

Using Machine Learning to Test Causal Hypotheses in Conjoint Analysis*

Dae Woong Ham^{†1}, Kosuke Imai^{‡1,2}, and Lucas Janson^{§1}

¹*Department of Statistics, Harvard University*

²*Department of Government, Harvard University*

August 24, 2022

Abstract

Conjoint analysis is a popular experimental design used to measure multidimensional preferences. Researchers examine how varying a factor of interest, while controlling for other relevant factors, influences decision-making. Currently, there exist two methodological approaches to analyzing data from a conjoint experiment. The first focuses on estimating the average marginal effects of each factor while averaging over the other factors. Although this allows for straightforward design-based estimation, the results critically depend on the distribution of other factors and how interaction effects are aggregated. An alternative model-based approach can compute various quantities of interest, but requires researchers to correctly specify the model, a challenging task for conjoint analysis with many factors and possible interactions. In addition, a commonly used logistic regression has poor statistical properties even with a moderate number of factors when incorporating interactions. We propose a new hypothesis testing approach based on the conditional randomization test to answer the most fundamental question of conjoint analysis: Does a factor of interest matter *in any way* given the other factors? Our methodology is solely based on the randomization of factors, and hence is free from assumptions. Yet, it allows researchers to use any test statistic, including those based on complex machine learning algorithms. As a result, we are able to combine the strengths of the existing design-based and model-based approaches. We illustrate the proposed methodology through conjoint analysis of immigration preferences and political candidate evaluation. We also extend the proposed approach to test for regularity assumptions commonly used in conjoint analysis. An open-source software package is available for implementing the proposed methodology.¹

*We thank Naoki Egami for advice. Imai thanks the Alfred P. Sloan Foundation for partial support (Grant # 2020–13946). Ham and Janson were partially supported by a CAREER grant from the National Science Foundation (Grant # DMS-2045981).

[†]Email: daewoongham@g.harvard.edu

[‡]Email: imai@harvard.edu, URL: <https://imai.fas.harvard.edu>

[§]Email: ljanson@fas.harvard.edu, URL: <http://lucasjanson.fas.harvard.edu>

¹The proposed methodology is implemented via an open-source software R package *CRTConjoint*, available through the Comprehensive R Archive Network (<https://cran.r-project.org/package=CRTConjoint>).

1 Introduction

Conjoint analysis, introduced more than half a century ago (Luce and Tukey, 1964), is a factorial survey-based experiment designed to measure preferences on a multidimensional scale. Under a commonly used “forced-choice” design, respondents are presented with two alternative profiles of randomly selected attributes (e.g., candidates with different traits or products with different features). They are then asked to choose their preferred profile. Conjoint analysis has been extensively used by marketing firms to determine desirable product characteristics (e.g., Bodog and Florian, 2012; Green, Krieger and Wind, 2001). Recently, it has gained popularity among social scientists (Hainmueller, Hopkins and Yamamoto, 2014; Raghavarao, Wiley and Chitturi, 2010) who are interested in studying individual preferences concerning elections (e.g., Ono and Burden, 2018), immigration (e.g., Hainmueller and Hopkins, 2015), employment (e.g., Popovic, Kuzmanovic and Martic, 2012), and other issues.

When analyzing conjoint experiments, the *design-based* approach, pioneered by Hainmueller, Hopkins and Yamamoto (2014), has been by far the most popular among social scientists. The main advantage of this nonparametric approach is its simplicity—it uses the difference-in-means estimator or linear regression to infer the average marginal component effect (AMCE) of each factor by averaging over the distribution of other factors. However, because the AMCE makes inferences about the marginal effects averaged over all the other factors, it may fail to capture important interactions. This is potentially problematic given that practitioners tend to use AMCE-based confidence intervals that are narrow and contain zero to conclude that a factor has a weak causal effect (Hainmueller, Hopkins and Yamamoto, 2014; Hainmueller and Hopkins, 2015; Ono and Burden, 2018). However, a narrow AMCE-based confidence interval containing zero only implies that a factor has a weak *marginal* effect, not necessarily that its total causal influence is weak.

A possible solution is the *model-based* approach that ranges from traditional parametric regression models (McFadden, 1973; Green and Srinivasan, 1990; Campbell, Mhlanga and Lesschaeve, 2013) to more recent machine learning algorithms (Egami and Imai, 2019; de la Cuesta, Egami and Imai, 2022; Abramson et al., 2020; Bansak et al., 2020; Goplerud, Imai and Pashley, 2022). While this approach can efficiently estimate various quantities of interest, it has the potential drawback of model misspecification, producing biased inference. Although researchers can reduce model misspecification by adopting a more complicated model (e.g., adding interaction terms of increasing order), such an approach can substantially reduce statistical power and produce invalid p -values even when the model is correctly specified (Candès and Sur, 2018). While subgroup analysis, a common practice to analyze only a subset of the data, is simpler, it suffers from well-known problems of multiple testing and p -hacking, which are of serious concern in conjoint analysis given a large number of possible causal effects of interest (see, e.g., Szucs, 2016; Harking, 2021). Finally, the use of machine learning algorithms, which is becoming increasingly common, cannot yield even consistent estimates in high-dimensional settings without strong assumptions.

In this paper, we propose a new approach to analyzing data from conjoint analysis that combines the strengths of the existing design-based and model-based approaches (Section 3). Specifically, we show how to conduct assumption-free hypothesis testing based on the conditional randomization test (CRT; Candès et al., 2018). In the causal inference literature, the CRT has been used to test interference between units (Aronow, 2012; Athey, Eckles and Imbens, 2018) and in other causal applications such as genetic studies (Bates et al., 2020). Instead of focusing on a particular causal effect, we ask the most fundamental question of conjoint analysis: Does a factor of interest matter

in any way given the other factors? The proposed approach allows one to answer this question with greater statistical power than the AMCE by utilizing flexible machine learning algorithms but without making any assumption about the underlying causal structure (e.g., presence or absence of interaction effects, patterns of heterogeneous effects, or within-respondent correlation across profile comparisons). Despite its flexibility, the CRT has an attractive statistical property that the resulting p -values are exactly valid regardless of the sample size or the number of factors. We also extend the proposed approach to test whether specific factor levels of interest, rather than all levels of a factor, influence respondent preferences *in any way* (Section 3.4).

The proposed methodology can also test the validity of assumptions commonly invoked in conjoint analysis (Hainmueller, Hopkins and Yamamoto, 2014) (Section 3.5). First, we show how to test for the presence of the profile order effect. Under the forced conjoint design, for example, reversing the order in which two profiles are presented within each evaluation may change the choice of profile. Second, we test the assumption of no carryover effect, which states that each respondent's response only depends on the current profiles and is not affected by previous evaluations. This assumption may be violated if respondents learn over several profile evaluations. Third, we test the assumption of no fatigue effect, which precludes the possibility that as respondents evaluate more profiles, they get tired and answer questions differently (Bansak et al., 2018, 2019). Thus, the proposed hypothesis testing approach can serve as the first step of analyzing conjoint data without assumptions, complementing existing approaches that estimate causal quantities of interest.

After presenting the proposed methodology, we conduct simulation studies in Section 4 to show that the CRT can achieve a higher statistical power than the AMCE by exploiting machine learning algorithms to detect complex treatment interactions. For empirical illustration, we apply the proposed methodology to two existing conjoint analyses (Section 5). The first pertains to immigration preferences among United States (U.S.) citizens. While some researchers contend that U.S. citizens generally prefer high-skilled immigrants regardless of their countries of origin, others have suggested that racially prejudiced respondents discriminate against non-European immigrants (Hainmueller and Hopkins, 2015; Newman and Malhotra, 2019). By combining machine learning algorithms with the CRT, we find that respondents do differentiate according to whether immigrants are from Mexico or European countries. The second application considers the role of gender in candidate evaluation (Ono and Burden, 2018). The original analysis found that voters discriminate between male and female candidates only for presidential elections but not for Congressional elections. However, a recent study suggests that this finding about congressional candidates may critically depend on the distribution of other characteristics such as partisanship and policy positions (de la Cuesta, Egami and Imai, 2022). We apply the proposed methodology and show that gender does play a statistically significant role in voter evaluation of Congressional candidates.

We emphasize that the CRT complements other existing methods by offering a useful test of whether a factor of interest matters at all. The test can be useful even when the AMCE fail to detect any statistically significant result.

2 Empirical Application

In this section, we briefly describe two empirical applications concerning the role of ethnocentrism in immigration preferences and gender discrimination in political candidate evaluations. We also outline the limitations of the commonly used approach based on the AMCE to motivate the proposed

methodology. In Section 5, we revisit these applications and apply our hypothesis testing approach.

2.1 Role of Country of Origin in Immigration Preference

Immigration is one of the most contentious issues in the United States today. A large body of literature investigates how cultural, economic, and racial factors shape public attitudes towards immigration (see, e.g., Hainmueller and Hopkins, 2014, and references therein). In an influential study, Hainmueller and Hopkins (2015) use a conjoint analysis to empirically examine the immigrant characteristics favored or disfavored by U.S. citizens. The survey was fielded between December 2011 and January 2012 on a nationally representative sample of U.S. citizens. The study used the forced-choice design, in which each respondent was presented with a pair of hypothetical immigrant profiles and asked which immigrant they would “personally prefer to see admitted to the United States.” Each of 1,396 respondents rated 5 pairs of profiles.

An immigrant profile consists of nine factors—prior trips to the U.S., reason for application, country of origin, language skills, profession, job experience, employment plans, education level, and gender, each of which has multiple levels (see Table 2 of Appendix D as well as the original article for details). Most factors are independently and uniformly randomized across their levels with the exception of two restrictions to avoid implausible pairs. First, immigrant profiles that list *escape persecution* as the “reason of immigration” can only have *Iraq*, *Sudan*, or *Somalia* as their “country of origin”. Second, a high-skill “profession” such as *financial analyst*, *research scientist*, *doctor*, and *computer programmer* is possible only if the “education level” is at least *2 years of college*. This restricted randomization scheme induces dependencies between these factors that must be properly accounted for when analyzing the data. The survey also contains information about respondents’ age, education, ethnicity, gender, and ethnocentrism. The study contains a random sample of 14,018 profiles.

In this study, Hainmueller and Hopkins estimate the AMCE, which represents the marginal effect of a factor of interest averaging over the other factors. Based on the statistically insignificant estimates for the AMCEs of the “country of origin” factor for Mexico and European countries (reproduced in Figure 1), they conclude that “despite media frames focusing on low-skilled, unauthorized immigration from Mexico, there is little evidence of penalty specific to Mexicans.” (p. 539). The authors obtain these estimates by fitting a linear regression model, where the outcome variable indicates whether the profile is selected and the predictors are the nine randomized factors. To account for the restricted randomization, they also include two sets of interaction terms, one between “country of origin” and “reason of immigration” and the other between “profession” and “education level”. To obtain the estimated AMCE of *Germany*, for example, Hainmueller and Hopkins take the main effect of *Germany* (the baseline is *India*) and then add it to the average of all the interaction terms between *Germany* and the “reason of immigration” factor. Clustered standard errors are computed by clustering on each respondent to account for dependency within a respondent.

Despite this overall finding, the AMCE-based approach may mask relevant interactions and heterogeneous treatment effects. Indeed, Hainmueller and Hopkins conduct a subgroup analysis and find that the “country of origin” factor has statistically significant interactions with the respondents’ ethnocentrism. They measure ethnocentrism using the feeling thermometer score (ranging from 0 to 100) for the respondent’s own groups minus the average feeling thermometer across the other groups. In addition, Newman and Malhotra (2019) reanalyze the same dataset and estimate three-




Figure 1: The estimated Average Marginal Component Effects (AMCEs) of immigrants’ countries of origin in the Hainmueller and Hopkins (2015) study. The plot shows the estimated AMCEs for *France*, *Germany*, *Mexico*, and *Poland*, which represent the average differences in the estimated probability of choosing an immigrant profile with a specific level of the “country of origin” factor, marginalizing other attributes. The baseline factor level is *India*, and the 95% confidence intervals are also shown.

way interactions among respondents’ ethnocentrism, “country of origin”, and “profession”. The authors compute the AMCEs of high-skilled immigrants (baseline of *janitor*) separately for each country of origin and respondent’s ethnocentric group. They find that these AMCEs are different between Mexican and European immigrants when compared among highly ethnocentric respondents (see Figure 1 in Newman and Malhotra, 2019, for further details).

In this paper, we apply the proposed hypothesis testing approach to testing whether or not immigrants from Mexico and those from Europe are viewed differently in *any way* while controlling for all the other experimental factors as well as the respondent characteristics. The rejection of this null hypothesis would mean that the country of origin of an individual plays a statistically significant role in some United States citizens’ preferences about that individual’s immigration to the United States.

2.2 Role of Gender in Political Candidate Evaluation

Perhaps the most common political science application of conjoint analysis is the measurement of voters’ candidate preferences. Recently, several scholars have used conjoint analysis to study the role of gender discrimination in candidate evaluation (e.g., Ono and Burden, 2018; Teele, Kalla and Rosenbluth, 2018). We revisit the study by Ono and Burden (2018) which examines whether voters prefer candidates of one gender over those of another after controlling for other candidate characteristics.² The study is based on a sample of voting-eligible adults in the U.S. collected in March 2016 and also uses the forced-choice conjoint design. The following 13 factors are independently and uniformly randomized across their levels: gender, age, race, family, experience in public office, salient personal characteristics, party affiliation, policy area of expertise, position on national security, position on immigrants, position on abortion, position on government deficit, and favorability among the public (see Table 3 in Appendix D and the original article for details). The survey also contains information about the respondents’ educational background, gender, age, region, social

²This study treats gender as a binary factor with levels *Male* and *Female*.




Figure 2: The estimated Average Marginal Component Effect (AMCE) of candidate's gender in the Ono and Burden (2018) study. We present the estimates for congressional candidates (left) and presidential candidates (right). The 95% confidence intervals are also shown.

class, partisanship, political interest, and ethnocentrism. There were 1,583 respondents each given 10 tasks, resulting in 15,830 observations, half of which were for congressional candidates and the other half for presidential candidates.

The original study yields a negative estimate of the AMCE of female candidates, relative to male counterparts, for presidential candidates. However, the estimated AMCE of female candidates is not statistically distinguishable from zero for congressional candidates, based on a simple t -test for the coefficient of *Male* from the linear regression with cluster standard errors. This finding led to the authors' conclusion that gender discrimination "is limited to presidential rather than congressional elections" (p. 583). Figure 2 reproduces these AMCE estimates. Like the immigration example, the authors fit a linear regression with all fourteen factors as predictors to obtain these estimates.

In this paper, we use the CRT to formally test whether the gender of congressional candidates matters in *any way* for voters' preferences while controlling for the other candidate characteristics. The rejection of this null hypothesis would indicate that gender does matter even for congressional candidates.

2.3 Limitations of Existing Approaches

Although the AMCE is a useful causal quantity of interest and can be easily and reliably estimated, it is not free of limitations. The AMCE is a marginal effect based on two types of averaging: (1) averaging over the distribution of other attributes, and (2) averaging over the responses (and hence respondents). Recall that in the standard causal inference setting with a binary treatment, a zero average treatment effect does not necessarily imply zero treatment effect for everyone. The treatment may benefit some and harm others, and these positive and negative effects can cancel out through averaging. The AMCE suffers from a similar problem, potentially masking important causal heterogeneity if there are interactions among attributes and/or between attributes and respondent characteristics. In the immigration conjoint experiment described above, for example, the overall AMCE estimates suggest little difference across countries of origin. And yet Hainmueller and Hopkins show that respondents with high ethnocentrism may have a certain preference over certain countries of origin when compared to those with low ethnocentrism.

Additionally, although an AMCE-based confidence interval that does not contain zero represents evidence that the factor matters, a narrow AMCE-based confidence interval that contains zero only implies that a factor has a weak *marginal* effect. Nevertheless, practitioners tend to use narrow AMCE-based confidence intervals that contain zero to conclude that a factor does not matter. For example, Hainmueller, Hopkins and Yamamoto (2014) conclude that the “candidates’ income does not matter much” and that the “candidates’ racial and ethnic backgrounds are even less influential”, based on the AMCE-based confidence interval for income and ethnicity (p. 19). Similarly, Ono and Burden conclude that gender effects are “limited to only Congressional candidates” (p. 3), based on AMCE-based confidence interval of gender for Congressional candidate.

Although the AMCE is popular in conjoint analysis, especially among political scientists, there also exist model-based approaches to flexibly estimate potentially any quantity of interest. In particular, logistic regression remain a popular model-based alternative in conjoint analysis (McFadden, 1973; Green and Srinivasan, 1990; Campbell, Mhlanga and Lesschaeve, 2013) especially in marketing research.³ Despite the flexibility of logistic regression, model misspecification remains a significant challenge. Although researchers may add more interactions to account for all possible effects, such an approach can reduce statistical power and more importantly lead to invalid p -values (Candès and Sur, 2018). We show in Appendix J Figure 9 that using logistic regression and accounting for all two-way interactions to reduce model misspecification can easily lead to invalid p -values.

