

Information, perceptions, and electoral behaviour of young voters: A randomised controlled experiment[☆]

Bruno P. Carvalho^{a,b}, Cláudia Custódio^c, Benny Geys^d, Diogo Mendes^e, Susana Peralta^{f,*}

^a Department of Economics, Universidad Carlos III de Madrid, Spain

^b ECARES, Université Libre de Bruxelles, Belgium

^c Imperial College London, CEPR, ECGI, United Kingdom

^d BI Norwegian Business School, Norway

^e Stockholm School of Economics, Swedish House of Finance, Sweden

^f Nova School of Business and Economics, Universidade NOVA de Lisboa, Portugal

ARTICLE INFO

JEL classification:

D72

D83

H11

H70

Keywords:

Elections

First-time voters

Performance information

Randomised controlled trial

Negativity bias

ABSTRACT

The way people absorb and process politically relevant information is central to their subsequent political behaviour (in terms of turnout and vote choice). Nonetheless, little is known about how *young* voters – who might be more impressionable than more experienced voters – respond to the provision of such information. In this article, we design a between-subject randomised controlled trial that exposes a sample of university students to positive, neutral or negative information about central government performance before the 2017 Portuguese local elections. We find that young voters update their perceptions more when exposed to negative news. This negativity bias is stronger for first-time voters. We also find that negative information significantly affects turnout of initially undecided young voters. Our results imply that sensitivity to information is heterogeneous and that some young voters may be prone to manipulation through the provision of negative news.

1. Introduction

Assessing whether and how distinct segments of the population react to information about government performance is important. For instance, it can help identify population groups prone to changing their opinions and behaviour when exposed to fake news (Allcott and Gentzkow, 2017; King et al., 2017), and may also enable more targeted and cost-effective information campaigns (Nickerson, 2007a; Bergh et al., 2021). As such, it can have critical implications for electoral democracy and political accountability (DellaVigna and Gentzkow, 2010). In this paper, we focus specifically on how *young* voters react to the provision of political information. Young voters' political behaviour has often been found to differ from the rest of the electorate (Holbein and Hillygus, 2020), either because they turn out less (Smets and

Van Ham, 2013; Foos et al., 2023) or because their preferences diverge from those of older voters (Becker et al., 2017; Geys et al., 2022). While considerable academic attention has recently been awarded to youth voter mobilisation (Nickerson, 2007b; Bhatti et al., 2017; Bergh et al., 2021; Bergh and Christensen, 2022; Foos et al., 2023), much less is known about whether and how young voters process and use political information. We address this question using a randomised controlled trial, which allows us to evaluate the impact of information provision on young voters' perception of incumbent performance as well as their subsequent voting behaviour.

From a theoretical perspective, we take inspiration from recent work arguing that voters deal with information analogous to a Bayesian updating process (DellaVigna and Gentzkow, 2010; Arias et al., 2022;

[☆] This work was funded by Fundação para a Ciência e Tecnologia, Portugal (UIDB/00124/2020, UIDP/00124/2020 and Social Sciences Data Lab, PINFRA/22209/2016), by POR Lisboa and POR Norte (Social Sciences Data Lab, PINFRA/22209/2016). Financial support by MICIN/AEI/10.13039/501100011033 (CEX2021-001181-M), and Comunidad de Madrid (EPUC3M11 (V PRICIT), H2019/HUM-5891), is also gratefully acknowledged. Bruno P. Carvalho and Diogo Mendes also acknowledge the funding provided by Fundação para a Ciência e Tecnologia, Portugal (SFRH/BD/71542/2010, PD/BD/105722/2014, PTDC/IIM-FIN/4177/2014, PTDC/EGE-OGE/28603/2017). We are very grateful for the comments of two anonymous referees, Miguel Portela, the late (and missed) Leib Helland, Jon Fiva, Pedro Freitas and Pedro Vicente, and to a team of researchers who helped us implement the project, led by Matilde Grácio and Sara Almeida. Susana Peralta wrote this paper while visiting *Banco de Portugal*, whose hospitality is gratefully acknowledged. The usual disclaimer applies.

* Corresponding author.

E-mail address: peralta@novasbe.pt (S. Peralta).

¹ Such a negativity bias is discussed, *inter alia*, in Lau (1985) and Baumeister et al. (2001).

Broockman and Kalla, 2022). That is, they adjust their prior beliefs about the likelihood of a certain hypothesis – such as the high or low ‘quality’ of an electoral candidate, or the benefit of turning out to vote – when new information becomes available. Crucially, this new information is often argued to gain in importance when individuals’ prior beliefs are weaker and when less information was available beforehand (Zaller, 1992; DellaVigna and Gentzkow, 2010; Arias et al., 2022). Building on these insights, we hypothesise that such a diminishing marginal influence of information makes young voters particularly persuadable by new information when they vote for the first time or remain undecided close to Election Day (Bergh and Christensen, 2022). Reflecting that individuals are generally more responsive to negative rather than positive information,¹ we thereby expect that young voters display an asymmetric reaction to positive and negative information. As a result, we predict that negative information provided to undecided young voters and neophytes will have the strongest impact on perceptions of incumbent performance as well as subsequent political decisions (i.e. turnout and voting behaviour).

To assess these theoretical propositions, we designed a between-subject experiment that induced exogenous variation in participants’ exposure to information (based on actual articles from mainstream press outlets). Specifically, we exposed 1,799 university students to news articles about central government performance in the week before the 2017 local elections in Portugal. The news had either a positive or a negative tone, and covered different policy domains (including public finance, employment, health care, education, road safety and pensions). A control group was exposed to a neutral news treatment about a non-endangered Portuguese dog breed. Our analysis then proceeds in three steps. We first examine whether and to what extent participants updated their perceptions about the incumbent central government in response to the treatment. Then, we assess whether the treatment affected self-reported turnout and voting behaviour in the 2017 local elections. Finally, we study two sub-groups – first-time voters and undecided voters (i.e., unresolved one week before the election) – to provide novel insights about how respondents from these groups may react to (and use) the information provided differently.²

Our main findings are as follows. First, we find that providing negative (positive) information induces an expected downward (upward) revision of government performance evaluations in the specific policy areas covered by the news pieces. The performance update is equivalent to about one-third of a standard deviation of the mean baseline performance perception, which reflects a substantial impact. This confirms that our information treatment was relevant and perceived correctly by respondents. We then examine individuals’ perceptions of the incumbent government’s overall performance. The results indicate that the provision of negative information has a significant and robust (across specifications) effect on young voters’ overall perception of the incumbent. The magnitude of the effect is at least -0.094 (on a scale from -1 to 1), and corresponds to one-fourth of a standard deviation in this measure. However, we do not find analogous results for positive information (the effect size is less than 0.03). This more pronounced sensitivity to negative information confirms the presence of a *negativity bias* in young voters.

Second, we find that our information treatment has no average treatment effect upon respondents’ turnout or vote choice in the 2017 local elections, which we show is not due to the possibility of treatment dilution (Angrist, 2006). Third, we illustrate that this average treatment

effect masks considerable heterogeneity across respondent groups. The reaction to negative information is smaller for individuals who report a higher interest in politics, while inexperienced (i.e., first-time, eligible) voters exhibit a more pronounced negativity bias and seem to dismiss positive information. This pattern is consistent with a diminishing marginal influence of information upon individuals’ decision-making process (Zaller, 1992; DellaVigna and Gentzkow, 2010; Broockman and Kalla, 2022). Regarding voting behaviour, we find that negative information decreases the likelihood of initially undecided voters casting a blank vote by 13 percentage points, and increases their probability of abstention or voting for opposition parties by just over 30 percentage points. These results extend the findings of previous work showing that negative campaigning may reduce voter turnout (Lau and Rovner, 2009) and influence voter preferences (Kendall et al., 2015).

Our analysis provides three key contributions. The first contribution relates to how electoral decisions are affected by information provision. The potential for information provision to have a heterogeneous impact on government performance perceptions – and subsequent political behaviour – among distinct segments of the population has received only limited attention thus far. To the best of our knowledge, we are first to look at this issue for young voters. Second, a rapidly growing literature studies how young voters can(not) be mobilised to register and/or vote using traditional as well as modern mobilisation tools (Nickerson, 2007b; Bhatti et al., 2017; Bergh et al., 2021; Bergh and Christensen, 2022; Foos et al., 2023). We extend this literature by moving beyond the (important) issue of youth voter mobilisation, and address how information provision affects young voters’ perceptions of political actors and their subsequent voting behaviour. As such, we take one further step towards a broader and more encompassing view of the drivers of youth political activity. Finally, although the existence of negativity bias is well-established across a wide number of contexts, we show that this bias is most prevalent among inexperienced young voters. This is a significant observation given that these voters are much more susceptible to and persuadable by news provision. From a public policy perspective, this observation strongly suggests a need to protect specific subsets of young voters against disruptive and negative (fake) news campaigns.

