

Improving the External Validity of Conjoint Analysis: The Essential Role of Profile Distribution

Brandon de la Cuesta¹, Naoki Egami² and Kosuke Imai³

¹ Postdoctoral Research Fellow, King Center on Global Development, Stanford University, Palo Alto, CA 94305, USA. Email: brandon.delacuesta@stanford.edu, URL: <https://brandonelacuesta.com>

² Assistant Professor, Department of Political Science, Columbia University, New York, NY 10027, USA. Email: naoki.egami@columbia.edu, URL: <https://naokiegami.com>

³ Professor, Department of Government and Department of Statistics, Harvard University, 1737 Cambridge Street, Institute for Quantitative Social Science, Cambridge, MA 02138, USA. Email: imai@harvard.edu, URL: <https://imai.fas.harvard.edu>

Abstract

Conjoint analysis has become popular among social scientists for measuring multidimensional preferences. When analyzing such experiments, researchers often focus on the average marginal component effect (AMCE), which represents the causal effect of a single profile attribute while averaging over the remaining attributes. What has been overlooked, however, is the fact that the AMCE critically relies upon the distribution of the other attributes used for the averaging. Although most experiments employ the uniform distribution, which equally weights each profile, both the actual distribution of profiles in the real world and the distribution of theoretical interest are often far from uniform. This mismatch can severely compromise the external validity of conjoint analysis. We empirically demonstrate that estimates of the AMCE can be substantially different when averaging over the target profile distribution instead of uniform. We propose new experimental designs and estimation methods that incorporate substantive knowledge about the profile distribution. We illustrate our methodology through two empirical applications, one using a real-world distribution and the other based on a counterfactual distribution motivated by a theoretical consideration. The proposed methodology is implemented through an open-source software package.

Keywords: causal inference, conjoint analysis, factorial experiments, external validity

1 Introduction

Conjoint analysis is a factorial survey experiment that is designed to measure multidimensional preferences. In a typical application, respondents are presented with a pair of hypothetical profiles whose attributes are randomly selected, and are then asked to choose their preferred profile. Examples of such profiles include political candidates (e.g., Teele, Kalla, and Rosenbluth, 2018), immigrants (e.g., Hainmueller and Hopkins, 2015), and public policies (e.g., Ballard-Rosa, Martin, and Scheve, 2017). Although it has been extensively used in marketing research (e.g., Green, Krieger, and Wind, 2001; Marshall and Bradlow, 2002), conjoint analysis has quickly gained popularity in political science due to its wide applicability and relative simplicity (Hainmueller, Hopkins, and Yamamoto, 2014). Indeed, as shown in Figure 1, the number of major political science journal articles that utilize conjoint analysis has increased dramatically over the last 5 years.

The most commonly used quantity of interest in conjoint analysis is the average marginal component effect (AMCE), which represents the causal effect of changing one attribute of a profile while averaging over the distribution of the remaining profile attributes (Hainmueller *et al.*, 2014). Because conjoint analysis often involves many attributes, averaging over their distribution makes the interpretation of causal effects simpler and more practical than conditioning on their specific

Political Analysis (2022)
vol. 30: 19–45
DOI: 10.1017/pan.2020.40

Published
14 January 2021

Corresponding author
Naoki Egami

Edited by
Jeff Gill

© The Author(s) 2021. Published
by Cambridge University Press
on behalf of the Society for
Political Methodology

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

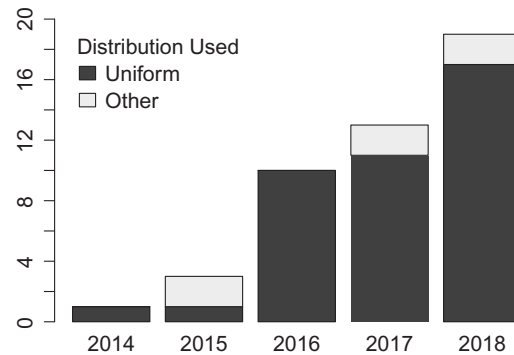


Figure 1. Recent growth of conjoint analysis and use of the uniform distribution for randomization in Political Science journal articles. Darker (lighter) fill represents the proportion of articles in which all the factors are randomized with the uniform (other) distribution. 88% of all reviewed articles use the uniform distribution. The plot is based on a review of articles published in political science journals from 2014 to 2018. See Supplemental Appendix A for the information about how the review was conducted.

values. For example, a researcher may be interested in the AMCE of candidate's gender that averages over the distribution of other candidate characteristics such as age, education, race, and policy positions. Thus, the definition of the AMCE critically depends on the distribution used to average over profile attributes.

Unfortunately, while this point is theoretically understood, in practice little attention has been paid to the choice of this distribution. As Figure 1 demonstrates, nearly 90% of the existing conjoint analyses use the uniform distribution. The problem is that the resulting estimate of the AMCE, which we call the *uniform* AMCE (uAMCE), gives equal weights to all conjoint profiles even when some of them are unrealistic from a substantive point of view. Ignoring the distribution of profiles fundamentally contradicts the key promise of conjoint analysis that the provision of information about several profile attributes makes the choice task realistic for respondents (Hainmueller, Hangartner, and Yamamoto, 2015). In fact, if other attributes do not systematically affect respondents' evaluation of the main attribute of interest, then one could simply elicit preferences over each attribute separately, making a conjoint experiment unnecessary. Therefore, conjoint analysis is beneficial precisely when we expect multiple attributes to jointly affect human decision making, and this is also the exact setting where the choice of profile distribution affects estimates of the AMCE the most.

In this paper, we study how the choice of profile distribution affects the conclusions of conjoint analysis. We define the *population* AMCE (pAMCE), which averages over the distribution of profile attributes in a target population of interest. Unlike the uAMCE, which is based on the uniform distribution, the pAMCE accounts for the relative frequency with which each profile occurs in the target population. This target profile distribution should be chosen according to the substantive interests of each study, similar to the choice of a target population of respondents in traditional survey sampling. The choice of distribution may be based on (1) real-world data, such as the characteristics and policy positions of actual politicians, or (2) a counterfactual distribution of theoretical interest. For each of the two scenarios, we provide empirical applications. We show that the difference between the uAMCE and pAMCE is large when the target profile distribution differs from uniform and when there exists interaction between the main attribute of interest and other attributes.

We propose two new strategies to estimate the pAMCE. The first approach, which we call *design-based confirmatory analysis*, incorporates the target profile distribution in the design stage (Section 4.1). We introduce three experimental designs that differ in terms of data requirements and necessary assumptions. In the most natural design, which we term joint population randomization, we propose randomizing conjoint profiles according to their target profile distribution

rather than the uniform. We then use a nonparametric estimator of the pAMCE, which can be computed using a weighted linear regression. This is a straightforward generalization of a widely used regression estimator (Hainmueller *et al.*, 2014).

Our second approach, *model-based exploratory analysis*, takes into account the target profile distribution at the analysis stage, after randomizing profiles and collecting data (Section 4.2). This approach is useful in estimating the pAMCE when researchers have to randomize profiles based on distributions different from the target profile distribution, such as the uniform. We propose fitting a flexible two-way interaction model and estimating the pAMCE as a weighted average of coefficients. Although this approach yields less precise estimates than the design-based confirmatory analysis, we discuss how to use regularization methods to partially recoup the loss of statistical efficiency (Egami and Imai, 2019).

One potential challenge of incorporating the target profile distribution is that the joint distribution of all attributes is difficult to obtain in some applications. For example, in a conjoint experiment of immigrant profiles, it may not be feasible to obtain the joint distribution of the (potentially many) attributes of immigrants that researchers wish to study. Recognizing this practical data constraint, we propose the marginal population randomization design, which only requires the knowledge of each factor's marginal distribution. Here, researchers randomize each factor independently with its marginal distribution. While this design requires a stronger assumption of no three-way or higher-order interactions, we provide a method to test its validity empirically. We also discuss how researchers can combine marginal distributions and partial joint distributions among several factors to relax this assumption.

The concern for unrealistic profiles is not new. In fact, researchers often remove a set of unusual profile combinations (e.g., doctors without college degree). Unfortunately, avoiding extreme cases is not sufficient for estimating the pAMCE. While some have begun to use unequal probabilities when randomizing profiles to partially address this concern (e.g., Hainmueller *et al.*, 2015; Huff and Kertzer, 2018; Leeper and Robison, 2018),¹ an overwhelming majority of researchers still use the uniform distribution without theoretically motivating it.² The substantive implication of this choice is that the resulting estimates of the AMCE are externally valid only when there is no interaction between attributes or when the uniform is the theoretically relevant profile distribution. Even though scholars have clearly discussed the importance of distributions used to randomize profiles (Hainmueller *et al.*, 2014),³ there currently exists no systematic way to incorporate the target profile distribution into the estimation of the AMCE. The proposed methodology directly addresses this problem by developing new experimental designs and estimation strategies. We note that our focus is on the external validity of conjoint *profiles*, and is distinct from another important issue of representativeness of *respondents* in survey experiments (see, e.g., Mutz, 2011; Mullinix *et al.*, 2015; Coppock, Leeper, and Mullinix, 2018; Miratrix *et al.*, 2018).

We illustrate the proposed methodology using two empirical applications. First, we reanalyze a conjoint experiment of political candidates by Ono and Burden (2019). The primary goal of the study is to estimate the effect of a candidate's gender on voter choice. The original study estimates the uAMCE of being female and finds that women candidates face discrimination in presidential but not in congressional elections. Specifying the target profile distribution to be the 115th U.S. Congress, we estimate the pAMCEs separately for Republican and Democratic legislators. We show

- 1 See also Barnes, Blumenau, and Lauderdale (2019) who point out that traditional conjoint experiments fail to generate realistic budget tradeoffs when studying public attitudes towards government spending.
- 2 Less than 4% of existing conjoint studies theoretically motivate distributions used for randomization in the article's main text. See Supplemental Appendix A, for additional information and a description of how these values were calculated.
- 3 Hainmueller *et al.* (2014) write "the choice of [population distribution] is important. It should always be made clear what weighting distribution of the treatment components was used in calculating the AMCE, and the choice should be convincingly justified. In practice, we suggest that the uniform distribution over all possible attribute combinations be used as a default, unless there is a strong substantive reason to prefer other distributions." (p. 12)

that the null effect of gender found in the original analysis for Congressional candidates is due to the large number of unrealistic profiles produced by the uniform distribution. Once we average profiles according to their real-world distributions, we recover a different result: women face a disadvantage when they run for Congress as Republicans but have an advantage when they run as Democrats. We also demonstrate that the uAMCE and pAMCE are similar for Presidential candidates because there exists little interaction between the main attribute of interest and other attributes within this subgroup.

As is the case for our first application, in many conjoint analyses, there exist natural target profile distributions, for which we can collect relevant data. In some cases, however, it might be impractical to gather corresponding real-world distributions (e.g., conjoint analysis of refugee profiles in Bansak, Hainmueller, and Hangartner, 2016). Alternatively, researchers may be interested in counterfactual profiles of theoretical interest, which may be rare or even absent in the real world. For example, Ballard-Rosa *et al.* (2017) examines a variety of hypothetical tax policy proposals that are infeasible in the real world politics, but are nonetheless essential in testing the authors' theoretical argument. Importantly, even in these scenarios, the AMCE estimates do depend on the choice of profile distribution. Thus, it is essential to use the proposed pAMCE framework to systematically investigate the sensitivity of the AMCE estimates to alternative theoretically relevant profile distributions.