Consequently, a consensus among researchers has emerged that flexible machine learning algorithms are necessary for capturing these causal interactions (de la Cuesta, Egami and Imai, 2022; Abramson et al., 2020; Bansak et al., 2020; Goplerud, Imai and Pashley, 2022). Yet machine learning algorithms, despite their flexibility, cannot yield consistent estimates in high-dimensional settings without strong assumptions. In addition, statistical inference in small samples remains a challenge (Dezeure et al., 2015; Chernozhukov et al., 2017; Imai and Li, 2021). Our goal is to address these problems through an assumption-free approach based on the conditional randomization test.

3 The Proposed Methodology

In this section we describe the proposed methodology based on the Conditional Randomization Test (CRT; Candès et al., 2018). We show how to apply the CRT to conjoint analysis in order to test whether a factor of interest matters, without making any assumptions. We discuss various test statistics that can be used with the CRT and several useful extensions for conjoint analysis.

3.1 Notation and Setup

For concreteness, we focus on the forced-choice conjoint design, under which a respondent is asked to choose one of two profiles. Our methodology is general and can be extended to other designs. Let n be the total number of respondents. As is often done in practice, suppose that each respondent evaluates J pairs of profiles, yielding a total of nJ responses (for notational simplicity, we assume the same number of evaluations for each respondent).

³Although a hierarchical modeling approach remains another popular model-based alternative in conjoint studies, we do not consider it here because it is based on a Bayesian framework rather than frequentist approach taken in this paper (Andrews, Ansari and Currim, 2002).

We use $Y_{ij} \in \{0, 1\}$ to represent the binary outcome variable for evaluation j by respondent i , which equals 1 for selecting the left profile and 0 for choosing the right profile. Although for convenience we use “left” and “right” to distinguish two profiles under each evaluation, the profiles do not necessarily have to be placed side by side on the actual survey platform. We use the following $nJ \times 1$ stacked vector representation for this outcome variable $\mathbf{Y} = [\mathbf{Y}_1; \mathbf{Y}_2; \dots; \mathbf{Y}_n]$, where $\mathbf{Y}_i = [Y_{i1}; Y_{i2}; \dots; Y_{iJ}]$ of dimension $J \times 1$ denotes the outcome variable for respondent i . We use $[a_1; a_2; \dots; a_n]$ to denote a vertical stacking of vectors or matrices a_1, a_2, \dots, a_n . We often observe some characteristics of the respondents, and we use \mathbf{V}_i to denote a $J \times r$ -dimensional matrix of r pre-treatment covariates for respondent i that are repeated across J rows.

Next, let p represent the total number of attributes or factors used for each conjoint profile. We use a scalar $X_{ij\ell}^L \in \{1, 2, \dots, K_\ell\}$ to denote the value of the ℓ th factor of interest for evaluation j by respondent i , where the superscript distinguishes the factors for the left (L) and right (R) profiles, and $K_\ell \geq 2$ is the total number of factor levels for factor ℓ . We use $\mathbf{X}_{ij}^L = [X_{ij1}^L; \dots; X_{ijq}^L]$ to denote a q -dimensional column vector, containing all q factors of interest for the left profile for respondent i in the j th evaluation where $q \leq p$. We define \mathbf{X}_{ij}^R similarly for the right profile. In addition, we use $\mathbf{X}_{ij} = [\mathbf{X}_{ij}^L; \mathbf{X}_{ij}^R]$ as a column vector of length $2q$ to represent the main factors of interest from two profiles together. Lastly, the remaining $(p - q)$ factors are denoted by $\mathbf{Z}_{ij} = [\mathbf{Z}_{ij}^L; \mathbf{Z}_{ij}^R]$, where each term is similarly defined. For example in the immigration conjoint experiment, if the main factor of interest is “country of origin”, the other factors include “education” and “profession”.

As done for the outcome variable, we stack all evaluation-specific factors to define respondent-level factor matrices, which are further combined to yield the factor matrix $\mathbf{X} = [\mathbf{X}_1; \mathbf{X}_2; \dots; \mathbf{X}_n]$ and $\mathbf{Z} = [\mathbf{Z}_1; \mathbf{Z}_2; \dots; \mathbf{Z}_n]$ of dimension $nJ \times 2q$ and $nJ \times 2(p - q)$, respectively, where $\mathbf{X}_i = [\mathbf{X}_{i1}^T; \mathbf{X}_{i2}^T; \dots; \mathbf{X}_{iJ}^T]$ and $\mathbf{Z}_i = [\mathbf{Z}_{i1}^T; \mathbf{Z}_{i2}^T; \dots; \mathbf{Z}_{iJ}^T]$ are matrices of dimension $J \times 2q$ and $J \times 2(p - q)$, respectively. Lastly, we also stack all respondent characteristics $\mathbf{V} = [\mathbf{V}_1; \mathbf{V}_2; \dots; \mathbf{V}_n]$ of dimension $nJ \times r$.

Finally, we use $\mathbf{Y}(\mathbf{x}, \mathbf{z})$ to denote the nJ -dimensional vector of the potential outcomes when $\mathbf{X} = \mathbf{x}$ and $\mathbf{Z} = \mathbf{z}$. This notation implies that we avoid assuming the no interference effect in the Stable Unit Treatment Value Assumption (SUTVA) since our vector of potential outcomes is a function of the entire set of treatments \mathbf{X} and \mathbf{Z} (Rubin, 1990). We assume a super-population framework, where the potential outcomes $\mathbf{Y}(\mathbf{x}, \mathbf{z})$ are assumed to be drawn from a population of infinite size. In Appendix B, we discuss how our framework is related to a finite-population framework, which is the basis of Fisher’s randomization test. In conjoint analysis, the profile attributes are randomized according to a known distribution, $P(\mathbf{X}, \mathbf{Z})$. Our framework is general, allowing for any randomization distribution. For example, some may randomize each factor independently using complete randomization whereas others may induce dependency among factors by removing a certain set of attribute combinations. In general, the randomization of the factors implies the following independence relation,

$$\mathbf{Y}(\mathbf{x}, \mathbf{z}) \perp\!\!\!\perp (\mathbf{X}, \mathbf{Z}) \text{ for all } \mathbf{x} \in \mathcal{X}, \text{ and } \mathbf{z} \in \mathcal{Z}, \quad (1)$$

where we use \mathcal{X} and \mathcal{Z} to represent the support of \mathbf{X} and that of \mathbf{Z} , respectively (see Chapter 3.6 of (Imbens and Rubin, 2015)).

3.2 The Conditional Randomization Test

The Conditional Randomization Test (CRT) is an assumption-free approach that combines design-based inference with flexible machine learning algorithms. For ease of presentation, we first introduce the CRT without incorporating the respondent characteristics \mathbf{V} and then return in Section 3.6 to show how \mathbf{V} can be incorporated in all the proposed methods. The CRT allows us to examine whether the factors of interest \mathbf{X} change the response \mathbf{Y} while holding the other factors \mathbf{Z} constant. Specifically, we test the following null hypothesis,

$$H_0 : \mathbf{Y}(\mathbf{x}, \mathbf{z}) \stackrel{d}{=} \mathbf{Y}(\mathbf{x}', \mathbf{z}) \text{ for all } \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \text{ and } \mathbf{z} \in \mathcal{Z}, \quad (2)$$

where we use $\stackrel{d}{=}$ to denote distributional equality. As a reminder, H_0 states that our entire *vector* of potential outcomes are equal in distribution for any values of \mathbf{X} . Our alternative hypothesis states that \mathbf{X} affects \mathbf{Y} in some way while keeping \mathbf{Z} unchanged. This is formalized as,

$$H_1 : \mathbf{Y}(\mathbf{x}, \mathbf{z}) \stackrel{d}{\neq} \mathbf{Y}(\mathbf{x}', \mathbf{z}) \text{ for some } \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \text{ and } \mathbf{z} \in \mathcal{Z}. \quad (3)$$

We emphasize that the null hypothesis defined in Equation (2) implies the absence of any causal effects involving the main factor(s) of interest. For example, the null hypothesis is false if \mathbf{X} affects \mathbf{Y} for any individual respondent or subgroup of respondents. Similarly, the null hypothesis does not hold if \mathbf{X} influences \mathbf{Y} only when \mathbf{Z} takes a certain set of values. Thus, the null hypothesis precludes any heterogeneous or interaction effects as well.

Contrast this hypothesis test formulation with that of the standard AMCE-based analysis, which asks whether each factor of interest \mathbf{X}_i matters *on average*. More specifically, Hainmueller, Hopkins and Yamamoto assume each individual's potential outcome is only a function of its own profile task, i.e., $Y_{ij}(\mathbf{X}, \mathbf{Z}) = Y_{ij}(\mathbf{X}_{ij}, \mathbf{Z}_{ij})$, and computes the marginal importance of \mathbf{X}_i by averaging each individual potential outcome over \mathbf{Z}_i as well as the respondents, which are assumed to be exchangeable, leading to the following null hypothesis,

$$H_0^{\text{AMCE}} : \mathbb{E}\{Y_{ij}(\mathbf{x}, \mathbf{Z}_{ij})\} = \mathbb{E}\{Y_{ij}(\tilde{\mathbf{x}}, \mathbf{Z}_{ij})\}, \quad (4)$$

where \mathbf{x} and $\tilde{\mathbf{x}}$ are the specified values of the main factors and the expectation is taken over \mathbf{Z}_{ij} (other factors) and the respondents. As briefly explained in Section 2.3, the limitation of the AMCE-based approach is that averaging over other factors can mask important causal interaction and heterogeneity.

We now establish the equivalence between the null hypothesis about the potential outcomes defined in Equation (2) and the conditional independence relation among observed variables. This result allows us to use the CRT, which is a general assumption-free methodology for testing conditional independence relations in designed experiments (Candès et al., 2018). We state this result as the following theorem whose proof is given in Appendix A.

Theorem 3.1 (Equivalence). The null hypothesis defined in Equation (2) is equivalent to the following conditional independence hypothesis under the randomization assumption of Equation (1),

$$H_0^{\text{CRT}} : \mathbf{Y} \perp\!\!\!\perp \mathbf{X} \mid \mathbf{Z}.$$

Algorithm 1: Conditional Randomization Test (CRT)

Input: Data $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, test statistic $T(\mathbf{x}, \mathbf{y}, \mathbf{z})$, total number of re-samples B , conditional distribution $\mathbf{X} \mid \mathbf{Z}$;

for $b = 1, 2, \dots, B$ **do**

 └ Sample $\mathbf{X}^{(b)}$ from the distribution of $\mathbf{X} \mid \mathbf{Z}$ conditionally independently of \mathbf{X} and \mathbf{Y} ;

Output: p -value $:= \frac{1}{B+1} \left[1 + \sum_{b=1}^B \mathbb{1}\{T(\mathbf{X}^{(b)}, \mathbf{Y}, \mathbf{Z}) \geq T(\mathbf{X}, \mathbf{Y}, \mathbf{Z})\} \right]^4$

The CRT is a general assumption-free methodology in designed experiments that produces exact p -values without asymptotic approximation. The CRT combines the advantages of both design-based and model-based approaches by enabling the use of any test statistic, including ones based on complex machine learning (ML) algorithms, while making no modeling assumptions. In the conjoint analysis literature, researchers have used traditional regression modeling (e.g., Barone, Lombardo and Tarantino, 2007; Hauber et al., 2016; McFadden, 1973) and more recently modern ML algorithms (e.g., Egami and Imai, 2019; Abramson et al., 2020). However, the validity of these analyses critically depends on modeling assumptions, parameter tuning, and/or asymptotic approximation. In contrast, the CRT assumes nothing about the conditional distribution of the outcome \mathbf{Y} given (\mathbf{X}, \mathbf{Z}) . Indeed, it does not even require the data to be independently or identically distributed, a property which we use later to test carryover and profile order effects. The only requirement is the specification of the conditional distribution of \mathbf{X} given \mathbf{Z} , which is readily available from the experimental design of conjoint analysis. Although the power of the CRT critically depends on the test statistic, the CRT always controls type 1 error no matter what the true model is. This contrasts with other model based approaches that require modeling assumptions to be valid (see Appendix J for more details).

Algorithm 1 summarizes the general procedure used to compute the exact p -value for the CRT. Note that if \mathbf{X} and \mathbf{Z} are independently randomized, as is often the case, one can simply sample $\mathbf{X}^{(b)}$ from the marginal distribution of \mathbf{X} . If, on the other hand, certain combinations of \mathbf{X} and \mathbf{Z} values (e.g., doctor without a college degree in the immigration conjoint experiment) are excluded, then we must use the appropriate conditional distribution of \mathbf{X} given \mathbf{Z} . Critically, Algorithm 1 is valid for complicated experimental designs, so long as one can sample from the conditional distribution \mathbf{X} given \mathbf{Z} . The CRT can be computationally intensive since it requires computing the test statistic T a total of $B + 1$ times. However these computations can easily be parallelized. Furthermore, recent works (Tansey et al., 2018; Liu et al., 2020) have shown that certain test statistic constructions also alleviate the need for these computations. In Appendix I, we detail several tricks that can be used to dramatically reduce the computation time when implementing the CRT. For the main application results presented in Section 5 (first column of Table 1), we note that the parallelized computational time was approximately six minutes with 50 cores to calculate each p -value with $B = 2,000$. Our software package makes it easy for practitioners to use multiple cores and provides a step-by-step instruction for using many cores on Amazon Web Services.⁵

⁴We add one to the numerator and denominator so that the distribution of the p -value is stochastically dominated by the uniform distribution as suggested by (Candès et al., 2018).

⁵The detailed instructions and example use cases can be found in a vignette of our open-source software package in <https://cran.r-project.org/web/packages/CRTConjoint/vignettes/CRTConjoint.html>.

The p -value of the CRT is valid⁶ regardless of sample size and test statistic (Candès et al., 2018). To see this, it suffices to recognize that under the null hypothesis all $B + 1$ test statistics, $T(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, $T(\mathbf{X}^{(1)}, \mathbf{Y}, \mathbf{Z})$, \dots , $T(\mathbf{X}^{(B)}, \mathbf{Y}, \mathbf{Z})$, are exchangeable given (\mathbf{Y}, \mathbf{Z}) . While any test statistic produces a valid p -value under the CRT, the choice of test statistic determines the statistical power. We now turn to this practically important consideration.

3.3 Test Statistics

To obtain a powerful test statistic that does not mask important interactions, we consider a test statistic based on the Lasso logistic regression with hierarchical interactions, or HierNet (Bien, Taylor and Tibshirani, 2013). We note that if researchers wish to target main effects without considering interactions, the F -statistic from a standard linear regression of the response \mathbf{Y} on \mathbf{X} is a reasonable test statistic to use. HierNet allows for the regularization of all possible two-way interaction terms while respecting their hierarchy. Specifically, HierNet constrains the two-way interaction effects to be smaller in magnitude than their corresponding main effects. For example, this implies that a two-way interaction effect will be set to zero if its relevant main effects are all zero. A stricter regularization on the interactions is desirable because the space of possible two-way interactions is large and grows quadratically, and we expect many to be indistinguishable from zero. Lastly, when fitting HierNet, we use the dummy variable encoding (i.e., each factor level is represented by its own dummy variable) but do not omit the baseline level. We can fit this overparameterized model because of the regularization of HierNet. The primary advantage of this approach is that the results are no longer dependent on the choice of baseline levels (Egami and Imai, 2019).

We begin by considering the simplest case where we have a single main factor of interest \mathbf{X} ($q = 1$). Without loss of generality, we assume that this is the first factor among the total of p factors. There are two types of interaction effects to consider (de la Cuesta, Egami and Imai, 2022). First, a within-profile interaction effect represents the interaction between one level of the main factor and another level of a different factor within the same profile. Second, a between-profile interaction effect represents the factor interaction between two profiles (left versus right) that are being compared under the forced choice design.

Our proposed test statistic is based on the sum of relevant squared main and interaction effects after subtracting their respective means,

$$T_{\text{HierNet}} = \underbrace{\sum_{k=1}^{K_1} (\hat{\beta}_k - \bar{\beta})^2}_{\text{main effects}} + \underbrace{\sum_{\ell=2}^p \sum_{k=1}^{K_1} \sum_{k'=1}^{K_\ell} (\hat{\gamma}_{1\ell kk'} - \bar{\gamma}_{1\ell k'})^2}_{\text{within-profile interaction effects}} + \underbrace{\sum_{\ell=1}^p \sum_{k=1}^{K_1} \sum_{k'=1}^{K_\ell} (\hat{\delta}_{1\ell kk'} - \bar{\delta}_{1\ell k'})^2}_{\text{between-profile interaction effects}}, \quad (5)$$

where $\hat{\beta}_k$ is the estimated main effect coefficient for the k th level of our factor of interest \mathbf{X} with $\bar{\beta}$ denoting the average of these estimated main effect coefficients, and $\hat{\gamma}_{1\ell kk'}$ and $\hat{\delta}_{1\ell kk'}$ represent the estimated within-profile and between-profile interaction effect coefficients between the k th level of the factor of interest \mathbf{X} and the k' th level of the ℓ th factor, respectively. Similar to the main effects, $\bar{\gamma}_{1\ell k'}$ and $\bar{\delta}_{1\ell k'}$ denote the averages of their corresponding estimated interaction effect coefficients. We do not consider third or higher order interactions because of the typical sample size in a conjoint

⁶That is, under H_0 , $P(p\text{-value} \leq \alpha) \leq \alpha$ for all $\alpha \in [0, 1]$.

experiment and a lack of powerful methods to detect such interactions. However, in Section 5 we illustrate how to incorporate third order interactions when prior substantive knowledge is available.

