Conjoint Experiment with a Multinomial Outcome: Americans' Stereotypes about Immigrants from Five World Regions

Kirill Zhirkov

University of Virginia

[Draft. Please don't cite or circulate]

Abstract

Currently, researchers use only two kinds of outcomes in conjoint analysis: binary choices and rating scales. In this paper, I describe a parametric version of conjoint analysis based on multinomial logistic regression. It departs from the original nonparametric strategy but allows estimating average marginal component effects (AMCEs) with more than two potential choices in a conjoint task. I apply the multinomial conjoint design to a substantively interesting topic: Americans' stereotypes about immigrants from different world regions. Using data from a nationally diverse sample, I demonstrate that immigrants from regions other than Europe are believed to speak worse English and, with the exception of Asians, break the rules more often. These results show why non-European immigrants can face more opposition in the United States, even if meritocratic norms with regard to immigrant admission are shared by most Americans.

Keywords: conjoint analysis, immigration, multinomial modeling, stereotypes, survey experiments

Conjoint survey experiments, despite being introduced to the political methodology toolkit less than a decade ago (Hainmueller, Hopkins, and Yamamoto 2014), are increasingly popular in the discipline. In conjoint-experimental designs, respondents are presented with potential choice options described using several attributes with randomized values. Observed respondents' choices allow researchers to simultaneously assess the independent causal effects of different factors on the decisions of interest, such as vote choices. Recent studies have demonstrated external validity of conjoint experiments (De la Cuesta, Egami, and Imai 2022; Hainmueller, Hangartner, and Yamamoto 2015; Jenke et al. 2021), showed their resilience to problems like satisficing (Bansak et al. 2018, 2021), developed strategies to estimate effects' heterogeneity (Horiuchi, Smith, and Yamamoto 2018; Zhirkov 2022), and proposed improvements in results' interpretation and presentation (Abramson, Kocak, and Magazinnik 2022; Bansak et al. 2022; Ganter 2021; Leeper, Hobolt, and Tilley 2020).

One aspect of the conjoint design that is surprisingly absent from existing methodological literature is the nature of outcome variables. Up to now, conjoint analyses in the discipline have almost exclusively used only two kinds of dependent variables: binary and interval. One particularly important kind of outcomes that can be used in conjoint experiments, but have not been addressed in methodological literature so far, are multinomial variables.[1] They represent a straignforward extension of the binary choice-based conjoint design but allow researchers to define multiple response options (beyond just "chosen" vs. "not chosen") that do not have to be meaningfully ordered.

Multinomial outcomes can be useful in many applications of conjoint analysis, especially those that deal with stereotypes rather than preferences (Flores and Schachter 2018; Myers, Zhirkov, and Lunz Trujillo 2022). Consider a recent conjoint study on

[1] Throughout the paper, I use the term "multinomial" to refer to polytomous nominal (unordered categorical) random variables: ones with more than two potential outcome values that do not have a meaningful rank order. Statistically, this usage is somewhat imprecise because it refers to a specific (single-trial) case of the more general multinomial distribution. I use this terminology to keep the parallel between the outcome and the estimation method (multinomial logistic regression).

stereotypes about U.S. parties (Goggin, Henderson, and Theodoridis 2020). With a binary choice outcome, such experiments are limited to contrasting stereotypes about Democrats and Republicans. The ability to use a multinomial outcome, instead, would allow researchers to asses stereotypes about independents—who are known to be qualitatively different from partisans (Klar and Krupnikov 2016). It also opens additional opportunities, such as measuring stereotypes about factions within the two parties (Clarke 2020; Groenendyk, Sances, and Zhirkov 2020). For instance, one can be able to explore the stereotypes about democratic socialists and MAGA conservatives as well as moderate Democrats and Republicans in a single conjoint experiment.

In this paper, I present a strategy to estimate the quantities of interest in conjoint experiments that use multinomial outcomes. It is based on the well-known multinomial logistic regression that can be easily implemented in all modern statistical software. I discuss how this strategy differs from the original conjoint approach: most importantly, it requires making parametric assumptions about the error structure. These assumptions, however, do not interfere with the main benefit of conjoint experiments: the causal interpretation of estimated effects. I also demonstrate a practical application of the multinomial conjoint design with an emphasis on effects' interpretation and presentation.

My application concerns a substantively important case from the literature on politics of immigration: Americans' stereotypes about immigrants from different world regions. An influential experimental study demonstrates the "hidden consensus" in immigration preferences: U.S. respondents, independently of partisanship, choose to admit immigrants who can contribute to the national economy (Hainmueller and Hopkins 2015). Follow-up experiments, however, show that these preferences are not completely origin-neutral as similarly skilled immigrants from Latin America and Europe are evaluated differently (Newman and Malhotra 2018). Overall, Americans seem to have hierarchical preferences regarding immigrant origins with Europeans at the top (Zhirkov

2021a). However, existence of such a hierarchy is compatible with widespread meritocratic preferences if immigrants from some regions are seen as having the desirable characteristics whereas immigrants from other regions are not.

In an original conjoint experiment, I explore Americans' stereotypes about immigrants from five world regions: Africa, Asia, Europe, Latin America, and the Middle East. Interestingly, I find no differences in stereotypes about education levels. At the same time, Europeans are believed to speak better English than immigrants from other world regions and, with the exception of Asians, be more law-abiding. This result shows why Americans can prefer European immigrants even if their preferences on immigration are broadly meritocratic. More generally, my analysis demonstrates how conjoint experiments with multinomial outcomes can help to answer important questions in political science.

Multinomial Conjoint Analysis: Motivation

Despite prominent advances in various aspects of conjoint-experimental methodology, conjoint designs currently used in political science almost exclusively rely on the same two types of outcomes as the original study that introduced the method to the discipline: binary, also known as choice outcomes, and interval, also known as rating outcomes (Hainmueller, Hopkins, and Yamamoto 2014).[2] In choice-based conjoint experiments, respondents are presented with two or more alternatives and asked to choose the one they prefer. With rating-based outcomes, respondents give a numerical score to each presented alternative that corresponds to their preference strength. The popularity of conjoint designs with these outcomes is not surprising. Obtaining the main causal quantities of interest, the average marginal component effects (AMCEs), does not require making any parametric assumptions. As a result, AMCEs can be estimated using an OLS regression, a simple and

[2] One important exception concerns conjoint experiments with ordered categorical outcomes (Hanretty, Lauderdale, and Vivyan 2020; Sen 2017). However, existing studies simply apply an ordered logistic regression to such outcomes without discussing implications of the departure from the nonparametric setup and the OLS-based estimation strategy.

easily interpretable method. Also, most substantive topics that political scientists study using conjoint experiments are compatible with either binary or interval outcomes.

However, answering some research questions in political science requires conjoint designs with more complicated outcomes. Those are cases when the range of choices of interest cannot be reduced to the "choose vs. not choose" dichotomy. Importantly, such scenarios are qualitatively different from the simple increase of the number of options in conjoint experiments. For instance, to more realisitically explore behavior in multiparty elections respondents can be offered three or four candidate profiles to choose from, instead of two. Still, there exist substantively interesting phenomena that would require changing the outcome structure rather than simply increasing the number of options to choose from. Even in a two-party election, one can abstain (in effect, deliberately refuse to make a choice) rather than vote for one of the candidates.

This problem is particularly relevant for conjoint experiments that measure stereotypes rather than preferences. After all, standard conjoint experiments measure the direction and (relative) strength of respondents' tastes. For instance, do voters prefer candidates who promise to increase spending or cut taxes, and how does this effect compare to position on abortion? Answering this question does not require knowing the specific party someone would vote for—if anything, introduction of party labels may contaminate the result by introducing additional considerations such as partisan identity. As a result, conjoint experiments with nameless candidates, whom respondents can choose or rate, is a straightforward way to elicit preferences even in multiparty democracies.

In the case of conjoint experiments that explore stereotypes the logic is opposite. This is an increasingly popular application of the method (Flores and Schachter 2018; Goggin, Henderson, and Theodoridis 2020), that has been recently validated on the example of stereotypes about welfare recipients in the United States (Myers, Zhirkov, and Lunz Trujillo 2022). In such designs, respondents are also shown profiles of hypothetical individuals with several randomized attributes—but asked to categorize them rather than choose the preferred option. And just as observed choices in standard conjoint analysis are used to infer preferences, observed guesses about category memberships in classification conjoint experiments allow researchers to measure stereotypes, or associations of certain attributes with different social groups (Hilton and Von Hippel 1996).

In these applications, the ability to use multinomial outcomes is particularly helpful. Since stereotypes are always group-specific, conjoint experiments aimed at measuring them have to employ real-world group labels as the outcome in categorization tasks. And the number of groups or categories of interest in such experiments are often greater than two. For instance, a proper investigation of racial stereotypes in the United States would require categorizing presented profiles as white, Black, Asian, Hispanic, an so on. The binary choice outcome, instead, limits the researcher to a simple—and, probably, less natural and less informative—comparison between, say, whites and nonwhites.

