HOW DO CAMPAIGNS SHAPE VOTE CHOICE? MULTICOUNTRY EVIDENCE FROM 62 ELECTIONS AND 56 TV DEBATES*

CAROLINE LE PENNEC AND VINCENT PONS

We use two-round survey data from 62 elections in 10 countries since 1952 to study the formation of vote choice, beliefs, and policy preferences and assess how televised debates contribute to this process. Our data include 253,000 observations. We compare the consistency between vote intention and vote choice of respondents surveyed at different points before, and then again after, the election, and show that 17% to 29% of voters make up their mind during the final two months of campaigns. Changes in vote choice are concomitant to shifts in issues voters find most important and in beliefs about candidates, and they generate sizable swings in vote shares. In contrast, policy preferences remain remarkably stable throughout the campaign. Finally, we use an event study to estimate the impact of TV debates, in which candidates themselves communicate with voters, and of shocks such as natural and technological disasters which, by contrast, occur independently from the campaign. We do not find any effect of either type of event on vote choice formation, suggesting that information received throughout the campaign from other sources such as the media, political activists, and other citizens is more impactful. *JEL Codes:* D72, D83, P00.

* For suggestions that have improved this article, we are grateful to Alberto Alesina, David Broockman, Ernesto Dal Bó, Ruth Dassonneville, Rafael Di Tella, Andrew Eggers, Benjamin Enke, Aaron Erlich, Frederico Finan, Alexander Fouirnaies, Gregory Martin, Benjamin Marx, Sendhil Mullainathan, Benjamin Olken, Torsten Persson, Nicolas Sauger, Andrei Shleifer, Edoardo Teso, and Matthew Weinzierl, as well as seminar participants at Stanford University, Houston University, Université Libre de Bruxelles, the 2018 Annual Meeting of the American Political Science Association, the 2019 Galatina Summer Meetings, the 2021 Annual Meeting of the American Economic Association, the 2021 IEB Workshop on Political Economy, the Quebec Political Economics Webinar Series, the 2021 LACEA-LAMES Annual Meeting, the 2022 RIDGE Political Economy Workshop, and the 2022 NBER Political Economy Summer Institute. We thank Sudeshna Barman, Ashley Liu, Jack Oster, and Irina Vlasache for great research assistance. Vincent Pons is the cofounder of a company in Europe, eXplain (https://explain.fr/). Besides this position, he has no relevant financial interests that relate to the research described in this paper. Caroline Le Pennec has no relevant financial interests that relate to the research described in this paper.

The Quarterly Journal of Economics (2023), 703–767. https://doi.org/10.1093/qje/qjad002.
Advance Access publication on March 2, 2023.

I. Introduction

Elections are the distinctive feature of representative democracies. In theory, citizens can use their vote to hold incumbents accountable and choose new representatives that are competent and aligned with their preferences (Ferejohn 1986; Besley and Coate 1997). In practice, elections only generate such desirable outcomes insofar as voters rely on adequate information to make their choice (e.g., Delli Carpini and Keeter 1997; Ferraz and Finan 2008). Whether that is the case can have dramatic social and economic consequences, since the quality and ideology of elected officials tend to affect the policies they implement (e.g., Pettersson-Lidbom 2008; Besley, Montalvo, and Reynal-Querol 2011). It is thus critical to assess which information voters use in their choice.

One long-standing view is that the weeks immediately preceding elections represent a crucial period, during which the electorate is flooded with information (e.g., Holbrook 1996). Elections often feature new parties and candidates, and even the propositions and track records of those who have stood for office before may have changed since then. Campaign information can help voters assess incumbents, compare the qualities and positions of all candidates, and reconsider their own policy preferences. Although candidates expend great effort to communicate with voters, interpersonal discussions and the coverage of the campaign by the media also provide rich information. An alternative view is that campaigns have minimal effects as most people have decided whom to vote for long in advance, based on group identities and party attachments (e.g., Lazarsfeld, Berelson, and Gaudet 1944).

We disentangle these contrasting views using two-round surveys in 62 elections around the world since 1952. Our analysis is guided by the idea, consistent with both views, that campaigns' influence will largely depend on the relative importance voters give to their long-standing partisan leanings versus election-specific information on the party platforms and the candidates of the day.

This article makes three distinct contributions. First, we use a novel method to determine the fraction of people that form their vote choice in the last two months before an election, and measure heterogeneity over time and across countries and voters with varying levels of preexisting knowledge and party attachments. Second, we explore whether changes in vote intentions are driven by changes in voters' beliefs about candidates or by other mechanisms, such as changes in their policy preferences or in issue salience. Third, we use an event

study to assess how TV debates—which provide rich information on the candidates—contribute to vote choice formation.

Existing studies have generally sought to isolate the effect on vote choice of one particular source of information, such as newspapers or television (e.g., DellaVigna and Kaplan 2007; Martin and Yurukoglu 2017), or one specific type of campaign communication, such as field visits or TV ads (e.g., Kalla and Broockman 2018; Pons 2018). By contrast, our first set of results relate to the overall impact of information received during campaigns.[1] The more election-specific information voters incorporate in their choice of candidate, the more likely they should be to make up their mind during the campaign. To estimate the share of the electorate who form their vote choice only shortly before elections, we assembled a data set of nationally representative surveys conducted around 62 elections from 1952 to 2017 in 10 countries: Austria, Canada, Germany, Italy, the Netherlands, New Zealand, Sweden, Switzerland, the United Kingdom, and the United States. The data set includes a total of 253,000 observations. All of the surveys entailed interviewing a new set of people every day prior to the election to elicit their vote intention, then surveying them again after the election to record their actual vote choice. By comparing voters' responses in the two rounds, we can ascertain whether they had already settled on their final choice by the time of the pre-electoral survey without having to rely on their recollection of the date when they formed their decision, unlike previous research. In addition, most of our data come from surveys that allocated respondents' survey date randomly, facilitating the interpretation of outcome differences over time.

We find that the fraction of people with identical vote declaration pre- and postelection increases by 17 percentage points over the 60 days leading up to the election, from a baseline of 71%. On the last day before the election, 12% of voters still do not know (or will not say) whom they will vote for or state a different vote intention than their ultimate choice. This brings the total fraction of voters making up their mind during the final two months of campaigns to between 17% (if none of these voters surveyed on the last day are really last-minute deciders) and 29% (if all of them are). In a given election, younger and less educated

1. As is common in the literature, we use "campaign" interchangeably with "electoral season" to designate the period preceding an election and to refer to all factors that may influence voters in that period, including candidates' campaigns and factors beyond their control.

voters are more influenced by campaign information, and voters who identify strongly with a party less so. Across elections, the influence of campaigns on vote choice has been relatively stable for the last 70 years, but it varies substantially from one country to another. Notably, vote choice consistency increases less in the two months preceding the election in the United States than in all the other countries in the sample, even though American campaigns spend much more money to communicate with voters. One possible explanation is that the U.S. two-party system is characterized by strong partisan attachments, making vote choices less malleable than in multiparty settings.

The increase in individual vote choice consistency is concomitant with aggregate trends in the relative strength of competing candidates. We compute each candidate's daily predicted vote share based on the vote intentions of respondents surveyed on that day, and compare it to their final vote share measured in the postelectoral survey. The total distance between predicted and final vote shares decreases by about 5 percentage points over the last 60 days before the election, indicating that vote choice formation during the campaign season can change elections' outcomes. We argue that our results are primarily driven by information acquired and/or processed during the final two months of campaigns and therefore that they provide a good measure of this information's overall impact on vote choice formation.

Our second set of results sheds light on the mechanisms through which information affects people's vote choice. To the extent that changes in vote intention are due to election-specific information on competing candidates and on parties' current platforms, one may first expect changed beliefs about candidates to be an important mediating factor. In fact, using a set of questions asked in both the pre- and postelectoral surveys, we find that the consistency in beliefs concerning candidates' quality and the issues they stand for increases over the campaign. These types of beliefs seem to matter more than beliefs about the relative chances of the contenders that could lead to strategic voting.

Second, voters may be primed by campaign communication to think about certain policy issues, and they may change their vote choice after reassessing their distance with the different candidates based on these issues. We find that the consistency between the issue that respondents consider the most important in the pre- and postelectoral survey increases by more than half of the increase in vote choice consistency.

Third, information shared during the electoral season may affect voter choice by changing their policy preferences. However, using policy questions asked again both pre- and postelection, we do not find any increase in the consistency in policy preferences expressed before and after an election. Furthermore, we show that increases in daily mean issue salience consistency and belief consistency are associated with strong increases in mean vote choice consistency, contrasting with the lack of significant relationship between the latter and policy preference consistency.

Our third set of results provides evidence on the relative importance of different sources of information. Given the mediating effects of beliefs about candidates, a plausible hypothesis is that changes in these beliefs, and in turn, changes in vote intentions, occur as a result of events in which the candidates themselves communicate with voters, such as televised debates. TV debates between candidates for president or prime minister are now part of the electoral cycle in many countries and have a strong apparent potential to inform voters: they give direct and simultaneous exposure to candidates and allow voters to compare their policy positions and performance. Debates draw larger audiences than any other campaign event, and they can also influence nonwatchers through subsequent discussions, social media posts, and media commentaries.

We use an event study approach pooling 56 TV debates in 31 elections and seven countries of our sample. We do not find any significant effect of TV debates on individual consistency between vote intention and vote choice. The fact that we investigate effects on vote choice consistency rather than vote intentions means that we would even uncover effects going in opposite directions for different debates (e.g., with some debates benefitting the incumbent and others a challenger) or different people (e.g., with some voters rallying behind one candidate and others their opponent), which could otherwise remain undetected—yet we find none. Furthermore, at the aggregate level, debates do not significantly affect the distance between predicted and final vote shares. Our null effects are precisely estimated. Considering the 95% confidence intervals, we find that on average, a TV debate contributes no more than 3% of the total increase in vote choice consistency over the final two months of campaigns and 2% of the total decrease in the distance to final vote shares.

Remarkably, we do not find that debates contribute to vote choice formation for any group of voters—including those who

report watching them and those most likely to form their vote choice shortly before the election—or when focusing on types of debates that could be expected to be more effective: the first debate held during the campaign or debates held in close races, fluctuating races, or multiparty systems. Finally, debates do not affect the predicted vote share even of lesser-known candidates, who benefit from the campaign the most overall.

After documenting the null effects of TV debates, we ask whether elections are swayed by shocks such as natural and technological disasters, which, by contrast, occur independently from the campaign. We focus on 27 disasters that occurred before elections in the sample and are included in the EM-DAT International Disasters Database. Using the same event study design, we do not find that disasters contribute to vote choice formation more than TV debates in the elections included in the sample.

From the null effects of TV debates and disasters, we conclude that the type of information that affects voter choice the most is likely neither information directly provided by candidates nor shocks exogenous to the campaign and on which candidates do not have any control, but information provided throughout the campaign by third parties. As the election approaches, we observe that voters receive a growing stream of information from the media, campaign activists, and discussions with family members, friends, and coworkers. Although candidates can try to influence these sources of information (e.g., by communicating with the media), they do not control them entirely.

Our article makes three important contributions to the large literature on the effects and drivers of persuasive communication (DellaVigna and Gentzkow 2010).

First, we estimate the impact on vote choice of a major type of partisan communication, on which the existing evidence is not conclusive: TV debates. A large number of studies explore the effects of TV debates by focusing on a unique election or a small number of races and comparing individual vote intentions, aggregate polls shares, or betting odds before and after debates (Shaw 1999; Shaw and Roberts 2000; Hillygus and Jackman 2003). Although many studies find modest or null effects, others conclude that debates truly matter (Benoit, Hansen, and Verser 2003; McKinney and Carlin 2004; Birdsell 2017). However, these studies' simple pre/post difference designs fail to control for underlying trends. By contrast, we take advantage of the large number of debates that took place in the periods covered by our surveys and of the variation in their timing to flexibly control

for the time to the election. This novel strategy provides more reliable estimates of debates' impact. The fact that our event study includes debates held in many elections and countries also increases our statistical power and the external validity of our estimates well beyond that of any preexisting work.

This study is related to recent experimental evidence on a different type of debate: nontelevised debates opposing parliamentary candidates in low-income democracies (Bidwell, Casey, and Glennerster 2020; Brierley, Kramon, and Ofosu 2020). Scarce political information characterizing these studies' contexts may help explain the substantial effects on vote choices they find. Using randomized experiments to measure the impact of presidential or prime-ministerial TV debates has proven more difficult. Mullainathan, Washington, and Azari (2009) encourage a random selection of New York City voters to watch the final 2005 mayoral election debate and do not find any significant effect on opinions about candidates but acknowledge that subsequent discussions and media commentaries may explain this null result. Instead, Fridkin et al. (2007) use a lab experiment to measure the effect of watching live the final 2004 U.S. presidential TV debate and of the media's instant analysis after it. Measuring participants' immediate reactions, the authors report large effects on candidates' evaluations. In contrast, we find null effects on vote intentions one to three days afterward, suggesting that debates' effects quickly fade away.

While most existing research on vote choice seeks to isolate the impact of a specific source of information, our second contribution is to provide an estimate of their overall influence in the last two months before an election. We build on Wlezien and Erikson (2002) and Jennings and Wlezien (2016), who show that polls become increasingly predictive of actual results as the election comes closer. Our finding that the distance between predicted and final vote shares decreases over time replicates this result in our set of elections. Using individual-level two-round surveys instead of aggregate polls enables us to determine the fraction of voters who arrive at their final choice during the campaign, which is generally larger than the reduction in the distance between predicted and final vote shares; compare the patterns of vote choice formation across different types of voters; and investigate the mechanisms through which information affects vote choice.

Most prior work studying the timing of vote decisions with individual-level data uses respondents' recall of the date when

they made their decision or their declared level of certainty about their vote intention (Chaffee and Choe 1980; Fournier et al. 2004). But voters surveyed before the election may not know how they will respond to information that is yet to come. Postelection recalls are also prone to error, due to people failing to remember when they made their decision or not consciously recording this moment. A smaller set of studies including Henderson and Hillygus (2016) define the time of decision as the date from which panel respondents select the same candidate across all subsequent interviews. While these studies, like ours, are based on the comparison of respondents' answers over time, they use data limited to a single election and cannot provide daily estimates of vote choice consistency, which are our main object of investigation. In addition, they focus on the level of consistency, which may be biased by misreporting (see Section III.A), rather than its change.

Finally, we build on prior work showing that information received by voters can affect their vote choice by changing their beliefs, whether on candidates or on the state of the economy (Gelman and King 1993; Kendall, Nannicini, and Trebbi 2015). Recent studies also document instances of people changing their actual preferences based on campaign interactions (Minozzi et al. 2015). Overall, we show that campaign information tends to affect vote choices and election outcomes by changing beliefs and issue salience more than policy preferences.

The remainder of the article proceeds as follows. Section II describes our data. Sections III and IV study the formation of vote choice, beliefs, policy preferences, and issue salience in the campaign. Section V estimates the impact of debates and disasters on these outcomes, and Section VI concludes.

II. Data

II.A. Campaign Surveys

We assembled a new data set of nationally representative surveys conducted around 62 elections in 10 countries from 1952 to 2017. The data come from the American National Election Studies (1952 to 2016), the Canadian Election Studies (1988 to 2015), the British Election Studies (2001 to 2016), the New Zealand Election Studies (1996 to 2002), the Dutch Parliamentary Election Studies (1998 to 2006), the National Annenberg Election Surveys (2000 to 2008), the German Longitudinal Election Studies (2009 to 2017), the Swiss Electoral Studies (2011 and 2015), the Italian

National Election Studies (2013), the Austrian National Election Studies (2013), and the Swedish National Election Studies (2014).[2] We keep all respondents surveyed 60 days before the election or less, as only a few surveys started earlier. Integrating the responses collected with independent questionnaires into a common empirical framework marks an important effort.

A few surveys cover multiple elections because multiple offices were on the ballot on the same day (e.g., president and member of Congress in the United States) or because voters can cast multiple ballots (e.g., Germany's first and second votes). We define each of these offices or ballots as a separate election. Conversely, the 2000, 2004, and 2008 U.S. presidential elections are covered by the American National Election Studies (ANES) and the National Annenberg Election Surveys (NAES). Online Appendix Table B.1 shows the full list of elections, their date, type, voting rule, and key features of the corresponding surveys. Twenty-seven percent of the elections were for a president, 58% for a lower house, 5% for an upper house, 5% for governor, 2% for the European Parliament, and 3% on referenda. Seventy-six percent used the plurality rule and 24% the proportional rule. We refer indifferently to the individual candidates competing in plurality elections and party lists competing in proportional elections as "candidates."

