Contagious populists: The impact of election information shocks on populist party preferences in Germany

Gerling, Lena; Kellermann, Kim Leonie

European Journal of Political Economy 72 (2022) 102098

Contents lists available at ScienceDirect

European Journal of Political Economy

journal homepage: www.elsevier.com/locate/ejpe

Contagious populists: The impact of election information shocks on populist party preferences in Germany

Lena Gerling, Kim Leonie Kellermann *

University of Münster, Center for Interdisciplinary Economics, Scharnhorststraße 100, 48151 Münster, Germany

ARTICLE INFO

JEL classification:
C26
D74
P16

Keywords:
Populism
Contagion effects
Information shocks
Social desirability bias

ABSTRACT

This study analyzes the role of social contagion in populist party support. Although the emergence of the German right-wing populist AfD was accompanied by controversial debates about the social acceptability of its nationalist program, electoral support has followed a clear upward trend. We analyze the impact of information shocks with respect to aggregate-level support for the AfD on individual vote intentions. Unexpectedly high aggregate support for a populist party may indicate a higher social acceptance of its platform and reduce the social desirability bias in self-reported vote intentions. Consequently, the likelihood that an individual will reveal an AfD vote intention increases. We test this mechanism in an event-study approach, exploiting quasi-random variation in survey interviews conducted around the time of German state elections. We define election information shocks as deviations of actual AfD vote shares from pre-election polls and we link these shocks to an individual's likelihood of reporting an AfD vote intention in subsequent survey interviews. Our results suggest that exposure to higher-than-expected AfD support significantly increases the probability of reporting an AfD vote intention by up to 2.7 percentage points. Testing alternative mechanisms, we find that this increase is in fact driven by reduced reputational concerns associated with expressing populist support.

1. Introduction

Established party systems in both Europe and the U.S. have recently been rattled by the fast-growing success of right-wing populist platforms.[1] In Germany, the rise of the *Alternative for Germany* (*Alternative fuer Deutschland*, AfD) party gained momentum in a series of state elections before the party entered the federal parliament in 2017 with a vote share as high as 12.6 percent. Scientific research has suggested that this rapid spread of right-wing populist support throughout the electorate may have been driven by social contagion related to peer pressure and social compliance (Coate and Conlin, 2004; DellaVigna et al., 2012). For example, recent studies have observed that the victory of Donald Trump triggered contagion effects in reported anti-immigrant attitudes both within and outside of the U.S. (Giani and Méon, 2021; Bursztyn et al., 2020). Comparable effects of interpersonal contagion have also played a role in protests, such as the 2011 London riots (Barbera and Jackson, 2020; Zeitzoff, 2017) and historic protest events (Aidt et al., 2021; Caprettini and Voth, 2020).

We analyze the existence of similar contagious dynamics in publicly stated support for the German AfD. We apply a quasi-experimental event-study design which exploits variation in vote intentions reported in surveys conducted closely around the time

* Corresponding author.
E-mail addresses: lena.gerling@uni-muenster.de (L. Gerling), kim.kellermann@wiwi.uni-muenster.de (K.L. Kellermann).

[1] Several recent studies deal with the determinants of the rise of populism, both with respect to economic and non-economic factors. See, among others, Algan et al. (2017), Noury and Roland (2020), Schaub and Morisi (2020), Guriev (2018), Guriev and Papaioannou (2020), Rodrik (2018) and Giuliano and Wacziarg.

https://doi.org/10.1016/j.ejpoleco.2021.102098
Received 15 January 2021; Received in revised form 29 July 2021; Accepted 8 August 2021
Available online 14 August 2021

of German state elections held between 2013 and 2017. According to the effect of the well-known *social desirability bias*, individuals may falsify their true vote intention if they believe that their political attitude is not socially acceptable. Hence, right-wing party preferences are systematically underreported in personal interviews (DellaVigna et al., 2012; Hainmueller and Hangartner, 2013). However, if AfD sympathizers observe larger aggregate support for originally controversial attitudes, they may be encouraged to reveal their genuine vote intentions. We capture new, unanticipated information about levels of social support by means of *election information shocks*, measured by the deviation of state election outcomes from pre-election poll predictions.

Our results provide systematic evidence that larger-than-expected AfD vote shares in state elections raise subsequently reported AfD vote intentions among individuals living in *other* states by up to 2.7 percentage points. This is a sizable effect, corresponding to an increase of 36 percent in publicly expressed AfD support. In contrast, we find no significant effect for the negative election information shocks in our sample. Exploring the underlying mechanisms, we show that the dominant driving force behind the contagious effect of shocks in observable populist support is the alleviation of reputational concerns.

We provide innovative evidence on social contagion, a yet understudied driving force of populist preferences, and make several contributions to the related literature. The small number of earlier studies on populist contagion, which mostly focus on the 2016 U.S. federal election, examine whether Donald Trump's rise in popularity affected xenophobic attitudes and hate crimes in and outside the United States (Giani and Méon, 2021; Bursztyn et al., 2020; Müller and Schwarz, 2020). We extend this research in multiple ways. First, we exploit a unique set of election information shocks besides the Trump phenomenon. In doing so, to the best of our knowledge, we are the first to investigate populist contagion in a multiparty system. Transferring this issue to another political system provides an intriguing insight into the global prevalence of populist contagion. Second, a vital contribution of our study is that we find contagion effects for a party that did not win the election and did not participate in government. This result bears important implications for real-world politics since it implies that a single, powerful event, like the Trump win, is not required to challenge democratic systems. Instead, we observe a more subtle form of populist contagion, which induces gradual changes in perceived acceptability of populism over time. Thus, while each increase in populist vote intentions may not be considered far-reaching, in sum, they may still change the political landscape of a country. Third, we focus on vote intentions rather than general political attitudes. While we cannot observe actual voting decisions, we provide a suggestive intuition of how changing social norms may induce a self-reinforcing trend of populist electoral success.

From a methodological point of view, we add to the existing literature by providing a concise theoretical framework which illustrates how observed social preferences feed back on individual political behavior. In order to empirically test this effect, we operationalize election information shocks as the vehicle for interpersonal information exchange. Finally, a key contribution lies in the explicit distinction of the relevant mechanisms of populist contagion and their analysis using observational data.

The German federal system, in combination with the emergence of a far right-wing platform, provides a unique setup well suited to analyze the effect of information shocks on individual attitudes. When founded in 2013, the AfD mainly promoted Euroskepticism and fiscal conservatism. Yet in 2015, when the European *refugee crisis* peaked, the party made a strong shift to the right, focusing almost exclusively on immigration with outright xenophobic elements – a near taboo in Germany. This programmatic transformation was rewarded in subsequent elections (see Fig. B.1). While voters were not fully confident in placing the AfD at the extreme right of the political spectrum in 2013, the party's narrow focus on immigration after 2015 led to a strong public perception as a far-right party.[2] The programmatic shift represents a potential stimulus for professed AfD supporters to be socially stigmatized. The AfD therefore constitutes an ideal example for testing whether and how shocks in aggregate-level support shape self-reported preferences (see Fig. 1).

For these self-reported vote intentions, we rely on individual-level data from the German *Politbarometer* survey. Importantly, we are interested in *reported* vote intentions. When revealing a vote intention to another person, individuals consider the likelihood that their response may be socially unacceptable or condemned. Self-reported vote intentions thus provide an ideal opportunity to study the effect of aggregate-level support on such reputational concerns. Our dataset provides a comprehensive sample of repeated cross-sections, covering vote intentions as well as politically relevant individual characteristics. To test the link between social and individual behavior, we have to identify sudden events that reveal reliable information about aggregate shifts in populist support. Strikingly, opinion polls prior to the German state elections had pronounced difficulty in correctly predicting AfD vote shares. As reported in appendix Fig. B.2, pre-election polls based on interviews systematically overestimated or underestimated the actual AfD vote share in a number of state elections held between 2013 and 2017. Therefore, these deviations capture information shocks about the *actual* level of social acceptance of the AfD (Giani and Méon, 2021). We thus define an *election information shock* as the deviation between the AfD vote share in any given state election and the most recent federal opinion poll prior to it.

Election outcomes and individual preferences within the same constituency, however, may not be independent (Manski, 1993, 2000). To address this problem, we relate state-level election outcomes to *federal* vote intentions, and we disregard all interviewees residing in the state where the election in question was held. We control for additional, potentially confounding effects at the national or state level by exploiting variation in reported preferences in the two polls closest to the election date. We compare a treatment group of respondents interviewed in the first survey after the election to a control group interviewed right before the election. Within this narrow election window, structural factors can be assumed to be sufficiently constant.

Applying entropy balancing and including state-of-residence as well as election-window fixed effects, we show that our treatment effect is robust across several econometric specifications. In addition, we provide a number of heterogeneity tests. Here, we show that the effect of positive election information shocks on self-reported vote intentions is boosted by intensive media consumption

[2] Fig. B.2 in the appendix shows AfD ratings from the German Longitudinal Election Study.

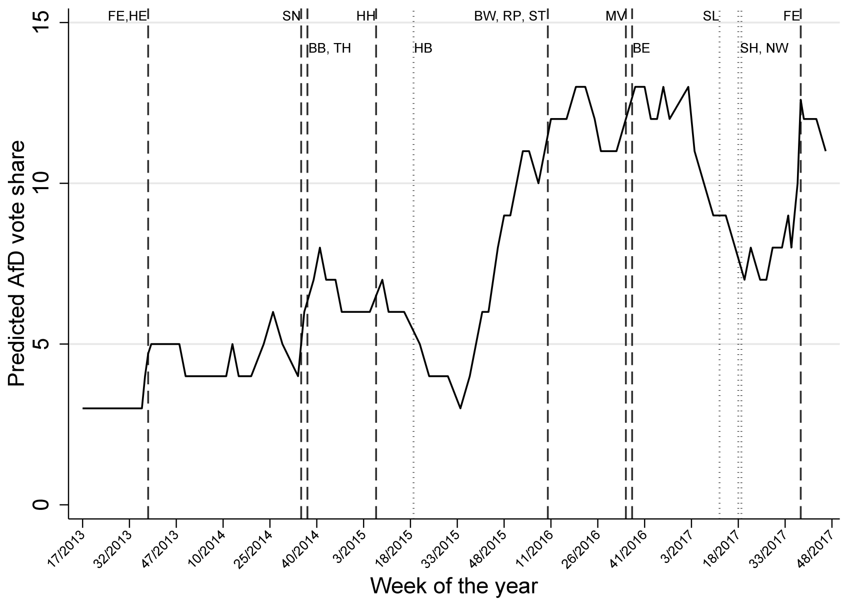

Fig. 1. Predicted AfD vote shares and state elections.
Note: Federal-level AfD polls and dates of state elections are shown in chronological order from 2013 to 2017. We exclude two additional state elections which took place in this period, in Bavaria in 2013 and in Lower Saxony in 2017. As they occurred either shortly before or after the federal elections, respectively, they do not fall into this particular election cycle. HE = Hesse, SN = Saxony, BB = Brandenburg, TH = Thuringia, HH = Hamburg, HB = Bremen, BW = Baden-Wuerttemberg, RP = Rhineland-Palatinate, ST = Saxony-Anhalt, MV = Mecklenburg-Hither Pomerania, BE = Berlin, SL = Saarland, SH = Schleswig-Holstein, NW = North Rhine-Westphalia, FE = Federal elections.

regarding the AfD, geographic proximity to the election state and low initial support in the individual's state of residence. Furthermore, individuals in former West Germany respond more strongly to a positive election information shock than respondents in former East Germany.