These design requirements are not readily compatible with any of the standard outcomes commonly used in conjoint experiments. As discussed above, they obviously do not work with binary choice outcomes. The same is true for ordinal outcomes, which have been used in preference-based conjoint experiments (Hanretty, Lauderdale, and Vivyan 2020; Sen 2017), because group labels cannot be meaningfully ordered. Interval outcomes can work in theory, but their application in practice is related to several problems. First, they are simply inefficient: a respondent has to rate each presented profile in terms of likeliness of belonging to each of the groups or categories of interest. It means that a conjoint experiment that features, say, ten profiles and five potential outcome categories will include the total of 50 rating questions. And this number grows geometrically as the numbers of rated profiles and/or social groups of interest increase.

Second, interval outcomes are subject to a common problem in survey methodology known as differential item functioning (DIF; Aldrich and McKelvey 1977; Brady 1985; Hare et al. 2015). Here, it means that respondents do not use the Likert-style rating scale similarly: for instance, some can give consistently higher likeliness scores to all groups. As a result, individual rating scores are incomparable, and their aggregation to calculate AMCEs for stereotypes about different groups can be problematic.[3] Therefore, an ideal outcome in such conjoint experiments should not rely on respondents' ability to estimate the absolute likeliness of a profile's belonging to certain groups or categories. Instead, relative likeliness should be inferred from observed categorizations.

Multinomial Conjoint Analysis: Formalization

I suggest that conjoint analysis can be easily extended to cases in which respondents are asked to categorize profiles into multiple mutually exclusive classes. Most intuitively, such designs can be applied to tasks that involve classifying hypothetical individuals into more than two social groups. However, conjoint experiments with multiple unordered choices can be encountered in other applications as welll: for instance, choosing among potential actions in a hypothetical international crisis. Outcomes of such conjoint tasks are familiar to methodologists: they are described using the categorical distribution that has well-studied statistical properties. Moreover, there exists a popular method to predict probabilities of these outcomes given a set of predictors (in this case, conjoint attribute values): multinomial logistic regression.

A common motivation for the multinomial logistic regression is the latent utility model. Respondents, when faced with a choice from several options, select the one associated with the highest utility. Even though these utilities cannot be calculated directly, relative contributions of different factors can be inferred from observed choices.

[3] The incomparability problem can be addressed using a method like anchoring vignettes (King et al. 2004; King and Wand 2007). However, it will not solve the inefficiency problem—on the contrary, introduction of anchoring vignettes will further increase the conjoint task length.

The same logic—I call it the latent likeliness model—can be easily applied to problems that involve categorizations rather than choosing the most preferred option.

Consider a conjoint experiment in which respondents indexed $i \in \{1, \dots, I\}$ are presented with profiles indexed $j \in \{1, \dots, J\}$. Each profile is described in terms of attributes indexed $k \in \{1, \dots, K\}$. These attributes have potential values indexed $h \in \{1, \dots, H^k\}$, with the total numbers of potential values varying across attributes. Respondents are asked to categorize each presented profile into mutually exclusive groups indexed $m \in \{1, \dots, M\}$. Let y_{ij} be a multinomial outcome reflecting observed categorization of profile j by respondent i. Let x_{ijkh} be a binary indicator that equals 1 if the value of attribute k of profile j presented to respondent i is h and 0 otherwise.

Assume that each respondent i has a set of unobserved real-valued scores, y^*_{ijm}, such that each score represents the likeliness of profile j to belong to group m and depends on the profile's attribute values:

$$y^*_{ijm} = \alpha_{jm} + \sum_{k=1}^{K} \sum_{h=2}^{H^k} \beta_{khm} x_{ijkh} + \varepsilon_{ijm} = l_{ijm} + \varepsilon_{ijm}, \tag{1}$$

where α_{jm} is the group-specific intercept, β_{khm} is the contribution of value h of attribute k to the stereotype about group m, ε_{ijm} is the error term, and l_{ijm} is the non-error part of the likeliness function. To achieve identification, the effect of the first value of each attribute is not estimated ($\beta_{k1m} = 0$).[4]

When completing the categorization task, for each presented profile the respondent selects the group associated with the greatest likeliness value:

$$y_{ij} = m \quad \iff \quad y^*_{ijm} = \max(y^*_{ij1}, \dots, y^*_{ijM}). \tag{2}$$

Then, if the error terms ε_{ijm} are independent and have the standard Gumbel distribution, also known as the standard type-I generalized extreme value distribution, the expression

[4] Another constraint, necessary for the estimation of multinomial logistic regression, is that all coefficients for the baseline group are assumed to be zero ($\alpha_{j1} = 0$ and all $\beta_{kh1} = 0$).

for probability that respondent i categorizes profile j as belonging to group m becomes (McFadden 1973):

$$\Pr(y_{ij} = m) = \frac{\exp(l_{ijm})}{\sum_{m'=1}^{M} \exp(l_{ijm'})}. \tag{3}$$

This expression is effectively a statement of the multinomial logistic regression, a common statistical model whose parameters can be obtained using maximum likelihood estimation implemented in all popular software packages.

The model setup presented above has some differences from the standard conjoint design (Hainmueller, Hopkins, and Yamamoto 2014). The most fundamental of them is the necessity to make parametric assumptions about the error structure. If error distributions deviate from the assumed Gumbel family, parameter estimates will be biased. At the same time, this is an extremely common statistical assumption made in all observational studies that employ multinomial logistic regression and/or its extensions.

The other aspects of multinomial conjoint analysis are less consequential—they concern design of the task as well as presentation and interpretation of the results. First, since each profile has to be categorized into one of the multiple mutually exclusive groups, it is most intuitive to present only one conjoint profile per survey page. As a result, such tasks can take more time than standard conjoint experiments with paired profiles. However, it can ultimately be benefitial since paired forced-choice conjoint designs can bias the results due to composition effects (Ganter 2021).

Second, as it is the case with multinomial logistic regressions that use observational data, interpretation and presentation of results are not straightforward. Effects in multinomial conjoint experiments, besides being estimated for each group of interest, are no longer expressed in terms of choice probability. Fortunately, political methodologists has developed rather detailed guidelines for the presentation of the other quantities of interest, such as predicted probabilities (King, Tomz, and Wittenberg 2000). When applied to conjoint experiments specifically, these quantities are known as marginal means (Leeper, Hobolt, and Tilley 2020). However, since the number of estimated marginal means equals the number of attribute values times the number of potential outcome categories, it is advisable to keep the number of values per attribute to a reasonable minimum. Finally, multinomial conjoint design does not allow obtaining individual marginal component effects (IMCEs) since their estimation is based on the nonparametric assumption and uses the OLS method (Zhirkov 2022).

Despite these differences from the standard setup, experiments with multinomial dependent variables keep most of the important benefits of conjoint analysis. First and foremost, since conjoint attribute values are independently randomized under the full control of the researcher, estimated effects have causal interpretation (Hainmueller, Hopkins, and Yamamoto 2014). Second, these effects simultaneously measure the two most meaningful aspects of preferences or stereotypes under consideration: direction and intensity (Bansak et al. 2022, Myers, Zhirkov, and Lunz Trujillo 2022).[5] Finally, additional methodological benefits of conjoint experiments—that include external validity (Hainmueller, Hangartner, and Yamamoto 2015; Jenke et al. 2021), resilience to satisficing (Bansak et al. 2018, 2021), and reduction of social desirability bias (Horiuchi, Markovich, and Yamamoto 2022)—should also apply to designs with multinomial outcomes.

Application: Stereotypes about Immigrant Origins

Here, I apply multinomial conjoint analysis to a substantively interesting research question: Americans' stereotypes about immigrants from different world regions. It makes a good application of the proposed design and estimation for several reasons. First, the combination of necessity to explore multiple stereotype dimensions simultaneously and

[5] As recently demonstrated, under certain conditions conjoint AMCEs can miss the preference or stereotype held by the majority of the population (Abramson, Kocak, and Magazinnik 2022). However, individual-level distributions of conjoint effects show that the directions of estimated AMCEs, as soon as they are significant, and median preferences coincide (Zhirkov 2022).

sensitivity of immigration as a topic (that leads to social desirability concerns) makes the conjoint experiment the optimal design choice. Second, this conjoint experiment measures stereotypes rather than preferences—meaning that the response options should feature real world regions as the potential categorization choices. Finally, it has to consider multiple regions of origin (such as Africa, Asia, Europe, Latin America, and the Middle East), thus being incompatible with standard binary choices.

Besides its suitability for purely methodological reasons, such an experiment can contribute to the literature on politics of immigration. Influential experimental studies suggest that Americans, as well as citizens of other industrial democracies, have broadly meroticratic preferences on immigration. Specifically, they are ready to admit immigrants who have valuable skills, speak good English, and follow the rules—whereas country of origin has little effect (Hainmueller and Hopkins 2015; also see Valentino et al. 2019). At the same time, Americans do differentiate between more and less desirable immigrant groups (Zhirkov 2021a). Most prominently, attitudes toward Hispanics impact whites' opinions on immigration and featuring immigrants from Latin America in the media boosts nativism (Brader, Valentino, and Suhay 2015; Branton et al. 2011; Valentino, Brader, and Jardina 2013). According to more recent evidence, immigrants from the Middle East can be another group that provokes particularly strong opposition (Konitzer et al. 2019).

