

Party Manifestos and Voters' Multidimensional Policy Preferences: Identification via a Conjoint Experiment*

Yusaku Horiuchi[†] Daniel M. Smith[‡] Teppei Yamamoto[§]

First Draft: July 6, 2015

This Draft: September 20, 2016

*Earlier versions of this paper were presented at the 2015 Meeting of the Midwest Political Science Association, the 2015 Meeting of the American Political Science Association, and seminars at Dartmouth, Penn, MIT, Keio, Harvard, Doshisha, and Stanford. We thank Mayumi Fukushima for her excellent research assistance, and Faisal Ahmed, Eric Beerbohm, Simon Chauchard, Kyle Dropp, Noam Gidron, Karen Jusko, Joshua Kertzer, Ellis Krauss, Kuniaki Nemoto, Brendan Nyhan, Steven Reed, Kay Shimizu, Arthur Spirling, Luke Swaine, Seiki Tanaka, and other conference and seminar participants for helpful feedback.

[†]Professor of Government and Mitsui Professor of Japanese Studies, Department of Government, Dartmouth College. Email: yusaku.horiuchi@dartmouth.edu, URL: <https://sites.dartmouth.edu/horiuchi/>

[‡]Assistant Professor, Department of Government, Harvard University. Email: danielm-smith@fas.harvard.edu, URL: <https://sites.google.com/site/danielmarkhamsmith>

[§]Associate Professor, Department of Political Science, Massachusetts Institute of Technology. Email: tepei@mit.edu, URL: <http://web.mit.edu/tepei/www>

Abstract

Representative democracy necessitates the aggregation of multiple policy issues by parties into competing bundles of policies, or “manifestos,” which are then evaluated holistically by voters in elections. This aggregation process complicates our understanding of the *multidimensional* policy preferences underlying a voter’s *single* choice of a party or candidate. We address this problem with a novel conjoint experiment based on actual party manifestos. By juxtaposing sets of issue positions as hypothetical manifestos and asking respondents to choose one, our study identifies the effects of specific positions on voters’ overall assessment of manifestos, heterogeneity in preferences across respondents, and the popularity ranking of manifestos. We illustrate the approach with a field experiment conducted during Japan’s 2014 House of Representatives election. Our analysis uncovers important discrepancies between voters’ preferences and the portrayal of the election by politicians and the media, underscoring the potential danger of inferring voters’ policy preferences only from election outcomes. (150 words)

Keywords: election, representation, manifesto, conjoint analysis, Japan

Word count: 11,022 (excluding the Title Page, Abstract, and Online Appendix)

1 Introduction

Representative democracy entails the aggregation of multiple policy issues by political parties into competing programs, or “manifestos,” which are then evaluated holistically by voters in elections.¹ Indeed, one of the key functions of elections in representative democracies is preference aggregation (for useful reviews, see Powell, 2007; Dewan and Shepsle, 2011). In a direct democracy (e.g., a single-issue referendum), each voter is presented with alternative policies for a given issue and then chooses the policy he or she prefers most. As a result, election results directly reveal the distribution of policy preferences among voters for the specific issue. Direct democracy for all policy decisions, however, is infeasible for populations of large sizes. Thus, modern democracy in most cases is representative democracy, and parties serve the important function of aggregating diverse policies and simplifying alternatives for voters.

Importantly, however, the nature of this preference aggregation process means that voters may not support all of the policy positions of the parties who ultimately get their votes. Voters typically make a *single* choice (i.e., of a particular party or candidate) that represents *multiple* policy positions. For example, a voter may choose a party that simultaneously proposes a tax increase, the allocation of a larger budget for education, and a less aggressive foreign policy, without necessarily supporting each of these components of the party’s policy program.

The problem of incomplete mappings between policy preferences and vote choice raises some important normative concerns about the functioning of representative democracy. For example, party leaders may deploy electoral campaign strategies that selectively emphasize policies that voters do not consider salient. Even worse, they may take positions that their

¹There is a large literature in political science examining the contents of party manifestos; see the website of the Manifesto Project (<https://manifestoproject.wzb.eu/>). For examples of innovative recent approaches to coding parties’ policy positions based on actual electoral manifestos, see Benoit et al. (2016) and Gabel and Huber (2000).

most ardent supporters oppose. After observing election outcomes, leaders may claim that winning a majority of votes or seats grants them a mandate to implement all of their policy proposals (e.g., Kramer, 1977; Conley, 2001; Fowler and Smirnov, 2007), and this may even be considered a “responsible” course of action (Manin, Przeworski and Stokes, 1999). Indeed, comparative research indicates a remarkable degree of consistency between what parties promise in pre-election manifestos and what parties do once in office (e.g., Klingemann, Hofferbert and Budge, 1994; Royed, 1996; Thomson, 2001). However, voters may actually have considerably different preferences on some policy issues within a manifesto.

A fundamental challenge for both political scientists and policymakers is, therefore, to understand and identify the multidimensional policy preferences behind vote choices in representative democracies. In this study, we propose the use of conjoint analysis as an effective strategy to address this challenge. Conjoint analysis is a method that has been used in marketing research for many years, but has only recently been refined and adopted for use in political science (Hainmueller, Hopkins and Yamamoto, 2014). By juxtaposing sets of multiple issue positions as hypothetical party manifestos (policy bundles) and asking respondents to choose the most preferred hypothetical party, conjoint analysis allows researchers to identify the effects of specific policy positions on voters’ overall assessment of manifestos, the degree of heterogeneity in preferences across respondents, and the popularity ranking of manifestos. Crucially, this approach improves upon standard measurement methods for policy preferences used in public opinion surveys in important ways, which we will describe. We illustrate our approach with a field experiment conducted during Japan’s 2014 House of Representatives election. To the best of our knowledge, this is the first study to use conjoint analysis to investigate voters’ holistic evaluation of parties’ policy manifestos as presented to voters in the context of an actual election.

Our goal is to highlight the utility of conjoint analysis for analyzing voters’ multidimensional policy preferences in an electoral context. We argue that conjoint analysis is a particularly effective approach for this purpose. At the same time, we stress that we are *not*

proposing that conjoint analysis should be used for the purpose of predicting overall election outcomes. Electoral outcomes in representative democracies are not only an incomplete reflection of voters' multidimensional policy preferences, but are also affected by non-policy factors such as candidates' personal attributes and the electoral institutions in use. The goal of our proposed approach is to isolate the policy preferences that cannot be identified by analyzing actual election outcomes or directly measuring vote intentions.

In what follows, we first introduce contextual information of our empirical case. In Section 3, we point out the limitations of standard survey questions and propose our alternative research design which addresses them, followed by a discussion of our statistical methodology. Section 4 reports the results of our analysis. The ruling coalition led by Prime Minister Shinzo Abe of the Liberal Democratic Party (LDP) framed the election around his economic policies (collectively dubbed "Abenomics"). After winning a resounding majority, Abe claimed a mandate to continue his economic agenda. However, our results reveal that voters' support for Abenomics was no higher than for the policy positions of other parties contesting the election. Indeed, our analysis suggests that the LDP's policy manifesto was actually among the least preferred bundles of policies offered to voters. In Section 5, we conduct several validity checks on our findings. The concluding section summarizes what we can learn from conjoint analysis about the 2014 election, and discusses how the use of this method can contribute to future studies of electoral politics and representation.

2 Japan's 2014 House of Representatives Election

Our empirical case comes from the Japanese House of Representatives election held on December 14, 2014. In mid-November, 2014, Prime Minister Abe suddenly signaled his intention to dissolve the House of Representatives. Many voters and commentators questioned why an election was necessary, as the LDP and its coalition partner, Komeito, controlled a two-thirds majority in the chamber, and still had two more years left in the legislative

term before another election would be constitutionally required (Pekkanen, Reed and Smith, 2015). Abe publicly justified the snap election as necessary to seek voter approval for his government's decision to delay a controversial increase in the consumption tax (VAT), scheduled to jump from 8% to 10% in 2015. More generally, the election was touted as a public referendum on the set of Abe's economic policies known as "Abenomics," which include the so-called "three arrows" of aggressive monetary policy, flexible fiscal policy, and structural reforms to Japan's system of regulation to encourage private investment and growth. The first two arrows were launched in 2013 and met with some success (Ito, 2013). By the autumn of 2014, however, the third arrow still remained in Abe's quiver, up against considerable opposition from the LDP's traditional support base in the over-protected agriculture and health industries (Noble, 2015).

Voters turned out in record-low numbers (53% of eligible voters) and the LDP-led ruling coalition again won a crushing victory over a fragmented opposition, with over 68% of the seats. The mixed-member electoral system combines 295 single-member districts (SMDs) allocated by plurality rule and 180 seats separately allocated to parties using closed-list proportional representation (PR) in eleven regional districts. The LDP won 76% of the SMD seats and 38% of the PR seats, while Komeito won 3% of the SMD seats and 14% of the PR seats. The Democratic Party of Japan (DPJ), the largest of the opposition parties and the governing party from 2009–2012, won just 13% of SMD seats and 19% of PR seats. After the election, Prime Minister Abe publicly stated that the LDP's victory represented an approval of him and his economic policy agenda by the electorate. The media similarly portrayed the LDP's victory as a voter endorsement of Abenomics (e.g., *Yomiuri Shimbun*, December 15, 2014).

But was the election truly a referendum on Abenomics? Each of the main parties that contested the election took positions on Abe's economic policies, to be sure, but parties also presented voters with different positions on a range of other issues, including nuclear energy, constitutional revision, and whether or not the size of the National Assembly (Diet)

should be reduced. How can we know whether the LDP's victory reflects voters' preferences for Abenomics or any of the other policy positions taken in the election? Because parties represent multiple policy positions, voters' evaluations of those parties can be quite complex. For example, since the triple disaster of the earthquake, tsunami, and nuclear crisis on March 11, 2011, some voters may be opposed to restarting nuclear power reactors (contrary to the LDP's position), but generally in favor of Abenomics. Other voters may care more about the Abe Cabinet's recent decision to reinterpret Japan's pacifist constitution to allow for collective self-defense. How did they evaluate parties with conflicting positions on the issues they cared about?

3 Experimental Design and Statistical Methodology

The case of Japan's 2014 election exemplifies the fundamental problem in analyzing voter preferences in representative democracies, where the process of preference aggregation by parties makes it difficult to understand voters' multidimensional preferences from the simple observation of election outcomes. In this section, we argue that conjoint analysis provides an effective means to address this problem. We first discuss why standard survey questions may have important limitations as a tool for understanding the multidimensional preferences underlying vote choice. We then describe the survey experiment we conducted in the days leading up to the 2014 election. Finally, we explain the statistical methodology we employ in the analysis of our survey data.

3.1 Limitations of Standard Survey Questions

In addition to analyzing actual election outcomes, researchers commonly rely on surveys administered before and/or after an election to investigate voters' preferences on specific policy issues. In such surveys, respondents are typically asked whether or how much they prefer a position on a particular policy issue over other positions, as well as how they rank

policy issues in order of priority. For example, just after the 2014 election, a Nikkei Research poll asked Japanese voters a number of simple questions about policies, such as “Do you support the Abe administration’s decision to delay the increase of the consumption tax to 10% until 2017?” and gave them four options: “Yes,” “No,” “Can’t say either way,” and “Don’t know.” Later in the survey, respondents were asked “Which policy issue do you think the Abe administration should prioritize?” and allowed respondents to choose as many options as they liked from ten issues ranging from economic policy to foreign relations.² Similar survey questions are common in the American National Election Survey (ANES) and the Comparative Study of Electoral Systems (CSES) surveys, all aimed at understanding voters’ policy preferences and vote choice.

Despite their ubiquitous use in pre- and post-election surveys, such standard survey questions are limited in important ways as a means to analyze the multidimensional policy preferences underlying voters’ decisions. First, standard questions ask for a respondent’s preferences and priority ranking without putting them in the context of a vote choice. A commonly shared concern among survey researchers is that answers to survey questions can fluctuate considerably depending on how the questions are asked (e.g., Kahneman, Slovic and Tversky, 1982). In particular, asking about a respondent’s opinions without putting them in a relevant context can drastically change the way he or she mentally processes the questions and provides an answer. Indeed, many researchers, particularly in the field of economics, are generally skeptical of survey-based research on preferences because of various framing effects (e.g., Bertrand and Mullainathan, 2001). In this regard, standard survey questions are likely to suffer from distortion because they ask about issue preferences and preference rankings independently of voting intentions.³

²*Chousa Kekka 2014-12*: Nikkei Research poll conducted December 24, 2014. English translation by the authors.

³It is important to note that even conjoint analysis does not fully address the potential problem of using stated preferences, as opposed to revealed preferences, to study political behavior. However, recent empiri-

Second, standard questions artificially separate respondents’ implicit utility calculus behind their vote choices into two components: priority ranking of issues and preference ordering between positions on each issue. In actual voting decisions, these two steps are inseparably connected in voters’ minds. For example, a voter may evaluate a particular policy position differently depending on the other policies with which it is bundled in the proposed party platform. This would be plausible if, for example, a voter valued programmatic coherence of party manifestos. Thus, only after making assumptions about the psychology of preference aggregation does it become possible to infer a voter’s overall utility from any given policy bundle based on questions asked and measured separately in two different types of survey questions. In this sense, standard survey questions have no guarantee of uncovering the underlying multidimensional preferences of respondents and how these preferences are related to their vote choices in an election.

To see this more formally under the standard framework of random utility models (McFadden, 1973), consider a voter who is casting a vote in an election. At stake in this election are three issues, which we label a , b , and c . On these three issues, suppose that there are P_a , P_b , and P_c policy positions proposed by the parties fielding candidates in the election, respectively, and that the voter gains the utility of U_{ij} from position j on issue $i \in \{a, b, c\}$, where $j \in \{1, \dots, P_i\}$. Assume further that the voter aggregates the three component utilities into his or her overall utility from voting for a party according to parameters attached to those components. We refer to these parameters as “weights” and denote them by β_i for $i \in \{a, b, c\}$. Making no further assumptions, we can write the voter’s overall utility from party $k \in \{1, \dots, K\}$ as,

$$U_k = f \left(\beta_a, \beta_b, \beta_c, \sum_{j=1}^{P_a} A_{kj} U_{aj}, \sum_{j=1}^{P_b} B_{kj} U_{bj}, \sum_{j=1}^{P_c} C_{kj} U_{cj}, \gamma, \varepsilon_k \right), \quad (1)$$

cal evidence suggests that stated preferences identified through conjoint analysis can accurately reproduce revealed preferences based on behavioral data, at least in certain contexts involving public decision-making (Hainmueller, Hangartner and Yamamoto, 2015).

where A_{kj} , B_{kj} and C_{kj} are binary variables indicating whether or not party k proposes position j on policies a , b , and c such that $\sum_{j=1}^{P_a} A_{kj} = \sum_{j=1}^{P_b} B_{kj} = \sum_{j=1}^{P_c} C_{kj} = 1$ respectively, γ represents a vector of other parameters for the systematic component of the utility function, and ε_k is a stochastic component. Note that equation (1) entails no assumption about the form of the (random) utility function, other than the fact that it takes the three component utilities (U_{ij}) and weights (β_i) as inputs into the systematic component.

