



No. 58

I4R DISCUSSION PAPER SERIES

Yellow Vests, Pessimistic Beliefs, and Carbon Tax Aversion (2022): A Comment

Nicholas Rivers

Matt Woerman

Kareman Yassin

September 2023

I4R DISCUSSION PAPER SERIES

I4R DP No. 58

Yellow Vests, Pessimistic Beliefs, and Carbon Tax Aversion (2022): A Comment

Nicholas Rivers¹, Matt Woerman², Kareman Yassin¹

¹University of Ottawa/Canada

²University of Massachusetts Amherst/USA

SEPTEMBER 2023

Any opinions in this paper are those of the author(s) and not those of the Institute for Replication (I4R). Research published in this series may include views on policy, but I4R takes no institutional policy positions.

I4R Discussion Papers are research papers of the Institute for Replication which are widely circulated to promote replications and meta-scientific work in the social sciences. Provided in cooperation with EconStor, a service of the [ZBW – Leibniz Information Centre for Economics](#), and [RWI – Leibniz Institute for Economic Research](#), I4R Discussion Papers are among others listed in RePEc (see IDEAS, EconPapers). Complete list of all I4R DPs - downloadable for free at the I4R website.

I4R Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

Editors

Abel Brodeur
University of Ottawa

Anna Dreber
Stockholm School of Economics

Jörg Ankel-Peters
RWI – Leibniz Institute for Economic Research

Yellow Vests, Pessimistic Beliefs, and Carbon Tax Aversion (2022): A comment

Nicholas Rivers, Matt Woerman, and Kareman Yassin*

July 21, 2023

Abstract

[Douenne and Fabre \(2022\)](#) implement a representative survey following the Yellow Vests movement in France that started in opposition to the carbon tax in 2018. They find that a majority of French citizens would oppose a carbon tax and dividend program with proceeds paid equally to each adult. The authors further find that respondents have pessimistic beliefs about several aspects of the policy. They then show how informational treatments cause respondents to update these beliefs, and they finally estimate the causal effect of these beliefs on support for the policy.

In this note, we focus on the second section of this paper: the causal effects of feedback on beliefs. Based on elicited household characteristics, [Douenne and Fabre \(2022\)](#) estimate whether each household “wins” or “loses” from the carbon tax and dividend reform. They provide this binary (win vs. lose) information to households and subsequently ask households to evaluate whether they believe they would financially benefit from the policy. By exploiting the discontinuity in win vs. lose feedback, they assess the degree to which feedback affects subjective beliefs, finding that a household that is told it will “win” as a result of the reform increases its subjective belief that it will not lose by about 25 percentage points. The subset of households that is part of the Yellow Vests movement, however, revises its subjective belief of not losing upwards by only 10 percentage points after being told that it will “win” from the carbon tax reform. Conversely, households who initially support the tax increase this belief by 41 percentage points when told they will “win.”

In this note we replicate this second section of the paper—the causal effects of feedback on beliefs—using the processed data provided by the authors. We successfully replicate the average treatment effect, but we find that the heterogeneous treatment effects may be biased due to model misspecification. While our results support the conclusion that these estimated effects depend on a household’s attitudes toward the policy, we find that the source of heterogeneity differs. Further, we note two changes to the analysis that we believe are appropriate (which do not affect the conclusions drawn): first, some (1.8%) of observations in the dataset appear to be misclassified—wrongly coded as if a household would “lose” when in fact they would “win”—and second, the main causal analysis is based on a regression discontinuity design, but does not include standard components of such a design (e.g., a RD plot, optimal selection of bandwidth, density analysis, placebo tests). We update the design to address both of these points. We find results that generally support the main conclusions of [Douenne and Fabre \(2022\)](#), but we urge caution when interpreting the heterogeneous treatment effects.

* Authors: Rivers: University of Ottawa. E-mail: nrivers@uottawa.ca. Woerman: University of Massachusetts Amherst. Email: mwoerman@umass.edu. Yassin: University of Ottawa. Email: kyass035@uottawa.ca; corresponding author. We thank Abel Brodeur and organizers of the Ottawa Replication Games for the opportunity to replicate this paper. We have received no payment for this replication, and we have no conflict of interests to report.