One potential explanation for the coexistense of meritocratic preferences and an origin hierarchy concerns stereotypes about immigrants from different world regions. Immigrants from certain places can be preferred to others if their origins are associated with desirable attributes. However, the existence of divergent stereotypes about immigrants from different parts of the world has not been directly established. There exists research on stereotypes about immigrants vs. natives (Blinder 2015; Lutz and Bitschnau 2022; Zhirkov 2021b), but not about immigrants from specific countries or regions. One rare and important exception is the finding that white Americans seem to associate immigrants from Latin America with illegal status—but the underlying study investigates stereotypes about legality, not origins (Flores and Schachter 2018). Evaluations of Hispanic immigrants also more strongly depend on their skills and perceived transgressions (Newman and Malhotra 2018; Hartman, Newman, and Bell 2014), but these findings provide only indirect evidence for the presence of corresponding negative stereotypes. In addition, there is almost no research on stereotypes about Asian immigrants, another group that is important for Americans' opinion on immigration (Citrin et al. 1997; Malhotra, Margalit, and Mo 2013), as well as about immigrants from Africa and the Middle East.

In this paper, I implement a multinomial conjoint experiment to explore Americans' stereotypes about immigrants from five world regions: Africa, Asia, Europe, Latin America, and the Middle East. Its main goal is to understand whether there exist systematic differences in perceived attributes of immigrants from different regions. Methodologically, I use this design to demonstrate how researchers can estimate and present the quantities of interest in conjoint experiments with multinomial outcomes—and how such experiments can be used to address substantively important questions in political science.

Experimental Design

Respondents for the online survey-experimental study were U.S. adults recruited in August 2022 on the Lucid Theorem platform. It has been shown to yield samples comparable to national probability benchmarks, such as the American National Election Studies, in terms of the key demographics (Coppock and McClellan 2019). The final sample in my experiment consisted of 1,979 respondents. Its characteristics were the following. Mean age was 45.1 years. Gender ratio was 48.5 male to 51.5 female. College education was reported by 45.1% of respondents. Median income was $40,000 to $44,999. In terms of race and ethnicity, 68.6% of respondents identified as non-Hispanic whites. Finally, 36.3% of the sample were Democrats, 30.7% were Republicans, and 33% were independents.

Table 1. Attributes for immigrant profiles in conjoint experiment

Attribute	Values
Age	*Young:* 20–39
	Older: 40–59
Gender	Male
	Female
Education	*Less than college:* Elementary school, Middle school, High school
	Some college or higher: 2-year college, 4-year college, Graduate degree
English proficiency	Fluent English
	Broken English
Government benefits	*Benefits:* Food stamps, Medicaid, Supplemental income, Housing assistance
	No benefits
Prior trips to U.S.	*Has a violation:* Overstayed visa, Unauthorized
	No violation: No, On a visa
Police record	*Has a record:* Assault, Drug possession, Theft
	No record

Note. Age values (in years) were randomly chosen from the specified intervals

In the conjoint experiment, respondents were presented with profiles of hypothetical immigrants and asked to guess which world region each immigrant came from. There were five categorization options for each profile: Africa, Asia, Europe, Latin America, and the Middle East. Each respondent was asked to categorize the total of 20 profiles, and only one profile was shown per a separate survey page.

Each profile was described in terms of seven attributes: age, gender, education, English proficiency, government benefits, prior trips to U.S., and police record. The order of attributes was randomized between respondents but kept constant for each individual respondent. All attribute values were randomized independently with uniform distributions (i.e., all potential values of each attribute had equal probabilities of being presented). The only exceptions were government benefits and police record. For these attributes, each profile had a 50% probability of being described as having no benefits or no record, respectively. For the remaining 50% of cases, the specific labels for benefits and police

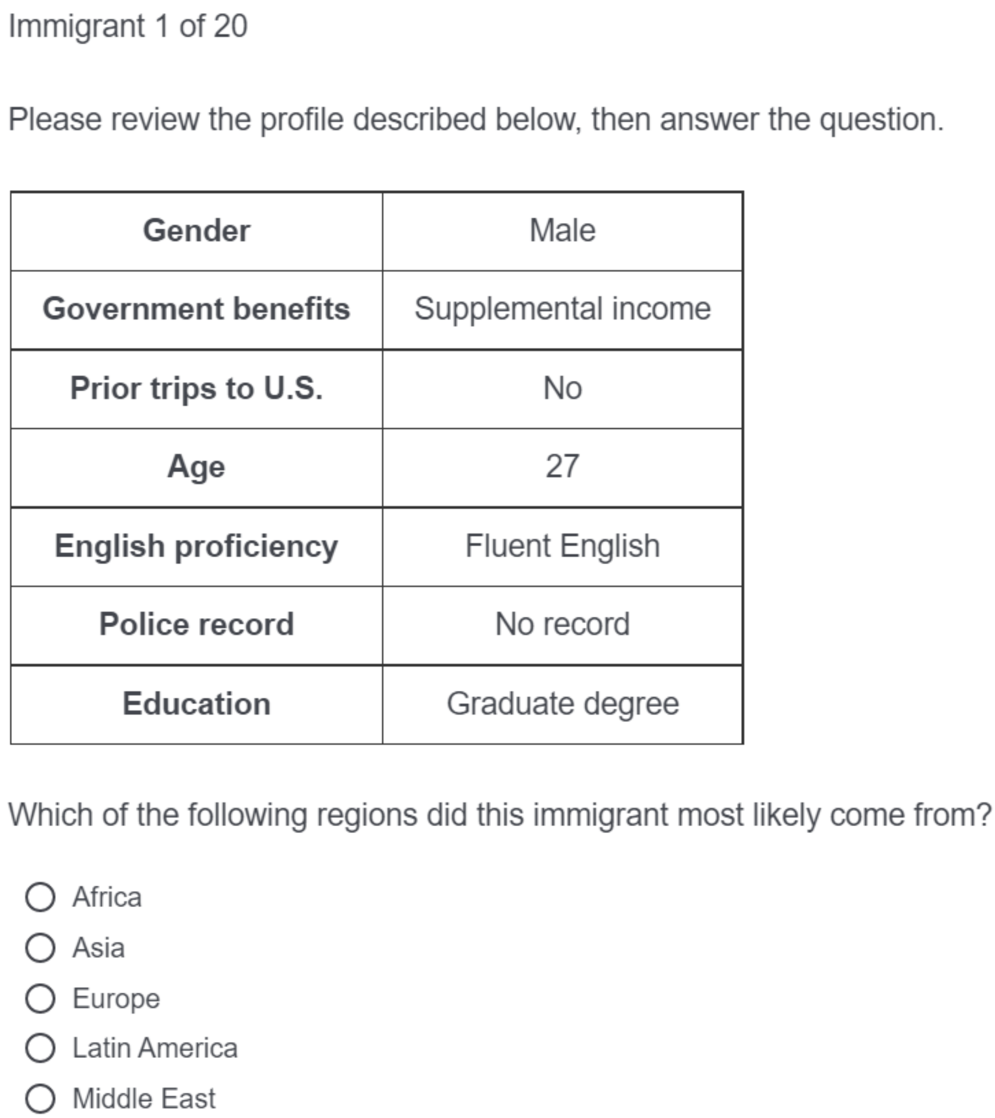
Immigrant 1 of 20

Please review the profile described below, then answer the question.

Gender	Male
Government benefits	Supplemental income
Prior trips to U.S.	No
Age	27
English proficiency	Fluent English
Police record	No record
Education	Graduate degree

Which of the following regions did this immigrant most likely come from?

- Africa
- Asia
- Europe
- Latin America
- Middle East

Figure 1. Sample conjoint profile as shown to respondents

records had euqal probabilities of being presented. See Table 1 for the full list of profile attributes potential values and Figure 1 for a sample profile as presented to respondents.

The survey also included a battery of items that measured respondents' level of ethnocentrism (Bizumic and Duckitt 2012). See Table S1 in Supplementary Material for the questions and response options.

Results

Respondents in the conjoint experiment categorized the total of 39,568 hypothetical immigrant profiles.[6] Of them, 16.2% were guessed as coming from Africa, 15.8% from Asia, 25% from Europe, 28.2% from Latin America, and 14.8% from the Middle East. I start the analysis from estimating a multinomial logistic regression that predicts profile categorizations on the basis of specific attribute values. For theoretical reasons (the importance of the contrast between European and non-European immigrants), I choose Europe as the baseline region of origin. I also aggregate some attribute values in order to keep the reasonable number of estimated effects. Following the standard practice of conjoint analysis (Hainmueller, Hopkins, and Yamamoto 2014), standard errors are clustered on the level of respondents.