Our estimation strategy benefits from the fact that state elections in Germany are staggered, providing us with multiple treatment events which are random with respect to the date of closely surrounding polls. This strengthens our supposition that the observed effect of election information shocks on self-reported vote intentions is in fact causal. In addition, we find no significant effects when conducting placebo tests for comparable time frames without elections, suggesting that our treatment effect is not driven by a general trend in AfD vote shares. Taking effect size into account, we find that larger positive election information shocks lead to a relatively stronger increase in self-reported AfD vote intentions.

Yet, social contagion may simply be the result of bandwagoning, in the sense that individuals may just be supporting a well-performing party to *join the winning side*. In order to rule out this mechanism, as well as strategic motives of expressing a controversial vote intention to promote political change, we conduct multiple tests. Neither do we find a robust pattern of contagion for non-populist parties nor for a group of respondents without reputational concerns. Our findings are, thus, most likely driven by reduced social stigmatization associated with openly expressing populist support.

The remainder of this paper is organized as follows. Section 2 provides an overview of the related literature and lays out our theoretical framework. Section 3 elaborates on our identification strategy and sample. Estimation results are presented in Section 4. Section 5 concludes.

2. Case selection and theoretical framework

2.1. The AfD as a right-wing populist party

The AfD was originally founded in 2013 as a special-issue party opposing the EU's financial-support policies in the aftermath of the *Euro crisis*. Initially, the AfD did not fully embrace the typical right-wing, nationalist tendencies and populist rhetoric (Arzheimer, 2015). However, the party underwent a fundamental transformation in the course of the so-called *refugee crisis* in 2014/15 when the nationalist wing took over leadership. Since then, the AfD has been decidedly emphasizing its position against the immigration-friendly policies of the Merkel administration (Berbuir et al., 2015; Salzborn, 2016).

Not only has the AfD drawn closer to comparable European parties, such as the French *Rassemblement National* and the Austrian *FPÖ*, some AfD members have also been linked to right-wing extremist organizations (Salzborn, 2016; Berbuir et al., 2015). Various party officials have sparked controversies by attending far-right rallies or using clearly xenophobic language and expressions from the Nazi era in public statements. Therefore, unsurprisingly, the legitimacy of support for the AfD has inspired a controversial debate. The party is often stigmatized as anti-democratic and unelectable (Berbuir et al., 2015). Given its national history, Germany's political landscape is very sensitive to movements featuring nationalist or xenophobic attitudes (Cantoni et al., 2019; Rydgren, 2005; Mudde, 2004). Therefore, openly sympathizing with such movements is often associated with social proscription.

Despite crossing these boundaries, there has been a clear upward trend in AfD election results. In the 2013 federal election, the AfD just failed to pass the threshold of 5 percent of valid votes required in order to enter the parliament. However, the following year, the AfD had great successes in the state elections in Brandenburg, Saxonia and Thuringia and entered all three state parliaments with vote shares of roughly 10 percent or higher.[3] By 2017, the party had won seats in 14 out of 16 state parliaments. It entered the federal parliament in 2017 as the largest opposition party with a vote share of 12.6 percent.

The pronounced discrepancy between how AfD supporters perceive the party compared to other citizens is an interesting issue.[4] While the topic has not yet been investigated by much in-depth research, some German research institutions provide descriptive insights. Interestingly, AfD supporters perceive the party to be less right wing than non-supporters do (Bergmann and Diermeier, 2017; Niedermayer and Hofrichter, 2016). As suggested by Bergmann and Diermeier (2017), one reason why AfD voters deemphasize their party's stigmatized position in in-person interviews may be due to a social desirability bias. In line with this, Friess and Neu (2018) observe more moderate attitudes among AfD supporters in face-to-face conversation than in anonymous online chats, particularly with respect to nationalism and anti-immigrant attitudes. While AfD supporters may achieve certain ends by emphasizing their controversial, anti-mainstream position, the notions outlined above rather suggest that, actually, AfD sympathizers fear potential reputational losses. Interestingly, Schmitt-Beck (2016) finds that poll consumption right before an election increases the probability that an individual will vote AfD.

2.2. Social contagion in political behavior

Yet, why would individuals copy their fellow voters' behavior? From the related literature, we have identified two main explanations for social contagion in political behavior, *bandwagon motives* and *reputational concerns*.

Regarding the former, social contagion operates through the bandwagon channel if individuals decide to support a well-performing candidate or party in order not to "waste" their vote on an unpopular electoral option (Dahlgaard et al., 2016; Klor and Winter, 2018; van der Meer et al., 2016; Barnfield, 2020). In that sense, individuals derive utility from *being on the winning side*. In run-off elections, this usually means that voters support the leading candidate, and thus the likely winner (Morton et al., 2015; van der Meer et al., 2016; Barnfield, 2020); but empirical research from multiparty systems shows that bandwagoning also exists with respect to a party that is simply gaining support in opinion polls, as compared to competitors that register decreasing support (Dahlgaard et al., 2016; van der Meer et al., 2016; Barnfield, 2020). This is an important notion with regard to our setting. While the AfD received increasing support in the German multiparty system, it was never in contention for government participation. Still, voters may have considered the party's upward trend worthwhile to join in.

Related to this idea, copying fellow voters' political behavior may also yield an informational cascade. In order to reduce the prohibitively high costs of collecting information on the quality of political contestants (Downs, 1957), individuals can simply adopt the behavior of other individuals on the assumption that they have collected and correctly evaluated the necessary information (Bikhchandani et al., 1992, 1998; Cao et al., 2011). Some theoretical work, however, has examined why some voters persist in supporting trailing parties. Piketty (2000) and Castanheira (2003) highlight the *communicative* function of voting decisions. Voting for parties with no chance of winning may be a useful way to inform fellow voters about one's true preferences or to exert pressure on winning parties to adopt some of the trailing parties' policies.

The second explanation for contagion in political behavior is of higher relevance to our context. Revealing certain political preferences may be associated with *reputational concerns* if voters fear social rejection or stigmatization upon voicing an unpopular opinion to an audience. This is particularly the case if the supported party or its ideas are generally considered unacceptable or even dangerous. Voters then want to conceal their party preference in order to avoid being publicly acknowledged as a supporter. Following Granovetter (1978) and Kuran (1987, 1989), individuals only adopt a potentially controversial behavior if the number of fellow citizens behaving this way exceeds an individual threshold level.

Transferring this idea to our example of right-wing populism is straightforward. If social advocacy for controversial populist positions is low, populist supporters would likely be subject to social rejection. Expressing a populist vote intention is associated with a reputational loss, so that voters harboring such a preference conceal it. However, growing support signals a change in public perception of the populists' electability. A larger number of supporters serve as a protective cloak under which individuals can voice their advocacy with less risk of reputational loss. Consequently, an individual's readiness to openly express a right-wing populist vote intention increases.

Empirical evidence of social contagion in political behavior has been found in various contexts, both at the individual (Huckfeld and Sprague, 1991; Straits, 1990; Glaser, 1959; Nickerson, 2008) and at the aggregate level (Morton et al., 2015; Khalil et al., 2019; Klor and Winter, 2018; Dahlgaard et al., 2016). Yet, these studies mainly observe bandwagon effects. Our approach, however, is in line with a number of recent studies that focus on reputation effects and social image concerns. Bursztyn et al. (2020) exploit the 2016 Trump win and experimentally show that, if people care about approval from a broader audience, right-wing populist electoral success increases the share of individuals who choose a xenophobic action. The authors design a theoretical model in

[3] Official election results at the federal and the state level have been obtained from the German Federal Statistical Office.

[4] See the online supplementary material for respective survey results also based on *Politbarometer* data.

which choosing a xenophobic action provides an individual that harbors xenophobic attitudes with utility from compliance with one's preferences and disutility from reputational concerns. Some xenophobes may falsify their true preferences if they observe a generally low acceptance of xenophobic actions. However, public signals such as populist right-wing electoral success suggest that the share of xenophobes in the audience is larger than expected and increases the share of individuals who express a respective attitude.

In line with this, Giani and Méon (2021) show that the unexpected Trump victory in the 2016 U.S. federal election increased the expression of xenophobic attitudes in Europe. Müller and Schwarz (2020) observe contagious effects of xenophobic Trump tweets that induce more real-life hate crime against Muslims. Huang and Low (2017) find more general effects of the Trump win in a lab experiment. After the election, participants exhibited a more aggressive, less cooperative and even more sexist behavior, which the authors trace back to the election having altered behavioral norms.

We build upon the related studies by Bursztyn et al. (2020) and Giani and Méon (2021), however, we explicitly distinguish the bandwagon channel and the reputation channel in our theoretical model. Furthermore, we derive testable predictions regarding social contagion in right-wing populist support through the reputation channel. Then, we test these predictions in a unique quasi-experiment that has not been exploited before. The main difference between support for a populist party compared to support for mainstream parties is that higher support for a populist party impacts reported preferences through both bandwagon and reputation channels. First, higher support signals better electoral performance, encouraging voters to join the winning side, just like for any other party. Second, higher support signals greater social acceptability of the populist platform, which stimulates supporters to reveal previously concealed preferences. However, for parties with an uncontroversial program, the reputation channel does not exist.

The above-mentioned studies suggest that behavioral convergence exists even without face-to-face communication. Thus, media consumption is a conceivable mediator for vote contagion (Aidt et al., 2021; Baudains et al., 2013). Although Dahlgaard et al. (2016) state that "it is almost impossible to avoid polls", Faas et al. (2008) emphasize that the information generated by polls only influences political behavior if individuals actually receive the information. The fact that media reports – specifically reports with a political bias – influence individual political views, has been observed in a number of studies (Durante et al., 2019; Durante and Knight, 2012; DellaVigna and Kaplan, 2007; Dewenter et al., 2019; Benesch et al., 2019; Boomgarden and Vliegenthart, 2009). A positive presentation of populist leaders also promotes sympathy among the electorate (Bos et al., 2011; Lubbers et al., 2002; van der Brug et al., 2000; Durante et al., 2019).

In the following section, we present a simple formalization that gives further structure to these mechanisms through which we expect election information shocks to affect individual vote intentions.

2.3. Theoretical mechanisms

Let $i = 1, 2, \ldots, N$ be an eligible voter and $j = 1, \ldots, J \geq 2$ political parties competing for voter support. To simplify the argument, we set $J = 2$ without loss of generality. Each individual receives utility from publicly stating a vote intention, e. g., in survey interviews. In this way, the only motivation for voters to report a vote intention is to be publicly known as a professed supporter.[5]

Individual utility from publicly revealing a vote intention comprises of three components. First, the individual party bias θ_i captures actual policy preferences and values. As $J = 2$, θ_i equals the utility surplus i receives when supporting j instead of $-j$ based on i's sincere political convictions (Kuran, 1989; Bursztyn et al., 2020).