To build this data set, we searched for all electoral surveys around the world that satisfy three criteria. First, they must survey respondents twice: once before the election, to elicit their vote intention, and once afterward, to record their ultimate choice. We observe 253,000 pre-election vote intentions (including people who say they do not know whom they will vote for) from 217,000 unique respondents, and postelectoral responses for 201,000 (80%) of these observations.[3] The median length between the election and the postelectoral survey was 14 days on average. Second, surveys must interview a new set of respondents every day until the election and record the corresponding date. Third, respondents surveyed on different dates must be as similar as possible.

2. The full list of links at which the surveys can be downloaded and the corresponding references are available in Online Appendix B.1.

3. The fraction of respondents surveyed twice should not be read as a success rate in resurveying respondents. Indeed, while most surveys attempt to reach all respondents surveyed before the election a second time afterward, others only attempt to resurvey a subset of pre-election respondents, bringing the fraction down. Conversely, a few surveys only release data for respondents successfully surveyed twice, bringing the fraction up.

To satisfy the third criterion, most of our sample comes from rolling cross sections—surveys that allocate each respondent's survey date randomly. This design implies that the set of respondents surveyed on any particular day can be treated as an independently drawn random sample and it reduces the risk that answers from respondents surveyed on different dates differ because of differences in their characteristics (Johnston and Brady 2002). To increase statistical power, we complemented our sample with surveys that were not designed as rolling cross sections but are statistically close to daily random sampling. Specifically, we include surveys that do not show too large imbalances in pairwise comparisons of daily respondents' observable characteristics (see Online Appendix A.1 for additional details).

Our key variables of interest are respondents' pre- and postelection vote declaration. We further use questions on policy preferences, issue salience, and beliefs about candidates that were asked in the same way before and after the election, allowing us to use the same specifications as for the formation of vote choice. We identified 46 questions from 12 surveys that recorded the policy preferences of a total of 106,000 respondents, and 76 questions from 11 surveys that elicited the beliefs of 112,000 respondents on the quality and policy positions of competing candidates. The full list of these questions is available in Online Appendix Tables B.3 and B.4. To measure changes in issue salience, we use open-ended questions in 12 surveys asking a total of 61,000 respondents which issue they find the most important in this election. We rank all possible answers in all surveys under 10 categories: economic policy, social policy, foreign policy, public safety, civil rights, moral values, institutions, politics, electoral issues, and other issues.

Finally, we keep the following covariates for heterogeneity and other analyses and standardize them across surveys: respondents' education, age, gender, income, and employment status, which are recorded by the vast majority of surveys, as well as their consumption of different media, the party they identify with, the strength of their party identification, their propensity to watch TV debates and read polls, whether they have recently been contacted or visited by a party, and how frequently they have discussions about politics, when available (Online Appendix Table B.2). To construct all our variables homogeneously across surveys, we follow a set of common rules, detailed in Online Appendix B.2.

II.B. Complementary Data

We supplement the survey data with information that we collected from separate sources including ParlGov and the Manifesto Project on competing candidates' party, their incumbency status, and whether they were on the ballot for the first time (Online Appendix Table B.5).

In addition, we systematically searched for the existence and dates of all TV debates between presidential or prime-ministerial candidates during the periods covered by the surveys. We used and cross-checked the following sources: academic papers, TV channels archives, newspaper articles, and Wikipedia. The full list of debates included in the analysis is in Online Appendix Table B.6.

Finally, we used the EM-DAT International Disasters Database to identify natural disasters (e.g., wildfires and floods) and technological disasters (e.g., industrial and transportation accidents) that occurred before the elections in the sample. We provide more information on EM-DAT in Online Appendix B.3 and show the full list of disasters included in the analysis in Online Appendix Table B.7.

III. The Formation of Vote Choice

Information released during campaigns may help voters decide whom to vote for and, in turn, increase the predicted vote shares of some candidates at the expense of others. Therefore, to assess how much campaign information matters, we use two distinct metrics: one based on individual-level data, and measuring how many people make up their mind in the last two months of the campaign, the other assessing how much aggregate vote shares change during this period.

Importantly, our goal is to measure the overall impact of all campaign information, not the quantity of it. This information may come from many sources, ranging from campaign organizations' messaging to discussions with family and friends. Whereas previous work has sought to isolate the effect of specific types of information, studying the timing of vote choice formation enables us to account for all information shared during campaigns, without having to observe which particular pieces of information people receive.

Our approach would overestimate the impact of campaigns if some individuals made up their mind during the electoral season for other reasons than the information they receive. To provide

evidence on the predominant role of campaign information, we check whether vote choice formation is concentrated when and where the campaign is most intense and salient. Furthermore, campaign information should be expected to matter more for voters who have less preexisting political knowledge and less for those who identify strongly with a party. Similarly, one may expect campaign effects to be stronger in countries with multiparty systems, where partisan attachments are weaker and vote choices may be more malleable than in countries with a two-party system like the United States. To test these predictions, we compare campaign effects across voter types and across countries.

We conclude this section by addressing possible threats to the validity of our results and discussing alternative interpretations.

III.A. Individual Vote Choice Formation

The fraction of people who decide which candidate to vote for in the last weeks before an election is difficult to estimate directly. Indeed, it is hard for voters to assess the likelihood that they will stick to their vote intention, ex ante, or to recall the exact date they made up their mind, ex post. We overcome this issue with a novel method using questions recording vote intention and vote choice, which are easier to respond to, and comparing a respondent's answers to both.

Formally, we define vote choice consistency as a dummy equal to 1 if the respondent's pre- and postelection vote declaration coincide and 0 if they differ or if the respondent said they did not know whom they would vote for in the first survey. Voters receiving a stream of information and incorporating it into their evaluation of candidates should be expected to show increasing consistency between their pre-election vote intention and their vote choice over time for two reasons. First, voters surveyed later in the campaign will have received more information. Accordingly, their posterior on candidates will be more precise, making them more likely to state a vote intention and less likely to change it afterward.[4]

4. Competing parties may provide conflicting information, leading some voters to remain uncertain about their choice until late in the campaign. However, even this type of information may strengthen others' confidence in their choice. For instance, voters who initially lean toward a certain candidate may feel more certain about their choice after that candidate responds to a competitor's attacks. Furthermore, pieces of information from sources other than parties and candidates may be less likely to cancel each other out.

Second, the later voters are given the pre-electoral survey, the less time there is for them to receive new information liable to change their vote intention afterward.

We estimate the share of respondents surveyed on any day who will vote according to their intention with the following OLS specification:

$$C_{it}^{e} = \sum_{t=-60}^{-1} \beta_t D_t + \alpha^e + W_{it}^{'}\lambda + u_{it}^{e}, \quad (1)$$

where C_{it}^{e} is the vote choice consistency of respondent i, surveyed for the first time t days before election e, D_t's are 60 fixed effects indicating the number of days relative to the election, α^e are election fixed effects, and W_{it} is a vector of controls. W_{it} includes fixed effects for the day of the week in which the pre-electoral survey took place and for the number of days separating the postelectoral survey from the election and, in some specifications, sociodemographic characteristics.[5]

The key coefficients of interest are the β_t's. We center all control variables around their mean value at $t = -1$ and do not include a constant, so that β_{-1} is equal to the outcome's sample average among respondents surveyed one day before the election and, for any $t \neq 1$, β_t is the (conditional) expected outcome for respondents surveyed t days before. Our sample includes all respondents surveyed both before and after the election who said that they intended to vote in the first survey and reported that they actually voted and gave a vote choice declaration in the second. Standard errors are clustered at the survey level.[6]

We plot the β_t coefficients against time in Figure I. We find that 60 days before the election, 71% of voters state a vote intention corresponding to their final vote choice, suggesting that they vote based on earlier information or along party lines. The fraction of people with identical pre- and postelection vote

5. In place of election fixed effects, we include two separate fixed effects (or survey × election fixed effects) for U.S. elections covered by ANES and NAES to also control for survey effects.

6. Our results are robust to allowing for correlation of the error terms with the wild cluster bootstrap procedure, as shown in Online Appendix C.1, and to clustering the standard errors at the level of the election date, as shown in Online Appendix C.2. Respondents in the 2008 wave of both the ANES and the NAES are then included in the same cluster, for instance.

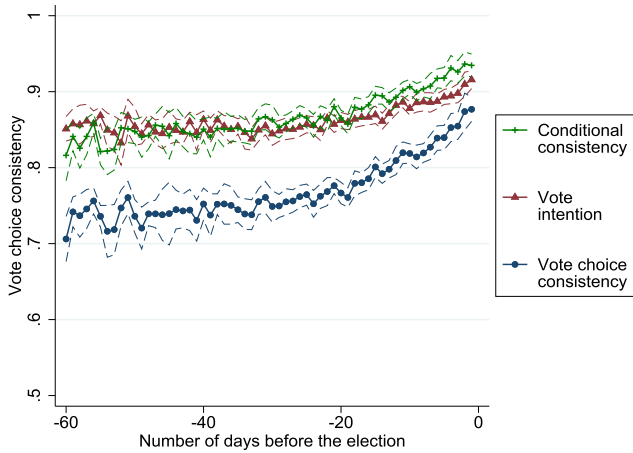

FIGURE I

Individual Vote Choice Consistency, Vote Intention, and Conditional Consistency

We show point estimates and 95% confidence intervals from specifications in the form of equation (1), regressing vote choice consistency ($N = 200{,}916$), vote intention ($N = 253{,}489$), and conditional vote choice consistency ($N = 178{,}176$) on 60 fixed effects indicating the number of days relative to the election and control variables listed in the text.

declarations increases to 88% during the final two months of campaigns. The 12% of voters surveyed on the last day before the election whose vote intention and vote choice remain different are of two types: half of them still do not know (or will not say) whom they will vote for, and the other half state a vote intention but later report a different vote choice.

Next we estimate the following equation to measure the daily average increase in vote choice consistency and test whether the trend is linear or convex:

$$C_{it}^{e} = \beta t + \delta t^{2} + \alpha^{e} + W_{it}^{'}\lambda + u_{it}^{e}, \tag{2}$$

where t is defined as negative the number of days separating the pre-electoral survey from the election, so that higher values of t indicate closer proximity to the election. Convexity should be expected if the stream of information available to voters increases as the election gets closer, either due to increasing demand (by voters eager to make up their mind) or supply (by candidates, the media, friends and family members, etc.). The results are reported in Table I. We find that each additional day increases

TABLE I

VOTE CHOICE CONSISTENCY AND DISTANCE TO FINAL VOTE SHARES

	Vote choice consistency		Vote intention		Conditional consistency		Distance between predicted and final vote shares	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Time	0.0025***	0.0061***	0.0010***	0.0031***	0.0018***	0.0038***	−0.0008***	−0.0013***
	(0.0003)	(0.0008)	(0.0002)	(0.0005)	(0.0003)	(0.0007)	(0.0002)	(0.0004)
Time^2		0.0001***		0.0000***		0.0000***		−0.0000
		(0.0000)		(0.0000)		(0.0000)		(0.0000)
Observations	200,916	200,916	253,489	253,489	178,176	178,176	3,125	3,125
R^2	0.070	0.071	0.072	0.073	0.041	0.041	0.520	0.523
Mean at day −1	0.8768	0.8768	0.9157	0.9157	0.9346	0.9346	0.0324	0.0324
Election fixed effects	x	x	x	x	x	x	x	x
Individual controls	x	x	x	x	x	x		
Aggregate controls							x	x

Notes. Standard errors clustered at the survey level are in parentheses. Time is defined as negative the number of days separating the pre-electoral survey from the election. In columns (1)–(6), we use one observation per respondent and estimate a specification in the form of equation (2). In columns (7) and (8), we use one observation per election per day and estimate a specification in the form of equation (5). Individual controls include fixed effects for the day of the week in which the pre-electoral survey took place and, in columns (1), (2), (5), and (6), fixed effects for the number of days separating the postelectoral survey from the election. Aggregate controls include fixed effects for the day of the week in which the pre-electoral survey took place and the average number of days separating the postelectoral survey from the election. In this table and others, respondents included in the sample (and, therefore, sample sizes) vary across outcomes, as explained in the text. ***, **, * indicate significance at 1%, 5%, and 10%, respectively.

vote choice consistency by 0.25 percentage point on average, an estimate that is significant at the 1% level (column (1)), and that the increase in consistency follows a convex pattern, with a significant acceleration in the last weeks preceding the election (column (2)). This pattern confirms that the process of vote choice formation is concentrated shortly before the election, when the intensity and salience of the campaign are at their peak.[7]

Our survey data may suffer from self-reporting biases. We address this concern in two ways. First, the fact that we define vote choice consistency by comparing the intention and ultimate choice of the same person eliminates any bias present in both declarations. Most expressions of survey demand effects and of social acceptability bias likely fall in this category.

Second, one may still be concerned by the possibility that some voters misreport their vote intention while reporting their actual vote choice, or the reverse. In particular, voters may misreport their vote choice because they forgot, or out of the desire to say they voted for the winner (Wright 1993), leading to an inflated fraction of inconsistent voters measured on any day. However, vote choice misreporting should not bias our estimate of the change in consistency over time. To see why, first note that taking the difference between the vote choice consistency of respondents surveyed at different dates eliminates any constant level of misreporting in and of itself. Misreporting could still vary over time: for instance, response accuracy could decrease with the time between the election and the postelectoral survey, which is in turn correlated with the date of the pre-electoral survey (Online Appendix Figure A.2). We address this possibility by controlling flexibly for the postelectoral survey lag. Conditional on the dummies for the number of days separating the postelectoral survey from the election, included in equations (1) and (2), vote choice misreporting should be uncorrelated with the timing of the pre-electoral survey, and the estimated change in consistency should be fully accurate.[8]

7. Furthermore, Online Appendix Figure A.1 replicates Figure I for a period that is twice as long, by including respondents surveyed more than 60 days before the election in 18 surveys in our original sample and in "precampaign" surveys that preceded 6 of our pre-electoral surveys (listed in Online Appendix Table B.8). We do not observe any increase in vote choice consistency between 120 and 60 days before the election. The daily change in vote choice consistency in this period is nonsignificant and 10 times smaller than during the last 60 days before the election (Online Appendix Table A.1).

8. Some forms of vote choice misreporting may depend on the distance between the pre- and the postelectoral survey more than the distance between the election

In sum, our method insulates the estimated 17 percentage point increase in vote choice consistency from 60 days to 1 day before the election from multiple plausible sources of reporting bias. In contrast, we may overestimate the share of voters who remain inconsistent on the day before the election if some of them misreport their vote choice but not their vote intention. We conclude that the fraction of voters forming their vote choice during the last two months before the election is between 17% (if none of the 12% of inconsistent voters surveyed the last day before the election are really last-minute deciders) and 29% (if all of them are).

It is instructive to compare these estimates with the fraction of voters who self-report making up their mind during the campaign. We identified postelectoral questions asking respondents to recall the timing of their decision in 45 of the 62 elections in our sample. Overall, 48% of the voters in these elections said that they decided whom to vote for during the campaign, which is much larger than our estimates. While our estimates indicate strong campaign effects, trusting voters' recollection of the moment when they formed their decision would exaggerate the importance of campaigns.[9]

III.B. Stating a Vote Intention and Conditional Consistency

The increase in vote choice consistency can result either from an increased fraction of people stating any vote intention or from increased vote choice consistency conditional on stating one. Indeed, campaign information may help undecided voters pick a candidate and change the vote intention of those who already

and the postelectoral survey. For instance, some respondents may try to give a postelectoral answer identical to their pre-electoral answer, instead of truthfully reporting whom they voted for. But providing identical answers may be easier when the two surveys were closer to each other. Therefore, in Online Appendix Figure A.3, we replace the dummies for the postelectoral survey lag with dummies for the number of days separating the two surveys. Reassuringly, we observe a similar increase in vote choice consistency over time as in Figure I.