Our second application, which is based on Peterson (2017), considers precisely those research settings where no natural target population exists or counterfactual profiles are of theoretical interest. Peterson (2017) examines how the amount of information about candidates alters the importance of copartisanship. By randomizing how much information voters receive, the author finds that the copartisan effect is weaker when they are shown additional information on policy positions and candidate attributes. We revisit this finding by applying the proposed methodology. We build three theoretically relevant counterfactual distributions that simulate high, medium, and low-information environments. We then show that the reduction in the effect of copartisanship is driven by the outsized influence of candidates' positions on abortion and deficit spending. While the original findings are based on a specific information environment, the proposed pAMCE framework enables the systematic investigation of their robustness.

2 Motivating Empirical Applications

Before presenting the proposed methodology, we describe a conjoint analysis that will motivate and illustrate the methodology proposed in this paper. We provide two empirical applications. The first application (Ono and Burden, 2019) is a common type of conjoint analysis based on profiles of politicians, which we use to demonstrate how to incorporate a real-world distribution of politicians' characteristics. In the second application (Peterson, 2017), we illustrate the importance of considering alternative profile distributions even in settings where no natural real-world distribution exists. We show how to systematically examine counterfactual profile distributions motivated by theoretical considerations.

2.1 The Effect of Candidate's Gender on Voter Choice

Scholars have long been interested in the conditions under which female candidates face obstacles to being elected (McDermott, 1997). A primary focus of the literature has been on whether a bias against female candidate is the result of taste-based or statistical discrimination (see, e.g., Arrow, 1998). While the taste-based discrimination argument implies that voters dislike the idea of having female candidates in office *per se*, the statistical discrimination hypothesis contends that voters, rightly or wrongly, associate female politicians with certain political backgrounds and policy preferences, and this association in turn shapes their vote choice. Under the statistical discrimination hypothesis, the provision of sufficient information about politicians beyond their

Table 1. Factors and levels used in Ono and Burden (2019). All factors are independently and uniformly randomized with levels in each factor shown with equal probability.

Factors	Levels
Gender	Male, Female
Race	Asian, Black, Hispanic, White
Age	36, 44, 52, 60, 68, 76
Family	Divorced, Never married, Married (no children), Married (2 children)
Experience	None, 4 years, 8 years, 12 years
Expertise	Economic policy, Education, Environmental issues, Foreign policy, Health care, Public safety
Character Trait	Compassionate, Honest, Intelligent, Knowledgeable, Leadership, Empathetic
Party	Republican, Democrat
Immigration Policy	Favors guest worker program, Opposes guest worker program
Security Policy	Strong military, Cut defense spending
Abortion Policy	Pro-choice, Neutral, Pro-life
Deficit Policy	Increase taxes, Take no action, Reduce spending
Favorability Rating	34%, 43%, 52%, 61%, 70%

gender should eliminate the bias against female politicians. If, on the other hand, voters are engaging in taste-based discrimination, they will disfavor female candidates even when other attributes are known.

In a recent study, Ono and Burden (2019) use a conjoint analysis to study the effects of candidate’s gender on vote choice. The authors test the aforementioned hypotheses by varying the gender of candidates and other factors such as partisanship. As in a typical conjoint analysis, respondents were asked to choose one of the two hypothetical political candidates, each of whom has the following factors: three demographic characteristics (age, race, gender), six political background (family life, years in office, area of expertise, partisanship, favorability rating, character trait), and four policy preferences (positions on abortion, immigration, national security and deficit reduction). In addition to attributes of the candidates, the original authors also randomly vary the office being sought at the candidate pair level; whether they run for President or Congress.

In Table 1, we summarize the levels of each factor used in this study. Each of 1,583 respondents evaluates 10 pairs of candidate profiles, indicating which one of the two profiles they prefer. Following the conventional conjoint analysis, all factors are independently randomized according to the uniform distribution so that each profile is equally likely. Under this uniform randomization design, the authors estimate the AMCE of candidate being female relative to male, marginalizing all other attributes, to be -1.25 percentage points (95% CI = $[-2.36, -0.19]$). This result implies that that female candidates suffer from a small disadvantage. The authors suggest that, because the conjoint analysis also presents other relevant information about politicians, this negative estimate represents evidence of taste-based rather than statistical discrimination. Importantly, Ono and Burden (2019) finds that the overall effect is driven by presidential candidates and there is little gender effect on congressional candidates. In particular, the estimated AMCE of being female is only -0.09 percentage points ($[-1.71, 1.48]$) for congressional candidates. On the other hand, the authors find a large negative effect of -2.42 percentage points ($[-3.96, -0.88]$) for presidential candidates. These findings led to the conclusion that discrimination against female candidates exists mostly in presidential elections rather than congressional elections.

Table 2. Factors used in Peterson (2017). Each respondent completed three choice tasks with each task containing two profiles. The full sample includes 1,059 respondents and 6,354 profiles. The design randomizes the number of factors shown to the respondent, which factors are shown, and the levels of each selected factor. The candidate's partisanship is always shown.

Factor	Levels
Age	28, 34, 40, 40, 46, 52, 58, 62, 68, 74
Gender	Male, Female
Race	Asian American, Black, Hispanic, White
Education	No college degree, AA (community college), BA from state university/small college/Ivy league
Profession	Business owner, Car dealer, Doctor, Farmer, High school teacher, Lawyer
Family	Never married, Married with 0/1/2 children, Divorced with 0/1/2 children
Military service	Served in U.S. military, No military service
Party	Democrat, Republican
Abortion stance	Never permissible, Permissible only when mother in danger, Always permissible
Spending stance	Large decrease, Small decrease, No change, Small increase, Large increase

2.2 The Effect of Information Environment on Partisan Voting

The study of copartisanship in the United States has long shown that voters demonstrate a strong preference for candidates of their own party (Campbell *et al.*, 1960). Although the importance of copartisanship is widely accepted, researchers disagree about its underlying mechanisms. Some argue that voters' support for parties is deeply rooted (Bartels, 2000). As a result, voters may use motivated reasoning when making decisions about which candidates to support, assessing information as favorable as possible given their partisan attachments (Bolsen, Druckman, and Cook, 2014; Druckman, 2014). Others argue that partisan cues mainly serve as substitutes for relevant information such as political background and policy preferences (Lau and Redlawsk, 2001; Bullock, 2011).

To adjudicate between these two theories, Peterson (2017) uses a conjoint analysis to estimate the extent to which the amount of information presented to voters conditions the importance of partisan cues. Respondents are asked to choose one of the two hypothetical candidates that vary along ten dimensions such as age, gender, race, and policy positions. These factors and their levels are given in Table 2.

A key feature of this study is that the randomization occurs in three steps. First, the author randomly selects the number of attributes to be presented to a respondent. The primary factor of interest, candidate party, is always shown, but the remaining nine factors are randomized to be shown or not shown. In particular, the number of additional factors is randomized to be 1, 3, 5, 7, or 9. In the second step, he then randomly chooses the selected number of factors from the remaining nine attributes. Finally, as in a typical conjoint analysis, levels are randomly chosen within each selected factor.

Under this design, Peterson (2017) examines how the effect of copartisanship changes with the amount of information about candidates respondents possess. The original analysis finds that showing more information greatly reduces the effect of copartisanship, suggesting that partisanship partially serves as substitutes for other relevant information. The author also extends this analysis by investigating which factor plays an outsized role in reducing the effect of

copartisanship. This analysis shows that the information about a candidate's position on abortion policy and deficit spending diminish the effect of copartisanship more than demographic features such as race and gender.

3 Causal Quantities of Interest

In this section, we consider causal quantities of interest in conjoint analysis. We first show that most existing conjoint analyses implicitly estimate the *uniform* average marginal component effect (uAMCE), which gives equal weights to all conjoint profiles. Unfortunately, the profile distribution in the real world is likely to be far from uniform. Therefore, we consider an alternative quantity, the *population* average marginal component effect (pAMCE) that directly incorporates the knowledge about the target profile distribution. We discuss the conditions under which the pAMCE differs from the uAMCE.

3.1 The Setup

Following the setup of Hainmueller *et al.* (2014), consider a conjoint analysis with a total of N respondents. In the experiment, each respondent, indexed by $i \in \{1, \dots, N\}$, completes K choice (or rating) tasks, and for a given task, a respondent chooses one of J profiles (or rate each of them). A conjoint profile is composed of L attributes represented by the corresponding L factors, where each factor ℓ has a total of D_ℓ levels. For example, the conjoint analysis of Ono and Burden (2019) has $N = 1,583$ respondents who are assigned to $K = 10$ tasks of choosing one of $J = 2$ candidates. Candidates differ in $L = 13$ factors and the levels of each factor are given in Table 1; for example, $D_1 = 2$ and $D_2 = 4$ where the first and second factors represent gender and race, respectively.

We denote the j th profile presented to respondent i in the k th task by a profile vector \mathbf{T}_{ijk} of length L . The ℓ th element of this vector represents the ℓ th factor of the profile, which takes one of D_ℓ levels, that is, $T_{ijk\ell} \in \{0, 1, \dots, D_\ell - 1\}$. For example, if for the first respondent, the first attribute of the first profile in the first task is male, then we have $T_{1111} = 0$.

Using the potential outcomes framework (Neyman, 1923; Rubin, 1974), let $Y_{ijk}(\mathbf{t})$ represent the potential outcome for respondent i when the stacked vector of J profiles $\mathbf{T}_{ik} = \mathbf{t}$ are presented to respondent i as the k th task. When the outcome is choice-based, only one of J potential outcomes for task k by respondent i takes the value of one whereas the other $J - 1$ potential outcomes are equal to zero. In contrast, when the outcome is rating-based, each outcome Y_{ijk} corresponds to the rating of profile j given by respondent i in the k th task.

This notation is based on the stable unit treatment value assumption (Cox, 1958; Rubin, 1990). In particular, we assume no carryover effect, implying that the outcome of a task is not affected by the same respondent's previous tasks (Hainmueller *et al.*, 2014). In addition, it is often assumed that the position of profiles does not affect the outcome (Hainmueller *et al.*, 2014). Under these assumptions, the potential outcome $Y_{ijk}(\mathbf{t})$ can be simplified as $Y_{ik}(\mathbf{t})$ because respondents would reveal the same outcomes regardless of positions of profiles j .

Under this framework, we review the definition of the AMCE originally proposed by Hainmueller *et al.* (2014). The AMCE represents the average causal effect of changing levels within each factor while averaging over other factors. For example, we might be interested in estimating the effect of a candidate's gender, averaging over the distribution of the other candidate characteristics such as age, ideology, and policy positions.

