


LETTER

Trading Liberties: Estimating COVID-19 Policy Preferences from Conjoint Data

Felix Hartmann¹ , Macartan Humphreys^{2,3,4}, Ferdinand Geissler¹, Heike Klüver¹ and Johannes Giesecke¹

¹Department of Social Science, Humboldt University of Berlin, 10117 Berlin, Germany; ²Department of International Economics, Government and Business, Copenhagen Business School, DK-2000, Frederiksberg, Denmark; ³Institutions and Political Inequality Research Unit, WZB Berlin Social Science Center, 10785 Berlin, Germany; ⁴Department of Political Science, Columbia University, New York, NY 10027, USA

Corresponding author: Felix Hartmann; Email: f.hartmann@hu-berlin.de

(Received 16 May 2022; revised 22 December 2022; accepted 25 December 2022)

Abstract

Survey experiments are an important tool to measure policy preferences. Researchers often rely on the random assignment of policy attribute levels to estimate different types of average marginal effects. Yet, researchers are often interested in how respondents *trade-off* different policy dimensions. We use a conjoint experiment administered to more than 10,000 respondents in Germany, to study preferences over personal freedoms and public welfare during the COVID-19 crisis. Using a pre-registered structural model, we estimate policy ideal points and indifference curves to assess the conditions under which citizens are willing to sacrifice freedoms in the interest of public well-being. We document broad willingness to accept restrictions on rights alongside sharp heterogeneity with respect to vaccination status. The majority of citizens are vaccinated and strongly support limitations on freedoms in response to extreme conditions—especially, when they vaccinated themselves are exempted from these limitations. The unvaccinated minority prefers no restrictions on freedoms regardless of the severity of the pandemic. These policy packages also matter for reported trust in government, in opposite ways for vaccinated and unvaccinated citizens.

Keywords: survey experiments; utility theory; structural modeling; conjoint analysis; preferences

Edited by: Jeff Gill

1. Introduction

The idea that citizens have indispensable rights lies at the heart of liberal democracies. However, when major crises such as the outbreak of a pandemic, a natural disaster, or a terrorist attack hit, the question arises whether and how much governments can limit civil rights in the interest of public welfare and security (Ackerman 2006; Dragu 2011).

The COVID-19 pandemic has highlighted the trade-offs that governments have to make. In order to contain the spread of the virus, governments have limited individual rights in an unprecedented way ranging from obligations to wear facial masks to general lockdowns. Mandatory policies, though more likely to produce effects than advisories, vary in the support they receive and the levels of compliance they enjoy (Betsch *et al.* 2020). One of the basic justifications for representative democracy is that it leads to policies that citizens support. Thus, it is a central question whether citizens support restrictions on individual rights in the interest of public welfare. Government policies that are not responsive to citizen preferences could undermine trust in government and decrease compliance (Bargain and Aminjonov 2020).

Previous research has arrived at mixed findings about the extent to which citizens are committed to fundamental civil liberties (Chong 1993; Graham and Svobik 2020) though studies have shown that citizens are willing to trade individual liberties for security in light of an external threat (e.g., Davis and Silver 2004). In the context of COVID-19, survey evidence suggests that citizens with higher health-related insecurity are more willing to sacrifice civil liberties (Alsan *et al.* 2020).

We contribute to this debate by examining under which conditions citizens support restricting freedoms and how such restrictions affect trust in political institutions. A conjoint experiment allows us to explore how citizens' preferences for civil rights restrictions along the dimensions of stringency and universality vary by the severity of the crisis. Critically, the factorial design allows us not only to estimate the marginal effect of each policy dimension, but also their *interactions* with crisis severity. Using a pre-registered structural model, we estimate policy ideal points and indifference curves for representative members of different subgroups. Unlike earlier studies (e.g., Alsan *et al.* 2020), we are able to directly model the willingness to *trade-off* individual freedoms and public welfare. This use of a factorial experiment to represent multidimensional policy preferences relates most closely to Abramson, Koçak, and Magazinnik (2022) and Ganter (2021). We describe the ways we extend their approaches below. To facilitate implementation, we develop an open source routine, `cjEuclid`, in R.¹

Our results show that there is a greater acceptance of more stringent policies as severity increases. Nevertheless, there is sharp heterogeneity in ideal points: vaccinated citizens strongly support stringent policies in response to extreme conditions while the small share of unvaccinated citizens are against restrictions in *all* conditions. Vaccinated citizens differentiate between vaccinated and unvaccinated fellow-citizens and are most likely to support restrictions for unvaccinated people only.