9. At the individual level, mistakes go in both directions, which could further lead us to mischaracterize the types of voters most influenced by campaigns. Using our main sample, we find that 11% of the voters who report that they knew whom they would vote for before the campaign had still not converged to their final vote choice by the time of their pre-electoral interview, 60 days or less before the election. Conversely, 66% of the respondents interviewed as part of the surveys starting more than two months before the election, and who said that they made their vote choice during the campaign, were already consistent by the time of their pre-electoral interview, 120 to 60 days before the election.

had a candidate in mind. Accordingly, we estimate equation (1) using as outcome a dummy equal to 1 if the respondent states a vote intention and 0 otherwise. We then restrict the sample to respondents stating a vote intention and use a dummy equal to 1 if their final vote choice corresponds.[10]

As shown in Figure I, the probability of stating a vote intention and conditional vote choice consistency increase in the last two months before the election, up to 92% and 93%, respectively. Furthermore, both outcomes follow an increasing and convex pattern similar to vote choice consistency (Table I, columns (3)–(6)).

We derive two key insights from these results. First, the increase in vote choice consistency shown in Section III.A could be generated by multiple processes of information aggregation, including one where people's vote choice is entirely swayed by the latest information. However, voters' increasing likelihood to express a vote intention suggests that they accumulate information and only state a vote intention when they feel sufficiently confident. Consistent with this view, people's certainty about their vote intention increases as well (Online Appendix Figure A.4). Second, the increase is larger for conditional consistency than for stating a vote intention, suggesting that vote choice formation during the campaign is driven by a decrease in the fraction of voters changing their mind at least as much as by a decrease in the fraction of undecided voters.

III.C. Heterogeneity across Countries and Election Years

To test whether the increase in vote choice consistency measured in Section III.A varies across countries and over time, we estimate equation (2) for each election separately, without the quadratic term. Figure II plots each election-specific daily increase in vote choice consistency against election year, along with country-level linear fits.

The extent to which people form their vote choice during the last two months before the election has been remarkably stable

10. Our sample for the first outcome includes all respondents surveyed before the election who said that they intended to vote. We do not control for fixed effects for the number of days separating the postelectoral survey from the election when using this outcome. For the second outcome, our sample includes all respondents surveyed before and after the election who said that they intended to vote and stated a vote intention in the first survey and who reported that they actually voted and gave a vote choice declaration in the second.

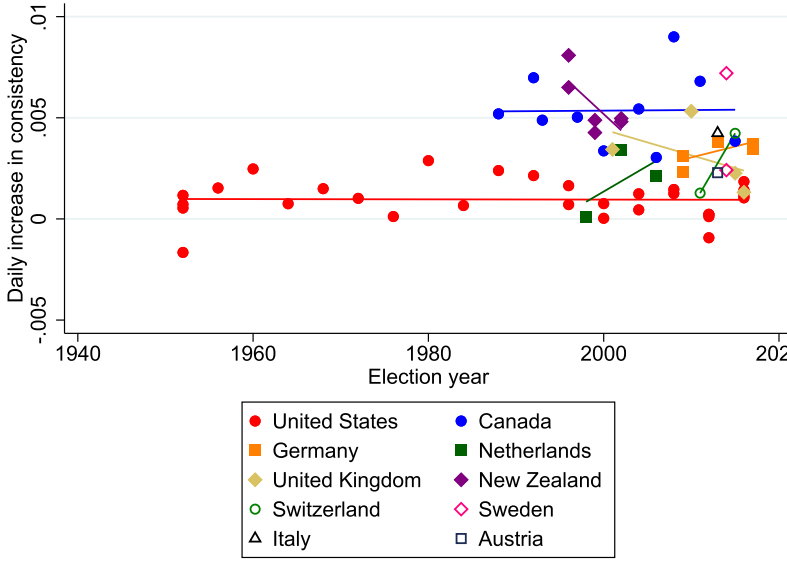

FIGURE II

Increase in Vote Choice Consistency across Countries and over Time

We show point estimates from election-level regressions of consistency on time against election year. Each point comes from a separate regression. Each regression controls for fixed effects for the day of the week in which the pre-electoral survey took place and for the number of days between the postelectoral survey and the election. Elections covered by two different surveys are represented by two different point estimates. We also show country-level linear fits of the point estimates, estimated by regressing the point estimates on election year, for all countries with two election years or more. $N = 65$.

over time in the United States and in Canada, the two countries with the largest number of elections in the sample (44% and 16%, respectively, accounting for 28% and 13% of respondents). It has decreased slightly in some countries with fewer elections (New Zealand and the United Kingdom), but increased slightly in others (Germany, the Netherlands, and Switzerland). Overall, the propensity to form one's vote choice in the campaign season has been relatively stable for the past 70 years, suggesting that campaigns continue to matter as much as before.[11] This constancy

11. Differences in average daily increase in vote choice consistency within countries and over time could partly be affected by differences in survey length. Indeed, our main specifications keep all respondents surveyed 60 days before the election or less, but some surveys start later than that, and we have shown in

is all the more striking as campaign methods have undergone major changes in this period, including the long decline and recent revival of strategies focusing on mobilizing of nonvoters rather than persuading of active voters (Panagopoulos 2016), new types of media have emerged, and ideological polarization has risen in many countries (Boxell, Gentzkow, and Shapiro 2022).

Second, even though campaigns are universally relevant, the size of the daily change in consistency differs substantially across countries. In particular, vote choice consistency increases less in the last two months before the election in the United States than in all other countries. To investigate the size and statistical significance of this difference, we use the following specification:

$$C_{it}^{e} = \beta t + \gamma \Omega^{e} t + \alpha^{e} + W_{it}^{'} \lambda + u_{it}^{e}, \tag{3}$$

where Ω^{e} is a dummy equal to 1 for the United States. As shown in Online Appendix Table A.2, column (1), the fraction of voters forming their vote choice during the electoral season in U.S. elections remains positive and significantly different from zero, but it is substantially lower (by about two-thirds) than in other countries. The fact that electoral campaigns spend much more money per capita in the United States than in other countries makes this result particularly puzzling. We do observe a larger increase in vote choice consistency by about 50% (which is significant at the 10% level) in U.S. swing states, where the campaign intensity is higher, than in nonswing states (Online Appendix Table A.3, column (1)). However, even in swing states, vote choice formation during the campaign remains lower than in other countries.

A possible interpretation is that the vast majority of U.S. elections are bipartisan, making vote choices less malleable, while all other countries have multiparty systems. In multiparty settings, a larger number of candidates are on the ballot, sometimes including candidates of new parties, which campaign information helps voters learn about. The diversity of candidates and frequent changes in the party system result in weaker partisan attachments, and voters can change their vote choice by switching

Table I that vote choice consistency increases faster shortly before the election. To address this concern, we identify the minimum number of days covered by any survey, country by country, and then restrict the sample to the respondents surveyed on or after that country-specific minimum. Reassuringly, Online Appendix Figure A.5, obtained after this homogenization, shows stable patterns over time, similar to Figure II.

between candidates with similar ideologies without having to cross the aisle.

Since the U.S. two-party system stems in part from using the plurality rule, an alternative interpretation is that cross-country differences in vote choice formation come from differences not between party systems but between plurality and proportional elections. For instance, switching sides could be more common in proportional elections if contests between party lists are less conflictual and polarized than face-offs between individual candidates. The presence in the sample of countries such as Canada and the United Kingdom, which use plurality voting like the United States yet have multiparty settings, enables us to test this hypothesis. We do not find any significant difference between plurality and proportional rule elections, in specifications also controlling for time interacted with the U.S. dummy (Online Appendix Table A.2, column (3)).

Other factors may drive the difference between the United States and other countries, including the fact that U.S. voters can start forming their choice during lengthy primary elections. Regardless of the exact reason, this result suggests that lessons from U.S. studies on drivers of voter behavior and electoral results, which account for most of the existing literature, may not extend to other contexts.

III.D. Heterogeneity across Voters

Voters who identify strongly with a party are likely to vote for this party's candidate in any case and be less influenced than others by campaign information. We should thus expect their vote choice consistency to be high 60 days before the election and to show little increase afterward. By contrast, voters with low preexisting knowledge on candidates or on the state of the economy may show a larger increase in vote choice consistency before the election than those who are more knowledgeable if they are more affected by the information they receive in this period, as Bayesian updating would predict.

To test these predictions, we compare the timing of vote choice formation for voters with different strengths of partisan attachments as well as different age and education levels, which are two strong correlates of political informedness (Angelucci and Prat 2021). We also test for differences along three additional

dimensions available across most surveys and known to predict vote choices: gender, income, and employment status.

We use the following specification separately for each characteristic:

$$C_{it}^{e} = \beta t + \Omega_i + \gamma \Omega_i t + \alpha^{e} + W_{it}^{'}\lambda + u_{it}^{e}, \tag{4}$$

where Ω_i is a dummy variable equal to 1 if respondent i is a "type-a" voter—defined as male, above the median age of that survey's respondents, college educated, above the median income, not employed, or strongly identifying with a party—and zero if they are "type-b"—female, below median age, not college educated, below median income, employed, or not identifying with any party or only weakly so. γ measures the differential increase in C_{it}^{e} for type-a voters over time.

We find that vote choice consistency increases substantially for all groups during the last 60 days before the election, but the increase is faster for younger voters and those without a college degree and slower for those who strongly identify with a party (Online Appendix Table A.4, columns (1)–(6)). These differences remain significant (at the 1% level) in a specification including all characteristics and their interaction with the time trend (column (7)). As predicted, other characteristics equal, voters with strong party identification are initially more consistent than those with weak or no party identification, and their vote choice consistency increases by 0.13 percentage point less per day on average (or 52% of the average daily change).

Conversely, voters without a college degree and younger voters show lower initial levels of vote choice consistency than college degree holders and older voters, but any additional day increases their consistency by an additional 0.05 and 0.07 percentage point, respectively, on average (20% and 28% of the average daily change). Much of the differential increase in vote choice consistency for younger voters is driven by voters under 25 years: this group's vote choice consistency increases by an additional 0.06 percentage point on average, compared with other voters below median age (Online Appendix Table A.5). The fact that young voters are more susceptible to the influence of campaigns echoes previous studies showing that their political behavior and attitudes are more easily influenced by external factors (Neundorf and Smets 2017; Cantoni and Pons 2022).

Finally, we find that the larger increases in vote choice consistency of younger and less educated voters as well as those without strong party identification are driven by larger changes in both the probability of stating a vote intention and vote choice consistency conditional on stating one (Online Appendix Table A.6).[12]

III.E. Convergence to Final Vote Shares

Changes in individual vote intentions may partly compensate for each other: voters switching from intending to vote for candidate A to actually voting for candidate B will not affect aggregate vote shares and the outcome of the election if an equal number of voters follow the opposite trajectory. However, voters receiving the same information may update their vote intentions in the same direction, for instance, toward candidates who prove to be more competent or to defend propositions closer to their preferences, generating broad shifts in candidate support. We should then expect increased individual consistency between vote intention and vote choice to be concomitant with a convergence from predicted vote shares to final vote shares.

We use vote intentions and vote choices reported in the pre- and postelectoral surveys to compute $\widetilde{V}^e_{ct}$, the predicted vote share of candidate c in election e among respondents surveyed at time t, and V^e_{ct}, the candidate's final vote share among the same respondents.[13] We define the overall distance between predicted and final vote shares as $\Delta V^e_t = \frac{1}{2}\sum_c \left|\widetilde{V}^e_{ct} - V^e_{ct}\right|$, which corresponds to the minimal share of voters who had to change their vote intention after the pre-electoral survey to explain the difference between predicted and final vote shares.

We measure changes in this outcome with a specification in the form of equation (1), but using only one observation per election per day instead of one observation per individual

12. Using voters' own recollection of the date when they formed their decision fails to fully capture the heterogeneity of campaign effects across voter types. For instance, less educated voters are not more likely than college-educated respondents to report that they decided whom to vote for during the campaign (Online Appendix Table A.7).

13. Our sample includes all respondents surveyed before and after the election who said that they intended to vote and stated a vote intention different from voting blank or null in the first survey and who reported that they actually voted and gave a vote choice declaration different from voting blank or null in the second.

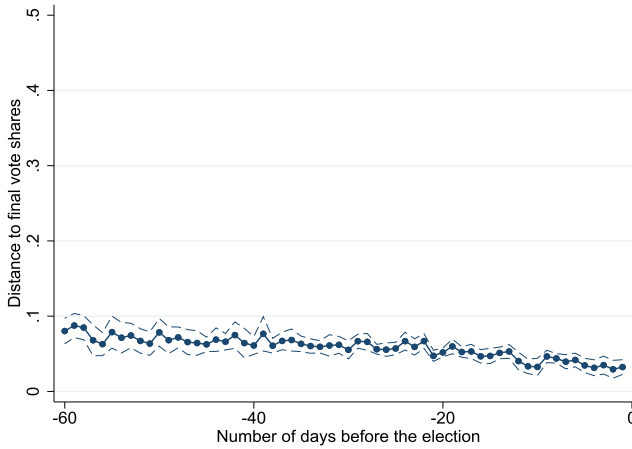

FIGURE III

Distance between Predicted and Final Vote Shares

We show point estimates and 95% confidence intervals from a specification in the form of equation (5), regressing the distance between predicted and final vote shares on 60 fixed effects indicating the number of days relative to the election and control variables listed in the text ($N = 3{,}125$).

response:

$$\Delta V_t^e = \sum_{t=-60}^{-1} \beta_t D_t + \alpha^e + W_t^{e'} \lambda + u_t^e, \tag{5}$$

where W_t^e includes pre-electoral survey day-of-the-week fixed effects, the average postelectoral survey lag among respondents who received the pre-electoral survey at time t, and, in some specifications, their average sociodemographic characteristics.[14]

We plot the β_t coefficients on Figure III. The overall distance between predicted and final vote shares is more than halved, from 8 percentage points on average 60 days before the election to 3.2 percentage points the day before.

14. To give more weight to vote shares measured more precisely, we weight each observation by $\frac{N_t^e}{N_t}$, where N_t is the total number of respondents surveyed at time t and N_t^e is the subset of these respondents surveyed for election e. We obtain a noisier but very similar graph when we do not weight observations by $\frac{N_t^e}{N_t}$ (Online Appendix Figure A.6).

III.F. Threats to the Validity of Our Results

By construction, our sample only includes respondents willing to answer surveys. In addition, we generally do not observe actual turnout and rely on self-reported participation to restrict our analysis to active voters. These two limitations generate possible concerns for the external and internal validity of our results.

A first concern relates to the validity of our results beyond survey respondents. The fact that all our surveys are nationally representative should alleviate this concern somewhat. The representativeness of our surveys is reflected in a strong 0.97 correlation between actual aggregate vote shares and vote shares computed based on survey respondents. In addition, around two-thirds of the surveys provide weights to increase the representativeness of their samples. We do not use these weights in our baseline specifications, because weighting schemes vary substantially across surveys and some surveys do not include any weights, but our main results remain very similar when taking weights into account (Online Appendix C.3).

A second, symmetric concern is that our sample may include individuals who should actually be excluded from it: respondents who said that they were going to vote before the election and that they did vote after the election, but actually abstained. Three surveys in our sample, the 1980, 1984, and 1988 ANES, enable us to identify nonvoters miscategorized as actual voters because they recorded both self-reported turnout and actual turnout, based on official voting records. We complemented these surveys with the 2008, 2012, and 2016 CCES surveys, which also validate respondents' turnout.[15]

Combining these six surveys, we find that turnout overreporters are 3 percentage points more likely to be inconsistent between their vote intention and vote choice than actual voters, on average. However, the inclusion of these individuals in our sample should not affect the change in vote choice consistency during the last two months, which is our main result. Indeed, we do not find a systematic correlation between the characteristics predicting overreported turnout and those predicting

15. The latter surveys are poorly suited to study vote choice formation over time and thus not included in our main sample, because the observable characteristics of their respondents vary tremendously from one day to the next, but they are useful for assessing the possible biases related to turnout overreporting.

increased vote choice consistency. We estimate the following specification:

$$T_{it}^{e} = \Omega_i'\rho + \sum_{t=-60}^{-1} \beta_t D_t + \alpha^e + W_{it}'\lambda + u_{it}^{e}, \tag{6}$$

where T_{it}^{e} is a dummy equal to 1 if the respondent indicated that they intended to vote and that they voted but in reality abstained, Ω_i is a vector of dummies indicating voter characteristics, and the other variables are defined as previously. Younger respondents, whom we found to show a larger increase in vote choice consistency in Online Appendix Table A.4, are more likely to overreport turnout, but less educated respondents and those with weak or no party identification, who are also more likely to converge to their final choice during the campaign, are less likely to be turnout overreporters (Online Appendix Table A.8). Moreover, the daily increase in vote choice consistency among respondents from ANES 1980, 1984, and 1988 is strikingly similar whether we include turnout overreporters or only validated voters (Online Appendix Table A.9).