This test statistic can be easily generalized to the setting where there is more than one factor of interest ($q > 1$). In such a case, we simply compute Equation (5) for each factor of interest, and then sum the resulting values to arrive at the final test statistic. T_{HierNet} aims to capture any differential effects the levels of \mathbf{X} have on the response through their main effects and relevant interaction effects. For example, suppose that \mathbf{X} is “gender” in the political candidate conjoint experiment (Section 2.2). Under H_0 , we would expect all main effects and any interaction effects of *male* and *female* to be roughly equivalent, thus making T_{HierNet} close to zero. However, suppose that *male* candidates with a certain “education level” were favored more than *female* candidates with a certain “education level”. Then, we would expect these interactions to differ, making T_{HierNet} further from zero.

We use cross-validation⁷ to choose the value of HierNet’s tuning parameter, which controls the degree of regularization. In addition, when the sample size and the number of factors are large, fitting HierNet can be computationally demanding. To alleviate this issue, we propose computational speedups of the HierNet test statistics, which are detailed in Appendix I. In particular, we drop \mathbf{X} when fitting the HierNet tuning parameter via cross-validation. Since this computationally expensive step does not depend on \mathbf{X} , we do not need to re-run it for each \mathbf{X}^b .

So far, we have constructed our test statistic as if there is no profile order effect. This implies that the effects of each factor do not depend on whether it belongs to the left or right profile. Formally, we have imposed the following symmetry constraints in our HierNet test statistic,

$$\hat{\beta}_k = \hat{\beta}_k^L = -\hat{\beta}_k^R, \quad \hat{\gamma}_{1\ell k k'} = \hat{\gamma}_{\ell\ell' k k'}^L = -\hat{\gamma}_{\ell\ell' k k'}^R, \quad \hat{\delta}_{\ell\ell' k k'} = -\hat{\delta}_{\ell'\ell k k'}, \quad (6)$$

where the superscripts L and R denote the left and right profile effects, respectively. $\hat{\delta}_{\ell\ell' k k'}$ denotes the between profile interaction between the k th level of factor ℓ in the left profile with the k' th level of factor ℓ' in the right profile.⁸ The signs of the estimated coefficients reflect the fact that the response variable \mathbf{Y} is recorded as 1 if the left profile is chosen and as 0 if the right profile is selected. These constraints reduce the dimension of parameters to be estimated by half.

Importantly, the validity of the proposed tests does not depend on whether the assumption of no profile order effect holds. Through simulations, Figure 4 of Appendix E shows that these constraints can significantly increase statistical power when there is no profile order effect. To incorporate this symmetry constraint, we append another copy of the dataset below the original data set, where the appended copy is identical to the original dataset except that the order of left and right profiles is flipped and the response variable is transformed as $\mathbf{1} - \mathbf{Y}$ before fitting HierNet (see Appendix E for details). In Section 3.5, we show how to use the CRT for testing the validity of the assumption of no profile order effect.

Because the validity of the CRT does not depend on modeling assumptions, one can incorporate a variety of assumptions into test statistics. In general, test statistics have a greater statistical power

⁷The CRT remains valid if used with cross-validation so long as the resampled test statistics based on \mathbf{X}^b similarly uses cross-validation. This ensures exchangeability.

⁸Equation (6) also implies that the between-profile interactions in Equation (5) for the same factor obey $\hat{\delta}_{\ell\ell' k k'} = -\hat{\delta}_{\ell\ell' k' k}$ for any factor ℓ and levels k, k' . In particular this implies that $\hat{\delta}_{\ell\ell k k} = 0$, i.e., between-profile interactions of the same factor and same level are zero, while $\hat{\delta}_{\ell\ell' k k'}$ are counted twice in Equation (5), i.e., between-profile interactions of the same factor and levels k and k' are counted twice in Equation (5).

if the assumptions hold in the true (unknown) data generating process. Therefore, as much as possible the choice of test statistic should reflect researchers’ substantive knowledge as illustrated in one of our empirical analyses (see Section 5.2).

3.4 Generalization of the Null Hypothesis and Test Statistic

Researchers are often interested in testing only a few levels of interest as opposed to testing the whole factor. Yet, simply dropping the observations that correspond to those irrelevant factor levels can lead to a loss of statistical power. An advantage of the formulation described below is that we can retain all observations including those whose factor levels are irrelevant, which can improve statistical power. For example, suppose we are interested in how respondents differentiate immigrants from *Mexico* and *Germany*. If the way in which respondents differentiate between immigrants from *Mexico* and those from *China* is different from how they distinguish between immigrants from *Germany* and those from *China*, then this implies that the respondents are viewing immigrants from *Mexico* differently than those from *Germany*. Therefore, detecting any differences for even the irrelevant levels may help improve the statistical power.

Here, we generalize the null hypothesis and test statistic, given in Equations (2) and (5), so that the methodology can accommodate any combinations of factor levels. We introduce a coarsening function h that groups factor levels of interest while assigning other factor levels to themselves. Formally, this coarsening function is defined as $h : \mathcal{X} \mapsto \tilde{\mathcal{X}}$, where $|\mathcal{X}| \geq |\tilde{\mathcal{X}}|$. Thus, for our aforementioned immigration example, h will assign the same value to immigrants from *Mexico* and *Germany* while leaving all other combinations mapped to different values.

Under this setup, we can test the null hypothesis that specific levels within \mathbf{X} do not affect the potential outcome in any way. Formally,

$$H_0^{\text{General}} : \mathbf{Y}(\mathbf{x}, \mathbf{z}) \stackrel{d}{=} \mathbf{Y}(\mathbf{x}', \mathbf{z}) \text{ for all } \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \text{ such that } h(\mathbf{x}) = h(\mathbf{x}') \text{ and } \mathbf{z} \in \mathcal{Z}. \quad (7)$$

The condition $h(\mathbf{x}) = h(\mathbf{x}')$ enables the comparison of the factor levels of interest alone. Additionally, H_0 is a special case of H_0^{General} when the coarsening function h is the identity function. Finally, applying the same argument as the one used to prove Theorem 3.1, it can be shown that H_0^{General} is equivalent to the following conditional independence relation,

$$\mathbf{Y} \perp\!\!\!\perp \mathbf{X} \mid h(\mathbf{X}), \mathbf{Z}. \quad (8)$$

To test this null hypothesis, we first fit the same HierNet with the main effects and two-way interaction effects. To incorporate the coarsening function h , our test statistic takes the same form as the one given in Equation (5) but is based only on the estimated coefficients that correspond to the factor levels of the group induced by the h function, i.e., *Mexico* and *Germany* in the above example. Appendix F.1 contains further details of testing the general null hypothesis and the corresponding CRT algorithm. In Section 5.1, we also provide an example of applying H_0^{General} that contains more details on this test statistic. Under our framework, we do not need to drop observations that have irrelevant factor levels, thus increasing statistical power as mentioned above. Finally, Appendix C details how to further generalize H_0^{General} when a researcher is interested in grouping factor levels, i.e., combining levels *France*, *Germany*, and *Poland* into one level *Europe* when testing H_0^{General} . Although coarsening via h also involved “grouping” levels, the grouping described in Appendix C aggregates factor levels to allow comparison between higher-level categories while the coarsening

function h allows us to focus our hypothesis test only on differences between a subset of factor levels of interest.

3.5 Testing the Regularity Assumptions of Conjoint Analysis

To further demonstrate the flexibility of the CRT, we also show how to use the CRT for testing the validity of several commonly made assumptions of conjoint analysis. Although we were interested in *rejecting* H_0 above, we are interested in *accepting* the null hypothesis for the hypotheses presented in this section. Therefore, we propose test statistics that are designed to be reasonably powerful for general settings in conjoint analysis.

Profile Order Effect. The assumption of no profile order effect states that changing the order of profiles, i.e., left versus right, does not affect the actual profile chosen (since the value of \mathbf{Y} corresponds to whether the left or right profile is chosen, \mathbf{Y} should be recoded as $\mathbf{1} - \mathbf{Y}$ when the profile order is changed). We denote the potential outcome $Y_{ij}(\mathbf{x}_{ij}^L, \mathbf{x}_{ij}^R, \mathbf{z}_{ij}^L, \mathbf{z}_{ij}^R)$, which is now a function of left and right profiles. Although not necessary, we assume here no interference between responses for notational clarity (see Appendix F.2 for the general case). Lastly, we use \mathcal{X}_{ind} and \mathcal{Z}_{ind} to denote the support of $[\mathbf{x}_{ij}^L; \mathbf{x}_{ij}^R]$ and that of $[\mathbf{z}_{ij}^L; \mathbf{z}_{ij}^R]$ respectively, representing the support of factors used in each individual’s evaluation (hence “ind” in the subscript).

We formally state the assumption of no profile order effect as the following null hypothesis that reordering of the left and right profiles has no effect on the adjusted response:

$$H_0^{\text{Order}} : Y_{ij}(\mathbf{x}_{ij}^L, \mathbf{x}_{ij}^R, \mathbf{z}_{ij}^L, \mathbf{z}_{ij}^R) \stackrel{d}{=} 1 - Y_{ij}(\mathbf{x}_{ij}^R, \mathbf{x}_{ij}^L, \mathbf{z}_{ij}^R, \mathbf{z}_{ij}^L), \text{ for all } i, j, [\mathbf{x}_{ij}^L; \mathbf{x}_{ij}^R] \in \mathcal{X}_{\text{ind}}, [\mathbf{z}_{ij}^L; \mathbf{z}_{ij}^R] \in \mathcal{Z}_{\text{ind}}.$$

We modify the HierNet test statistic in Equation (5) with the same HierNet fit on $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ but without enforcing the constraints in Equation (6) as,

$$\begin{aligned} T_{\text{HierNet}}^{\text{Order}}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) &= \sum_{\ell=1}^p \sum_{k=1}^{K_{\ell}} (\hat{\beta}_{\ell k}^L + \hat{\beta}_{\ell k}^R)^2 + \sum_{\ell=1}^p \sum_{\substack{\ell'=1 \\ \ell' \neq \ell}}^p \sum_{k=1}^{K_{\ell}} \sum_{k'=1}^{K_{\ell'}} (\hat{\gamma}_{\ell \ell' k k'}^L + \hat{\gamma}_{\ell \ell' k k'}^R)^2 \\ &\quad + \sum_{\ell=1}^p \sum_{\ell'=1}^p \sum_{k=1}^{K_{\ell}} \sum_{k'=1}^{K_{\ell'}} (\hat{\delta}_{\ell \ell' k k'} + \hat{\delta}_{\ell' \ell k' k})^2. \end{aligned}$$

Since the symmetry constraints given in Equation (6) must hold under H_0^{Order} , a large value of this test statistic indicates a potential violation of the null hypothesis. To conduct the CRT for testing H_0^{Order} , we resample and recompute our test statistics. Appendix F.2 provides details about the testing procedure.

Carryover Effect. Researchers also often rely on the assumption of no carryover effect (Hainmueller, Hopkins and Yamamoto, 2014). The assumption states that the order of the J evaluations each respondent performs has no effect on the outcomes. This assumption is violated, e.g., if respondents use information from their previous evaluations when assessing a given pair of profiles. To test this carryover effect, we assume no interference across respondents but consider potential interference across evaluations within each respondent.

Let $\mathbf{x}_{i,1:(j-1)}$ represent all the profile attributes that were presented to respondent i from the first evaluation to the $(j - 1)$ th evaluation. Then, the potential outcome can be written as a function of both current and previous profiles, i.e., $Y_{ij}(\mathbf{x}_{i,1:(j-1)}, \mathbf{z}_{i,1:(j-1)}, \mathbf{x}_{ij}, \mathbf{z}_{ij})$ for $j \geq 2$, where we assume no interference between respondents but allow interference within a respondent. Our null hypothesis is that, for a given evaluation $j \geq 2$, the response Y_{ij} is independent of all the previous profiles conditional on the current profiles:

$$H_0^{\text{Carryover}} : Y_{ij}(\mathbf{x}_{i,1:(j-1)}, \mathbf{z}_{i,1:(j-1)}, \mathbf{x}_{ij}, \mathbf{z}_{ij}) \stackrel{d}{=} Y_{ij}(\mathbf{x}'_{i,1:(j-1)}, \mathbf{z}'_{i,1:(j-1)}, \mathbf{x}_{ij}, \mathbf{z}_{ij}),$$

where for all $i \geq 1, j \geq 2$, $\mathbf{x}_{i,1:(j-1)}, \mathbf{x}'_{i,1:(j-1)} \in \mathcal{X}_{\text{ind}}^{j-1}$, $\mathbf{z}_{i,1:(j-1)}, \mathbf{z}'_{i,1:(j-1)} \in \mathcal{Z}_{\text{ind}}^{j-1}$, $\mathbf{x}_{ij} \in \mathcal{X}_{\text{ind}}$, $\mathbf{z}_{ij} \in \mathcal{Z}_{\text{ind}}$ with $\mathcal{X}_{\text{ind}}^{j-1}$ and $\mathcal{Z}_{\text{ind}}^{j-1}$ denoting the support of $\mathbf{x}_{i,1:(j-1)}$ and that of $\mathbf{z}_{i,1:(j-1)}$, respectively.

We test this null hypothesis by using a test statistic that targets whether the immediately preceding evaluation affects the current evaluation. We believe targeting the lag-1 effect in the test statistic is reasonable because if a carryover effect exists, respondents are likely to be affected most by the immediately preceding evaluation. For example, if respondents believe that they have placed too much weight on profiles' professions in the previous evaluation, they might decide to rely on the current profiles' "country of origin" factor more than its "profession" factor in order to balance across evaluations. Under this scenario, we would expect a significant interaction between previous profiles' "profession" factor and current profiles' "country of origin" factor.

We modify the test statistic given in Equation (5) in the following way. Suppose that J is even (if J is odd, simply consider $J - 1$ evaluations). We first define a new response vector $\mathbf{Y}_i^* = [Y_{i2}; Y_{i4}; \dots; Y_{iJ}]$ by taking every other evaluation. Similarly, we can define new factors of interest $\mathbf{X}_i^* = [[\mathbf{X}_{i1}; \mathbf{Z}_{i1}]^\top; [\mathbf{X}_{i3}; \mathbf{Z}_{i3}]^\top, \dots; [\mathbf{X}_{i,J-1}; \mathbf{Z}_{i,J-1}]^\top]$ and a new set of conditioning variables $\mathbf{Z}_i^* = [[\mathbf{X}_{i2}; \mathbf{Z}_{i2}]^\top; [\mathbf{X}_{i4}; \mathbf{Z}_{i4}]^\top, \dots; [\mathbf{X}_{iJ}; \mathbf{Z}_{iJ}]^\top]$. We then fit HierNet with the new response $\mathbf{Y}^* = [\mathbf{Y}_1^*; \mathbf{Y}_2^*; \dots; \mathbf{Y}_n^*]$ on $(\mathbf{X}^*, \mathbf{Z}^*)$, where $\mathbf{X}^* = [\mathbf{X}_1^*; \mathbf{X}_2^*; \dots; \mathbf{X}_n^*]$ and \mathbf{Z}^* is defined similarly.

For this particular scenario, HierNet will estimate all main effects and interaction effects of \mathbf{Z}^* , and the interaction effects between \mathbf{X}^* and \mathbf{Z}^* , which are of primary interest. To increase the power of the test, we set all main and interaction effects of \mathbf{X}^* to zero since we do not expect the previous profile alone to impact the respondent's choice. Furthermore, we also do not expect the interaction effects between \mathbf{X}^* and \mathbf{Z}^* to differ, depending on the ordering (left versus right) of the relevant factors. Therefore, we enforce all these interaction effects to have equal magnitude as done in Equation (6) (see Appendix E for further details). This leads to the following test statistic,

$$T_{\text{HierNet}}^{\text{Carryover}}(\mathbf{X}^*, \mathbf{Y}^*, \mathbf{Z}^*) = \sum_{\ell=1}^p \sum_{\ell'=1}^p \sum_{k=1}^{K_\ell} \sum_{k'=1}^{K_{\ell'}} \hat{\gamma}_{\ell\ell'kk'}^2,$$

where $\gamma_{\ell\ell'kk'}$ represents the coefficient of an interaction term between the k th level of the ℓ th factor of the profile used in the previous evaluation and the k' th level of the ℓ' th factor of the profile used in the current evaluation. Appendix F.2 explains how to resample the test statistic in this setting.