Second, i's utility increases if j is particularly successful on the political stage and generally receives high support. This relates to the bandwagon effect as described in Section 2.2. v_j denotes the current share of society members who *actually* support party j. Yet, importantly, i is unable to observe v_j. As we are considering a situation where reputational concerns may play a role, people may conceal their true vote intention in public or express support for another party for strategic reasons. Therefore, i bases their publicly expressed vote intention for j on perceived support $\mu_{i,j} = q_i v_j \in [0, 1]$. By $q_i \geq 0$, we denote the individual perception bias of the true value v_j (Bursztyn et al., 2020). Depending on, for instance, i's amount and type of media consumption, their peer groups or level of interest in politics, $\mu_{i,j}$ may overestimate or underestimate the true value of v_j (DellaVigna and Kaplan, 2007).

We denote i's utility $p_{i,j}$ from j's popularity as

$$p_{i,j}(\mu_{i,j}) = a_i \mu_{i,j} \quad \text{with } a_i \geq 0. \tag{1}$$

As $\mu_{i,j}$ serves as a signal of j's probability of winning an upcoming election and its perceived aptitude to govern, utility from expressing a vote intention for j increases in $\mu_{i,j}$ (Bikhchandani et al., 1992; Bursztyn et al., 2020).

Third, and particularly relevant in this context, there may be a utility loss $l_{i,j}$ from openly supporting a party that promotes a socially or morally controversial program. Such a program may advocate for ideas that violate basic social values such as equality, freedom of religion or the respectful treatment of minorities. Let $\epsilon_j \in [0, 1]$ denote the share of controversial, socially unacceptable policies in j's program. If $\epsilon_j = 0$, the platform is free from controversial ideas. Assume that party $-j$ is an established mainstream party with a program which may not be in line with each voter's political convictions but that is generally considered acceptable, $\epsilon_{-j} = 0$. Yet, j is a challenger party promoting a radical and populist platform with $\epsilon_j > 0$. Let the reputational loss $l_{i,j}$ be a linear function of ϵ_j. The higher ϵ_j, the less acceptable is j's program and the higher is $l_{i,j}$.

[5] This reported vote intention does not necessarily align with actual votes. When actually voting, citizens may take into account the effective consequences of their choices (Downs, 1957; Riker and Ordeshook, 1968), an element which we neglect.

However, if i believes that a share $\mu_{i,j}$ of fellow voters openly supports party j, i may therefore conclude that they do not have to fear social rejection from these voters who exhibit the potentially inappropriate behavior of supporting j themselves. This means that the reputation loss only occurs with a probability of $1-\mu_{i,j}$. The larger $\mu_{i,j}$ is, the lower is the expected reputation loss (Kuran, 1989; Bursztyn et al., 2020). Therefore, we can denote $l_{i,j}$ as

$$l_{i,j}(\mu_{i,j}) = b_i\epsilon_j(1-\mu_{i,j}) \quad \text{with } b_i \geq 0. \tag{2}$$

The additive utility of individual i when expressing a vote intention for party j then reads

$$\begin{aligned} U_{i,j} &= p_{i,j}(\mu_{i,j}) - l_{i,j}(\mu_{i,j}) + \theta_i \\ &= a_i\mu_{i,j} - b_i\epsilon_j(1-\mu_{i,j}) + \theta_i, \qquad a_i, b_i \geq 0. \end{aligned} \tag{3}$$

The parameters a_i and b_i capture the individual utility attributed to the bandwagon and reputation channels, respectively. Individual i now compares total utilities from publicly supporting either j or $-j$ and chooses the option which promises the higher utility. That is, i reports a vote intention for j if

$$\begin{aligned} & U_{i,j} > U_{i,-j} \\ \Leftrightarrow\ & a_i\mu_{i,j} - b_i\epsilon_j(1-\mu_{i,j}) + \theta_i > a_i\mu_{i,-j}. \end{aligned} \tag{4}$$

Proposition 1. *Let* $\mu^*_{i,j} = \frac{a_i\mu_{i,-j}+b_i\epsilon_j-\theta_i}{a_i+b_i\epsilon_j}$. *Then*
(a) i publicly expresses support for j if $\mu_{i,j} > \mu^*_{i,j}$, i. e., if the perceived level of support for j in society is sufficiently high.
(b) i publicly expresses support for $-j$ if $\mu_{i,j} \leq \mu^*_{i,j}$, i. e., if the perceived level of support for j in society is low.

As j is the controversial option, i will only express a respective vote intention if $\mu_{i,j} > \mu^*_{i,j}$, i. e., if i perceives themselves as living in a society with a generally high support for the populist party (Bursztyn et al., 2020; Kuran, 1989).[6]

We can now illustrate the effect of a shock to $\mu_{i,j}$. Suppose that we start in a situation where $U_{i,j} < U_{i,-j}$, i. e. $\mu_{i,j}$ is too low for i to publicly side with the populist party. Now suppose that $\mu_{i,j}$ is updated due to, e. g., a local-level election. Let σ be a publicly observable, positive shock to $\mu_{i,j}$ so that we have a new perceived aggregate support share $\mu'_{i,j} = \mu_{i,j} + \sigma$.

Proposition 2. *Suppose that $\mu_{i,j} < \mu^*_{i,j}$ and thus $U_{i,j} < U_{i,-j}$. Now, i receives a signal $\sigma > 0$ to update her perception of aggregate support for j so that $\mu'_{i,j} = \mu_{i,j} + \sigma$. Let* $\sigma^*_i = \frac{a_i\mu_{i,-j}+b_i\epsilon_j-\theta_i}{a_i+b_i\epsilon_j} - \mu_{i,j}$.
(a) i decides to express a vote intention for j if $\sigma > \sigma^*_i$.
(b) i decides to stay with $-j$ if $\sigma \leq \sigma^*_i$.

We here observe an asymmetric effect on the reported vote intention. While small shocks do not alter i's publicly expressed vote intention, i is willing to switch to the populist party if σ is sufficiently large. The shock here changes revealed support through both the bandwagon and reputation channels. First, a positive shock increases j's governing probability and thereby i's utility from party popularity. Second, at the same time, the shock reduces the reputational loss, which also increases overall utility. The same asymmetric effect can be observed for a negative shock $\sigma < 0$. If i is a populist supporter, and perceived aggregate support for the populist party decreases, i will rather choose to express a vote intention for the mainstream party $-j$ if $\sigma < \sigma^*_i$ but will stick with j if $\sigma \geq \sigma^*_i$.

The main difference between our approach and previous, comparable studies (Kuran, 1989; Bursztyn et al., 2020; Castanheira, 2003) is the existence of the reputation channel, which is only relevant in the case of potentially controversial platforms. Therefore, we can compare the results in propositions 1 and 2 to a situation where there is no reputational loss. Such a case does not limit to j being a non-populist party. The reputational loss would also disappear for individuals with $b_i = 0$, meaning that they simply do not care about what their fellow voters think of their political statements. Likewise, a perception $\mu_{i,j} = 1$ would eliminate the effect of the reputation channel.

Proposition 3. *Suppose that there is no reputation loss for i, i. e., $U_{i,j} = a_i\mu_{i,j} + \theta_i$ and $U_{i,j} < U_{i,-j}$.*
(a) i switches to a populist vote intention if a positive shock $\sigma > 0$ exceeds $\sigma^{**}_i = \frac{a_i\mu_{i,-j}-\theta_i}{a_i} - \mu_{i,j}$.
(b) $\sigma^{**}_i > \sigma^*_i$ if $\mu_{i,-j} > a_i - \theta_i$.

As reputational concerns play no role here, a larger shock σ^{**}_i may be required to make i publicly express a vote intention for j than is the case when reputational concerns are a factor. While this may seem counterintuitive, the explanation is quite straightforward. If i has reputational concerns, a sufficiently large shock implies a twofold utility gain by increasing utility from party popularity and, at the same time, reducing the reputational loss. If the reputation channel is eliminated, however, overall gains can only be realized through an increase in utility from party performance, which then has to be higher: a larger positive shock is required. In particular, large shocks are necessary when the perceived aggregate support for the mainstream party $\mu_{i,-j}$ is high. For a shock $\sigma \in (\sigma^*_i, \sigma^{**}_i]$, we can expect to observe a positive effect on an individual's report of their vote intention for j when reputational concerns are in play, but we expect no effect if they are not.

In the following, we link these theoretical considerations to the empirical analysis of individual vote intentions.

[6] Note that $\mu^*_{i,j}$ is an individual threshold which we cannot determine numerically.

3. Empirical approach and data

3.1. Database

We are studying the period between the two most recent federal elections in Germany, which occurred on September 22, 2013 and September 24, 2017. At the state level, 14 out of 16 elections were held during this regular election cycle.[7] Data on individual vote intentions has been taken from the *Politbarometer* surveys (Forschungsgruppe Wahlen, 2019). To obtain the data, the *Forschungsgruppe Wahlen* (*Election Research Group*) conducts telephone interviews at two- to three-weeks intervals with approximately 1,250 respondents per survey round. The survey team applies a rigorous sample selection methodology based on randomly generated household phone numbers and the members' birthdays. The survey dates are usually fixed at the beginning of the year, but their frequency may be increased around important events. In the lead-up to federal elections, interviews are conducted on a weekly basis. Applying sample weights ensures that each survey contains a sample of individuals which is representative of the eligible German voting population.

The questions cover different topics related to current political issues, where a core set of questions is asked in every round. Foremost, respondents are asked about their readiness to participate in a hypothetical federal election the following Sunday and what their respective vote intention would be (the so-called *Sunday Question*). The exact wording of the question is: "If there were federal elections next Sunday, which party would you vote for?" We use the answer to this question to construct our dependent variable capturing the self-reported individual vote intention. In addition, the surveys document a wide range of demographic and socio-economic characteristics.

The *Politbarometer* has two advantages crucial to our research design. First, the surveys are not used to collect information with regard to state-level elections. For this purpose, separate surveys around the election in question are performed in the respective state. Therefore, we can assume that the interview date of a given individual in the *Politbarometer* is random with respect to state-level election dates. Second, interviews are conducted via telephone. Thus, there is a personal interaction between the interviewer and the interviewee. Although the survey results are, of course, published anonymously, the interviewer gains access to personal information about the participant, such as, gender, age or job. The interviewer even knows the interviewee's phone number and place of residence, which completely eliminates anonymity. Therefore, the interaction can be considered close enough to trigger any potential social desirability bias.

Even though the sample drawn in each survey round is representative, aggregates of vote intentions reported in the interviews do not necessarily match actual election outcomes. Deviations relate to short-term events, indecisive voters or false statements, whether made wittingly or unwittingly; such false statement again links back to reputational concerns in personal interactions. Therefore, the polls that are eventually published are calculations made based on the raw interview data, using a predictive model which includes additional assumptions about voting persistence and macro-economic fundamentals and trends. Importantly, in our analysis, we use the raw interview data for individual vote intentions. Throughout the paper, we use the term *survey* to refer to this raw interview data. In contrast, we use the term *federal-level poll* when referring to the estimate based on the last national survey conducted before the state election in question. In addition, we use *state-level polls* to refer to the most recent forecast based on a state-specific survey, which is usually published two to three days before the election in question.