We turn to threats to the results' internal validity that could ensue from changes in the composition of our daily samples over the electoral season. As indicated in Section II, one of the criteria we used to decide which surveys to include in our sample was that the sociodemographic characteristics of their respondents surveyed on different dates be as similar as possible. This should reduce the risk of sample selection bias. Yet some of these variables show slight imbalance over time (Online Appendix Table A.10). Reassuringly, our findings are nearly identical when controlling for them (Online Appendix C.4). Moreover, we check that our results are robust to restricting the sample to surveys designed as rolling cross sections, accounting for 69% of our observations, in which the pre-electoral interview date is allocated randomly across respondents (Online Appendix C.5). Beyond these general robustness checks, we address two specific concerns, related respectively to changes over time in turnout intentions in the pre-electoral survey, and in the fraction of respondents who are successfully resurveyed in the postelectoral survey.

First, our sample is restricted to respondents who said that they intended to vote in the pre-electoral survey, and these

respondents' proportion and type may change over time.[16] As shown in Online Appendix Figure A.7 and Online Appendix Table A.11 (column (1)), the share of respondents who say that they intend to vote increases by about 6 percentage points over the 60 days leading up to the election. This pattern is interesting in itself. It may reflect the mobilizing effects of campaigns (Hillygus 2005), as well as an increased propensity to overreport voting as the election comes closer. But it leads to increases over time in the fraction of respondents entering in the samples used to measure changes in vote choice consistency and in the likelihood of stating a vote intention (Online Appendix Table A.11, columns (2) and (3)).

Respondents who only say that they intend to vote (and enter the sample) if they are surveyed close to the election can reasonably be expected to be less interested in politics and less consistent in their vote intention, on average. In particular, our data show that turnout overreporters, who may account for part of the increase in turnout intentions over time, are less consistent than actual voters. Therefore, if anything, we should expect changes in the composition of daily samples to bias the estimated increase in vote choice consistency downward. In fact, the observed increase in vote choice consistency is slightly larger when we alleviate changes in sample composition by keeping, instead of dropping, respondents who state that they are unlikely to vote, in surveys that record their vote intention (Online Appendix Table C.19, column (1)). More generally, Online Appendix C.6 shows that our results are qualitatively very similar when including unlikely voters.

Second, the sample we use for individual vote choice consistency and distance between predicted and final vote shares includes all respondents surveyed before and after the election. A possible concern is that the reinterview rate (the fraction of respondents surveyed in the pre-electoral survey who also appear in the postelectoral survey) is slightly lower for respondents who received the pre-electoral survey closer to the election (Online Appendix Table A.11, column (4)). However, the decrease in reinterview rate is much lower than the increase in vote choice consistency (−0.06 percentage point per day, on average, against 0.25),

16. Excluding respondents who stated that they were unlikely to vote in the pre-electoral survey ensures symmetry with postelectoral surveys, where nonvoters are usually not asked whom they would have voted for if they had voted, as well as homogeneity across surveys: while some surveys ask likely nonvoters whom they would vote for if they did vote, others only record the vote intention of respondents intending to vote.

so the former cannot explain much of the latter. Furthermore, reassuringly, the type of respondents who are not reinterviewed does not change over time. Using a specification in the form of equation (4) and a dummy equal to 1 if the respondent is reinterviewed as the outcome, we find that the coefficients on sociodemographic characteristics interacted with the distance between the pre-electoral survey and the election are all close to null and nonsignificant (Online Appendix Table A.12). Finally, we do not find any systematic pattern across surveys: while some surveys show a negative trend in reinterview rate, others show a positive trend, including within the same series of surveys. As shown in Online Appendix Table C.23, excluding the quartile of surveys with the largest decreases in reinterview rate yields a daily increase in vote choice consistency of the same magnitude as in Table I.

III.G. Alternative Interpretations

The increase in vote choice consistency and the convergence to final vote shares during the last two months before an election can only be used to assess the effect of information received by voters during this period if they are primarily driven by this information. All our results are consistent with this interpretation, and several heterogeneity analyses provide direct support for it, particularly the faster increase in vote choice consistency in U.S. swing states, in which campaign activities tend to be concentrated; during the final weeks before the election, when the campaign is most intense; and for younger and less educated voters, who have less preexisting information. Below, we discuss a complementary interpretation and state the case against three alternative ones.

1. *Cognitive Costs of Processing Information.* Beyond reflecting the influence of new information, changes in vote intentions during campaigns may plausibly also be due to delays in incorporating existing information into vote decisions. In this view, processing information and making one's vote choice takes time and effort, and the fraction of voters who have paid this cognitive cost increases over time. Then, the increase in vote choice consistency and the convergence to final vote shares shown above should be interpreted as an upper bound on the impact of information received during the campaign itself.

We see this interpretation as complementary to our main one. In both interpretations, vote choice consistency increases as

voters incorporate information in their choice of candidate. This view contrasts with the rival long-standing view holding that vote choices are based on partisan identification and ideology. Whether the information changing people's mind is received during the last 60 days before the election or earlier is more second order.

Furthermore, while postponing one's vote choice until shortly before the election could simply be a form of procrastination, it could also be a rational decision by voters who expect to receive useful information to the very last day. Conversely, information received during the campaign may resonate with prior information (Gennaioli and Shleifer 2010). For instance, voters may update their evaluation of a candidate negatively if new information about an unfolding crisis contradicts that candidate's earlier predictions. The increased consistency observed in the last few weeks should then be attributed jointly to people receiving new information and spending more time to process it in this period.

Finally, it is unlikely that delays in information processing could alone fully account for our results. Indeed, the next section provides evidence of convergence for outcomes that do not require that people exert cognitive skills but are likely to be affected by new information, namely, beliefs about candidates, as well as the perceived importance of different issues.

2. *Activation of Preexisting Partisan Loyalties.* An alternative interpretation is that much of the vote choice formation that we observe comes from the activation of preexisting partisan loyalties during the campaign, rather than the influence of information. For instance, voters may wait until later in the campaign to acknowledge that they will vote for the candidate of the party they have long felt close to. Although such a thought process is unlikely to explain voters switching between candidates, it could in principle contribute to increase the fraction of voters stating a vote intention (instead of saying they are undecided).

However, based on this interpretation, we should expect a large fraction of voters who make up their mind during the campaign to have strong partisan attachments and to end up voting for the party which they identified with from the beginning. The fact that the probability of stating a vote intention and vote choice consistency increase less for voters who strongly identify with a party, as shown in Section III.D, provides evidence against the first part of this prediction. We also reject the second part. We compare respondents' actual vote with the party which they say

that they identify with and estimate a specification in the form of equation (4), where we define people who voted for the party they identify with as type-a voters and those who voted for a different party as type-b. We find that the increase in vote choice consistency is lower by 64% or 57% among the former group, whether respondents' party of identification is defined based on their post- or pre-electoral answers (Online Appendix Table A.13).[17] These results suggest that vote choice formation results less from activating one's preexisting partisan ties than severing them.

3. *Information-Free Shocks.* Another alternative interpretation is that increased vote choice consistency stems from a stream of shocks containing no relevant information but affecting vote choice, for instance, by altering voters' utility function. Voters surveyed closer to an election may be more likely to vote for the candidate they announced simply because fewer shocks will hit them between their survey and the election.

This interpretation may be appealing for its simplicity, yet it is at odds with the observed convergence between predicted and final vote shares. Random shocks affecting different voters should cancel each other out in the aggregate. To account for the convergence to final vote shares, one would need to assume that shocks are correlated across voters. Shocks affecting the electorate at large may of course exist, and they play a central role in models such as probabilistic voting (Lindbeck and Weibull 1987). However, it is difficult to see how they could change the views of a large fraction of voters without containing any relevant information.[18]

17. A possible concern with defining party identification based on postelectoral answers is that voters may indicate that they identify with the party to which their vote choice converged during the campaign season, even though they did not identify with that party beforehand. Using pre-electoral answers yields another concern: party identification recorded before the election may change due to campaign information, leading the sample of people categorized as having voted for the party they identified with ex ante to also vary over time. If party identification did change due to the campaign, we should expect the consistency in this outcome to increase over time: voters surveyed later in the campaign would be more likely to have already converged to their final party identification. Reassuringly, party identification remains stable during the last two months before the election (Online Appendix Figure A.8).

18. In fact, many versions of the probabilistic voting model use political scandals or economic downturns, which are rich in information about candidates' probity and about the competence of the incumbent, as examples of aggregate shocks (e.g., Galasso and Nannicini 2017).

4. *Signals about the State of the World.* An alternative interpretation for our results is that voters do not make up their mind before the election based on information provided as part of the campaign but based on another type of information: signals about the underlying state of the world, which would have been released even absent the election. While such signals do not contain direct information on candidates, they could affect voters' views on their relative suitability to the current situation. For instance, voters may rally behind economically savvy candidates if they learn that the global economy is slowing down, affecting their country's economic outlook. In that case, it would be inaccurate to attribute the increase in vote choice consistency and decreased distance to final vote shares taking place in the last two months before the election to campaign effects.

However, signals about the state of the world are released all year long, not just during campaigns. Yet the convergence from vote intentions to vote choices mostly takes place shortly before the election, when the electoral campaign is most intense. Furthermore, natural and technological disasters are an example of signals about the state of the world which, although exogenous to the campaign, could plausibly increase the appeal of some candidates. But anticipating the results shown in Section V, we do not find any effect of such shocks.

IV. Formation of Beliefs, Policy Preferences, and Issue Salience

Information can affect vote choice through multiple mechanisms. To the extent that a campaign's influence hinges on the importance voters give to election-specific information about the candidates on the ballot versus their long-standing partisan leanings, one may expect the large increase in vote choice consistency shown in Section III to result from voters updating their beliefs about candidates' issue positions and quality. This may lead them to favor a new candidate that they find more competent or whom they find themselves ideologically closer to than they thought initially.

We compare the importance of this mechanism with three others. First, voters' own policy preferences might change over the course of the campaign, estranging them from certain candidates. Second, campaign information may increase the salience of some issues and prime voters to evaluate candidates based on them.

Finally, polls may clarify the identity of the front-runners and lead strategic voters to desert candidates with little chance of winning.

IV.A. Formation of Policy Preferences

To investigate the role played by policy preferences, we test whether the formation of vote choice is mirrored by changes in this outcome. We define individual i's preference consistency on question q as:

$$C_{it}^{q} = 1 - \frac{\left|\widetilde{A}_{it}^{q} - A_{it}^{q}\right|}{\overline{A^{q}} - \underline{A^{q}}},$$

where $\widetilde{A}_{it}^{q}$ (resp. A_{it}^{q}) is the respondent's answer to the question before (resp. after) the election and $\overline{A^{q}} - \underline{A^{q}}$ is the range of possible answers. This normalization allows us to use questions with different answer scales in the same regression.

We estimate equations (1) and (2) using preference consistency as the outcome.[19] As shown in Figure IV, Panel A and Online Appendix Table A.14, Panel A, policy preferences remain remarkably stable in the last two months before an election. The probability of stating a policy preference does not change over time either. Simply put, the persuasive communication voters receive in the campaign season does not alter their policy preferences.

The fact that vote choice formation during the campaign is not mirrored by a similar process of policy preference formation indicates, of course, that the latter is unlikely to explain the former. It also suggests that the reverse relationship, whereby voters adjust their policy preferences to their choice of candidate (Lenz 2012), does not play an important role in the elections we study. Furthermore, this result supports the assumption of stable preferences, which is a cornerstone of most models of electoral competition.

19. Our sample includes all respondents surveyed before and after the election who stated a policy preference in the second survey, and we replace election fixed effects with question fixed effects. Unlike vote choice, the accuracy of reported policy preferences should not depend on the time lapsed since the election. A more likely source of variation in policy preference changes is the distance between the pre- and postelectoral survey. Therefore, in this specification and in all regressions in this section, we control for dummies indicating the number of days separating the pre- and postelectoral survey instead of the distance between the latter and the election.

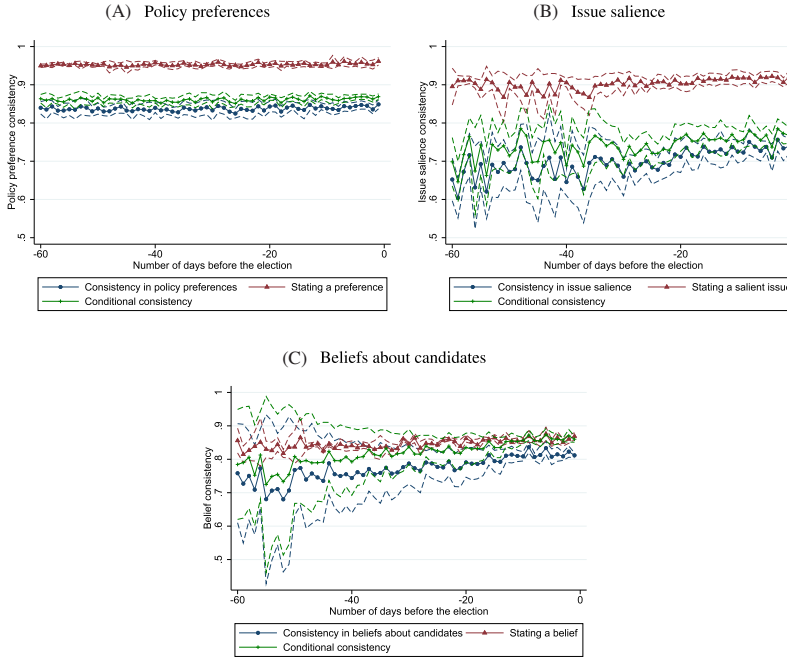

FIGURE IV

Consistency in Policy Preferences, Issue Salience, and Beliefs about Candidates

We show point estimates and 95% confidence intervals from specifications in the form of equation (1), regressing consistency in policy preferences (Panel A), issue salience (Panel B), and beliefs about candidates (Panel C) on 60 fixed effects indicating the number of days relative to the election and control variables listed in the text. We use one observation per respondent per policy question in Panel A, one observation per respondent in Panel B, and one observation per respondent per belief question in Panel C. In each panel, we consider three outcomes: consistency (N = 228,562; 46,108; and 478,039); stating a preference (resp. a salient issue or a belief) (N = 330,843; 60,713; and 809,037); and conditional consistency (N = 222,785; 44,049; and 440,771).

IV.B. Changes in Issue Salience

We use a similar method to explore the role of priming. We define issue salience consistency as a dummy equal to 1 if the respondent mentions an issue that they consider to be the most important in the pre-electoral survey and if they provide the same answer in the postelectoral survey.

As shown in Figure IV, Panel B, consistency in issue salience increases by 9 percentage points during the two months before voting. This outcome's daily increase is significant at the 5% level

and equal to more than half the daily increase in vote choice consistency shown in Table I (Online Appendix Table A.14, Panel B, columns (1) and (2)). It is driven by increases in the probability of stating an important issue and in consistency conditional on stating an issue (columns (3)–(6)).

These results indicate that electoral season information increases the salience of some issues while decreasing the importance of others, and they suggest that priming mechanisms may contribute to vote choice formation.

IV.C. Changes in Beliefs about Candidates

We finally investigate whether increased vote choice consistency in the period leading to an election is also driven by changes in beliefs voters hold about candidates.

Using questions recording beliefs about candidates' issue positions and quality, we find that the average daily increase in belief consistency in the last two months before the election is about two-thirds of the daily change in vote choice consistency but that it falls short of statistical significance (Figure IV, Panel A and Online Appendix Table A.14, Panel C, columns (1) and (2)).[20] The increase in the probability of stating a belief over time (instead of responding "I don't know") is estimated more precisely and significant at the 1% level (columns (3) and (4)). We observe similar patterns when we distinguish the beliefs voters hold about candidates' issue positions and about their quality: increases in respondents' likelihood to state a belief about candidates' issue positions and about their quality are of comparable magnitudes and are both significant at the 1% level (Online Appendix Table A.15).

