two applicants was either male or female. Respondents were then provided two résumés and they were asked to rate their likelihood of recommending each applicant for an interview. They were also asked to rate each applicant on a series of characteristics related to competence, warmth, and polish. As in the audit study, the order in which these résumés were shown and the résumé formats that were used were randomized and counterbalanced.

Résumé Design

Two résumé templates were used and randomly assigned to each highbrow and lowbrow résumé in a pair. The templates varied moderately in their occupational histories, educational backgrounds, font, wording, and formatting. I used two templates to ensure that the manipulation of cultural taste was not detected. Both templates listed Bachelor of Arts degrees from small and similarly ranked liberal arts colleges with high acceptance rates in Michigan or Ohio. They also listed approximately two years of employment experience as an administrative or office assistant at a catering company or ticket agency. Both listed a series of computer skills, including proficiency in MS Word, Adobe, QuickBooks, and Mac OSX.

Template #1 used Times New Roman font and listed time intervals for each position in a column on the right-hand side of the résumé. Template #2 used Cambria font and listed time

intervals for each position in a column on the left-hand side of the résumé. I also varied the wording of job descriptions and subheadings across the two templates as well as the contact information each provided. Each template included a home address and both templates' addresses were within a few blocks of one another in the same zip code. I chose zip codes that were relatively affordable for the city and were centrally located within the city, since a short commute time was listed as qualification for some jobs. Despite these differences between the two templates, the assignment of a template to any given résumé had no independent effect on the résumé's outcomes. The assignment of résumé templates to highbrow and lowbrow résumés was also randomized and counterbalanced to ensure that it was not confounded with the assignment of cultural taste.

Lastly, Google email accounts and Google Voice accounts were created for each applicant name in each city so that employers could contact the applicant via phone or email and I could track employers' responses to the résumés. Each email address included the applicant's name followed by three random numbers. In each city, all phone numbers included the same area code. For each phone number, I set up a voicemail account with an outgoing message. I developed four short scripts and hired male and female U.S. residents on Mechanical Turk to provide a recording of himself or herself reading one of these scripts. I then set up active voicemail accounts using the highest quality male recordings for "Brad Miller" and "Scott Ryan" and the highest quality female recordings for "Sarah Miller" and "Hannah Ryan." All voice actors had American English, non-Southern accents.

References

Fiske, Susan T., Amy J. Cuddy, Peter Glick, and Jun Xu. 2002. "A Model of (Often Mixed) Stereotype Content: Competence and Warmth Respectively Follow from Perceived Status and Competition." *Journal of Personality and Social Psychology* 82(6): 878-902.

Kahneman, Daniel and Amos Tversky. 1972. "Subjective probability: A judgment of representativeness." *Cognitive Psychology* 3(3): 430-454.

Paolacci, Gabriele, Jesse Chandler, and Panagiotis G. Ipeirotis. 2010. "Running experiments on Amazon Mechanical Turk." *Judgment and Decision Making, 5*(5): 411-419.