1 Introduction

Using a representative sample of 3,002 French households, [Douenne and Fabre \(2022\)](#) designed and implemented a survey to understand attitudes towards a carbon tax policy and how beliefs about it develop. In the survey, they suggested a policy that would increase the carbon tax by 50€/tCO₂, with the tax revenue redistributed uniformly to each adult.¹ Survey results show that only about 10% of respondents approve of the reform and 70% disapprove of it; the rest do not know or do not want to answer. In the first section of the paper, the authors show that respondents have pessimistic beliefs about the policy's environmental effectiveness, progressivity, and effects on purchasing power. In the second section, they show how an informational treatment causes respondents to update these beliefs. In the third section, they estimate the causal effect of these beliefs on support for the policy reform. In this note, we focus on the second section of this paper: the causal effects of feedback on beliefs.

To study the belief formation process and to examine the importance of self-interest and fairness motives in accepting the reform, [Douenne and Fabre \(2022\)](#) begin by eliciting the respondent's initial belief about whether they expect to financially "win" or "lose" from the policy reform, their subjective net gain in purchasing power, whether the reform would be effective in reducing pollution and combating climate change, and whether they approve of the reform or not. To half of the sample selected at random, the authors then provide an informational treatment about the effect of the reform on the household's budget. The authors ask the household about their energy equipment and use, which is then combined with official household survey data to simulate that household's net gain due to the reform (i.e., dividend less carbon tax costs). Each household is then given customized feedback on whether they are predicted to "win" or "lose" due to the reform, but not the magnitude of the win or loss. The authors tell the respondents—and show in their paper—that these predictions are correct in five out of six cases. Following this customized informational treatment, the survey again asked respondents about their beliefs regarding whether they would win or lose from implementation of the policy and about their approval of the proposed policy reform.²

To study this causal effect of feedback on beliefs, where this informational treatment is not random but customized to each household, the authors explain on page 98 that "The binary win/lose feedback is a variable that jumps from 0 to 1 when our continuous estimation of respondents' net gains exceeds the zero threshold." Thus, they were able to estimate the threshold effect around zero net gain using an empirical strategy similar to a regression discontinuity design. They focus on a subsample of respondents whose simulated net gain is less than 50€ per year in absolute value, and they estimate the discrete effect of crossing that zero net gain threshold. Their results on page 98 states that "Everything else being equal, after receiving the feedback, respondents' predicted winners are 27 p.p. more likely to believe they do not lose than respondents' predicted losers." However, they also find important sources of heterogeneity: "among respondents who initially accepted the tax and dividend policy . . . the effect of the win feedback on the belief not to lose goes from 27 p.p to 41 p.p., while for those supportive of the Yellow Vests . . . it goes down to less than 10 p.p."

In this note, we investigate the replicability of these causal estimates of feedback on beliefs. The

¹Later in the survey, the authors propose an alternative tax reform of 50€/CO₂ with the revenue redistributed based on a respondent's position in the income distribution. This reform is outside our scope of interest with respect to replication.

²The authors also have a different informational treatment on the progressiveness of the tax that is outside our replication scope of interest.

replication data and code are freely provided by the authors and available online. We use their processed dataset (named `s` in their R code files) and re-code the estimation of the causal effects of feedback on beliefs. Note that the authors further analyze how beliefs affect support for policy reforms; these additional analyses are beyond the scope of our replication exercise, so we make no claims about their replicability.

2 Replication

Our replication strategy is to use the processed data provided by the authors to replicate the results on the causal effect of information feedback on a respondent’s belief of winning, including heterogeneity in this effect. We provide data visualization as well as re-code the main estimating equations.

Our focus is on Equation 2 and Table 4 of [Douenne and Fabre \(2022\)](#).³ Equation 2 describes a regression discontinuity design. The running variable is the simulated gain (or loss) from the implementation of the policy reform, which is denoted $\hat{\gamma}$ (the variable is referred to as `simule_gain` in the dataset). This simulated gain, $\hat{\gamma}$, is not disclosed to respondents during the survey. Instead, the survey discloses binary information on whether a household is expected to win or lose from implementation of the carbon tax reform, but not the magnitude of the gain or loss; this dummy variable is denoted $\hat{\Gamma}$ (the dummy variable is referred to as `simule_gagnant` in the dataset). Thus, the research design is a sharp regression discontinuity with a cutoff at $\hat{\gamma} = 0$.