The standard survey design amounts to measuring the components of equation (1) separately in the hope that it might help to identify the overall utility U_k for each voter. For example, a question asking whether respondents support a policy position proposed by the government can be thought of as a measurement on U_{ij} which identifies its value up to its ranking relative to $U_{ij'}$ where $j' \neq j$. A question that asks respondents to rank issues in terms of priority can be regarded as measuring β_i , again up to their relative ordering. However, it is evident from equation (1) that attempts to make inferences about the overall utility U_k based on these partial measurements are futile without additional assumptions about the form of $f(\cdot)$.

For example, suppose that the analyst is willing to assume that the overall utility is a linear combination of the component utilities, and the weights on the component utilities are given by their order of priority. Under this assumption, equation (1) can now be written as,

$$U_k = \alpha + \beta_a \sum_{j=1}^{P_a} A_{kj} U_{aj} + \beta_b \sum_{j=1}^{P_b} B_{kj} U_{bj} + \beta_c \sum_{j=1}^{P_c} C_{kj} U_{cj} + \varepsilon_k, \quad (2)$$

where $\alpha = \gamma$ in equation (1). Note that equation (2) takes the form of the standard linear random utility model. Therefore, once the overall utility is specified, it can then be mapped onto the observed vote choice as follows,

$$\Pr(Y_k = 1 \mid A, B, C) = \Pr(U_k \geq U_{k'} \forall k' \in \{1, \dots, K\} \mid A, B, C) \quad (3)$$

where Y_k is a binary indicator of choosing party k and A , B and C represent vectors of the policy position indicators for all K parties.

Although this approach affords researchers a straightforward way to analyze voter preferences, it is fraught with risks of misleading inferences. First and foremost, there is typically no theoretical justification for the particular functional form assumed in equation (2). For example, voters who value programmatic coherence of parties may evaluate policy positions in light of positions on other policies in the manifesto. Such contingent effects must be included in the utility model as interaction terms, which equation (2) overlooks. Second, even if the researcher somehow manages to specify the utility model correctly, there is no guarantee that the standard types of survey questions can validly measure its components (U_{ij} and β_j), particularly when these questions are asked outside the context of vote choice. Indeed, typical questions only measure component preferences and weights in terms of their relative ranking with one another, so the researcher must make another set of strong assumptions to complete the mapping from the measurements to the overall utility function behind the single vote choice.

3.2 Our Experimental Design

We propose the use of conjoint analysis as a more effective alternative. Conjoint analysis is a survey experiment method in which respondents are presented with multiple hypothetical *profiles* randomly generated by juxtaposing sets of multiple *attributes*, each of which has multiple *levels*, and are then asked to choose the profile they would prefer most (Hainmueller, Hopkins and Yamamoto, 2014). In our experiment, we generated hypothetical *parties* confronting various important *policy issues* in an election, and taking various *positions* on those issues. We then asked respondents to choose the party they would prefer most. In recent years, there has been an increasing use of conjoint analysis in several substantive research areas in political science, such as attitudes toward immigrants and political candidates (e.g., Hainmueller and Hopkins, 2014; Franchino and Zucchini, 2014), Eurozone bailout policies (Bechtel, Hainmueller and Margalit, 2014), and global climate change cooperation (Bechtel and Scheve, 2013). As far as we are aware, we are the first to use conjoint analysis in the

context of an actual ongoing electoral campaign to study voters’ policy preferences.

We designed and implemented our experiment as follows. In the run-up to the start of the campaign on December 2, we pored through each of the five major national daily newspapers (*Asahi*, *Yomiuri*, *Mainichi*, *Sankei*, and *Nikkei*) to determine which issues were being discussed by the media. Upon publication of the official party manifestos, we examined each to decide on the final set of policy issues that we judged to be most prominent in the election. The major issues of the campaign were (1) the consumption tax increase, (2) employment policy, (3) monetary and fiscal policy, (4) economic growth strategy, (5) nuclear energy (specifically, whether or not to restart nuclear power reactors), (6) TPP (Trans-Pacific Partnership) trade agreement membership, (7) the right to collective self-defense, (8) constitutional revision, and (9) National Assembly seat reduction. Issues 2–4 are the most directly related to Prime Minister Abe’s economic policy agenda, Abenomics.

For each of these nine issues, we generated 3–4 distinct positions that succinctly but accurately summarized the actual positions of all nine major parties that contested the election: LDP, Komeito, DPJ, Japan Innovation Party (JIP), Party for Future Generations (PFG), People’s Life Party (PLP), Social Democratic Party (SDP), and Japan Communist Party (JCP). We excluded the policy positions of minor and fringe parties, as well as independents. Some of the parties held the same basic positions on certain issues. For example, the LDP and Komeito campaigned predominantly on a common platform as partners in the coalition government. However, only two parties (SDP and JCP) were coded as having identical positions on all of the nine main issues.⁴ The complete set of policy positions, and their correspondence with the actual party manifestos, is shown in Section 4.

In the survey experiment, we presented each respondent with a table containing two hypothetical party manifestos with positions on the nine policy issues (randomly ordered for each respondent), and asked “*Imagine, hypothetically, that the following two parties were nominating candidates in this general election. Which party would you support? Even if you*

⁴SDP and JCP did have different positions on some minor issues that we did not include in the study.

仮に、次のような公約を掲げた2つの政党が今回の総選挙で候補者を擁立していると想定してください。あなたは、どちらの政党を支持しますか。もし、どちらを支持するかははっきりとは言えない場合でも、どちらか一方、あえていえば支持する方を選んでください。

	政党1	政党2
雇用政策	年功序列を撤廃して 労働市場の流動化を促す	多様な働き方を認め、正規・ 非正規を問わず雇用を拡大
金融財政政策	大胆な金融緩和と機動的な 財政出動によりデフレ脱却	過度の金融緩和や円安、 公共事業のバラマキを是正
TPP	TPPへの参加反対	参加して積極的に自由化を推進
議員定数削減	議員定数削減を実現する	選挙制度調査会の答申を尊重し、 よりよい選挙制度改革に取り組む
集団的自衛権	閣議決定のみに基づく行使の 容認には反対	閣議決定のみに基づく行使の 容認には反対
成長戦略	地方産業・中小企業の 活性化による成長実現	農業・医療など 岩盤規制を打破
憲法改正	現行憲法条文のいかなる変更にも 反対。平和憲法を守る	現行憲法の基本原理を維持した上で 必要な条文を追加
消費増税	期限を決めずに延期	2017年4月に10%にし、 軽減税率を導入
原発再稼働	安全基準に合格すれば 認める	責任ある逃避計画など 厳しい条件で容認

どちらを支持するか

政党1	政党2
-----	-----

Figure 1: Example of conjoint table shown to respondents (in Japanese). See the main text for an English translation of the question text. The column headers say “Party 1” and “Party 2.” English translations for the row labels (issues, randomly ordered) and the table contents (positions, randomly sampled as examples) appear on the left side of Figure 2, along the vertical axis. The text below the table says “Which do you support?” followed by the two party options in grey boxes.

are not entirely sure, please indicate which of the two you would be most inclined to support.”

For each issue, one of the 3–4 policy positions was randomly assigned to each party. Each respondent then registered his or her preference for one of the two hypothetical parties, and this exercise was repeated five times. An example of the type of conjoint table (in Japanese) viewed by respondents in the survey experiment is shown in Figure 1. Each column is a hypothetical party (Party 1 and Party 2); each row lists the parties’ positions on each of the nine policy issues. After the five conjoint exercises, we asked each respondent a number of questions about their social demographic background, political ideology, vote intention in the SMD and PR tiers of the election, and support for the Abe Cabinet.

The advantage of our conjoint analysis over standard survey questions is twofold. First, conjoint analysis jointly measures, based on a response to a single exercise, how much each issue is relevant for voters in choosing a hypothetical party and which position on each issue is more preferable compared to other positions.⁵ As discussed in Section 3.1, an important drawback of the standard approach is that the analyst must make an untestable assumption about how the preferences on different policy dimensions and their relative weights are aggregated into a single utility value leading to a vote choice. Conjoint analysis, on the other hand, makes such potentially arbitrary assumptions unnecessary by using each respondent’s observed response (i.e., the selection of the most preferred policy bundle) as the basis for making inferences about his or her multidimensional policy preferences.

More formally, Hainmueller, Hopkins and Yamamoto (2014) show that the fully-randomized

⁵It is important to note that our conjoint analysis shows which issue is more “relevant” for respondents when making their choice of party, but does not necessarily show which issue is truly “important” in their lives. A certain issue may be important for respondents, but they may regard two or more parties’ positions on that issue as almost indistinguishable, and thus less relevant in choosing a specific party. Parties may also strategically make some important issues (for voters) less salient during a campaign, and try to mobilize voters by highlighting other issues that are less important for voters but effective for emphasizing ideological differences (e.g., Budge and Farlie, 1983; Petrocik, 1996). The media may also play a role in making some issues more prominent than others.

conjoint analysis nonparametrically identifies the average marginal component effect (AMCE) for each of the attribute levels on the probability of choosing a profile. What this means is that conjoint analysis allows researchers to estimate how much a policy position, on average, increases or decreases a respondent’s utility for choosing a bundle containing that policy position.⁶ In the random utility framework we introduced above, the AMCE for position 1 as opposed to position 2 on issue A is equal to the following quantity,

$$\Pr(Y_k = 1 \mid A_{1k} = 1) - \Pr(Y_k = 1 \mid A_{2k} = 1), \quad (4)$$

or equivalently,

$$\Pr(U_k \geq U_{k'} \forall k' \in \{1, \dots, K\} \mid A_{1k} = 1) - \Pr(U_k \geq U_{k'} \forall k' \in \{1, \dots, K\} \mid A_{2k} = 1). \quad (5)$$

Thus, conjoint analysis allows us to identify the effect of a policy position on the relative magnitude of the overall utility behind a voter’s party preference. Moreover, since the AMCE is estimated on the same scale for each issue, analysts can easily compare the effect sizes across different issues and make inferences about the relative weights of the issues when a respondent chooses his or her most preferred party. It is crucial to note that this identification is achieved without any assumption about the underlying form of a voter’s utility calculus, in contrast to the approach based on standard survey questions discussed in Section 3.1.

Second, our conjoint approach forces respondents to evaluate policy packages as a whole and to make a choice, just as they would do in the real election. Conjoint-like tables of parties’ actual policy positions are often found in pre-election newspaper coverage of party manifestos, particularly in multiparty parliamentary democracies where voters must compare and choose from among many alternatives. Indeed, some of the newspapers we used to identify the major issues of the campaign had already presented the contrasting positions of the parties on major issues in a simple conjoint-like table for readers. Such pre-election

⁶See Hainmueller, Hopkins and Yamamoto (2014) for the precise definition of the AMCE and more discussion about the meaning of the quantity.

newspaper reports, and even party campaign materials, condense information about policy positions for voters to evaluate. The fact that Japanese voters were probably exposed to similar tables with the same kind of condensed information in the real world enhances the external validity of our analysis.

Our survey was implemented by Research Now, which recruited respondents online during the period between December 4 (two days into the campaign) and the morning of December 14 (when polls opened). Our sample of 1,951 respondents is not a probability sample, but is roughly representative of the general population in terms of observed demographic characteristics of age, gender, prefecture of residence, income, and educational background. In our main empirical analysis, we correct for the observed imbalance between our online sample of respondents and the target population of Japanese voting-age adults in terms of these key demographic covariates using post-stratification weights obtained via entropy balancing (Hainmueller, 2012).⁷ The results are qualitatively identical with or without the use of these weights.

3.3 Statistical Methodology

We employ three different types of techniques for analyzing the data from our conjoint survey experiment in order to extract quantities that speak to our interests.⁸ First, we are interested in identifying the relative salience of the nine policy issues, and respondents' preferences for positions on these issues. In other words, we ask: *Which policy issues are more relevant in respondents' choice over manifestos, and which specific positions on those issues are most positively or negatively evaluated by the respondents?*

To answer these questions, we follow the approach proposed by Hainmueller, Hopkins and Yamamoto (2014) and estimate the AMCE of the policy positions using their regression-based estimator. Hainmueller, Hopkins and Yamamoto (2014) show that, when attribute

⁷Due to nonresponse, the sample size slightly drops to 1,922 in the analyses where we use the weights.

⁸We provide further details of our statistical methodology in Online Appendix A.

levels are randomized independently from one another, the ordinary least squares (OLS) estimates of the coefficients from the linear regression of the choice indicator on the set of dummy variables for the levels of the attributes provide unbiased and consistent estimates of the AMCEs. We follow this procedure and report confidence intervals that are robust to the correlation of preferences within each respondent.

Second, although the AMCEs provide meaningful information about the overall average effects of the policy positions on respondents' preferences over manifestos, they may mask important heterogeneity in these effects across different types of respondents. For example, a policy position may have zero average effect on choice probability either because every respondent is indifferent about the policy or because there are two equally sized groups of respondents who each strongly like and dislike the position. The identification of effect heterogeneity is indeed crucial for understanding the preferences behind vote choices in multiparty elections, since parties' policy programs are often targeted toward particular segments of the electorate. Our next question is therefore: *How does the effect of each policy position vary systematically across groups of respondents?*

To answer this question, we use a hierarchical Bayes approach to model the variability in the AMCEs of the policy positions as functions of respondents' partisanship. Specifically, we use each respondent's vote intention in the PR portion of the upcoming election as the grouping variable. The PR vote intention is arguably the better measure for this purpose than the SMD vote intention, since the former can be regarded as closer to the sincere party preferences of respondents, whereas the latter is more likely to be affected by non-policy factors specific to the individual candidates running in the respondents' home districts, such as personal characteristics. Moreover, all major parties run lists in PR, but not all major parties nominate a candidate in every SMD.⁹ Our quantities of interest, therefore, are the means of respondent-varying AMCEs for the party groups. The variation in these quantities

⁹We show the results estimated using the SMD vote intention in Online Appendix B. The results are qualitatively the same.

will be small and centered around the overall AMCE estimate from the first analysis if there is little heterogeneity in preferences across the party groups about the policy position. On the contrary, the variation will be large if the party groups have more heterogeneous policy preferences.