2. Theoretical background and hypotheses

How – if at all – do voters process and use political information? While many theoretical arguments have been proposed in a rich and varied literature, we take recent work viewing voter reactions to information “in a manner analogous to Bayesian reasoning” as our point of departure (Broockman and Kalla, 2022, p.2). A central premise in Bayesian information processing models is that individuals update their priors (or initial beliefs) about the likelihood of something being true or false when new information becomes available to them. This is not meant to imply that we believe voters actively use Bayes’ rule to update their beliefs. Nevertheless, a process whereby individuals partially – albeit potentially only very minimally – adjust their stance when new information becomes available seems to match well with empirical observations across a range of academic disciplines (DellaVigna and Gentzkow, 2010).

Crucially, a core prediction of Bayesian-inspired models is that new information is deemed more effective or persuasive when individuals’ prior beliefs are weaker and when less information was available beforehand (DellaVigna and Gentzkow, 2010; Arias et al., 2022). This is reminiscent of Zaller’s (1992)’s argument that information will have a lower effect on people with more considerations already in their ‘consideration pool’ (see also Bergh and Christensen, 2022; Broockman and Kalla, 2022). In practical terms, this line of argument implies that information will be characterised by a diminishing marginal effect when more of it is already available. As people learn more about politics (and politicians), any additional information will thus become less influential. This is important since young voters voting the first

² The provision of information as a treatment in a political setting may raise ethical concerns. To this point, we relied exclusively on information provided in the general media, which in Portugal is mostly perceived as independent from political parties and organisations. Furthermore, our research adheres to all ten Principles for Human Subjects Research approved by the APSA Council (see also section A of the Online Appendix) and obtained approval from the ethics committee at [anonymised for review].

time or remaining undecided close to Election Day are likely to have built up less prior political knowledge compared to more experienced young voters (Bergh and Christensen, 2022). As such, we maintain that they will be particularly persuadable by the provision of (new) information. This leads to our first hypothesis:

Hypothesis 1. Young voters update their perceptions of government performance more after exposure to information when they vote for the first time or remain undecided close to Election Day.

Naturally, not all information can be expected to be treated equally. A large literature illustrates, for instance, that individuals (and organisations) are generally more responsive to negative compared to positive performance information (Lau, 1985; Baumeister et al., 2001; Nielsen and Moynihan, 2017; Hong, 2019). Such a *negativity bias* is particularly prevalent in situations where people are actively encouraged to compare across alternative outcomes (such as political candidates in an election). The reason is that explicitly comparative settings shift “individuals’ relative attention towards potential threats rather than opportunities in the payoff distribution” (Kuehnhanss et al., 2017, p.1010). Applying this line of reasoning to our setting, we expect that young voters in general will display an asymmetric reaction to positive and negative information. Nonetheless, when combined with **Hypothesis 1**, we can further specify this prediction. Indeed, combining negativity bias with a higher sensitive to news among undecided young voters and neophytes suggests that the latter subsets of young voters will display a particularly pronounced impact of negative information provision. This leads to our second and third hypotheses:

Hypothesis 2. Young voters update their perceptions of government performance more when exposed to negative rather than positive news.

Hypothesis 3. Young first-time or undecided voters are most responsive to negative rather than positive news.

Finally, changing beliefs need not necessarily become reflected in changing behaviour. This relationship will depend on the elasticity of behaviour with respect to beliefs, which is likely to be lower when initial beliefs are stronger (DellaVigna and Gentzkow, 2010). In other words, when people are more certain about their initial positions, more will be required to move them away from this position and any change in beliefs will have a weaker behavioural impact. In our setting, this implies that any change in young voters’ perceptions of government performance is not guaranteed to affect subsequent political decisions (i.e. turnout and voting behaviour). Yet, given our discussion thus far (esp. **Hypotheses 1** and **3**), any such effects are most likely to materialise among undecided young voters and neophytes. They not only are expected to show a stronger response to new information, but their weaker priors also suggest a higher elasticity of behaviour with respect to beliefs. This leads to our fourth and final hypothesis:

Hypothesis 4. Young first-time or undecided voters are most likely to adjust their political decisions (i.e. turnout and voting behaviour) following exposure to (negative) news.

3. Institutional setting and experimental design

3.1. Institutional setting

The Portuguese government is organised into three levels: central government, municipalities, and civil parishes.³ Municipalities are responsible for local public services such as education, healthcare

³ In addition, the Azores and Madeira archipelagos have elected regional governments.

facilities, and amenities including parks and transportation. The municipal government is composed of an executive branch – the Town Council – and a legislative branch – the Municipal Assembly. All local government representatives are elected simultaneously every four years in nationwide local elections (normally two years after the national elections). All citizens aged 18 and above are eligible to vote.

At the time of our study (2017), there were five main political parties in Portugal that run in both central and local government elections. *Bloco de Esquerda* (BE) and *Coligação Democrática Unitária* (CDU) are the left-most political parties, *Partido Socialista* (PS) is a socialist party on the centre-left, while to the right we have *Partido Social-Democrata* (PSD) and *Centro Democrático e Social* (CDS). Together, these parties obtained 86% of all votes for the Town Council in 2017 (the remainder was won by other parties, independent candidates and groups of citizens), and PS won 159 out of the 308 municipalities. Following the 2015 parliamentary election, PS became the incumbent central government party, with the parliamentary support of the remaining left-wing parties (BE and CDU). Therefore, at the time of our experiment, participants had enough time to form beliefs about the performance of the incumbent central government.

3.2. Experimental design

The 2017 Portuguese local elections took place on October 1st. We conducted three survey rounds between September 18th and October 6th. (i) In the baseline survey, two weeks before the election, we collected information about demographic background, political preferences, awareness and interest, and planned voting behaviour. (ii) In the treatment survey, one week before the election, we randomly exposed participants to factual information bundles about central government performance. This coincided with the campaign period, in which respondents were likely to be exposed to other external information (for a similar approach, see Nickerson, 2007b; Bhatti et al., 2017; Bergh et al., 2021; Foos et al., 2023).⁴ (iii) In the follow-up survey, we collected self-reported voting behaviour (turnout and vote choice) in the week after the election. Additional details can be found in Appendix A, where Figure A.1 shows a timeline, and Table A.1 summarises the information collected in each survey.⁵ We implemented the experiment at two universities in Lisbon (Nova SBE and ISCAL).⁶ All surveys were administered paper-based, in a classroom, either at the beginning or the end of a lecture. The implementation protocol included a set of instructions that was read aloud, before the questionnaires were distributed. Participants were invited to take part in a study about the upcoming election and could opt out by not completing the survey(s). They were asked to construct a unique, anonymous identifier based on a sequence of numbers from their birth date and phone numbers, used in all survey rounds to allow merging. We provided no monetary incentives.

Since most of our sample were undergraduate students and the legal voting age in Portugal is 18 years old, about 41% of respondents were first-time voters and few would have developed strong voting experience. This is important to test for heterogeneous effects of information provision related to previous voting experience.⁷

⁴ The official campaign period starts two weeks before the election date and ends two days prior to the election, as defined by Law 1/2001.

⁵ The experiment was pre-registered at <https://www.socialscienceregistry.org/trials/2539>. As indicated in this registration, we targeted a sample of young voters who had never voted before as well as more experienced voters, such as to allow heterogeneity analysis across these two subsamples.

⁶ At Nova SBE, each survey was implemented in a different round as described above. Owing to the academic calendar at ISCAL, we collapsed the baseline and treatment surveys into one, implemented in the week preceding the election.

⁷ Testing for heterogeneity induces concerns regarding sample sizes and statistical power. Appendix Table A.2 shows that our sample size enables us

We should note at this point that while classroom experiments such as ours have been used to gain novel insights about young voters' political behaviour (e.g., [Bennion and Nickerson, 2016](#)), university students are obviously not representative of young people in general. They are, for instance, characterised by higher income and education levels and having parents with higher education and employment levels (more details below), all of which make them more likely to become politically active than young voters outside universities ([Smets and Van Ham, 2013](#)). Hence, one might worry about the external validity of a student sample and its ability to generalise to the population of young voters. A large-scale field experiment where young people are recruited also beyond the university setting would be more beneficial from this perspective. Yet, the potential ethical implications of such a field experiment would also be much larger, particularly when the expectation is that people update their beliefs after news provision and adjust their electoral behaviour on this basis. In the end, we opted for a student sample as politically attentive/interested young individuals offer a best-case scenario to study the potential impact of information provision on their political choices and behaviours (for a similar argument in a different setting, see [Arceneaux and Nickerson, 2009](#); [Bergh et al., 2021](#); [Bergh and Christensen, 2022](#)). Null effects in our setting would indeed make it highly unlikely that any effects materialise among less politically interested young people.