2. Research Design

We fielded a preregistered experiment embedded in a nationally representative survey in Germany. We recruited 10,525 respondents between 8 and 22 September 2021 (for details, see the Supplementary Material). We use a $3 \times 3 \times 3$ factorial design, with each respondent shown two different vignettes. The two vignettes differ on two policy dimensions (policy stringency and universality), but agree on the third dimension, which provides a background condition (pandemic severity).

First, we describe the status quo in terms of the severity of the pandemic at the time of the survey. Second, we randomly assign subjects to consider one of three alternative levels of *severity* of the pandemic: a worsening of the situation, a sharp worsening or a dramatic worsening.

We then presented two different proposals on how to counteract such a development. First, we randomly assign an attribute describing the *stringency* of restrictions. Second, we assign an attribute to capture variation in the *universality* of restrictions. See Table 1 for details on these variations. Last, respondents were asked their preferred choice between two proposals (0/1), their rating of each of the two proposals (0–10), their feeling of trust toward the (federal) government, in case, one proposal would be implemented (0–10; where 0 is “No trust at all” and 10 is “Complete trust”), and—for the unvaccinated—their likelihood of taking a vaccination under each proposal (0–10). Each respondent received two packages of two vignettes successively.

2.1. Estimation

Our primary analysis examines the average marginal effects of conditions on support for policy and our secondary analysis uses the experiment to assess citizens' “optimal” policies under different conditions. For the primary analysis, we regress outcomes on the three dimensions together with all interactions. All variables are centered on 0 allowing for ease of interpretation of linear terms as average

¹The replication code and data for this article can be found at Hartmann *et al.* (2023). The package can be located here: <https://github.com/macartan/cjEuclid>.

Table 1. Experimental design.

Factor	Level
Pandemic severity	(0) <i>Moderate worsening</i> (7-day-incidence 150, intensive care bed occupancy 80%)
	(1) <i>Sharp worsening</i> (7-day-incidence 300, intensive care bed occupancy 90%)
	(2) <i>Dramatic worsening</i> (7-day-incidence 800, intensive care bed occupancy 100%)
Policy stringency	(0) <i>Least restrictions</i> (masks)
	(1) <i>Moderate restrictions</i> (plus limitations on social events)
	(2) <i>Most restrictions</i> (plus broader limitations on movements)
Policy universality	(0) <i>Most exemptions</i> (restrictions do not apply to vaccinated, recovered, or tested citizens)
	(1) <i>Some exemptions</i> (restrictions do not apply to vaccinated or recovered citizens)
	(2) <i>Fewest exemptions</i> (restrictions apply to all citizens)

effects. We estimate marginal effects using an OLS regression with individual level fixed effects and heteroskedasticity-robust standard errors:

$$\begin{aligned}
 Y_{it} = & \beta_0 + \beta_1 Z^1 + \beta_2 Z^2 + \beta_3 Z^3 + \\
 & \beta_3 Z^1 Z^2 + \beta_4 Z^1 Z^3 + \beta_5 Z^2 Z^3 + \beta_6 Z^1 Z^2 Z^3 + \\
 & u_i + \epsilon_{it},
 \end{aligned} \tag{1}$$

where Y is a binary variable measuring the policy choice in each round ($t \in \{1, 2\}$), each described by three factors Z^1 (severity), Z^2 (stringency), and Z^3 (universality). Note that in this demeaned regression, the constant captures the average outcome across all conditions and the coefficient on Z^1 captures the average effect of Z^1 across other conditions (additional interaction terms do not need to be added for this average effect because at the mean of the (demeaned) other conditions, these additional terms are zero; see also Lin 2013).