Fatigue Effect. Researchers may be concerned that a respondent performing a large number of conjoint evaluations may experience the "fatigue effect," resulting in a declining quality of responses. Recently, Bansak et al. (2018) conducted an empirical study, in which they examine how the pattern of responses depends on the number of evaluations each respondent performs. While

a typical conjoint experiment asks each respondent to carry out 5 to 10 evaluations, the authors increase this number up to 30 evaluations. They find that the results are robust to increasing the number of evaluations. Here, we show how to use the CRT to formally test the presence of the fatigue effect.

Similar to the carryover effect, we investigate whether there is a fatigue effect within each respondent’s potential outcome $Y_{ij}(\mathbf{x}_{ij}, \mathbf{z}_{ij})$, where we again assume no interference effect as done when testing no profile order effect. We test the following null hypothesis that the potential outcome is unaffected if respondent i evaluated the same pair of profiles $(\mathbf{x}_{ij}, \mathbf{z}_{ij})$ but at a later or earlier evaluation $j' \neq j$:

$$H_0^{\text{Fatigue}} : Y_{ij}(\mathbf{x}_{ij}, \mathbf{z}_{ij}) \stackrel{d}{=} Y_{ij'}(\mathbf{x}_{ij}, \mathbf{z}_{ij}) \text{ for all } i, j, j', \mathbf{x}_{ij} \in \mathcal{X}_{\text{ind}}, \text{ and } \mathbf{z}_{ij} \in \mathcal{Z}_{\text{ind}}.$$

We propose a similar HierNet test statistic that reflects a scenario where respondents will only pay attention to a shrinking number of factors as they rate more profiles. In this case, we would expect interactions between the factors and the evaluation order index $\mathbf{F} = (F_1, F_2, \dots, F_n)$, which represents an nJ -dimensional integer vector with $\mathbf{F}_i = (1, 2, \dots, J)$ for all $i = 1, 2, \dots, n$. Again, for the sake of statistical power, we impose the absence of profile order effects on HierNet as done in Equation (5). Our proposed test statistic is the following from a HierNet fit of \mathbf{Y} on $(\mathbf{X}, \mathbf{Z}, \mathbf{F})$,

$$T_{\text{HierNet}}^{\text{Fatigue}}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{F}) = \sum_{\ell=1}^p \sum_{k=1}^{K_\ell} \hat{\gamma}_{\ell k}^2,$$

where $\hat{\gamma}_{\ell k}$ represents the coefficient of an interaction term between \mathbf{F} and level k of factor ℓ . Appendix F.2 shows how to resample and recompute the test statistics to test H_0^{Fatigue} .

3.6 Incorporating Respondent Characteristics

In conjoint experiments, researchers often expect factors of interest to interact strongly with respondent characteristics (Ono and Burden, 2018; Hainmueller and Hopkins, 2015; Newman and Malhotra, 2019). It is possible to exploit this fact when applying the CRT by directly incorporating respondent characteristics, \mathbf{V} , into the test statistic. Doing so can substantially increase the statistical power.

We incorporate respondent characteristics in the CRT procedure by appending \mathbf{V} to \mathbf{Z} and holding both (\mathbf{Z}, \mathbf{V}) constant. Since respondent characteristics are not randomized factors, unlike \mathbf{Z} , \mathbf{V} is not guaranteed to be independent of the potential outcomes. We can test the following causal null hypothesis that conditions upon \mathbf{V} ,

$$H_0 : \mathbf{Y}(\mathbf{x}, \mathbf{z}) \stackrel{d}{=} \mathbf{Y}(\mathbf{x}', \mathbf{z}) \mid \mathbf{V} \text{ for all } \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \text{ and } \mathbf{z} \in \mathcal{Z}. \quad (9)$$

Theorem 3.1 can be easily extended to show that this null hypothesis is equivalent to the conditional independence relation $\mathbf{Y} \perp\!\!\!\perp \mathbf{X} \mid \mathbf{Z}, \mathbf{V}$. Algorithm 1 also stays the same except we sample \mathbf{X}^b from the distribution of $\mathbf{X} \mid (\mathbf{Z}, \mathbf{V})$ and our test statistic is now a function of $(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{V})$.

A major benefit of incorporating respondent characteristics is the ability to capture respondent characteristic interactions into the test statistic. Consequently, we incorporate an additional predic-

tor \mathbf{V} when fitting HierNet and modify the HierNet test statistic in Equation (5) as,

$$\begin{aligned}
T_{\text{HierNet}} = & \underbrace{\sum_{k=1}^{K_1} (\hat{\beta}_k - \bar{\beta})^2}_{\text{main effects}} + \underbrace{\sum_{\ell=2}^p \sum_{k=1}^{K_1} \sum_{k'=1}^{K_\ell} (\hat{\gamma}_{1\ell kk'} - \bar{\gamma}_{1\ell k'})^2}_{\text{within-profile interaction effects}} \\
& + \underbrace{\sum_{\ell=1}^p \sum_{k=1}^{K_1} \sum_{k'=1}^{K_\ell} (\hat{\delta}_{1\ell kk'} - \bar{\delta}_{1\ell k'})^2}_{\text{between-profile interaction effects}} + \underbrace{\sum_{m=1}^r \sum_{k=1}^{K_1} \sum_{w=1}^{L_m} (\hat{\xi}_{1mkw} - \bar{\xi}_{1mw})^2}_{\text{respondent interaction effects}}, \tag{10}
\end{aligned}$$

where $\hat{\xi}_{1mkw}$ represents the interaction between the k th level of our factor of interest \mathbf{X} and the w th level of the m th factor of a respondent characteristic. If a respondent characteristic is a numeric variable, we could either coarsen it into a factor variable or directly include it in the model as a numeric variable.⁹ Again, $\bar{\xi}_{1mw}$ denotes the average of these estimated interaction coefficients. Similar to Equation (10), we add the additional constraint that $\hat{\xi}_{\ell mkw}^L = \hat{\xi}_{\ell mkw}^R = \hat{\xi}_{\ell mkw}$, where the superscript L and R similarly denotes the left and right profile effects. Finally, we modify $T_{\text{HierNet}}^{\text{Order}}$ similarly by adding $\sum_{\ell=1}^p \sum_{m=1}^r \sum_{k=1}^{K_\ell} \sum_{w=1}^{L_m} (\hat{\xi}_{\ell mkw}^L + \hat{\xi}_{\ell mkw}^R)^2$ to the original test statistic in $T_{\text{HierNet}}^{\text{Order}}$ to account for respondent characteristic interactions. For all empirical applications in Section 5, we incorporate \mathbf{V} and use the modified test statistics to test H_0 and the no profile order effect.

4 Simulation Studies

A primary advantage of the CRT is that it can yield powerful statistical tests by incorporating machine learning algorithms to capture complex interactions in high dimensions. The CRT achieves this while maintaining the finite sample validity of the resulting p -values. In this section, we conduct simulation studies to show that the CRT with the HierNet test statistic can be substantially more powerful than the AMCE-based test.

For simplicity, our simulation setting assumes one evaluation for each respondent ($J = 1$). We show additional simulations with multiple evaluations with respondent random effects in Figure 7 of Appendix G.4. In addition, there is only one main factor of interest \mathbf{X} ($q = 1$) and ten other factors \mathbf{Z} and no \mathbf{V} . Each factor is assumed to have two levels and is independently and uniformly randomized. To clearly separate main and interaction effects, we use the sum-to-zero constraint by coding each binary factor as $(-0.5, 0.5)$. Our response model, $\Pr(Y_i = 1 \mid X_i, \mathbf{Z}_i)$, follows a logistic regression that includes main effects for (\mathbf{X}, \mathbf{Z}) and within-profile and between-profile interaction effects between \mathbf{X} and \mathbf{Z} and among \mathbf{Z} . Our setup also has no profile order effects.

Our goal is to examine the difference in statistical power between the CRT and AMCE-based tests in the presence of interactions. Thus, we vary the size of the interaction effect as well as the number of non-zero interaction effects between \mathbf{X} and \mathbf{Z} . Throughout all of our simulations, we use an effect size of 0.1 for all main effects for (\mathbf{X}, \mathbf{Z}) and an effect size of 0.05 for all interactions among \mathbf{Z} and fix the sample size to $n = 3,000$. We also make all non-zero interactions between \mathbf{X} and \mathbf{Z} equal in size for all simulations. As shown in Appendix G.2, the power of the CRT remains

⁹HierNet standardizes all variables when performing the fit. Therefore, even if a numeric variable is on a different scale, all estimated coefficients remain comparable.




Figure 3: The figure shows how the power of the CRT and AMCE-based tests varies as the size of interaction effects (left plot) or the number of non-zero interaction effects (right) increases. The AMCE-based test (red circles) is based on the t -test from the estimated regression coefficient. The CRT uses the HierNet test statistic given in Equation (5). The sample size is $n = 3,000$. Finally, the standard errors are negligible with a maximum value of 0.016.

identical if there are heterogeneous interaction effects with varying sizes. Appendix G provides further details of our simulation setup.

For each simulated data set, we test the null hypothesis of no causal effect using the CRT and AMCE-based approaches. We then compute the proportion of times each test rejects the null hypothesis. For the CRT, we use the HierNet test statistic given in Equation (5) obtained by fitting the response \mathbf{Y} on our predictors (\mathbf{X}, \mathbf{Z}) and enforcing the constraints shown in Equation (6). For the AMCE, we use the t -test based on the estimated regression coefficient of \mathbf{X} obtained through fitting the response \mathbf{Y} on a single predictor \mathbf{X} .¹⁰

The left plot of Figure 3 shows how the statistical power of each test varies as the size of interaction effects between \mathbf{X} and \mathbf{Z} increases. The number of non-zero interaction effects between \mathbf{X} and \mathbf{Z} is fixed at six. For ease of interpretation, we plot the percentage of total outcome variance explained by the interaction effects between \mathbf{X} and \mathbf{Z} on the x -axis.¹¹ In our setup, the total variance represents the outcome variance explained by all main and interaction effects under the latent representation of the logistic regression model (see Appendix G for details). Consistent with our theoretical expectation, the CRT (blue triangles) becomes more powerful than the AMCE-based test (red circles) as the interaction size increases. For example, when the interaction size is strong enough to account for about 30% of the total variance, the CRT is approximately 20 percentage points more powerful than the AMCE-based test. When there is no interaction effect, the CRT is

¹⁰Although this is a valid procedure to compute the AMCE estimate for \mathbf{X} , practitioners typically compute the AMCE estimates of all factors (\mathbf{X}, \mathbf{Z}) simultaneously with a single linear regression of \mathbf{Y} on all (\mathbf{X}, \mathbf{Z}) . Figure 6 of Appendix G.3 shows that the power of the AMCE remains indistinguishable when using all factors (\mathbf{X}, \mathbf{Z}) in a single linear regression. Lastly, although our goal is to compare the CRT with the AMCE, we acknowledge that practitioners may also use the omnibus F -test for testing interactions by including all the two-way interactions. We show in Appendix J that such an approach leads to inflated p -values, thus we omit this as a baseline comparison here.

¹¹The x -axis ticks correspond to interaction sizes of 0, 0.025, 0.05, 0.075, 0.1, and 0.125, respectively. For example, the 20% point on the x -axis refers to an interaction size of 0.05.

only slightly less powerful (by about 3 percentage points) than the AMCE-based test.

The right plot of Figure 3 shows how the power of the tests change as one varies the number of non-zero interaction effects. The size of interaction effects is fixed to 0.06, around half the size of the main effect. We find that as expected, the CRT becomes more powerful than the AMCE-based test as the number of interaction increases. For example, when there are twelve interactions the CRT is approximately 10 percentage points more powerful than the AMCE-based test. Even when there is no interaction effect at all, the loss of statistical power is minimal. Appendix E presents additional simulation results, showing that the use of no profile order constraints given in Equation (6) increases the power of test.

5 Application to Conjoint Experiments

In this section, we apply the proposed CRT to the two conjoint studies introduced in Section 2.

5.1 Immigration Preferences and Ethnocentrism

We begin our analysis of the immigration conjoint experiment by testing whether respondents differentiate between immigrants from Mexico and those from European countries. We use the same data set as the one used in Hainmueller and Hopkins (2015). This gives us a total sample of 6,980 observations with $n = 1,396$ respondents each rating $J = 5$ tasks. Our main factor of interest \mathbf{X} is the “country of origin” variable. Since we are only interested in testing how respondents differentiate Mexican and European candidates, we use the generalized hypothesis H_0^{General} defined in Equation (7) and coarsen the three levels — *Germany*, *France*, and *Poland* — into one level called *Europe* (see Appendix C for a formal treatment). Furthermore, the h function in H_0^{General} takes the “country of origin” variable and maps the relevant levels of *Mexico* and *Europe* to one output and the remaining levels to other unique outputs. We include all the other randomized factors and respondent characteristics (see Section 2.1) as \mathbf{Z} and \mathbf{V} , respectively, except the ethnocentrism variable, which is only measured for a subset of respondents. We incorporate this variable at the end of this section.

We fit HierNet using \mathbf{Y} as the response and our main factor \mathbf{X} , other randomized factors \mathbf{Z} , and respondent characteristics \mathbf{V} as the predictors. We then compute the test statistic given in Equation (10) while imposing the implied no profile order effect constraints given in Equation (6) (with the constraints applied to the respondent characteristic interactions too). As mentioned briefly in Section 3.4, we slightly modify this test statistic by only using the estimated coefficients for *Mexico* and *Europe* while ignoring the other coefficients.

As shown in the left upper cell of Table 1, the CRT p -value of this test statistic is 0.042, providing evidence that respondents differentiate immigrants from *Mexico* and *Europe*. For comparison, we also compute the p -value based on the estimated AMCE of being from *Mexico* compared to being from *Europe*. We apply a commonly used linear regression approach described in Section 2.1 to compute this p -value. Specifically, we first fit a linear regression model using “country of origin”, “reason of immigration”, and their interaction as predictors to account for the restricted randomization (Hainmueller, Hopkins and Yamamoto, 2014). The standard errors are clustered by respondent. We then use the F -test of the linear equality constraint that implies the null hypothesis under the

	CRT	AMCE	Profile order effect	Carryover effect	Fatigue effect
Immigration	0.042	0.27	0.80	0.12	0.45
Gender	0.026	0.93, 0.40	0.15	0.97	0.66

Table 1: The p -values based on the Conditional Randomization Test (CRT) and the Average Marginal Component Effect (AMCE) Estimation. The two rows represent the different applications. The first two columns present the p -values from the HierNet-based CRT and AMCE-based test statistics. The first row presents the p -values for testing whether the immigrant’s “country of origin” (*Mexico* or *Europe*) matters for immigration preferences while the second row tests whether candidate’s “gender” matters for voters’ preferences of Congressional candidates, respectively. The second AMCE-based p -value for “gender” corresponds to a fair comparison with the CRT-based p -value by additionally testing the “gender” interaction with the candidate’s “party affiliation” (*Democratic* or *Republican*). The remaining columns report the p -values for testing no profile order effect, no carryover effect, and no fatigue effect for the respective application.

linear model. As shown in the upper cell of the second column of Table 1, the resulting p -value for the difference between *Mexico* and *Europe* is 0.27, which is statistically insignificant.

The above result suggests that the CRT may be able to capture complex interactions and yield greater statistical power than the AMCE-based test. The two largest interactions in the observed test statistic are within-profile interactions between “country of origin” and “education” and between “country of origin” and “prior trips to U.S.” factors, which included whether or not the immigrant entered the U.S. illegally. Thus, we next assess the degree to which the interaction effects account for this difference in statistical power. To do this, we use a Lasso logistic regression without interaction terms where we only include the main effects of $(\mathbf{X}, \mathbf{Z}, \mathbf{V})$. The CRT p -value of using only the relevant levels of the main effects of \mathbf{X} as the test statistic is 0.082, which is somewhat larger than the p -value based on HierNet test statistic. This suggests that interactions play some role in yielding a more powerful test than the AMCE-based approach.

Hainmueller and Hopkins suggest that respondents do not differentiate between immigrants from *Mexico* and those from *Europe* based on the results of their main analysis (see Figure 1, which replicates this analysis). However, they also conduct a subgroup analysis and find that “country of origin” has statistically significant interaction(s) with the respondent’s ethnocentrism through a subgroup analysis (see also Newman and Malhotra, 2019, for related findings). Thus, we now repeat the same analysis as above except that we include this ethnocentrism variable as an additional respondent characteristic in \mathbf{V} . Note that unlike the original analysis, we do not dichotomize this variable and use the original continuous scale. Since ethnocentrism is only measured for white and black respondents, the number of total respondents is reduced to $n = 1,135$. Despite this reduction in sample size, the inclusion of the ethnocentrism variable produces the p -value of 0.019, which is smaller than the p -value of the analysis without this variable. As expected, the largest interaction in the observed test statistic involves the ethnocentrism variable. All together, our analysis provides evidence that respondents differentiate immigrants from *Mexico* and *Europe*.

Lastly, we use the CRT to test the three commonly made regularity assumptions of conjoint analysis: no profile order effect, no carryover effect, and no fatigue effect. The last three columns in Table 1 present the p -values from the various tests described in Section 3.5. We find no evidence that these assumptions are violated in the immigration conjoint experiment (the first row). In particular, the fact that we do not detect profile order effects suggests that imposing the symmetry constraint as done in Equation (6) likely improves power.