3.2. Empirical specification and sample

To test the effect of state-level elections on self-reported vote intentions, we use an event-study design with a quasi-randomized treatment and a repeated cross-section of surveyed individuals. Our empirical approach closely relates to strategies applied in Depetris-Chauvin et al. (2020), Giani and Méon (2021), Mikulaschek et al. (2020). Treatment is defined as individual exposure to an election information shock, which is a deviation of the state-level election result from the previously reported polls. While we use the term *treatment* to refer to the *average treatment effect* (ATE) of an election information shock throughout the paper, the reader should keep in mind that we actually estimate an *intention to treat* (ITT) effect. Treated individuals in our setting are, in principle, subject to the election information shock, but some individuals may not know or care about it (Muñoz et al., 2019).

We aim to estimate the short-term effect of information shocks on self-reported federal vote intentions. To that end, we exploit the fact that our data contains information about the calendar *week of survey* (WoS) for each individual (also referred to as the survey round). We match this information with the calendar *week of election* (WoE). Following Mitra et al., we define for each election e an *election window* (δ_e) which identifies the period in which the surveys for the WoE are conducted. δ_e is given by

$$\delta_e = [\underline{t}, \bar{t}]\, with \begin{cases} \bar{t} = min\{t\}_{t>t_e}, & WoE_e = t_e \\ \underline{t} = max\{t\}_{t<t_e}, & WoE_e = t_e. \end{cases}$$

The upper boundary $\bar{t}$ refers to the first survey after the election, constituting the treatment group. $\underline{t}$ refers to the most recent survey before the election, in which the respective control group is interviewed.

Our estimation model then reads as follows:

$$Y_{i,r,t,e} = \beta\sigma_{t,e} + \Gamma' X_i + \Lambda' Z_{r,e} + \gamma_e + \theta_r + \varepsilon_{i,r,t,e}, \tag{5}$$

[7] Two states, Bavaria and Lower Saxony, held elections before or after the federal election cycle. In Bavaria, elections were on September 15, 2013; in Lower Saxony, elections were on January 20, 2013 and October 15, 2017.

Table 1
Election windows.

State	Election date	AfD vote share	Pre-election poll	Type of shock	WoE	WoS control group	WoS treatment group
SN	Aug 31, 2014	9.7	4.0	+	35	34	36
HH	Feb 15, 2015	6.1	6.0	+	7	5	9
HB	May 10, 2015	5.5	6.0	–	19	16	21
BW		15.1	10.0	+			
RP	Mar 13, 2016	12.6	10.0	+	10	7	11
ST		24.3	10.0	+			
MV	Sep 4, 2016	20.8	11.0	+	35	32	38
BE	Sep 18, 2016	14.2	11.0	+	37		
SL	Mar 26, 2017	6.2	9.0	–	12	10	14
SH	May 7, 2017	5.9	8.0	–	18	17	20
NW	May 14, 2017	7.4	8.0	–	19		

Note: Pre-election poll refers to the last federal-level poll before the indicated election date. SN = Saxony, HH = Hamburg, HB = Bremen, BW = Baden-Wuerttemberg, RP = Rhineland-Palatinate, ST = Saxony-Anhalt, MV = Mecklenburg-Hither Pomerania, BE = Berlin, SL = Saarland, SH = Schleswig-Holstein, NW = North Rhine-Westphalia.

where i denotes the individual, t denotes the week of survey, r denotes the state of residence and e refers to the election. $\varepsilon_{i,r,t,e}$ is an individual-level error term. Our dependent variable $Y_{i,r,t,e}$ is a binary indicator that equals 1 if an individual reports an AfD vote intention in the Sunday Question and 0 if the individual reports that they intend to vote for a party other than the AfD, that they do not intend to vote or that they do not know what party to vote for.[8]

$\sigma_{t,e}$ captures the treatment, that is, the exposure to an election information shock. If respondent i is interviewed in the first survey after a state election e and hence $t = \bar{t}$, then $\sigma_{t,e} = 1$. If respondent i is interviewed in the last survey before the election day and hence $t = \underline{t}$, then $\sigma_{t,e} = 0$. Our key parameter of interest is thus β which captures the average difference in an individual's likelihood to report an AfD vote intention right after a state election compared to shortly before the election.

When defining the treatment and control groups, a few points merit careful attention. First, we test the effect of a state election in r only for respondents in all *other* states $-r$. We choose this sample restriction for two reasons. The first is that we take into account what Manski (1993) calls the "reflection problem". It is a priori unclear whether the observed AfD vote share has an impact on individual preferences within the same area or whether the average vote share is simply the aggregation of individual preferences in this area. Hence, the exogeneity of our treatment may be questionable when including respondents of the election state. In contrast, we can assume that election outcomes in one state are not driven by vote intentions reported in surveys closely before or after an election in *other* states, after controlling for election-window fixed effects. The second is that respondents from the election state(s) may report a vote intention, e. g., in order to affect coalition negotiations or a policy agenda right after a state election. Such strategically motivated responses would conflate our hypothesized treatment effect that stems from reputational concerns. As shown in Table C.3, vote intentions reported by respondents from the election state(s) are in fact partly driven by strategic motivations. Therefore, we leave respondents from the election state(s) out of the baseline analysis. However, as a robustness check, we also provide evidence for the full sample and we find no qualitative difference from the baseline results.

Second, each individual is assigned to only one election window, either to the control group or the treatment group. For our identification strategy to be valid, we need to ensure that treated individuals are affected by the election information shock in question, while individuals in the control group are not, at least in the short term. In some instances, two or more elections take place on the same day or follow each other too quickly, so that election windows overlap. In Schleswig-Holstein and North Rhine-Westphalia, for instance, elections took place on two subsequent Sundays (May 7 and May 14, 2017), corresponding to calendar weeks 18 and 19. Survey data is available for calendar weeks 17 and 20, narrowly enclosing these elections. Therefore, we treat these elections as one election window. In this case, treated individuals are exposed to both elections, while non-treated individuals are interviewed before the first election. Furthermore, some election windows overlap such that the treatment group for the first election would at the same time be the control group for the second election. This is the case for Saxony (election held on August 31, 2014) and for Brandenburg and Thuringia (elections held on September 14, 2014). We therefore drop the latter elections from our sample, because the control group may be affected by the earlier election in the adjacent state of Saxony.[9] In a final step, we identify other potentially confounding events (Muñoz et al., 2019). To that end, we control for elections held at other administrative levels (e. g., federal elections, European elections). This process leaves us with seven election windows that cover a total of eleven single elections and a sample of 20,861 individual-level observations. Individuals interviewed in weeks which do not fall into any election window are dropped. An overview of the election windows is provided in Table 1. Summary statistics are provided in appendix table Table A.1.

In order to determine the direction and magnitude of an election information shock, we calculate the absolute difference between an AfD state election outcome and the most recent federal-level poll before the respective election, as depicted in Fig. B.2(a). This is

[8] We deem these latter two manifestations of particular importance for capturing a mobilizing effect.

[9] To ensure that we do not overestimate the treatment effect for Saxony, we drop all individuals residing in Brandenburg and Thuringia from the treatment group in this election window.

a straightforward approach, which allows us to include all available elections in our study period without making any pre-selection with respect to the informative value of single shocks. Yet, this strategy may also present some difficulties. First, in line with our theoretical approach, we may not observe any significant effects if absolute deviations of election outcomes from opinion polls are too small. While the effect we are interested in may still exist, we would be unable to observe it. Moreover, especially small *absolute* deviations may be less informative of the party performance than *relative* deviations from the poll.[10] In order to account for the possibility that our strategy ascribes white noise to our hypothesized treatment effect, we employ a number of alternative definitions of shocks, such as relative deviations from the pre-election poll in percent, shocks expressed in standard deviations from the pre-election poll and the mean difference from the pre-election poll for election windows with more than one election. We also show separate treatment effects for each election window to rule out that our results are driven by less informative, small shocks.

Regarding the choice of the opinion poll, an alternative strategy would be to compare state-level election outcomes to the last pre-election poll at the state level (see Fig. B.2(b)). Yet, this approach presupposes that state polls are consumed by all individuals in the treatment group. While state-election *outcomes*, especially those that stand out relative to the national average, are usually intensively covered by major media outlets, public interest outside the state in question is usually low in the lead-up to the election. One may therefore question whether individuals residing in other states actually consume the information about state-specific election polls. In fact, appendix Fig. B.3 supports these doubts. It displays the relative frequencies of Google searches of the respective state name in the four weeks around the election, showing sharp peaks only on election day. Hence, state-level politics are not very salient before the election, at least in other states. Individuals from other states are supposedly more likely to compare the information from state elections to general AfD support. Moreover, German state elections have been observed to function as *barometer elections*, that is, they often reflect changes in voter preferences at the federal level (Anderson and Ward, 1996; Jeffery and Hough, 2001, 2003), which further supports our approach.[11] We validate our argument by performing several estimations using state-level polls with no qualitatively different results.

In order to identify an unbiased estimate $\hat{\beta}$ of the treatment effect, four critical assumptions have to hold. First, respondents in the treatment and control group should not systematically differ with respect to individual-level characteristics. Therefore, we include a vector of individual-level controls, X_i, including demographic factors (gender, age, marital status and highest level of education) and socio-economic factors (employment status).[12] We also include a vector of political attitudes because vote intentions are likely correlated with other political views, comprising one's party choice in the last federal election, self-placement on the left–right scale, satisfaction with the government and the importance of the immigration issue. In addition, we also add control variables at the state-of-residence level to explicitly capture the impact of regional variation in economic and social issues on individual political behavior. $Z_{r,e}$ includes per capita household income, the population share of foreigners, per capita crimes and the unemployment rate. These covariates vary over states and election windows.

Since the *Politbarometer* surveys are designed to be representative of the voting-age population, the distribution of covariates in the treatment and control group is generally very similar, as is confirmed in a number of covariate balance tests (see appendix Table B.1). However, there may be slight but critical differences. Therefore, we apply the entropy balancing approach, as in Hainmueller (2012). With the use of this matching procedure, weights are determined and assigned to the observations in the control group so that their covariate distribution matches the covariate distribution in the treatment group (Hainmueller, 2012; Giani and Méon, 2021). For our analysis, we balance the first three moments of the covariate distributions. In addition, we apply an even more demanding balancing strategy by constructing entropy weighting schemes at the election-window level. Weights are applied such that a treated individual in election window δ_e is compared to control individuals from the same δ_e, matching on the state of residence as well as the demographic and socio-economic covariates. This strategy allows us to control for a large fraction of unobserved heterogeneity in socio-economic characteristics, state-specific voting patterns, and general trends in AfD support.[13] Applying the described matching approaches, we are left with negligible covariate differences between the control and treatment group.[14]

A second important assumption for identification is that state-level elections actually disseminate novel and unexpected information about aggregate-level AfD support (Muñoz et al., 2019). As demonstrated above, the specific case of the AfD is well-suited to identify such effects because of its young but controversial history combined with pronounced uncertainty about the party's true aggregate support. Yet, if respondents in the control group anticipated state-level election results, our estimate would be biased downwards. Referring to the above-mentioned lack of interest in state-specific political discussion before the election, we feel confident that respondents in other states were unable to anticipate election outcomes. In the unlikely case that such anticipation existed, we interpret our results as conservative estimates of the true treatment effect.