3. Results

Figure 1 presents the results. The key quantities of interest are the interaction terms: the extent to which citizens' preferences depend on conditions. We begin, however, by examining the average effects of each factor. We see first that, *on average*, no outcomes are responsive to background pandemic severity: citizens do not base assessments of policies in general, or trust in government on pandemic severity (though as will see their support for particular policies do depend on severity). Citizens, on average, are also not strongly responsive to policy stringency, all else equal, though there is evidence of a weak preference for less stringent policies. On average, citizens prefer less universal policies—that is, they prefer policies that allow exemptions for the vaccinated or the tested—more universal policies (which do not reward vaccination) are also associated with marginally lower likelihoods of vaccination among the unvaccinated.²

Turning to the interactions, there is a clear evidence that citizens are more supportive of more stringent policies as conditions worsen. This is the most important result of our analysis. We see from the interaction between severity and stringency that citizens strongly prefer more severe policies when conditions are bad; and in these cases, greater stringency is associated with greater willingness to vaccinate. The severity–universality interaction shows that citizens are still less supportive of universal restrictions (with fewest exemptions) when conditions are bad, though these interactions are generally

²See Supplementary Material for AMCEs (Section J of the Supplementary Material), refreshment sample analyses (Section K of the Supplementary Material), and fitted values (Section L of the Supplementary Material).

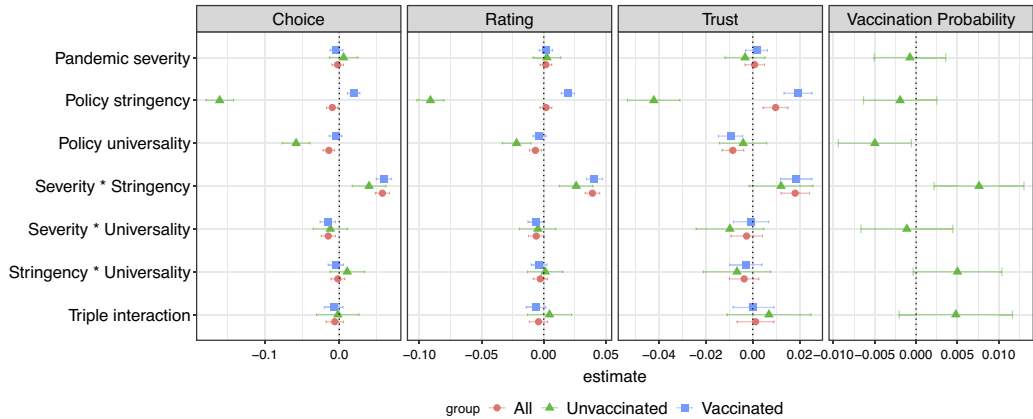


Figure 1. Average marginal effects, with two-way interactions and individual fixed effects; bars represent 95% confidence intervals. All variables are centered on zero. The coefficients on the severity * stringency and severity * universality interactions capture the extent to which respondents put more weight on severity and universality as a function of pandemic severity.

weaker and not significant at the 0.05 level for trust. We see no evidence for interactions between the stringency and universality of conditions or for three-way interactions.

3.1. Marginal Effects: Unvaccinated Citizens

The subgroup effects reported in Figure 1 show that there is stronger aversion to more stringent policies among unvaccinated citizens as well as a stronger reaction against more universal policies—the latter driven specifically by an aversion to policies that exempt only those who are vaccinated or recovered. Unvaccinated citizens prefer that they themselves are also exempted from restrictions if they can show a negative test. They show also greater tolerance for restrictions when conditions are severe, but the interaction is not strong enough to result in preferences for greater restrictions in any condition. Likewise, opposition to vaccination is weakest when stringent policies are put in place in severe conditions, but the effects are substantively small. Whereas stringency is associated with greater trust on average, it is associated with lower trust for the unvaccinated at all levels of pandemic severity. The willingness to vaccinate is also not responsive to background conditions.

3.2. Policy Ideal Points

It is difficult to identify citizens’ optimal policies from the marginal effects discussed above and to characterize how ideal points on one dimension depend on values on other dimensions. To address this problem, we posit a utility function that represents underlying preferences and estimate parameters of the utility function that in turn allow us to calculate ideal points and the slopes of indifference curves. More agnostically, this approach provides a way to summarize experimentally identified quantities in a way that makes them interpretable using a model of preferences. This approach contrasts for instance to that in Ganter (2021) which does not assume a utility model, but also restricts statements to averages of contributions of factors to choices.³

To characterize preferences, we assume citizens’ preferences can be represented by a general linear quadratic function of the form $u(z) = a + q'z - z'rz$, where z is an n -dimensional vector of attributes, r is a $n \times n$ symmetric parameter matrix, and q is an n -dimensional parameter vector. Note that we

³More specifically, preferences in Ganter (2021) are defined over levels of factors rather than over packages, as is the norm in decision theory and trade-offs across dimensions are not considered.

presuppose here continuous preferences over the policy space (though we do not presuppose convex preferences).