We complement this evidence by comparing changes in vote choice consistency across supporters of different types of candidates. This approach is less direct, but it has the advantage of using all our data, not just surveys recording voters' beliefs both before and after the election. In addition, it enables us to

20. Questions recording beliefs about candidates are of two types. Some ask voters to select one of the candidates, for example, the candidate that talks the most about a particular issue. Belief consistency is then defined as a dummy equal to 1 if the respondent provides the same answer in the pre- and postelectoral survey. Other questions survey voters about a particular candidate, for example, how competent this candidate is. Belief consistency is then defined similarly as preference consistency, using the range of possible answers as the denominator. Our estimates pool both types of questions. Again, we replace election fixed effects with question fixed effects.

distinguish information and beliefs about candidates' positions and quality from beliefs about their relative chances (e.g., based on polls). The first type of information can be expected to benefit the candidates about whom voters have less prior knowledge, including candidates of new parties and small candidates, as well as any candidate challenging the incumbent. If changes in related beliefs are responsible for vote choice formation, people who eventually vote for lesser-known candidates should be more likely to make up their minds during the electoral season, and these candidates should see their predicted vote share increase over time. Instead, information on the relative chances of the contenders should lead strategic voters to rally to the strongest candidates, increasing these candidates' predicted vote shares and overall vote share concentration.[21] We test these opposite predictions.

First, we compare individual vote choice formation between voters who end up voting for well-established candidates and for initially lesser-known candidates. We estimate specifications in the form of equation (4), where we define as type-a people voting for a challenger, a small candidate, or the candidate of a new party; and as type-b people voting for the party that won the last election, an initially strong candidate (with an average predicted vote share larger than 10% in the first five days of the survey), or for a party that had competed in earlier elections. Vote choice consistency begins lower and increases faster among those who eventually vote for challengers, small candidates, and new parties (Online Appendix Table A.17).[22] These differential increases remain statistically significant in a specification controlling for all candidate types as well as their interaction with the time trend, with or without sociodemographic controls (columns (4) and (5)).

Second, we compare changes in the predicted vote share of different types of candidates in the 60 days leading up to the election, using the individual likelihood to vote for these candidates

21. We define voting strategically as voting based on likely outcomes of the election rather than expressively, for one's favorite candidate. Outside of strategic considerations, information on candidates' chances may affect voters if they use it as a signal of quality or if they desire to vote for the winner (Granzier, Pons, and Tricaud 2021).

22. The fractions of voters who report voting for a challenger, a small candidate, or a new party in the postelectoral survey are stable over time, suggesting that our results are not driven by changes in sample composition (Online Appendix Table A.16).

as the outcome, the same specification and sample of elections, and excluding respondents who do not state a vote intention or announce that they will cast a blank or null vote. We find that the support for challengers and small candidates increases as the election gets closer (Online Appendix Table A.18, columns (1)–(4)). We also measure changes in the following index of vote share concentration: $M_t^e = \sum_c \left(\widetilde{V}_{ct}^e\right)^2$. As shown in columns (7) and (8) and in Online Appendix Figure A.9, vote share concentration decreases, on average, in the last two months before an election.

Finally, changes in beliefs about candidates' relative chances should be more consequential for vote choice under plurality rule than under proportional rule, where the incentives to be strategic are weaker. But Online Appendix Table A.2 does not show a larger change in vote choice consistency in plurality elections.

These results converge to support the view that vote choice changes during campaigns are driven by changed beliefs about candidates' positions and quality more than about their chances. This conclusion echoes recent evidence showing that voters often behave expressively (Pons and Tricaud 2018).

IV.D. Intermediation Analysis

The fact that consistency in beliefs and issue salience, but not consistency in policy preferences, increase during the campaign indicates that changes in the two former outcomes (and not in the latter) may contribute to vote choice formation. To examine the relationship between these variables more directly, we compute mean vote choice consistency in each election on each day, and regress it on mean consistency in beliefs, issue salience, and policy preferences.

Formally, we estimate the following regression model:

$$(7)\quad \overline{C_t^e} = \rho_P \overline{C_t^{e,P}} + \rho_S \overline{C_t^{e,S}} + \rho_B \overline{C_t^{e,B}} + \sum_{t=-60}^{-1} \beta_t D_t + \alpha^e + W_t^{e'}\lambda + u_t^e,$$

where $\overline{C_t^e}$ is the mean vote choice consistency among respondents surveyed t days before election e and $\overline{C_t^{e,P}}$ (resp. $\overline{C_t^{e,S}}$ and $\overline{C_t^{e,B}}$) is the mean consistency in policy preferences (resp. in issue salience and in beliefs about candidates). This specification separates the influence of changes in preferences, beliefs, and issue salience consistency on vote choice consistency from two other important

sources of correlation. First, individual characteristics such as age and education may affect vote choice consistency and consistency in beliefs, preferences, and issue salience. The ensuing correlation should not affect our point estimates because our regression is at the day level and our daily samples are generally balanced on these characteristics. The same applies to other individual characteristics such as interest in politics, which we do not observe but which we can expect to be balanced over time in the rolling cross sections. Second, mean vote choice consistency and consistency in the regressors may increase over time as the election gets closer and more information becomes available, independently of any impact of the latter on the former. The 60 fixed effects indicating the number of days relative to the election, which are included in equation (7), control for this second source of correlation.

It remains that, in addition to capturing the impact of preferences, beliefs, and issue salience consistency on vote choice consistency, our coefficients of interest ρ_P, ρ_B, and ρ_S may also reflect the impact of information shocks deviating from the average 60 daily fixed effects and affecting vote choice consistency and consistency in our regressors, but whose effect on the former outcome is not mediated by the latter.

The results are reported in Table II, Panel A. We do not observe any significant relationship between consistency in policy preferences and in vote choice (column (1)). Instead, increases in mean issue salience and belief consistency are both associated with large increases in mean vote choice consistency (columns (2) and (3)). Both estimates are significant at the 5% level. They remain statistically significant (at the 5% and 10% level) and of very similar magnitude (0.25 and 0.13 percentage point) in a specification including all three variables (column (4)).

Similarly, Panel B shows that 1 percentage point increases in the likelihood of stating an important issue or a belief about candidates are associated with 0.20 and 0.24 percentage point increases in the likelihood of stating a vote intention, which are significant at the 1% level, but the probability of stating a policy preference is uncorrelated with this outcome.

While these point estimates do not necessarily represent causal evidence, they do support the conclusion that changes in beliefs about candidates and priming mechanisms are likely

TABLE II
DRIVERS OF VOTE CHOICE FORMATION

	(1)	(2)	(3)	(4)
Panel A: Mean vote choice consistency				
Mean consistency in policy preferences	−0.103			−0.138
	(0.114)			(0.120)
Mean consistency in issue salience		0.255**		0.247**
		(0.095)		(0.095)
Mean consistency in beliefs about candidates			0.134**	0.126*
			(0.064)	(0.064)
Observations	3,129	3,129	3,129	3,129
R^2	0.807	0.809	0.808	0.810
Election fixed effects	x	x	x	x
Fixed effects for number of days to election	x	x	x	x
Aggregate controls	x	x	x	x
Panel B: Mean probability of stating a vote intention				
Mean probability of stating a policy preference	0.005			−0.027
	(0.151)			(0.149)
Mean probability of stating a salient issue		0.200***		0.203***
		(0.053)		(0.052)
Mean probability of stating a belief on candidates			0.234***	0.237***
			(0.076)	(0.075)
Observations	3,144	3,144	3,144	3,144
R^2	0.840	0.840	0.841	0.841
Election fixed effects	x	x	x	x
Fixed effects for number of days to election	x	x	x	x
Aggregate controls	x	x	x	x

Notes. Standard errors clustered at the survey level are in parentheses. We estimate specifications in the form of equation (7), using one observation per election per day. In Panel A, mean vote choice consistency is computed based on all respondents surveyed before and after the election who said that they intended to vote in the first survey and reported that they actually voted and gave a vote choice declaration in the second. Mean consistency in policy preferences, mean consistency in issue salience, and mean consistency in beliefs about candidates are computed based on all respondents surveyed before and after the election who stated a policy preference, a salient issue, or a belief about candidates in the second survey. In surveys including multiple policy preferences or beliefs questions, consistency in policy preferences and consistency in beliefs are averaged at the respondent level before taking the mean across respondents surveyed on a given day. In Panel B, mean probability of stating a vote intention is computed based on all respondents surveyed before the election who said that they intended to vote. Mean probability of stating a policy preference, a salient issue, or a belief about candidates are computed based on all respondents surveyed before the election. We weight each observation by the number of respondents it was constructed from, relative to the overall number of respondents surveyed at the same relative time to the election. Aggregate controls include fixed effects for the day of the week in which the pre-electoral survey took place and the average number of days separating the postelectoral survey from the election. We also control for pre-electoral survey day-of-the-week fixed effects, the average postelectoral survey lag among respondents who received the pre-electoral survey at time t, and three dummy variables indicating whether $\overline{C_t^{e,P}}$, $\overline{C_t^{e,S}}$, and $\overline{C_t^{e,B}}$ are missing. ***, **, * indicate significance at 1%, 5%, and 10%, respectively.

to contribute to the formation of vote choice but that policy preferences do not.[23]

23. We note a small discrepancy between the unweighted (baseline) and weighted versions of these results. As shown in Online Appendix Table C.10, with survey weights, mean consistency in policy preferences also has a significant

These results are consistent with prior evidence established by Lenz (2012) regarding the drivers of changes in vote intentions, except for one important difference: that study does not find that changes in issue salience matter. A possible explanation is that Lenz (2012) does not directly measure the importance that voters give to different issues as we do, but infers it from the strength of the association between people's views on that issue and their vote intention. Furthermore, he explores specific shocks that can be expected to increase the salience of a certain issue, which makes it easier to study the causal relationship with changes in vote choice but also limits the range of priming effects that can be investigated. Our approach may capture the effects of changes in issue salience, which take multiple weeks or even the entire campaign to fully materialize.

V. Impact of TV Debates

Finally, we investigate which of three distinct sources of information are responsible for the formation of vote choice.

Given the mediating role of beliefs about the quality and positions of candidates established in Section IV, a first plausible hypothesis is that changes in these beliefs and, in turn, some of the changes in vote intentions occur as a result of communication from candidates themselves, especially during TV debates. Direct messages from charismatic candidates may be more persuasive than indirect communication organized by their campaign, for example, with rank-and-file canvassers knocking on doors on their behalf. Beyond TV debates, candidates can also communicate with voters through campaign rallies and ads in various media. However, TV debates are more salient, and they attract much larger audiences than rallies.[24] Overall, 71% of respondents report watching TV debates in our sample, and official TV ratings that we found for 43 out of the 56 TV debates indicate that debate

effect on mean vote choice consistency. However, the result that policy preference consistency does not increase during the campaign remains unchanged, so our conclusion that the observed increase in vote choice consistency is unlikely to be driven by changes in policy preferences holds.

24. For instance, while Donald Trump's 2016 rallies attracted unusually large crowds, the number of participants in these rallies reached a few tens of thousands at most, against an estimated 84 million for the first debate pitting him against Hillary Clinton.

watchers account for an average of 22% of the population.[25] Furthermore, voters may find TV debates more informative than scripted rally speeches or TV ads because debates subject all candidates to the same exercise and test their ability to respond in real time to moderators' unexpected questions and opponents' arguments. On the other hand, voters may deem any information provided by candidates about themselves, including through debates, not credible. Indeed, models of Bayesian persuasion suggest that individuals may not be convinced by the information they receive if they believe that the sender has incentives to manipulate it (DellaVigna and Gentzkow 2010; Kamenica and Gentzkow 2011).

Second, we compare TV debates' effects with the effects of shocks which, instead, occur independently from the campaign and are entirely outside of candidates' control, such as natural and technological disasters. The literature exploring effects of such factors using local variation (Cole, Healy, and Werker 2012) emphasizes that these factors may change vote intentions for irrational reasons (Achen and Bartels 2016) but that they can also provide valuable information on candidates. For instance, voters may learn about incumbents' preparedness and competence from the damages caused by unexpected natural disasters and from the policies announced and implemented in response to them (Ashworth, Bueno de Mesquita, and Friedenberg 2018), and they may adjust their beliefs about other candidates based on statements they issue in reaction to these shocks.

Third, voters may be influenced by information they receive as part of the campaign but that candidates do not directly control, including messages from the media, political activists, and discussions with friends or family members. Voters may perceive these third parties as more credible and less likely to manipulate information than candidates, generating larger effects.[26]

25. The sources of TV ratings are provided in Online Appendix B.4. The discrepancy between self-reports and official ratings may be explained by several factors. First, TV ratings and the estimate based on our surveys are about different objects, watching a specific debate versus watching debates in general. Second, TV ratings include all viewers above 2 years old in the United States, and above 14 in other countries, whereas self-reports are only based on survey respondents above 18 and who intend to vote in the upcoming election. Third, elections with a larger fraction of debate watchers had more respondents. Fourth, TV ratings may underestimate the number of actual debate viewers if they fail to account for those watching in public settings, and survey respondents may overreport watching debates to signal their interest in politics.

26. The prediction that receivers respond more when messages are credible has received empirical support; for instance, in Chiang and Knight (2011).

Indeed, it may be more costly for them to lie (Cho and Kreps 1987); they often share common interests with the voters they are trying to persuade (Crawford and Sobel 1982); and their information is more incomplete, making selective transmission more difficult (Fischer and Stocken 2001). Candidates may be able to influence the information provided by these sources, but only to a certain extent. Indeed, many media outlets remain reluctant to become the mouthpiece for a specific camp, and voters debating the merits of the different candidates will likely misremember some of their propositions. Even political activists campaigning on candidates' behalf may convey a different message to voters than the campaign's official talking points (Enos and Hersh 2015).

Online Appendix Figure A.10 and Online Appendix Table A.19 show the change in the fraction of voters who report getting information frequently from newspapers, TV, radio, and the Internet, having seen election polls recently, discussing politics frequently with others, and having been contacted or visited by a party recently. All these outcomes build up over the electoral season, suggesting that the corresponding sources of information may contribute to the concomitant increase in vote choice consistency. On the other hand, the slow and continuous convergence of vote intentions to final vote choices observed across elections does not preclude the possibility that discrete events such as TV debates and natural disasters, taking place at different times in each race, play a decisive role by changing the mind of many voters. We use an event study approach to estimate their effect.

V.A. Estimation Strategy

Holding TV debates before national elections is the norm in a growing number of countries. After the first presidential TV debate between Richard Nixon and John F. Kennedy in 1960, TV debates were held in each U.S. election beginning in 1976, and the practice quickly spread to other countries. Debates' ubiquity warrants a multicountry event study to examine their impact.

In some elections, multiple debates take place. We exclude debates held less than three days from one another, to be able to estimate effects up to three days after.[27] This leaves us with a total of 56 debates. Debates in our sample were held between 5 and 44 days before the election, with an average of 24 days before

27. The results are robust to an extended model excluding debates held less than five days from one another (Online Appendix C.8).

(Online Appendix Figure A.11). They are concentrated in the period when vote choice consistency increases the fastest, making them as plausible a driver of vote choice formation as the sources of information shown in Online Appendix Figure A.10. The full list of debates is included in Online Appendix Table B.6, along with the following information: whether the debates featured candidates for president or prime minister, their date, and the time to the election.

An observation is a respondent × debate × election. A few debates affect several elections, namely, the first and second votes in Germany, and the electoral vote and party vote in New Zealand. In addition, for each debate, our estimation uses all respondents in the corresponding survey. As a result, the same response is included multiple times when multiple debates were held before an election. In total, our sample includes 331,000 observations. We cluster standard errors at the debate level to adjust for the correlation between the error terms of all observations related to the same debate. This clustering also accommodates for the fact that some debates preceding the 2000, 2004, and 2008 U.S. presidential elections are covered both by the ANES and the NAES.[28]

Our main specification is as follows:

(8)
$$Y_{it}^{d} = \sum_{k=-3}^{-1} \mu_k + \sum_{k=1}^{3} \mu_k + \mu_{4-} + \mu_{4+} + \sum_{t=-1}^{-60} \beta_t D_t + \alpha^d + W_{it}^{'}\lambda + u_{it}^{d},$$

where Y_{it}^{d} is the outcome for respondent i, surveyed t days before the election corresponding to debate d, μ_k $(-3 \leqslant k \leqslant 3)$ are dummies indicating the number of days relative to the debate, μ_{4-} and μ_{4+} are dummies equal to 1 for respondents surveyed four days or more before or after the debate, respectively, and α^d are debate × election fixed effects.[29]

The key coefficients of interest are μ_1, μ_2, and μ_3, which measure the impact of debates one to three days after, relative to the

28. As in Section III, we also check the robustness of our results to allowing for correlation of the error terms with the wild cluster bootstrap procedure (Online Appendix Tables C.3 and C.4) and to clustering the standard errors at the level of the election date (Online Appendix Tables C.7 and C.8).