Coefficients from the multinomial logistic regression are presented in **??**. Presented effects measure the direction and (relative) strength of group-specific stereotypes. The most consistent of them is that immigrants who speak broken (as opposed to fluent) English are uniformly more likely to be categorized as non-Europeans. Some other effects, however, deviate from the pattern of non-European immigrants being seen more negatively than European ones. For instance, immigrants from Asia are seen as less likely to both receive government benefits and have police record than Europeans. An ability to uncover such nuanced effects across different origin groups is an important benefit of the multinomial outcome compared with a binary one.[7] Surprisingly, respondents do not really use education levels (the main proxy for skill sets) when categorizing immigrant origins—suggesting that the corresponding stereotypes are weak or nonexistent. If anything, immigrants from Latin America are seen as more educated than Europeans.

[6] Some respondents ended up categorizing fewer than 20 profiles.

[7] A formal likelihood-ratio test confirms presence of significant differences in effects across the non-European origins and, thus, necessity for a multinomial specification ($\chi^2_{21} = 436.1$, $p < .001$).

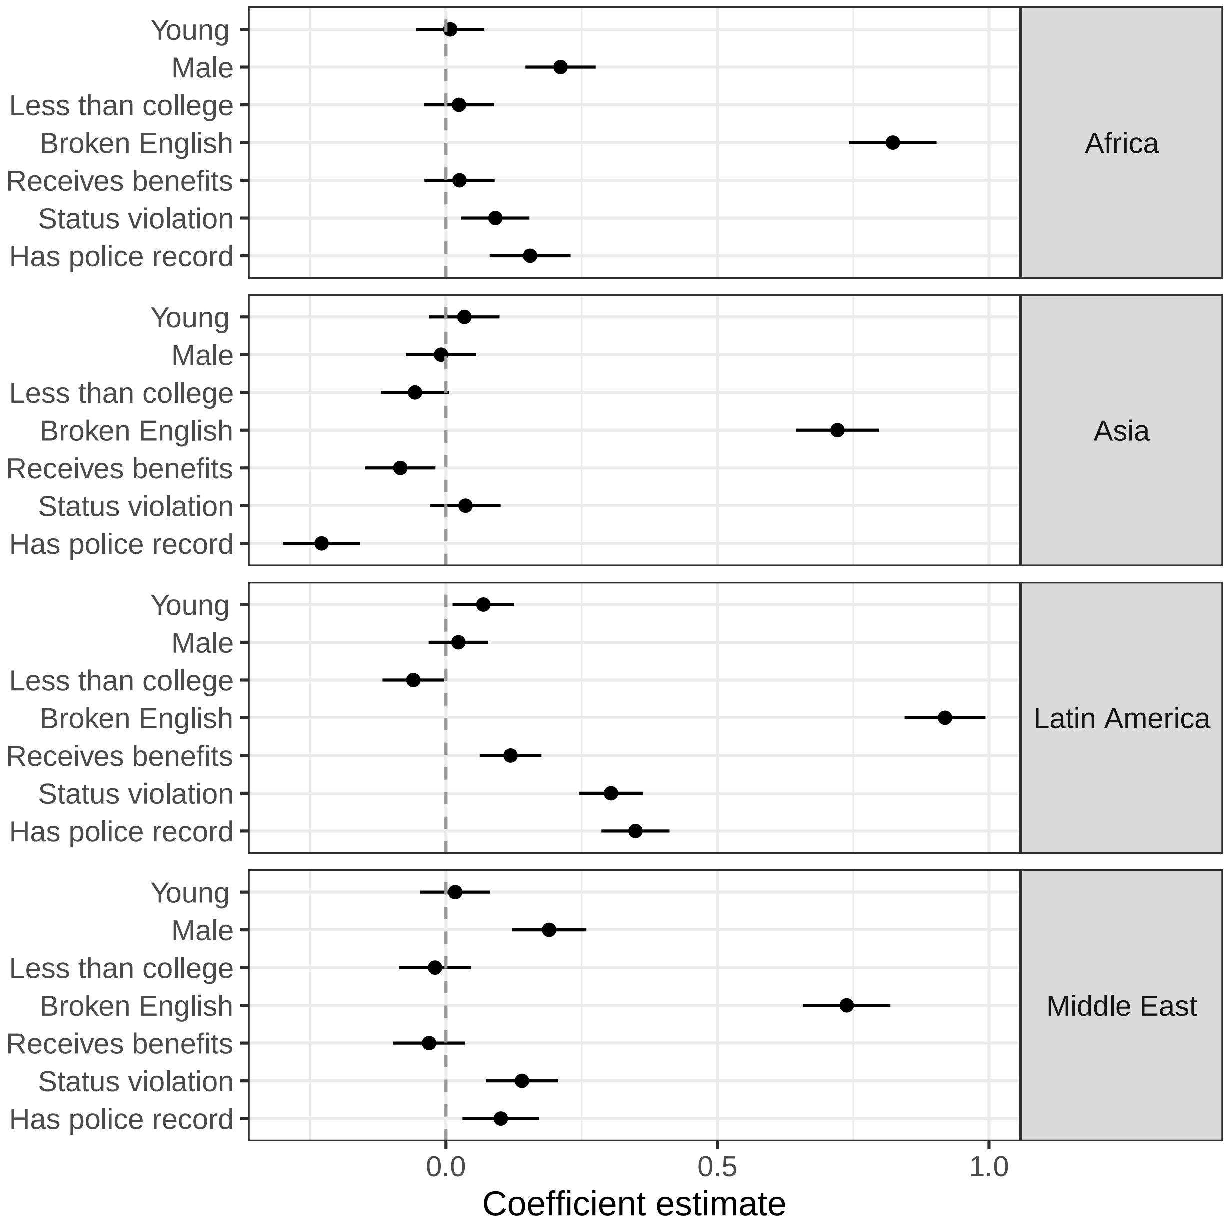

Figure 2. Coefficients from the multinomial logistic regression
Note. $N = 39,568$. Europe is the baseline outcome category.

Also, overall stereotypes do not differ across respondents' levels of ethnocentrism (see Figure S1 in Supplementary Material).

This presentation of results from multinomial conjoint experiments has a problem related to the nature of the estimated regression coefficients. They are not expressed in

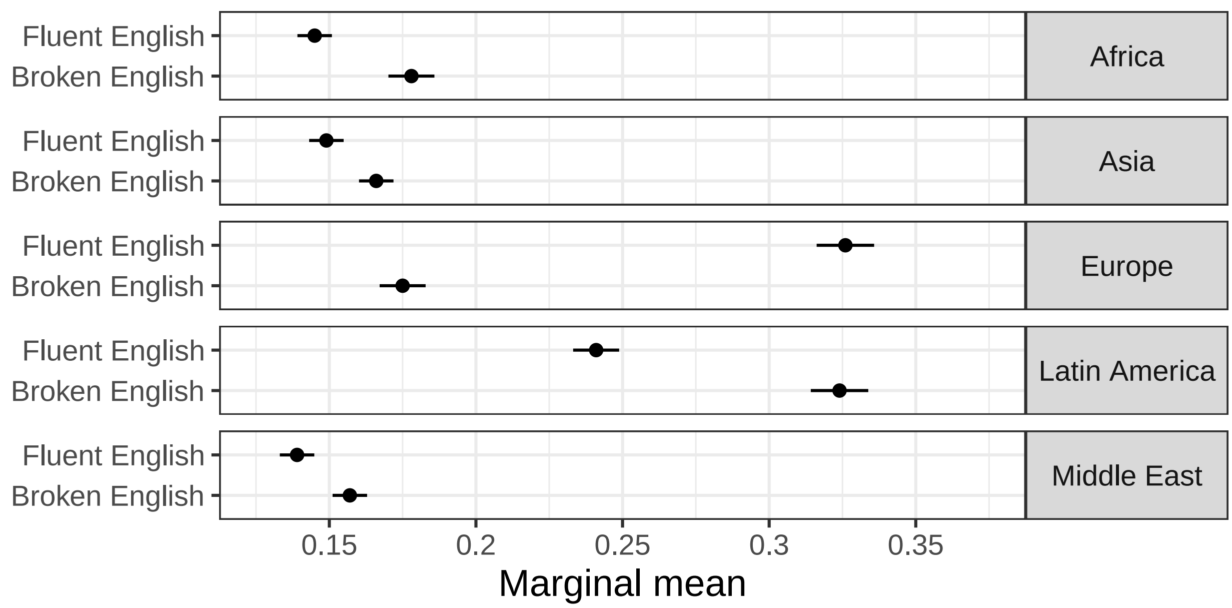

Figure 3. Marginal means for the English proficiency attribute

terms of categorization probabilities, which are the quantities of theoretical interest. Fortunately, there is a class of quantities other than treatment effects that can be used to present the results of a conjoint experiment in an intuitive way: marginal means (Leeper, Hobolt, and Tilley 2020).

At the same time, marginal means still have to be computed for each combination of an attribute value and a categorization option. Therefore, for space considerations Figure 3 presents only marginal means for English proficiency as the most consequential attribute.[8] Unlike regression coefficients, marginal means highlight the effects of different attribute values on predicted probabilities of specific categories. They show that, when presented on the probability scale, the differences are largest for Europe and Latin America. An immigrant described as speaking fluent English is guessed to be from Europe 32.6% of the time and to be from Latin America 24.1% of the time, on average. For

[8] I define it on the basis of pseudo-R^2 statistics (see Table S2 in Supplementary Material). A similar method has been used in the literature before (Jenke et al. 2021).