3.2.1. Information treatment

We employed a between-subjects design, relying on six different policy areas – public finance, healthcare, pensions, education, youth employment, and road safety – to reduce the likelihood that our results are influenced by a specific one. Each participant was randomly exposed to either positive or negative information about central government performance in two of these policy areas. The treatment consisted of two actual news articles, on two different policy areas, to increase the likelihood of the treatment being informative for respondents. The articles were selected from mainstream outlets – two newspapers and one news radio – according to the following criteria: (i) recentness, to guarantee timeliness, relevance and accuracy; (ii) availability of positive and negative news on the same policy area from the same outlet, with similar dates, to the extent possible. [Table 2](#) lists the sources for each article and the respective date, ranging from October 14, 2016, to September 24, 2017.

The effectiveness of the treatment hinges upon (i) the subjects' trust on the news outlets; (ii) the perception of the information as fair and unbiased; and (iii) the subjects' interpretation of the treatments as positive (resp., negative). In terms of (i), note that the selected outlets are among the most trusted in Portugal, with scores of 7.19/10 for *Expresso* (the highest amongst newspapers in Portugal), 7.07/10 for *Diário de Notícias*, and 7.04/10 for *Público*, according to *The Reuters Institute at the University of Oxford* ([Newman et al., 2019](#)). The most trusted outlet is included in *all* the treatments.

Regarding (ii), we note that the variety of outlets tackles possible concerns of different political leanings. Moreover, Portuguese newspapers are notable for a very low level of ideological bias. The average response to the question “The political orientation of the most prominent journalists is well-known to the public” in the *European Media Systems Survey 2013* is equal to 3.6 in Portugal, 1.14 standard deviations below the mean of 5.3 for the 34 countries surveyed (on a scale of 1 to 10). This data also shows that *Expresso* and *Diário de Notícias* are perceived by the respondents as locating in the centre of the political spectrum, i.e., they get average scores between 5.3 and

6.2 (where 0 means left and 10 means right), aligned with the average value of this variable, which is 5.6 for the 34 countries.⁸ We further highlight that Portuguese newspapers never publicly endorse political candidates (unlike in, say, the UK or US).⁹

Finally, to confirm that the negative and positive treatments were perceived as such by the subjects, (iii), we conducted a survey in 2022 among students at the same university. Each respondent was presented with a random draw of 4 to 5 articles used in the original experiment. Subjects were then asked to rank the performance of the government in office at the time of the news articles in each specific policy area covered in the articles (on a scale between –10 and 10, with a neutral option). The results confirm that the news articles are perceived as expected, and that the absolute value of negative evaluations is equivalent to that of the positive evaluations. A more detailed account of this validation survey and the respective results is presented in [Appendix B](#), and it confirms that the news articles are perceived as expected. Moreover, the absolute value of negative evaluations is in line with that of the positive ones.¹⁰

We created six positive, six negative and one neutral information bundle (13 versions in total), whose composition in terms of policies is detailed in [Table 1](#). Using twelve different information bundles reduces the likelihood of contagion across different treatment groups, an important concern given that the surveys were administered in class, and control and treatment subjects sit side-by-side in close physical space (see [Section 3.2.2](#)). This approach also mitigates the possibility that the source of the material or any perceived difference in the degree of negativity/positivity of a specific news item influences the information treatment effects. The control group articles (with similar length and format as the treatment ones) were policy neutral, about a non-endangered Portuguese dog breed not targeted by public policies (i.e. *Serra da Estrela*). In all cases, news articles were edited to avoid salient visual differences across bundles, and they included the headline (in larger bold font), an abridged version of the original text, and a graph or a picture. In the first article, we included a graph created from actual data to support the information conveyed in the text, visually similar across positive and negative bundles. The second article included a picture related to the policy area, which was the same for the positive and negative treatments; see examples in [Figures A.4 and A.5](#) in the appendix.

We treat participants with information about the *central* (as opposed to *local*) government at the time of *local* elections. This choice aims at mitigating the risk of dilution, since general elections involve the dissemination of large quantities of information about central government policies. Even so, this information remains relevant to voters for at least three reasons. First, politicians' partisanship correlates in terms of expected policies across different government levels ([Geys and Vermeir, 2014](#); [Schönhage and Geys, 2020](#)). Second, voters may use local elections to show (dis)approval of the central government ([Marien et al., 2015](#)). Third, [Ashworth et al. \(2018\)](#) show that incumbents' electoral outcomes are affected by shocks outside their control – such as hurricanes or policy outcomes at other levels of government – which provide rational voters with opportunities to infer information about incumbents. We return to this in our validity checks in [Section 4.3](#).

3.2.2. Sampling and randomisation

All surveys were administered by the research team in 71 classes (attended by 4 to 120 students). In-class implementation allowed us to conduct the field work in a short time period around the election.

to detect a 5 percentage point change in young voters' opinions about central government performance with 77% (64%) certainty at 90% (95%) confidence level. Similar tests were also performed with comparable results for our other outcome variables (i.e. voter turnout and party support; full details upon request).

⁸ Publicly available data, retrieved from <https://www.mediasystemsineuro.pe.org/emss2013.htm> ([Popescu et al., 2013](#)).

⁹ The relative lack of polarisation of Portuguese news outlets, when compared to other European countries, is confirmed by [Santana Pereira and Nina \(2016\)](#), [Magalhães \(2009\)](#) and [Graça \(2017\)](#).

¹⁰ We thank an anonymous referee for suggesting this survey.

Table 1
Percentage of respondents by survey type.

Group	Survey no.	News subjects		% of respondents	
		1st article	2nd article	by survey	by group
T. Negative	1	Public Finance	Pensions	5.84	32.91
	2	Public Finance	Youth Employment	5.50	
	3	Public Finance	Road Safety	4.95	
	4	National Health Service	Youth Employment	5.67	
	5	National Health Service	Education	5.11	
	6	National Health Service	Road Safety	5.84	
T. Positive	7	Public Finance	Pensions	5.67	33.24
	8	Public Finance	Youth Employment	5.84	
	9	Public Finance	Road Safety	5.89	
	10	National Health Service	Youth Employment	4.84	
	11	National Health Service	Education	5.17	
	12	National Health Service	Road Safety	5.84	
Control	13	Portuguese Dog Breed		–	33.85

Notes: Breakdown of news topics covered in the informational treatment bundles for all surveys. The last two columns report the percentage of respondents per survey type and by treatment group, respectively.

Table 2
Sources and dates of news treatment.

Group	Subject	Outlet	Date
T. Negative	Public Finance	Expresso	July 13 2017
	Pensions	Expresso	November 20 2016
	Youth Employment	Expresso	February 8 2017
	Road Safety	Diário de Notícias	July 3 2017
	National Health Service	Expresso	September 24 2017
	Education	TSF	September 8 2017
T. Positive	Public Finance	Expresso	July 12 2017
	Pensions	Expresso	October 26 2016
	Youth Employment	Expresso	December 19 2016
	Road Safety	Diário de Notícias	April 20 2017
	National Health Service	Expresso	October 14 2016
	Education	TSF	September 6 2017
Control	Portuguese Dog Breed	Público	October 14 2015

Notes: Breakdown of news topics covered in the informational treatment bundles for all surveys.

Although the baseline and follow-up surveys are relatively straightforward, the treatment survey required additional care. To guarantee randomisation at the individual level, we organised the treatment surveys in blocks following the sequence: one control, two random positive, one control, and two random negative treatment surveys. These blocks of six surveys were distributed in class, row-by-row, with the aim of ensuring adequate randomisation (balancing).¹¹

To assess the ex-post validity of the randomisation protocol, we first analyse the distribution of answers per survey type. This is shown in Table 1. As expected, the three groups comprise about one-third of the sample, and the distribution by survey type is close to uniform. Second, we conduct a series of comparison tests by treatment groups (see Fig. 1). We find no statistically significant differences in background characteristics across the three groups (positive and negative treatments, and control group), including for the ex-ante perception of general government performance (as defined in more detail below). The balance also holds for the subsamples of first-time and initially undecided voters (see Appendix Figure C.2).

In spite of this successful randomisation, one possible concern with in-class university-based experiments is that students may share newly obtained knowledge with their peers — inside as well as outside the classroom. For instance, students who received the treatment survey earlier may discuss this with peers in other classes (some of whom may participate in the treatment on later days). Such within- and across-class spillovers may skew estimates of treatment effects. A number

¹¹ At least two alternative approaches can be envisaged: randomisation at the class level (unfeasible owing to the very different class sizes), or making use of seat maps to alternate survey types (unfeasible owing to the absence of pre-defined seats in either school).

of aspects within our design help mitigate such concerns. First, such spillovers are much less likely across the two schools used in our experiment. Hence, similarity of findings across both samples would provide at least suggestive evidence that intra-university spillovers may only be a major concern when such spillovers are of closely comparable magnitude in both settings (which seems unlikely given the different nature of both schools). Second, as mentioned, our sample includes widely differing class sizes, which could affect both within-class spillovers (as smaller classes may have closer student interactions) and extra-class spillovers (as larger classes have more students who could inform peers). Hence, similarity of findings across class sizes would again be suggestive that spillovers are not a major concern (unless they are of closely comparable magnitude across class sizes) – we tackle this in a robustness check. Finally, since credibly detecting and quantifying potential spillover effects is notoriously difficult, we should also point out that they are most likely to lead to attenuation bias in our setting. The reason is that respondents would at least to some extent become exposed to information about both positive and negative outcomes. Hence, we would under-estimate the effects of interest.