2.1 Misclassification

We first assess the relationship between the running variable ($\hat{\gamma}$) and the treatment indicator ($\hat{\Gamma}$), and we show the results in Figure 1 and Table 1. We find that some of the observations in the dataset are misclassified—the definitions of $\hat{\gamma}$ and $\hat{\Gamma}$ are inconsistent with the description in the text. In Figure 1, some observations with $\hat{\gamma} < 0$ are given a code of $\hat{\Gamma} = 1$. These households are coded as if they were told they would “win” as a result of the reform, even though $\hat{\gamma} < 0$, suggesting that they would lose. Likewise, some households are coded as if they were told that they would “lose” from the reform, even though $\hat{\gamma} > 0$. Table 1 shows the number and type of misclassified observations. In total 56 out of 3,002 observations (1.8%) received an incorrect coding of $\hat{\Gamma}$. In the analysis that follows later, [Douenne and Fabre \(2022\)](#) drop observations with simulated gains or losses greater than 50€ ($|\hat{\gamma}| > 50\text{€}$). For the subset of included observations, almost 7% of observations are misclassified. Later we show that this misclassification does not have a substantial impact on results.

2.2 Data visualization

The estimation of the causal effect of feedback on beliefs relies on a regression discontinuity. When using such an approach, it is standard to visualize data using a regression discontinuity (RD) plot. For example, [Cattaneo and Titiunik \(2022\)](#) state: “Although a graphical illustration cannot replace formal econometric analysis, it is often helpful to display the general features of the study, to gauge whether the discontinuity observed at the cutoff is unusual with respect to the overall shape of the regression function, and to have an overall understanding of the data to be used in the subsequent formal analysis.”

³We do not attempt to replicate the other analyses of [Douenne and Fabre \(2022\)](#) or make any claims about their replicability.

[Douenne and Fabre \(2022\)](#) did not provide a regression discontinuity plot to support the econometric estimates, and so we produce one here as Figure 2. We follow advice from [Cattaneo and Titiunik \(2022\)](#): “The typical strategy is to fit a polynomial of order four, but this should be modified as needed. The local means are constructed by partitioning the support of the score into disjoint bins and computing sample averages of the outcome for each bin, using only observations whose score value falls within that bin; each local mean is plotted against the bin mid-point.” While [Cattaneo and Titiunik \(2022\)](#) caution against substituting graphical analysis for formal econometric analysis, we note that the regression discontinuity plot is indicative of a treatment effect of approximately 20 percentage points—generally consistent with the findings reported in Table 4 of [Douenne and Fabre \(2022\)](#).⁴

While [Cattaneo and Titiunik \(2022\)](#) suggest including all observations in an RD plot, we note that the running variable $\hat{\gamma}$ is highly skewed—see Figure 3—with only a handful of observations where $\hat{\gamma} < -300\text{€}$. As a result, we reproduce the RD plot with a sample restriction of $\hat{\gamma} > -300\text{€}$. Again, the discontinuity in the data is clearly evident in the visualization.

2.3 Regression model

We next aim to replicate the estimation of Equation 2 of [Douenne and Fabre \(2022\)](#), which they present in column 1 of Table 4 on page 99. Table 2 provides our results of re-estimating this regression and alternative specifications using the processed data. The dependent variable in all specifications pertains to respondents’ beliefs regarding their gain after receiving feedback. It takes a value of 1 if respondents believe they will not experience losses, and 0 otherwise. Thus, the coefficient of interest can be interpreted as the causal effect of feedback on the probability that a household believes they will not lose from the proposed policy reform.

Column 1 of Table 2 is our attempt to exactly replicate (using our own code) the main result in Column 1 of Table 4 of [Douenne and Fabre \(2022\)](#). Like [Douenne and Fabre \(2022\)](#), we restrict the sample to observations within 50€ of $\hat{\gamma} = 0$, use the provided sample weights, include the set of covariates as described in the paper,⁵ and use $\hat{\Gamma}$ (`simule_gagnant` in the code) as the treatment variable. We find that the magnitude of our estimated coefficient closely aligns with the results obtained by [Douenne and Fabre \(2022\)](#). Column 2 presents the same specification as column 1, but without including the covariates.⁶ The qualitative finding is not significantly perturbed by dropping covariates.

Subsequently, we address the influence of misclassified observations, which we describe above. To do so, we set the treatment variable to be equal to one for households with $\hat{\gamma} > 0$, corresponding to the sharp RD design described by [Douenne and Fabre \(2022\)](#). Column 3 of Table 2 presents the estimation results with the inclusion of covariates, while column 4 reports the results without covariates. Notably, we find that the presence of misclassified observations does not significantly alter the qualitative outcomes of the analysis, although the value of the point estimate decreases modestly.