5.2 Role of Gender in Candidate Evaluation

For the gender conjoint experiment, we test whether or not the gender of Congressional candidates matters in voter preferences. We use the same data as the one used in Ono and Burden (2018). The Congressional dataset consists of 7,915 observations with 5 tasks performed by each of $n = 1,583$ respondents. Our main factor of interest \mathbf{X} is a binary variable representing *male* or *female*. In addition, we use the remaining 12 randomized factors \mathbf{Z} and all the respondent characteristics \mathbf{V} (see Section 2.2). We test the main null hypothesis H_0 introduced in Section 3.2.

As mentioned in Section 3.3, the use of substantive knowledge can improve the power of the test. To demonstrate this, we leverage the Presidential candidate dataset from the same conjoint experiment to find the strongest interaction with the gender of candidates. We then include this interaction term as an additional main effect in HierNet when computing the test statistic given in Equation (10).¹² By including it as a main effect, HierNet applies less shrinkage on this interaction term. In addition, HierNet will consider potential three-way interactions involving this interaction term and other variables in \mathbf{Z} . The power will be greater if strong interactions in the Presidential candidate data are also present in the Congressional candidate data.

To find the strongest interaction in the Presidential candidate dataset, we obtain a CRT p -value for each variable in (\mathbf{Z}, \mathbf{V}) with a test statistic that focuses on the interaction strength for the corresponding variable under consideration. Specifically, the test statistic uses Lasso logistic regression with all main effects of $(\mathbf{X}, \mathbf{Z}, \mathbf{V})$ and an additional interaction between \mathbf{X} and one variable from (\mathbf{Z}, \mathbf{V}) (see Appendix H for more details). We choose the variable with the lowest p -value as the strongest interaction. The Presidential data shows that the candidate’s “party affiliation” (*Democratic* or *Republican*) had the most significant interaction with their “gender”. Appendix H contains further details and a robustness check by repeating the same analysis but choosing the variable with the second lowest p -value as the additional main effect.

As shown in the second row of Table 1, the CRT p -value using the HierNet test statistic is 0.026, showing that gender may matter even for Congressional candidates. We find the largest two interactions in the observed test statistic were two three-way interactions: one between “gender”, “party affiliation”, and “respondent’s political interest” and the other between “gender”, “party affiliation”, and the “respondent’s party affiliation”. This result is consistent with the findings in de la Cuesta, Egami and Imai (2022), which suggests the existence of higher order interactions involving party affiliation. We assess the role of interaction effects using the same procedure above in the immigration example. The resulting CRT p -value from a Lasso logistic regression with only the main effects of $(\mathbf{X}, \mathbf{Z}, \mathbf{V})$ and the additional interaction with “gender” and “party affiliation” is 0.15, suggesting that the other interactions were also helpful in detecting significance.

For comparison, we compute the p -value based on the estimated AMCE of “gender” for Congressional candidates as presented in Figure 2. Similar to the common strategy used for the immigration conjoint experiment, we fit a linear regression model using “gender” as the sole predictor while clustering standard errors by respondents. We find that the p -value is 0.89. However, since the CRT leveraged the Presidential candidate data to up-weight the interaction with “party affiliation”, we also create a fair comparison by obtaining analogous p -values for the AMCE. To do this,

¹²Our test statistic $T_{\text{HierNet}}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{V})$ is not only a function of the Congressional dataset $(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{V})$ but also the entire presidential dataset (\mathbf{P}) . Therefore, we must now hold all $(\mathbf{Z}, \mathbf{V}, \mathbf{P})$ fixed in the resampling procedure. However, this does not change Algorithm 1 and the resulting p -value remains valid because $\mathbf{X} \mid (\mathbf{Z}, \mathbf{V}, \mathbf{P})$ is still an independent fair coin flip between the levels of *male* and *female*.

we again fit a linear regression using “gender” but with an additional main effect of “party affiliation” and the interaction of “gender” and “party affiliation”. We then report the p -value from a F -test for both the main effect of “gender” and the interaction with “party affiliation”. The resulting p -value is 0.40, showing that the AMCE-based result remains statistically insignificant. Finally, the last three columns in the second row show no evidence that the regularity assumptions are violated for this conjoint experiment.¹³

6 Concluding Remarks

Conjoint analysis is a popular methodology for analyzing multi-dimensional preferences and decision making. In this paper, we propose an assumption-free approach for conjoint analysis based on the conditional randomization test (CRT). The proposed methodology allows researchers to test whether a set of factors of interest matter at all without assuming a statistical model. We also extend the proposed methodology to test for differential effects for any combination of factor levels and other regularity assumptions commonly invoked in conjoint analysis like the profile order effect. Unlike the standard AMCE analysis, the CRT can avoid masking important interactions due to averaging over other factors. When constructing CRT test statistics, researchers can use machine learning algorithms and/or domain knowledge to detect complex interactions among factors, without making modeling assumptions. The CRT is easy to implement and provides exact (i.e., non-asymptotic) p -values that are valid even in high dimensions. We believe that this flexibility combined with its assumption-free nature makes the CRT a powerful tool for conjoint analysis. The CRT can complement the existing methods like the AMCE analysis by providing a useful way to examine whether a factor of interest matters at all.

¹³We only test the no profile order effect assumption for Congressional candidates because this is the data relevant to the research question. However, for testing the carryover effect and fatigue effect, we use the full dataset including the Presidential candidates in order to increase power.

References

- Abramson, Scott, Korhan Kocak, Asya Magazinnik and Anton Strezhnev. 2020. Improving Preference Elicitation in Conjoint Designs using Machine Learning for Heterogeneous Effects. Technical Report. The Annual Summer Meeting of the Society for Political Methodology.
- Andrews, Rick L., Asim Ansari and Imran S. Currim. 2002. “Hierarchical Bayes versus Finite Mixture Conjoint Analysis Models: A Comparison of Fit, Prediction, and Partworth Recovery.” Journal of Marketing Research 39:87–98.
URL: <https://doi.org/10.1509/jmkr.39.1.87.18936>
- Aronow, Peter M. 2012. “A General Method for Detecting Interference Between Units in Randomized Experiments.” 41:3–16.
- Athey, Susan, Dean Eckles and Guido W. Imbens. 2018. “Exact P-values for Network Interference.” Journal of the American Statistical Association 113:230–240.
- Bansak, Kirk, Jens Hainmueller, Daniel Hopkins and Teppei Yamamoto. 2020. “Using Conjoint Experiments to Analyze Elections: The Essential Role of the Average Marginal Component Effect (AMCE).” SSRN Electronic Journal.
- Bansak, Kirk, Jens Hainmueller, Daniel J. Hopkins and Teppei Yamamoto. 2018. “The Number of Choice Tasks and Survey Satisficing in Conjoint Experiments.” Political Analysis 26:112–119.
- Bansak, Kirk, Jens Hainmueller, Daniel J. Hopkins and Teppei Yamamoto. 2019. “Beyond the breaking point? Survey satisficing in conjoint experiments.” Political Science Research and Methods.
- Barone, Stefano, Alberto Lombardo and Pietro Tarantino. 2007. “A weighted logistic regression for conjoint analysis and Kansei engineering.” Quality and Reliability Engineering International 23:689 – 706.
- Bates, Stephen, Matteo Sesia, Chiara Sabatti and Emmanuel Candès. 2020. “Causal inference in genetic trio studies.” Proceedings of the National Academy of Sciences 117:24117–24126.
URL: <https://www.pnas.org/content/117/39/24117>
- Bien, Jacob, Jonathan Taylor and Robert Tibshirani. 2013. “A lasso for hierarchical interactions.” Ann. Statist. 41:1111–1141.
- Bodog, Simona and G.L. Florian. 2012. “Conjoint Analysis in Marketing Research.” Journal of Electrical and Electronics Engineering 5:19–22.
- Campbell, Benjamin L, Saneliso Mhlanga and Isabelle Lesschaeve. 2013. “Consumer Preferences for Peach Attributes: Market Segmentation Analysis and Implications for New Marketing Strategies.” Agricultural and resource economics review 42:518–541.
- Candès, Emmanuel J. and Pragma Sur. 2018. “The phase transition for the existence of the maximum likelihood estimate in high-dimensional logistic regression.” arXiv: Methodology.

- Candès, Emmanuel, Yingying Fan, Lucas Janson and Jinchi Lv. 2018. “Panning for Gold: Model-X Knockoffs for High-dimensional Controlled Variable Selection.” Journal of the Royal Statistical Society: Series B 80:551–577.
- Chernozhukov, Victor, Mert Demirer, Esther Duflo and Iv’an Fernández-Val. 2017. Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments. Papers No. 1712.04802. arXiv.org.
URL: <https://ideas.repec.org/p/arx/papers/1712.04802.html>
- 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 30:19–45.
- Dezeure, Ruben, Peter Bühlmann, Lukas Meier and Nicolai Meinshausen. 2015. “High-Dimensional Inference: Confidence Intervals, p -Values and R-Software hdi.” Statist. Sci. 30:533–558.
URL: <https://doi.org/10.1214/15-STS527>
- Egami, Naoki and Kosuke Imai. 2019. “Causal Interaction in Factorial Experiments: Application to Conjoint Analysis.” Journal of the American Statistical Association 114:529–540.
- Fisher, Ronald A. 1935. The Design of Experiments. London: Oliver and Boyd Chapter II.
- Goplerud, Max, Kosuke Imai and Nicole E. Pashley. 2022. Estimating Heterogeneous Causal Effects of High-Dimensional Treatments: Application to Conjoint Analysis. Technical Report. arxiv:2201.01357.
- Green, Paul, Abba Krieger and Yoram Wind. 2001. “Thirty Years of Conjoint Analysis: Reflections and Prospects.” Interfaces 31:S56–S73.
- Green, Paul E. and V. Srinivasan. 1990. “Conjoint Analysis in Marketing: New Developments with Implications for Research and Practice.” Journal of Marketing 54:3–19.
- Hainmueller, Jens and Daniel J. Hopkins. 2014. “Public Attitudes Toward Immigration.” Annual Review of Political Science 17:225–249.
- Hainmueller, Jens and Daniel J. Hopkins. 2015. “The Hidden American Immigration Consensus: A Conjoint Analysis of Attitudes toward Immigrants.” American Journal of Political Science.
- Hainmueller, Jens, Daniel J. Hopkins and Teppei Yamamoto. 2014. “Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments.” Political Analysis 22:1–30.
- Harking, Andrade. 2021. “Cherry-Picking, P-Hacking, Fishing Expeditions, and Data Dredging and Mining as Questionable Research Practices.” Frontiers in Psychology 16.
- Hauber, A. Brett, Juan Marcos González, Catharina G.M. Groothuis-Oudshoorn, Thomas Prior, Deborah A. Marshall, Charles Cunningham, Maarten J. IJzerman and John F.P. Bridges. 2016. “Statistical Methods for the Analysis of Discrete Choice Experiments: A Report of the ISPOR Conjoint Analysis Good Research Practices Task Force.” Value in Health 19:300 – 315.

- Imai, Kosuke and Michael Lingzhi Li. 2021. “Experimental Evaluation of Individualized Treatment Rules.” Journal of the American Statistical Association.
- Imbens, Guido W. and Donald B. Rubin. 2015. Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction. Cambridge University Press.
- Liu, Molei, Eugene Katsevich, Lucas Janson and Aaditya Ramdas. 2020. “Fast and Powerful Conditional Randomization Testing via Distillation.” arXiv preprint arXiv:2006.03980.
- Luce, R.Duncan and John W. Tukey. 1964. “Simultaneous conjoint measurement: A new type of fundamental measurement.” Journal of Mathematical Psychology 1:1 – 27.
- McFadden, Daniel. 1973. Frontiers of Econometrics (eds. P. Zarembka). New York: Academic Press Chapter Conditional Logit Analysis of Qualitative Choice Behavior, pp. 105–142.
- Newman, Benjamin J. and Neil Malhotra. 2019. “Economic Reasoning with a Racial Hue: Is the Immigration Consensus Purely Race Neutral?” The Journal of Politics 81:153–166.
- Ono, Yoshikuni and Barry C. Burden. 2018. “The Contingent Effects of Candidate Sex on Voter Choice.” Political Behavior.
- Popovic, Milena, Marija Kuzmanovic and Milan Martic. 2012. “Using Conjoint Analysis To Elicit Employers’ Preferences Toward Key Competencies For A Business Manager Position.” Management - Journal for theory and practice of management 17:17–26.
- Raghavarao, D., J.B. Wiley and P. Chitturi. 2010. Choice-based conjoint analysis: Models and Designs. Chapman and Hall/CRC.
- Rubin, Donald B. 1990. “Comments on “On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9” by J. Splawa-Neyman translated from the Polish and edited by D. M. Dabrowska and T. P. Speed.” Statistical Science 5:472–480.
- Szucs, Denes. 2016. “A Tutorial on Hunting Statistical Significance by Chasing N.” Frontiers in Psychology 7:1444.
URL: <https://www.frontiersin.org/article/10.3389/fpsyg.2016.01444>
- Tansey, Wesley, Victor Veitch, Haoran Zhang, Raul Rabadan and David Blei. 2018. “The Holdout Randomization Test: Principled and Easy Black Box Feature Selection.”
- Teele, Dawn Langan, Joshua Kalla and Frances Rosenbluth. 2018. “The Ties That Double Bind: Social Roles and Women’s Underrepresentation in Politics.” American Political Science Review 112:525–541.

A Proof of Theorem 3.1

Proof. We first prove that if $\mathbf{Y} \perp\!\!\!\perp \mathbf{X} \mid \mathbf{Z}$, then $\mathbf{Y}(\mathbf{x}, \mathbf{z}) \stackrel{d}{=} \mathbf{Y}(\mathbf{x}', \mathbf{z})$ for any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ and $\mathbf{z} \in \mathcal{Z}$.

$$\begin{aligned} P(\mathbf{Y} \mid \mathbf{Z} = \mathbf{z}) &= P(\mathbf{Y} \mid \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}) \\ &= P(\mathbf{Y}(\mathbf{x}, \mathbf{z}) \mid \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}) \\ &= P(\mathbf{Y}(\mathbf{x}, \mathbf{z}) \mid \mathbf{Z} = \mathbf{z}) \\ &= P(\mathbf{Y}(\mathbf{x}, \mathbf{z})), \end{aligned}$$

where the third and fourth equalities follow from the randomization of \mathbf{X} and that of \mathbf{Z} , respectively. Similarly, we can show $P(\mathbf{Y} \mid \mathbf{Z} = \mathbf{z}) = P(\mathbf{Y}(\mathbf{x}', \mathbf{z}))$, thus we have shown $P(\mathbf{Y}(\mathbf{x}, \mathbf{z})) = P(\mathbf{Y}(\mathbf{x}', \mathbf{z}))$ for any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ and $\mathbf{z} \in \mathcal{Z}$.

To the prove the other direction, we want to show that $\mathbf{Y}(\mathbf{x}, \mathbf{z}) \stackrel{d}{=} \mathbf{Y}(\mathbf{x}', \mathbf{z})$ implies $\mathbf{Y} \perp\!\!\!\perp \mathbf{X} \mid \mathbf{Z}$, or equivalently $P(\mathbf{Y} \mid \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}) = P(\mathbf{Y} \mid \mathbf{X} = \mathbf{x}', \mathbf{Z} = \mathbf{z})$ for any value of $\mathbf{z} \in \mathcal{Z}$ and any value of $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$.

$$\begin{aligned} P(\mathbf{Y}(\mathbf{x}, \mathbf{z})) &= P(\mathbf{Y}(\mathbf{x}, \mathbf{z}) \mid \mathbf{Z} = \mathbf{z}) \\ &= P(\mathbf{Y}(\mathbf{x}, \mathbf{z}) \mid \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}) \\ &= P(\mathbf{Y} \mid \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}), \end{aligned}$$

where the first two equalities follow from the randomization of \mathbf{Z} and that of \mathbf{X} , respectively. The same argument shows $P(\mathbf{Y}(\mathbf{x}', \mathbf{z})) = P(\mathbf{Y} \mid \mathbf{X} = \mathbf{x}', \mathbf{Z} = \mathbf{z})$. Finally, because $\mathbf{Y}(\mathbf{x}, \mathbf{z}) \stackrel{d}{=} \mathbf{Y}(\mathbf{x}', \mathbf{z})$ we have that $P(\mathbf{Y}(\mathbf{x}, \mathbf{z})) = P(\mathbf{Y}(\mathbf{x}', \mathbf{z}))$, implying $P(\mathbf{Y} \mid \mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z}) = P(\mathbf{Y} \mid \mathbf{X} = \mathbf{x}', \mathbf{Z} = \mathbf{z})$ for any value of $\mathbf{z} \in \mathcal{Z}$ and any value of $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$. \square

B Relation to Finite-Population Inference

In Section 3, we introduced H_0 under the super-population framework. Here, we consider the relationship between H_0 and Fisher's sharp null of no treatment effect (Fisher, 1935). Under the finite-population framework where the potential outcomes are fixed and the randomness comes only from the randomization of treatment assignment, if we assume no interference between units, H_0 reduces to testing $Y_{ij}(\mathbf{x}_{ij}, \mathbf{z}_{ij}) = Y_{ij}(\mathbf{x}'_{ij}, \mathbf{z}_{ij})$ for all i, j and all possible values of $\mathbf{x}_{ij}, \mathbf{x}'_{ij} \in \mathcal{X}$ and $\mathbf{z}_{ij} \in \mathcal{Z}$. Under the super-population framework, if we make the same no-interference assumption, H_0 reduces to testing $Y_{ij}(\mathbf{x}_{ij}, \mathbf{z}_{ij}) \stackrel{d}{=} Y_{ij}(\mathbf{x}'_{ij}, \mathbf{z}_{ij})$ for all i, j and all possible values of $\mathbf{x}_{ij}, \mathbf{x}'_{ij}, \mathbf{z}_{ij}$, where the potential outcomes are assumed to be drawn from a population. These two null hypotheses are not equivalent even though the CRT can test H_0 under both the finite-population and super-population frameworks (as proven in Theorem 3.1). This is because the distribution of $\mathbf{Y}(\mathbf{x}, \mathbf{z})$ can still be equal to $\mathbf{Y}(\mathbf{x}', \mathbf{z})$ even if $Y_{ij}(\mathbf{x}, \mathbf{z})$ is different than $Y_{ij}(\mathbf{x}', \mathbf{z})$ for some i, j .