Third, a crucial step in identifying the multidimensional preferences behind vote choice via conjoint analysis is to map the estimated AMCEs back onto preferences among bundled policy options. It is of particular interest to estimate the popularity of the policy bundles actually proposed by the parties fielding candidates in the concurrent real election and see how they fare against one another. That is, we ask: *What is the relative popularity of the actual policy bundles proposed by parties?*

To answer this question, we first estimate the choice probability of each unique policy bundle by modeling the choice indicator as a function of the policy positions that occurred in the data. An important consideration here is that the specification for our model should be flexible enough to accommodate non-additive effects of policy positions. To this end, we model respondents' party choice as a linear function of all policy position indicators and their pairwise interaction terms.¹⁰ This model, however, requires a careful estimation strategy because it contains a large number of regressors.¹¹ Fitting the model simply via OLS would run the risk of overfitting and poor predictive performance. To address the issue of potential overfitting, we employ the regularization technique called ridge regression (Hoerl and Kennard, 1970).¹² After obtaining the coefficient estimates, we estimate the

¹⁰An even more flexible specification for the utility function would include higher-order interaction terms. However, we find empirically that such models perform worse for our data in terms of their estimated prediction errors. We also test our estimation method against various other techniques and find that it performs as well as any other method, while retaining its flexibility in functional form specifications. Details of our model selection procedure are provided in Online Appendix A.3.

¹¹With our experimental design, the total number of coefficients to be estimated equals 198, a large number even with our relatively large sample of 1,951 respondents each completing five choice tasks.

¹²Compared to OLS, our estimation strategy tends to penalize large coefficient estimates for the interaction

ranking of each unique policy bundle by enumerating all possible combinations of policy positions and calculating their predicted choice probabilities. We also obtain the confidence intervals for these ranks via the block bootstrap at the respondent level, again to allow for the possibility of intra-respondent correlation of preferences. Based on these estimated rankings, we evaluate how the bundles that correspond to the actual party manifestos fare against one another.

4 Results

We now present the key findings from our empirical analysis. As discussed in Section 3.3, we evaluate (1) the average effects of a range of policy positions on respondents’ choice of policy manifesto, (2) the potential underlying heterogeneity across groups, and (3) the relative ranking of actual party manifestos.

4.1 Average Effects of Policy Positions on Overall Preferences

We start with the analysis of average preferences between hypothetical manifestos for the respondents in our sample as a whole. The results are shown in Figure 2. The figure presents the estimated AMCEs for non-LDP policy positions (solid circles) along with cluster-robust 95% confidence intervals (horizontal bars). The estimates are grouped into the nine policy issues we included in our conjoint experiment, as indicated by the labels on the left and also by the colors of the plotted objects. On each policy issue, the LDP’s position is indicated by a solid circle on the zero line without an interval estimate. The correspondence between the

terms and favor estimates closer to zero. Such “shrinkage estimators” are known to perform better in terms of prediction errors than non-regularized estimators such as OLS, when a feature space is high-dimensional and the data matrix is relatively sparse (Hastie, Tibshirani and Friedman, 2009), as is found to be true in our application (see Section 4.3). Thus, our estimation strategy allows for a flexible specification for the underlying utility function behind respondents’ choice of party manifestos, while retaining predictive performance and interpretability (see also footnote 10).

policy positions and actual party manifestos is also indicated by the party acronyms inside of the square brackets in the labels for the levels of the attributes.

For example, on the *Consumption Tax* issue, the actual policy position of the PFG (“Delay until other reforms are made”) is estimated to increase the choice probability for a manifesto by about 2 percentage points, compared to the baseline position of “Delay the tax increase until April 2017 and reduce other tax rates” which was proposed by the LDP and Komeito in the actual election manifestos. The 95% confidence interval for this estimate, however, is $[-0.4, 4.5]$ and contains zero. On the other hand, the other two positions (“Delay the tax increase indefinitely” and “Stop the tax increase and reduce the existing tax”) are found to be even more popular on average, with the AMCEs estimated to be about 4.4 and 4.3 percentage points with 95% confidence intervals of $[1.8, 7.1]$ and $[1.7, 7.0]$, respectively.

A striking pattern in the estimates shown in Figure 2 is that the estimated AMCEs tend to be more often positive than negative on many policy dimensions, indicating that the LDP’s policy positions were not as popular as one might have expected from the actual election results. Indeed, of the eleven policy positions where support is estimated to be significantly different (at the .05 level) from support for the LDP’s position, nine turn out to be more popular than the LDP’s position on the corresponding policy dimension. For example, on the issue of *Nuclear Power*, the position of “Restart[ing] nuclear reactors if proven safe” proposed by both parties in the government coalition is decidedly less popular than the other two positions taken by the opposition parties (“Restart nuclear reactors only under strict safety guidelines” and “Do not restart nuclear reactors”), which have the estimated AMCEs of 4.7 and 3.4 percentage points with 95% confidence intervals of $[2.3, 7.1]$ and $[0.8, 6.0]$, respectively. Moreover, the only two positions that are estimated to be significantly less popular than the LDP’s (“Oppose joining TPP” and “Oppose any reduction of proportional representation seats”) were proposed by opposition parties on the extreme left, rather than the LDP’s main competitors in the election, the DPJ and JIP.

Another remarkable finding is that the three policy issues that constitute Abenomics

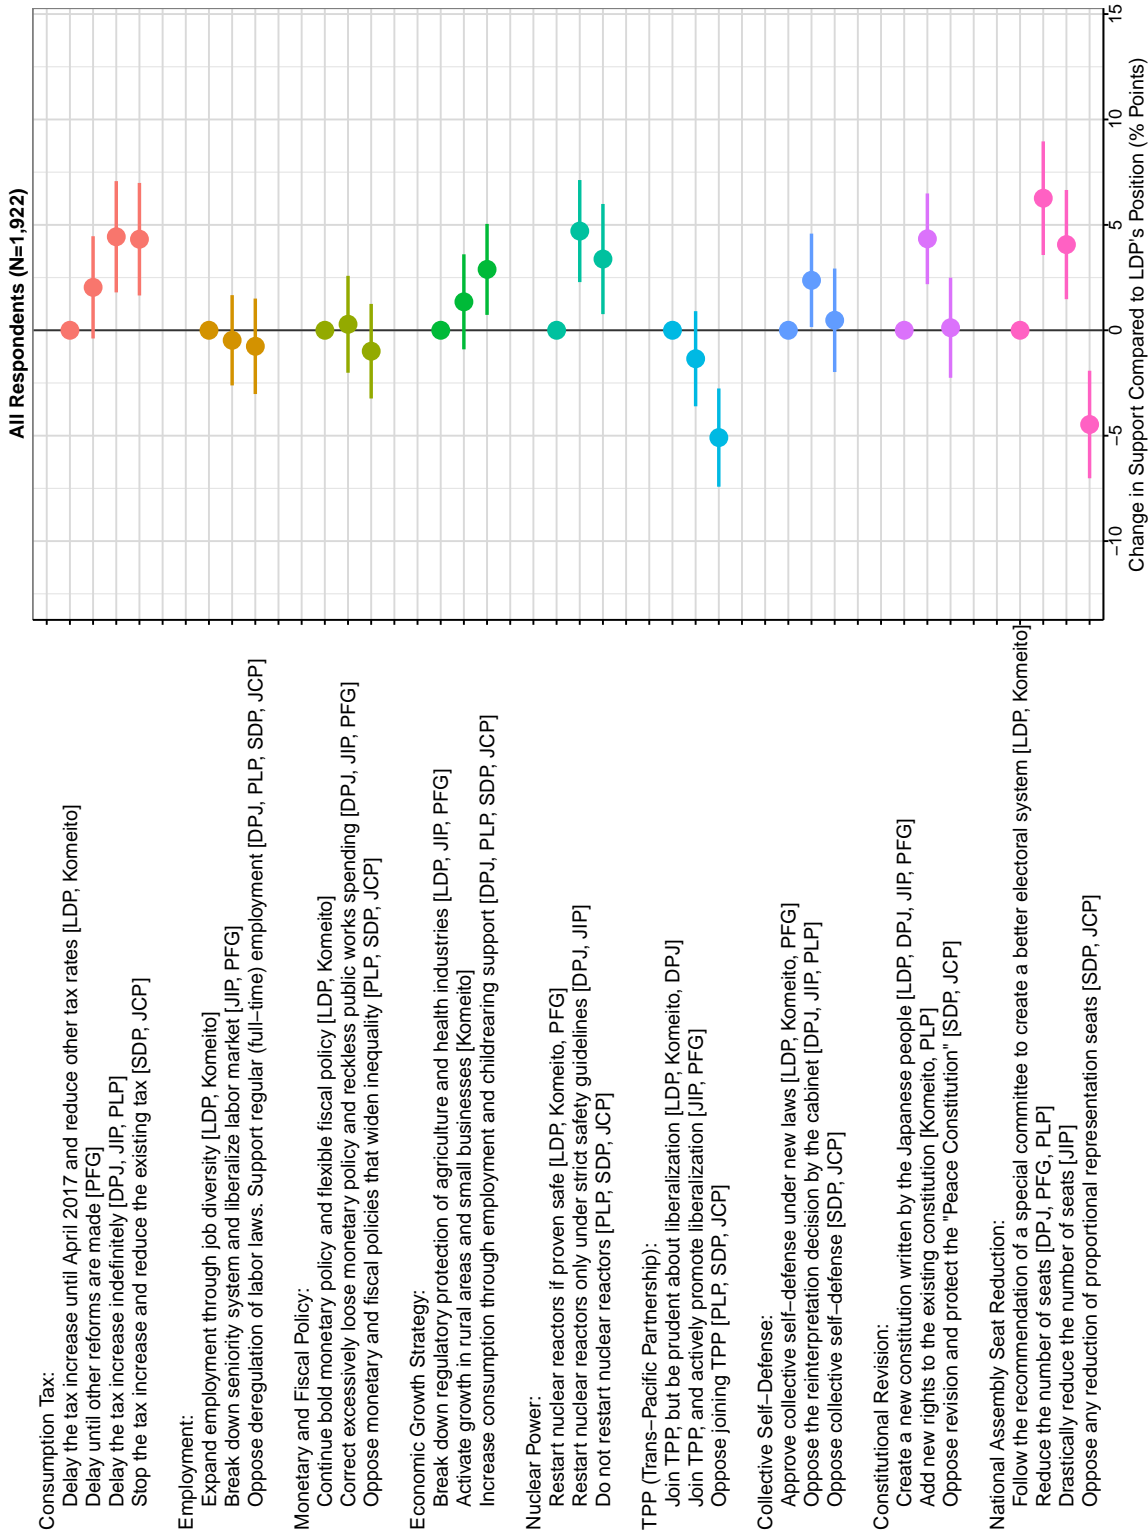


Figure 2: Average Effects of Policy Positions on Respondents' Preference for a Hypothetical Party Manifesto. Each solid circle in the plot represents the estimated average marginal component effect (AMCE) of a policy position on a respondent's probability of choosing a hypothetical manifesto containing that position, compared against a manifesto with the baseline (i.e., LDP's) position on that policy. The horizontal bars represent 95% confidence intervals robust to clustering at the respondent level.

— *Employment, Monetary and Fiscal Policy*, and *Economic Growth Strategy* — are found to have only small effects on the choice probability, if any, while other, non-economic issues such as *Constitutional Revision* and *National Assembly Seat Reduction* are estimated to have large impacts on respondents’ preferences over manifestos. Indeed, all but one of the policy positions that either explicitly support or oppose Abe’s three arrows are estimated to have effects that are statistically indistinguishable from zero. The only exception is the *Economic Growth Strategy* position of “Increase consumption through employment and childrearing support” proposed by opposition parties such as the DPJ, PLP, JCP and SDP, which is estimated to have a small but significantly positive effect (2.9 percentage points, 95% CI [0.7, 5.0]). This finding runs directly counter to the popular belief that the election was a single-issue election that amounted to a referendum on Abenomics.

4.2 Underlying Heterogeneity

Next, we analyze the degree of heterogeneity in respondents’ preferences with respect to the policy positions in our survey. As discussed in Section 3.3, the overall AMCEs reported in Figure 2 represent the effects of the policy positions averaged across all respondents and may, therefore, hide important variation in their preferences. Figure 3 shows that this is indeed the case for some of the policy issues in our experiment. As an illustration, we present the results for two of the twenty policy positions (see Online Appendix B for the full set of results). On the left, we show the estimated AMCE for one of the *Employment* positions (“Break down seniority system and liberalize labor market”) compared against the LDP’s baseline position (“Expand employment through job diversity”) for each of the eleven respondent groups based on their intended PR vote choice, along with the 95% posterior credible intervals. On the right, we show similar estimates for one of the *Constitutional Revision* positions (“Oppose revision and protect the ‘Peace Constitution’” compared against “Create a new constitution written by the Japanese people”). Even though the overall AMCEs for these two policy positions are both very close to zero (see Figure 2), our estimates in Figure 3

Effect Variation by Vote Intention (PR)

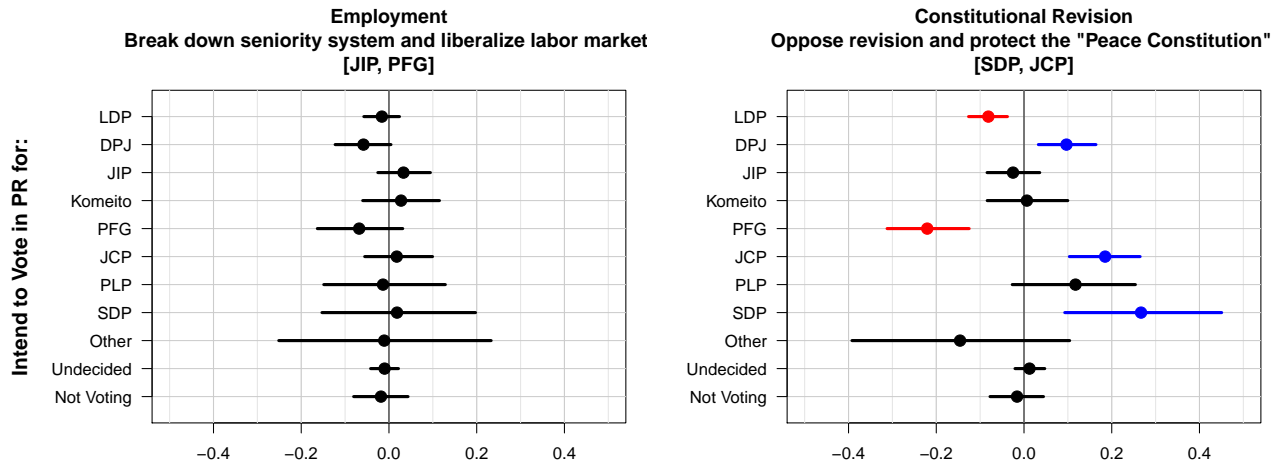


Figure 3: Heterogeneity in Effects of Policy Positions by Intended PR Vote Choice. The figures summarize the posterior distributions of effects of two policy positions (indicated at the top of each plot), modeled as functions of respondents' intended vote choice in the PR portion of the upcoming election (indicated on the vertical axis). The solid circles represent the posterior means of the group-specific AMCEs and the horizontal bars represent 95% credible intervals. The positive (negative) effects that have 95% credible intervals not overlapping with zero are highlighted in blue (red).

reveal a striking contrast in terms of heterogeneity underlying those null average effects.