Table 2. Tests of attributes' overall importance

Attribute	χ^2_4	p-value
Age	7.2	.127
Gender	72.3	$< .001$
Education	10.1	.039
English proficiency	676.9	$< .001$
Government benefits	46.7	$< .001$
Prior trips to U.S.	122.0	$< .001$
Police record	270.1	$< .001$

Note. Results of likelihood-ratio tests.

someone with broken English the corresponding probabilities are 17.5% and 32.4%, respectively. The differences in probabilities for Africa, Asia, and the Middle East are also significant but much smaller in magnitude—likely, due to the floor effect.[9] Proper estimation of effects even for groups that are relatively rare choices (and thus are subject to floor effects) is another important benefit of the multinomial specification.[10]

Researchers using multinomial conjoint experiments can also be interested in understanding whether each specific attribute is relevant or not for profile categorizations. Unlike in most other regression-based methods, such inferences cannot be made on the basis of significance tests for single multinomial logit coefficients—even when an attribute has only two potential values. Instead, one has to implement a joint likelihood-ratio test that all choice-specific coefficients associated with a certain attribute are simultaneously zero (Long 1997, 161). Results of such tests for the seven attributes in my experiment are presented in Table 2. They show that immigrants' gender, English proficiency, government benefits, prior trips to U.S., and police record all impact categorizations. Age does not have an effect whereas education is only marginally significant.

[9] For marginal means for the other attributes, see Figures S2–S7 in Supplementary Material.

[10] Still, researchers should keep this aspect of the design in mind. It is probably unadvisable to increase the number of categorization options to a degree when mean probabilities become too low.

Discussion and Conclusion

In this paper, I have described how applied researchers in political science can design and analyze conjoint experiments with multinomial outcomes: more than two choice options that cannot be meaningfully ordered. The quantities of interest from such conjoint tasks can be estimated using a well-known method: multinomial logistic regression. This estimation requires making parametric assumptions about the error structure—but preserves causal interpretation of effects and keeps other benefits of the conjoint method. I have also documented some important differences in presentation and interpretation of results from multinomial conjoint experiments. To illustrate practical application of the proposed design, I have carried out an original conjoint experiment intended to measure Americans' stereotypes about immigrants from five world regions: Africa, Asia, Europe, Latin America, and the Middle East.

Besides providing a methodological demonstration, presented analysis makes a substantive contribution to the literature on politics of immigration. Specifically, I find that Americans associate immigrants from regions other than Europe with speaking poor English and, with the exception of Asians, being rule-breakers. These two attributes are known to provoke opposition to immigration among natives (Enos 2014; Newman, Hartman, and Taber 2012; Wright, Levy, and Citrin 2016). Fluent English and respect for laws can signal integration into the American society (Levy and Wright 2020; Ostfeld 2017), or be seen as an economic asset.[11] Independently of the specific mechanism, stereotypes linking immigrants from outside of Europe to undesirable attributes can lead Americans to oppose non-European immigration—even if overall preferences on immigration are meritocratic.

[11] Although in some contexts immigrants who speak poor English but make a sincere effort are viewed most favorably (Hopkins 2015).

At the same time, my results do not show notable stereotypes that connect different immigrant origins with higher or lower education levels. This is an important finding since education can be seen as the closest proxy for skills that are known to strongly impact immigration preferences (Hainmueller and Hopkins 2015; Valentino et al. 2019). However, this is just one possible interpretation of this result. For instance, respondents can see language as a more important cue for economic contribution than education. Americans can also discount non-Western education credentials as an indicator of valuable skills. An additional test that can be implemented in future experimental studies is manipulating immigrants' occupations as a proxy for skill.

Another interesting finding concerns the association between immigrant origins and reliance on government benefits. I demonstrate that the corresponding stereotypes are not uniform: immigrants from Latin America are seen as more reliant on benefits then Europeans whereas immigrants from Asia less so. This result can explain a recent controversy in the literature that concern "immigrationization" of welfare in the United States (Garand, Xu, and Davis 2017; Levy 2021). Do Americans increasingly see welfare recipients as immigrants? Results of my conjoint experiment suggest that these perceptions likely vary across origins: some immigrant groups are seen as more likely to receive government benefits than others.

A separate question deals with the role of race in Americans' stereotypes about immigrant origins. The importance of racism and ethnocentrism in U.S. public opinion is well documented (Cramer 2020; Kinder and Kam 2009). So, are positive stereotypes about European immigrants indicate presence of racial prejudice? My results do not provide a definite answer. Some stereotypes indeed follow the "white vs. nonwhite" dichotomy: for instance, all non-European immigrant groups are believed to speak worse English. Others, however, do not since Asian immigrants are seen as more law-abiding than Europeans. Also, these stereotypes do not depend on respondents' ethnocentrism. Future conjoint experiments can directly explore Americans' stereotypes about different racial groups—and, again, multinomial design would be helpful in such studies.

As a potential avenue for future methodological research, researchers can apply advanced multinomial modeling techniques to conjoint experiments. The standard multinomial logistic regression makes a common but rather restrictive assumption regarding independent distributions of errors across different outcomes, known as the independence of irrelevant alternatives or IIA. This assumption is relaxed in other types of models, such as mixed logit (Glasgow 2001), that can also be more computationally demanding.

Overall, by describing how researchers can design and analyze conjoint experiments with multinomial outcomes, this paper makes one more addition to the growing flexibility of this powerful method. Using the example of stereotypes about different immigrant origins in the United States, I also demonstrate how multinomial conjoint experiments can be used to address substantively important questions in political science.

References

Abramson, Scott F., Korhan Kocak, and Asya Magazinnik. 2022. "What Do We Learn about Voter Preferences from Conjoint Experiments?" *American Journal of Political Science*. Published ahead of print. https://doi.org/10.1111/ajps.12714.

Aldrich, John H., and Richard D. McKelvey. 1977. "A Method of Scaling with Applications to the 1968 and 1972 Presidential Elections." *American Political Science Review* 71 (1): 111–30.

Bansak, Kirk, Jens Hainmueller, Daniel J. Hopkins, and Teppei Yamamoto. 2018. "The Number of Choice Tasks and Survey Satisficing in Conjoint Experiments." *Political Analysis* 26 (1): 112–19.

Bansak, Kirk, Jens Hainmueller, Daniel J. Hopkins, and Teppei Yamamoto. 2021. "Beyond the Breaking Point? Survey Satisficing in Conjoint Experiments." *Political Science Research and Methods* 9 (1): 53–71.

Bansak, Kirk, Jens Hainmueller, Daniel J. Hopkins, and Teppei Yamamoto. 2022. "Using Conjoint Experiments to Analyze Election Outcomes: The Essential Role of the Average Marginal Component Effect." *Political Analysis*. Published ahead of print. https://doi.org/10.1017/pan.2022.16.

Bizumic, Boris, and John Duckitt. 2012. "What Is and Is Not Ethnocentrism? A Conceptual Analysis and Political Implications." *Political Psychology* 33 (6): 887–909.

Blinder, Scott. 2015. "Imagined Immigration: The Impact of Different Meanings of 'Immigrants' in Public Opinion and Policy Debates in Britain." *Political Science Research and Methods* 9 (1): 53–71.

Brader, Ted, Nicholas A. Valentino, and Elizabeth Suhay. 2015. "What Triggers Public Opposition to Immigration? Anxiety, Group Cues, and Immigration Threat." *American Journal of Political Science* 52 (4): 959–78.

Brady, Henry E. 1985. "The Perils of Survey Research: Inter-Personally Incomparable Responses." *Political Methodology* 11 (3/4): 269–91.

Branton, Regina, Erin C. Cassese, Bradford S. Jones, and Chad Westerland. 2011. "All Along the Watchtower: Acculturation Fear, Anti-Latino Affect, and Immigration." *Journal of Politics* 73 (3): 664–79.

Citrin, Jack, Donald P. Green, Christopher Muste, and Cara Wong. 1997. "Public Opinion Toward Immigration Reform: The Role of Economic Motivations." *Journal of Politics* 59 (3): 858–81.

Clarke, Andrew J. 2020. "Party Sub-Brands and American Party Factions." *American Journal of Political Science* 64 (3): 452–70.

Coppock, Alexander, and Oliver A. McClellan. 2019. "Validating the Demographic, Political, Psychological, and Experimental Results Obtained from a New Source of Online Survey Respondents." *Research and Politics* 6 (1). https://doi.org/10.1177/2053168018822174.

Cramer, Katherine. 2020. "Understanding the Role of Racism in Contemporary US Public Opinion." *Annual Review of Political Science* 23: 153–69.

De la Cuesta, Brandon, Naoki Egami, and Kosuke Imai. 2022. "Improving the External Validity of Conjoint Analysis: The Essential Role of Profile Distribution." *Political Analysis* 30 (1): 19–45.