4. Data and methodology

4.1. Data

After discarding respondents who did not answer all three survey rounds, our sample comprises 1,799 participants.¹² The final sample

¹² The attrition rate between survey rounds is low (13% at Nova SBE and 7% at ISCAL), and we find no evidence of differential attrition rates

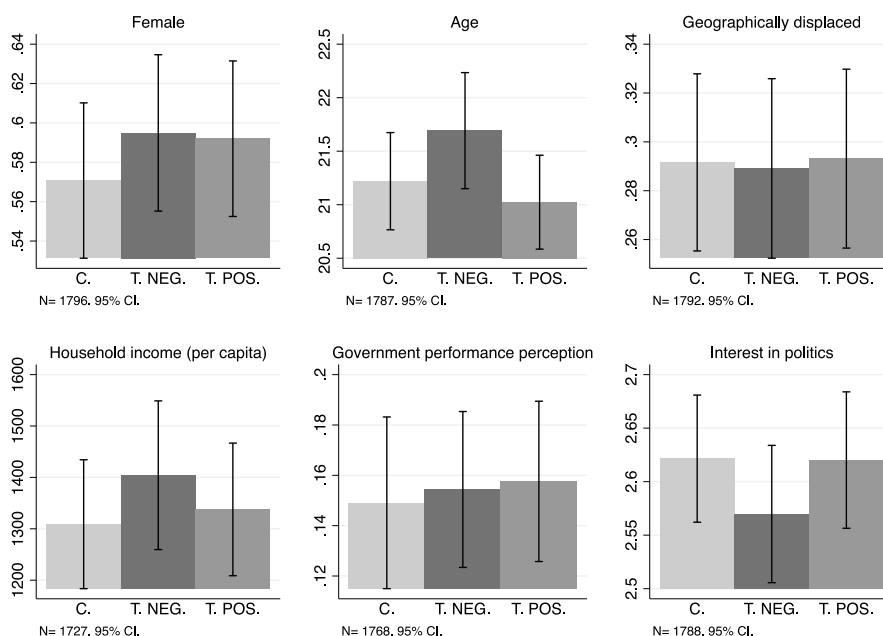


Fig. 1. Balance Tests.

Notes: Each graph plots the average (and its 95% confidence interval) of a covariate by treatment group. Female and Geographically displaced are dummy variables; Age and Household income are discrete counts of their corresponding continuous variables. General government performance (baseline) is a 5-level scale, coded between -1 and 1, that measures respondents' ex-ante central government performance perception. Interest in politics was rated on a 4-level Likert-type scale.

includes subjects from 170 of the 308 Portuguese municipalities (see Appendix Figure C.1). Table 3 presents descriptive statistics across socio-demographics (Panel A), political knowledge and interest (Panel B), and electoral behaviour in the 2017 local elections (Panel C). For a more detailed list of variables, see Table C.1 in Appendix.

When compared to the 2015 Portuguese Electoral Study (EEPT) (Lobo et al., 2015), our subjects are similar in political left-right positioning and interest in politics. Yet, they are younger, richer, live in larger households, and have more educated parents, who are less likely to be unemployed. Such differences decrease substantially when we restrict the EEPT sample to respondents aged under 35 (as in our sample). Importantly, self-reported voting behaviour in our sample is in line with that observed in the EEPT (full details in Appendix Table C.3).

Most of the sample is composed of Bachelor students with an average age of 21.3 years. The age distribution highlights a small percentage of respondents not eligible to vote in the 2017 local elections (5%) and a large percentage (47%) of inexperienced voters (i.e. below the legal voting age at the time of the last previous election — Presidential election in 2016). This implies that just over 41% of our sample is a first-time voter in the 2017 local elections. Approximately 60% of the sample is female, and most respondents come from well-educated households. Panel B shows that respondents had, on average, a positive overall perception of the government before the information treatment. This confirms contemporaneous polling by pollsters *Eurosondagem* (August 2017; sample size 1,011; 49% rated prime minister performance as good, while 30% felt the same about the government as a whole) and *Aximage* (September 2017; sample size 597; almost 50% believed that the government performance had exceeded their expectations).

Most participants (95%) were aware that the local election was about to take place, and 76% planned to vote at the time of the baseline survey, while 18% were undecided.¹³ This provides important variation

across treatment and control groups, nor between the positive and negative information treatments.

¹³ We collected voting intentions for those who intended to vote ('Which candidate do you plan to vote for?') and also for the remaining participants

to assess how initially undecided voters responded to the provision of information. The last four rows of Panel C describe respondents' self-reported electoral behaviour in the 2017 local elections, collected during the last survey round. We observe that 64% of respondents claimed to have voted. This is above the official national turnout rate of 55%, and amounts to a self-reporting bias of circa 9%, which is below the average of 13% observed in a global sample of post-election surveys (Selb and Munzert, 2013). We also find that 45% of those who voted (or would have voted) chose one of the parties holding the majority in the national government.

4.2. Identification and outcome variables

We evaluate whether information provision (i) changes respondents' perception about central government performance (policy-specific and general) and (ii) affects respondents' subsequent voting behaviour.

The following baseline specification addresses the first goal:

$$Performance_i = \alpha + \beta_1 T. neg_i + \beta_2 T. pos_i + \gamma X_i + \varepsilon_i \quad (1)$$

The dependent variable – performance_i – measures respondent *i*'s subjective assessment of government performance both in general terms and in specific policy areas. General government performance was assessed using a scale from 'Very Unsatisfactory' to 'Very Satisfactory', without the possibility of a neutral answer, converted into a [-1, 1] range. These questions were included in both the baseline and treatment surveys, so that we can study the individual-level change in perceived government performance induced by the information treatment.¹⁴ Performance in *specific policy areas* was also surveyed before

(‘If you did vote, which candidate would you vote for with the highest probability?’). In both cases, respondents could choose a ‘Prefer not to answer’ option.

¹⁴ These questions are shown in Figure A.2. The baseline ISCAL questionnaire included the option ‘I don't know/cannot evaluate’, which we coded as zero. We show in the robustness checks that our results are unchanged by removing these observations. A more detailed discussion is provided in Subsection E.1 of the appendix.

Table 3
Descriptive statistics.

PANEL A: socio-demographic characteristics					
	Obs.	Mean	St. dev.	Min	Max
Female	1,796	0.59	0.49	0	1
Age	1,787	21.31	5.97	16	54
Number of people in household	1,727	3.58	1.15	1	11
Monthly per capita household income	1,727	1349	1624	111	17,677
At least one parent attained higher education	1,748	0.46	0.5	0	1
At least one parent is civil servant	1,760	0.27	0.44	0	1
At least one parent is unemployed	1,799	0.16	0.37	0	1
Geographically displaced	1,792	0.29	0.45	0	1
PANEL B: Interest and knowledge in politics					
	Obs.	Mean	St. dev.	Min	Max
Interest in politics	1,788	2.6	0.78	1	4
Registered with political party	1,774	0.06	0.24	0	1
Political spectrum: Left	1,771	0.22	0.42	0	1
Political spectrum: Centre	1,771	0.14	0.35	0	1
Position on spectrum influenced by:					
Own ideas	1,716	0.6	0.49	0	1
Family and friends' opinion	1,716	0.4	0.49	0	1
Opinion about party leaders/politicians	1,716	0.22	0.41	0	1
Voted in last election (2016)	1,784	0.39	0.49	0	1
Not eligible to vote in last election (2016) (under 18)	1,784	0.47	0.5	0	1
Government performance perception (baseline)	1,768	0.15	0.4	-1	1
Government performance perception (post-treatment)	1,731	0.16	0.4	-1	1
PANEL C: 2017 Electoral behaviour					
	Obs.	Mean	St. dev.	Min	Max
Aware of election	1,787	0.95	0.21	0	1
Intends to vote in election (turnout, baseline)	1,799	0.76	0.42	0	1
Not eligible to vote (under 18)	1,799	0.05	0.22	0	1
Allowed to vote for first time (first election)	1,797	0.41	0.49	0	1
Intends to vote for support parties (baseline)	1,328	0.32	0.46	0	1
Intends to vote for other parties (baseline)	1,328	0.33	0.47	0	1
Intends to vote blank (baseline)	1,328	0.12	0.33	0	1
Undecided for whom to vote (undecided)	1,681	0.18	0.39	0	1
Voted in election	1,797	0.64	0.48	0	1
Voted/would have voted for:					
Support parties	1,452	0.45	0.5	0	1
Other parties	1,452	0.4	0.49	0	1
Blank	1,452	0.15	0.36	0	1

Notes: Based on data collected by the authors during the three survey rounds.

and after the information treatment. In the second round, subjects were asked whether and how much the information changed their performance perception in the surveyed policy area using a five-point scale ranging from ‘Yes, it improved a lot’ to ‘Yes, it worsened a lot’ (including a neutral alternative), as shown in Figure A.3 in appendix.