29. We include separate fixed effects for distinct elections affected by the same debate. We also include two separate fixed effects for U.S. debates covered both by ANES and NAES.

omitted category μ_0.[30] As this specification makes clear, we do not use elections without debates as our counterfactual. The mere fact of having debates may change the kinds of candidates chosen by parties, the overall amount and type of information provided during the campaign, and how much voters pay attention to it. Our estimates do not capture such general equilibrium mechanisms potentially affecting all voters. Instead, we estimate direct effects of debates on voters surveyed afterward relative to those surveyed beforehand. This enables us to assess the extent to which TV debates contribute to the increase in vote choice consistency observed in the corresponding elections. Our estimates capture effects of debates as well as effects of subsequent media commentaries and discussions, on both debate watchers and nonwatchers. In Section V.D, we separate the effects on these two groups.

Importantly, the fact that debates took place at different times in different elections allows us to control flexibly for the number of days relative to the election, with the 60 daily fixed effects D_t. This is critical to disentangle the effect of debates from the underlying time trends shown in Section III. In addition, the vector W_{it} controls again for day-of-the-week and postelectoral survey lag fixed effects and, in some specifications, for sociodemographic characteristics.

Our identifying assumption is that conditional on all these controls, and conditional on having a TV debate during our observation window, the date of the debate is uncorrelated with the outcome. In addition, we assume that any pretrend before the fourth day preceding the debate or any impact after the fourth day following it are accurately captured by the fixed effects μ_{4-} and μ_{4+}.[31]

30. We use the day of the debate as the reference group because debates take place in the evening. Therefore, the vast majority of respondents surveyed on that day are surveyed before the debate. The exact time of the interview is available for 4,095 respondents surveyed on the day of 26 different debates. We find that only 16% of them were surveyed after the debate started.

31. Our results pointing to the lack of increase in vote choice consistency and the lack of decrease in the distance to final vote shares are robust to an alternative specification that does not require this assumption because it uses a sample restricted to a balanced panel of observations for each of the three days preceding and following each debate and excluding all respondents surveyed before or after (Online Appendix C.9). The drawback of that specification is that we can only control for election fixed effects instead of debate × election fixed effects because collinearity prevents estimating a full set of debate fixed effects, time fixed effects, and fixed effects for days relative to the debate.

There are three important potential threats to our identification strategy. First, systematic differences in the characteristics of respondents surveyed before and after debates would violate the identifying assumption and could lead us to mistakenly attribute to debates changes in outcomes originating in sample composition differences. The fact that most of our surveys are rolling cross sections, which allocate respondents' survey date randomly, alleviates this risk. It remains that debates, like other campaign events, may affect the characteristics of people willing to answer the survey. To address this concern, Online Appendix Table A.20 reports balance checks for voter characteristics as well as watching debates. Out of 80 differences, 9 are significant at the 10% level, 5 at the 5% level, and none at the 1% level, which is in line with what would be expected.

A second potential risk arises if unexpected shocks occurring on the same day or immediately before or after the debate bias our estimates. This risk is important for existing studies that use pre/post difference designs and focus on a unique debate or a few debates only. In our case, such shocks would only violate the identifying assumption if they were systematically correlated with debates' dates. Given the large number of debates in the study, and conditional on the daily fixed effects and other controls, this should not be the case.

The third potential violation of the identifying assumption comes from the fact that, of course, debates do not happen unexpectedly. Their dates are known long in advance, so candidates and the media may strategically time their communication around them. This could generate continuous trends in outcomes around debates, which the predebate dummies μ_{-3}, μ_{-2}, and μ_{-1} allow us to test for. However, these dummies would not capture changes only taking place after the debate. One possibility is that the amount of information increases (or that voters pay significantly more attention to it) after the debate, biasing our estimates upward. Given our mostly null results, changes susceptible to bias our estimates downward would be more concerning. Downward bias could occur if candidates decreased the intensity of their campaign, and if media decreased their coverage thereof, after the debate (e.g., because they anticipate debate-related information will lower the returns of any other type of communication), or if voters decreased their media consumption.

Online Appendix Table A.21 tests for changes in media consumption and partisan communication around debates. Columns

(1)–(4) use dummies for getting information frequently from newspapers, TV, radio, and the Internet as outcomes. None of the predebate or postdebate dummies are significant, providing evidence of stable media consumption around debates. Columns (5) and (6) show no significant effect on the probability of having seen election polls recently or of discussing politics frequently with others. Finally, we obtain only nonsignificant coefficients when using dummies for having been contacted and having been visited by a party as outcomes (columns (7) and (8)), except for a small decrease in party contact the day after the debate, significant at the 10% level. Although these results support our identifying assumption, we note that the object of all questions—either overall media consumption or having been contacted by a party recently, not just on the day of the survey—limits the power of these tests.

The next two subsections measure mean effects of debates on individual and aggregate outcomes. We then explore potential sources of heterogeneity in debate impact.

V.B. Debates' Effects on Individual Outcomes

We first measure the effect of TV debates on our main outcome, C_{it}, the individual consistency between vote intention and vote choice.[32]

This outcome comparing pre- and postelectoral survey responses is well suited to our event-study design. We would not be able to measure the impact of TV debates using only postelectoral responses, since vote choices reported by all respondents may reflect debates' influence. Instead, if debates do help voters decide between candidates, we should expect the fraction of people stating a vote intention identical to their eventual vote to be higher among those who answered the pre-electoral survey right after the debate than those surveyed right before.

32. Once again, our sample includes all respondents surveyed before and after the election who said that they intended to vote in the first survey and reported that they actually voted and gave a vote choice declaration in the second. TV debates marginally affect selection into this sample, with one positive coefficient three days after the debate, significant at the 10% level (Online Appendix Table A.22, column (3)). However, column (1) shows no significant effect on turnout intention, and Online Appendix Tables C.21 and C.22 show that our main findings are robust to including unlikely voters in the sample, alleviating the concern that our results may be biased by differential sample selection. Furthermore, column (4) shows that the postelectoral reinterview rate is balanced around debates.

Furthermore, previous studies estimating the effects of debates and other campaign events have used pre-electoral survey responses, but focusing instead on vote choice consistency should help uncover effects which these outcomes could miss.

To see why, first note that the types of candidates benefiting from TV debates may vary across elections and debates. Using vote intentions as outcome, these effects could get netted out when pooling multiple debates in the same event study. For instance, suppose that in each debate a unique candidate—the incumbent in half of the debates and a challenger in the other half—wins over some voters from the other side. Debates truly change the course of every race, yet overall effects measured on intending to vote for the incumbent would be null. Using as outcomes intention to vote for left-wing versus right-wing candidates or for outsiders versus front-runners could generate similarly misleading null effects. Instead, vote choice consistency, which in this example increases following each debate (because voters persuaded by debate winners are found inconsistent if surveyed before but consistent if surveyed after), would show a positive effect.

Second, the effects of debates may also vary across voters in an election. Once again, effects benefitting different candidates could be netted out using traditional outcomes. Consider a debate increasing the likelihood that voters of opposing sides all express support for the nominee of their preferred party. Our measure of individual vote choice consistency would capture this effect even if the net impact on individual vote intention and predicted aggregate vote shares were null. The ability to detect any type of effects is desirable in general and especially useful here: because it works against finding a null, it only makes our mostly null results more trustworthy.

We report the coefficients on the μ_k dummies indicating the number of days relative to the debate in Table III, column (1), and plot them in Figure V. We do not observe any pretrend in vote choice consistency in the three days preceding debates. The dummies for the days following debates are also all close to zero and nonsignificant. On average, debates decrease individual vote choice consistency by a nonsignificant 0.7 percentage point in the three following days. Estimates of overall vote choice formation during the final two months of campaigns shown in Section III

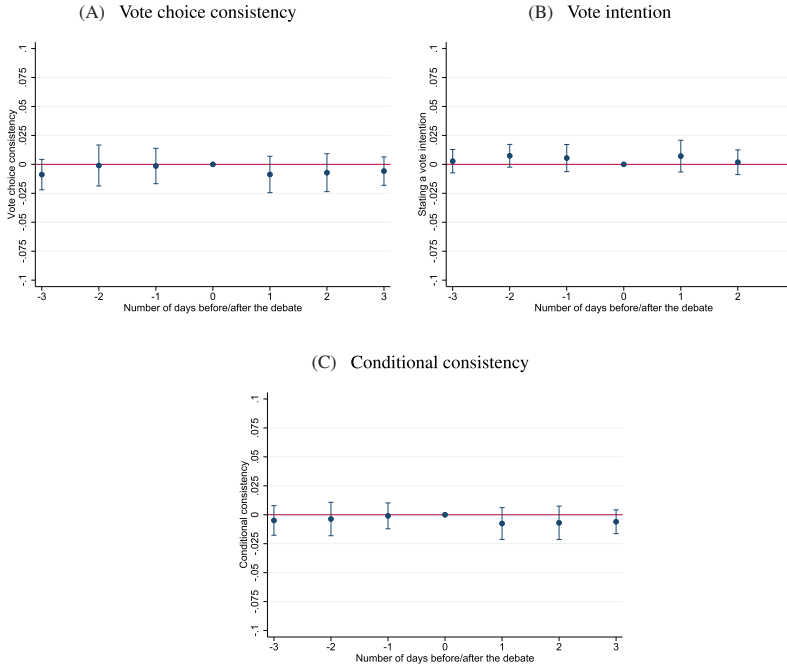

FIGURE V

Debates' Effects on Vote Choice Consistency

We show point estimates and 95% confidence intervals from specifications in the form of equation (8), regressing vote choice consistency, vote intention, and conditional vote choice consistency on dummy variables for being surveyed one, two, or three days before the debate, as well as dummies for being surveyed one, two, or three days after the debate. We also include dummies for being surveyed four days or more before or after the debate, respectively, and omit the dummy for being surveyed on the day of the debate. We control for debate × election fixed effects and effects for the number of days relative to the election and for the day of the week in which the pre-electoral survey took place. In Panels A and C, we also control for fixed effects for the number of days separating the postelectoral survey from the election. Standard errors are adjusted for clustering at the debate level. $N = 263{,}681$; 330,621; and 240,826, respectively.

provide a useful benchmark to interpret this result. Considering the upper bound of the 95% confidence interval, we can reject any impact of debates on vote choice consistency higher than 0.5 percentage point at the 5% level, which corresponds to 3% of the overall 17 percentage point increase over the electoral season

TABLE III
DEBATES' EFFECTS ON VOTE CHOICE CONSISTENCY

	Vote choice consistency		Vote intention		Conditional consistency	
	(1)	(2)	(3)	(4)	(5)	(6)
Before −3	−0.014**	−0.015**	0.003	0.002	−0.015***	−0.016***
	(0.006)	(0.007)	(0.005)	(0.005)	(0.005)	(0.006)
−3	−0.009	−0.010	0.003	0.003	−0.005	−0.006
	(0.007)	(0.007)	(0.005)	(0.005)	(0.006)	(0.007)
−2	−0.001	−0.002	0.007	0.007	−0.004	−0.004
	(0.009)	(0.009)	(0.005)	(0.005)	(0.007)	(0.007)
−1	−0.001	−0.002	0.005	0.005	−0.001	−0.002
	(0.008)	(0.008)	(0.006)	(0.006)	(0.006)	(0.006)
+1	−0.009	−0.009	0.007	0.007	−0.008	−0.008
	(0.008)	(0.008)	(0.007)	(0.007)	(0.007)	(0.007)
+2	−0.007	−0.008	0.002	0.001	−0.007	−0.008
	(0.008)	(0.008)	(0.005)	(0.005)	(0.007)	(0.007)
+3	−0.006	−0.006	0.007	0.007	−0.006	−0.006
	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
After +3	−0.002	−0.003	0.008	0.008	−0.005	−0.005
	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)

TABLE III
CONTINUED

	Vote choice consistency		Vote intention		Conditional consistency	
	(1)	(2)	(3)	(4)	(5)	(6)
Observations	263,681	263,681	330,621	330,621	240,826	240,826
R^2	0.072	0.077	0.069	0.076	0.042	0.044
Mean, day of the debate	0.811	0.811	0.896	0.896	0.887	0.887
Number of debates	56	56	56	56	56	56
Debate × election fixed effects	x	x	x	x	x	x
Fixed effects for number of days to election	x	x	x	x	x	x
Individual controls	x	x	x	x	x	x
Sociodemographic controls		x		x		x
Linear combination of estimates						
Average predebate dummies −3, −2, and −1	−0.004	−0.005	0.005	0.005	−0.003	−0.004
	(0.007)	(0.007)	(0.004)	(0.004)	(0.005)	(0.005)
Average postdebate dummies 1, 2, and 3	−0.007	−0.008	0.005	0.005	−0.007	−0.007
	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)

Notes. Standard errors clustered at the debate level are in parentheses. We use one observation per respondent and estimate specifications in the form of equation (8). Individual controls include fixed effects for the day of the week in which the pre-electoral survey took place and, in columns (1), (2), (5), and (6), fixed effects for the number of days separating the postelectoral survey from the election. Sociodemographic controls include education (dummies indicating above high school education and college degree), gender, age, income quartiles, and employment status. The mean values of the three predebate dummies and of the three postdebate dummies are also reported, along with their standard errors. ***, **, * indicate significance at 1%, 5%, and 10%, respectively.

(Figure I).[33] These results are robust to controlling for respondents' sociodemographic characteristics (column (2)).[34]

We next measure the impact of debates on stating any vote intention[35] and on vote choice consistency conditional on stating one. We do not observe any pretrend in the likelihood to state a vote intention, and the impact on this outcome is not significant on any day after the debate (Figure V, Panel B and Table III, columns (3) and (4)). The average effect of debates on the likelihood to state any vote intention in the three following days is positive but small and not significant in either specification shown in columns (3) and (4). None of the three predebate or postdebate dummies are significant for conditional vote choice consistency either, and the effect of debates on this outcome in the three following days is negative and nonsignificant (Figure V, Panel C and Table III, columns (5) and (6)).

Finally, as shown in Online Appendix Table A.23 as well as in Online Appendix Figure A.12, none of the relative days dummies are significant when using consistency in issue salience or in beliefs about candidates as the outcome, suggesting that debates do not affect the factors found to be likely contributors to vote choice formation in Section IV. Effects on policy preference consistency are not significant either.

All our point estimates measure the impact of debates in the short run. If our estimates were positive and significant, we could be worried about potential subsequent reversion to the

33. The increase in vote choice consistency during the final two months of campaigns is nearly identical in elections with a TV debate as in the full sample.

34. A potential concern is that our null effects may hide the fact that debates increase the choice consistency of some while decreasing that of others. This scenario is unlikely but not impossible. Consider a debate in which a charismatic candidate seduces both voters from their own party and from a rival party. The first group was previously intending to vote for other candidates but they now intend to vote for their party's candidate and will stick to this choice until the election, so that their vote choice consistency increases after the debate. Further assume, by contrast, that the second group of voters from the rival party only temporarily depart from their intention to vote for their candidate, so their vote choice consistency temporarily decreases after the debate. Overall, changes in vote intentions among both groups of voters would lead to null effects on vote choice consistency. However, they would increase the predicted vote share of the charismatic candidate immediately after the debate, which would be captured by the second aggregate outcome examined in Section V.C.

35. This outcome is defined on the sample of respondents who said that they intended to vote. Online Appendix Table A.22, column (2) shows that TV debates do not affect selection into this sample.