Expanded, utility can be written as a function of linear terms, quadratic terms, and two-way interactions of attributes:

$$u(z) = a + \sum_j \left(q_j z_j - r_{j,j} z_j^2 - \sum_{k \neq j} r_{j,k} z_j z_k \right). \tag{2}$$

This linear quadratic model has a direct relationship with the generalized Euclidean model used in the spatial theory of voting (see, e.g., Enelow and Hinich 1984). Specifically, if r is positive semi-definite, then estimated utility is consistent with “Euclidean” (spatial) preferences, with an ideal point and utility declining with distance from the ideal. A common representation in this case is $u(z) = k - (z - z^*)' r (z - z^*)$, where z^* is an n -dimensional vector of ideal points and r , the same matrix as before, captures the weighting of dimensions and the correlation of preferences across dimensions.⁴

The utility function in Equation (2) is a generalization of that found in Abramson *et al.* (2022) which presupposes circular indifference curves. In particular, if r is an identity matrix—implying not just that r is positive semi-definite but also that all dimensions have equal salience and there is no dependence across dimensions, then Equation (2) reduces to $u(z) = a + \sum_j (q_j z_j - z_j^2)$ which may be rewritten as the familiar expression seen in Abramson *et al.* (2022) $u(z) = b - \sum_j (z_j^* - z_j)^2$, where $z_j^* = q_j/2$ and $b = a - \sum_j (q_j^2/4)$.

If, however, the implied weighting matrix is not positive semi-definite, then the implied preferences are not Euclidean. This does not preclude the representation of preferences via the general linear quadratic model in Equation (2). In this case, we still have indifference curves, and ideal points (now on the edge of the policy space), however, utility is not guaranteed to decline monotonically with distance from ideals.

With the more general expression, policy ideal points on a dimension can depend on the situation on another dimension. For instance, the ideal severity of restrictions can be higher when the future risk is more extreme, allowing us to learn about, rather than assume, the structure of trade-offs. We can imagine all parameters varying across all individuals in which case a regression of ratings on attributes with two-way interactions and quadratic terms provides an estimate of the average values of parameters. In practice, we use Equation (2) as our regression equation (with each interaction term entering once not twice as in Equation (2)) and calculate average parameter values for different subgroups. In all, with n dimensions, the regression estimates $\left(1 + n + n + \frac{n(n-1)}{2} \right)$ terms—a constant, n linear terms, n quadratic terms, and $n(n-1)/2$, two-way interactions; this corresponds exactly to the number of terms in the utility function: the constant, n values for q , n points on the diagonal of r , and $n(n-1)/2$ of diagonal terms.

From these parameter estimates, we can represent estimated policy ideal points (blue dots in Figure 2) and implied indifference curves (ellipses). We can also calculate the implied weighting matrices and assess whether they are positive semi-definite (we report these matrices in Section G of the Supplementary Material).

One advantage of the structural representation is that it provides estimates of ideal points, which are not directly readable from Figure 1. Individuals might be sensitive to policy changes without substantial

⁴These two representations are transformations of each other when r is invertible. In particular, we have that if $k - (z - z^*)' r (z - z^*) = k - z^* r z^* - z' r z + 2z^* r z = a + q' z - z' r z$, then $a = k - z^* r z^*$ and $q' = 2z^* r$ and so $z^* = (q' r^{-1})/2$. By the same token, we can see that policy ideal points on dimension j are found by taking the first derivative of utility (Equation (2)) with respect to z_j , setting this to zero and solving for z_j^{**} . For instance, for $j = 1$, this yields

$$z_1^{**}(z_2, z_3) = \frac{q_1 - 2r_{1,2}z_2 - 2r_{1,3}z_3}{2r_{11}},$$

which depends on z_2 and z_3 .