C Grouping Factor Levels

In this appendix, we further detail how to test H_0^{General} when the analyst is interested in grouping multiple factor levels. For example, in the immigration conjoint application (Section 5.1), we

wish to test whether respondents differentiate immigrants from Mexico and those from Europe where there are three distinct levels for countries from Europe: *Germany*, *France*, and *Poland*. We now formally show how to group these levels up into one category *Europe* and test the hypothesis H_0^{General} .

We introduce a coarsening function c that takes q factors of interest \mathbf{X} as input and transforms them to a new set of grouped factors $\bar{\mathbf{X}}$. Formally, this function is defined as $c : \mathcal{X} \mapsto \bar{\mathcal{X}}$ where $\bar{\mathcal{X}}$ represents the support of the grouped factors $\bar{\mathbf{X}}$ with $|\bar{\mathcal{X}}| \geq |\mathcal{X}|$. For the above example, we define the function c on the “country of origin” factors such that the *Mexico* level takes one value whereas the *France*, *Germany*, and *Poland* levels all take another value: $Europe \in \bar{\mathbf{X}}$. All other factor levels are mapped to different values in $\bar{\mathbf{X}}$.

Next, we define the outcome for each level of newly transformed factor levels. Given the coarsening function c defined above, we introduce the marginalized potential outcome variable $\bar{\mathbf{Y}}(\bar{\mathbf{x}}, \mathbf{z})$, which averages over the distribution of original factor levels that are grouped. Formally, this new outcome variable has the following mixture structure,

$$\bar{\mathbf{Y}}(\bar{\mathbf{x}}, \mathbf{z}) = \frac{\sum_{\mathbf{x}' \in \mathcal{X}} \mathbf{1}\{c(\mathbf{x}') = \bar{\mathbf{x}}\} \mathbf{Y}(\mathbf{x}', \mathbf{z}) P(\mathbf{X} = \mathbf{x}' | \mathbf{Z} = \mathbf{z})}{\sum_{\mathbf{x}' \in \mathcal{X}} \mathbf{1}\{c(\mathbf{x}') = \bar{\mathbf{x}}\} P(\mathbf{X} = \mathbf{x}' | \mathbf{Z} = \mathbf{z})}, \quad (11)$$

where $\mathbf{z} \in \mathcal{Z}$, $\bar{\mathbf{x}} \in \bar{\mathcal{X}}$, and $P(\mathbf{X} = \mathbf{x} | \mathbf{Z} = \mathbf{z})$ represents the conditional distribution of \mathbf{X} given \mathbf{Z} used in the experiment. For example, if we group three European countries—*France*, *Germany*, and *Poland*—and create one new factor level *Europe*, then its marginalized potential outcome will be a mixture distribution of the original potential outcomes for the three countries weighted by their known randomization probabilities conditional on the other factors.

Furthermore, the previously introduced coarsening function h now takes the newly grouped up factor $\bar{\mathbf{X}}$ and maps it to the new coarsened factor, i.e., $h : \bar{\mathcal{X}} \mapsto \tilde{\mathcal{X}}$. Consequently, our updated generalized null hypothesis is,

$$\bar{H}_0^{\text{General}} : \bar{\mathbf{Y}}(\bar{\mathbf{x}}, \mathbf{z}) \stackrel{d}{=} \bar{\mathbf{Y}}(\bar{\mathbf{x}}', \mathbf{z}) \text{ for all } \bar{\mathbf{x}}, \bar{\mathbf{x}}' \in \bar{\mathcal{X}}, \text{ such that } h(\bar{\mathbf{x}}) = h(\bar{\mathbf{x}}') \text{ and } \mathbf{z} \in \mathcal{Z}. \quad (12)$$

Finally, it can also be shown by applying the same argument as the one used to prove Theorem 3.1 that $\bar{H}_0^{\text{General}}$ is equivalent to the following conditional independence relation,

$$\bar{\mathbf{Y}} \perp\!\!\!\perp \bar{\mathbf{X}} \mid h(\bar{\mathbf{X}}), \mathbf{Z}. \quad (13)$$

To test this null hypothesis, we keep the original test statistic T_{HierNet} shown in Equation (5) under the same symmetry constraints given in Equation (6) except that we use $\bar{\mathbf{X}}$ in place of \mathbf{X} to account for coarsening based on the function c .

D Data Description

Tables 2 and 3 present all the factors and their respective levels used in the immigration conjoint experiment (Section 2.1) and the gender candidate conjoint experiment (Section 2.2), respectively.

Factor for Immigration Conjoint Experiment	Factor Levels
Education level	No Formal Education, Fourth Grade, Eight Grade, High School, Two Years College, College Degree, Graduate Degree
Gender	Female, Male
Country of origin	Germany, France, Mexico, Philippines, Poland, India, China, Sudan, Somalia, Iraq
Language	Fluent English, Broken English, Tried to speak English but unable to, Spoke through an interpreter
Reason for application	Reunite with family members, Seek better job, Escape political/religious persecution
Profession	Gardener, Waiter, Nurse, Teacher, Child Care Provider, Janitor, Construction Worker, Financial Analyst, Research Scientist, Doctor, Computer Programmer
Job experience	No job training or prior experience, One to two years, Three to five years, More than five years
Employment plan	Has a contract with a U.S. employer, Does not have a contract with a U.S. employer, but has done job interviews, Will look for work after arriving in the U.S., Has no plans to look for work at this time
Prior trips to the U.S.	Never been to the U.S., Entered the U.S. once before on a tourist visa, Entered the U.S. once before without legal authorization, Has visited the U.S. many times before on tourist visas, Spent six months with family members in the U.S.

Table 2: All nine randomized factors and their respective levels in the conjoint experiment used in Hainmueller and Hopkins (2015).

Factor For Gender Conjoint Experiment	Factor Levels
Gender	Male, Female
Age	36, 44, 52, 60, 68, 76
Race/ethnicity	White, Black, Hispanic, Asian American
Family	Single (never married), Single (divorced), Married (no child), Married (two children)
Experience in public office	12 years, 8 years, 4 years, No experience
Salient personal characteristics	Strong leadership, Really cares about people like you, Honest, Knowledgeable, Compassionate, Intelligent
Party affiliation	Republican, Democrat
Policy area of expertise	Foreign policy, Public safety (crime), Economic policy, Health care, Education, Environmental issues
Position on national security	Cut military budget and keep U.S. out of war, Maintain strong defense and increase U.S. influence
Position on immigrants	Favors guest worker program, Opposes guest worker program
Position on abortion	Pro-choice, Pro-life, Neutral
Position on government deficit	Reduce through tax increase, Reduce through spending cuts, Does not want to reduce
Favorability rating among public	34%, 43%, 52%, 61%, 70%

Table 3: All thirteen randomized factors and their respective levels in the conjoint experiment used in Ono and Burden (2018).

E Enforcing The No Profile Order Effect

We detail here a way to enforce the no profile order effect constraints in Equation (6) for general test statistics. We show through simulations that these constraints, when they hold, can substantially improve statistical power.

Under these constraints, switching the "left" and "right" profile order does not change the value of the test statistic. We now formalize this intuition. First denote $D \in \mathbb{R}^{nJ \times (2p+r+1)}$ as the regression data matrix composed of $(\mathbf{X}, \mathbf{Z}, \mathbf{V}, \mathbf{Y})$. Let $s_E : \mathbb{R}^{nJ \times (2p+r+1)} \rightarrow \mathbb{R}^{nJ \times (2p+r+1)}$ denote a function that takes a data matrix as an input and swaps the "left" and "right" profile order for rows E of the data matrix, where $E \subset \{1, 2, \dots, nJ\}$. For example, suppose we swap just the first row and we denote $(\tilde{\mathbf{X}}, \tilde{\mathbf{Z}}, \tilde{\mathbf{V}}, \tilde{\mathbf{Y}})$ as the columns for the output $s_{\{1\}}(D)$. Then $\tilde{\mathbf{X}}_{11}^L = \mathbf{X}_{11}^R$, $\tilde{\mathbf{X}}_{11}^R = \mathbf{X}_{11}^L$, $\tilde{\mathbf{Z}}_{11}^L = \mathbf{Z}_{11}^R$, $\tilde{\mathbf{Z}}_{11}^R = \mathbf{Z}_{11}^L$, and $\tilde{Y}_{11} = 1 - Y_{11}$ and all remaining rows remain identical as the original data D . The function applies no swap to the respondent characteristic \mathbf{V} ($\tilde{\mathbf{V}} = \mathbf{V}$).

We introduce a new data matrix D^c that appends the original data matrix with a data matrix that swaps the profile order for all rows. Formally, $D^c = [D; s_{\{1,2,\dots,nJ\}}(D)]$. Table 4 shows an example of D^c with \mathbf{X} as gender, \mathbf{Z} as party affiliation, and \mathbf{V} as the respondent's age. Conceptually, D^c aims to destroy all information about the profile order since the "left" and "right" profile are now indistinguishable, thus ensuring any test statistic that uses D^c will respect the no profile order effect constraints. The following lemma formally states this result.

Row Number	Gender ^L	Gender ^R	Party ^L	Party ^R	Respondent Age	Y
1	Male	Female	Democrat	Republican	27	1
⋮	⋮	⋮	⋮	⋮	⋮	⋮
nJ + 1	Female	Male	Republican	Democrat	27	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮

Table 4: Visual example of D^c data matrix with \mathbf{X} as gender, \mathbf{Z} as party affiliation, and \mathbf{V} as the respondent’s age. With a slight abuse of notation Gender^L denotes the left profile’s gender values. Gender^R, Party^L, and Party^R are defined similarly.

Lemma E.1. Let $T(\cdot)$ be a row invariant test statistic,¹⁴ then for any $E \subset \{1, 2, \dots, nJ\}$ we have that $T(D^c) = T(s_E(D)^c)$, where $s_E(D)^c = [s_E(D); s_{1,2,\dots,nJ}(s_E(D))]$.

Proof. The lemma holds because $D^c = s_E(D)^c$ up to row permutations. Using the assumption that $T(\cdot)$ is a row invariant test statistic we obtain the equality. Many algorithms like Random Forest have random initializations in the process so $T(D^c) = T(s_E(D)^c)$ may only hold up to distributional equality. \square

Lemma E.1 states that the test statistic will remain invariant to any relabelling of “left” and “right” profiles as long the test statistic is a function of the D^c data matrix. All regression based algorithms will respect exact row invariance. Lemma E.1 allows practitioners to build any test statistic, even ones that do not have natural “left” and “right” coefficients, while still enforcing the no profile order effect. In particular, we enforce the constraints in Equation (6) by using D^c as the input in HierNet.

Simulations. We now show that imposing the no profile order effect constraints in Equation (6) can substantially increase statistical power under the same simulation setup in Section 4 when the no profile order effect assumption is satisfied. To evaluate the power gain, we also fit HierNet without imposing the constraints in Equation (6) by using the original data D . We use $T_{\text{HierNet}}(\mathbf{X}, \mathbf{Y}, \mathbf{Z})^L + T_{\text{HierNet}}(\mathbf{X}, \mathbf{Y}, \mathbf{Z})^R$ as the test statistic, where $T_{\text{HierNet}}(\mathbf{X}, \mathbf{Y}, \mathbf{Z})^L$ is the same as $T_{\text{HierNet}}(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ in Equation (5) but all coefficients correspond to their respective estimates for the left profile. We similarly define $T_{\text{HierNet}}(\mathbf{X}, \mathbf{Y}, \mathbf{Z})^R$.

Figure 4 shows that imposing the no profile order effect constraints can significantly increase power. For example, the power of HierNet without the constraints (purple squares) is roughly equal or smaller than that of even the AMCE (red circles) when the interaction effect accounts for 20% of the variance or when there are as many interactions as twelve. Furthermore, we see the power of using HierNet that imposes the constraints (blue triangles) is consistently higher than that of the HierNet without the constraints.

Enforcing No Profile Order Effect in the Carryover Effect Test Statistic. As mentioned in Section 3.5, we enforce the no profile order effect constraints when conducting the test of no carryover effect. We detail here how to implement this.

¹⁴More formally we say $T(D)$ is row invariant if $T(D) = T(\pi(D))$ for any possible π , where π denotes a possible permutation of the rows of D .

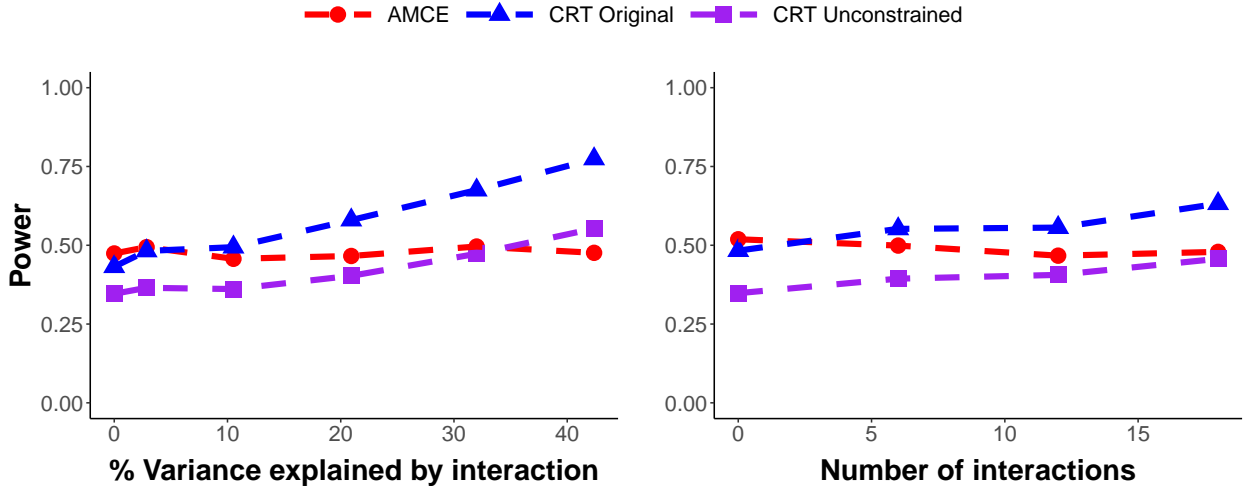


Figure 4: This figure represents the power gain from imposing the no “Profile Order Effect” constraints in Equation (6) under the same simulation setup in Figure 3. We keep the AMCE and the original HierNet statistical power curves (red circles and blue triangles, respectively) and add the new unconstrained HierNet test statistic (purple squares).