Figure 5 visually represents Equation 2 of [Douenne and Fabre \(2022\)](#) but excludes covariates, which is

⁴Note that we do not consider covariates in generating the RD plot.

⁵We exclude three covariates that [Douenne and Fabre \(2022\)](#) include: an indicator if the respondent’s tax acceptance is “NSP” (“ne sais pas” or “do not know”), an indicator if the respondent believes the proposed policy would be progressive, and an indicator if the respondent believes the proposed policy would reduce pollution and mitigate climate change. These covariates contain some missing values, so we cannot include them in the regression and retain the full sample of respondents.

⁶In a regression discontinuity design, covariates are typically used to increase the precision of estimates but should not alter the estimated effects ([Calonico et al. 2019](#)).

comparable to our column 2 of Table 2. This model assumes that, other than the discrete change at the discontinuity cutoff, the same global quadratic polynomial holds on both sides of the discontinuity cutoff. The figure illustrates a noticeable discontinuity at the cutoff, which corresponds to an estimated net gain of zero, indicating that beliefs change when a household is told they will gain from the proposed policy reform.

2.4 Regression discontinuity design

The regression given in Equation 2 and the estimates reported in Table 4 of Douenne and Fabre (2022) are roughly based on a regression discontinuity design, however, best practices in RDD estimation—as described by Cattaneo and Titiunik (2022)—are not incorporated in their research design. Figure 5 is a visual depiction of the estimation approach underlying Equation 2 and Table 4. Three features are important to note:

1. The regression includes a quadratic polynomial in the running variable that is restricted to be identical on both sides of the cutoff. In contrast, it is best practice to estimate a local linear polynomial that can vary on each side of the cutoff.
2. The “bandwidth” for the regression is 50€, which appears to be an *ad hoc* choice. In contrast, it is best practice to use a data-driven approach to select the optimal bandwidth in a RDD, rather than use a subjective choice that increases researcher degrees of freedom.
3. It is best practice to describe the weighting of observations within the bandwidth, and often a triangular weighting kernel is applied, which gives higher weights to observations near the cutoff. In contrast, the author do not discuss weights in their regression—other than survey weights—from which we infer the regression uses a uniform kernel.

We next re-estimate the coefficients in Table 4 of Douenne and Fabre (2022) while following the best practices in regression discontinuity designs (Cattaneo and Titiunik 2022). We use the R package `rdr robust` (Calonico et al. 2015).

Table 3 presents results. In each case, the running variable is $\hat{\gamma}$ with a sharp RD cutoff of $\hat{\gamma} = 0$. Each column of this table reports results for a different RDD specification. Within each column, we report the conventional RDD estimate, the bias-corrected RDD estimate, and the bias-corrected RDD estimate with robust confidence interval, respectively.

Column 1 of Table 3 reports results using a triangular kernel, an optimal bandwidth that minimizes mean-squared error (MSE), local linear polynomials, and no covariates. These estimates most closely reflect the best practices described above. We find results that are consistent with our Table 2 and Table 4 of Douenne and Fabre (2022).

We also estimate five additional RDD specifications that incorporate alternative researcher choices. Here we describe how those specifications differ from column 1. Column 2 uses a uniform kernel and a bandwidth of 50€, as in Douenne and Fabre (2022). Column 3 uses a uniform kernel. Column 4 estimates a local quadratic polynomial of the running variable. Column 5 includes the set of covariates described above. Column 6 drops the misclassified observations that we described above.

We generally find consistent results across all of these different RDD specifications. In some cases, the magnitude of the treatment effect is modestly smaller, but still positive and statistically significant.

Treatment effects lose statistical significance only in column 4 when using a local quadratic polynomial, which is not considered a best practice.

In conclusion, use of the `rdrobust` package in R to estimate this RDD—including the best practices in this research design—yields results that are consistent with our earlier findings and the results of [Douenne and Fabre \(2022\)](#). The use of a triangular kernel with MSE-optimized bandwidth, the inclusion of covariates, and the exclusion of misclassified observations did not significantly alter the estimated effects. This finding reaffirms the robustness and applicability of the main results, supporting their generalizability and reliability. To summarize, our results do not differ qualitatively from the main results of [Douenne and Fabre \(2022\)](#), providing confidence in the reliability and applicability of the main findings.