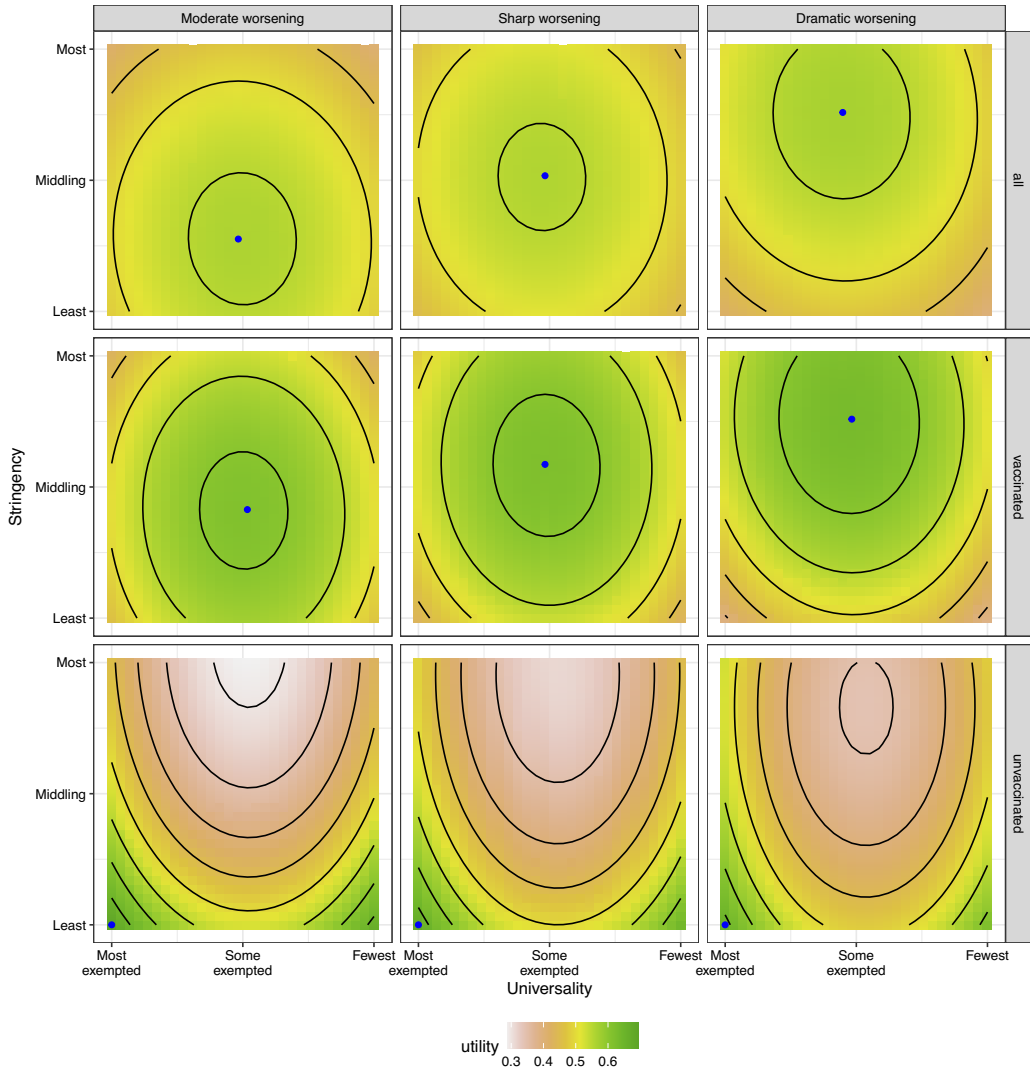


Figure 2. Ideal points and trade-offs. Blue points mark the ideal policy and contours indicate utility loss with distance from ideal points. Contours can be interpreted as indifference curves. Note that for the unvaccinated group, the level sets are non-convex and ideal points are on the edge of the policy space.

change in their actual policy ideal points, for instance, if these are at the corner of a policy space. We show these first for a “representative citizen”—that is, a citizen whose utility is given by the *average* values of the parameters in Equation (2).⁵

The upper panel of Figure 2 presents the estimated ideal points and indifference curves for the representative citizen. These are generated by calculating fitted values for all policy profiles using Equation (2). Note that, following our assumption of continuous preferences over the policy space, we