The binary variables T.Neg_{*i*} and T.Pos_{*i*} indicate the treatment received by subject *i* (the reference category is the control group). All specifications include class-level dummies to account for implementation idiosyncrasies, including the precise date. Some specifications include the baseline value of the outcome variable, a vector of individual socio-demographic and political variables *X_i* – indicators for the university, gender, whether the student lived away from her parents (‘geographically displaced student’), parental unemployment, self-identification as Catholic, self-position on left–right political spectrum, ability to determine this position, reliance on family and friends to form own opinion, perception of politics as ‘too complicated’, belief that the politician in power makes a difference to policy outcomes, together with age, household size, monthly net income, and interest in politics — as additional controls.¹⁵ Standard errors are robust to heteroskedasticity.

¹⁵ Including covariates in randomised controlled trials may improve the efficiency of the estimate of treatment effects, via the reduction of the residual variance (Robinson and Jewell, 1991).

Next, we analyse voting behaviour using the following reduced-form specification:

$$\text{Voting}_i = \alpha + \beta_1 \text{T. neg}_i + \beta_2 \text{T. pos}_i + \gamma Z_i + \eta_i \tag{2}$$

where voting_{*i*} measures self-reported turnout or vote choice (including blank votes), recorded during the post-election survey. The treatment variables T. neg_{*i*} and T. pos_{*i*} are defined as before. The vector of controls *Z_i* covers pre-treatment survey desired turnout and vote choice, government performance perception, and the set of socio-demographic and political variables in Eq. (1). Standard errors are robust to heteroskedasticity.

4.3. Validity of the experiment

Before we proceed to the results, we present a set of preliminary analyses that confirm the validity of our experimental setup.

The first validity check pertains to the extent to which national-level policies may be relevant during local elections (Geys and Vermeir, 2014; Marien et al., 2015; Daniele et al., 2020). We asked respondents in the post-election survey to identify the determinants of their vote choice (multiple answers allowed). The list of alternatives (Appendix Table D.1) included several reasons indicating a link between local and central politics: 25% (22%) of the respondents indicate they are willing to use their local vote to support (penalise) the central government, and

Table 4
Policy-specific government performance perception.

	Specific gov. perf.				
	(1)	(2)	(3)	(4)	(5)
T. neg.	-0.113*** [-8.35]	-0.113*** [-8.28]		-0.110*** [-8.23]	-0.117*** [-8.20]
T. pos.	0.131*** [10.09]		0.130*** [10.04]	0.132*** [10.34]	0.117*** [8.62]
Gov. perf. (baseline)				0.065*** [4.62]	0.058*** [3.85]
Controls	No	No	No	No	Yes
Observations	1,736	1,159	1,169	1,710	1,503
Adjusted R-squared	0.156	0.050	0.082	0.164	0.156
Coef. equality (F-Test)	282.1	–	–	278.81	224.38
Coef. equality (p-value)	0.000	–	–	0.000	0.000

Notes: The outcome variable is the government performance perception associated with the policy areas covered in the treatment. In columns 2 and 3, we restrict the sample to compare separately the Negative and Positive treatment groups with the Control group. All regressions include dummies for the class of implementation; standard errors are robust and *t*-statistics are reported in brackets. Controls are detailed in Section 4.2. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

38% to reward the local candidate with the best connections to the central government. This confirms that central government performance is relevant to local outcomes for a sizeable share of our respondent pool.¹⁶

The second validity check pertains to respondents' processing the information provided in our experiment. In particular, we show that the information treatment was effective since it leads voters to update their performance perception on the covered policy areas. More precisely, Table 4 displays different specifications of Eq. (1), using the average perception of government performance in the two specific policy areas included in the treatment as outcomes.¹⁷ Columns 2 and 3 exclude the respondents receiving positive and negative treatment, respectively, while (4) includes respondents' pre-treatment performance perception, and in (5) we add further socio-demographic and political controls. All specifications confirm that our information bundles were effective in changing participants' perception of the performance of the government in the policy areas featured in the news article in the expected direction.

It is also important to confirm that prior negative opinions about the policy topic do not drive the result of the experience — this is done in columns (1) and (2) of Table 5.¹⁸ We compare the negative with the positive treatment, and exclude the control group because its information is unrelated with specific policy areas. A further validity check exploits the idea that updating based on the provided information should be lower among those with higher prior interest in the policy areas addressed in the information bundles, i.e., more knowledgeable voters should be less responsive to the treatments. Columns (3) and (4) of 5 show that *more* informed young voters (those who report having a higher interest in politics) indeed react *less* to negative information. We also find descriptive evidence that more informed voters are less likely to report having the intention to vote and end up not voting.¹⁹

Taken together, these results show that our sample contains respondents that are willing to use local elections to send a message to the

central government, and that update their perceptions in line with the information provided in the experiment (suggesting our treatment was effective). We now move to our main results.

5. Results

5.1. Government general performance perception

We evaluate whether the positive and negative news treatments changed the perception of central government general performance pre- to post-treatment. On average, the unconditional change in general government performance perceptions is 0.059 (resp., -0.081) for subjects who received the positive (resp., negative) treatment. We observe a small increase (0.03) in the control group, suggesting that external events may have positively affected the general perception of the government between the two rounds. However, any general development between surveys does not account for the observed differences across our positive and negative treatments.

The regression results from estimating Eq. (1) in Table 6 show that perceptions shifted substantially more in the treatment groups. Column 1 reports the results without controls, column 2 (3) excludes the positive (negative) treatment groups from the sample, column 4 controls for respondents' pre-treatment perception of general government performance, and column 5 adds further socio-demographic and political controls. Finally, column 6 uses the *change* in the performance measure as the outcome variable.

The positive information treatment effect is positive but small (at most 0.033), and only statistically significant in some specifications. This contrasts with the large and statistically significant negative treatment effect of between -0.094 and -0.121. Note that this is not driven by a lack of information updating resulting from the positive treatment, as shown in Section 4.3. The coefficients for $T.Neg_i$ in columns 1 to 5 are equal to approximately one-fourth of the standard deviation of pre-treatment general government performance perceptions, which is a meaningful impact. These results confirm our Hypothesis 2.

These results provide evidence of a *negativity bias* in general government performance updating among young voters. As shown in Online Appendix Section E.4, this conclusion continues to hold even after standard corrections for multiple hypothesis testing.²⁰ The greater weight placed on negative – relative to positive – information has previously been reported in, for instance, (Lau, 1985). The fact that we used six positive and six (corresponding) negative information treatments

¹⁶ We also ran Eq. (2) including indicator variables for respondents' willingness to reward (resp., punish) the central government during local elections. The results show a statistically significant association, which further confirms the relevance of central government performance during local elections in Portugal. Results are available upon request.

¹⁷ Results are robust to considering only the policy area shown last, or creating a panel of the two answers for each individual, as shown in Table E.1 in appendix.

¹⁸ An observation in these regressions is a pair (individual, article), i.e., there are two observations per respondent.

¹⁹ This is shown in Table C.2, where we compare socio-demographics and political preference variables of respondents who reported they intended to vote in the baseline survey, but state in the follow-up they have not voted.

²⁰ We thank an anonymous referee for the suggestion to implement such validation checks.

Table 5
Prior opinions, interest in politics and performance perception.

	Opinion		Interest	
	Specific gov. perf.		Specific gov. perf.	
	(1)	(2)	(3)	(4)
T. neg.	-0.237*** [-8.68]	-0.226*** [-7.81]	-0.346*** [-10.00]	-0.331*** [-8.90]
T. neg. x Neg. opinion (subject specific)	-0.019 [-0.60]	-0.020 [-0.59]		
T. neg. x High interest (subject specific)			0.131*** [2.60]	0.120** [2.26]
Neg. opinion (subject specific)	-0.074*** [-3.23]	-0.073*** [-2.95]		
High interest (subject specific)			-0.021 [-0.59]	-0.015 [-0.39]
Second article	-0.009 [-0.66]	-0.007 [-0.49]	-0.060** [-2.32]	-0.053** [-1.99]
Gov. perf. (baseline)		0.050** [2.45]		0.078** [2.17]
Controls	No	Yes	No	Yes
Observations	2,211	1,929	616	555
Adjusted R-squared	0.149	0.156	0.193	0.190

Notes: The outcome variables are government performance perception associated with the policy areas covered in the treatment. Negative opinion and High interest are (subject-specific) dummy variables equal to 1 when a respondent reports a negative pre-treatment performance perception (or high interest) for the policy area of the treatment. Second article is equal to 1 when the specific performance perception refers to an article presented as the second piece of information in the treatment bundle. All regressions include dummies for the class of implementation; standard errors are robust and *t*-statistics are reported in brackets. Controls are detailed in Section 4.2. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6
General government performance perception.