2.5 Placebo tests

In Table 4 we present placebo tests with alternative dependent variables, rather than the one of interest. To validate the main results, we should not find significant estimates because the placebo outcomes chosen are not theoretically or empirically expected to be related to the treatment indicator, feedback on net gains, which changes discretely at $\hat{\gamma} = 0$. We test four placebo outcomes: an indicator if the respondent is male, an indicator if the respondent is young (below 30 years old), an indicator if the respondent is old (greater than 65 years old), and household size. We find no effect on any of these outcomes at the running variable cutoff of $\hat{\gamma} = 0$, indicating no other changes at the cutoff that could confound the main results. Thus, Table 4 supports the main results of the paper that there is a causal effect of feedback on beliefs.

We also use alternative running variable cutoff values as another placebo test. Because the feedback given to households changed discretely only at the running variable cutoff of $\hat{\gamma} = 0$, other cutoffs should not yield a regression discontinuity with statistically significant treatment effects on household beliefs. Figure 6 shows that most placebo cutoffs generate no treatment effect. Interestingly, however, cutoffs around -60€ also generate treatment effect estimates that are roughly as large as the main estimate.

2.6 Heterogeneous treatment effects

Another key result of [Douenne and Fabre \(2022\)](#) is that the effect of feedback is heterogeneous and varies based on the respondent's pre-existing attitudes about the policy reform, which the authors report in Table 4 on page 99. For example, respondents who initially accepted the reform are more likely to change their beliefs about not losing from the reform, while respondents who support the Yellow Vests are less likely to do so. We finally aim to replicate these heterogeneous treatment effects.

To estimate this heterogeneity, [Douenne and Fabre \(2022\)](#) add to Equation 2 an interaction between the treatment variable $\hat{\Gamma}$ and each heterogeneity indicator. They do not, however, interact the running variable $\hat{\gamma}$ with each heterogeneity indicator, which implicitly assumes the quadratic relationship between the running variable and the outcome of interest does not exhibit any heterogeneity. [Carril et al. \(2018\)](#) describe how this assumption, if violated empirically, can bias subgroup treatment effect estimates in an RDD.

To further investigate the existence of heterogeneity, we first generate an RD plot for each subgroup considered by [Douenne and Fabre \(2022\)](#). Figure 7 gives the RD plots by tax acceptance, Figure 8 gives

the RD plots by Yellow Vests support, and Figure 9 gives the RD plots by initial belief of loss. Comparing the two subgroups in each figure provides graphical evidence of heterogeneous treatment effects at the cutoff of $\hat{\gamma} = 0$. It also shows, however, that on each side of the cutoff, the relationship between the running variable and the outcome may also exhibit heterogeneity. Thus, as cautioned by Carril et al. (2018), the polynomials of the running variable must also be interacted with each heterogeneity indicator to ensure subgroup RDD estimates are not biased.

We next replicate the heterogeneity results of Douenne and Fabre (2022) and estimate alternative specifications that correct for any potential bias. We focus our replication on these results presented in column 3 of their Table 4; we do not consider alternative definitions of these heterogeneity variables or alternative outcome variables. We report these results in Table 5. Column 1 corresponds to column 3 of Table 4 in Douenne and Fabre (2022), showing that we closely replicate their results when using a comparable specification. In column 2, we add interactions with the global quadratic of the running variable; in column 3, we further replace the global quadratic with local linear polynomials; and in column 4, we replace the uniform kernel with a triangular kernel.

Estimates from the latter three regressions differ importantly from the results of Douenne and Fabre (2022) and our column 1. We find that, for households that initially accepted the proposed policy reform, the feedback treatment increases their probability of believing they will not lose by an additional 38–62 percentage points (p.p.), as compared to only 15 p.p. in column 1. Additionally, we estimate that among supporters of Yellow Vests, this treatment effect is 8–16 p.p. smaller, but these estimates are imprecise and not statistically different from zero at conventional levels of significance; these results contrast with the potentially biased results in column 1 that indicate treatment effects are 16 p.p. smaller—and statistically significant—for Yellow Vests supporters.