DEFINITION 1 (Average Marginal Component Effect (Hainmueller *et al.*, 2014)). *The average causal effect of changing factor ℓ from level t_0 to t_1 for a given profile while averaging over the*

other factors is given by,

$$\tau_\ell(t_1, t_0; \Pr(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k})) = \sum_{(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) \in \mathcal{T}} \mathbb{E} [Y_{ik}(t_1, \mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) - Y_{ik}(t_0, \mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k})] \times \Pr(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}),$$

where $\mathbf{t}_{ijk,-\ell}$ represents an $(L - 1)$ dimensional vector representing the levels of all factors except for factor ℓ of the j th profile in the k th task completed by respondent i , $\mathbf{t}_{i,-j,k}$ denotes the levels of all factors for the remaining profiles other than profile j , and \mathcal{T} is the support of $\Pr(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k})$. Finally, the expectation is over a random sample of the respondents and task positions.

At its core, the AMCE averages not only across respondents but also across conjoint profiles, such as political candidates. We show below that this marginalizing distribution over profiles plays an essential role in conjoint analysis.

3.2 The Uniform Average Marginal Component Effect

The definition of the AMCE clearly shows that the use of different profile distributions $\Pr(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k})$ can lead to substantively different conclusions (Hainmueller *et al.*, 2014). Nevertheless, in practice, little attention is paid to the choice of this profile distribution. In particular, most existing conjoint analyses use the uniform distribution, in which each factor is independently and uniformly randomized, making each conjoint profile equally likely. We call the resulting quantity as the *uniform* average marginal component effect (the uAMCE).

DEFINITION 2 (Uniform Average Marginal Component Effect). *The uniform average causal effect of changing factor ℓ from level t_0 to t_1 for a given profile while marginalizing the other factors is given by,*

$$\tau_\ell^u(t_1, t_0) = \tau_\ell(t_1, t_0; \Pr^u(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k})) = \sum_{(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) \in \mathcal{T}^u} \mathbb{E} [Y_{ik}(t_1, \mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) - Y_{ik}(t_0, \mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k})] \Pr^u(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}),$$

where $\Pr^u(\cdot)$ denotes the uniform distribution and \mathcal{T}^u is the support of $\Pr^u(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k})$.

The central problem of the uAMCE is that it equally weights all profiles regardless of how realistic they are. Because any AMCE represents a weighted average of causal effects across all profiles used in the experiment, the estimates partially based on unrealistic profiles may yield misleading findings. The problem is not entirely new. In fact, users of conjoint experiments are often concerned about unrealistic profiles and remove highly unlikely profiles (e.g., Hainmueller *et al.*, 2014). Although this restricted randomization can eliminate extreme cases (e.g., doctors without college degree), the overall distribution of profiles may still be far away from a target profile distribution. Given that one of the core advantages of conjoint experiments is to mimic real-world decision making process (Hainmueller *et al.*, 2015), it is critical to define causal quantity of interest that reflects a target population.

3.3 The Population Average Marginal Component Effect

To improve the external validity of conjoint analysis, we consider the *population* AMCE (pAMCE), which marginalizes factors over the target population distribution of profiles rather than the uniform distribution. This target population of profiles depends on the substantive context of each application, similarly to survey research where a target population of respondents must be

specified. This can be obtained from a real world data set on the attributes of actual politicians as in the case of Ono and Burden (2019) study (our first application). Alternatively, it can be a counterfactual distribution of theoretical interest that, for example, represents a different information environment for voters as in the Peterson (2017) study (our second application). Formally, we define the pAMCE as follows.

DEFINITION 3 (Population Average Marginal Component Effect). *The population average causal effect of changing factor ℓ from level t_0 to t_1 for a given profile while marginalizing the other factors is given by,*

$$\begin{aligned} \tau_\ell^*(t_1, t_0) &= \tau_\ell(t_1, t_0; \Pr^*(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k})) \\ &= \sum_{(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) \in \mathcal{T}^*} \mathbb{E} \left[Y_{ik}(t_1, \mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) - Y_{ik}(t_0, \mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) \right] \Pr^*(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}), \end{aligned}$$

where $\Pr^*(\cdot)$ denotes the target population distribution and \mathcal{T}^* is the support of $\Pr^*(\mathbf{T}_{ijk,-\ell}, \mathbf{T}_{i,-j,k})$.

The distinction between the uAMCE and the pAMCE is simple and yet important. While the uAMCE marginalizes other factors over the uniform distribution, the pAMCE averages them over the target population distribution of profiles. Therefore, the pAMCE appropriately weights profiles according to the frequency with which they occur in the target distribution. Formally, we can characterize the difference between these two quantities as follows,

$$\begin{aligned} &\tau_\ell^*(t_1, t_0) - \tau_\ell^U(t_1, t_0) \\ &= \sum_{(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) \in \mathcal{T}^* \cup \mathcal{T}^U} \mathbb{E} \left[\{ Y_{ik}(t_1, \mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) - Y_{ik}(t_0, \mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) \} \right. \\ &\quad \left. - \{ Y_{ik}(t_1, \mathbf{t}'_{ijk,-\ell}, \mathbf{t}'_{i,-j,k}) - Y_{ik}(t_0, \mathbf{t}'_{ijk,-\ell}, \mathbf{t}'_{i,-j,k}) \} \right] \\ &\quad \times \{ \Pr^*(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) - \Pr^U(\mathbf{t}_{ijk,-\ell}, \mathbf{t}_{i,-j,k}) \}. \end{aligned} \tag{1}$$

This difference between the uAMCE and the pAMCE has two components. The first term quantifies the average causal interaction effect between the factor of interest and all the other factors including those of other profiles (Egami and Imai, 2019). For example, the effect of being female relative to male might be larger for white candidates than black candidates. The second term represents the difference between the uniform and the target profile distributions. Therefore, the difference between the uAMCE and the pAMCE is large when the causal effect of factor ℓ interacts with other factors and when the target profile distribution is far away from the uniform distribution.

3.4 Empirical Illustrations

Using the two studies introduced in Section 2, we empirically illustrate the importance of target profile distributions. For the first application, there exists a natural real-world profile distribution that can be used to estimate the pAMCE. Using data on the characteristics of actual politicians, we construct a distribution of profiles that more accurately reflects what real-world politicians look like. We show that this distribution is strikingly different from uniform. In our second application, we demonstrate how the pAMCE can be useful even when there exists no natural real-world profile distribution for which data can be collected. Specifically, we analyze theoretically relevant counterfactual distributions and systematically investigate how empirical findings change according to the choice of profile distributions.

3.4.1 *The Use of Real-world Distributions.* As in the vast majority of conjoint analyses, Ono and Burden (2019) randomize factors independently by choosing each level with equal probability. This produces a uniform distribution in which all attribute combinations are equally likely. While the uniform distribution is commonly used in applications of the conjoint analysis, the corresponding real-world distribution of attributes are rarely uniform.

Indeed, the uniform randomization produces highly unusual profiles. For example, two-thirds of Republican candidate profiles will have abortion positions of “neutral” or “pro-choice.” The difference between this distribution and the one that of actual Republican politicians is stark. Using a legislator scorecard produced by the National Right to Life Council, a conservative nonprofit that advocates for pro-life policies, only 1 of the 296 Republican legislators (0.33%) could be classified as pro-choice and only 2 (0.67%) as neutral. A similar pattern emerges for Democrats. Two-thirds of presented candidate profiles take a value of neutral or pro-life, yet similarly low percentages of Democratic politicians hold those positions.

The case of abortion position may be especially dramatic, but the real-world distributions of nearly all of the attributes presented in Table 1 differ markedly from the uniform distribution. As a target profile distribution, we use data of actual legislators in the 115th Congress and compute the real-world joint distribution of 12 of the 13 attributes examined in Ono and Burden (2019). We do not produce the distribution of the *Trait* attribute due to its highly subjective quality and thus keep the uniform distribution for it. Because party is strongly correlated with nearly all remaining attributes, we consider the target profile distributions of Republican and Democrat politicians separately. Supplemental Appendix B includes details about the construction of this joint distribution.

Figure 2 shows that the marginal distributions of actual politicians’ characteristics (gray bars for Democrats, shaded bars for Republicans) differ substantially from the uniform distribution (white bars). In the case of the gender, which is the focus of the original analysis, neither the Republican nor Democratic distributions resemble the uniform: only 10.2% of Republicans and 32.2% of Democratic legislators are female. We find a similar pattern for the remaining attributes. The difference is most pronounced for the attributes that are likely to be salient to subjects, such as race and major policy positions. This suggests that the uAMCE may significantly differ from the pAMCE.

Finally, we note that the original experiment considers hypothetical political candidates. Thus, the ideal target profile distribution would be the real-world distribution of the attributes for all candidates, not only for elected legislators. Unfortunately, because the original conjoint experiment was not designed with fidelity to the real-world distribution in mind, there are many factors for which it is not possible to gather corresponding real-world distributions using data from all candidates. As a result, we use politicians in the 115th Congress as our main target profile distribution, for whom we were able to collect real-world distributions for most factors (as visualized in Figure 2).

In Section 5.1, we consider the robustness of the pAMCE estimates by replacing profile distributions of race, gender, and experience, based on publicly available candidate-level datasets. We also consider different theoretically relevant profile distributions on policy dimensions. Even when it is infeasible to collect the real-world distribution of all factors for all candidates, it is critical to take into account more realistic profile distributions and improve the external validity of conjoint analysis.

3.4.2 *The Use of Counterfactual Distributions.* Peterson (2017) is primarily interested in how the effect of copartisanship changes according to the amount of other relevant information about candidates. Therefore, our analysis focuses on the first two steps of the original randomization—randomizing the number of factors to show and then randomly selecting which factor to present given the




Figure 2. Experimental and target profile distributions of factors in Ono and Burden (2019). We compare the uniform distribution used in the original experiment and two real-world distributions of politicians' characteristics and policy positions; Republican and Democrat legislators.

selected number of factors to be shown. Because each randomization uses the uniform distribution, every factor is equally likely to be shown. In particular, the marginal probability of each factor being shown is a little above 50% (see Figure 3). If researchers use the widely used linear regression estimator (Hainmueller *et al.*, 2014), the resulting estimate of the AMCE represents the causal effect of copartisanship while averaging over low, medium, and high information environments.

Rather than averaging over different information environments that have distinct substantive meanings, we may be interested in investigating how the pAMCE depends on different information environments. In particular, we consider two counterfactual distributions: a low information environment in which subjects observe each factor (other than copartisanship) only 20% of the time, and a high information environment in which each factor is observed 80% of the time. Figure 3 compares these low- and high-information counterfactual distributions to the one used in the original analysis. As the figure demonstrates, these low and high-information environments differ substantially from the medium-information environment produced by the original design. This suggests that the AMCE estimate based on the conventional regression estimator may differ from the pAMCE s based on the two counterfactual distributions representing specific information environments of theoretical interest. The framework of the pAMCE is




Figure 3. Original and counterfactual distributions of factors in the information experiment (Peterson, 2017). We compare the distribution used in the experiment and two counterfactual distributions of information environment.

essential to systematically assess how the AMCE estimates might change under different profile distributions.