Enos, Ryan D. 2014. "Causal Effect of Intergroup Contact on Exclusionary Attitudes." *Proceedings of the National Academy of Sciences of the United States of America* 111 (10): 3699–704.

Flores, Rene D., and Ariela Schachter. 2018. "Who Are the 'Illegals'? The Social Construction of Illegality in the United States." *American Sociological Review* 83 (5): 839–68.

Ganter, Flavien. 2021. "Identification of Preferences in Forced-Choice Conjoint Experiments: Reassessing the Quantity of Interest." *Political Analysis*. Published ahead of print. https://doi.org/10.1017/pan.2021.41.

Garand, James C., Ping Xu, and Belinda C. Davis. 2017. "Immigration Attitudes and Support for the Welfare State in the American Mass Public." *American Journal of Political Science* 61 (1): 146–62.

Glasgow, Garrett. 2001. “Mixed Logit Models for Multiparty Elections.” *Political Analysis* 9 (2): 116–36.

Goggin, Stephen N., John A. Henderson, and Alexander G. Theodoridis. 2020. “What Goes with Red and Blue? Mapping Partisan and Ideological Associations in the Minds of Voters.” *Political Behavior* 42: 985–1013.

Groenendyk, Eric, Michael W. Sances, and Kirill Zhirkov. 2020. “Intraparty Polarization in American Politics.” *Journal of Politics* 82 (4): 1616–20.

Hainmueller, Jens, Dominik Hangartner, and Teppei Yamamoto. 2015. “Validating Vignette and Conjoint Survey Experiments against Real-World Behavior.” *Proceedings of the National Academy of Sciences of the United States of America* 112 (8): 2395–400.

Hainmueller, Jens, and Daniel J. Hopkins. 2015. “The Hidden American Immigration Consensus: A Conjoint Analysis of Attitudes toward Immigrants.” *American Journal of Political Science* 59 (3): 529–48.

Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto. 2014. “Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments.” *Political Analysis* 22 (1): 1–30.

Hanretty, Chris, Benjamin E. Lauderdale, and Nick Vivyan. 2020. “A Choice-Based Measure of Issue Importance in the Electorate.” *American Journal of Political Science* 64 (3): 519–35.

Hare, Christopher, Armstrong, II, David A., Ryan Bakker, Royce Carroll, and Keith T. Poole. 2015. “Using Bayesian Aldrich-McKelvey Scaling to Study Citizens’ Ideological Preferences and Perceptions.” *American Journal of Political Science* 59 (3): 759–74.

Hartman, Todd K., Benjamin J. Newman, and C. Scott Bell. 2014. "Decoding Prejudice Toward Hispanics: Group Cues and Public Reactions to Threatening Immigrant Behavior." *Political Behavior* 36 (1): 143–66.

Hilton, James L., and William von Hippel. 1996. "Stereotypes." *Annual Review of Psychology* 47: 237–71.

Hopkins, Daniel J. 2015. "The Upside of Accents: Language, Inter-group Difference, and Attitudes toward Immigration." *British Journal of Political Science* 45 (3): 531–57.

Horiuchi, Yusaku, Zachary Markovich, and Teppei Yamamoto. 2022. "Does Conjoint Analysis Mitigate Social Desirability Bias?" *Political Analysis* 30 (4): 535–49.

Horiuchi, Yusaku, Daniel M. Smith, and Teppei Yamamoto. 2018. "Measuring Voters' Multidimensional Policy Preferences with Conjoint Analysis: Application to Japan's 2014 Election." *Political Analysis* 26 (2): 190–209.

Jenke, Libby, Kirk Bansak, Jens Hainmueller, and Dominik Hangartner. 2021. "Using Eye-Tracking to Understand Decision-Making in Conjoint Experiments." *Political Analysis* 29 (1): 75–101.

Kinder, Donald R., and Cindy D. Kam. 2009. *Us against Them: Ethnocentric Foundations of American Opinion.* Chicago: University of Chicago Press.

King, Gary, Christopher J. L. Murray, Joshua A. Salomon, and Ajay Tandon. 2004. "Enhancing the Validity and Cross-Cultural Comparability of Measurement in Survey Research." *American Political Science Review* 98 (1): 191–207.

King, Gary, Michael Tomz, and Jason Wittenberg. 2000. "Making the Most of Statistical Analyses: Improving Interpretation and Presentation." *American Journal of Political Science* 44 (2): 347–61.

King, Gary, and Jonathan Wand. 2007. “Comparing Incomparable Survey Responses: Evaluating and Selecting Anchoring Vignettes.” *Political Analysis* 15 (1): 46–66.

Klar, Samara, and Yanna Krupnikov. 2016. *Independent Politics: How American Disdain for Parties Leads to Political Inaction.* New York: Cambridge University Press.

Konitzer, Tobias B., Shanto Iyengar, Nicholas A. Valentino, Stuart Soroka, and Raymond M. Duch. 2019. “Ethnocentrism versus Group-specific Stereotyping in Immigration Opinion: Cross-national Evidence on the Distinctiveness of Immigrant Groups.” *Journal of Ethnic and Migration Studies* 45 (7): 1051–74.

Leeper, Thomas J., Sara B. Hobolt, and James Tilley. 2020. “Measuring Subgroup Preferences in Conjoint Experiments.” *Political Analysis* 28 (2): 207–21.

Levy, Morris. 2021. “Once Racialized, Now ‘Immigrationized’? Explaining the Immigration–Welfare Link in American Public Opinion.” *Journal of Politics* 83 (4): 1275–91.

Levy, Morris, and Matthew Wright. 2020. *Immigration and the American Ethos.* Cambridge: Cambridge University Press.

Long, J. Scott. 1997. *Regression Models for Categorical and Limited Dependent Variables.* Thousand Oaks, CA: Sage Publications.

Lutz, Philipp, and Marco Bitschnau. 2022. “Misperceptions about Immigration: Reviewing Their Nature, Motivations and Determinants.” *British Journal of Political Science.* Published ahead of print. https://doi.org/10.1177/2053168018822174.

Malhotra, Neil, Yotam Margalit, and Cecilia Hyunjung Mo. 2013. “Economic Explanations for Opposition to Immigration: Distinguishing between Prevalence and Conditional Impact.” *American Journal of Political Science* 57 (2): 391–410.

McFadden, Daniel L. 1973. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics*, edited by Paul Zarembka, 105–42. New York: Academic Press.

Myers, C. Daniel, Kirill Zhirkov, and Kristin Lunz Trujillo. 2022. "Who Is 'On Welfare'? Validating the Use of Conjoint Experiments to Measure Stereotype Content." *Political Behavior*. Published ahead of print. https://doi.org/10.1007/s11109-022-09815-0.

Newman, Benjamin J., Todd K. Hartman, and Charles S. Taber. 2012. "Foreign Language Exposure, Cultural Threat, and Opposition to Immigration." *Political Psychology* 33 (5): 635–57.

Newman, Benjamin J., and Neil Malhotra. 2018. "Economic Reasoning with a Racial Hue: Is the Immigration Consensus Purely Race Neutral?" *Journal of Politics* 81 (1): 153–66.

Ostfeld, Mara. 2017. "The Backyard Politics of Attitudes Toward Immigration." *Political Psychology* 38 (1): 21–37.

Sen, Maya. 2017. "How Political Signals Affect Public Support for Judicial Nominations: Evidence from a Conjoint Experiment." *Political Research Quarterly* 70 (2): 374–93.

Valentino, Nicholas A., Ted Brader, and Ashley E. Jardina. 2013. "Immigration Opposition Among U.S. Whites: General Ethnocentrism or Media Priming of Attitudes About Latinos?" *Political Psychology* 34 (2): 149–66.

Valentino, Nicholas A., Stuart N. Soroka, Shanto Iyengar, Toril Aalberg, Raymond Duch, Marta Fraile, Kyu S. Hahn, Kasper M. Hansen, Allison Harell, Marc Helbling, Simon D. Jackman, and Tetsuro Kobayashi. 2019. "Economic and Cultural Drivers of Immigrant Support Worldwide." *British Journal of Political Science* 49 (4): 1201–26.

Wright, Matthew, Morris Levy, and Jack Citrin. 2016. “Public Attitudes Toward Immigration Policy Across the Legal/Illegal Divide: The Role of Categorical and Attribute-Based Decision-Making.” *Political Behavior* 38 (1): 229–53.

Zhirkov, Kirill. 2021a. “Social Dominance Orientation and Differential Affect toward Immigrant Origin Groups: Evidence from Three Immigration-Receiving Countries.” *International Journal of Intercultural Relations* 85: 170–83.

Zhirkov, Kirill. 2021b. “Who Are ‘the Immigrants’? Beliefs About Immigrant Populations and Anti-Immigration Attitudes in the United States and Britain.” *Social Science Quarterly* 102 (1): 228–37.

Zhirkov, Kirill. 2022. “Estimating and Using Individual Marginal Component Effects from Conjoint Experiments.” *Political Analysis* 30 (2): 236–49.