Specifically, for the employment policy, the estimated group-specific average effects are distributed tightly around the overall AMCE with little variability across party groups. Indeed, all eleven estimated average effects are indistinguishable from zero at the conventional .05 significance level, as indicated by their posterior intervals overlapping with the zero line. This implies that respondents were uniformly indifferent between the employment policy position put forth by the government (LDP and Komeito) and that proposed by the opposition (JIP and PFG) regardless of their own partisanship.

In contrast, the respondents were sharply divided on their preferences about the protectionist position on the Constitution put forth by the SDP and JCP compared to the revisionist position proposed by the LDP, DPJ, JIP and PFG. The estimated average effects for the position are found to be highly variable depending on respondents' partisanship. Two respondent groups supporting conservative parties (LDP and PFG) exhibit large and

statistically significant preferences against the protectionist position, while supporters of the more liberal opposition parties (DPJ, JCP and SDP) indicate equally large, significant preferences in favor of the position. Note that on this issue, the DPJ’s core supporters appear to disagree with the party’s actual position. On the whole, the respondents are estimated to be indifferent between revising and protecting the current constitution *on average*, as indicated by the estimated AMCE of zero in Figure 2. However, our analysis of the underlying heterogeneity reveals the sharp divergence in preferences with respect to these positions across respondents with different vote intentions. Other issues with considerable heterogeneity across groups include nuclear power and collective self-defense (see Online Appendix B). The large heterogeneity across groups for the issues of constitutional revision and collective self-defense are in line with the conventional wisdom that security policy is the most salient cleavage in Japanese politics (e.g., Otake, 1999).

4.3 Ranking of Party Manifestos

The final crucial step in analyzing respondents’ multidimensional policy preferences based on a conjoint experiment is to estimate the relative popularity of the policy bundles as a whole against one another. In the actual election, the LDP received 33% of the vote in the PR tier of the electoral system, which, as we discussed earlier, is generally viewed as a more sincere expression of voter preferences. The other major parties’ PR vote shares were as follows: DPJ (18%), JIP (16%), Komeito (14%), JCP (11%), PFG (3%), SDP (2%), PLP (2%).¹³ How do these actual vote shares compare to the estimated preferences for the parties’ manifestos?¹⁴

Figure 4 shows the results of this analysis for our survey respondents. In each plot, we

¹³These percentages do not sum to 100 due to the exclusion of very minor parties.

¹⁴Importantly, the actual PR vote shares closely match the distribution of the PR vote intention variable among the respondents of our survey (see Figures C.1 and C.2 in Appendix C). This also bolsters our confidence about the representativeness of our study sample, at least for the purpose of our analysis.

Ranks of Actual Party Manifestos

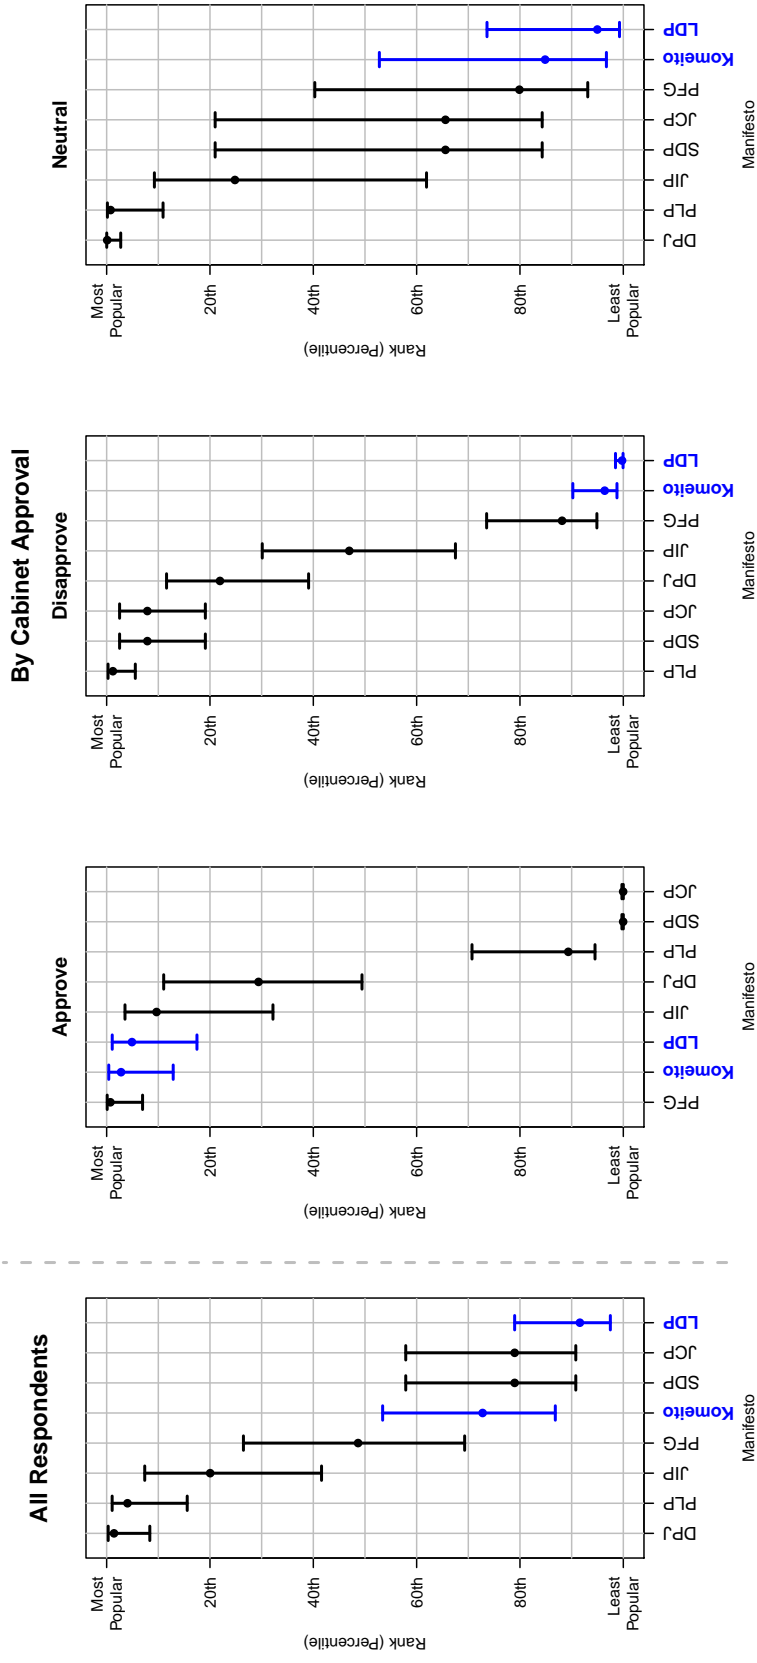


Figure 4: Ranking of Actual Manifestos by Popularity. In each plot, the solid circles represent the estimated ranks of the policy bundles (in percentiles) that correspond to the actual party manifestos in the upcoming election, as indicated on the horizontal axis, among all of the 34,992 possible combinations of the policy positions. The vertical bars represent the 95% confidence intervals based on the block bootstrap at the respondent level. The parties in the government coalition are indicated by blue, bold labels on the horizontal axis. The leftmost plot shows the estimates for all respondents. The three remaining plots show results for the respondents who approved of, disapproved of, and were neutral about the performance of the Abe Cabinet.

estimate the ranks of the seven actual party manifestos among the $4^2 \times 3^7 = 34,992$ possible policy bundles and show their percentiles along with the 95% block bootstrapped confidence intervals.¹⁵ The leftmost plot presents estimated ranks for the sample of all respondents. The estimates indicate that the LDP’s manifesto was highly unpopular; indeed, it is estimated to be the *least* popular bundle among the seven actual manifestos, with the estimated percentile rank of 91.6 and 95% confidence interval of [79.0, 97.5]. On the other hand, the DPJ’s manifesto is found to be the most popular among the bundles corresponding to the actual party manifestos, with the estimated percentile rank of 1.4 and 95% confidence interval of [0.3, 8.4]. This finding would seem to cast strong doubt on any interpretation of the election outcome as giving a popular mandate to the LDP-Komeito coalition government to carry out their proposed policy platform.

The remaining three plots in Figure 4 show the results for the respondents stratified by their approval of the Abe Cabinet’s performance. This categorization is based on a question we included in the survey after the conjoint exercises, in which respondents were asked “Do you approve or disapprove of the performance of the Abe Cabinet?” Respondents were asked to choose between “Approve,” “Disapprove,” and “Can’t say either way.” The respondents in our survey are roughly equally divided into these three strata. Among the 620 respondents who approve of the Abe Cabinet’s performance, the Komeito and LDP manifestos are estimated to be the second and third most popular, respectively, behind the more right-wing PFG.¹⁶ The percentile ranks of the Komeito and LDP manifestos are estimated to be 2.8 and 4.9, with 95% confidence intervals of [0.4, 12.9] and [1.1, 17.5], respectively. In contrast, the 634 respondents who disapprove of Abe’s performance are found to strongly dislike both

¹⁵Note that the policy positions for SDP and JCP are identical on the issues we included (see also footnote 4), which makes the estimates of their manifesto ranks exactly the same in every plot and leaves us with a total of seven unique policy bundles.

¹⁶This is not surprising because the PFG’s policies are similar to the LDP’s, and also because the PFG was formed by conservative former LDP members. Enthusiastic supporters of Abe, who is often seen as a right-wing prime minister, may have thus responded more positively to the PFG’s manifesto than the LDP’s.

of the government parties' manifestos, with the estimated percentile ranks of 99.8 and 96.4 (with 95% confidence intervals of [98.5, 99.9] and [90.2, 98.8]) for the LDP and Komeito, respectively. Finally, but importantly, the results for the remaining 664 respondents who neither approve nor disapprove of the cabinet's performance most closely replicated the preferences of the entire sample, in that the DPJ's manifesto was preferred the most while the LDP's was disliked the most.

These latter results suggest a possible clue to an interesting puzzle — how the LDP and its coalition partner won the election despite a lack of voter approval for the coalition's policies. We find that those 664 respondents who were neutral about Abe's performance were more likely to be either undecided about their vote choice or not intending to vote in the upcoming election than the rest of the sample (57.8% vs. 35.9%, with the p-value of < 0.000 for a two-sided t-test). There is ample evidence in the political science literature that voters who are ambivalent about candidates or party choices in an election are less likely to turn out than voters with clear partisan preferences (for a review, see Smets and van Ham, 2013). Given the historic low turnout (53%), it is likely that many of these undecided and unlikely voters stayed home. Had turnout been higher, our results suggest that the DPJ and other opposition parties might have won more votes and seats. In other words, the low turnout likely helped the LDP win the election despite the unpopularity of its policies.¹⁷

Ultimately, explaining all of the institutional and contextual factors that contributed to the actual outcome of the election is beyond the scope of this paper, and it is important to emphasize that the goal of our conjoint analysis is not to explain the overall election outcome. Rather, our primary objective is to isolate the multidimensional policy preferences of voters from their actual vote choices, which are also influenced by considerations not directly connected to parties' policy manifestos and therefore not necessarily explained by

¹⁷This interpretation of the role of turnout in the outcome of the 2014 election is also consistent with explanations in the existing literature based on empirical analyses of observed vote and abstention patterns across districts (Scheiner, Smith and Thies, 2015).

policy preferences alone. An advantage of conjoint analysis is that it can help to reveal such discrepancies between voters' policy preferences and the actual election results.

5 Validity Checks

Evidence from standard public opinion surveys can be distorted due to various sources of response bias. Although our study obviates some of these concerns by design, there are still several possible ways through which our results may not reflect the true preferences of Japanese voters. In this section, we conduct a series of validity checks to address such possible vulnerabilities.

5.1 Respondent Fatigue and Satisficing

An important concern about the validity of conjoint experiments is respondent fatigue. Any survey can yield invalid results if respondents answer questions without paying close attention. It may be argued that conjoint experiments are particularly prone to this type of bias due to the difficulty of the tasks involved. To complete the choice tasks, respondents must process information presented in the conjoint tables and synthesize it in their minds before making a decision. Of course, voters in the real world are also tasked with processing information provided by parties, and may rely on simple heuristic shortcuts (e.g., Popkin, 1991; Lupia and McCubbins, 1998). Nevertheless, since more cognitively demanding tasks in survey experiments are more likely to induce fatigue and cause effort-saving strategies among respondents than simple questions (Krosnick, 1999; Mutz, 2011), it is important to check for any sign of satisficing behavior in the data.