4 The Proposed Methodology

In this section, we propose two new approaches to estimate the pAMCE. First, we show how to conduct a *design-based confirmatory analysis*, in which we incorporate target profile distributions when designing experiments. In contrast, the second approach—a *model-based exploratory analysis*—takes into account target distributions after randomizing profiles. This latter approach is useful in estimating the pAMCE from existing conjoint experiments that have randomized profiles with distributions different from the target population.

4.1 Design-Based Confirmatory Analysis

The proposed design-based confirmatory analysis consists of new experimental designs and their associated estimators of the pAMCE. We describe each in turn.

4.1.1 Experimental Designs. We introduce three experimental designs; the joint population randomization design, the marginal population randomization design, and the mixed randomization design. While all experimental designs allow for the consistent estimation of the pAMCE, they differ in terms of data requirements and assumptions.

We begin with the *joint population randomization design*. In this design, researchers randomize profiles according to their target profile distribution.

DEFINITION 4 (Joint Population Randomization Design).

$$\Pr^R(\mathbf{T}_{ik} = \mathbf{t}) = \Pr^*(\mathbf{T}_{ik} = \mathbf{t}) \quad \text{for all } \mathbf{t} \in \text{support of } \mathbf{T}_{ik} \text{ and for all } i \text{ and } k, \quad (2)$$

where $\Pr^R(\cdot)$ denotes the distribution used for randomization and $\Pr^*(\cdot)$ represents the target profile distribution.

This design is simple and intuitive since it directly incorporates the target profile distribution into randomization. The main advantage is that the design allows for nonparametric estimation of the pAMCE using a weighted difference-in-means estimator described in the next section.

While the joint population randomization design enables nonparametric estimation, it requires the knowledge of the joint distribution of profile attributes. In practice, this requirement might be difficult to satisfy for many applications. An alternative design that relaxes this stringent data requirement is the *marginal population randomization design*. Under this design, researchers randomize each factor independently according to its marginal profile distribution of the target population.

DEFINITION 5 (Marginal Population Randomization Design).

$$\Pr^R(T_{ijk\ell} = t) = \Pr^*(T_{ijk\ell} = t) \quad \text{for all levels } t \text{ and for all } i, j, k, \ell. \quad (3)$$

For example, we randomize three factors {Gender, Race, Education} independently with each marginal distribution, $\Pr^*(\text{Gender})$, $\Pr^*(\text{Race})$, and $\Pr^*(\text{Education})$, respectively, rather than using the joint distribution $\Pr^*(\text{Gender}, \text{Race}, \text{Education})$.

The main advantage of this approach is that it only requires information about separate marginals of the target profile distribution. Gathering data on marginal distributions is likely to be easier in most contexts. In fact, some researchers have begun to incorporate marginal distributions of the target profile population in their research (see Leeper and Robison, 2018). Another significant benefit is that we can estimate the pAMCE using simple difference-in-means under this design. In practice, this means that researchers can estimate the pAMCE using a linear regression because factors are independent of each other.

The marginal population randomization design is not free of limitations. In particular, without further assumptions, this design estimates the approximate pAMCE where we only partially capture the target profile distribution. Nevertheless, compared to the uAMCE, this design already greatly improves the external validity of conjoint analysis. Indeed, a similar approximation strategy is often used in other contexts, including survey research, in which sampling weights are computed using population marginals, and causal inference with observational data, in which observed covariates are balanced only with respect to their marginal means.

What assumption is required for the consistent estimation of the pAMCE only with separate marginal distributions rather than the joint distribution of profile attributes? It turns out that we only need to assume the absence of three-way or higher order interactions among factors. Suppose that there are three factors Gender, Race, and Education, and they have two-way interactions; Gender \times Race, Gender \times Education, and Race \times Education. In this case, a simple difference-in-means estimator is still consistent for the pAMCE so long as there exists no three-way or higher order interaction such as Gender \times Race \times Education. It is important to emphasize that the marginal population randomization design allows for the existence of any two-way interaction, which often captures the strongest interaction in many applications.

There are several ways to address concerns about the assumption of no three-way or higher-order interaction. First, researchers can extend this marginal population randomization design by incorporating the partial joint distributions. Suppose that the joint distribution $\Pr^*(\text{Race}, \text{Education})$ is available while all other factors are randomized independently according to their separate marginal distributions. In this case, we can consistently estimate the pAMCE of Gender via a weighted difference-in-means (see Section 4.1.2) even when there exists the three-way interaction Gender \times Race \times Education if the joint distribution of Race and Education is incorporated into randomization. In general, if we incorporate the joint distributions of M factors, the consistent estimation of the pAMCE is possible even if there exist $(M + 1)$ -way interactions involving these factors. Finally, we can test the assumption of no three-way and higher-order interactions using the standard F -test (see Section 4.2.4).

As the final design, we introduce the *mixed randomization design*, which can yield more efficient estimates when researchers are interested in only a small number of factors (e.g., one or two) and view the remaining factors as background information they control for. For this design, we first separate L factors into two types $\mathbf{T} = \{\mathbf{T}^M, \mathbf{T}^C\}$; (1) *main factors* of interest \mathbf{T}^M , for which researchers wish to estimate the pAMCE, and (2) *control factors* \mathbf{T}^C , which are included as the background information. The distinction between the main and control factors is essential because there is a statistical tradeoff; as the number of the main factors increases, the estimation of the pAMCE becomes less precise. Under the mixed randomization design, we randomize the main factors of interest based on the uniform distribution and the control factors based on their target profile distribution.

DEFINITION 6 (Mixed Randomization Design).

$$\begin{aligned} \text{Main factors : } & \Pr^R(T_{ijk\ell} = t) = \frac{1}{D_\ell} \text{ for all levels } t \text{ in factor } \ell \in \mathcal{M} \text{ and for all } i, j, k \\ \text{Control factors : } & \Pr^R(\mathbf{T}^C_{ik} = \mathbf{t}) = \Pr^*(\mathbf{T}^C_{ik} = \mathbf{t}) \text{ for all } i \text{ and } k. \end{aligned} \tag{4}$$

For example, as in the original study (Ono and Burden, 2019), suppose researchers are primarily interested in estimating the pAMCEs of factor Gender and use the other 12 factors as control factors. Under the mixed design, we randomize Gender using uniform while randomizing the remaining factors based on their target profile distribution.

This design has two primary advantages. First, by prespecifying a small number of main factors at the design stage, researchers can increase the research transparency and credibility in the same way that preregistration does (Blair *et al.*, 2019). Second, under the mixed randomization design, we can often estimate the pAMCEs of the main factors more efficiently than under the two alternative designs. In fact, we show that when researchers have a single main factor, the mixed randomization design is optimal under the assumption of no cross-profile interaction effects (see Supplemental Appendix D.3). In contrast, when multiple factors are of interest, the comparison of statistical efficiency across the three designs gives an inconclusive answer (see Section 4.1.3 for the sample size formula).

4.1.2 *The Weighted Difference-in-Means Estimator.* We introduce a general weighted difference-in-estimator that is consistent for the pAMCE under all three experimental designs described above. We then show how this general estimator can be simplified under some designs.

Formally, the weighted difference-in-means estimator of the pAMCE can be written as follows (Hájek, 1971),

$$\widehat{\tau}_\ell^*(t_1, t_0) = \frac{\sum_{i=1}^N \sum_{j=1}^J \sum_{k=1}^K 1\{T_{ijk\ell} = t_1\} w_{ijk\ell} Y_{ijk}}{\sum_{i=1}^N \sum_{j=1}^J \sum_{k=1}^K 1\{T_{ijk\ell} = t_1\} w_{ijk\ell}} - \frac{\sum_{i=1}^N \sum_{j=1}^J \sum_{k=1}^K 1\{T_{ijk\ell} = t_0\} w_{ijk\ell} Y_{ijk}}{\sum_{i=1}^N \sum_{j=1}^J \sum_{k=1}^K 1\{T_{ijk\ell} = t_0\} w_{ijk\ell}}, \tag{5}$$

where the weights are defined as,

$$w_{ijk\ell} = \frac{1}{\Pr^R(T_{ijk\ell} | \mathbf{T}_{ijk,-\ell}, \mathbf{T}_{i,-j,k})} \times \frac{\Pr^*(\mathbf{T}_{ijk,-\ell}, \mathbf{T}_{i,-j,k})}{\Pr^R(\mathbf{T}_{ijk,-\ell}, \mathbf{T}_{i,-j,k})}. \tag{6}$$

The weights equal the product of two terms. The first term represents the randomization distribution of $T_{ijk\ell}$ given all the other factors $\{\mathbf{T}_{ijk,-\ell}, \mathbf{T}_{i,-j,k}\}$, whereas the second term is the ratio between the target profile distribution of $\{\mathbf{T}_{ijk,-\ell}, \mathbf{T}_{i,-j,k}\}$ and their randomization distribution. Therefore, the weights are greater for observations that are more prevalent in the target profile

distribution than in the randomization distribution. We prove the consistency of this estimator in Supplemental Appendix D.1.

Under the joint population randomization design, the second term of the weights is equal to one and thus, weights are simplified as follows,

$$w_{ijk\ell}^{\text{Joint}} = \frac{1}{\Pr^*(T_{ijk\ell} \mid \mathbf{T}_{ijk,-\ell}, \mathbf{T}_{i,-j,k})}$$

Under the marginal population randomization design, both the first and second terms are canceled out and hence, weights are equal to one for all observations. Therefore, simple difference-in-means is consistent for the pAMCE under the assumption of no three-way or higher-order interaction.

RESULT 1 ((Estimation under Marginal Population Randomization Design)). Under the assumption of no three-way or higher-order interaction, the following simple difference-in-means estimator is consistent for the pAMCE after randomizing profiles according to the marginal population randomization design (Equation (3)).

$$\frac{\sum_{i=1}^N \sum_{j=1}^J \sum_{k=1}^K 1\{T_{ijk\ell} = t_1\} Y_{ijk}}{\sum_{i=1}^N \sum_{j=1}^J \sum_{k=1}^K 1\{T_{ijk\ell} = t_1\}} - \frac{\sum_{i=1}^N \sum_{j=1}^J \sum_{k=1}^K 1\{T_{ijk\ell} = t_0\} Y_{ijk}}{\sum_{i=1}^N \sum_{j=1}^J \sum_{k=1}^K 1\{T_{ijk\ell} = t_0\}} \xrightarrow{p} \tau_{\ell}^*(t_1, t_0) \quad (7)$$

This difference-in-means estimator can be computed by regressing Y_{ijk} on an intercept and $\mathbf{X}_{ijk\ell}$ with regression where $\mathbf{X}_{ijk\ell}$ is a vector of $(D_{\ell} - 1)$ dummy variables for the levels of $T_{ijk\ell}$ excluding the baseline level t_0 . Then, this difference-in-means estimator equals the estimated coefficient on the dummy variable for the level t_1 of factor ℓ (Greene, 2011; Hainmueller *et al.*, 2014). We provide the proof in Supplemental Appendix D.2.

Finally, under the mixed randomization design, while weights do not have a simple expression, we can use the general weighted difference-in-means estimator given in Equation (5).

In practice, the proposed weighted difference-in-means estimator can be computed via a weighted linear regression model.⁴ Since the weighted linear regression is used only to compute the nonparametric weighted difference-in-means estimator, no additional modeling assumption is imposed. This weighted regression estimator generalizes the regression estimator proposed in Hainmueller *et al.* (2014).