Supplementary Material

Contents

Table S1. Full ethnocentrism battery

(Preamble) Below are items that relate to different cultures and ethnic groups. Work quickly and record your first reaction to each item. There are no right or wrong answers. Please indicate the degree to which you disagree or agree with each item.
(1) In most cases, I like people from my culture more than I like others.
(2) I don't think I have any particular preference for my own cultural or ethnic group over others. (reversed)
(3) The world would be a much better place if all other cultures and ethnic groups modelled themselves on my culture.
(4) The values, way of life, and customs of most other cultures are probably just as good as those of my own. (reversed)
(5) Our culture would be much better off if we could keep people from different cultures out.
(6) I like the idea of a society in which people from completely different cultures, ethnic groups, and backgrounds mix together freely. (reversed)
(7) We need to do what's best for our own people, and stop worrying so much about what the effect might be on other peoples.
(8) We should always show consideration for the welfare of people from other cultural or ethnic groups even if, by doing this, we may lose some advantage over them. (reversed)

Note. Respondents were randomly presented with four statements, one from each following pair: (1) and (2), (3) and (4), (5) and (6), (7) and (8). Answers were given on a 7-point Likert-type scale from 1 = *Strongly disagree* to 7 = *Strongly agree*

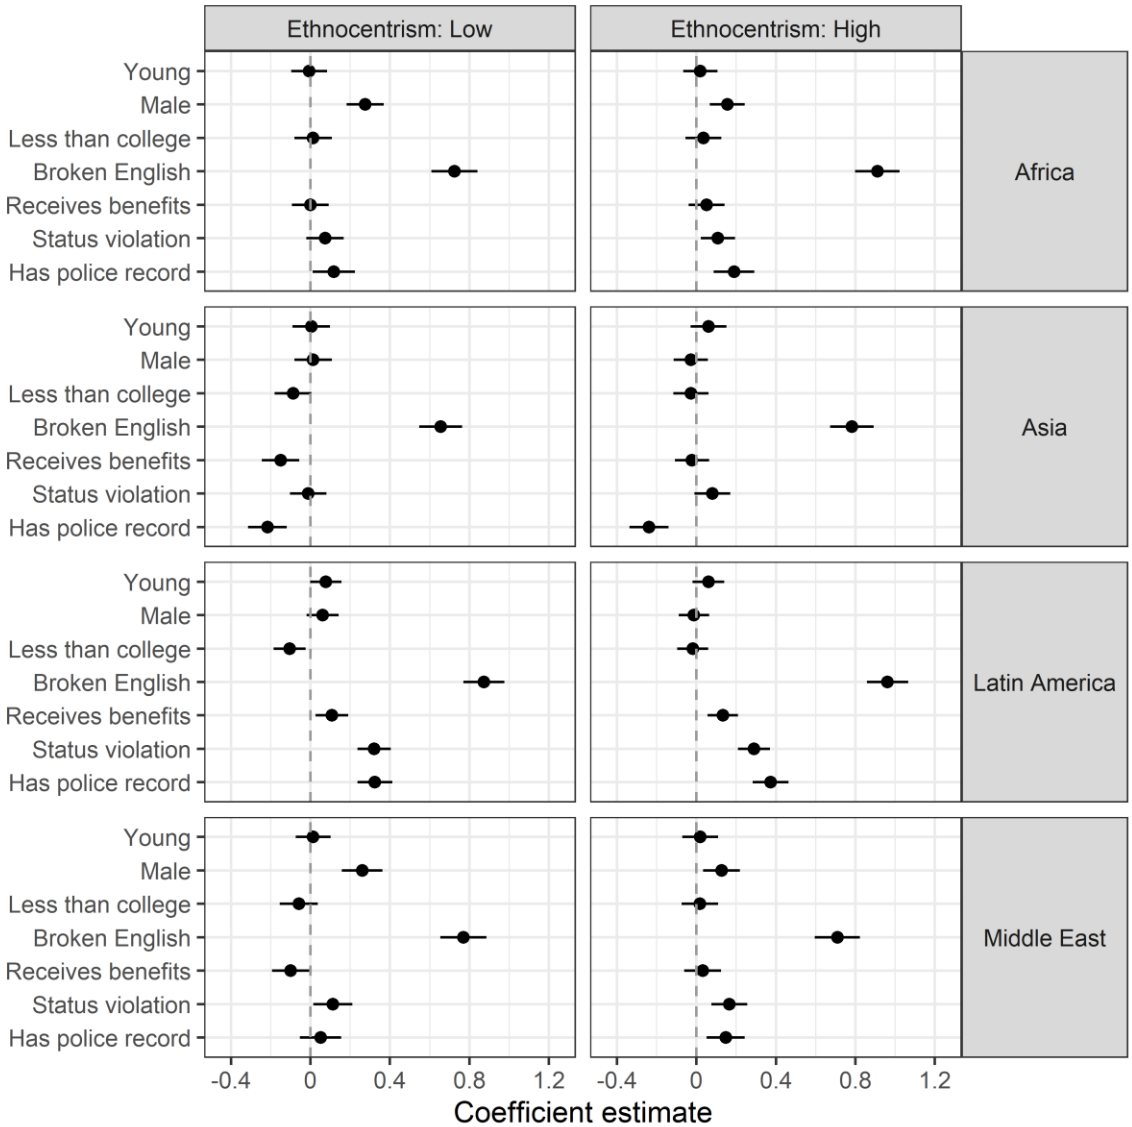

Figure S1. Coefficients from the multinomial logistic regression by respondents' levels of ethnocentrism

Table S2. Relative importance of the attributes

Attribute	Pseudo-R^2
Age	< .001
Gender	< .001
Education	< .001
English proficiency	.010
Government benefits	< .001
Prior trips to U.S.	.001
Police record	.003