We take two approaches to address this concern. First, following the suggestion of Hainmueller, Hopkins and Yamamoto (2014), we test whether policy positions have different effects on preferences depending on where in the conjoint table they are shown to each respondent. If the respondents are cognitively overloaded with the information provided in

the conjoint tasks, they might only pay attention to the part of each task that is easiest to process (e.g., the top of the table). If this occurs, then the policy positions should exhibit “order effects.” We test for this possibility by estimating conditional AMCEs for each policy position given its row position and checking whether they differ from one another via a series of t-tests. Figure C.3 in Appendix C plots the resulting p-values against the quantiles of the uniform distribution on the unit interval, indicating the two distributions align very closely. This means that variation in the AMCEs across row positions is no greater than what might occur by chance. Indeed, only 5.6% of the p-values (9 out of 160) are smaller than the standard threshold of 0.05, almost exactly what is expected when the null hypothesis of “no order effect” is true.

Second, we investigate whether the effects of policy positions vary across the five sequential tasks performed by the respondents. The goal here is to detect whether respondents experience fatigue over the course of completing the tasks. If respondents feel that the conjoint tasks are overly demanding, the quality of their responses may degrade as they start paying less attention to the tasks and mechanically clicking on response options. Figure C.4 in Appendix C presents the results of an analysis that is similar to our order effect analysis, except that we look at the effect variation over task counts. Again, we find no evidence of effect heterogeneity across the conditions: the distribution of the p-values is almost indistinguishable from a uniform distribution, and only 6.3% (5 out of 80) of the p-values are smaller than the 0.05 threshold.¹⁸ Thus, we find no evidence that our results suffer from response bias due to cognitive overload or fatigue among the respondents.

¹⁸We also tested the effect variation across task counts by (1) comparing the AMCEs between the first two tasks and the last two tasks and (2) testing the statistical significance of the interaction between each policy position and the task count as a linear regressor via OLS. Neither indicated evidence of effect variation across tasks.

5.2 Contamination with Party Identification

An additional validity concern specific to the context of our study is that respondents might suspect a connection between the policy bundles presented to them and the actual parties in the election. In other words, a respondent may react to a particular policy position (such as restarting nuclear reactors) not because he or she supports the policy *per se*, but because he or she supports the party (in this case, LDP, Komeito or PFG) with which the policy is associated in the respondent’s mind. Such a possibility poses a challenge to our identification strategy because our goal is to isolate voters’ *policy* preferences from other determinants of vote choice in the actual election, such as feelings of attachment to parties that derive from non-policy sources.

Although this remains a possibility, it is worth noting that some aspects of our study design make it unlikely that respondents will attach specific party labels to the policy bundles presented to them. First, we deliberately emphasized the hypothetical nature of the manifestos on the introductory screen presented to respondents before they saw the five conjoint tables, as well as in a sentence added above each table (see Section 3.2 for the exact question wording). Coupled with our evidence (in Section 5.1) that the respondents likely paid sufficient attention to the questions, we expect that most, if not all, of them clearly understood that the parties in the tables were hypothetical.

Second, due to the full randomization of policy positions, only a minuscule fraction of the bundles exactly corresponded to the actual party manifestos. Indeed, more than 96% of the bundles that were generated for the conjoint tasks were “unrealistic hybrids” in the sense that they contained at least one position from both the conservative governing party (LDP) and the most left-wing opposition parties (JCP, SDP and PLP). Estimating the AMCEs of the policy positions on the subsample of tasks that we consider to be highly unlikely to be associated with real-world parties produces results that are nearly identical to the estimates using the entire sample reported in Figure 2.¹⁹ That is, respondents appear to evaluate policy

¹⁹We define highly unlikely bundles as those that contain at least one position from both the LDP and

bundles in the same manner whether or not the bundles plausibly correspond to actual party manifestos. This suggests that our experiment captures voters' policy preferences that are not contaminated by extraneous considerations such as non-policy party identification.

6 Conclusion

The aggregation of multiple policies into alternative and competing programs for voters to evaluate in elections is a key function of political parties in representative democracies. However, because of this process of preference aggregation into manifestos, election outcomes alone do not provide complete information on the underlying distribution of voter's multidimensional preferences, or how preferences for specific policies influence voters' choices between alternative parties. Moreover, standard survey methods employed in most pre- and post-election surveys are inadequate at approximating the type of decision-making process that voters actually undertake in evaluating alternative party choices.

We have argued that conjoint analysis is a useful tool for revealing voters' multidimensional policy preferences, particularly when coupled with information from the actual party manifestos presented to voters in elections. We have illustrated this approach with a field experiment conducted during the 2014 House of Representatives election in Japan. In sum, the results of our three sets of analysis suggest important implications for how the LDP's victory in the election should be interpreted. First, our findings indicate that for many of the nine major policy issues in the election, the LDP's policy positions were less popular than those of the opposition. Second, the economic policies of Prime Minister Abe ("Abenomics") were less relevant than other non-economic policies for respondents in choosing their pre-

the leftist opposition parties (JCP, SDP and PLP) on the four policy areas we found to be divisive in our main analysis (*Consumption Tax*, *Nuclear Power*, *Collective Self-Defense*, and *Constitutional Revision*). The subsample constitutes approximately 74% of all possible bundles. We present the full results of the analysis in Figure C.5 in Appendix C.

ferred policy bundles. This is surprising given that Abe and the media framed the election as a referendum on Abenomics. Our analysis of the degree of heterogeneity in preferences confirmed this result; namely, respondents' preferences with respect to economic policies were similar regardless of which party they intended to support in the actual election. In contrast, some of the non-economic issues, such as constitutional revision and the approval of collective self-defense, showed substantial heterogeneity among respondents depending on their vote intention. Finally, the LDP's policy bundle was one of the least popular bundles. The policy bundle proposed by the largest opposition party, the DPJ, was actually the most popular.

By identifying voters' multidimensional policy preferences via a conjoint experiment, our study reveals an important shortcoming in the existing ways of analyzing elections in representative democracies. Indeed, there is substantial evidence that many scholars and the media misinterpreted Japanese voters' policy preferences in the election. Our analysis reveals that voters neither thought that Abenomics was an important issue at stake, nor preferred the LDP's other policy proposals. Instead, our results suggest that the LDP's victory was arguably due to various factors *other than the party's policy proposals*. Based on these results, it may be problematic that after the election, the LDP continued to pursue the same (unpopular) policies on which it campaigned. An advantage of using conjoint analysis in the context of an actual election campaign is that it can help to better illuminate this type of discrepancy between voters' multidimensional policy preferences and the parties that ultimately get their votes. This discrepancy between the distribution of voters' multidimensional preferences for policies and the distribution of votes (and then seats) among parties is an important, yet under-investigated, problem — not only in Japan, but in all representative democracies.

We believe that conjoint analysis based on actual party manifestos is a better approach to understanding the multidimensional policy preferences of voters in representative democracies, and argue that our approach should be built into future election surveys. To estimate

the gap between each voter's multidimensional policy preferences and his or her actual vote choice, it would be ideal to design a two-wave panel study, in which each voter is asked to do conjoint exercises during the campaign period and then to report his or her vote choice. How this gap varies under diverse electoral settings (such as pure SMD and PR systems) in other democracies is another important question for future researchers to investigate.

References

- Bechtel, Michael M., Jens Hainmueller and Yotam M. Margalit. 2014. “Exploring the Multidimensionality of Preferences for International Redistribution: The Case of the Eurozone Bailouts.” Available at SSRN: <http://ssrn.com/abstract=2163594> or <http://dx.doi.org/10.2139/ssrn.2163594>.
- Bechtel, Michael M. and Kenneth F. Scheve. 2013. “Mass support for global climate agreements depends on institutional design.” *Proceedings of the National Academy of Sciences* 110(34):13763–13768.
- Benoit, Kenneth, Drew Conway, Benjamin E. Lauderdale, Michael Laver and Slava Mikhaylov. 2016. “Crowd-sourced Text Analysis: Reproducible and Agile Production of Political Data.” *American Political Science Review* 110(02):278–295.
- Bertrand, Marianne and Sendhil Mullainathan. 2001. “Do People Mean What They Say? Implications for Subjective Survey Data.” *American Economic Review* 91(2):67–72.
- Budge, Ian and Dennis J. Farlie. 1983. *Explaining and Predicting Elections: Issue Effects and Party Strategies in Twenty-Three Democracies*. London: George Allen and Unwin.
- Conley, Patricia Heidotting. 2001. *Presidential Mandates: How Elections Shape the National Agenda*. Chicago: The University of Chicago Press.
- Dewan, Torun and Kenneth A. Shepsle. 2011. “Political Economy Models of Elections.” *Annual Review of Political Science* 14:311–330.
- Fowler, James H. and Oleg Smirnov. 2007. *Mandates, Parties, and Voters: How Elections Shape the Future*. Philadelphia: Temple University Press.
- Franchino, Fabio and Francesco Zucchini. 2014. “Voting in a Multi-dimensional Space: A Conjoint Analysis Employing Valence and Ideology Attributes of Candidates.” *Polit-*

ical Science Research and Methods FirstView Article / November 2014, pp 1 - 21 DOI: 10.1017/psrm.2014.24, Published online: 28 October 2014.

Gabel, Matthew J. and John D. Huber. 2000. "Putting Parties in Their Place: Inferring Party Left-Right Ideological Positions from Party Manifestos Data." *American Journal of Political Science* 44(1):94.

Hainmueller, Jens. 2012. "Entropy Balancing: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies." *Political Analysis* 20(1):25–46.

Hainmueller, Jens and Daniel J. Hopkins. 2014. "The Hidden American Immigration Consensus: A Conjoint Analysis of Attitudes toward Immigrants." *American Journal of Political Science* 00(0):1–20.

Hainmueller, Jens, Daniel J. Hopkins and Teppei Yamamoto. 2014. "Causal Inference in Conjoint Analysis: Understanding Multi-dimensional Preferences via Stated Preference Experiments." *Political Analysis* 22(1):1–30.

Hainmueller, Jens, Dominik Hangartner and Teppei Yamamoto. 2015. "Validating Vignette and Conjoint Survey Experiments against Real-World Behavior." *Proceedings of the National Academy of Sciences of the United States of America* 112(8):2395–2400.

Hastie, Trevor, Robert Tibshirani and Jerome Friedman. 2009. *The elements of statistical learning: Data Mining, Inference, and Prediction*. 2 ed. Springer.

Hoerl, Arthur E. and Robert W. Kennard. 1970. "Ridge regression: Biased estimation for nonorthogonal problems." *Technometrics* 12(1):55–67.

Ito, Takatoshi. 2013. "Abenomics: Early Success and Prospects." *Japan Spotlight* pp. 4–7.

Kahneman, Daniel, Paul Slovic and Amos Tversky. 1982. *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press.

- Klingemann, Hans-Dieter, Richard I. Hofferbert and Ian Budge. 1994. *Parties, Policies, and Democracy*. Boulder, CO: Westview Press.
- Kramer, Gerald H. 1977. "A Dynamic Model of Political Equilibrium." *Journal of Economic Theory* 16:310–334.
- Krosnick, Jon A. 1999. "Survey Research." *Annual Review of Psychology* 50(1):537–567.
- Lupia, Arthur and Mathew D. McCubbins. 1998. *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* Cambridge: Cambridge University Press.
- Manin, Bernard, Adam Przeworski and Susan C. Stokes. 1999. Elections and Representation. In *Democracy, Accountability, and Representation*, ed. Adam Przeworski, Susan C. Stokes and Bernard Manin. Cambridge University Press pp. 29–54.
- McFadden, Daniel. 1973. Conditional Logit Analysis of Qualitative Choice Behavior. In *Frontiers in Econometrics*, ed. P. Zarembka. New York: Academic Press pp. 105–142.
- Mutz, Diana C. 2011. *Population-based Survey Experiments*. Princeton, NJ: Princeton University Press.
- Noble, Gregory W. 2015. Abenomics in the 2014 Election: Showing the Money (Supply) and Little Else. In *Japan Decides 2014: The Japanese General Election*, ed. Robert Pekkanen, Steven R. Reed and Ethan Scheiner. New York, NY: Palgrave Macmillan pp. 155–169.
- Otake, Hideo. 1999. *Nihon Seiji no Tairitsujiku*. Tokyo: Chuko Shinsho.
- Pekkanen, Robert J., Steven R. Reed and Daniel M. Smith. 2015. Japanese Politics Between the 2012 and 2014 Elections. In *Japan Decides 2014: The Japanese General Election*. New York, NY: Palgrave Macmillan pp. 9–21.
- Petrocik, John R. 1996. "Issue Ownership in Presidential Elections, with a 1980 Case Study." *American Journal of Political Science* 40(3):825–850.

- Popkin, Samuel. 1991. *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns*. Chicago, IL: University of Chicago Press.
- Powell, G. Bingham Jr. 2007. Aggregating and Representing Political Preferences. In *The Oxford Handbook of Comparative Politics*, ed. Charles Boix and Susan C. Stokes. New York: Oxford University Press pp. 653–677.
- Royed, Terry J. 1996. “Testing the Mandate Model in Britain and the United States: Evidence from the Reagan and Thatcher Eras.” *British Journal of Political Science* 26(1):45–80.
- Scheiner, Ethan, Daniel M. Smith and Michael F. Thies. 2015. The 2014 Japanese Election Results: The Opposition Cooperates, but Fails to Inspire. In *Japan Decides 2014: The Japanese General Election*, ed. Robert J. Pekkanen, Steven R. Reed and Ethan Scheiner. New York, NY: Palgrave Macmillan pp. 22–38.
- Smets, Kaat and Carolien van Ham. 2013. “The embarrassment of riches? A meta-analysis of individual-level research on voter turnout.” *Electoral Studies* 32:344–359.
- Thomson, Robert. 2001. “The programme to policy linkage: The fulfilment of election pledges on socio-economic policy in the Netherlands, 1986–1998.” *European Journal of Political Research* 40:171–197.

Online Appendix

A Details of the Statistical Methodology

In this appendix, we describe the details of the statistical methods we employed for the analysis of our conjoint survey data.