4.1.3 *Effective Sample Size.* When using the proposed weighting estimator, it is important to compute the effective sample size (ESS) to determine the statistical efficiency of each design prior to conducting an experiment. We use Monte Carlo simulation by randomizing profiles according to a specific design and then compute the ESS as follows (Kish, 1965),

$$ESS = \frac{(\sum_{ijk} w_{ijk\ell})^2}{\sum_{ijk} w_{ijk\ell}^2} \quad (8)$$

When weights are equal to one for every observation, the ESS is equal to the total sample size NJK . As weights diverge from one, the ESS becomes smaller. Using ESS, we can easily compute

4 As before, the weighted difference-in-means estimator defined in Equation (5) can be computed by regressing Y_{ijk} on an intercept and $\mathbf{X}_{ijk\ell}$ with weights $w_{ijk\ell}$ where $\mathbf{X}_{ijk\ell}$ is a vector of $(D_{\ell} - 1)$ dummy variables for the levels of $T_{ijk\ell}$ excluding the baseline level t_0 . Then, the weighted difference-in-means estimator equals the estimated coefficient on the dummy variable for the level t_1 of factor ℓ .

the following standard error multiplier between any two designs,

$$\sqrt{\frac{\text{ESS under one design}}{\text{ESS under another design}}}, \tag{9}$$

which quantifies the expected ratio of standard error that would result under one design over that under another design. By computing the ESS and the standard error multiplier at the design stage, researchers can choose an experimental design that most efficiently estimates the pAMCEs. Note that since weights are different for each pAMCE, we must compute these statistics separately.

4.2 Model-Based Exploratory Analysis

When researchers incorporate the target profile distribution at the design stage, the above approach estimates the pAMCEs without bias. In some cases, however, we may wish to explore the pAMCEs of various factors using a conjoint experiment that has been fielded using profile distributions different from the target population. This is especially important when there exists no natural target profile distribution, leading to the use of the uniform randomization. Even in such cases, it is essential to examine the robustness of the AMCE estimates to alternative profile distributions that are of theoretical interest. To do so, we introduce a model-based estimator. While it requires additional modeling assumptions, this approach is useful for exploratory and sensitivity analyses. We also provide diagnostic tools for relevant modeling assumptions in Supplemental Appendix E.

4.2.1 Latent Utility Model. We begin by introducing a latent utility model that allows all two-way interactions. Specifically, we assume that the latent utility for each profile is a function of the main effect of each factor, the two-way interactions between all the factors, and the two-way interaction of the same factor between the two profiles within a given pair (e.g., the effect of age of one profile may depend on the age of the other profile). The modeling assumption is violated if three-way or higher order interaction effects exist. Although we believe that in most practical settings this assumption approximately holds, we offer a simple model specification test in Section 4.2.4.

Formally, our latent utility model of respondent i for profile j when compared against profile j' in the k th task is defined as follows,

$$\begin{aligned} \tilde{Y}_{ijk}(\mathbf{T}_{ijk}, \mathbf{T}_{ij'k}) = & \tilde{\alpha} + \sum_{\ell=1}^L \mathbf{X}_{ij\ell}^T \tilde{\beta}_{\ell} + \sum_{\ell=1}^L \sum_{\ell' \neq \ell} (\mathbf{X}_{ij\ell} \times \mathbf{X}_{ij\ell'})^T \tilde{\gamma}_{\ell\ell'} - \sum_{\ell=1}^L \mathbf{X}_{ij'\ell}^T \tilde{\beta}_{\ell} \\ & - \sum_{\ell=1}^L \sum_{\ell' \neq \ell} (\mathbf{X}_{ij'\ell} \times \mathbf{X}_{ij'\ell'})^T \tilde{\gamma}_{\ell\ell'} + \sum_{\ell=1}^L (\mathbf{X}_{ij\ell} \times \mathbf{X}_{ij'\ell})^T \tilde{\delta}_{\ell\ell} + \tilde{\epsilon}_{ijk}, \end{aligned} \tag{10}$$

where $\mathbf{X}_{ij\ell}$ is a vector of $(D_{\ell} - 1)$ dummy variables for the levels of $T_{ij\ell}$ excluding the baseline level and \times represents the cartesian product operator, for example, $(\mathbf{X}_{ij\ell} \times \mathbf{X}_{ij\ell'})^T \tilde{\gamma}_{\ell\ell'} = \sum_{d=1}^{D_{\ell}-1} \sum_{d'=1}^{D_{\ell'}-1} X_{ij\ell d} X_{ij\ell' d'} \tilde{\gamma}_{\ell\ell' d d'}$. The coefficients $\tilde{\beta}_{\ell}$ denote the main effects of factor ℓ , while the coefficients $\tilde{\gamma}_{\ell\ell'}$ indicate two-way interactions between the two factors ℓ and ℓ' . Finally, the coefficients $\tilde{\delta}_{\ell\ell}$ represent two-way interactions between factor ℓ across the two profiles j and j' . Under the assumption of no profile-order effects, the effects of factors in profile j and those in profile j' are symmetric. This is why the effect of $\mathbf{X}_{ij\ell}$ is $\tilde{\beta}_{\ell}$ and that of $\mathbf{X}_{ij'\ell}$ is $-\tilde{\beta}_{\ell}$. Similarly, the effect of $\mathbf{X}_{ij\ell} \times \mathbf{X}_{ij\ell'}$ is $\tilde{\gamma}_{\ell\ell'}$ while that of $\mathbf{X}_{ij'\ell} \times \mathbf{X}_{ij'\ell'}$ is $-\tilde{\gamma}_{\ell\ell'}$.

As in the conventional latent utility model, we do not directly observe the latent utility. Instead, we observe the choices made by respondents. Each respondent is assumed to choose profile j

when its latent utility is higher than the latent utility of the other profile j' , that is,

$$Y_{ik}(\mathbf{T}_{ijk}, \mathbf{T}_{ij'k}) = \begin{cases} 1 & \text{if } \tilde{Y}_{ijk}(\mathbf{T}_{ijk}, \mathbf{T}_{ij'k}) > \tilde{Y}_{ij'k}(\mathbf{T}_{ijk}, \mathbf{T}_{ij'k}), \\ 0 & \text{otherwise.} \end{cases}$$

There are many ways to connect the latent utility model to the choice outcome model. For example, when we assume the error term follows the type I extreme value distribution, we obtain the well-known conditional logit model (McFadden, 1974). For the ease of interpretation, we rely on the following linear probability model (Egami and Imai, 2019),

$$\begin{aligned} & \Pr(Y_{ik} = 1 \mid \mathbf{T}_{ijk}, \mathbf{T}_{ij'k}) \\ &= \left\{ \tilde{Y}_{ijk}(\mathbf{T}_{ijk}, \mathbf{T}_{ij'k}) - \tilde{Y}_{ij'k}(\mathbf{T}_{ijk}, \mathbf{T}_{ij'k}) \right\} + 0.5 \\ &= \alpha + \sum_{\ell=1}^L (\mathbf{X}_{ijk\ell} - \mathbf{X}_{ij'k\ell})^\top \beta_\ell + \sum_{\ell=1}^L \sum_{\ell' \neq \ell} (\mathbf{X}_{ijk\ell} \times \mathbf{X}_{ijk\ell'} - \mathbf{X}_{ij'k\ell} \times \mathbf{X}_{ij'k\ell'})^\top \gamma_{\ell\ell'} + \sum_{\ell=1}^L (\mathbf{X}_{ijk\ell} \times \mathbf{X}_{ij'k\ell})^\top \delta_{\ell\ell} \end{aligned} \tag{11}$$

where the coefficients have direct connections to the latent utility model given in Equation (10), that is, $\beta_\ell = 2\tilde{\beta}_\ell$, $\gamma_{\ell\ell'} = 2\tilde{\gamma}_{\ell\ell'}$ and $\delta_{\ell\ell} = 2\tilde{\delta}_{\ell\ell}$. We estimate this linear probability model via ordinary least squares by regressing Y_{ijk} on an intercept, the difference in the main terms for all the factors, the difference in the interaction terms for all the two-way interactions, and the interaction terms across profiles for all the factors.

This model does not impose the linearity assumption because each level of a given factor enters the model as a separate dummy variable. The model also allows for all two-way interactions between and across profiles. Therefore, the key assumption is the absence of three-way or higher order interactions, which can be easily relaxed at the expense of statistical efficiency.

4.2.2 *Estimation of the Population AMCE.* Using the above linear probability model, we can estimate the pAMCE as a weighted average of the estimated coefficients,

$$\widehat{\tau}_\ell^*(t_1, t_0) = \widehat{\beta}_{\ell 1} + \sum_{\ell' \neq \ell} \sum_{d=1}^{D_{\ell'}-1} \widehat{\gamma}_{\ell 1 \ell' d} \Pr^*(T_{ijk\ell'} = d) + \sum_{d=1}^{D_\ell-1} \widehat{\delta}_{\ell 1 \ell d} \Pr^*(T_{ijk\ell} = d). \tag{12}$$

where the marginal distributions are used as weights. Thus, under the two-way interactive linear probability model, we only need to collect the marginal distributions of the target profile population $\Pr^*(T_{ijk\ell} = d)$. This greatly relaxes data requirements in practice.

As we saw earlier, when there is no interaction between or across factors, the uAMCE equals the pAMCE. That is, when $\widehat{\gamma}_{\ell 1 \ell' d} = \widehat{\delta}_{\ell 1 \ell d} = 0$, we have $\widehat{\tau}_\ell^*(t_1, t_0) = \widehat{\tau}_\ell^U(t_1, t_0) = \widehat{\beta}_{\ell 1}$. In addition, it is straightforward to estimate the difference between the uAMCE and the pAMCE,

$$\begin{aligned} \widehat{\text{Diff}} &= \widehat{\tau}_\ell^*(t_1, t_0) - \widehat{\tau}_\ell^U(t_1, t_0) \\ &= \sum_{\ell' \neq \ell} \sum_{d=1}^{D_{\ell'}-1} \widehat{\gamma}_{\ell 1 \ell' d} \{ \Pr^*(T_{ijk\ell'} = d) - \Pr^U(T_{ijk\ell'} = d) \} + \sum_{d=1}^{D_\ell-1} \widehat{\delta}_{\ell 1 \ell d} \{ \Pr^*(T_{ijk\ell} = d) - \Pr^U(T_{ijk\ell} = d) \}. \end{aligned}$$

Thus, as mentioned earlier, the difference is large when there exist significant interactions and when the target profile distribution is far away from the uniform distribution. Finally, we can

decompose this difference as the sum of components due to different factors,

$$\widehat{\text{Diff}} = \sum_{\ell'=1}^L \widehat{\text{Diff}}_{\ell'} = \sum_{\ell'=1}^L \sum_{d=1}^{D_{\ell'}-1} \widehat{\gamma}_{\ell' \ell' d} \{ \text{Pr}^*(T_{ijk \ell'} = d) - \text{Pr}^{\text{U}}(T_{ijk \ell'} = d) \}. \tag{13}$$

Through this decomposition, researchers can unpack the origin of the difference between the uAMCE and the pAMCE.