[10] For instance, in Saxony the AfD obtained 9.7 percent of votes compared to a federal poll result of 4.0 percent, whereas in Saxony-Anhalt the party obtained 24.3 percent of votes compared to a poll of 10 percent. While in Saxony-Anhalt the absolute deviation was much larger than in Saxony (14.3 vs. 5.7 percent), the AfD result exceeded the federal poll by about 43 percent in both elections.

[11] Individuals from other states could learn about state-level pre-election polls retrospectively if they consult media coverage on election day. If the election result of a party comes as a surprise, media reports may disseminate this information. However, this does not necessarily mean that individuals refrain from comparing this electoral performance to federal-level pre-election support. If the election result greatly exceeds the latest federal poll but individuals learn that it only slightly exceeds the state-level pre-election poll, our observed effect would only be biased downwards.

[12] The respective status groups are included as a set of dummies with full-time employment representing the reference group. Apart from that, the respondent can indicate being in school, in part-time employment, marginally employed, unemployed, in vocational training or retired.

[13] We do not match on political attitudes, as these variables could, in principle, be affected by the treatment. We elaborate on this issue in Section 4.2.3.

[14] We also present covariate imbalance statistics for each election window in the online supplementary material.

Third, there should be no unobserved macro-level events or trends that coincide with state-level elections (Muñoz et al., 2019). We include election-window fixed effects (γ_e) to account for the possibility that respondents in different election windows were exposed to unobserved time-variant macro-level effects, such as general differences in political demand across the electoral cycle. We also include state-of-residence fixed effects (θ_r) to control for unobserved heterogeneity across regions.

Fourth and finally, our model errors are likely correlated across individuals. To assess statistical inference, we cluster $\epsilon_{i,r,t,e}$ at different levels of aggregation and compare both clustered and bootstrapped standard errors. In particular, we expect within-cluster correlation at the level of the survey round because treatment is assigned at this level. Furthermore, we also expect within-cluster correlation at the regional level because existing evidence suggests profound differences in voting patterns across German states. Since the appropriate way of clustering in our case is not a priori clear, we follow Cameron and Miller (2015) and cluster at progressively higher levels. For our baseline estimates, we cluster the standard errors at the intersection of an individual's state of residence and the survey round. Furthermore, we also cluster separately by state of residence and survey round. However, since the number of clusters shrinks with an increasing level of aggregation (with 16 regions and 14 survey rounds), reliable inference becomes more challenging. This is the case in our setting since (a) the inference parameter is the coefficient of a treatment dummy variable, (b) the treatment is assigned at the cluster level and (c) there are only few treated clusters (MacKinnon, 2019). In such cases, using clustered standard errors could lead to over- or under-rejection of the null hypothesis (Cameron and Miller, 2015). To address this issue, we apply the so-called *wild cluster restricted bootstrap* (WCRB) procedure to estimate our standard errors, using the *boottest*-command in STATA (Roodman et al., 2019).

By means of the wild bootstrap procedures we are able to draw valid statistical inference, even when the number of clusters is small (Cameron et al., 2008; Cameron and Miller, 2015).[15] MacKinnon and Webb (2017, 2018) demonstrate that the wild cluster bootstrap with the null hypothesis imposed performs reasonably well when the number of treated clusters is not too small relative to the total number of clusters. These conditions are fulfilled in our case. Out of our 14 survey rounds, seven rounds (50 percent) are treated. Additionally, Cameron and Miller (2015) show that with only few clusters, inference can be improved by applying the 6-point distribution proposed by Webb (2013) instead of the common 2-point Rademacher distribution. MacKinnon and Webb (2017) also show that the restricted bootstrap procedure leads to more conservative *p*-values than the unrestricted approach and that it tends to moderately under-reject. Thus, the WCRB should be the most reliable and conservative approach in our case.[16]

We now use the above-outlined estimation strategy in order to test the following predictions from our theoretical model:

i. A positive (negative) election information shock σ increases (reduces) the likelihood that a respondent reports an AfD vote intention iff the shock is sufficiently large, i.e., if $|\sigma| > |\sigma_i^*|$ (Proposition 2).
ii. For a sufficiently large shock $|\sigma| > |\sigma_i^*|$, the effect on the individual vote intention increases in absolute shock magnitude.
iii. Without reputational concerns, larger shocks are required to produce a change in reported vote intentions (Proposition 3). Thus, for a shock $\sigma \in (\sigma_i^*, \sigma_i^{**}]$, we should observe no effect on reported vote intentions if
   (a) $\epsilon_j = 0$, i.e., we consider vote intentions for a non-populist party.
   (b) $b_i = 0$, i.e., the respondent does not care about a potential reputational loss from supporting party j.

4. Empirical results

4.1. Baseline results

In order to test predictions (i) and (ii), we proceed in two steps. First, as in (i), we investigate the effect of mere exposure to an election information shock by means of a sub-sample analysis. To that end, we split our sample into positive and negative election information shocks. As reported in Table 1, state-level vote shares positively deviated from federal polls in Saxony, Brandenburg, Thuringia, Hamburg, Baden-Wuerttemberg, Rhineland-Palatinate, Saxony-Anhalt, Mecklenburg-Hither Pomerania and Berlin. In contrast, state-level vote shares negatively deviated from federal polls in Bremen, Schleswig-Holstein, North Rhine-Westphalia and Saarland. In a second step, as in (ii), we use the absolute difference between the realized and predicted AfD vote shares to measure the intensity of an election information shock and test its effect on vote intentions.

Table 2 presents our baseline results from the sub-sample analysis. The variable *post-election* captures the treatment σ, i. e., whether a respondent is exposed to an election information shock. In panel A, we include election windows with positive election information shocks. In panel B, we examine election windows with negative shocks. We perform step-by-step OLS regressions in models (1) through (5). In model (1), beside the treatment effect, we only include state-of-residence and election-window fixed effects. In model (2), we add demographic and socio-economic covariates, while in model (3) we also control for individual political attitudes. In model (4), we add macroeconomic controls for the state of residence. In model (5), we apply the above-outlined entropy balancing method. Since our dependent variable is binary, we provide average marginal effects of a logit estimation in column (6). Regarding statistical inference, we report cluster-robust standard errors at the intersection of an individual's state of residence and the survey round in square brackets, with the respective *p*-values reported below. In addition, we report the *p*-values for the null hypothesis that $\beta = 0$ with WCRB standard errors clustered at the state-of-residence level and survey-round level, respectively.

[15] For a recent empirical application of this method, see also Gehring and Schneider (2018).

[16] See MacKinnon (2019) for a critical discussion of the wild bootstrap methods and alternative recent developments to approach the problem of fewer clusters. In Section II of the online supplementary material, we compare the WCRB estimates with alternative approaches that are recommended in the case of fewer clusters. The results show that the WCRB estimates consistently lie within the range of different estimates.

Table 2
Election information shocks and self-reported vote intentions: Baseline results.

DV: AfD vote intention	OLS (1)	OLS (2)	OLS (3)	OLS (4)	OLS (5)	Logit (6)
Panel A: Positive shocks						
Post-election	0.0260	0.0265	0.0244	0.0244	0.0207	0.0267
	[0.00583]	[0.00597]	[0.00582]	[0.00572]	[0.00465]	[0.00558]
Cluster-robust SE, *p*-value	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
WCRB by state of residence, *p*-value	0.002***	0.002***	0.002***	0.002***	0.002***	0.002***
WCRB by survey round, *p*-value	0.004***	0.004***	0.005***	0.004***	0.013**	0.001***
R^2	0.02	0.03	0.26	0.26	0.25	0.34
Observations	11,290	11,211	9,909	9,909	9,909	9,868
Panel B: Negative shocks						
Post-election	−0.0047	−0.0034	−0.0029	−0.0029	−0.0034	−0.0020
	[0.00447]	[0.00437]	[0.00470]	[0.00442]	[0.00392]	[0.00426]
Cluster-robust SE, *p*-value	0.293	0.443	0.537	0.515	0.386	0.633
WCRB by state of residence, *p*-value	0.484	0.606	0.695	0.695	0.588	0.750
WCRB by survey round, *p*-value	0.517	0.761	0.833	0.835	0.745	0.694
R^2	0.01	0.02	0.15	0.15	0.20	0.27
Observations	9,571	9,517	7,912	7,912	7,912	8,930
State-of-residence FE	Y	Y	Y	Y	Y	Y
Election-window FE	Y	Y	Y	Y	Y	Y
Demographics	N	Y	Y	Y	Y	Y
Socio-economics	N	Y	Y	Y	Y	Y
Political attitudes	N	N	Y	Y	Y	Y
State-level controls	N	N	N	Y	Y	Y
Entropy balancing	N	N	N	N	Y	N

Note: Estimates for OLS regressions in columns (1) to (5) and average marginal effects for logit regressions in column (6). Standard errors (SE) clustered by *state of residence* × *survey round* are reported in square brackets. The respective *p*-value is based on a standard Wald test under the null hypothesis that $\beta = 0$. *WCRB p-values* test the same hypothesis using the wild cluster restricted bootstrap (WCRB) with standard errors clustered by state of residence and by survey round, respectively. We apply the *boottest* command in STATA, using the 6-point distribution from Webb (2013). Election windows included in panel A are: SN, HH, (BW, RP, ST), (MV, BE). Election windows included in panel B are: HB, (SH, NW), SL. Demographics: age (18–70+, 10 cat.), age squared, gender (0–1), marital status (0–1), full set of dummies on education attainment (low, medium, high, in school). Socio-economics: full set of dummies on employment status (full time, part time, marginal, unemployed, in training, retired, other). Political attitudes: last vote AfD (0–1), self-positioning on left–right-scale (0–10), scaling of government performance (0–10), immigration perceived as most important issue (0–1). State-level controls: household income p.c., population share of foreigners, crimes p.c. (all measured in the pre-election year), unemployment rate (measured in the pre-election month). In models (1) to (4) and model (6), sample weights are used which are provided with the poll data. In model (5), matching weights from entropy balancing are applied based on the demographic and socio-economic covariates as well as the state of residence. R^2 reports the adjusted R-squared for models (1) to (5) and pseudo R-squared for model (6). In panel B, model (6), we do not include immigration attitude as a control because including this variable reduces the sample such that the WCRB methods are not feasible. The sign and magnitude of the coefficient estimate does not change upon inclusion of this covariate. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Inspecting our estimates for the sub-sample of positive shocks, we find a positive treatment effect that is statistically significant at the 1 percent level across specifications and levels of clustering in panel A of Table 2. In quantitative terms, the estimates indicate that the likelihood that an individual will report an AfD vote intention when interviewed shortly after a positive election information shock is about 2 to 2.7 percentage points higher than for respondents interviewed shortly before the election. In our preferred specification in column (5), we find that being exposed to a higher-than-expected AfD election outcome increases the propensity to report an AfD vote intention by about 2.1 percentage points. This is a sizable effect, corresponding to an increase of about 36 percent after a state election compared to the sample's average probability of reporting an AfD vote intention, which is 5.8 percent.