	General gov. perf.					Δ (General gov. perf.)
	(1)	(2)	(3)	(4)	(5)	(6)
T. neg.	-0.100*** [-4.13]	-0.094*** [-3.84]		-0.106*** [-6.17]	-0.116*** [-6.39]	-0.121*** [-6.11]
T. pos.	0.033 [1.47]		0.031 [1.37]	0.033** [1.98]	0.029* [1.65]	0.023 [1.15]
Gov. perf. (baseline)				0.666*** [26.79]	0.635*** [23.62]	
Controls	No	No	No	No	Yes	Yes
Observations	1,731	1,151	1,165	1,708	1,509	1,509
Adjusted R-squared	0.018	0.014	0.005	0.481	0.497	0.045
Coef. equality (F-Test)	32.11	-	-	67.5	63.78	51.6
Coef. equality (p-value)	0.000	-	-	0.000	0.000	0.000

Notes: The outcome variable is either general government performance perception post-treatment (columns 1 to 5) or change in general government performance perception (column 6), calculated as the difference between the post-treatment and baseline perception measures. In columns 2 and 3, we restrict the sample to compare separately the Negative and Positive treatment groups with the Control group. All regressions include dummies for the class of implementation; standard errors are robust and *t*-statistics are reported in brackets. Controls are detailed in Section 4.2. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

excludes the possibility that our results are influenced by the particularities of the information and policy areas in the different treatments. Furthermore, Fig. 1 shows that there are no observable differences across the positive and negative treatment (and control) groups, hence the findings cannot be the result of *ex-ante* differences across treatment groups.

5.2. Voting behaviour

We now examine the extent to which the exogenously induced update in performance perceptions translates into the decision to cast a vote, to vote for candidates from the parties supporting the government (PS, BE and CDU; henceforth ‘Support parties’), to vote for other candidates (PSD, CDS, other smaller parties or independent candidates; henceforth ‘Other parties’), or to cast a blank vote. The results are summarised in Table 7. We report two specifications, one which controls only for respondents’ pre-treatment government performance

perception, and a second one that includes socio-demographic and political controls and municipality-level fixed effects. All regressions include class dummies and the baseline vote intention.

Columns 1 and 2 of Table 7 show that neither treatment has a statistically significant impact on voter turnout, although the point estimates are negative. Information provision may have opposite effects on turnout: negative information, for instance, can keep individuals from voting if it harms their trust in politics and politicians, or increase the desire to vote if voting is perceived as a tool to change the status quo. Our non-significant results may therefore indicate the absence of a treatment effect, or be due to these effects cancelling each other out.²¹

²¹ The research design allows us to investigate which voters reported the intention to vote in the baseline survey and stated they had not voted in the follow-up survey. Table C.2 reports the difference in means in a set of socio-demographic and political preferences’ variables between these respondents

Table 7
Turnout and voting decision (intention to treat).

	Turnout		Voting decision					
			Support parties		Other parties		Blank	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T. neg.	-0.026 [-1.18]	-0.013 [-0.52]	0.018 [0.69]	0.010 [0.31]	0.012 [0.47]	0.025 [0.82]	-0.008 [-0.40]	-0.011 [-0.50]
T. pos.	-0.026 [-1.17]	-0.034 [-1.35]	0.012 [0.43]	0.006 [0.17]	0.021 [0.81]	0.013 [0.42]	-0.003 [-0.13]	0.006 [0.25]
Turnout (baseline)	0.688*** [34.11]	0.646*** [23.93]						
Vote support parties (baseline)			0.606*** [24.02]	0.520*** [14.19]				
Vote other parties (baseline)					0.634*** [24.82]	0.504*** [13.33]		
Vote blank (baseline)							0.602*** [14.64]	0.514*** [9.91]
Gov. perf. (baseline)	0.014 [0.62]	0.014 [0.53]	0.196*** [7.64]	0.139*** [4.49]	-0.159*** [-6.28]	-0.125*** [-4.38]	0.001 [0.07]	-0.011 [-0.46]
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Municipality dummies	No	Yes	No	Yes	No	Yes	No	Yes
Observations	1,768	1,545	1,191	1,074	1,191	1,074	1,191	1,074
Adjusted R-squared	0.381	0.421	0.427	0.454	0.475	0.521	0.342	0.344

Notes: The outcome variables are either a dummy variable equal to 1 when a respondent reported voting in the election (columns 1 and 2) or dummy variables equal to 1 when respondents reported a vote (intention) for parties in the central government (Support parties), for the opposition (Other parties) or for no one (Blank). The corresponding baseline variables (collected before the election) are included as controls. All regressions include dummies for the class of implementation; standard errors are robust and *t*-statistics are reported in brackets. Controls include municipal-level dummies and those detailed in Section 4.2. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In the remaining columns we turn to respondents' self-reported vote choice. We aggregate the reported party choice of both voters and non-voters, as we asked the latter who they would have voted for.²² Similar to turnout, we find no average treatment effect on the likelihood of voting for the parties supporting the central government, on the likelihood of voting for other parties, or on the probability of casting a blank ballot. As expected, individuals' baseline voting intention is positively correlated with the final decision, and a higher baseline government performance perception has a positive association with voting for the parties in the central government (columns 3 and 4), whereas it reduces the likelihood of voting for other parties (columns 5 and 6).

We report on a number of robustness checks in the Online Appendix. First, we show that our results are not driven by the assumptions used to handle neutral options in our government performance measures (Subsection E.1). Next, in Subsection E.2, we illustrate the consistency of our findings when excluding subsamples that might be considered influential ex-ante (e.g., we keep only respondents who believe ideology is shared among central and local politicians from the same party; we drop respondents from Nova SBE, non-voters, respondents from municipalities of the Portuguese islands or from municipalities with less than 10 individuals in our sample). We also show that the results are the same when we interact the treatment indicators with below-median class size dummies. If anything, the negativity bias is slightly more nuanced in small classes. The similarity of findings across class sizes is suggestive that spillovers are not a major concern.²³ In Subsection E.3, we find unchanged results when estimating our baseline regressions using (ordered) logit models.

and the remaining sample. We find individuals who first said they would vote and then did not on average have lower income and their parents are less likely to have higher education. They are more often geographically displaced students, with lower interest in politics.

²² Non-voters include participants who did not cast a vote on the election day, or those that were not in legal age to vote. The results are qualitatively similar when we drop non-voters (full details in the Appendix Subsection E.2).

²³ The median size class is 20 students

There are three possible explanations for the absence of a significant average treatment effect in Table 7. First, our news treatment is best seen as an intention to treat, which may be associated with substantial *treatment dilution* given that it was implemented during the electoral campaign. This would imply that our results are lower bound estimates. We deal with this issue by restricting our analysis to the respondents that changed their perception about government performance as a result of the treatment, and use treatment assignment to instrument the respondents' updated perceptions of general government performance (following Angrist, 2006). The detailed results are reported in Online Appendix Subsection F, and also fail to show statistically significant average treatment effects.

Second, it may be that voters in our setting do not reward/punish local candidates based on their perception of central government performance, but rather focus on local-level issues and candidates. As mentioned above, a considerable share of our respondents do in fact take into account the central government when deciding about their vote in local elections (Marien et al., 2015). Naturally, it could still be that this share is too small to generate a notable average treatment effect.

Finally, different types of voters may incorporate their updated central government perceptions differently into their local voting decisions. That is, the negligible average treatment effect in Table 7 may mask considerable heterogeneity across respondent groups. We consider such heterogeneity in the next section.

6. Undecided and inexperienced voters

We now explore heterogeneous effects in two subsamples: respondents who had not yet decided their vote in the week before the election ('Undecided') and first-time eligible voters ('First Election'), i.e., we test Hypotheses 1, 3, and 4. The former responded 'I do not know yet' in the baseline survey question about who they plan to vote for. The latter are identified based on their birth date and the timing of the January 2016 presidential election (the last election prior to the 2017 local elections). The variable 'First Election' equals 1 when respondents were legally entitled to vote for the first time in the 2017 election. Both types of respondents may be more sensitive to information, either because they need to decide between available alternatives, or because they have

Table 8
Undecided and first election voters and performance perception.