⁵We refer to a representative citizen rather than an average citizen because ideal points are nonlinear in parameters. Absent an assumption on homogeneity of parameters on the diagonal of r across subjects (as in Abramson *et al.* 2022), the implied ideal point constructed from average parameter values is not guaranteed to be equal to the average ideal point. Nevertheless, if we define ideal points on a subjectively rescaled space according to $\tilde{z}_j^* = z_j^* r_{jj}$, then \tilde{z}_j^* is indeed linear in parameters.

fit on a continuous space but estimates are derived from responses to the $3^3 = 27$ points in this space. Looking across columns, we see first that preferences are convex, consistent with Euclidean preferences.⁶ In addition, *ideal points* on policy stringency are strongly responsive to pandemic conditions. For the full sample, the ideal is within the policy range and shifts from around midway between least and moderately stringent when there is a moderate worsening to around midway between moderate stringency and severe stringency when there is a dramatic worsening of conditions. This is broadly in line with raw data patterns where in head-to-head comparisons, about 40% prefer the most stringent policies (relative to the middle category) when conditions are severe and 52% prefer the least stringent policies (relative to the middle category) when conditions have only moderately worsened. In contrast, estimated ideal points on the universality dimension are essentially unresponsive to conditions.

Looking at the *indifference curves*, we see that when conditions are only moderately poor, the representative citizen is approximately indifferent between the least stringent policies applied universally, and the most stringent policies applied with exemptions for vaccinated citizens. When there is dramatic worsening, they strictly prefer the most stringent policies (though always with exemptions for the vaccinated). With moderate worsening, unvaccinated citizens are approximately indifferent between moderate policies applied universally and weak policies targeted at them—that is, they are willing to accept more stringent policies if they are not targeted at them; their ideal, however, is for the least stringent and least universal policy.

These patterns mask heterogeneity across individuals, however. In Supplementary Material, we report on heterogeneity by party and by occupational group (Section F of the Supplementary Material). Here, we focus on vaccination status. The lower part of the figure represents the utilities for a representative member of the unvaccinated group. These report remarkable ideal policy invariance, with preferences for the least universal, least severe restrictions under all pandemic conditions. The representative unvaccinated citizen exhibits non-convex preferences on the policy space as we have defined it, preferring either that restrictions apply universally or that they are applied with broad exemptions, including for tested unvaccinated citizens.

Thus, while we can model their preferences using our linear quadratic model, they cannot be given the interpretation of Euclidean preferences. On the one hand, this highlights the flexibility of the procedure we use here: rather than imposing spatial preferences, we model preferences over the space and learn whether these can reasonably be represented by a Euclidean utility function. On the other hand, the inability to represent the preferences of the unvaccinated as spatial preferences can call our representation of the policy space into question. For instance, here the unvaccinated may focus less on the number exempted, but on the fact that in one policy option, *they* are specifically targeted. Further exploration of this possibility would benefit from data with a finer representation of the space than we have available to us.

On a technical note, we can see from the models that in this application, we cannot reject the null of separability across the two policy dimensions—that is, we cannot reject the null of no interaction between severity and universality for either the vaccinated or unvaccinated subgroup. We can however generally reject the null of no interaction between severity and each of the two policy dimensions. That is, preferences for each dimension depend on severity, but the dimensions do not interact strongly with each other. We can also reject the null of uniform salience across dimensions in eight out of nine tests—the same weight applied to both dimensions—though this needs to be understood in terms of the metric stipulated by our conditions. Finally, although we estimate parameters here assuming preferences are defined with respect to three dimensions, in fact subjects were evaluating policies and so it is natural to estimate preferences over policy *conditional* on severity. We do this in Section E of the Supplementary Material with consistent findings.

⁶Strictly, preferences over the two policy dimensions are convex. See Section G of the Supplementary Material for the estimated weighting matrices.

3.2.1. Caveats and External Validity

We acknowledge several caveats. First, the study measures stated preference using hypothetical scenarios. Nevertheless, preferences of citizens are important for public welfare evaluations and can be predictive of behavior (Hainmueller, Hangartner, and Yamamoto 2015). Second, by projecting policy options onto two dimensions, stringency and universality, we limit the features of policies that may affect support and compliance. Schmelz (2021) documents, for instance, greater opposition to mandatory use of a contact tracing app than to restrictions on movements. Thus, our findings on the stringency dimension in particular should be understood in terms of how this dimension was operationalized (Table 1). Third, the structural model analysis is model dependent—assuming, in particular, a generalized linear quadratic functional form which implies symmetries in preferences over a continuous space.