In contrast, we find no significant effect – if anything, we find a slight negative effect – of the exposure to lower-than-expected AfD vote shares on individually reported vote intentions, as displayed in panel B of Table 2. Importantly, statistical inference based on the different standard error estimates is highly consistent across specifications. Since we pool shocks of different magnitudes in Table 2, we cannot directly infer from these results whether negative election information shocks have no effect on vote intentions at all or whether our estimates are conflated by shocks that are too small to alter vote intentions, which would be in line with our theoretical predictions. Yet, given that the observed negative shocks are on average much smaller than the positive ones in our sample (1.8 vs. 7.5 percentage points[17]), the results point to a potential threshold effect: a shock in perceived aggregate AfD support alters self-reported vote intentions only if the observed shift is sufficiently large to provide an individual with a respective utility benefit.

We further investigate the link between the magnitude of shocks and vote intentions in Fig. 2. We measure *shock intensity* as the absolute difference (in percentage points) between the election outcome and the federal-level pre-election poll. Fig. 2(a) shows the effects of election information shocks at different intensities. The underlying model estimates a simple linear specification of the pooled sample, including *shock intensity* and its interaction with the treatment variable. In Fig. 2(b), we relax the linearity assumption

[17] Calculations are based on the absolute difference between the election outcome and the federal-level pre-election poll. See Table 1.

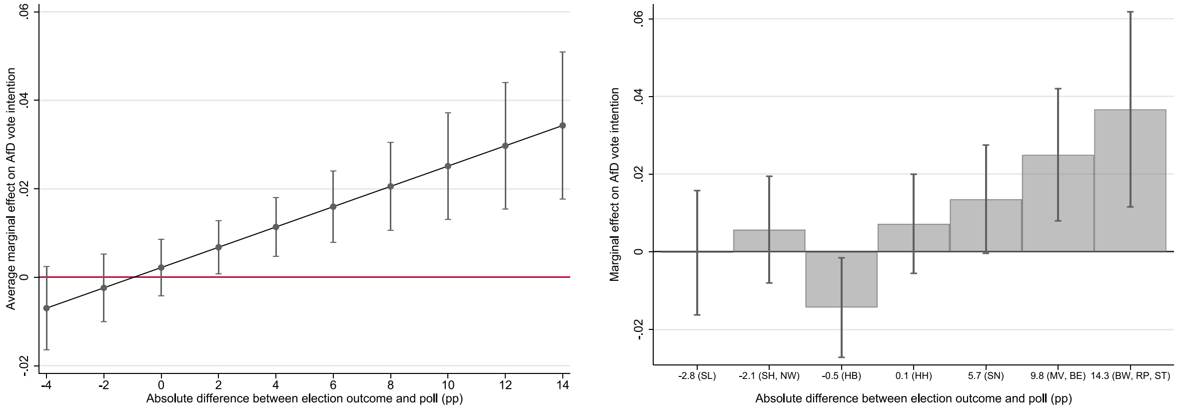

(a) Linear estimation

(b) Flexible estimation

Fig. 2. Shock intensity and vote intentions: Conditional treatment effects.
Note: Fig. 2(a): Conditional effects of *post-election* at different values of *shock intensity*. The conditional effects are calculated based on the specification of column (1) in appendix C.1. *shock intensity* measures the absolute difference (in percentage points) between the election outcome and the federal-level pre-election poll. For election windows with more than one election, we consider the largest shock in absolute terms. All other covariates are held constant at their means. Caps indicate 95% confidence intervals. Red horizontal line marks a marginal effect of 0. Fig. 2(b): Conditional effects of *post-election* for each election window. The conditional effects are calculated from estimating the specification of column (5) in Table 2 for the pooled sample, including a full set of election-window dummies and their interaction with *post-election* (Hamburg is the omitted baseline category). All other covariates are held constant at their means. Caps indicate 95% confidence intervals. Regression coefficients reported in appendix Table C.2.

and interact *post-election* with a full set of election-window dummies. This allows us to test whether there is a non-monotonous effect of election information shocks on vote intentions which differs between positive and negative shocks and which is indicative of potential threshold effects of shock size. For positive shocks, the figures show that the impact on AfD vote intentions indeed increases with shock intensity. In addition, there is suggestive evidence of a threshold effect for positive shocks. The coefficients turn statistically significant for a deviation of 2 percentage points or more from the pre-election poll.

For negative shocks, the evidence is less clear. We find a negative effect for the election in the state of Bremen, which is significant at the 5 percent level. This suggests that negative election surprises can, in principle, affect individual vote intentions. In addition, Fig. 2(b) also hints at a non-monotonous effect for negative shocks. Yet, if we assume that the absolute shock threshold for both positive and negative shocks is identical (or at least very similar), the observed negative shocks are simply not large enough to induce substantial shifts in vote intentions. Furthermore, Fig. 2(b) suggests that our results are not exclusively driven by earlier or later elections, as such a pattern would imply that our treatment effect simply reflects a time trend in overall AfD support.[18] Generally, the evidence thus supports predictions (i) and (ii). We investigate some alternative explanations throughout the remaining sections.

4.2. Robustness of baseline findings

4.2.1. Alternative definitions of election information shocks

For our baseline definition of a shock, we use the difference between the AfD vote share in a given state election and the most recent *federal-level* AfD poll. With this definition, we assume that respondents residing in states other than the election state refer to the federal AfD support when evaluating the state-level election outcome. Yet, this definition may conflate a surprise regarding the absolute strength of the AfD in a given state and a surprise regarding its relative strength in the election state compared to the national average.

To disentangle these effects, we propose an alternative definition of election information shocks, using the deviation of state-level election outcomes of the most recent pre-election *state-level* poll. Doing so affects both the assignment of election windows to the sub-samples of positive and negative shocks as well as the shock magnitude.[19] In Fig. 3(a), we present results for the link between shock intensity and vote intentions. Similar to our baseline results, we again find a significantly positive and increasing effect for positive shocks but no significant effect for negative shocks, which are again rather small in magnitude.

In addition, our strategy to define a shock as an absolute – rather than a relative – deviation of the election outcome from the most recent poll may raise concerns, as explained in Section 3. To investigate whether our results are sensitive to the chosen benchmark, we repeat our analysis in Fig. 3(b), expressing the difference between the election outcome and the federal-level poll in standard deviations of the poll. The results resemble our baseline estimates, again suggesting a threshold effect for positive shocks at around one standard deviation (2.31 percentage points). Beyond that, appendix Fig. C.1 displays further evidence that our results are robust to other alternative definitions of election information shocks.

[18] For additional results on the timing of shocks, see Section II of the online supplementary material.

[19] Details on shocks defined by deviations from state-level polls is given in Section II of the online supplementary material.

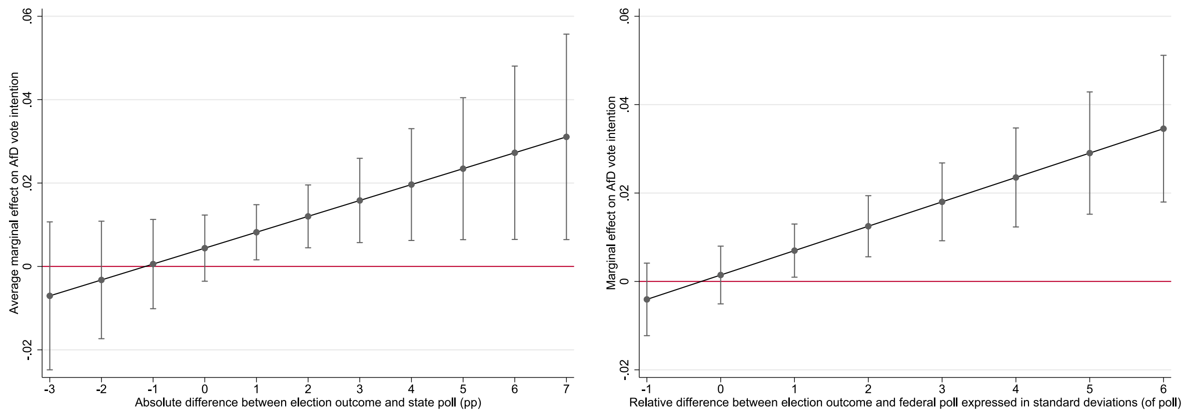

(a) Absolute difference from state-level polls.

(b) Relative shock intensity.

Fig. 3. Alternative definitions of shock intensity.
Note: Conditional effects of *post-election* at different values of *shock intensity*. The conditional effects are calculated based on the specifications of columns (3) and (5) in appendix Table C.1. Fig. 3(a): *shock intensity* measures the absolute difference (in percentage points) between the election outcome and the *state-level* pre-election poll. Fig. 3(b): *shock intensity* measures the absolute difference between the election outcome and the federal-level pre-election poll expressed in standard deviations of the poll. For election windows with more than one election, we consider the largest shock in absolute terms in both specifications. All other covariates are held constant at their means. Caps indicate 95% confidence intervals. Red horizontal line marks a marginal effect of 0.

Table 3
Results for placebo election information shocks.

DV: AfD vote intention	Placebo treatment I		Placebo treatment II	
	Positive shocks (1)	Negative shocks (2)	Positive shocks (3)	Negative shocks (4)
Counterfactual post-election	−0.0030	−0.0067	−0.0039	0.0006
	[0.00341]	[0.00431]	[0.00448]	[0.00360]
Cluster-robust SE, *p*-value	0.388	0.122	0.385	0.869
WCRB by state of residence, *p*-value	0.395	0.270	0.489	0.893
WCRB by survey round, *p*-value	0.477	0.230	0.699	0.517
Adj. R^2	0.26	0.21	0.24	0.21
Observations	10,467	7,913	10,263	7,915
State-of-residence FE	Y	Y	Y	Y
Election-window FE	Y	Y	Y	Y
Demographics	Y	Y	Y	Y
Socio-economics	Y	Y	Y	Y
Political attitudes	Y	Y	Y	Y
State-level controls	Y	Y	Y	Y
Entropy balancing	Y	Y	Y	Y

Note: Table displays estimates for OLS regressions. Standard errors (SE) clustered by *state of residence × survey round* are reported in square brackets. For *Placebo treatment I*, we use the two most recent survey rounds *before* the actual election. For *Placebo treatment II*, we use the two most recent survey rounds *after* the actual election. *Counterfactual* election windows for sample with positive shocks: SN, HH, (BW, RP, ST), (MV, BE). *Counterfactual* election windows for sample with negative shocks: HB, (SH, NW), SL. Control variables as in Table 2. State-level control variables have been adapted to the corresponding time periods when necessary. Matching weights from entropy balancing are based on the demographic and socio-economic covariates as well as the state of residence. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

4.2.2. Placebo treatment

A major threat to identification in our empirical model is the presence of unobserved time-varying factors at the macro level which may affect an individual's disposition towards reporting an AfD vote intention but which are not captured by the election-window fixed effects. To check for general trends in AfD support or other confounding events, we conduct placebo tests for *counterfactual* election windows. To this end, we identify windows of two subsequent surveys that are not affected by a state election but that may be driven by the same cyclical or macro-level factors as our examined elections. We apply two different placebo tests. First, we choose the two most recent surveys *prior* to an election window. This method implies that the control group from our baseline regression now becomes the treatment group in the placebo regression. Second, we run the same regressions using the first two surveys *after* the election window, i. e., the treatment group from our baseline regression becomes the control group in the placebo regression. Table 3 displays no significant effects of these counterfactual treatments. These results provide strong support for our claim that it is in fact the election information shock which shapes individually reported vote intentions.