Since \mathbf{X}^* represents the lag-1 profile values, we do not expect \mathbf{X}^* alone (without interactions with \mathbf{Z}^*) to influence respondents’ choice of the left versus right profile. Therefore, we force all main effects and interactions among \mathbf{X}^* to be zero in the following way. Let $\mathbf{X}_i^{*L} \in \mathbb{R}^{\frac{J}{2} \times p}$ denote the columns of \mathbf{X}_i^* that correspond to the left profile in \mathbf{X}_i^* . More formally, $\mathbf{X}_i^{*L} = [(\mathbf{X}_{i1}^L); (\mathbf{Z}_{i1}^L)]^\top; [(\mathbf{X}_{i3}^L); (\mathbf{Z}_{i3}^L)]^\top; \dots; [(\mathbf{X}_{i,J-1}^L); (\mathbf{Z}_{i,J-1}^L)]^\top$. If J is not even, consider up to $J - 1$ instead. We define $\mathbf{X}_i^{*R}, \mathbf{Z}_i^{*L}, \mathbf{Z}_i^{*R}$ similarly. We also define $\mathbf{X}^{*L} = [\mathbf{X}_1^{*L}; \dots; \mathbf{X}_n^{*L}] \in \mathbb{R}^{\frac{nJ}{2} \times p}$ with $\mathbf{X}^{*R}, \mathbf{Z}^{*L}, \mathbf{Z}^{*R}$ defined similarly. Then, we append copies of the following: $[(\mathbf{X}^{*R})^\top; (\mathbf{X}^{*L})^\top; (\mathbf{Z}^{*R})^\top; (\mathbf{Z}^{*L})^\top]^\top$; $[(\mathbf{X}^{*L})^\top; (\mathbf{X}^{*R})^\top; (\mathbf{Z}^{*R})^\top; (\mathbf{Z}^{*L})^\top]^\top$; $[(\mathbf{X}^{*R})^\top; (\mathbf{X}^{*L})^\top; (\mathbf{Z}^{*L})^\top; (\mathbf{Z}^{*R})^\top]^\top$ to the original data matrix $[(\mathbf{X}^{*L})^\top; (\mathbf{X}^{*R})^\top; (\mathbf{Z}^{*L})^\top; (\mathbf{Z}^{*R})^\top]^\top$, resulting in a total of $2nJ$ rows. Lastly, we also append copies of $[\mathbf{1} - \mathbf{Y}_i^*; \mathbf{1} - \mathbf{Y}_i^*; \mathbf{Y}_i^*]$ to the original response \mathbf{Y}_i^* . Appending the first copy of $[(\mathbf{X}^{*R})^\top; (\mathbf{X}^{*L})^\top; (\mathbf{Z}^{*R})^\top; (\mathbf{Z}^{*L})^\top]^\top$ to the original data matrix enforces the familiar constraint in Equation (6) using Lemma E.1. The remaining copies force all the main effects and interactions among \mathbf{X}^* to be zero by appealing to the same reasoning in Lemma E.1.

F CRT Procedure for Testing Extensions

In this appendix, we describe in further detail how to carry out all the resampling procedures for the tests introduced in Sections 3.4 and 3.5.

F.1 Testing H_0^{General}

When testing H_0^{General} , the resampling procedure is different than Algorithm 1 because Equation (8) forces us to hold $(h(\mathbf{X}), \mathbf{Z})$ constant rather than holding only \mathbf{Z} constant. The conditional distribution of $\mathbf{X} \mid (h(\mathbf{X}), \mathbf{Z})$ constrains \mathbf{X} to be randomized only within the factor levels of interest. For

Algorithm 2: Generalized H_0 : Testing H_0^{General}

Input: Data $(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{V})$, test statistic $T(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{v})$, $h(\mathbf{x})$, total number of re-samples B ;
for $b = 1, 2, \dots, B$ **do**
 Sample $\mathbf{X}^{(b)}$ from the distribution of $\mathbf{X} \mid (h(\mathbf{X}), \mathbf{Z}, \mathbf{V})$ conditionally independently of \mathbf{X}
 and \mathbf{Y} ;
Output: $p\text{-value} := \frac{1}{B+1} \left[1 + \sum_{b=1}^B \mathbb{1}\{T(\mathbf{X}^{(b)}, \mathbf{Y}, \mathbf{Z}, \mathbf{V}) \geq T(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{V})\} \right]$

concreteness, consider the immigration example in Section 5.1. In this case, $h(\mathbf{X})$ groups levels *Mexico* and *Europe* to one output while keeping all the other countries of origin the same. Therefore, when obtaining the resamples for the “country of origin” factor, we keep all countries except *Mexico* and *Europe* constant, i.e., countries such as *China* do not get re-randomized. For the entries corresponding to *Mexico* and *Europe*, we resample values *Mexico* and *Europe* with probabilities $(0.25, 0.75)$, respectively (since 3 countries make up Europe). We present this resampling procedure in Algorithm 2. Lastly, we remark that Algorithm 2 remains the same when testing $\overline{H_0}^{\text{General}}$ except we replace \mathbf{X} with $\overline{\mathbf{X}}$.

F.2 Testing the Regularity Assumptions

Profile Order Effect. We first describe testing the assumption of no profile order effect without imposing SUTVA. For any $E \subset \{1, \dots, nJ\}$, let $\mathbf{x}^{(E)}$ swap the left and right profile values in rows E of \mathbf{x} while leaving the remaining rows unchanged. We similarly define $\mathbf{z}^{(E)}$. Then, let $\mathbf{Y}^{(E)}(\mathbf{x}, \mathbf{z})$ flip the bits of (replace 1’s with 0’s and vice versa) the entries of $\mathbf{Y}(\mathbf{x}, \mathbf{z})$ corresponding to indices in E while leaving the remaining entries unchanged. For example, for simplicity, assume no \mathbf{Z} while $nJ = 3$, $E = \{1\}$, $\mathbf{x} = [(G, F); (P, G); (M, F)]$ (the left profile values come first), and $\mathbf{Y}(\mathbf{x}) = (1, 0, 1)$. Then $\mathbf{x}^{(E)} = [(F, G); (P, G); (M, F)]$ and $\mathbf{Y}^{(E)}(\mathbf{x}) = (0, 0, 1)$. The observed $\mathbf{Y}^{(E)}$ is defined similarly. We can formally state the null hypothesis of no profile order effect as follows:

$$H_0^{\text{Order}} : \mathbf{Y}(\mathbf{x}, \mathbf{z}) \stackrel{d}{=} \mathbf{Y}^{(E)}(\mathbf{x}^{(E)}, \mathbf{z}^{(E)}) \text{ for all } E \subset \{1, \dots, nJ\}, \mathbf{x} \in \mathcal{X}, \text{ and } \mathbf{z} \in \mathcal{Z}.$$

This null hypothesis states that for all possible reorderings of the left and right profiles there is no causal impact on which profile is chosen. For the resampling procedure, we hold the realized values of all (\mathbf{X}, \mathbf{Z}) constant while only resampling the subset E , i.e., drawing nJ independent Bernoulli coin flips to determine which of the nJ rows to include as part of E as described above. Algorithm 3 details the procedure to calculate the CRT p -value for testing H_0^{Order} . Lastly, to not enforce Equation (6) in $T_{\text{HierNet}}^{\text{Order}}$, we fit HierNet on the original data matrix $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ rather than D^c .

Carryover Effect. When testing the carryover effect, we need to hold the even numbered tasks \mathbf{Z}^* fixed while resampling all odd numbered tasks \mathbf{X}^* . Therefore, we resample all factors $(\mathbf{X}_{ij}, \mathbf{Z}_{ij})$ for $j = 1, 3, \dots, J - 1$ from the experimental distribution for all the factors while holding $(\mathbf{X}_{ij}, \mathbf{Z}_{ij})$ for $j = 2, 4, \dots, J$ constant. Algorithm 4 details the procedure to calculate the CRT p -value for testing $H_0^{\text{Carryover}}$.

Algorithm 3: CRT: Profile Order Effect

Input: Data $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$, test statistic $T(\mathbf{x}, \mathbf{y}, \mathbf{z})$, total number of re-samples B ;
for $b = 1, 2, \dots, B$ **do**
 Sample nJ independent Bernoulli(0.5) independently of $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$. E^b is the index set corresponding to values of 1 in the nJ Bernoulli's;
Output: p -value $:= \frac{1}{B+1} \left[1 + \sum_{b=1}^B \mathbb{1}\{T(\mathbf{X}^{(E^b)}, \mathbf{Y}^{(E^b)}, \mathbf{Z}^{(E^b)}) \geq T(\mathbf{X}, \mathbf{Y}, \mathbf{Z})\} \right]$

Algorithm 4: CRT: Carryover Effects

Input: Data $(\mathbf{X}^*, \mathbf{Y}^*, \mathbf{Z}^*)$, test statistic $T(\mathbf{x}^*, \mathbf{y}^*, \mathbf{z}^*)$, total number of re-samples B ;
for $b = 1, 2, \dots, B$ **do**
 Sample $(\mathbf{X}_{ij}^b, \mathbf{Z}_{ij}^b)$ from the distribution of $(\mathbf{X}_{ij}, \mathbf{Z}_{ij})$ independently of \mathbf{Y} for $i = 1, 2, \dots, n$ and $j = 1, 3, \dots, J-1$. Let $(\mathbf{X}_i^b)^* = [(\mathbf{X}_{i1}^b; \mathbf{Z}_{i1}^b)^\top; (\mathbf{X}_{i3}^b; \mathbf{Z}_{i3}^b)^\top; \dots; (\mathbf{X}_{i,J-1}^b; \mathbf{Z}_{i,J-1}^b)^\top]^\top$ and $(\mathbf{X}^b)^* = [(\mathbf{X}_1^b)^*; (\mathbf{X}_2^b)^*; \dots; (\mathbf{X}_n^b)^*]$;
Output: p -value $:= \frac{1}{B+1} \left[1 + \sum_{b=1}^B \mathbb{1}\{T((\mathbf{X}^b)^*, \mathbf{Y}^*, \mathbf{Z}^*) \geq T(\mathbf{X}^*, \mathbf{Y}^*, \mathbf{Z}^*)\} \right]$

Fatigue Effect. To carry out the CRT to test the fatigue effect, we re-sample only the task evaluation order index \mathbf{F} for each respondent uniformly from the set of all permutations on $\{1, \dots, J\}$, denoted as Π_J , while holding all the experimental factor values fixed. Algorithm 5 details the procedure to calculate the CRT p -value for testing H_0^{Fatigue} .

G Details of the Simulation Setup

We provide here a detailed description of our simulation setup. For simplicity, we focus on the setting in which each respondent only has one evaluation, i.e., $J = 1$. Appendix G.4 (see Figure 7) presents additional simulations that have multiple evaluations per respondent based on respondent random effects. In the setting where each respondent evaluates only one task, we treat each response as an independent observation and drop subscript j . We also assume that every factor (\mathbf{X}, \mathbf{Z}) is uniformly and independently randomized, implying that all treatment combinations are equally likely. As before, \mathbf{X} represents the main factors of interest while \mathbf{Z} denotes the other factors. For simplicity, we assume that all factors (\mathbf{X}, \mathbf{Z}) are binary with their success probabilities equal to 0.5, and we have one factor of interest ($q = 1$).

G.1 The Basic Setup

To clearly separate main and interaction effects, we use the sum-to-zero constraint by coding each binary factor as $(-0.5, 0.5)$. Our data generating process uses the following logistic regression model under the forced-choice design,

$$\Pr(Y_i = 1 \mid X_i, \mathbf{Z}_i) = \text{logit}^{-1} \left[\beta_X(X_i^L - X_i^R) + \beta_Z^\top(\mathbf{Z}_i^L - \mathbf{Z}_i^R) + 2\gamma^\top \{(X_i^L \mathbf{Z}_i^L) - (X_i^R \mathbf{Z}_i^R)\} + 2\delta^\top \{(X_i^L \mathbf{Z}_i^R) - (X_i^R \mathbf{Z}_i^L)\} + 2\tilde{\gamma}^\top \{(\mathbf{Z}_i^L \times \mathbf{Z}_i^L) - (\mathbf{Z}_i^R \times \mathbf{Z}_i^R)\} \right],$$

Algorithm 5: CRT: Fatigue Effect

Input: Data $(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{F})$, test statistic $T(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{f})$, total number of re-samples B ;

for $b = 1, 2, \dots, B$ **do**

 └ Sample \mathbf{F}^b uniformly from Π_J the set of all permutations on $\{1, \dots, J\}$;

Output: p -value $:= \frac{1}{B+1} \left[1 + \sum_{b=1}^B \mathbb{1}\{T(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{F}^b) \geq T(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{F})\} \right]$

where β_X and β_Z represent the coefficient vectors for the main effects of \mathbf{X} and \mathbf{Z} , respectively, and γ and δ denote the coefficient column vectors for the within-profile and between-profile interactions between \mathbf{X} and \mathbf{Z} , respectively. For simplicity, we omit between-profile interactions among \mathbf{Z} and consider only within-profile interactions among \mathbf{Z} with effect sizes $\tilde{\gamma}$. To facilitate interpretation, each interaction coefficient is multiplied by 2 because our encoding of interaction effects, e.g., $\{(X_i^L \mathbf{Z}_i^L) - (X_i^R \mathbf{Z}_i^R)\}$, results in three possible values $(-0.5, 0, 0.5)$.

This data generating process also implies the absence of profile order effects. Lastly, we note that the logistic regression has a latent variable, Y'_i , representation such that $Y_i = 1$ if $Y'_i > 0$ and 0 otherwise, where we let ϵ_i follow a standard logistic distribution and

$$Y'_i = \beta_X(X_i^L - X_i^R) + \beta_Z^\top(\mathbf{Z}_i^L - \mathbf{Z}_i^R) + 2\gamma^\top \{(X_i^L \mathbf{Z}_i^L) - (X_i^R \mathbf{Z}_i^R)\} + 2\delta^\top \{(X_i^L \mathbf{Z}_i^R) - (X_i^R \mathbf{Z}_i^L)\} + 2\tilde{\gamma}^\top \{(\mathbf{Z}_i^L \times \mathbf{Z}_i^L) - (\mathbf{Z}_i^R \times \mathbf{Z}_i^R)\} + \epsilon_i.$$

We consider settings in which \mathbf{X} has one main effect of $\beta_X = 0.1$, which is fixed for all simulations. In addition, \mathbf{Z} consists of ten factors with four non-zero main effects with a magnitude of 0.1 with alternating signs, which is fixed for all simulations, i.e., $\beta_Z = (0.1, -0.1, 0.1, -0.1, 0, \dots, 0)$. Lastly, we fix $\tilde{\gamma}^\top$ to have fifteen non-zero entries of 0.05, where the non-zero interactions are randomly chosen from all possible within-profile interactions among \mathbf{Z} . The remaining entries are all zero.

The sample size is fixed to $n = 3,000$ throughout the simulations. For each simulation, we generate within-profile and between-profile interactions between \mathbf{X} and \mathbf{Z} by randomly selecting the specified number of interactions from all possible interactions. The total number of non-zero interaction effects between \mathbf{X} and \mathbf{Z} varies from 0 to 18. The number of non-zero within-profile interactions between \mathbf{X} and \mathbf{Z} is kept identical to that of non-zero between-profile interactions between \mathbf{X} and \mathbf{Z} . We make all non-zero within-profile interactions positive and all non-zero between-profile interactions negative while fixing all non-zero interaction effects between \mathbf{X} and \mathbf{Z} to be equal in magnitude. We explore additional simulations in Appendix G.2 where there are heterogeneous interaction effects. We set $B = 200$ for all simulations presented in this paper.

We measure the total variance with respect to the latent representation of the logistic regression. Consequently, we define the variance explained by the interaction effects between \mathbf{X} and \mathbf{Z} and all other “remaining” effects (main effects of \mathbf{X} , \mathbf{Z} , and interactions among \mathbf{Z}) as,

$$\begin{aligned} \sigma_{\text{Interaction}}^2 &:= \mathbb{V}(2\gamma^\top \{(X_i^L \mathbf{Z}_i^L) - (X_i^R \mathbf{Z}_i^R)\} + 2\delta^\top \{(X_i^L \mathbf{Z}_i^R) - (X_i^R \mathbf{Z}_i^L)\}) \\ \sigma_{\text{Remaining}}^2 &:= \mathbb{V}(\beta_X(X_i^L - X_i^R) + \beta_Z^\top(\mathbf{Z}_i^L - \mathbf{Z}_i^R) + 2\tilde{\gamma}^\top \{(\mathbf{Z}_i^L \times \mathbf{Z}_i^L) - (\mathbf{Z}_i^R \times \mathbf{Z}_i^R)\}), \end{aligned}$$

where $\mathbb{V}(\cdot)$ denotes the variance of the respective random variable. The total variance is defined as $\sigma_{\text{Interaction}}^2 + \sigma_{\text{Remaining}}^2$. Since $\beta_X, \beta_Z, \tilde{\gamma}$ are fixed for all simulations, $\sigma_{\text{Remaining}}^2$ is fixed with a value of:

$$\sigma_{\text{Remaining}}^2 = 9 \times 0.1^2 \times 2 \mathbb{V}(X_i^L) + 15 \times 4 \times 0.05^2 \times 2 \mathbb{V}(X_i^L X_i^R) = 9 \times 0.1^2 \times \frac{1}{2} + 15 \times 4 \times 0.05^2 \times \frac{1}{8} = 0.06375,$$

where the first equality holds because all random variables ($X_i^L, X_i^R, \mathbf{Z}_i^L, \mathbf{Z}_i^R$) are independent, centered at zero, and identically distributed (element-wise identically distributed for the multivariate \mathbf{Z}_i^L and \mathbf{Z}_i^R).

Furthermore, given a signal size I and the number of interactions n_I (n_I denotes the total number of within-profile and between-profile interactions), we have that:

$$\sigma_{\text{Interaction}}^2 = 8I^2 \mathbb{V}(X_i^L X_i^R) n_I = \frac{I^2 n_I}{2}.$$

Finally, since the left plot of Figure 3 contains six interactions ($n_I = 6$), the x -axis is computed by $\frac{3I^2}{3I^2 + 0.06375}$, where I takes the following values (0, 0.025, 0.05, 0.075, 0.1, 0.125).