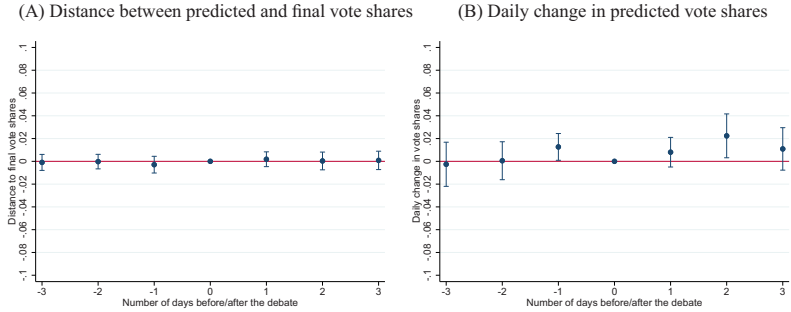

FIGURE VI

Debates' Effects on Aggregate Vote Shares

We show point estimates and 95% confidence intervals from specifications in the form of equation (9), using the distance between predicted and final vote shares and the daily change in predicted vote shares as outcomes. We control for debate × election fixed effects as well as fixed effects for the number of days relative to the election and for the day of the week in which the pre-electoral survey took place. We also control for the average number of days separating the postelectoral survey from the election in Panel A. $N = 3{,}802$ and 3,749, respectively. Other notes are as in Figure V.

mean. Instead, it seems unlikely that the null effects we measure shortly after the debate give way to large effects later on.

V.C. Debates' Effects on Aggregate Outcomes

Debates' lack of effect on individual vote choice consistency does not necessarily preclude effects on aggregate vote shares. Indeed, debates may lead some voters to change their views without fully converging on their final vote choice yet, and nonetheless reduce the distance to final vote shares.

We measure debates' effects on ΔV_t, the overall distance between predicted and final vote shares defined in Section III.E, with a specification using only one observation per debate per day:

(9)
$$\Delta V_t^d = \sum_{k=-3}^{-1} \mu_k + \sum_{k=1}^{3} \mu_k + \mu_{4-} + \mu_{4+} + \sum_{t=-1}^{-60} \beta_t D_t + \alpha^d + W_t' \lambda + u_t^d.$$

As shown in Figure VI and Table IV, columns (1) and (2), all pre- and postdebate relative-days dummies are close to zero and nonsignificant. The average effect on this outcome in the three days following debates is positive, small (0.1 percentage point),

and not statistically significant. Considering the lower bound of the 95% confidence interval, we can reject any impact lower than −0.1 percentage point at the 5% level, which corresponds to 2% of the overall 5 percentage point decrease in the distance between predicted and final vote shares over the electoral season (Figure III). This result is unchanged when controlling for sociodemographic variables (column (2)).[36]

Second, we test whether debates generate short-term shifts in aggregate vote intentions, which may be the case even if they do not contribute to the convergence to final vote shares. We define the overall daily change in predicted vote shares as $\delta V_t^d = \frac{1}{2}\sum_c \left|\widetilde{V}_{ct}^d - \widetilde{V}_{ct-1}^d\right|$, where $\widetilde{V}_{ct}$ is the predicted vote share of candidate c among time t respondents. δV_t corresponds to the minimal share of voters who had to change their vote intention to explain the difference between predicted vote shares' distributions at time t and $t-1$. We estimate equation (9) using this outcome, and we show the results in Figure VI, Panel B and Table IV, columns (3) and (4).[37] We find some evidence that debates increase the daily change in predicted vote shares: the dummy for the second day after the debate is significant at the 5% level, and the average effect in the three days following the debate is 1.4 percentage points, significant at the 10% level. The average of the three postdebate dummies is no longer significant but of similar magnitude in the specification including sociodemographic controls (column (4)). However, the dummy for the day preceding the debate is also positive and statistically significant. The postdebate increase may thus be driven in part by an unusually low change in vote shares on the day of the debate.

Debates' positive effects on the daily change in vote shares, together with their (nonsignificant) negative effects on individual vote choice consistency, suggest that, if anything, they move a small fraction of voters away from their final choice, in the short term. Overall, while debates may generate short-term shifts in vote shares, these do not contribute to the overall increase in vote choice consistency established in Section III.

36. As in Section III.E, we weight each observation by the number of t respondents it was constructed from, relative to the total number of respondents surveyed at the same time before or after the debate.

37. In this specification, we weight each observation by the number of t and $t-1$ respondents it was constructed from, relative to the total number of respondents surveyed at the same time before or after the debate.

TABLE IV

DEBATES' EFFECTS ON AGGREGATE VOTE SHARES

	Distance between predicted and final vote shares		Daily change in predicted vote shares	
	(1)	(2)	(3)	(4)
Before −3	0.005	0.006	0.009	0.007
	(0.004)	(0.004)	(0.006)	(0.006)
−3	−0.001	0.000	−0.003	−0.004
	(0.004)	(0.004)	(0.010)	(0.010)
−2	−0.000	0.000	0.001	−0.001
	(0.003)	(0.003)	(0.008)	(0.009)
−1	−0.003	−0.002	0.013**	0.012*
	(0.004)	(0.004)	(0.006)	(0.006)
+1	0.002	0.002	0.008	0.006
	(0.003)	(0.003)	(0.006)	(0.007)
+2	0.000	0.001	0.022**	0.021**
	(0.004)	(0.004)	(0.010)	(0.009)
+3	0.001	0.002	0.011	0.009
	(0.004)	(0.004)	(0.009)	(0.009)
After +3	0.002	0.002	0.006	0.005
	(0.003)	(0.004)	(0.005)	(0.005)

TABLE IV
CONTINUED

	Distance between predicted and final vote shares		Daily change in predicted vote shares	
	(1)	(2)	(3)	(4)
Observations	3,802	3,802	3,749	3,749
R^2	0.575	0.581	0.470	0.477
Mean, day of the debate	0.046	0.046	0.085	0.085
Number of debates	56	56	56	56
Debate × election fixed effects	x	x	x	x
Fixed effects for number of days to election	x	x	x	x
Aggregate controls	x	x	x	x
Sociodemographic controls		x		x
Linear combination of estimates				
Average predebate dummies −3, −2, and −1	−0.001	−0.000	0.004	0.002
	(0.003)	(0.003)	(0.007)	(0.007)
Average postdebate dummies 1, 2, and 3	0.001	0.002	0.014*	0.012
	(0.003)	(0.003)	(0.008)	(0.008)

Notes. Standard errors clustered at the debate level are in parentheses. We use one observation per election per day and estimate specifications in the form of equation (9). Aggregate controls include fixed effects for the day of the week in which the pre-electoral survey took place and, in columns (1) and (2), the average number of days separating the postelectoral survey from the election. Sociodemographic controls are averages of the sociodemographic variables included in Table III. Other notes are as in Table III. ***, **, * indicate significance at 1%, 5%, and 10%, respectively.

V.D. Heterogeneous Effects of Debates

The mostly null average effects of debates reported heretofore do not rule out the possibility that debates matter relatively more in certain conditions or for certain groups of voters. We first study the potential mediating influence of debates' timing and election type. Debates may affect vote choice more in contexts in which preexisting knowledge about the candidates is low or incentives to pay attention are high, for instance because the race is tight. To test this hypothesis, we compare the first debate of each race, when voters do not know much about the candidates, to debates taking place later, when the fraction of voters who have already arrived at their final choice is higher; close races to expected landslides; highly fluctuating races to more stable races; the U.S. bipartisan elections to multiparty elections, in which a larger fraction of voters form their vote choice shortly before the election (as shown in Section IV); and plurality rule to proportional rule elections.

Formally, we interact the relative days dummies with indicators of debate type or election type:

$$
\begin{aligned}
Y_{it}^{d} = & \left[\sum_{k=-3}^{-1}\mu_k + \sum_{k=1}^{3}\mu_k + \mu_{4-} + \mu_{4+}\right] \times \Omega^{d} \\
& + \left[\sum_{k=-3}^{-1}\eta_k + \sum_{k=1}^{3}\eta_k + \eta_{4-} + \eta_{4+}\right] \times \left(1 - \Omega^{d}\right) \\
& + \sum_{t=-1}^{-60}\beta_t D_t + \alpha^{d} + W_{it}^{'}\lambda + u_{it}^{d},
\end{aligned} \tag{10}
$$

where Ω is a dummy equal to 1 for "type-a" races or debates, for which effects may be expected to be larger, and 0 otherwise. We interact the μ_k's and η_k's with Ω and $(1 - \Omega)$, respectively, to directly test the null that neither type-a nor type-b debates have significant effects.

While a few pre- and postdebate relative days dummies are statistically significant, as would be expected given the large number of tests, we do not find any positive and significant effect on vote choice consistency in the three days following early or later debates or debates held in any subset of races we examine (Online Appendix Table A.24). This is true whether we consider the day-related dummies separately or take their average.

We use a similar method to study treatment impact heterogeneity on the overall distance between predicted and final vote shares. Considering the average of the three postdebate

dummies, we do not find any significant effect on the distance to final vote shares of any type of debate, or of debates held in any type of race (Online Appendix Table A.25). Only two individual postdebate dummies are negative and significant, corresponding to the effects of first debates and U.S. debates, two days after the debate (columns (1) and (4)). However, in both cases, the coefficients for one and three days after the debate and, again, the average of the three postdebate coefficients are nonsignificant. In addition, the lower distance between predicted and final vote shares observed two days after the debates is not mirrored by any substantial increase in individual consistency in either case (Online Appendix Table A.24, columns (1) and (4)).

While we explore heterogeneity along a large number of dimensions, we still may have failed to consider the one dimension that truly matters. Therefore, we make a final attempt using the following, more agnostic approach. We estimate a specification in the form of equation (8), where each relative-day dummy is interacted with a full set of debate indicators, yielding a specific set of coefficients μ_{4-}, μ_{-3}, ..., μ_{3}, and μ_{4+} for each debate. The mean values of the three debate-specific postdebate dummies μ_1, μ_2, and μ_3 are plotted in Online Appendix Figure A.13, for each debate separately and in ascending order. Since the number of observations corresponding to a specific debate and relative day is small, and we have a unique cluster for each debate, we do not report confidence intervals for the debate-specific estimates and refrain from interpreting these estimates individually. Rather, we are interested in the overall shape of the distribution and in the possible presence of outliers—that is, exceptional debates which, unlike the average event, may have changed the course of the corresponding election. We do not find evidence of such events. Instead, we observe that estimates are centered around zero (out of the 52 estimates, 30 are negative and 22 are positive) and smoothly distributed, and we do not detect any clear outlier. Furthermore, the effects of debates held in each country tend to span the full range, with both negative and positive values.

In Online Appendix Figure A.14, we repeat this exercise, using the distance between predicted and final vote shares as the outcome. Once again, the mean debate-specific effects are centered around zero, with no clear outlier. Overall, these two

figures strengthen our conclusion that debates' null effects hold across a large number of settings.[38]

We now explore treatment impact heterogeneity along voter characteristics with a specification in the form of equation (10). First and foremost, we measure effects separately for voters who report watching debates and those who do not. We do not separate watchers from nonwatchers based on information recorded in the pre-electoral survey, as this may generate different splits among people surveyed before and after the debate. Instead, we use postelectoral survey questions recording whether the respondents watched any of the debates held before that election. This information is available for half of the debates. Debates could also plausibly have larger effects on voters whom we found to be more likely to form their vote choice during the campaign in Section III.D: voters with weak or no party identification, who may be freer to switch candidates, and those without college education and young voters, who are likely to be less informed before the debate.

Using these four variables and our other sociodemographic characteristics as mediating factors, and considering average effects in the three days following debates, we do not find any significant and positive effect on vote choice consistency or any significant and negative effect on the convergence to final vote shares for any of the 14 subgroups of voters we consider (Online Appendix Tables A.26 and A.27). The only exception is a negative effect, significant at the 5% level, on the distance to final vote shares for debate watchers. However, point estimates of similar magnitude and identical sign on predebate dummies for this group suggest that this effect is spuriously driven by unusually high distance to final vote shares among debate watchers surveyed on the day of the debate. In addition, this pattern is not mirrored by an increase in vote choice consistency: on average, the sign of the effect on the latter outcome is negative for debate watchers (Online Appendix Table A.26, column (1)). Overall, we do not find any clear evidence that debates contribute to the process of vote choice formation for any type of voters.

38. Consistent with the lack of heterogeneous effects across debates, our main null results are robust to using the difference-in-differences estimator from de Chaisemartin and D'Haultfœuille (2020), which eliminates any bias coming from heterogeneous effects and from dynamic effects over time (Online Appendix C.10).

Finally, we test whether debates systematically benefit some candidates at the expense of others and, in particular, whether they contribute to the increase in the vote share of lesser-known candidates, shown in Online Appendix Table A.18. We first run a specification in the form of equation (8) for each type of candidate, using a dummy equal to 1 if the respondent intends to vote for them as the outcome. Then, we run a specification in the form of equation (9) to estimate the impact on vote share concentration. Debates affect significantly neither the predicted vote share of challengers, small candidates, or candidates of new parties nor the concentration of predicted vote shares (Online Appendix Table A.28).

V.E. Effects of Disasters

Because TV debates do not contribute to vote choice formation despite being the most salient campaign events and showcasing the candidates themselves, one may wonder whether the increase in vote choice consistency documented in Section III is due instead to shocks occurring independently from the campaign, outside of politicians' control. Using the EM-DAT International Disasters Database, we identified 27 natural and technological disasters that occurred before 15 elections in three countries of the sample (Canada, Germany, and the United States), and that started more than three days before and three days after another disaster in the same country.[39] We estimate their impact with our event study design including three predisaster and postdisaster dummies (as in equation (8)).

The comparison of point estimates for the days preceding and following the disasters does not reveal any systematic effect on vote choice formation. Specifically, we first consider the effects of disasters on vote choice consistency (Online Appendix Figure A.15, Panel A and Online Appendix Table A.29, columns (1) and (2)). While the dummy for the third day after the disaster is positive and significant at the 5% level, the average impact of disasters on vote choice consistency in the three following days is nonsignificant and close to the average of the three predisaster dummies. When we extend the analysis to a five-day window, we

39. The full list of disasters is provided in Online Appendix Table B.7. Fifty-one percent are storms, 15% floods, 15% transport accidents, 11% miscellaneous accidents (e.g., food contamination outbreaks), and 7% wildfires.

find that the effects four and five days after a disaster are small and nonsignificant (Online Appendix Table A.30).

Disasters do not have any clear effect on the probability to state a vote intention and on consistency conditional on stating a vote intention, either (Online Appendix Figures A.15, Panel B and A.15c and Online Appendix Table A.29, columns (3)–(6)). In the latter case, the dummy for the first day after a disaster starts is positive and significant at the 10% level, but the average of the three postdisaster dummies is nonsignificant and comparable to the average of the three predisaster dummies. Next, Online Appendix Figure A.16, Panel A and Online Appendix Table A.31, columns (1) and (2), show no significant effect of disasters on the distance between predicted and final vote shares: the dummies for the days preceding and following disasters are all close to zero and nonsignificant. Finally, all pre- and postdisaster relative-days dummies for the impact of disasters on the daily change in vote shares are positive and most of them are significant (Online Appendix Figure A.16, Panel B and Online Appendix Table A.31, columns (3) and (4)) suggesting that disaster days are characterized by unusually low fluctuations in vote intentions relative to neighboring days (rather than disasters having an actual impact on this outcome).

Overall, these results suggest that disasters do not contribute to vote choice formation more than TV debates, but they are imprecise and should be interpreted with caution, given the small number of events (less than half the number of TV debates).

VI. Conclusion

We study vote choice formation during campaigns, using 253,000 observations from two-round surveys in 62 elections around the world since 1952. Our method does not rely on people's own recollection of the date when they made up their mind, but instead on measuring the consistency between individuals' responses to pre- and postelectoral surveys. Focusing on this outcome also enables us to study the effects of specific events while allowing for the possibility that different voters are influenced in divergent ways. We examine TV debates and disasters, but studies measuring the effects of other types of events would benefit from using this method instead of considering outcomes such as vote intention, which only capture net effects.

Overall, the fraction of people who state a pre-election vote intention identical to their eventual vote choice increases by 17 percentage points on average in the 60 days leading up to the vote. This large increase in individual vote choice consistency is concomitant with a 5 percentage point reduction in the distance between predicted and final vote shares. In other words, voters who make up their mind in this period affect the electoral results. We provide suggestive evidence that changes in vote choice come from changed beliefs about candidates' positions or quality more than beliefs about their chances of victory, and that priming also contributes to vote choice formation. We do find more modest effects of campaigns on voters with strong party attachments and in the U.S. two-party system, echoing the view that long-standing partisan attachments can reduce voters' receptiveness to election-specific information. However, even in these cases, campaign effects remain sizable.