4. Conclusion

In this study, we examined under what conditions the citizens are willing to sacrifice individual civil rights in the interest of national security and public health. Combining data from a conjoint experiment with a structural model in which individuals are assumed to have preferences representable by a generalized linear quadratic utility function, we estimate policy ideal points, preference trade-offs across policies, and the sensitivity of ideal points to context.

Our results show that there is a broad acceptance of mandatory restrictions and that there is a greater acceptance of more stringent policies as severity increases. However, our findings highlight not just the broad support for restrictions in the general interest, but also the ways in which restrictive policies are contested. We document, in particular, the willingness of the majority to target restrictions against unvaccinated citizens who themselves are least supportive of any restrictions.

Acknowledgments. We would like to thank Flavien Ganter, the three anonymous reviewers, and the editor (Jeff Gill) for their thoughtful comments and discussion. Research for this contribution is part of the Cluster of Excellence “Contestations of the Liberal Script” (EXC 2055, Project-ID: 390715649), funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy.

Supplementary Material. For supplementary material accompanying this paper, please visit <https://doi.org/10.1017/pan.2023.25>.

Data Availability Statement. The replication code and data for this article can found at Hartmann *et al.* (2023).

Ethics Statement. Informed consent was obtained after the nature and possible consequences of the study were explained to survey participants. The analyses were pre-registered (<https://osf.io/4vvgf6/>) and obtained institutional review board approval at Humboldt University of Berlin.

References

- Abramson, S. F., K. Koçak, and A. Magazinnik. 2022. “What Do We Learn about Voter Preferences from Conjoint Experiments?” *American Journal of Political Science* 66 (4): 1008–1020.
- Ackerman, B. 2006. *Before the Next Attack: Preserving Civil Liberties in the Age of Terrorism*. New Haven, CT: Yale University Press.
- Alsan, M., et al. 2020. “Civil Liberties in Times of Crisis.” Technical report. National Bureau of Economic Research.
- Bargain, O., and U. Aminjonov. 2020. “Trust and Compliance to Public Health Policies in Times of COVID-19.” *Journal of Public Economics* 192: 104316.
- Betsch, C., et al. 2020. “Social and Behavioral Consequences of Mask Policies during the COVID-19 Pandemic.” *Proceedings of the National Academy of Sciences* 117 (36): 21851–21853.
- Chong, D. 1993. “How People Think, Reason, and Feel about Rights and Liberties.” *American Journal of Political Science* 37 (3): 867.
- Davis, D. W., and B. D. Silver. 2004. “Civil Liberties Vs. Security: Public Opinion in the Context of the Terrorist Attacks on America.” *American Journal of Political Science* 48 (1): 28–46.
- Dragu, T. 2011. “Is There a Trade-Off between Security and Liberty? Executive Bias, Privacy Protections, and Terrorism Prevention.” *American Political Science Review* 105 (1): 64–78.

- Enelow, J. M., and M. J. Hinich. 1984. *The Spatial Theory of Voting: An Introduction*. Cambridge: Cambridge University Press.
- Ganter, F. 2021. "Identification of Preferences in Forced-Choice Conjoint Experiments: Reassessing the Quantity of Interest." *Political Analysis* 31: 1–15.
- Graham, M. H., and M. W. Svolik. 2020. "Democracy in America? Partisanship, Polarization, and the Robustness of Support for Democracy in the United States." *American Political Science Review* 114 (2): 392–409.
- 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 of the United States of America* 112 (8): 2395–2400.
- Hartmann, F., et al. 2023. "Replication Data for: Trading Liberties: Estimating Covid Policy Preferences from Conjoint Data." <https://doi.org/10.7910/DVN/SVD9SV>. Harvard Dataverse, V1.
- Lin, W. 2013. "Agnostic Notes on Regression Adjustments to Experimental Data: Reexamining Freedman's Critique." *Annals of Applied Statistics* 7 (1): 295–318.
- Schmelz, K. 2021. "Enforcement may Crowd Out Voluntary Support for COVID-19 Policies, Especially Where Trust in Government is Weak and in a Liberal Society." *Proceedings of the National Academy of Sciences* 118 (1): e2016385118.