4.2.3 Regularization. The main drawback of the model-based exploratory analysis is its large estimation uncertainty. When there are many factors and each factor has several levels, the model with all two-way interaction effects can produce large standard errors. We consider regularization as a way to partially recoup this loss of statistical efficiency relative to the design-based confirmatory analysis. For example, the conjoint analysis of Ono and Burden (2019) contains 13 factors with a total of 49 levels. This means that the estimated pAMCE will be the weighted average of a large number of interaction terms. In such cases, a regularized regression approach can be effective in reducing estimation uncertainty.

In particular, we follow Egami and Imai (2019) and collapse levels within factors using a regularized regression. For instance, even though Ono and Burden (2019) use six levels for factor Age (36, 44, 52, 60, 68, 76 years old), not all the differences between the six levels may be relevant. It may be, for example, that the effects for the first three levels are indistinguishable from each other and can be collapsed into fewer levels (e.g., 36/44/52, 60/68, 76 years old). We can use a regularized regression to identify such coarsening patterns. By reducing the number of levels, the proposed regularized regression can improve efficiency and estimate the pAMCE more precisely.

Specifically, we estimate the pAMCE by collapsing levels while avoiding regularization bias through cross fitting (Chernozhukov *et al.*, 2018). We begin by randomly splitting data into two parts, training and test data. Using the training data, we first collapse levels within factors via the generalized lasso (Tibshirani and Taylor, 2011),

$$\hat{\beta} = \underset{\beta}{\text{argmin}} \sum_{k=1}^K \sum_{j=1}^J \sum_{i=1}^N \left(Y_{ijk} - \beta_0 - \sum_{\ell=1}^L \mathbf{x}_{ijk\ell}^T \beta_{\ell} \right)^2 + \lambda \sum_{\ell=1}^L \sum_{d=1}^{D_{\ell}-1} \pi_{\ell d} |\beta_{\ell d} - \beta_{\ell, d-1}|, \tag{14}$$

where we select tuning parameter λ using cross validation. By weighting according to effect size, the adaptive weights help regularize smaller effects more and larger effects less.⁵ Importantly, we do not shrink the coefficients $\beta_{\ell d}$ themselves and instead regularize their differences $|\beta_{\ell d} - \beta_{\ell, d-1}|$ so that we can collapse unnecessary levels (Egami and Imai, 2019). When levels are unordered, researchers can use an alternative penalty that regularizes all pairwise differences, that is, $\sum_{\ell=1}^L \sum_{d=0}^{D_{\ell}-1} \sum_{d' \neq d} |\beta_{\ell d} - \beta_{\ell, d'}|$.

Second, using the separate test data, we fit the proposed linear probability model with collapsed levels and then estimate the pAMCE based on the weighted average expression given in Equation (12). Because unnecessary levels are removed in the previous step, we can estimate the pAMCE more precisely. It is important that we collapse levels with the training data and estimate the pAMCE with the separate test data to remove bias due to regularization.

Finally, we flip the role of training and test data and repeat the two steps described above. We average the two estimates from each test data as the estimate of the pAMCE. For uncertainty estimates, we use the block bootstrap by sampling respondents with replacement. We implement the cross-fitting for each bootstrap replicate. Uncertainty estimates are calculated based on

5 Adaptive weights are defined as $\pi_{\ell d} = \sqrt{N_{\ell d} + N_{\ell, d-1}} / |\hat{\beta}_{\ell d}^{OLS} - \hat{\beta}_{\ell, d-1}^{OLS}|$ where $N_{\ell d}$ is the number of observations with $T_{ijk\ell} = t_d$ and $\hat{\beta}_{\ell d}^{OLS}$ is the OLS estimate of $\beta_{\ell d}$ (Gertheiss and Tutz, 2010).

the empirical distribution of the estimated pAMCE over the bootstrap sample. In Supplemental Appendix F, we provide simulation studies to show how much the proposed regularization method can improve efficiency without inducing bias.

4.2.4 *Assessing the Absence of Higher-Order Interaction.* The model introduced above (Equation (11)) as well as the marginal population randomization design (Equation (3)) assumes the absence of three-way or higher order interaction. We can directly test the assumed absence of three-way interaction by conducting the standard F -test. Specifically, we incorporate three-way interactions between three factors ℓ , ℓ' , and ℓ'' by adding $(\mathbf{X}_{ijk\ell} \times \mathbf{X}_{ijk\ell'} \times \mathbf{X}_{ijk\ell''})^T \boldsymbol{\zeta}_{\ell\ell'\ell''}$ to the two-way interactive model of Equation (11) where $\boldsymbol{\zeta}_{\ell\ell'\ell''}$ is a vector of coefficients for the three-way interactions. Then, we test the existence of this three-way interaction via F -test with the null hypothesis, $H_0 : \boldsymbol{\zeta}_{\ell\ell'\ell''} = 0$. When the statistical power of detecting three-way interaction effects is of concern, it is recommended to rely on the regularization approach described above.

4.3 Summary

Table 3 summarizes the methodologies introduced in this section in terms of required data and assumptions. Several points are worth emphasizing. First, if researchers expect the target profile distribution to differ from the uniform distribution and factors to interact with one another, we recommend that they use one of the proposed experimental designs. The design-based approach is considerably more efficient than the model-based approach.

Second, the choice of experimental designs largely depends on the availability of data about the target profile distribution although the sample size calculation can be conducted to compare the statistical efficiency of these designs. Ideally, researchers have the joint distribution, and hence are able to use the joint population randomization design. If only the marginal distributions are available, the marginal population randomization can be used at the cost of making an additional assumption about the absence of third or higher-order interactions. If large higher-order interaction effects are expected, incorporating partial joint distributions can relax the assumption. In addition, the mixed-randomization design is available if researchers are interested in testing hypotheses about one or two factors while controlling for other factors.

Finally, even when there exists no natural target profile distribution for which data can be collected, it is important to conduct the model-based approach to explore the robustness of the AMCE estimates to the choice of profile distributions. We recommend researchers systematically examine different counterfactual profile distributions motivated by a theoretical consideration (see our second example based on Peterson, 2017 in Section 5.2).

Table 3. Data requirements and assumptions of design-based and model-based approaches.

Approach	Data requirement	Assumption	Note
Design-based confirmatory analysis			
• Joint population randomization	Joint distribution over all profile attributes	None	
• Marginal population randomization	Marginal distributions of each profile attribute	Absence of three-way or higher order interaction	Relax the assumption with partial joint distributions
• Mixed randomization	Joint distribution over control factors	None	Efficient when focus on one or two main factors
Model-based exploratory analysis			
• Linear probability model	Marginal distributions of each profile attribute	Absence of three-way or higher order interaction	Relax the assumption with partial joint distributions




Figure 4. Estimates of the pAMCEs of being female in Ono and Burden (2019). We estimate the pAMCE for Republican and Democrat politicians. Even though an estimate of uAMCE is close to zero for congressional candidates, the pAMCE for congressional candidates under the Democrat distribution is large and positive.

5 Empirical Applications

We apply the proposed methodology to the empirical applications described in Section 2. For the first application, we find that two key conclusions regarding the effect of gender are due to the uniform distribution used in the original study. Estimating the pAMCE using the real-world profile distribution instead, we find that the effect of being a female candidate varies according to party and office they seek for. For the second application, we show how to systematically explore the pAMCE based on counterfactual distributions of theoretical interest.

5.1 The Effect of Candidate’s Gender on Voter Choice

In the first application, the primary quantity of interest is the AMCE of candidate’s gender on voter choice. Instead of the uniform distribution used in the original analysis, we estimate the pAMCE using the real-world distribution of elected politicians in the 115th Congress, as described in Section 3.4. In particular, we estimate this quantity separately using the distributions of Democratic and Republican politicians’ characteristics (see Figure 4).

5.1.1 Design-Based Analysis. We begin by performing the design-based confirmatory analysis proposed in Section 4.1. Because in the original study each attribute is randomized according to the uniform distribution, we conduct a simulation study to assess the performance of the marginal population randomization and mixed randomization designs.⁶ To do this, we first fit a linear regression model with all two-way interactions between the thirteen factors summarized in Table 1 and use this estimated model as the true data generating process. For the marginal population randomization design, we randomize each factor independently based on a marginal distribution of the target population. For the mixed randomization design, the primary factor of interest, that

⁶ As we propose in Section 4.1, researchers can directly conduct the design-based confirmatory analysis when researchers can incorporate target profile distributions in the design stage.

is, gender, is randomized according to the uniform distribution and the remaining factors are randomized using their target population distribution. We estimate the pAMCE via the simple difference-in-means under the marginal population design and the weighted difference-in-means under the mixed design. Standard errors are clustered by respondents. We repeat the same procedure 100 times and average over point estimates and standard errors.

The left plot in the top row of Figure 4 presents the results. First, we focus on the results based on the mixed randomization design. In contrast to the estimates of the uAMCE, we find that the pAMCE is estimated to be -2.20 percentage points (95% confidence interval = $[-3.30, -1.10]$) when using the distribution of Republicans and 2.64 percentage points ($[1.53, 3.74]$) for Democrats. We obtain similar results under the marginal population randomization design although the standard errors are slightly larger. Recall that under the uniform distribution, the estimated uAMCE of gender on vote choice is small and negative. This demonstrates that the estimated AMCE critically depends on the target distribution of candidates' attributes.

One key conclusion of the original study is that the negative effect of being female is found only for presidential candidates but not for congressional candidates. We revisit this finding by using the real-world politicians as the target profile distributions. In particular, we now conduct the design-based confirmatory analysis separately for congressional and presidential candidates. These results are presented in the top row of the second and third columns of Figure 4. Consistent with the original analysis, the estimated uAMCE of being female is -0.09 percentage points ($[-1.71, 1.48]$) for congressional candidates and -2.42 percentage points ($[-3.96, -0.88]$) for presidential candidates. For presidential candidates (the third plot in the top row), the pAMCE of being female is similar to the corresponding uAMCE for both Democratic and Republican distributions. Female presidential candidates face barriers compared to male candidates regardless of party. This result shows that the pAMCE and the uAMCE estimates can be similar even when the target profile distribution is far from uniform. This is because there exists little interaction between gender and other factors within this subgroup (see formal discussions in Section 3.3).

Interestingly, for congressional candidates (the second plot in the top row), the results of the uAMCE and pAMCEs diverge. The uAMCE implies that gender has little effect in congressional races. Yet a more realistic profile distribution suggests a more nuanced finding: women are disadvantaged when they run as Republicans and advantaged when they run as Democrats. Under the mixed randomization design, female Republican candidates are 1.98 percentage points ($[0.42, 3.54]$) less likely to be chosen than their male counterparts, while female Democratic candidates are 5.69 percentage points ($[4.13, 7.25]$) more likely to be chosen. The latter effect is large in substantive terms, equaling the effect of candidates' experience in office and their position on deficit reduction.