In summary, the heterogeneity results in Table 4 of Douenne and Fabre (2022) may be biased due to model misspecification. When we correct for this bias, we find that the feedback treatment is even more persuasive among households who already accept the reform, but we cannot reject that it has no differential effect on households that show opposition to the reform by supporting the Yellow Vests. These results support the general conclusion of Douenne and Fabre (2022) that the effect of this feedback depends on the respondent’s attitudes toward the reform, but we show the primary driver is acceptance of the reform and not the respondent’s support for the Yellow Vests movement.

3 Conclusion


In conclusion, our replication results of the RDD regressions are generally consistent with the comparable findings of Douenne and Fabre (2022), focusing on the causal effects of feedback on beliefs. The exclusion of “misclassified observations” and re-coding the regression using the `rdrobust` in package R does not significantly alter the main results. We do, however, find that the heterogeneous treatment effects may be biased. When we correct for this bias, we still find heterogeneous effects due to prior attitudes about the reform, but the exact source of heterogeneity differs. In general, we reaffirm the robustness and applicability of the main results, but we urge caution when interpreting the heterogeneous treatment effects.

References

- Calonico, S., Cattaneo, M. D., Farrell, M. H. and Titiunik, R.: 2019, Regression Discontinuity Designs Using Covariates, *The Review of Economics and Statistics* **101**(3), 442–451.
- Calonico, S., Cattaneo, M. D. and Titiunik, R.: 2015, The r journal: rdrobust: An r package for robust nonparametric inference in regression-discontinuity designs, *The R Journal* **7**, 38–51. <https://doi.org/10.32614/RJ-2015-004>.
- Carril, A., Cazor, A., Gerardino, M. P., Litschig, S. and Pomeranz, D.: 2018, Subgroup analysis in regression discontinuity designs.
- Cattaneo, M. D. and Titiunik, R.: 2022, Regression discontinuity designs, *Annual Review of Economics* **14**, 821–851.
- Douenne, T. and Fabre, A.: 2022, Yellow vests, pessimistic beliefs, and carbon tax aversion, *American Economic Journal: Economic Policy* **14**(1), 81–110.


4 Figures

Figure 1: Miscalclassification of observations




Notes: This figure depicts the relationship between the running variable (named `simule_gain` in the code and denoted by $\hat{\gamma}$) and the treatment indicator (named `simule_gagnant` in the code and denoted by $\hat{\Gamma}$). Although this research design is a sharp regression discontinuity—with treatment status changing discretely when the running variable crosses zero—56 survey respondents are not coded as such. 54 respondents were simulated to lose from the proposed policy ($\hat{\gamma} < 0$) but were assigned to be “winners” ($\hat{\Gamma} = 1$); 2 respondents were simulated to gain from the proposed policy ($\hat{\gamma} > 0$) but were assigned to be “losers” ($\hat{\Gamma} = 0$).

Figure 2: Regression discontinuity plot




Notes: This figure plots the relationship between the running variable (named `simule_gain` in the code and denoted by $\hat{\psi}$) and the outcome of interest—whether the survey respondent believes they will not lose from the proposed policy—for all survey respondents. The dots give average outcomes for bins of the running variable. The lines are fourth-order polynomials of the data and can vary on each side of the discontinuity cutoff.

Figure 3: Density of the running variable




Notes: This figure plots the density of the running variable (named `simule_gain` in the code and denoted by $\hat{\psi}$). Very few observations have a running variable value below -300.

Figure 4: RD plot closer look




Notes: This figure plots the relationship between the running variable (named `simule_gain` in the code and denoted by \hat{y}) and the outcome of interest—whether the survey respondent believes they will not lose from the proposed policy—for survey respondents with a running variable value of -300 or greater. The dots give average outcomes for bins of the running variable. The lines are flexible polynomials of the data and can vary on each side of the discontinuity cutoff.

Figure 5: Econometric model with global quadratic




Notes: This figure depicts the econometric model that [Douenne and Fabre \(2022\)](#) estimate using Equation 2 from page 98 but excluding covariates, which is comparable to our column 2 of Table 2. This model assumes that, other than the discrete change at the discontinuity cutoff, the same global quadratic polynomial holds on both sides of the discontinuity cutoff.

Figure 6: Placebo cutoffs




Notes: This figure plots the results of placebo tests using different regression discontinuity cutoff values. The blue dot and error bars correspond to the main specification with a threshold of zero. The lines on either side present the estimated treatment effects using alternative cutoff values, as given on the x-axis. Error bars and shaded areas give 95% confidence intervals. Few alternative cutoff values yield statistically significant treatment effects.