A.1 Average Marginal Component Effects

To obtain unbiased and consistent estimates of the AMCEs, we fit the following linear model to our data,

$$Y_{ijk} = \beta_0 + \sum_{l=1}^9 \sum_{d=2}^{D_l} \beta_{ld} X_{ldijk} + \varepsilon_{ijk}, \quad (6)$$

where $Y_{ijk} \in \{0, 1\}$ is the binary choice indicator for manifesto j in task k of respondent i , X_{ldijk} is the dummy variable for the d th position of policy l , β_{ld} is the corresponding coefficient, and ε_{ijk} represents the error term, which is statistically independent of the regressors due to the randomization of the attributes. Note that we index our nine policy issues by $l \in \{1, \dots, 9\}$ and the positions on policy l by $d \in \{1, \dots, D_l\}$, where D_l equals the total number of positions for policy l (e.g., $D_l = 4$ for $l = 1$, consumption tax policy) and $d = 1$ corresponds to the LDP's position, which is taken as the reference category. We then use the OLS estimates of β_{ld} as our estimates of AMCE for the d th position of policy l , with White cluster-corrected standard errors to account for within-respondent correlation of preferences.

A.2 Effect Heterogeneity

For the analysis of heterogeneous effects across groups of respondents, we extend the model in equation (6) by allowing the coefficients to vary across respondents, i.e.,

$$Y_{ijk} = \beta_{0i} + \sum_{l=1}^9 \sum_{d=2}^{D_l} \beta_{ldi} X_{ldijk} + \varepsilon_{ijk}, \quad (7)$$

where ε_{ijk} is now assumed to be an independently, identically and normally distributed random variable with mean zero. We then model the varying coefficients as functions of respondent-level covariates as follows,

$$\beta_i = \gamma W_i + \eta_i, \quad (8)$$

where $\beta_i = [\beta_{0i}, \beta_{12i}, \dots, \beta_{9D_9i}]^\top$, W_i is a vector of covariates for respondent i , γ is a matrix of respondent-level coefficients, and η_i is a vector of respondent-level error terms such that $\eta_i \sim \mathcal{N}(0, \Sigma)$. We use noninformative priors for the unmodeled parameters γ , σ and Σ , such that $\gamma_{mp} \sim \mathcal{N}(0, 10^6)$, $\sigma \sim \text{Unif}(0, 10)$ and $\Sigma \sim \text{IW}(I_M, M + 1)$ where $\gamma = [\gamma_1, \dots, \gamma_p, \dots, \gamma_P]$, $\gamma_p = [\gamma_{1p}, \dots, \gamma_{mp}, \dots, \gamma_{Mp}]^\top$, $M = \sum_{l=1}^9 (D_l - 1) + 1 = 21$, $P = 11$ and I_k denotes the identity matrix of dimension k . Our quantity of interest from the model is $\mu_w \equiv \mathbb{E}[\beta_i | W_i = w]$.

The model is fitted via a Gibbs sampler implemented on JAGS 3.4.0. We run four chains in parallel in order to assess convergence, with the parameters initiated at dispersed locations on the parameter space. We use a distinct pseudo-random number generator on each chain to avoid potential problems with the sampler. After 40,000 iterations on each chain, of which the first 20,000 are discarded as burn-in draws, the chains show adequate evidence suggesting convergence to the true posterior: The Gelman-Rubin (GR) diagnostic scores for the parameters of interest are no greater than 1.01, with the exception of one parameter which has the GR score of 1.02 as expected by chance given the large number of parameters. We subsequently thin the chains by retaining every tenth draw, leaving the total of 8,000 simulation draws for our analysis.

A.3 Ranking of Profiles

To obtain the predicted ranking of all possible hypothetical manifestos, we use the following linear model:

$$Y_{ijk} = \beta_0 + \sum_{l=1}^9 \sum_{d=2}^{D_l} \beta_{ld} X_{ldijk} + \sum_{l=2}^9 \sum_{l' < l} \sum_{d=2}^{D_l} \sum_{d'=2}^{D_m} \gamma_{ll'dd'} X_{ldijk} X_{l'd'ijk} + \varepsilon_{ijk}, \quad (9)$$

where $\gamma_{l'l'dd'}$ denotes an unknown coefficient for the interaction between the d th position of policy l and the d' th position of policy l' . We estimate the coefficients $[\beta_0, \beta_{12}, \dots, \beta_{9D_9}]$ with L_2 penalty to avoid overfitting. That is, our estimates minimize the following sum of squared residuals with a shrinkage penalty on the interaction terms:

$$\begin{aligned} & \sum_{i=1}^n \sum_{j=1}^2 \sum_{k=1}^5 \left\{ Y_{ijk} - \hat{\beta}_0 - \sum_{l=1}^9 \sum_{d=2}^{D_l} \hat{\beta}_{ld} X_{ldijk} - \sum_{l=2}^9 \sum_{l' < l} \sum_{d=2}^{D_l} \sum_{d'=2}^{D_m} \hat{\gamma}_{l'l'dd'} X_{ldijk} X_{l'd'ijk} \right\}^2 \\ & + \lambda \sum_{l=2}^9 \sum_{l' < l} \sum_{d=2}^{D_l} \sum_{d'=2}^{D_m} \hat{\gamma}_{l'l'dd'}^2, \end{aligned} \quad (10)$$

where λ is the tuning parameter that is chosen to minimize the mean squared prediction error obtained via ten-fold cross-validation, following the standard practice (Hastie, Tibshirani and Friedman, 2009).

The ridge penalty was selected based on a systematic comparison of empirical performance among 17 alternative model specifications and estimation techniques. These methods include: (1) OLS with no interaction term, (2) OLS with second-order interactions, (3) OLS with third-order interactions, (4) ridge regression with L_2 penalty on all model coefficients with no interaction, (5) ridge regression with second-order interactions, (6) ridge regression with third-order interactions, (7) ridge regression with L_2 penalty only on interaction terms with second-order interactions and (8) with third-order interactions, (9)–(13) the same set of specifications as (4)–(8) using LASSO (L_1) penalty, (14) Bayesian model averaging (BMA) over all possible predictor combinations with no interaction, (15) BMA with second-order interactions, (16) BMA with second-order interactions with zero prior on models including interaction terms but not their component main effects, and (17) BMA with second-order interactions with zero prior on models not including either of the main effects. We evaluated the performance of these 17 methods with their estimated mean squared prediction errors obtained via ten-fold cross-validation. The results indicate that our method (7) performs at least as well as any other method in the comparison set based on the chosen metric. It is worth noting that methods that utilize interaction effects with no regularization (2 and 3) are found to perform substantially worse than any of the regularized methods.

B Heterogeneity in Policy Preferences: Full Results

Figures B.1, B.2, and B.3 report the full set of results showing the heterogeneity in preferences for all twenty policy positions. Figures B.4, B.5, and B.6 show the results estimated using SMD vote intention rather than PR vote intention.

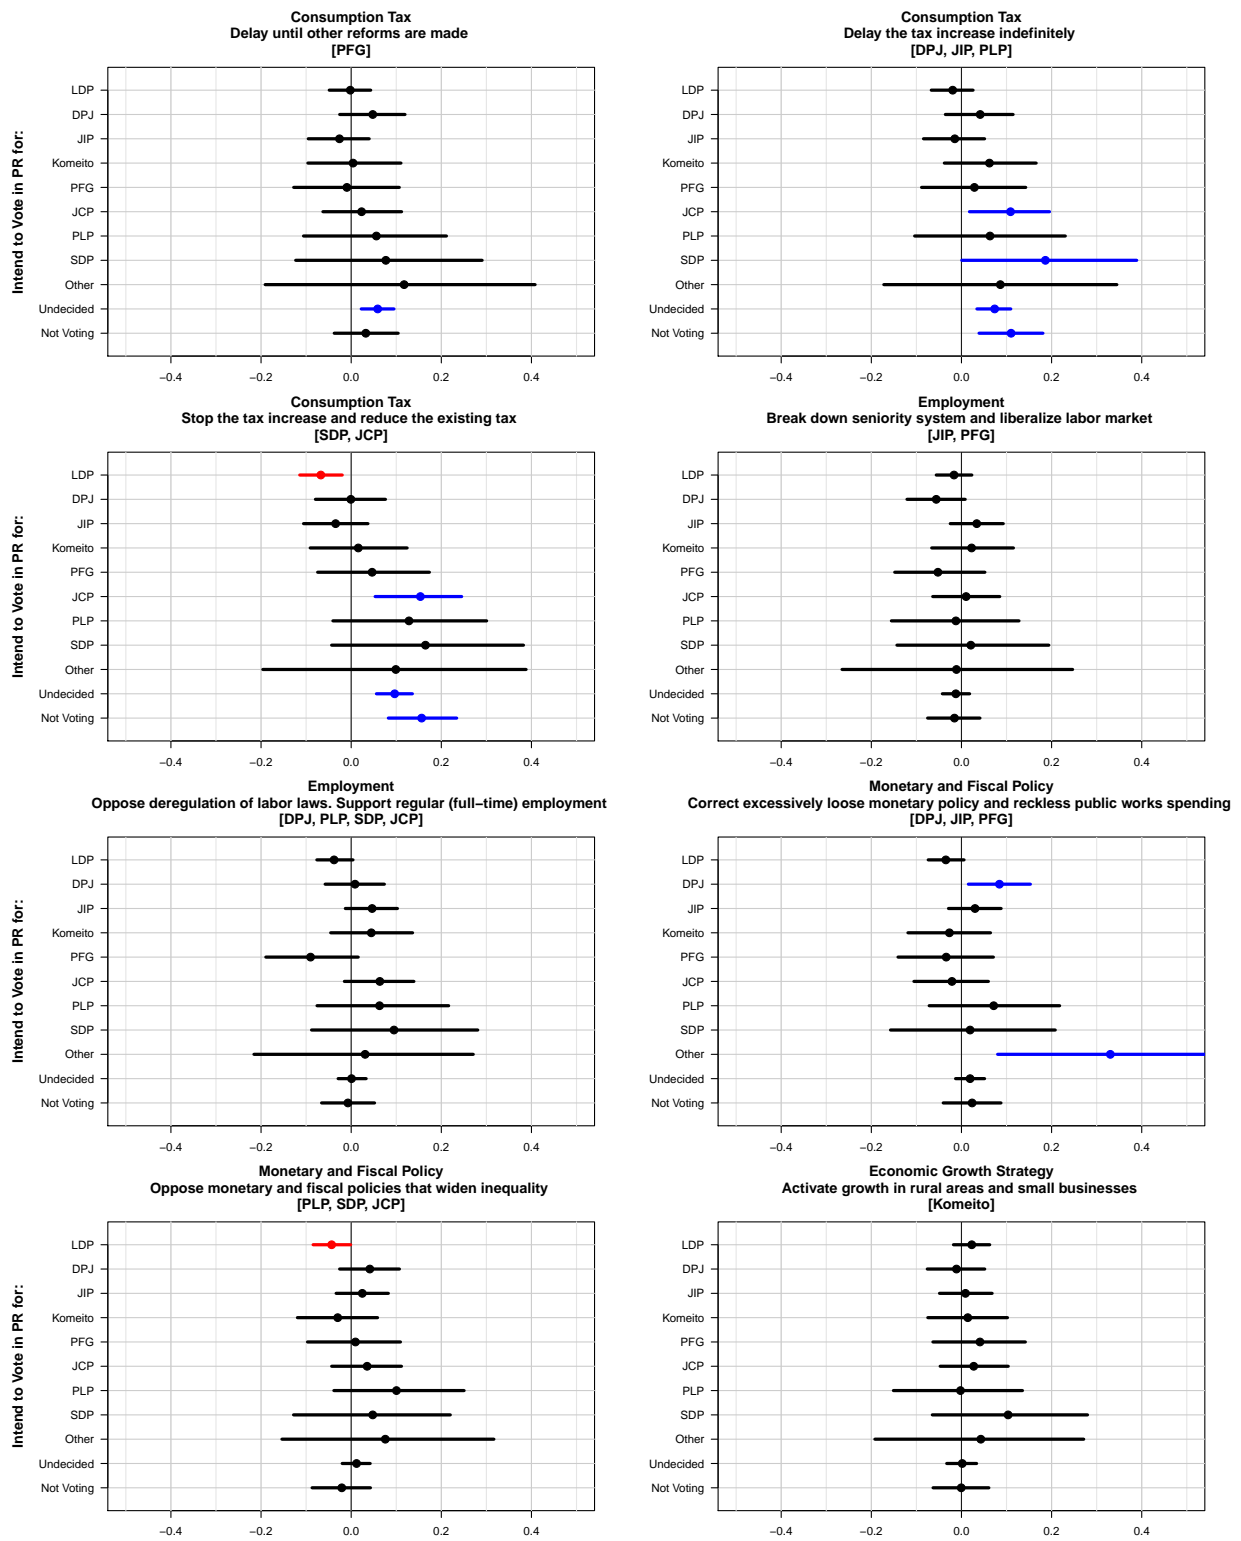


Figure B.1: Effect Heterogeneity by PR Vote Intention: Full Results (Part 1 of 3).