To calculate the statistical power of each test, we compute a p -value for each of 1,000 Monte Carlo data sets. For the CRT, we use the HierNet test statistic given in Equation (5) and impose the constraints in Equation (6), whereas we use the t -test based on the estimated regression coefficient of \mathbf{X} for the AMCE-based test. The AMCE-based analysis assumes no profile-order effect. Therefore, as suggested by Hainmueller, Hopkins and Yamamoto (2014), we run the linear regression by stacking all the left and right profiles, resulting in $2n$ rows. More formally, we have the new response $\mathbf{Y}^{\text{AMCE}} = [\mathbf{Y}; \mathbf{1} - \mathbf{Y}]$ regressed on $\mathbf{X}^{\text{AMCE}} = [\mathbf{X}^L; \mathbf{X}^R]$, where $\mathbf{X}^L = [X_1^L; X_2^L; \dots; X_n^L]$ and \mathbf{X}^R is defined similarly (note we have dropped the j subscript since $J = 1$). Finally, the standard errors are clustered by respondents as suggested by Hainmueller, Hopkins and Yamamoto. Since $J = 1$ for our simulation, we cluster the standard errors on each evaluation task, i.e., each unique cluster consists of the left and right profile for each task. We then compute the power as the proportion of p -values that are less than $\alpha = 0.05$.

G.2 Heterogeneous Interaction Size

As stated in Section 4, we fix all interaction sizes to be equal for every simulation. Here, we examine if the simulation results presented in Figure 3 change if there are heterogeneous interaction effects.

We compare the statistical power under two different data generating processes using the CRT with test statistic in Equation (5). We denote the original data generating process as the ‘‘homogeneous’’ scenario since all interaction sizes are equal under this scenario. We create an additional ‘‘heterogeneous’’ scenario that contains two unique varying interaction effects - one that is strong, I_s , and one that is weak I_w . To facilitate a fair comparison between the ‘‘heterogeneous’’ and ‘‘homogeneous’’ scenario, we force the total variance explained by the interactions to be equal under both scenarios. We assign all strong interaction effects to the within-profile interaction and all weak interaction effects to the between-profile interaction. Suppose there are only two non-zero interactions between \mathbf{X} and \mathbf{Z} . Since we impose $\sigma_{\text{Interaction}}^2$ to be equal under both the ‘‘homogeneous’’ and ‘‘heterogeneous’’ scenario, we have the following equation:

$$I_w^2 + I_s^2 = 2I^2,$$

where I is the interaction size for the ‘‘homogeneous’’ scenario.

The possible values of strong and weak effects lie on the circle of radius $\sqrt{2}I$ centered at the origin. We pick $I_s(I) = \sqrt{\frac{4}{3}}I$ and $I_w(I) = \sqrt{\frac{2}{3}}I$, i.e., the point corresponding to the thirty degree angle of the circle. The variance explained by the interaction is equal for both the ‘‘heterogeneous’’




Figure 5: The figure shows the power of the “heterogeneous” (light green circles) and “homogeneous” (blue triangles) scenario in the same simulation setting as in Figure 3.

and “homogeneous” scenario when there are two non-zero interactions. To create a power curve for the “heterogeneous” scenario analogous to the left plot in Figure 3, we create three interactions of size $I_s(I)$ for the within-profile interaction and three interactions of size $I_w(I)$ for the between-profile interaction. The two scenarios still maintain equal variance explained by the interaction effects because all random variables are independent and centered at zero. For the analogue of the right plot of Figure 3, we similarly keep the relative proportion of $I_w(I)$ and $I_s(I)$ fixed and increase the number of non-zero interactions to match the “homogeneous” scenario. For example, if there are twelve non-zero interaction effects, then there are six within-profile interactions of size $I_s(I)$ and six between-profile interactions of size $I_w(I)$ for the “heterogeneous” scenario.

Figure 5 shows that the power of the “heterogeneous” and “homogeneous” scenario is indistinguishable under the same simulation setting in Figure 3. This shows that we lose no generality by considering only the simple “homogeneous” scenario in the main simulations in Figure 3.

G.3 Additional AMCE Simulations

The AMCE computed in Figure 3 was based on a linear regression of \mathbf{Y} on \mathbf{X} , without \mathbf{Z} included among the predictors. Although this is valid and sufficient to compute the AMCE of \mathbf{X} , practitioners often compute the AMCE of all factors (\mathbf{X}, \mathbf{Z}) simultaneously with a single linear regression of \mathbf{Y} on all (\mathbf{X}, \mathbf{Z}) . We compute the power of the “long AMCE” that is based on the t -test for the estimated regression coefficient of \mathbf{X} obtained by regressing the response \mathbf{Y} on all (\mathbf{X}, \mathbf{Z}) . Figure 6 shows that the power of the “long AMCE” (orange squares) is indistinguishable from that of the original AMCE presented in Figure 3 (red circles).

G.4 Simulations with Multiple Tasks per Respondent

Appendix G.1 presents the simulation results where each respondent evaluates only one task ($J = 1$). Here, we consider a simulation setup where each respondent evaluates $J = 5$ tasks while still




Figure 6: The figure shows the power of the “long AMCE” that uses all (\mathbf{X}, \mathbf{Z}) in the linear regression fit (orange squares) and the original AMCE that uses only \mathbf{X} in the linear regression fit (red circles) in the same simulation setting as in Figure 3.

fixing the total sample size $nJ = 3,000$. We keep the same simulation setup as the one described in Appendix G.1 except that we allow each respondent to have a random effect $U_j \sim N(0, \sigma_{RE}^2)$. More formally, our new data generating process is,

$$\Pr(Y_{ij} = 1 \mid X_{ij}, \mathbf{Z}_{ij}) = \text{logit}^{-1} \left[\beta_X(X_{ij}^L - X_{ij}^R) + \beta_Z^\top(\mathbf{Z}_{ij}^L - \mathbf{Z}_{ij}^R) + 2\gamma^\top \{(X_{ij}^L \mathbf{Z}_{ij}^L) - (X_{ij}^R \mathbf{Z}_{ij}^R)\} + 2\delta^\top \{(X_{ij}^L \mathbf{Z}_{ij}^R) - (X_{ij}^R \mathbf{Z}_{ij}^L)\} + 2\tilde{\gamma}^\top \{(\mathbf{Z}_{ij}^L \times \mathbf{Z}_{ij}^L) - (\mathbf{Z}_{ij}^R \times \mathbf{Z}_{ij}^R)\} + U_j \right],$$

where U_j is the random effect for each respondent j . We keep all simulation parameters the same as that in Figure 3 except we use the above data generating process with $J = 5$ evaluation tasks and random effects with $\sigma_{RE}^2 = 0.1$ to produce Figure 7.

Although the CRT HierNet test statistic does not change with the addition of multiple respondent evaluations, the AMCE estimate must properly account for the respondent effect. As suggested by Hainmueller, Hopkins and Yamamoto (2014), we use the robust clustered standard errors clustered on respondents for the linear regression of \mathbf{Y} on \mathbf{X} and use the t -test based on the estimated regression coefficient of \mathbf{X} to produce the power curve (red). Figure 7 shows the results are similar to those shown in Figure 3, suggesting that our results are not sensitive to the number of evaluations per respondent.

H Leveraging Presidential Data in the Gender Application

We detail here how we leverage the Presidential candidate data to find the strongest interaction, as mentioned in Section 5.2. Because we are only interested in whether the additional interactions are significant, we run the Lasso logistic regression with main effects of $(\mathbf{X}, \mathbf{Z}, \mathbf{V})$ and one interaction between \mathbf{X} (“gender”) and one variable in (\mathbf{Z}, \mathbf{V}) , producing one CRT p -value for each variable in (\mathbf{Z}, \mathbf{V}) . For example, when including an interaction between “gender” and “profession”, we include




Figure 7: The figure shows the power of the AMCE (red circles) and the CRT (blue triangles) using the HierNet test statistic in Equation (5). We modify the simulation setting in Figure 3 by having each respondent evaluate $J = 5$ tasks with a total of $nJ = 3,000$ responses. Otherwise, the simulation setup remains identical to that in Figure 3. Each respondent has a random effect of $\sigma_{RE}^2 = 0.1$.

both within-profile and between-profile interactions between all levels of “gender” and “profession”. Consequently, the test statistic for testing the interaction between \mathbf{X} (“gender”) and factor ℓ of \mathbf{Z} is:

$$T_{\text{Presidential}}^{\text{Candidate Factor, } \ell} = \sum_{k=1}^K \sum_{k'=1}^{K_\ell} (\hat{\gamma}_{1\ell kk'} - \bar{\gamma}_{1\ell k'})^2 + \sum_{k=1}^K \sum_{k'=1}^{K_\ell} (\hat{\delta}_{1\ell kk'} - \bar{\delta}_{1\ell k'})^2. \quad (14)$$

The test statistic for testing the interaction between \mathbf{X} and factor m of the respondent characteristic \mathbf{V} is:

$$T_{\text{Presidential}}^{\text{Respondent Characteristic, } m} = \sum_{k=1}^K \sum_{w=1}^{L_m} (\hat{\xi}_{1mkw} - \bar{\xi}_{1mw})^2, \quad (15)$$

where all coefficient estimates refer to the same corresponding estimates in Equation (10) and $K = 2$ refers to levels *male* and *female*. Lastly, when running the CRT we similarly enforce the constraints in Equation (6) (along with the constraints on the respondent characteristics) by fitting the Lasso logistic regression on the appended D^c data matrix.

Table 5 shows the resulting p -value for each variable in (\mathbf{Z}, \mathbf{V}) . Many of the variables have p -values lower than 0.1. The variables such as “position on immigrants”, “position on abortion”, “position on government deficit”, and “position on national security”, are all related to the disparate views Democratic and Republican candidates may have, in line with “party affiliation” being the most significant. Even among the respondent characteristics, the respondent’s political ideology, a measure of how conservative or liberal the respondent is, is the most significant variable. We conduct a robustness analysis by repeating the analysis of Section 5.2 but using the second most significant variable, the “position on abortion” factor, to interact and include as the additional main effect in HierNet, though we note that “position on abortion” is quite a bit less significant than “party affiliation” with more than four times as large a p -value. The resulting p -value is 0.078. Although

Variable	p -value
Age	0.060
Race	0.31
Family	0.22
Experience in public office	0.30
Salient personal characteristic	0.45
Party affiliation	0.0049
Policy area of expertise	0.25
Position on national security	0.067
Position on immigrants	0.067
Position on abortion	0.022
Position on government deficit	0.032
Favorability rating among public	0.63
Respondent gender	1.00
Respondent education	1.00
Respondent age	1.00
Respondent class	0.41
Respondent region	1.00
Respondent race	1.00
Respondent partisanship	0.11
Respondent thought on Hillary Clinton	1.00
Respondent interest in politics	0.43
Respondent political ideology	0.042

Table 5: Resulting p -values using the CRT Lasso logistic regression with main effects of $(\mathbf{X}, \mathbf{Z}, \mathbf{V})$ and an additional interaction between \mathbf{X} and one variable in (\mathbf{Z}, \mathbf{V}) from the Presidential candidate data. The test statistic captures the interaction terms with each variable in (\mathbf{Z}, \mathbf{V}) as shown in Equation (14) and Equation (15), respectively. All respondent characteristic variables are labelled with “respondent”.

this is not as significant as the main analysis, it still provides suggestive evidence that gender plays a role in voting for Congressional candidates.

I Computational Details

The HierNet test statistic introduced in Equation (5) is powerful but can be computationally expensive because the CRT requires a total of $B + 1$ cross-validated HierNet fits. To address this problem, we speed up HierNet in three ways. First, we reduce the default convergence tolerance for the optimization algorithm in the HierNet package from 10^{-6} to 10^{-3} . Second, following the “distillation” idea introduced in (Liu et al., 2020), we cross-validate the sparsity parameter lambda only through a HierNet fit of \mathbf{Y} on \mathbf{Z} without involving any \mathbf{X} . Because (\mathbf{Y}, \mathbf{Z}) remains constant for all $B + 1$ fits, we only need one cross-validation fit. Lastly, we initialize the starting parameters in the optimization algorithm with one HierNet fit that is uniformly and randomly chosen from the $B + 1$ HierNet fits. Since we uniformly choose one out of $B + 1$ HierNet fits as the initialization, this procedure still satisfies the exchangeability needed for the CRT’s validity. Because many of the parameters estimated from the $B + 1$ different HierNet fits will likely be similar to each other,




Figure 8: This figure represents the power of the original faster HierNet test statistic (blue triangles) with the three computational speedups and the slower HierNet test statistic (black squares) without the computational speedups in the same simulation setting as in Figure 3.

the initialization likely saves computation time.

Although the above procedure significantly reduces computational complexity, practitioners may worry if there is a significant loss of power from this simplification. Consequently, we plot in Figure 8 the original HierNet power curve shown in Figure 3 that leverages the aforementioned three speed-ups (in blue) and the computationally slower HierNet power curve without the three speed-ups (in black). Figure 8 shows that the computational modifications have no significant impact on power.

J Inflated p -values for Logistic Regression

Although the AMCE is popular in conjoint analysis, especially among political scientists, there also exist model-based approaches. Logistic regression remains a popular model-based approach to conjoint analysis (McFadden, 1973; Green and Srinivasan, 1990; Campbell, Mhlanga and Lesschaeve, 2013). We explore in this section how this modeling approach can lead to invalid inference in conjoint analysis. When testing H_0 , researchers may want to account for not only the main effects of (\mathbf{X}, \mathbf{Z}) but also all two-way interactions, as done similarly in HierNet, to reduce model misspecification. Under this scenario, we show through simulations in Figure 9 that even reasonable sample sizes and dimensions of (\mathbf{X}, \mathbf{Z}) can lead to invalid p -values, i.e., the type 1 error is greater than the desired $\alpha \in [0, 1]$.

We use a similar simulation setting as the one in Appendix G.1 but simplify it further. Since we are interested in showing how the p -values obtained from a logistic regression may be invalid in general, we do not have “left” or “right” profiles but only one profile, leading to the following data generating process,

$$\Pr(Y_i = 1 \mid X_i, \mathbf{Z}_i) = \text{logit}^{-1} [\beta_X X_i + \beta_Z^\top \mathbf{Z}_i + \gamma^\top (X_i \mathbf{Z}_i) + \tilde{\gamma}^\top (\mathbf{Z}_i \times \mathbf{Z}_i)].$$




Figure 9: Inflated p -values from logistic regression. The left figure shows the proportion of p -values, obtained through a F -test from a logistic regression, less than $\alpha = 0.05$ when the number of other factors in \mathbf{Z} is (3, 5, 10, 11, 12, 13) and H_0 is true. The red dotted line at $\alpha = 0.05$ represents the expected proportion of p -values less than 0.05 if the p -values are valid. The right figure shows the histogram of 1,000 Monte-Carlo p -values when the number of other factors of \mathbf{Z} is 12. The sample size is $n = 5,000$ and each factor has four factor levels. All Monte Carlo standard errors are below 0.016.

All factors have four levels, and we similarly assume one factor of interest $q = 1$ while varying the number of other factors. Since we are interested in the behavior of the p -values under the null H_0 , we force all effects of X on the response to be zero, i.e., $\beta_X = \gamma = 0$. For simplicity we also make all effects of \mathbf{Z} zero, i.e., $\beta_Z = \tilde{\gamma} = 0$ and fix the sample size to $n = 5,000$. To reflect the researcher's desire to reduce model misspecification by accounting for all two-way interactions as HierNet does, we fit a logistic regression of \mathbf{Y} on all main effects and two-way interactions of (\mathbf{X}, \mathbf{Z}) . We then obtain a p -value for testing H_0 by an F -test that tests $\beta_X = \gamma = 0$. We obtain 1,000 Monte-Carlo p -values and plot the proportion of p -values less than $\alpha = 0.05$ in the left plot of Figure 9. We also vary the number of factors of \mathbf{Z} , which is shown in the x -axis of the left plot. On the right plot of Figure 9, we plot the histogram of the 1,000 p -values obtained when the number of factors of \mathbf{Z} is 12.

Under the null hypothesis, we expect any valid p -value to have type 1 error control, i.e., $P(p\text{-value} \leq \alpha) \leq \alpha$ for all $\alpha \in [0, 1]$. The left plot of Figure 9 shows that only five other factors of \mathbf{Z} is enough to cause the proportion of p -values less than $\alpha = 0.05$ to be noticeably inflated at 7%. The inflation becomes particularly apparent when there are twelve other factors of \mathbf{Z} , which causes the proportion of p -values less than 0.05 to be as high as 34%. The histogram on the right plot of Figure 9 visually shows how the p -values are clearly far from the expected uniform distribution and have an undesirable peak at zero, resulting in poor type 1 error control. This phenomenon is studied in (Candès and Sur, 2018) and arises because the p -values' validity in a logistic regression depends on a low-dimensional asymptotic result. We note that a conjoint analysis has typically more than ten factors, where each factor usually has more than three levels. Therefore, Figure 9 shows the potential dangers of using a model-based approach like the logistic regression to flexibly capture all interactions.