Table 4
Sample variations and sequential *g* estimation.

DV: AfD vote intention	All respondents		Mediator = Satisfaction with federal government		Mediator = Immigration as most important issue?	
	Positive shocks (1)	Negative shocks (2)	Positive shocks (3)	Negative shocks (4)	Positive shocks (5)	Negative shocks (6)
Post-election	0.0209	−0.0028	0.0211	−0.0028	0.0179	−0.0004
	[0.00455]	[0.00375]	[0.00484]	[0.00436]	[0.00468]	[0.00439]
Cluster-robust SE, *p*-value	0.000***	0.455	0.000***	0.529	0.000***	0.925
WCRB by state of residence, *p*-value	0.001***	0.634	0.003***	0.686	0.011**	0.951
WCRB by survey round, *p*-value	0.011**	0.780	0.005***	0.857	0.005***	0.935
Adj. R^2	0.25	0.19	0.03	0.02	0.03	0.02
Observations	11,641	8,456	9,909	7,912	9,909	7,912
State-of-residence FE	Y	Y	Y	Y	Y	Y
Election-window FE	Y	Y	Y	Y	Y	Y
Demographics	Y	Y	Y	Y	Y	Y
Socio-economics	Y	Y	Y	Y	Y	Y
Political attitudes	Y	Y	N	N	N	N
State-level controls	Y	Y	Y	Y	Y	Y
Entropy balancing	Y	Y	Y	Y	Y	Y

Note: Table displays estimates for OLS regressions. Standard errors (SE) clustered by *state of residence* × *survey round* are reported in square brackets. In columns (1) and (2), we add the respondents from the respective election state to the baseline sample. In columns (3) to (6), we follow the sequential *g* approach as outlined in Acharya et al. (2016). Coefficients are for the second-stage regression, capturing the ACDE of election information shocks on AfD vote intentions. Mediator variable is as indicated in the respective panel. The remaining variables from the *political attitudes* vector are treated as *intermediate confounders* (Acharya et al., 2016). Election windows for the samples of positive shocks: SN, HH, (BW, RP, ST), (MV, BE). Election windows for the samples of negative shocks: HB, (SH, NW), SL. Control variables as in Table 2. Matching weights from entropy balancing based on the demographic and socio-economic covariates as well as the state of residence. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

4.2.3. Sample variations and alternative specifications

In our baseline specification, we focus on vote intentions of respondents in states other than the election state in order to avoid problems of joint identification. In columns (1) and (2) of Table 4, we relax this restriction and add the respondents from the respective election state to our sample. The coefficient estimates closely resemble our baseline findings.[20] In Section II of the online supplementary material, we also show that our baseline results are robust to alternative bootstrap methods, alternative matching strategies and alternative data sources for pre-election polls.

We also test whether our results are consistent for other individual-level measures of AfD support. To this end, we use a respondent's general *party affinity* as an alternative dependent variable. The results confirm the proposed link between election information shocks and self-reported political attitudes but again only for positive shocks as shown in appendix Table C.4. With much caution, one may understand this result as tentative evidence of the fact that sudden shifts in observed aggregate-level support for the AfD also affect an individual's disposition to report a more persisting party preference.

Finally, the use of political attitudes as control variables in our baseline estimation may raise concerns about a potential *post-treatment bias* (Muñoz et al., 2019). That is, the obtained treatment effects may be biased if a political attitude variable is itself influenced by the treatment and impacts on the reported vote intention (Acharya et al., 2016; Imai et al., 2010). If so, it would be unclear whether a significant estimate for the election information shock represents the direct treatment effect or an indirect effect that operates through a political attitude variable as a mediator (Acharya et al., 2016).

This role of a mediator may particularly apply to two of our political attitudes: *satisfaction with the federal government* and *perception of immigration as the most important political issue*. The observation of a shock in support for a right-wing challenger party like the AfD may alter individual perceptions about how well the government performs and how pressing an issue immigration is, which, in turn, impacts on the individual vote intention. Thus, the question is whether there is a direct effect of election information shocks on AfD vote intentions, which represents the change in the social desirability bias associated with openly supporting the AfD, or whether we only observe the indirect effect of political attitudes.

Simply leaving political attitudes out of the analysis, however, would cause an omitted variable bias (Acharya et al., 2016; Imai et al., 2010). Our goal is, therefore, to separate the *average controlled direct effect* (ACDE) of the election information shock from the *average natural indirect effect* (ANIE) that the shock has on vote intentions through a political attitude variable as a mediator (Acharya et al., 2016). To disentangle the effects, we use the sequential *g* estimation as proposed in Acharya et al. (2016).[21] We show the results of our preferred baseline specification using the sequential *g* estimation in columns (3) to (6) of Table 4. In columns (3)

[20] Additionally, we repeat several other estimations for the sample comprising all respondents to show that the results are stable. The linear and flexible estimations displayed in Fig. 2 are repeated for the full sample in appendix Fig. C.2. We also conduct the placebo analysis and some of the mechanism tests from Section 4.3 for the full sample and we report the results in appendix Table C.3.

[21] We provide a more comprehensive explanation of how this approach works in Section II of the online supplementary material. For recent applications of the sequential *g* estimation, see for instance Depetris-Chauvin and Özak (2020), Carpena and Zia (2020), Brown et al. (2019), Moya and Carter (2019).

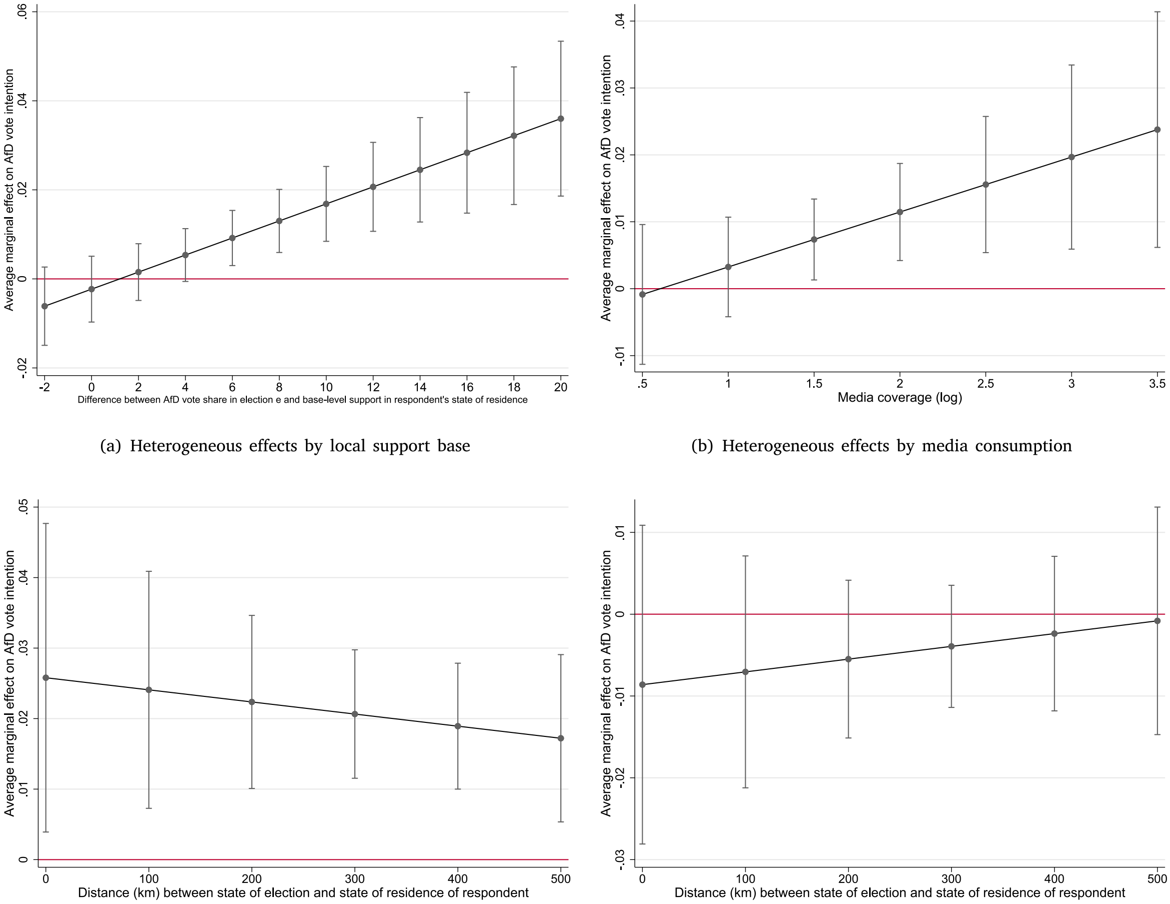

(a) Heterogeneous effects by local support base

(b) Heterogeneous effects by media consumption

(c) Heterogeneous effects by distance — positive shocks

(d) Heterogeneous effects by distance — negative shocks

Fig. 4. Heterogeneous effects of election information shocks on vote intentions.
Note: Conditional effects of *post-election* at different values of the respective interacted variable. The conditional effects are calculated based on the specifications of column (1), (5), (6) and (7) in appendix Table C.5. In sub Fig. 4(a), *local support base* measures the difference (pp) between the AfD vote share in state election e and the vote share that the AfD obtained in the state of residence of respondent i in the 2013 federal election. For state election windows with more than one election, we consider the largest vote share. In Fig. 4(b), *media consumption, log* measures relative frequencies (in logs) of weekly Google searches of *AfD* (or *Alternative fuer Deutschland*). For the control group, we take the value in the control WoS. For the treatment group, we use the value in the WoE to account for a direct link between elections and media consumption. In sub Figs. 4(c) and 4(d), *distance* measures the shortest distance between the capitals of the election state and the state of residence of the respondent. For election windows containing more than one election we consider the distance to the closest capital of the respondent's state of residence. In all specifications, the other covariates are held constant at their means. Caps indicate 95% confidence intervals. Red horizontal line marks a marginal effect of 0.

and (4), we treat *government satisfaction* as the potential mediator and in columns (5) and (6), the mediator variable is *immigration as the most important issue?*.