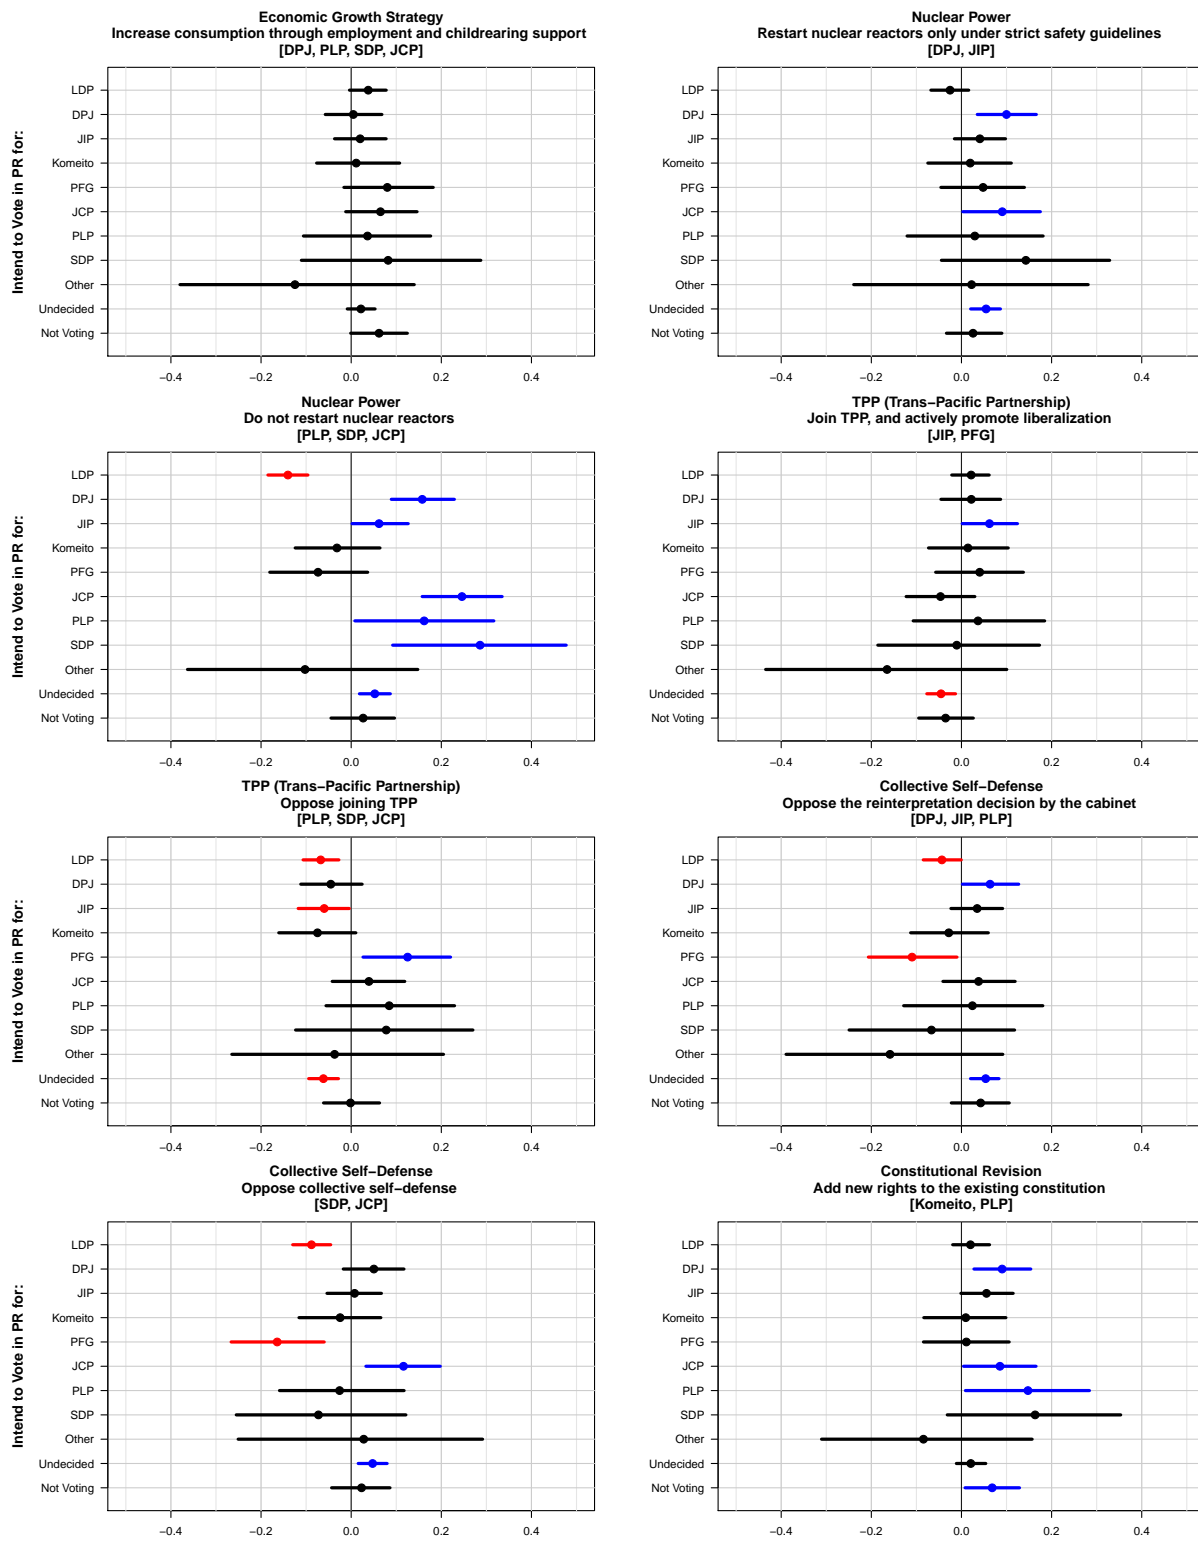


Figure B.2: Effect Heterogeneity by PR Vote Intention: Full Results (Part 2 of 3).

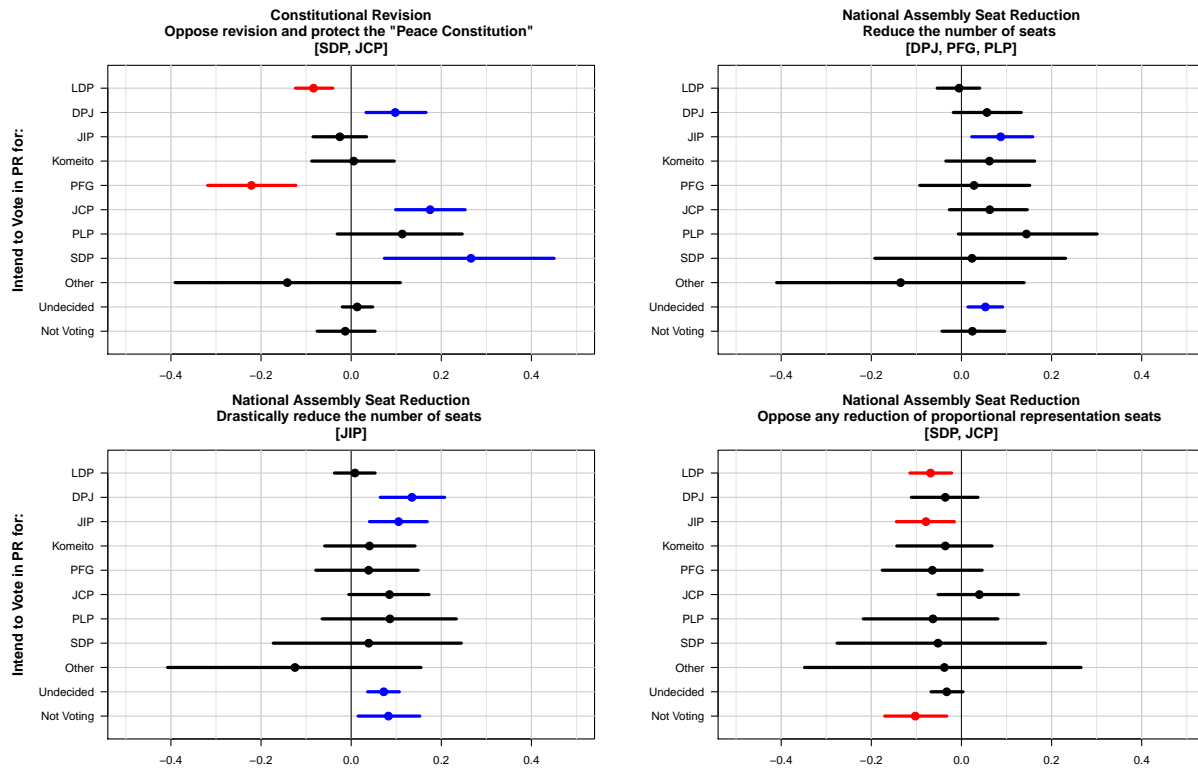


Figure B.3: Effect Heterogeneity by PR Vote Intention: Full Results (Part 3 of 3).

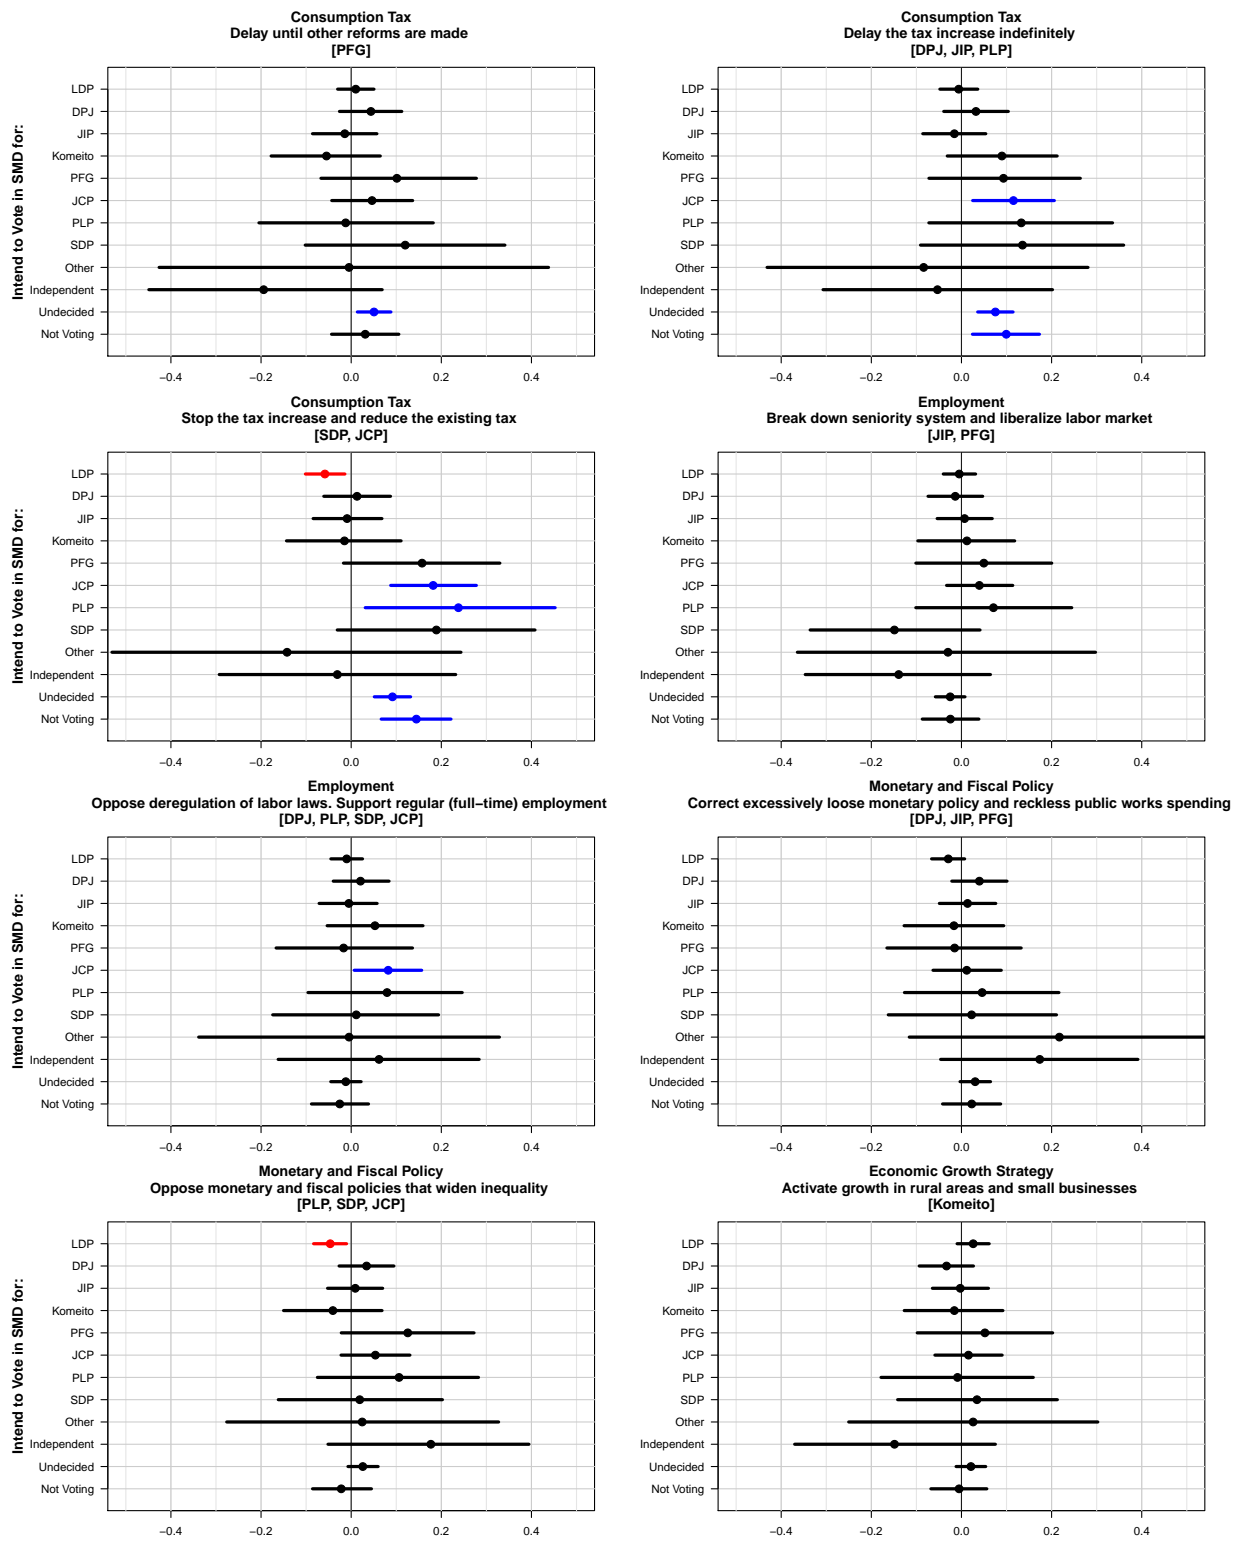


Figure B.4: Effect Heterogeneity by SMD Vote Intention: Full Results (Part 1 of 3).