	Specific government performance				General government performance			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T. neg.	-0.104*** [-6.70]	-0.112*** [-6.91]	-0.114*** [-6.55]	-0.133*** [-7.05]	-0.085*** [-3.06]	-0.114*** [-5.56]	-0.055* [-1.67]	-0.080*** [-3.36]
T. pos.	0.121*** [8.24]	0.102*** [6.70]	0.126*** [7.31]	0.112*** [6.09]	0.027 [1.03]	0.025 [1.24]	0.068** [2.17]	0.067*** [3.02]
T. neg. x Undecided	-0.050 [-1.41]	-0.030 [-0.82]			-0.027 [-0.44]	0.027 [0.54]		
T. pos. x Undecided	0.055 [1.48]	0.087** [2.22]			0.079 [1.42]	0.047 [1.14]		
T. neg. x First election			0.004 [0.15]	0.036 [1.24]			-0.109** [-2.25]	-0.085** [-2.25]
T. pos. x First election			0.010 [0.38]	0.014 [0.48]			-0.084* [-1.89]	-0.090** [-2.56]
Undecided	0.031 [1.31]	-0.004 [-0.17]			0.041 [0.99]	-0.014 [-0.49]		
First elec.			0.003 [0.13]	-0.006 [-0.27]			0.064* [1.75]	0.062** [2.02]
Gov. perf. (baseline)		0.054*** [3.55]		0.058*** [3.86]		0.632*** [23.13]		0.636*** [23.66]
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	1,629	1,423	1,735	1,503	1,624	1,430	1,730	1,509
Adjusted R-squared	0.166	0.168	0.154	0.155	0.016	0.500	0.020	0.499

Notes: The outcome variables are government performance perception associated with the policy areas covered in the treatment (Panel A) or in general (Panel B). Undecided is a dummy variable equal to 1 if respondents stated in the baseline survey that they had not yet decided upon their vote. First election is a dummy variable equal to 1 when respondents were legally entitled to vote for the first time in the local elections. All regressions include dummies for the class of implementation; standard errors are robust and *t*-statistics are reported in brackets. Controls are detailed in Section 4.2. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

no accumulated voting habits and less pre-existing knowledge (which increases the marginal impact of new information).²⁴

Table 8 shows the treatment effects on the performance perception updating across subsamples of undecided voters (columns 1, 2, 5 and 6) and first-time eligible voters (columns 3, 4, 7 and 8), by interacting the treatment assignment dummies $T.Neg_i$ and $T.Pos_i$ with dummies that identify the relevant members of each subsample. The first two columns show a statistically significant and positive coefficient for the interaction between the positive treatment and undecided voters with the full set of controls (column 2), providing some support for the hypothesis that undecided voters update their policy-specific performance beliefs more extensively after the provision of positive information. Conversely, although the coefficient of the interaction with the negative treatment has also the expected negative sign, it is not statistically significant. Hence, our results provide only partial support for stronger opinion updating among undecided voters. First-time voters update policy-specific performance assessments in much the same way as the rest of our respondents (i.e. none of the interaction terms in columns 3 and 4 are significant).

Columns 5 and 6 indicate no differential treatment effect on general government performance perception for the group of undecided voters. The findings are different for first-time voters, who seem to aggregate information in a different way. The provision of negative information translates more strongly into lower perceived government performance, whereas the effect of positive information on perceived general government performance is more diluted than in the rest of our respondents (columns 7 and 8). As expected, Online Appendix Section E.4 illustrates that the statistical significance of this result weakens when correcting for multiple hypothesis testing. Yet, the complete discounting of positive information among inexperienced voters remains statistically significant at conventional levels. Overall, therefore, our findings highlight that less experienced individuals attach somewhat

more value to negative information and significantly less value to positive information. As such, they appear to exhibit a particularly stringent kind of negativity bias. These two sets of results, combined, show that Hypotheses 1 and 3 are confirmed for the case of first-time (inexperienced) voters, but not for undecided ones.

We turn to the results pertaining to voting behaviour in Table 9. Undecided voters who received a negative treatment were significantly less likely to vote, whereas the positive treatment had no significant effect. This result contributes to an ongoing debate about whether negativity in electoral campaigns mobilises or demobilises voters (Lau and Rovner, 2009). Specifically, previous results that demobilisation arises when ‘a person is exposed to negativity after selecting a preferred candidate’ (Krupnikov, 2011, p.797) are questioned by our findings that it depresses turnout particularly among the undecided. This result partially confirms our Hypothesis 4: undecided voters do respond to the negative information treatment by decreasing their participation in the elections.

Turning to vote choice, undecided voters receiving any information treatment were less likely to cast a blank ballot, and those receiving a negative treatment are more likely to cast a vote for parties that do not support the central government. The impact on voting for support parties is not statistically significant, but the coefficients have the expected sign for both positive and negative treatments. Since these respondents did not react differently to the information provided, nor did they use it in a different way in the perception updating process (see Table 8), these results suggest that the difference lies in how participants use the information to determine their local vote choice. Again, a correction for multiple hypothesis testing leads to a general reduction in the observed *p*-values. Nonetheless, Online Appendix Section E.4 illustrates that we continue to find significant support at the 90% confidence level for the fact that respondents in the negative information treatment vote more for other parties. We repeat this exercise for first-time voters in Appendix Table G.1, but we do not find significant changes in turnout or vote choice when compared with the rest of the sample.

²⁴ Although our randomisation protocol did not target subgroups, balance tests in Appendix Figure C.2 display similar means across treatment groups for a wide range of socio-demographic and pre-treatment characteristics.

Table 9
Turnout and voting decision of undecided voters.

	Turnout		Voting decision					
			Support parties		Other parties		Blank	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T. neg.	-0.001 [-0.05]	0.013 [0.45]	0.018 [0.67]	0.021 [0.67]	-0.020 [-0.80]	-0.020 [-0.68]	0.012 [0.67]	0.005 [0.25]
T. pos.	-0.006 [-0.24]	-0.019 [-0.69]	-0.010 [-0.38]	-0.011 [-0.33]	-0.012 [-0.49]	-0.018 [-0.60]	0.022 [1.14]	0.031 [1.34]
T. neg. x Undecided	-0.128* [-1.94]	-0.127* [-1.80]	-0.038 [-0.45]	-0.105 [-1.10]	0.164** [2.08]	0.215** [2.55]	-0.123* [-1.83]	-0.109 [-1.53]
T. pos. x Undecided	-0.086 [-1.40]	-0.066 [-0.97]	0.050 [0.61]	0.034 [0.36]	0.128 [1.63]	0.125 [1.44]	-0.164** [-2.54]	-0.148** [-2.07]
Undecided	-0.001 [-0.03]	0.025 [0.56]	0.267*** [4.54]	0.261*** [3.88]	0.167*** [3.10]	0.139** [2.26]	0.255*** [4.98]	0.188*** [3.36]
Turnout (baseline)	0.701*** [31.00]	0.656*** [21.46]						
Vote support parties (baseline)			0.699*** [27.05]	0.607*** [15.80]				
Vote other parties (baseline)					0.722*** [28.75]	0.598*** [15.48]		
Vote blank (baseline)							0.636*** [15.76]	0.544*** [10.58]
Gov. perf. (baseline)	0.018 [0.79]	0.010 [0.38]	0.147*** [5.84]	0.111*** [3.67]	-0.143*** [-5.87]	-0.111*** [-4.02]	-0.006 [-0.33]	-0.012 [-0.53]
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Municipality dummies	No	Yes	No	Yes	No	Yes	No	Yes
Observations	1,659	1,466	1,191	1,074	1,191	1,074	1,191	1,074
Adjusted R-squared	0.360	0.400	0.467	0.484	0.516	0.558	0.380	0.361

Notes: The outcome variables are turnout (columns 1 and 2) or vote choice (columns 3 to 8), defined as in Table 7. The corresponding baseline variables are included as controls. Undecided is a dummy variable equal to 1 for respondents who had not decided upon their vote one week before the election. All regressions include dummies for the class of implementation; standard errors are robust and *t*-statistics are reported in brackets. Controls include municipal-level dummies and those detailed in Section 4.2. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7. Conclusion

Building on a Randomised Controlled Experiment implemented during the 2017 Portuguese local elections, we show that providing negative information about incumbent performance to young voters induces a substantial and significant downward revision of government performance perceptions. In contrast, providing positive information has a smaller and mostly insignificant impact. This is our first result, and it confirms our hypothesis that young voters are subject to negativity bias. Importantly, we confirm theoretical predictions that such a negativity bias is particularly strong for first-time voters, which could be of considerable use in identifying and protecting targets vulnerable to disruptive (negative) fake news campaigns. Conversely, undecided voters do not update information more strongly than the remaining ones, contrary to what we hypothesised. Turning to voting behaviour, we show that news provision has a significant effect on particularly those voters who remained undecided until a few days before the election. They became less likely to cast a blank vote, and, when receiving negative information, were more likely to abstain or vote for opposition parties. This result confirms our fourth hypothesis in what concerns undecided voters.