Figure 7: RD plots by initial tax acceptance




Notes: This figure plots the relationship between the running variable (named `simule_gain` in the code and denoted by \hat{y}) and the outcome of interest—whether the survey respondent believes they will not lose from the proposed policy—separately for respondents who initially accept the tax (left panel) and for respondents who initially do not accept the tax (right panel). The dots give average outcomes for bins of the running variable. The lines are flexible polynomials of the data and can vary on each side of the discontinuity cutoff.

Figure 8: RD plots by Yellow Vests support



Notes: This figure plots the relationship between the running variable (named `simule_gain` in the code and denoted by \hat{y}) and the outcome of interest—whether the survey respondent believes they will not lose from the proposed policy—separately for respondents who do not support Yellow Vests (left panel) and for respondents who support Yellow Vests (right panel). The dots give average outcomes for bins of the running variable. The lines are flexible polynomials of the data and can vary on each side of the discontinuity cutoff.

Figure 9: RD plots by initial belief of loss



Notes: This figure plots the relationship between the running variable (named `simule_gain` in the code and denoted by \hat{y}) and the outcome of interest—whether the survey respondent believes they will not lose from the proposed policy—separately for respondents who initially believe they do not lose (left panel) and for respondents who initially believe they will lose (right panel). The dots give average outcomes for bins of the running variable. The lines are flexible polynomials of the data and can vary on each side of the discontinuity cutoff.

5 Tables

Table 1: Classification of observations

$\hat{\gamma} > 0$	$\hat{\Gamma} = 1$	n
FALSE	FALSE	742
FALSE	TRUE	54
TRUE	FALSE	2
TRUE	TRUE	2204

Notes: This table displays the number of survey respondents that were correctly classified or misclassified as a “winner” or “loser” in the data. 54 respondents were simulated to lose from the proposed policy ($\hat{\gamma} < 0$) but were assigned to be “winners” ($\hat{\Gamma} = 1$); 2 respondents were simulated to gain from the proposed policy ($\hat{\gamma} > 0$) but were assigned to be “losers” ($\hat{\Gamma} = 0$). The remaining 2946 survey respondents were correctly classified.

Table 2: Re-estimating Equation 2

	(1)	(2)	(3)	(4)
Coded as predicted winner ($\hat{\Gamma} = 1$)	0.272 (0.059) [<0.001]	0.241 (0.056) [<0.001]		
Predicted winner ($\hat{\gamma} > 0$)			0.216 (0.066) [0.001]	0.234 (0.069) [<0.001]
Covariates	X		X	
Observations	757	757	757	757
R^2	0.269	0.081	0.258	0.073

Notes: This table reports results of estimating Equation 2 from page 98 of [Douenne and Fabre \(2022\)](#) using their processed data. In all specifications, the dependent variable takes a value of 1 if, after receiving feedback, the respondent believes they will not experience losses from the proposed policy; it takes a value of 0 otherwise. Columns 1 and 2 report results when using the treatment variable as it is coded in the data. Columns 3 and 4 report results when we reclassify treatment status to correctly correspond to the running variable, or predicted gains from the policy. Columns 1 and 3 include covariates, while columns 2 and 4 do not. Standard errors are in parentheses and p-values are in brackets.

Table 3: Estimating RDD using rdrobust package

	(1)	(2)	(3)	(4)	(5)	(6)
Conventional	0.232 (0.067) [<0.001]	0.238 (0.066) [<0.001]	0.241 (0.064) [<0.001]	0.157 (0.101) [0.117]	0.190 (0.080) [0.018]	0.282 (0.069) [<0.001]
Bias-Corrected	0.234 (0.067) [<0.001]	0.183 (0.066) [0.005]	0.261 (0.064) [<0.001]	0.131 (0.101) [0.193]	0.169 (0.080) [0.035]	0.287 (0.069) [<0.001]
Robust	0.234 (0.081) [0.004]	0.183 (0.102) [0.072]	0.261 (0.074) [<0.001]	0.131 (0.111) [0.240]	0.169 (0.095) [0.075]	0.287 (0.085) [<0.001]
Kernel	Triangular	Uniform	Uniform	Triangular	Triangular	Triangular
Bandwidth	mserd	Manual	mserd	mserd	mserd	mserd
Polynomial	Linear	Linear	Linear	Quadratic	Linear	Linear
Covariates					X	
Drop misclassified						X
Observations	872	757	805	878	556	793