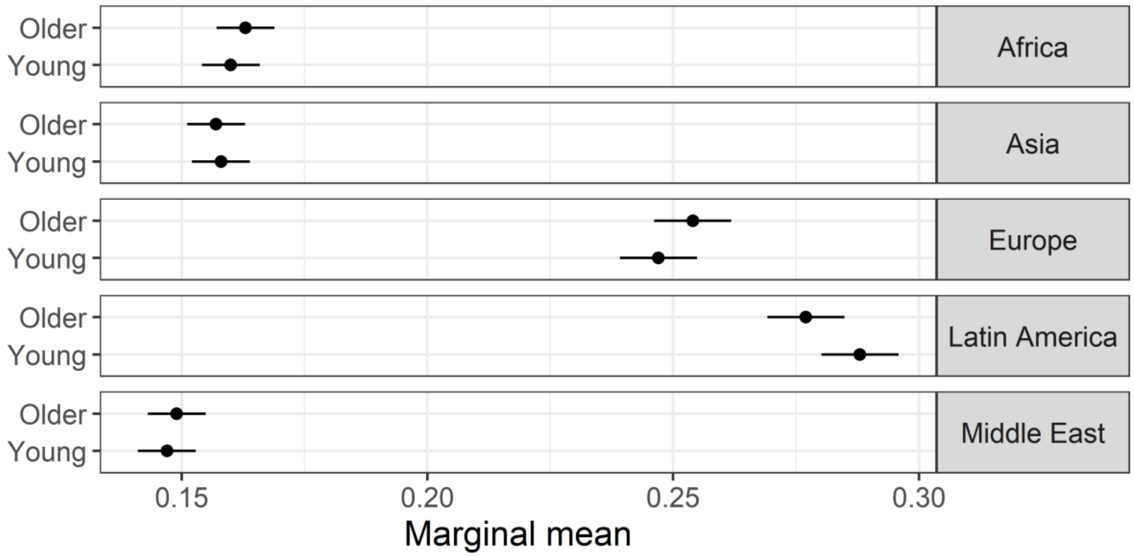

Figure S2. Marginal means for the age attribute

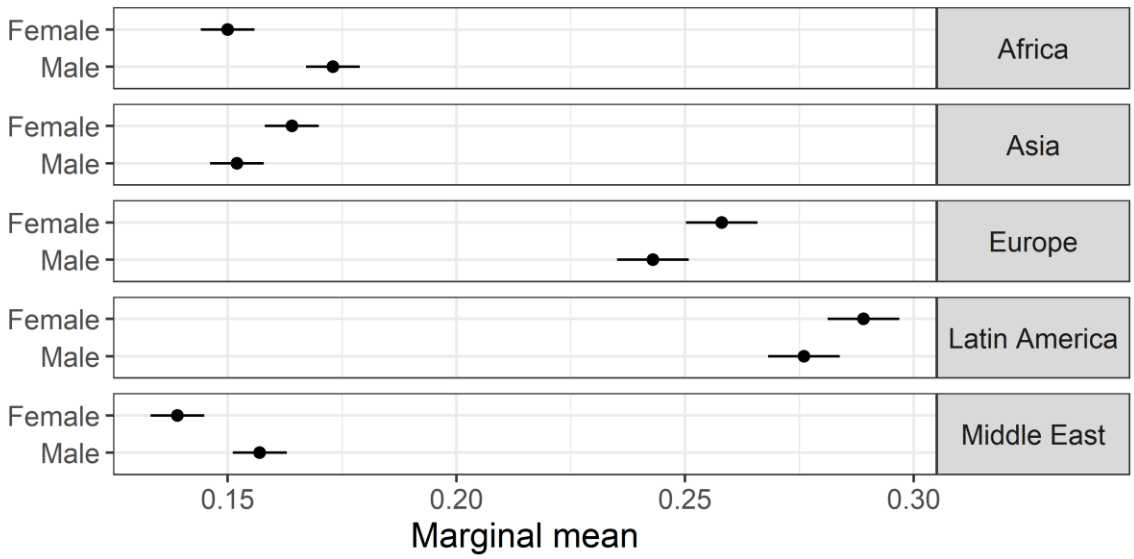

Figure S3. Marginal means for the gender attribute

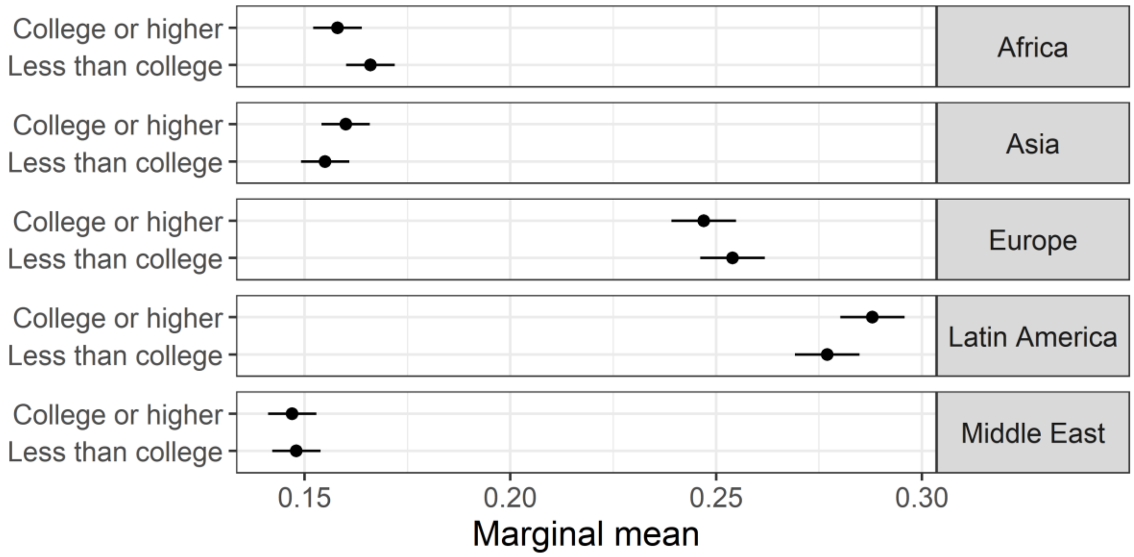

Figure S4. Marginal means for the education attribute

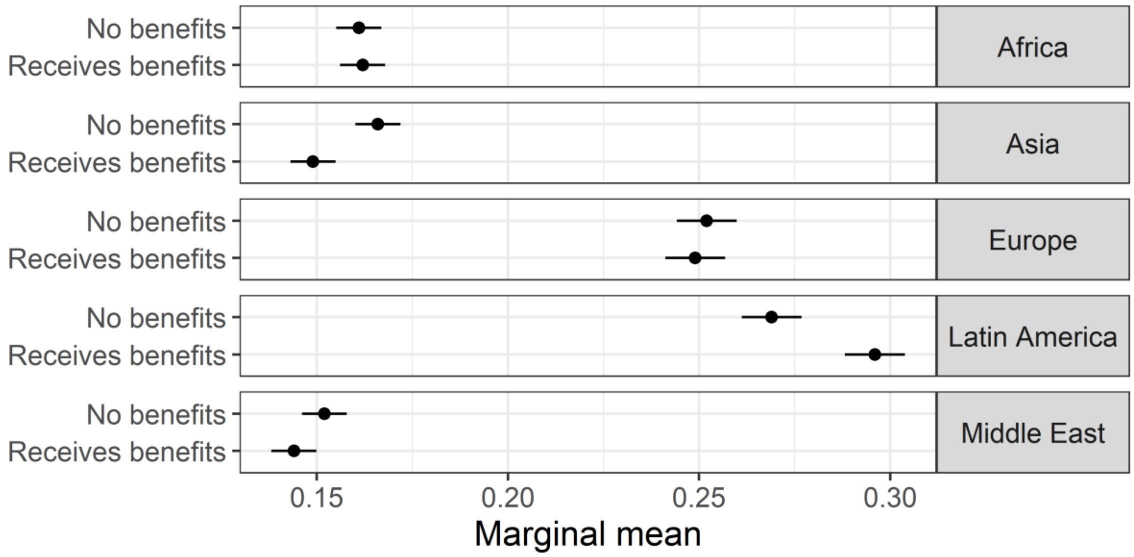

Figure S5. Marginal means for the government benefits attribute

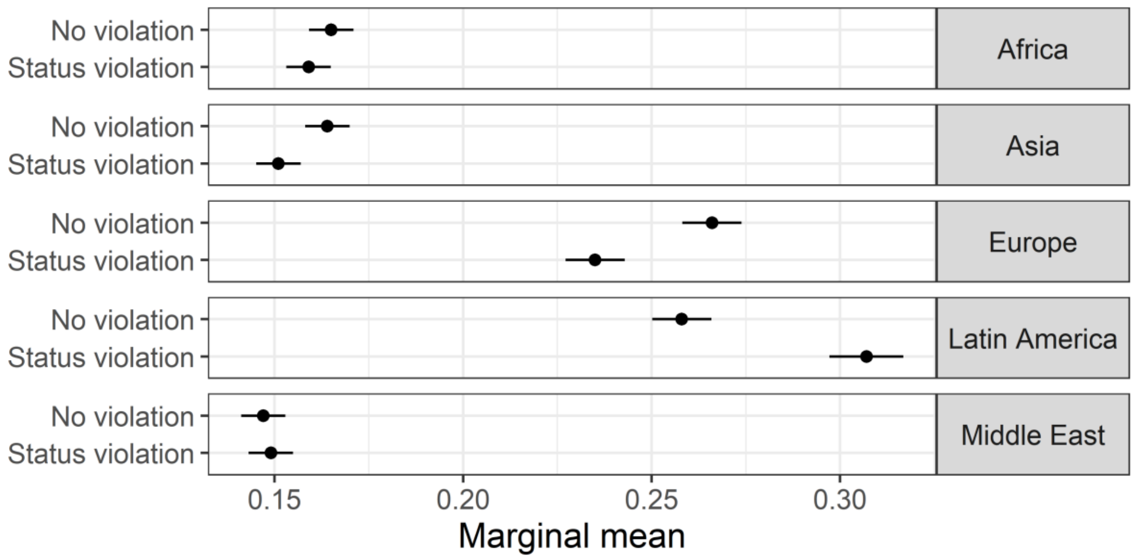

Figure S6. Marginal means for the prior trips to U.S. attribute

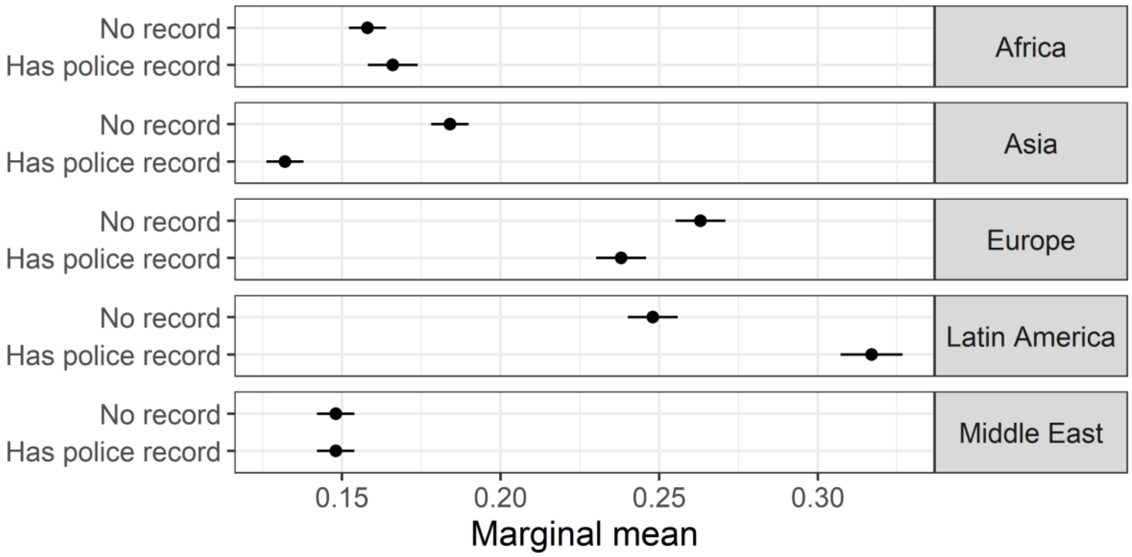

Figure S7. Marginal means for the police record attribute