5.1.2 Model-Based Analysis. Now, we illustrate the model-based exploratory analysis introduced in Section 4.2. This approach is useful especially when researchers are interested in exploring the pAMCE with conjoint experiments that have already been conducted using the uniform or any distributions different from the target distribution. First, we focus on estimating the pAMCE for both presidential and congressional candidates together as done in the original analysis. As explained in Section 4.2.3, we incorporate all two-way interactions among all the thirteen factors in Table 1 except for *Office* and then collapse levels within factors using the generalized lasso. Standard errors are calculated with 2,000 block bootstraps clustered by respondents.

As expected, the results are similar to those from the design-based analysis but with larger standard errors (see the left plot in the bottom row of Figure 4). The estimated pAMCE is -2.87 percentage points ($[-6.39, 0.63]$) when using the distribution of Republican politicians and 2.84 percentage points ($[-0.20, 5.87]$) for Democratic politicians. We also repeat the same analysis for presidential and congressional candidates separately (the second and third plots in the second

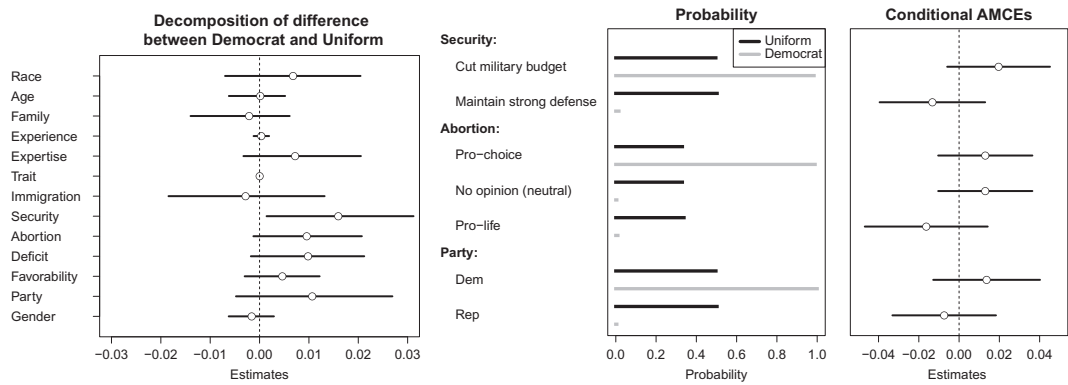


Figure 5. Decomposing the difference between the estimated uAMCE and pAMCE of being a female candidate. For congressional candidates, we compare the uAMCE and the pAMCE based on the Democrat distribution. The first plot decomposes the overall difference into each factor. The second and third factors investigate how effect heterogeneity and the difference in the profile distributions result in the difference in the uAMCE and the pAMCE.

row). As in the design-based confirmatory analysis, we find that female congressional candidates have a disadvantage when they run as Republicans and have an advantage when they run as Democrats.

The standard errors are much larger than those in the design-based confirmatory analysis because the uniform distribution used in the experiment is markedly different from the target profile distribution. This postadjustment of the large differences reduces the effective sample size. Since the model-based analysis marginalizes all the two-way interactions, the efficiency loss is especially severe when the number of factors and levels within each factor are large, as in this example. Although regularization recoups some of this efficiency loss, the design-based analysis yields smaller standard errors in such high dimensional designs.

To investigate the sources of the difference between the uAMCE and pAMCE, we apply the decomposition formula given in Equation (13). For the sake of illustration, we focus on the difference between the estimated uAMCE and pAMCE for congressional Democratic candidates (the second plot in the bottom row). The first plot of Figure 5 shows the results of this decomposition, with each row representing the difference attributable to a single factor. We find, for example, that the difference due to factor *Security Policy* is about 1.6 percentage points ([0.14, 3.12]), implying that the estimated AMCE increases by 1.6 percentage points when we use the distribution of Democratic politicians for *Security Policy* rather than the uniform distribution. Importantly, less than 20% of the overall difference is attributable to *Party*, meaning that we cannot estimate the pAMCE just by considering the interaction between *Gender* and *Party*. The results show that the difference between the uAMCE (−0.09 percentage points) and pAMCE (6.17 percentage points) can be attributed to a combination of many factors even though the contribution of each factor is not necessarily precisely estimated. Even if the difference due to each factor is small, when aggregated across many factors, the overall difference between the uAMCE and the pAMCE can be substantial. This result underscores the need to collect relevant data for as many factors as possible when building the target distribution.

To further unpack the source of this difference, we examine the *conditional average marginal component effect* (cAMCE). The cAMCE is the AMCE of the factor of interest—in this case, gender—conditional on the level of another factor.⁷ A difference in the cAMCEs across the levels of the second factor implies an interaction with the factor of interest. For example, a difference in the cAMCEs of *Gender* conditional on *Security Policy* would indicate an interaction between *Gender* and

7 Within each factor, the weighted sum of the cAMCEs—with weights equal to the population probabilities of each level—is equal to the pAMCE of interest.




Figure 6. Assess the existence of three-way interactions. We compare the pAMCE estimates from models that assume two-way interactions and that incorporate three-way interactions.

Security Policy. We can use the cAMCE to determine whether each factor's contribution to the difference between the uAMCE and the pAMCE (the first plot of Figure 5) is due to large interactions, large changes in distribution, or a combination of the two. If interactions between the primary factor of interest and secondary factors are responsible for most of the difference between the uAMCE and the pAMCE, even small changes in distribution will make the uAMCE different from the pAMCE.

The right two plots of Figure 5 visualize the distributions of three factors Security Policy, Abortion Policy, and Party alongside the cAMCEs of Gender conditional on each of the three factors. For example, the first row in the third plot presents the estimated cAMCE of being female relative to male, conditional on having the Security Policy factor equal to Cut military budget. Focusing on the Security Policy factor, we observe that although its cAMCEs are modest in size, the distribution for Democratic politicians differs substantially from the uniform distribution. Thus, the difference induced by the Security Policy factor is being driven primarily by distributional differences. Repeating this approach for each factor tells us whether the difference between the uAMCE and the pAMCE is a function of distributional changes or causal interactions.

As an important diagnostic, we evaluate the assumption of no three-way interactions. In particular, we incorporate three-way interactions between Gender, Party, and each of the four policy positions given that the difference between the uAMCE and the pAMCE is mostly attributable to those factors. Because we have information about the joint distribution of politicians' attributes, we use them when we marginalize over three-way interactions to estimate the pAMCE. Figure 6 shows that the pAMCE estimates based on the three-way interaction model are similar to those based on the two-way model both in terms of point estimates and standard errors. This result demonstrates that, even when researchers have access only to marginal distributions of the target profile population, it is possible to consistently estimate the pAMCEs by using the marginal population randomization design.

Finally, we examine the robustness of the pAMCE estimates based on the 115th Congress to alternative profile distributions based on the available candidate-level data rather than the data on elected politicians. Although these candidate-level data do not contain information for all factors, we can take into account a number of important candidates' characteristics. In particular, we use DIME data set (Bonica, 2015) and the Reflective Democracy (RefDem) dataset⁸ to obtain the profile distributions of three key variables (race, gender, and experience). We also use our substantive knowledge to investigate different theoretically relevant profile distributions on policy dimensions. We provide details of these alternative profile distributions in Supplemental Appendix C. We show that the pAMCE estimates are robust to these different profile distributions

8 This dataset is available at <https://wholeads.us/resources/for-researchers/>




Figure 7. The estimated population AMCEs of copartisanship in Peterson (2017). We estimate the pAMCE of being copartisan under three different distributions – a medium information distribution and the low and high information distributions.

that more accurately reflect the real-world distribution of political candidates (see Figure A3 in Supplemental Appendix C). These results imply that the difference between the pAMCE and uAMCE is mainly driven by the fact that the uniform distribution is far away from the actual distribution of politicians’ characteristics. In contrast, the difference between the distribution of attributes for elected politicians and that for candidates is relatively minor and has little impact on the empirical findings.

5.2 The Effect of Information Environment on Partisan Voting

In this section, we revisit a major finding of the original study that the importance of copartisanship declines as voters are given more information about candidates.

5.2.1 Design-Based Analysis. We begin with the design-based confirmatory analysis. To estimate the pAMCE, we use the marginal population randomization design by randomizing each factor according to the counterfactual distributions of interest. We also employ the mixed randomization design, retaining the uniform distribution for a primary factor of interest—copartisanship, in this case—and using the counterfactual distributions for all remaining factors. We rely on a weighted difference-in-means estimator (Equation (5)) and cluster standard errors by respondents.

The left plot of Figure 7 presents results of this analysis. Consistent with the original finding, the pAMCE of copartisanship is estimated to be the largest under the low information distribution (61.84 percentage points, [59.06, 64.62]) while the effect is the smallest under the high information distribution (38.39 percentage points, [35.13, 41.65]) using the mixed randomization design. Results are similar under the marginal population randomization design. Thus, the importance of copartisanship in subjects’ voting decisions can vary by more than 20 percentage points depending on the information environment.

5.2.2 Model-Based Analysis. Next, we estimate the same three quantities using model-based exploratory analysis. To do so, we run an unregularized linear regression using all two-way interactions between the ten factors described in Table 2. While regularization is generally preferred, the factors here are binary. Since the goal of regularization is to improve efficiency by collapsing levels of a factor that have similar effects, regularization is not needed in this case. Standard errors are based on 2,000 block bootstraps clustered by respondents.

The second plot in Figure 7 presents these results. As in the design-based confirmatory analysis, the pAMCE of copartisanship is the largest under the low information distribution (61.87

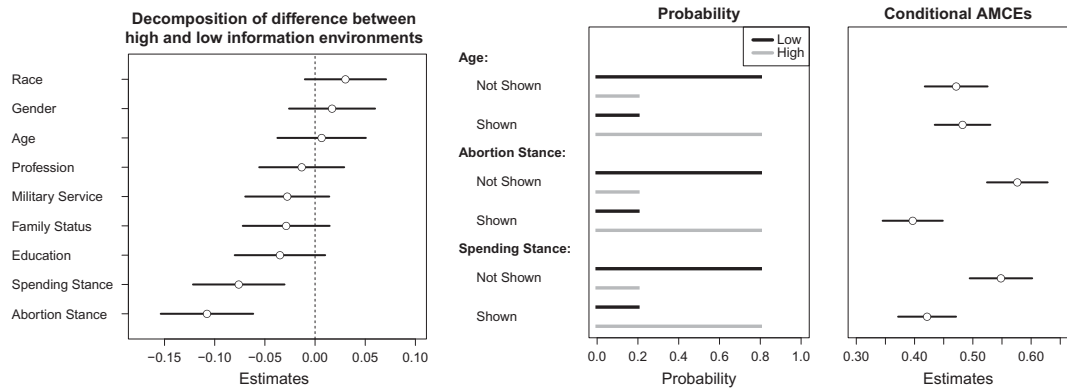


Figure 8. Decomposition of the difference between the two pAMCEs of copartisanship between high and low information environment. The left plot decomposes the overall difference into each factor. The difference is mainly due to the two factors, *Spending Stance* and *Abortion Stance*. The second and third plots investigate how the conditional AMCE and the difference in the profile distributions contribute to the difference.

percentage points, [57.34,66.40]) and the effect is the smallest under the high information distribution (38.21 percentage points, [33.69,42.73]). Although standard errors for the model-based exploratory analysis are larger than those of the design-based confirmatory analysis, the difference between them in this application is relatively small. This is due to the fact that the design in Peterson (2017) is low-dimensional, comprised only of binary factors.