Notes: This table reports results of the regression discontinuity that is similar to Equation 2 from page 98 of [Douenne and Fabre \(2022\)](#) and follows the best practices described by [Cattaneo and Titiunik \(2022\)](#). Each column reports a different regression discontinuity specification. For each specification, the rows report the conventional RDD estimate, the bias-corrected RDD estimate, and the bias-corrected RDD estimate with robust confidence interval, respectively. As described in the table, the specifications vary in their kernel, bandwidth selection, local polynomial order, inclusion of covariates, and inclusion of misclassified respondents. Standard errors are in parentheses and p-values are in brackets.

Table 4: Placebo outcomes

	(1)	(2)	(3)	(4)
Conventional	0.077 (0.064) [0.230]	-0.053 (0.055) [0.332]	-0.021 (0.060) [0.728]	-0.187 (0.207) [0.366]
Bias-Corrected	0.073 (0.064) [0.252]	-0.053 (0.055) [0.335]	-0.006 (0.060) [0.919]	-0.252 (0.207) [0.223]
Robust	0.073 (0.077) [0.342]	-0.053 (0.066) [0.425]	-0.006 (0.071) [0.932]	-0.252 (0.243) [0.299]
Kernel	Triangular	Triangular	Triangular	Triangular
Bandwidth	mserd	mserd	mserd	mserd
Polynomial	Linear	Linear	Linear	Linear
Observations	1217	1025	1011	1129

Notes: This table reports results for placebo outcomes using the regression discontinuity design in column 1 of Table 3. The placebo outcomes are an indicator for male, an indicator for young (age less than 30), an indicator for old (age of 65 or greater), and household size, respectively. All specifications use a triangular kernel with optimal bandwidth selection and local linear polynomials. Standard errors are in parentheses and p-values are in brackets.

Table 5: Estimating Equation 2 with heterogeneous effects

	(1)	(2)	(3)	(4)
Predicted winner ($\hat{\Gamma} == 1$)	0.256 (0.065) [<0.001]	0.180 (0.074) [0.015]	0.209 (0.086) [0.015]	0.220 (0.078) [0.005]
Initial tax acceptance (A^0)	0.061 (0.069) [0.379]	-0.065 (0.097) [0.506]	-0.163 (0.128) [0.204]	-0.272 (0.109) [0.013]
Yellow Vests supporter	0.013 (0.059) [0.818]	-0.090 (0.081) [0.268]	-0.075 (0.105) [0.476]	0.026 (0.091) [0.780]
$\hat{\Gamma} \times A^0$	0.150 (0.080) [0.061]	0.378 (0.124) [0.002]	0.460 (0.145) [0.002]	0.618 (0.126) [<0.001]
$\hat{\Gamma} \times$ Yellow Vests supporter	-0.162 (0.070) [0.020]	-0.084 (0.106) [0.430]	-0.081 (0.120) [0.502]	-0.155 (0.106) [0.143]
$\hat{\Gamma} \times G$	0.084 (0.080) [0.292]	0.043 (0.119) [0.719]	-0.017 (0.142) [0.907]	-0.013 (0.123) [0.914]
Kernel	Uniform	Uniform	Uniform	Triangular
Bandwidth	Manual	Manual	Manual	Manual
Polynomial	Global quadratic	Global quadratic	Local linear	Local linear
Interactions with polynomial		X	X	X
Observations	757	757	757	757
R^2	0.401	0.410	0.410	0.440

Notes: This table reports results of estimating the regression in column 3 of Table 4 from [Douenne and Fabre \(2022\)](#) and alternative specifications. All specifications include interactions with three sources of heterogeneity: Yellow Vests supporter, initial tax acceptance (A^0), and initial belief about gain or loss (G). Column 1 replicates column 3 of Table 4 from [Douenne and Fabre \(2022\)](#), which uses a global quadratic of the running variable and the sources of heterogeneity are interacted with only the treatment variable. Column 2 reports results when the heterogeneity indicators are interacted with both the treatment variable and the global quadratic of the running variable. Column 3 reports results similar to column 2 but with the global quadratic replaced by local linear functions of the running variable. Column 4 reports results similar to column 3 but with the uniform kernel replaced by a triangular kernel. Standard errors are in parentheses and p-values are in brackets.