Interestingly, while first-time voters react more strongly to information provision, it is the undecided voters who are more likely to adjust their political behaviour due to the information treatment.

These findings contribute to our understanding of youth political behaviour. A recent and vibrant academic literature has paid considerable attention to the effectiveness – or lack thereof – of youth voter mobilisation via social networks, internet, direct text messages, or social media (Nickerson, 2007b; Bhatti et al., 2017; Bergh et al., 2021; Bergh and Christensen, 2022; Foes et al., 2023). Yet, this literature is predominantly concerned with the (lack of) impact of Get Out The Vote campaigns, and does not engage with whether and how young voters process and use political information about, for instance,

incumbents’ performance in office. Clearly, however, it is critical to understand whether or not young – and presumably more impressionable – voters make use of performance information to update their beliefs about incumbents (DellaVigna and Gentzkow, 2010; Arias et al., 2022; Brockman and Kalla, 2022), and whether or not these updated beliefs translate into subsequent electoral choices. Our analysis takes one step in this direction by evaluating the causal impact of (positive or negative) information provision on young voters’ perception of government performance and subsequent voting behaviour. As such, our analysis extends existing work of youth voter mobilisation and allows us to gain a broader and more encompassing view of the determinants and drivers of youth political activity. From a more practical perspective, however, our results also highlight that government policies and activities to protect young voters against, for instance, disruptive and negative (fake) news campaigns should be targeted towards specific – though not always equally easily identifiable – subsets of the population.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data used in the paper was collected by the authors and is currently being used in another paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.electstud.2023.102625>.

References

- Allcott, H., Gentzkow, M., 2017. Social media and fake news in the 2016 election. *J. Econ. Perspect.* 31 (2), 211–236.
- Angrist, J.D., 2006. Instrumental variables methods in experimental criminological research: What, why and how. *J. Exp. Criminol.* 2 (1), 23–44.
- Arceneaux, K., Nickerson, D., 2009. Who is mobilized to vote? A re-analysis of 11 field experiments. *Am. J. Polit. Sci.* 53 (1), 1–16.
- Arias, E., Larreguy, H., Marshall, J., Querubin, P., 2022. Priors rule: When do malfeasance revelations help or hurt incumbent parties? *J. Eur. Econom. Assoc.* 20 (4), 1433–1477.
- Ashworth, S., Bueno de Mesquita, E., Friedenber, A., 2018. Learning about voter rationality. *Am. J. Polit. Sci.* 62 (1), 37–54.
- Baumeister, R.F., Bratslavsky, E., Finkenauer, C., Vohs, K.D., 2001. Bad is stronger than good. *Rev. General Psychol.* 5 (4), 323–370.
- Becker, S.O., Fetzer, T., Novy, D., 2017. Who voted for brexit? A comprehensive district-level analysis. *Econ. Policy* 32 (92), 601–651.
- Bennion, E., Nickerson, D., 2016. I will register and vote, if you teach me how: A field experiment testing voter registration in college classrooms. *PS: Polit. Sci. Polit.* 49 (4), 867–871.
- Bergh, J., Christensen, D.A., 2022. Getting out the vote in different electoral contexts: The effect of impersonal voter mobilization techniques in middle and high salience norwegian elections. *J. Elections Public Opinion Parties* 1–17.
- Bergh, J., Christensen, D.A., Matland, R., 2021. When is a reminder enough? Text message voter mobilization in a European context. *Polit. Behav.* 43 (3), 1091–1111.
- Bhatti, Y., Dahlgard, J.O., Hansen, J.H., Hansen, K., 2017. How voter mobilization from short text messages travels within households and families: Evidence from two nationwide field experiments. *Elector Stud* 50, 39–49.
- Broockman, D., Kalla, J., 2022. When and why are campaigns' persuasive effects small? Evidence from the 2020 US presidential election. *Am. J. Polit. Sci.*
- Daniele, G., Galetta, S., Geys, B., 2020. Abandon ship? Party brands and politicians responses to a political scandals. *J. Public Econ.* 184 (104172).
- DellaVigna, S., Gentzkow, M., 2010. Persuasion: Empirical evidence. *Annu. Rev. Econ.* 2 (1), 643–669.
- Foos, F., John, P., Unan, A., 2023. Progressive campaigns, social media ads and young voters: Null effects from the 2019 UK general election.
- Geys, B., Heggedal, T.-R., Sørensen, R., 2022. Age and vote choice: Is there a conservative shift among older voters? *Elector. Stud.* 78, 102485.
- Geys, B., Vermeir, J., 2014. Party cues in elections under multilevel governance: Theory and evidence from US states. *J. Eur. Econom. Assoc.* 12 (4), 1029–1058.
- Graça, F.V.S., 2017. A política e os media: O enviesamento da imprensa portuguesa em 2009 e 2015. ISCTE-IUL: Dissertação de Mestrado.
- Holbein, J., Hillygus, S., 2020. Making Young Voters: Converting Civic Attitudes Into Civic Action. Cambridge University Press.
- Hong, S., 2019. A behavioral model of public organizations: Bounded rationality, performance feedback, and negativity bias. *J. Public Admin. Res. Theory* 29 (1), 1–17.
- Kendall, C., Nannicini, T., Trebbi, F., 2015. How do voters respond to information? Evidence from a randomized campaign. *Amer. Econ. Rev.* 105 (1), 322–353.
- King, G., Schmeer, B., White, A., 2017. How the news media activate public expression and influence national agendas. *Science* 358 (November 10), 776–780.
- Krupnikov, Y., 2011. When does negativity demobilize? Tracing the conditional effect of negative campaigning on voter turnout. *Am. J. Polit. Sci.* 55 (4), 797–813.
- Kuehnhans, C., Murdoch, Z., Geys, B., Heyndels, B., 2017. Identity, threat aversion, and civil servants' policy preferences: Evidence from the European parliament. *Public Admin.* 95 (4), 1009–1025.
- Lau, R.R., 1985. Two explanations for negativity effects in political behavior. *Am. J. Polit. Sci.* 29 (1), 119–138.
- Lau, R.R., Rovner, I.B., 2009. Negative campaigning. *Annu. Rev. Political Sci.* 12, 285–306.
- Lobo, M., Magalhães, P., Gaspar, J., 2015. Estudo eleitoral Português, 2015.
- Magalhães, P., 2009. A qualidade da democracia em Portugal: A perspectiva dos cidadãos. Rep. SEDES.
- Marien, S., Dassonneville, R., Hooghe, M., 2015. How second order are local elections? Voting motives and party preferences in Belgian municipal elections. *Local Govern. Stud.* 41 (6), 898–916.
- Newman, N., Fletcher, R., Kalogeropoulos, A., Nielsen, R.K., 2019. Digital News Report 2019. Technical report, Department of Politics and International Relations, University of Oxford.
- Nickerson, D., 2007a. Does email boost turnout. *Q. J. Polit. Sci.* 2 (4), 369–379.
- Nickerson, D., 2007b. Quality is job one: Professional and volunteer voter mobilization calls. *Am. J. Polit. Sci.* 51 (2), 269–282.
- Nielsen, P., Moynihan, D., 2017. How do politicians attribute bureaucratic responsibility for performance? Negativity bias and interest group advocacy. *J. Public Admin. Res. Theory* 27 (2), 269–283.
- Popescu, M., Buzarnescu, B., Czikora, E., Gosselin, T., Marincea, A., Santana Pereira, J., Toka, G., 2013. European media systems survey 2013. In: Data Set. Bucharest: Median Research Centre.
- Robinson, L.D., Jewell, N.P., 1991. Some surprising results about covariate adjustment in logistic regression models. *Internat. Statist. Rev.* 59 (2), 227.
- Santana Pereira, J., Nina, S.R., 2016. A democracia nos media Portugueses: Pluralismo polí tico- partidário na televisão e na imprensa. In: *Democracia e Meios de Comunicação: Pluralismo, Liberdade de Expressão e Informação* (Bello, Enzo and Moura Ribeiro, Samantha S.). Lumen Juris, pp. 225–247.
- Schönhage, N.L., Geys, B., 2020. Party cues and incumbent assessments under multilevel governance. *Elector. Stud.* 69 (102260).
- Selb, P., Munzert, S., 2013. Voter overrepresentation, vote misreporting, and turnout bias in postelection surveys. *Elector. Stud.* 32 (1), 186–196.
- Smets, K., Van Ham, C., 2013. The embarrassment of riches? A meta-analysis of individual-level research on voter turnout. *Elector. Stud.* 32 (2), 344–359.
- Zaller, J., 1992. *The Nature and Origins of Mass Opinion*. Cambridge University Press, Cambridge MA.