After showing that copartisanship effects are indeed smaller when a larger number of candidate characteristics are shown, the author conducts the second analysis to unpack the mechanism by identifying which information is responsible for reducing the effect of copartisanship. To answer this question, he considers an extreme counterfactual distribution, in which only one factor (in addition to copartisanship) is shown to respondents and examines the difference in the pAMCE of copartisanship with and without this additional factor. The author repeats this analysis separately for each of the nine factors and finds that policy positions on spending and abortion result in the largest differences.

In our pAMCE framework, there is no need to consider each factor in isolation. Instead, we directly examine the sources of the difference in the pAMCE of copartisanship between the low and high information environments. To do so, we use the decomposition formula. The left plot of Figure 8 shows that the difference observed in Figure 7 is mainly driven by two factors, *Spending Stance* (−7.60 percentage points, [−12.16, −3.04]) and *Abortion Stance* (−10.77 percentage points, [−15.38, −6.16]). This result suggests that respondents use copartisanship mainly as a cue for policy stances on spending and abortion, consistent with the original findings.

Finally, we examine why these factors drive the difference in the copartisanship effect between the low and high information environments. The second and third plots of Figure 8 present the distribution and the cAMCEs of each factor. Taking *Spending Stance* as an illustration, we find that the cAMCE for *Shown* (the bottom estimate) is much smaller than for *Not Shown* (the second estimate from the bottom). There is a strong interaction between factors *Party* and *Spending Stance*, yielding the large difference of the pAMCE between the high and low information environments. In contrast, little difference exists in the cAMCEs of copartisanship conditional on *Age* (see the first and second estimates in the third plot). This is why the difference of the pAMCE due to *Age* is small (third estimate in the first plot).

6 Concluding Remarks

Over the last several years, conjoint analysis has become increasingly popular in political science. One advantage of conjoint analysis is its unique ability to help researchers systematically examine

various decision making processes faced by individuals in the real world. This attractive feature has boosted the external validity of empirical conclusions based on conjoint analysis.

Yet, little attention has been paid to the choice of the profile distribution used for randomization. While most researchers use the uniform distribution for convenience, this leads to a causal quantity—the *uniform* average marginal component (uAMCE) effect—that gives equal weights to all possible profiles, including those that rarely occur in the real world.

We address this problem by defining an alternative quantity of interest, the *population* average marginal component effect (pAMCE), using the target profile distribution based on substantive knowledge. We propose new experimental designs and estimation methods for inferring the pAMCE. We then illustrate their use with two empirical applications, one using a real-world distribution and the other based on a counterfactual distribution motivated by a theoretical consideration.

While we focus on the issues related to the distribution of profiles in conjoint analysis, our proposed methodology applies to any factorial experiments with many factors. Moreover, the importance of designing realistic interventions goes beyond conjoint analysis and survey experiments. Indeed, unlike the widely recognized issues related to the representativeness of the experimental sample, the realism of treatments is an essential yet under-appreciated element of external validity. We thus believe that the use of realistic treatments is essential in ensuring the theoretical and practical relevance of any experimental research.

Acknowledgments

We thank Jens Hainmueller, Dan Hopkins, Dean Knox, Shiro Kuriwaki, Thomas Leavitt, Erik Peterson, and Teppei Yamamoto for helpful comments and conversations.

Data Availability Statement

Replication code for this article has been published in Code Ocean, a computational reproducibility platform that enables users to run the code and can be viewed interactively at de la Cuesta, Egami, and Imai (2020a) or at <https://doi.org/10.24433/CO.9475665.v1>. A preservation copy of the same code and data can also be accessed via Dataverse at de la Cuesta *et al.* (2020b) or at <https://doi.org/10.7910/DVN/HVY5GR>.

Supplementary Material

For supplementary material accompanying this paper, please visit <https://doi.org/10.1017/pan.2020.40>.

Bibliography

- Arrow, K. J. 1998. "What Has Economics to Say about Racial Discrimination?" *The Journal of Economic Perspectives* 12(2):91–100.
- Ballard-Rosa, C., L. Martin, and K. Scheve. 2017. "The Structure of American Income Tax Policy Preferences." *The Journal of Politics* 79(1):1–16.
- Bansak, K., J. Hainmueller, and D. Hangartner. 2016. "How Economic, Humanitarian, and Religious Concerns Shape European Attitudes Toward Asylum Seekers." *Science* 354(6309):217–222.
- Barnes, L., J. Blumenau, and B. Lauderdale. 2019. "Measuring Attitudes towards Public Spending using a Multivariate Tax Summary Experiment." Technical report, University College London.
- Bartels, L. M. 2000. "Partisanship and Voting Behavior, 1952–1996." *American Journal of Political Science* 44(1):35–50.
- Blair, G., J. Cooper, A. Coppock, and M. Humphreys. 2019. "Declaring and Diagnosing Research Designs." *American Political Science Review* 113(3):838–859.
- Bolsen, T., J. N. Druckman, and F. L. Cook. 2014. "The Influence of Partisan Motivated Reasoning on Public Opinion." *Political Behavior* 36(2):235–262.
- Bonica, A. 2015. "Database on Ideology, Money in Politics, and Elections (DIME)." <https://doi.org/10.7910/DVN/O5PX0B>, Harvard Dataverse, V3.
- Bullock, J. G. 2011. "Elite Influence on Public Opinion in an Informed Electorate." *The American Political Science Review* 105(3):496–515.

- Campbell, A., P. Converse, W. Miller, and D. Stokes. 1960. *The American Voter*. Hoboken, NJ: Chicago University Press.
- Chernozhukov, V. et al. 2018. "Double Machine Learning for Treatment and Structural Parameters." *Econometrics Journal* 21:C1 – C68.
- Coppock, A., T. J. Leeper, and K. J. Mullinix. 2018. "Generalizability of Heterogeneous Treatment Effect Estimates Across Samples." *Proceedings of the National Academy of Sciences* 115(49):12441–12446.
- Cox, D. R. 1958. *Planning of Experiments*. Hoboken, NJ: Wiley.
- de la Cuesta, B., N. Egami, and K. Imai. 2020a. "Replication Data for: Improving the External Validity of Conjoint Analysis: The Essential Role of Profile Distribution." Code Ocean, V1. doi: <https://doi.org/10.24433/CO.9475665.v1>.
- de la Cuesta, B., N. Egami, and K. Imai. 2020b. "Replication Data for: Improving the External Validity of Conjoint Analysis: The Essential Role of Profile Distribution." <https://doi.org/10.7910/DVN/HVY5GR>, Harvard Dataverse, V1.
- Druckman, J. N. 2014. "Pathologies of Studying Public Opinion, Political Communication, and Democratic Responsiveness." *Political Communication* 31(3):467–492.
- Egami, N., and K. Imai. 2019. "Causal Interaction in Factorial Experiments: Application to Conjoint Analysis." *Journal of the American Statistical Association* 114(526):529–540.
- Gertheiss, J., and G. Tutz. 2010. "Sparse Modeling of Categorical Explanatory Variables." *The Annals of Applied Statistics* 4(4):2150–2180.
- Green, P. E., A. M. Krieger, and Y. Wind. 2001. "Thirty Years of Conjoint Analysis: Reflections and Prospects." *Interfaces* 31(3_supplement):56–73.
- Greene, W. H. 2011. *Econometric Analysis*. London: Pearson.
- Hainmueller, J., D. Hangartner, and T. Yamamoto. 2015. "Validating Vignette and Conjoint Survey Experiments against Real-World Behavior." *Proceedings of the National Academy of Sciences* 112(8):2395–2400.
- Hainmueller, J., and D. J. Hopkins. 2015. "The Hidden American Immigration Consensus: A Conjoint Analysis of Attitudes Toward Immigrants." *American Journal of Political Science* 59(3):529–548.
- Hainmueller, J., D. J. Hopkins, and T. Yamamoto. 2014. "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22(1):1–30.
- Hájek, J. 1971. "Comment on 'An Essay on the Logical Foundations of Survey Sampling, Part One.'" In *The Foundations of Survey Sampling*, edited by V. P. Godambe and D. A. Sprott, 236. New York: Holt, Rinehart, and Winston.
- Huff, C., and J. D. Kertzer. 2018. "How the Public Defines Terrorism." *American Journal of Political Science* 62(1):55–71.
- Kish, L. 1965. *Survey Sampling*. New York: John Wiley & Sons.
- Lau, R. R., and D. P. Redlawsk. 2001. "Advantages and Disadvantages of Cognitive Heuristics in Political Decision Making." *American Journal of Political Science* 45(4):951–971.
- Leeper, T. J., and J. Robison. 2018. "More Important, but for What Exactly? The Insignificant Role of Subjective Issue Importance in Vote Decisions." *Political Behavior* 42:239–259.
- Marshall, P., and E. T. Bradlow. 2002. "A Unified Approach to Conjoint Analysis Models." *Journal of the American Statistical Association* 97(459):674–682.
- McDermott, M. 1997. "Voting Cues in Low-Information Elections: Candidate Gender as a Social Information Variable in Contemporary United States Elections." *American Journal of Political Science* 41(1):270–283.
- McFadden, D. 1974. Conditional Logit Analysis of Qualitative Choice Behavior. In *Frontiers in Econometrics*, edited by P. Zarembka. New York: Academic Press.
- Miratrix, L. W., J. S. Sekhon, A. G. Theodoridis, and L. F. Campos. 2018. "Worth Weighting? How to Think About and Use Weights in Survey Experiments." *Political Analysis* 26(3):275–291.
- Mullinix, K. J., T. J. Leeper, J. N. Druckman, and J. Freese. 2015. "The Generalizability of Survey Experiments." *Journal of Experimental Political Science* 2(2):109–138.
- Mutz, D. C. 2011. *Population-Based Survey Experiments*. Princeton, NJ: Princeton University Press.
- Neyman, J. 1923. "On the Application of Probability Theory to Agricultural Experiments. Essay on Principles (with discussion). Section 9 (translated)." *Statistical Science* 5(4):465–472.
- Ono, Y., and B. C. Burden. 2019. "The Contingent Effects of Candidate Sex on Voter Choice." *Political Behavior* 41:583–607.
- Peterson, E. 2017. "The Role of the Information Environment in Partisan Voting." *The Journal of Politics* 79(4):1191–1204.
- Rubin, D. B. 1974. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology* 66(5):688.
- Rubin, D. B. 1990. "Comment: Neyman (1923) and Causal Inference in Experiments and Observational Studies." *Statistical Science* 5(4):472–480.
- Teele, D. L., J. Kalla, and F. Rosenbluth. 2018. "The Ties That Double Bind: Social Roles and Women's Underrepresentation in Politics." *American Political Science Review* 112(3):525–541.
- Tibshirani, R. J., and J. Taylor. 2011. "The Solution Path of the Generalized Lasso." *The Annals of Statistics* 39(3):1335 – 1371.