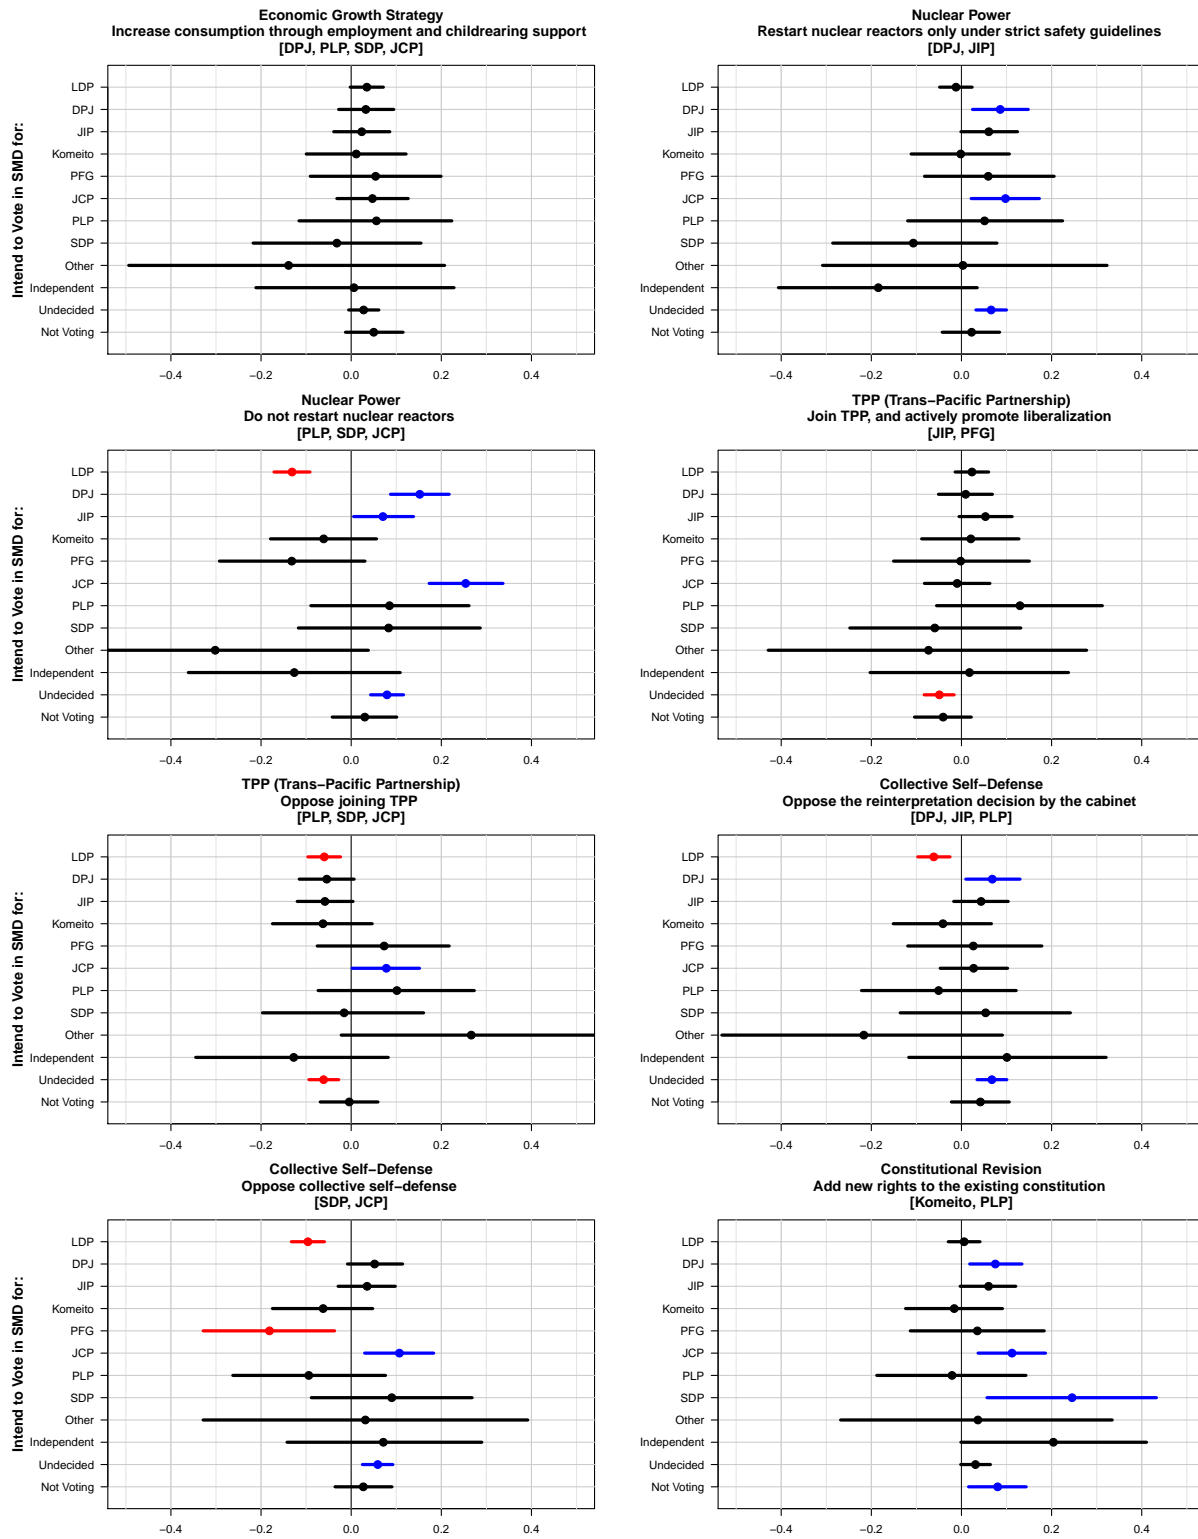


Figure B.5: Effect Heterogeneity by SMD Vote Intention: Full Results (Part 2 of 3).

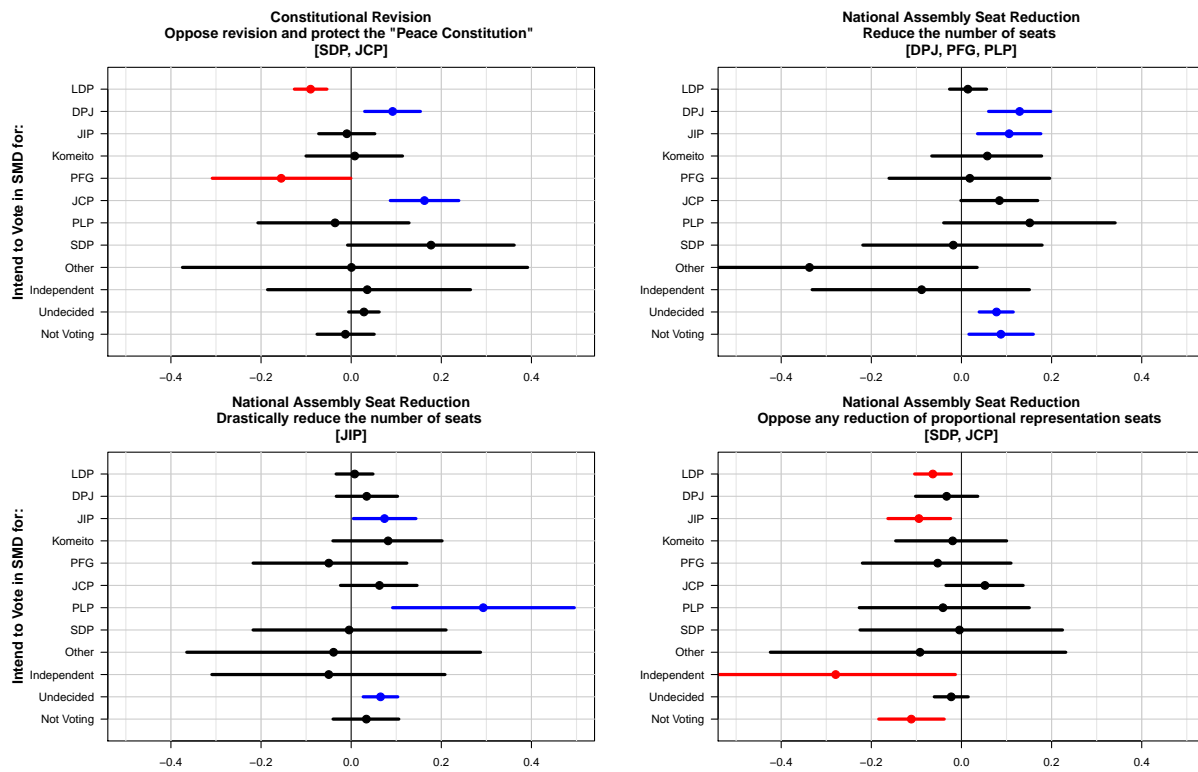


Figure B.6: Effect Heterogeneity by SMD Vote Intention: Full Results (Part 3 of 3).

C Results of Validity Checks

Figure C.1 shows the party vote shares in the PR tier in the 2014 Japanese House of Representatives election on the left and the distribution of the intended vote variable for the PR tier among our survey respondents, excluding those who chose either “Undecided” or “Not Intending to Vote.” The shares match closely despite the fact that our sample is not a probability sample from the population.

Figure C.2 includes all of the respondents and compares the distribution to the population of eligible voters as a whole, with the difference between the official recorded votes and the population size coded as “Abstain/Invalid.” Again, the distributions match quite closely. Even the proportion of abstained/invalid votes in the voting-age citizen population is similar to the proportion of respondents in the survey population who are undecided or do not intend to vote.

Figures C.3 and C.4 present the results of our validity checks in terms of respondent fatigue and satisficing. See Section 5 and the captions of the figures for a full description of the analyses. On the whole, the results provide no evidence of cognitive overload or fatigue effects among the respondents.

Figure C.5 shows the results of our analysis of policy bundles that respondents are highly unlikely to associate with actual parties in the election. See Footnote 19 for the exact definition of those bundles. The results are nearly identical to Figure 2, indicating that respondents choose bundles based on policy positions themselves instead of their guesses about which actual parties may correspond to the hypothetical bundles presented.

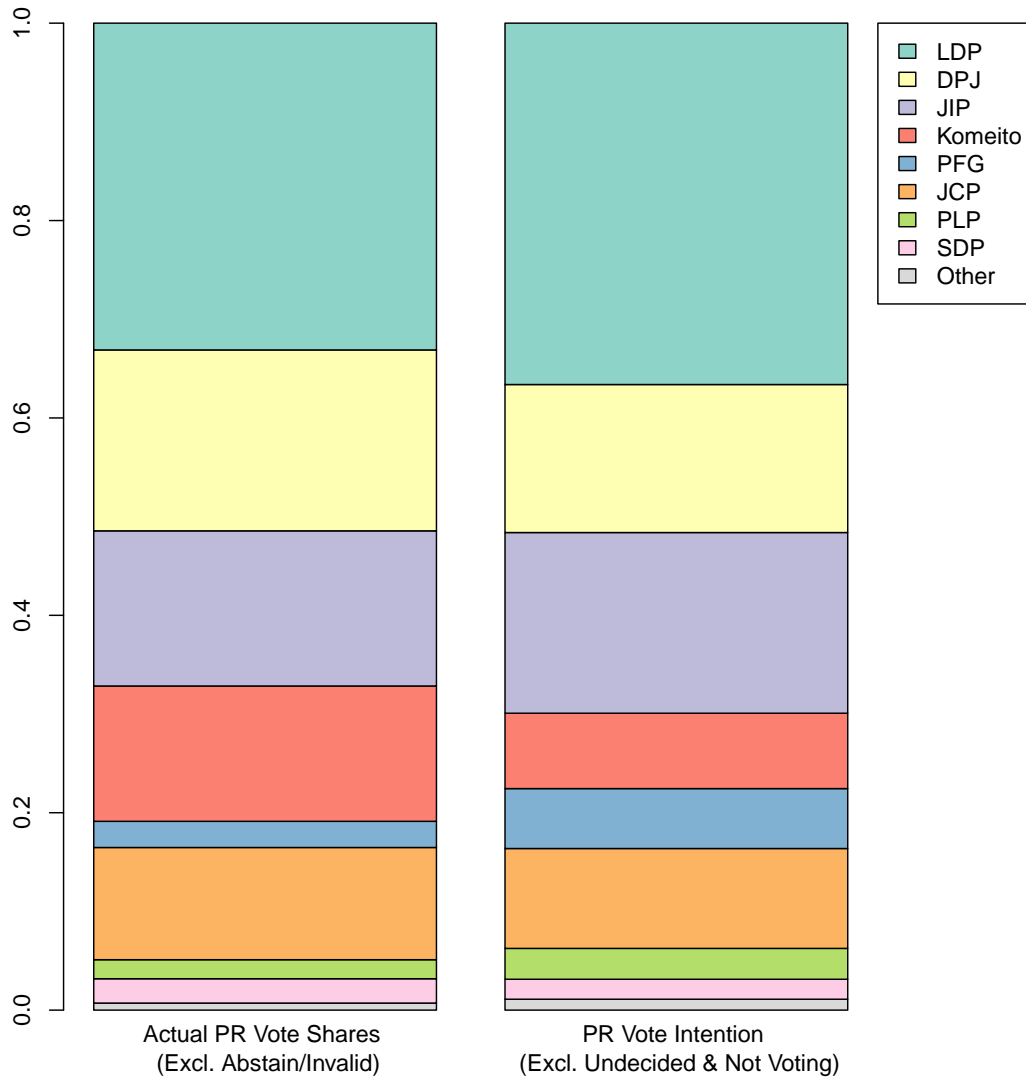


Figure C.1: Comparison between the Actual PR Vote Shares and the PR Vote Intention Variable (Excluding Undecided/Non-Voting Respondents).




Figure C.2: Comparison between the Actual PR Vote Shares and the PR Vote Intention Variable (Full Sample).




Figure C.3: Test of Attribute Order Effects. The vertical axis of the plot shows the two-sided p-values from the tests of no difference in the conditional AMCEs for attributes shown in the top row as opposed to another row of the conjoint table. The horizontal axis shows the uniform quantiles on the unit interval, which is the theoretical distribution of those p-values under the null of no difference. The p-values are obtained from a linear regression of the binary choice outcome on the policy position dummies, row position dummies, and their interactions (coefficients on the intercept and main effects are not included), with standard errors adjusted for clustering at the respondent level.




Figure C.4: Test of Effect Variation over Task Counts. The plot shows the results of analysis similar to Figure C.3, except that the p-values are calculated for interactions between the policy position dummies and the task count dummies.




Figure C.5: Average Effects of Policy Positions on Respondents' Preference for a Bundle that is Highly Unlikely to be an Actual Party Manifesto. The plot shows the same estimates as in Figure 2 for the subsample of highly unlikely bundles. See Footnote 19 for the exact definition of these bundles.