While our results support the view that campaigns have substantial effects on vote choice, it would be incorrect to infer that voters are swayed by just any information they receive in the electoral season. We were surprised to find that people's policy preferences are not affected by the campaign. In addition, we do not find any clear evidence that shocks occurring independently from the campaign, such as natural and technological disasters, have important effects on vote choice formation in the national elections we study. Most important, our event study finds that TV debates—for all the interest they generate, the large viewing audience they draw, and the many media commentaries they provoke—neither increase individual vote choice consistency nor reduce the distance to final vote shares. If anything, TV debates move a small fraction of voters away from their final vote choice, in the short term.

The fact that our sample includes data from 10 countries makes the external validity of our results unusually broad. This said, we note that all these countries are well-established democracies. The fraction of voters making up their mind during campaigns may be even larger in countries with younger democratic regimes, less stable party systems, and lower baseline levels of political information. Naturally, the effects of TV debates may be different in such countries.

Overall, our results suggest that even if voters sometimes seem relatively uninformed, their vote choice actually aggregates extensive information, beyond just debates, and that other sources are more effective. A possible interpretation is that voters

discard candidates' debate statements because they rationally expect them to be more biased than information coming from the media, discussions with other voters, and other third parties, or that voters only pay attention to statements aligned with their beliefs. An alternative interpretation is that the particular medium through which debates are broadcast is what matters: it is difficult for candidates to change people's minds, and this does not happen via TV or radio communication. This interpretation is consistent with the fact that the evidence on the effects of political ads diffused through these channels on vote choice is mixed (Gerber et al. 2011; Spenkuch and Toniatti 2018), whereas more personalized contacts such as door-to-door visits or town hall meetings tend to have large persuasive effects (e.g., Fujiwara and Wantchekon 2013; Pons 2018).

One implication is that candidates should focus on organizing these more effective activities if they want to increase their chances of winning. In the elections we study, only a minority of voters report having been contacted or visited by a party. Our results also have implications for the regulation of campaigns. Since the first presidential TV debate in the United States in 1960, there has been a continuous effort to diffuse this innovation to countries that have not adopted it yet and to improve the format of debates where they have become a tradition. Our results suggest that some of this energy may be better spent in reforming campaign regulations to ensure that all candidates have equal access to voters and in monitoring the most personal and tailored forms of partisan communication to improve the quality of information available to voters.

HEC Montréal, Canada
Harvard University, United States

Supplementary Material

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online.

Data Availability

The data underlying this article are available in the Harvard Dataverse, https://doi.org/10.7910/DVN/XMDFQO (Le Pennec and Pons 2022).

REFERENCES

Achen, Christopher H., and Larry M. Bartels, *Democracy for Realists: Why Elections Do Not Produce Responsive Government* (Princeton, NJ: Princeton University Press, 2016).

Angelucci, Charles, and Andrea Prat, "Is Journalistic Truth Dead? Measuring How Informed Voters Are about Political News," 2021, available at SSRN, http://dx.doi.org/10.2139/ssrn.3593002.

Ashworth, Scott, Ethan Bueno de Mesquita, and Amanda Friedenberg, "Learning about Voter Rationality," *American Journal of Political Science*, 62 (2018), 37–54. https://doi.org/10.1111/ajps.12334.

Benoit, William L., Glenn J. Hansen, and Rebecca M. Verser, "A Meta-Analysis of the Effects of Viewing U.S. Presidential Debates," *Communication Monographs*, 70 (2003), 335–350. https://doi.org/10.1080/0363775032000179133.

Besley, Timothy, and Stephen Coate, "An Economic Model of Representative Democracy," *Quarterly Journal of Economics*, 112 (1997), 85–114. https://doi.org/10.1162/003355397555136.

Besley, Timothy, Jose G. Montalvo, and Marta Reynal-Querol, "Do Educated Leaders Matter?" *Economic Journal*, 121 (2011), F205–F227, https://doi.org/10.1111/j.1468-0297.2011.02448.x.

Bidwell, Kelly, Katherine Casey, and Rachel Glennerster, "Debates: Voting and Expenditure Responses to Political Communication," *Journal of Political Economy*, 128 (2020), https://doi.org/10.1086/706862.

Birdsell, David S., "Political Campaign Debates," in *The Oxford Handbook of Political Communication,* Kate Kenski and Kathleen Hall Jamieson, eds. (Oxford: Oxford University Press, 2017).

Boxell, Levi, Matthew Gentzkow, and Jesse Shapiro, "Cross-Country Trends in Affective Polarization," *Review of Economics and Statistics*, (2022), https://doi.org/10.1162/rest_a_01160.

Brierley, Sarah, Eric Kramon, and George Kwaku Ofosu, "The Moderating Effect of Debates on Political Attitudes," *American Journal of Political Science*, 64 (2020), 19–37, https://doi.org/10.1111/ajps.12458.

Cantoni, Enrico, and Vincent Pons, "Does Context Outweigh Individual Characteristics in Driving Voting Behavior? Evidence from Relocations within the United States," *American Economic Review*, 112 (2022), 1226–1272. https://doi.org/10.1257/aer.20201660.

Chaffee, Steven H., and Sun Yuel Choe, "Time of Decision and Media Use During the Ford-Carter Campaign," *Public Opinion Quarterly*, 44 (1980), 53–69. https://doi.org/10.1086/268566.

Chiang, Chun-Fang, and Brian Knight, "Media Bias and Influence: Evidence from Newspaper Endorsements," *Review of Economic Studies*, 78 (2011), 795–820. https://doi.org/10.1093/restud/rdq037.

Cho, In-Koo, and David M. Kreps, "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102 (1987), 179–221. https://doi.org/10.2307/1885060.

Cole, Shawn, Andrew Healy, and Eric Werker, "Do Voters Demand Responsive Governments? Evidence from Indian Disaster Relief," *Journal of Development Economics*, 97 (2012), 167–181. https://doi.org/10.1016/j.jdeveco.2011.05.005.

Crawford, Vincent P., and Joel Sobel, "Strategic Information Transmission," *Econometrica*, 50 (1982), 1431–1451. https://doi.org/10.2307/1913390.

de Chaisemartin, Clément, and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 110 (2020), 2964–2996. https://doi.org/10.1257/aer.20181169.

DellaVigna, Stefano, and Matthew Gentzkow, "Persuasion: Empirical Evidence," *Annual Review of Economics*, 2 (2010), 643–649. https://doi.org/10.1146/annurev.economics.102308.124309.

DellaVigna, Stefano, and Ethan Kaplan, "The Fox News Effect: Media Bias and Voting," *Quarterly Journal of Economics*, 122 (2007), 1187–1234. https://doi.org/10.1162/qjec.122.3.1187.

Delli Carpini, Michael X., and Scott Keeter, *What Americans Know about Politics and Why It Matters* (New Haven, CT: Yale University Press, 1997).
Enos, Ryan D., and Eitan D. Hersh, "Party Activists as Campaign Advertisers: The Ground Campaign as a Principal-Agent Problem," *American Political Science Review*, 109 (2015), 252–278. https://doi.org/10.1017/S0003055415000064.
Ferejohn, John, "Incumbent Performance and Electoral Control," *Public Choice*, 50 (1986), 5–25. https://doi.org/10.1007/BF00124924.
Ferraz, Claudio, and Frederico Finan, "Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes," *Quarterly Journal of Economics*, 123 (2008), 703–745. https://doi.org/10.1162/qjec.2008.123.2.703.
Fischer, Paul E., and Phillip C. Stocken, "Imperfect Information and Credible Communication," *Journal of Accounting Research*, 39 (2001), 119–134. https://doi.org/10.1111/1475-679X.00006.
Fournier, Patrick, Richard Nadeau, André Blais, Elisabeth Gidengil, and Neil Nevitte, "Time-of-Voting Decision and Susceptibility to Campaign Effects," *Electoral Studies*, 23 (2004), 661–681. https://doi.org/10.1016/j.electstud.2003.09.001.
Fridkin, Kim L., Patrick J. Kenney, Sarah Allen Gershon, Karen Shafer, and Gina Serignese Woodall, "Capturing the Power of a Campaign Event: The 2004 Presidential Debate in Tempe," *Journal of Politics*, 69 (2007), 770–785. https://doi.org/10.1111/j.1468-2508.2007.00574.x.
Fujiwara, Thomas, and Leonard Wantchekon, "Can Informed Public Deliberation Overcome Clientelism? Experimental Evidence from Benin," *American Economic Journal: Applied Economics*, 5 (2013), 241–255. http://dx.doi.org/10.1257/app.5.4.241.
Galasso, Vincenzo, and Tommaso Nannicini, "Political Selection under Alternative Electoral Rules," *Public Choice*, 171 (2017), 257–281. https://doi.org/10.1007/s11127-017-0436-x.
Gelman, Andrew, and Gary King, "Why Are American Presidential Election Campaign Polls So Variable When Votes Are So Predictable?," *British Journal of Political Science*, 23 (1993), 409–451. https://doi.org/10.1017/S0007123400006682.
Gennaioli, Nicola, and Andrei Shleifer, "What Comes to Mind," *Quarterly Journal of Economics*, 125 (2010), 1399–1433. https://doi.org/10.1162/qjec.2010.125.4.1399.
Gerber, Alan S., James G. Gimpel, Donald P. Green, and Daron R. Shaw, "How Large and Long-Lasting Are the Persuasive Effects of Televised Campaign Ads? Results from a Randomized Field Experiment," *American Political Science Review*, 105 (2011), 135–150. https://doi.org/10.1017/S000305541000047X.
Granzier, Riako, Vincent Pons, and Clémence Tricaud, "The Large Effects of a Small Win: How Past Rankings Shape the Behavior of Voters and Candidates," NBER Working Paper no. 26599, 2021.
Henderson, Michael, and D. Sunshine Hillygus, "Changing the Clock: The Role of Campaigns in the Timing of Vote Decision," *Public Opinion Quarterly*, 80 (2016), 761–770. https://doi.org/10.1093/POQ%2FNFW027.
Hillygus, D. Sunshine, "Campaign Effects and the Dynamics of Turnout Intention in Election 2000," *Journal of Politics*, 67 (2005), 50–68. https://doi.org/10.1111/j.1468-2508.2005.00307.x.
Hillygus, D. Sunshine, and Simon Jackman, "Voter Decision Making in Election 2000: Campaign Effects, Partisan Activation, and the Clinton Legacy," *American Journal of Political Science*, 47 (2003), 583–596. https://doi.org/10.1111/1540-5907.00041.
Holbrook, Thomas M., *Do Campaigns Matter?* (Thousand Oaks, CA: SAGE Publications, 1996). https://dx.doi.org/10.4135/9781452243825.
Jennings, Will, and Christopher Wlezien, "The Timeline of Elections: A Comparative Perspective," *American Journal of Political Science*, 60 (2016), 219–233. https://doi.org/10.1111/ajps.12189.

Johnston, Richard, and Henry E. Brady, "The Rolling Cross-Section Design," *Electoral Studies*, 21 (2002), 283–295. https://doi.org/10.1016/S0261-3794(01)00022-1.
Kalla, Joshua L., and David E. Broockman, "The Minimal Persuasive Effects of Campaign Contact in General Elections: Evidence from 49 Field Experiments," *American Political Science Review*, 112 (2018), 148–166. https://doi.org/10.1017/S0003055417000363.
Kamenica, Emir, and Matthew Gentzkow, "Bayesian Persuasion," *American Economic Review*, 101 (2011), 2590–2615. https://doi.org/10.1257/aer.101.6.2590.
Kendall, Chad, Tommaso Nannicini, and Francesco Trebbi, "How Do Voters Respond to Information? Evidence from a Randomized Campaign," *American Economic Review*, 105 (2015), 322–353. http://dx.doi.org/10.1257/aer.20131063.
Lazarsfeld, Paul Felix, Bernard Berelson, and Hazel Gaudet, *The People's Choice: How the Voter Makes Up His Mind in a Presidential Campaign* (New York: Duell, Sloan and Pearce, 1944). https://doi.org/10.7312/laza93930.
Le Pennec, Caroline, and Vincent Pons, "Replication Data for: 'How Do Campaigns Shape Vote Choice? Multicountry Evidence from 62 Elections and 56 TV Debates'," (2022), Harvard Dataverse, https://doi.org/10.7910/DVN/XMDFQO.
Lenz, Gabriel S., *Follow the Leader? How Voters Respond to Politicians' Policies and Performance* (Chicago: University of Chicago Press, 2012). https://doi.org/10.1093/poq/nft014.
Lindbeck, Assar, and Jörgen W. Weibull, "Balanced-Budget Redistribution as the Outcome of Political Competition," *Public Choice*, 52 (1987), 273–297.
Martin, Gregory J., and Ali Yurukoglu, "Bias in Cable News: Persuasion and Polarization," *American Economic Review*, 107 (2017), 2565–2599. https://doi.org/10.1257/aer.20160812.
McKinney, Mitchell S., and Diana B. Carlin, "Political Campaign Debates," in *Handbook of Political Communication Research*, Lynda, Lee Kaid, ed, (London: Routledge, 2004), 203–234.
Minozzi, William, Michael A. Neblo, Kevin M. Esterling, and David M. J. Lazer, "Field Experiment Evidence of Substantive, Attributional, and Behavioral Persuasion by Members of Congress in Online Town Halls," *Proceedings of the National Academy of Sciences*, 112 (2015), 3937–3942. https://doi.org/10.1073/pnas.1418188112.
Mullainathan, Sendhil, Ebonya Washington, and Julia R. Azari, "The Impact of Electoral Debate on Public Opinions: An Experimental Investigation of the 2005 New York City Mayoral Election," in *Political Representation* Ian Shapiro, Susan C Stokes, Elisabeth Jean Wood and Alexander S Kirshner, eds. (Cambridge: Cambridge University Press, 2009), 324–341. http://dx.doi.org/10.1017/CBO9780511813146.014.
Neundorf, Anja, and Kaat Smets, *"Political Socialization and the Making of Citizens,"* Oxford Handbooks Online (2017), 1–28, https://doi.org/10.1093/oxfordhb/9780199935307.013.98.
Panagopoulos, Costas, "All about that Base: Changing Campaign Strategies in U.S. Presidential Elections," *Party Politics*, 22 (2016), 179–190. https://doi.org/10.1177/1354068815605676.
Pettersson-Lidbom, Per, "Do Parties Matter for Economic Outcomes? A Regression-Discontinuity Approach," *Journal of the European Economic Association*, 6 (2008), 1037–1056. https://doi.org/10.1162/JEEA.2008.6.5.1037.
Pons, Vincent, "Will a Five-Minute Discussion Change Your Mind? A Countrywide Experiment on Voter Choice in France," *American Economic Review*, 108 (2018), 1322–1363. https://doi.org/10.1257/aer.20160524.
Pons, Vincent, and Clémence Tricaud, "Expressive Voting and Its Cost: Evidence from Runoffs with Two or Three Candidates," *Econometrica*, 86 (2018), 1621–1649. https://doi.org/10.3982/ECTA15373.
Shaw, Daron R., "A Study of Presidential Campaign Event Effects from 1952 to 1992," *Journal of Politics*, 61 (1999), 387–422. https://doi.org/10.2307/2647509.

Shaw, Daron R., and Brian E. Roberts, "Campaign Events, the Media and the Prospects of Victory: The 1992 and 1996 US Presidential Elections," *British Journal of Political Science*, 30 (2000), 259–289. https://doi.org/10.1017/S0007123400000120.
Spenkuch, Jörg L., and David Toniatti, "Political Advertising and Election Results," *Quarterly Journal of Economics*, 133 (2018), 1981–2036. https://doi.org/10.1093/qje/qjy010.
Wlezien, Christopher, and Robert S. Erikson, "The Timeline of Presidential Election Campaigns," *Journal of Politics*, 64 (2002), 969–993. https://doi.org/10.1111/1468-2508.00159.
Wright, Gerald C., "Errors in Measuring Vote Choice in the National Election Studies, 1952–88," *American Journal of Political Science*, 37 (1993), 291–316. https://doi.org/10.2307/2111533.