The coefficients reflect the ACDE of the treatment and show that, if anything, our baseline estimates in Table 2 suffer from a very small quantitative post-treatment bias. In Section II of the online supplementary material, we further check the validity of our analysis by also repeating the placebo tests using the sequential g estimation and by treating the other two political attitude variables as potential mediators. The estimates remain qualitatively unchanged, so we feel confident that our results which include political attitudes do not suffer from a severe post-treatment bias.

4.2.4. Heterogeneous effects

Our baseline results suggest that the impact of election information shocks depends on the sign and the magnitude of the specific shock. In this section, we investigate whether there is additional effect heterogeneity, focusing on three potential sources: differences in electoral conditions, spatial-cultural proximity to the election state, and media consumption. Following up on the fact that treatment in our setting is rather an *intention to treat*, we argue that individuals are more likely to recognize the treatment if they live in AfD-friendly areas, in regions closer to the election state or if they are subject to more extensive media coverage about the party.

The informational content of an election may subjectively depend on the level of AfD support in a respondent's own region. An unexpectedly high AfD vote share in a state election may impact differently on a respondent living in a region where the AfD has performed well in previous elections and where AfD candidates are already part of the legislature at the local or state level, compared to a respondent in a region without a substantial support base. Therefore, we investigate how respondents react to the AfD's electoral performance depending on the level of party popularity in the respondent's state of residence. To this end, we construct a variable *local support base* that measures the difference between the AfD vote share in a given state election and the vote share that the AfD obtained in the federal election in 2013 in the respondent's state of residence.[22] Fig. 4(a) presents the conditional treatment effects at different values of this variable. The results indicate that individuals react more strongly to shocks in areas with low local support for the AfD, suggesting that our treatment effect in fact depends on the extent to which perceived AfD popularity changes upon observation of an election information shock. In appendix Table C.5, we also show results for other approaches that measure different electoral conditions.

As a second source of heterogeneity, we exploit geographic variation in AfD support. The largest positive shocks occurred in states in former East Germany, where the AfD has been particularly successful in mobilizing voters (see Table 1). Generally, the difference in voting patterns between the former eastern and western states may correspond to a division along socio-economic and cultural lines. Thus, when reflecting on dominant public views, individuals may be more strongly affected by changes in a region they consider their socio-cultural or spatially close peer group (Arbatli and Gomtsyan, 2019). This gives rise to the question of whether, on the one hand, respondents are more likely to identify with voting patterns in geographically close regions and, on the other hand, whether respondents react differently to a shock depending on whether they reside in either the former eastern or western part of the country. In Figs. 4(c) and 4(d), we show that the effect of a positive shock is in fact conditional on the geographic distance between the election state and the respondent's state of residence. As geographic distance likely reflects cultural distance, the results suggest that cultural proximity does indeed increase the responsiveness to positive election information shocks. Interestingly, we again find no similar pattern for negative shocks. This is an important observation as it allows us to exclude the specific location of negative shocks as an explanation for the insignificant effect.

In appendix Fig. C.4, we explicitly exploit the East–West divide by splitting both election windows and survey respondents according to their affiliation with former East or West Germany. The results on these dyadic regressions show that respondents residing in formerly West German states react more strongly to positive shocks in either part of the country. This pattern suggests that larger-than-expected AfD support specifically encouraged respondents in states with previously low AfD advocacy to reveal a respective vote intention, in line with Fig. 4(a).

Finally, we investigate whether the impact of shocks differs by the extent to which the AfD election results receive attention through aggregate media consumption. A critical assumption of our empirical setup is that individuals actually receive the information about election outcomes in other states. Therefore, one important transmission channel is media consumption. In order to measure this, we use the relative frequencies (in logs) of Google searches for the party name *AfD* (or *Alternative fuer Deutschland*) in the survey weeks under observation. This variable provides a proxy for the average level of interest in the AfD at the national level. Fig. 4(b) shows the respective conditional treatment effect. The effect of election information shocks on AfD vote intentions increases with the extent of media consumption, as expected. Individuals react more strongly to the information shock if the AfD was a frequent subject of media discussions.[23] This is a crucial result for the causal interpretation of the identified relationship as it suggests that the differences in AfD vote intentions between the treatment and control groups are in fact driven by information dissemination to individuals in other parts of the country.

Overall, the results presented in this section underpin our argument that German state elections affect individual AfD vote intentions by providing novel and unanticipated information about aggregate preferences. Importantly, the results confirm that the way in which individuals respond to election information shocks depends on the extent to which these individuals are (likely) affected by changed perceptions of the party's aggregate popularity. In addition, the evidence does not suggest that the lack of a significant treatment effect for negative shocks is driven by a relevant underlying source of heterogeneity. Rather, the results strongly support our presumption that the negative shocks in our sample are simply not large enough to affect individual vote intentions. From that we can draw the tentative, yet notable conclusion of a reinforcing upward trend in expressed populist support during our study period.

4.3. Mechanisms

4.3.1. Reputation vs. bandwagon channel

As outlined in Section 2.3, we argue that election information shocks can affect an individual's decision to report an AfD vote intention through both the bandwagon and the reputation channel. In order to investigate whether it is indeed reputational concerns

[22] Given our data structure, the AfD vote share from the earlier 2013 federal election is the only available source of geographic variation in the local AfD support base that we have. Since the *Politbarometer* contains no disaggregated information about a respondent's district or county, using earlier local-level election outcomes is not an option.

[23] Unfortunately, with our measure, we cannot account for the tonality of the respective media content. Yet, since we focus on media consumption around the time of state elections and since AfD election outcomes are official figures, we can assume that the AfD election outcome as such is provided by any media outlet found on Google. Thus, individuals learn about the same official figures, independently of whether this result may be evaluated as good or bad by a given media outlet.

that explain the observed contagious effect, we disentangle the channels by isolating situations where only one of the them is active. Here, we focus on shutting down the reputation channel and propose two empirical tests relating to prediction (iii).

First, our theoretical framework implies that the impact of a given election information shock for a party on the likelihood to report a respective vote intention is smaller if the policies promoted by this party are free from reputational concerns. Thus, we apply our baseline strategy to two arguably non-populist German parties, *The Greens* and the *Free Democratic Party* (FDP), a smaller liberal party. Both parties are located near the center (center-left and center, respectively) of the German party system and can be characterized as established, uncontroversial parties who have both served as junior coalition partners in federal governments in the past twenty years. At the same time, these parties have traditionally specialized in rather narrow programs (environmental policies in case of The Greens, liberal market policies in case of the FDP) and have repeatedly been subject to waves of party popularity depending on the salience of the promoted policy issues (Volkens et al., 2020; Bakker et al., 2019). The fact that party popularity has been subject to frequent fluctuations is important for our research design because it ensures that state-level elections actually reveal unanticipated information about aggregate support. Building on our theoretical predictions, we expect smaller and less consistent effects for The Greens and the FDP, as compared to the AfD, because, if anything, election information shocks affect vote intentions for these non-populist parties only through the bandwagon channel but not through the reputation channel.

In order to avoid that the recent electoral successes of the AfD simultaneously affect changes in votes for The Greens and the FDP and vice versa, we focus on the electoral cycle from 2009 to 2013 for the two non-populist parties. During this period, the AfD did not yet exist. To identify suitable election windows, we follow the procedure explained in Section 3.2. We measure the direction and intensity of election information shocks according to our baseline definition, i. e., taking the absolute difference between the state-level election outcome of The Greens (FDP) and the pre-election federal-level poll.[24]

Figs. 5(a) and 5(c) show the conditional treatment effects at different shock magnitudes for The Greens and the FDP. While the figures depict a positive association between shock intensity and vote intentions for either party, most effects are statistically insignificant, suggesting that bandwagon motives alone are not sufficient to explain the strong contagion effects we find for the AfD. In appendix Tables III.2 and III.3, we also provide coefficient estimates from a sub-sample analysis. While we find some statistical evidence that a larger-than-expected election outcome increases the share of respective vote intentions for The Greens and the FDP, these effects are much smaller, compared to our baseline results for the AfD, and they are less consistent.[25]

Figs. 5(b) and 5(d) depict the results from the more flexible shock intensity model where we interact our treatment variable with a full set of election-window dummies. Again, we do not find a systematic pattern between shock size and vote intentions for these non-populist parties. Contrasting this evidence with the pronounced and robust pattern for the AfD, it supports our supposition that reported vote intentions for the AfD are driven by both bandwagon motivations *and* reputational concerns.

We here choose to exploit the empirical evidence for The Greens and the FDP to test the suggested underlying mechanism. Nevertheless, for the sake of completeness, we provide further evidence on potential contagion effects for all other parties in the German parliament in our observation period. Section III in the online supplementary material presents descriptives, estimation results as well as interpretations of the observed effects. Specifically, we repeat the baseline estimation as in Table 2 and the robustness tests as in Tables 3 and 4 for the other parties. Similar to the results for The Greens and the FDP, we find single contagion effects, which can mostly be traced back to bandwagoning and time trends. The effects are smaller and not very robust, as compared to the effects we find for the AfD.

To further strengthen our argument, we apply a second strategy to test whether the effect of a specific shock in observable aggregate support for the AfD differs when we compare respondents with and without reputational concerns. From prediction (iii), it follows that the impact of a given shock on vote intentions is smaller for respondents that do not care about a potential reputational loss than for respondents who do. Unfortunately, we cannot directly measure reputational concerns. Therefore, we infer their existence from a number of survey questions. In a first step, we restrict our sample to respondents who report that they have not voted for the AfD in the second-to-last federal election in 2013. Since the AfD was still a marginal political platform at that time and had not yet adopted its strong right-wing stance, these respondents presumably have not voted for the AfD because of bandwagon motives rather than reputational concerns. To fully isolate the bandwagon channel, we add a second step in which we split this group of non-AfD voters from 2013 into respondents who state that the AfD is their preferred party and those who rank another party first. The underlying survey item asks respondents to rank a number of parties in descending order of assessed competence and appeal. We use this ranking to create the binary variable *AfD preferred party* that takes the value 1 if the respondent ranks the AfD first and 0 if they rank any other party first. Respondents who state that the AfD is their preferred party obviously have no reputational concerns and freely communicate their sympathy. Thus, if – among this group of respondents – the share who reports an AfD vote intention systematically differs before and after an election information shock, such differences should be entirely driven by bandwagon motives. In contrast, for respondents who state that they prefer another party, systematic differences in reported AfD vote intentions before and after an election information shock may be driven by both channels.

In columns (1) and (2) of Table 5, we interact the treatment variable with *AfD preferred party* for both positive and negative shocks. However, we find no statistically significant effect for the group of respondents without reputational concerns but find a positive effect, that is statistically significant at the 1 percent level, for the group of respondents with reputational concerns

[24] In Section III of the online supplementary material, we provide details on the election windows included and the calculation of shock intensities.

[25] Being exposed to a larger-than-expected vote share for The Greens (FDP) increases the likelihood that an individual will report a respective vote intention by about 1 percentage point. Compared to the sample's average probability to report a Green or an FDP vote intention of 11.6 and 6.8 percent, respectively, this translates into an increase of about 9 and 15 percent. In comparison, our baseline estimate for the AfD suggest an increase in the average probability to state a respective vote intention of 36 percent.

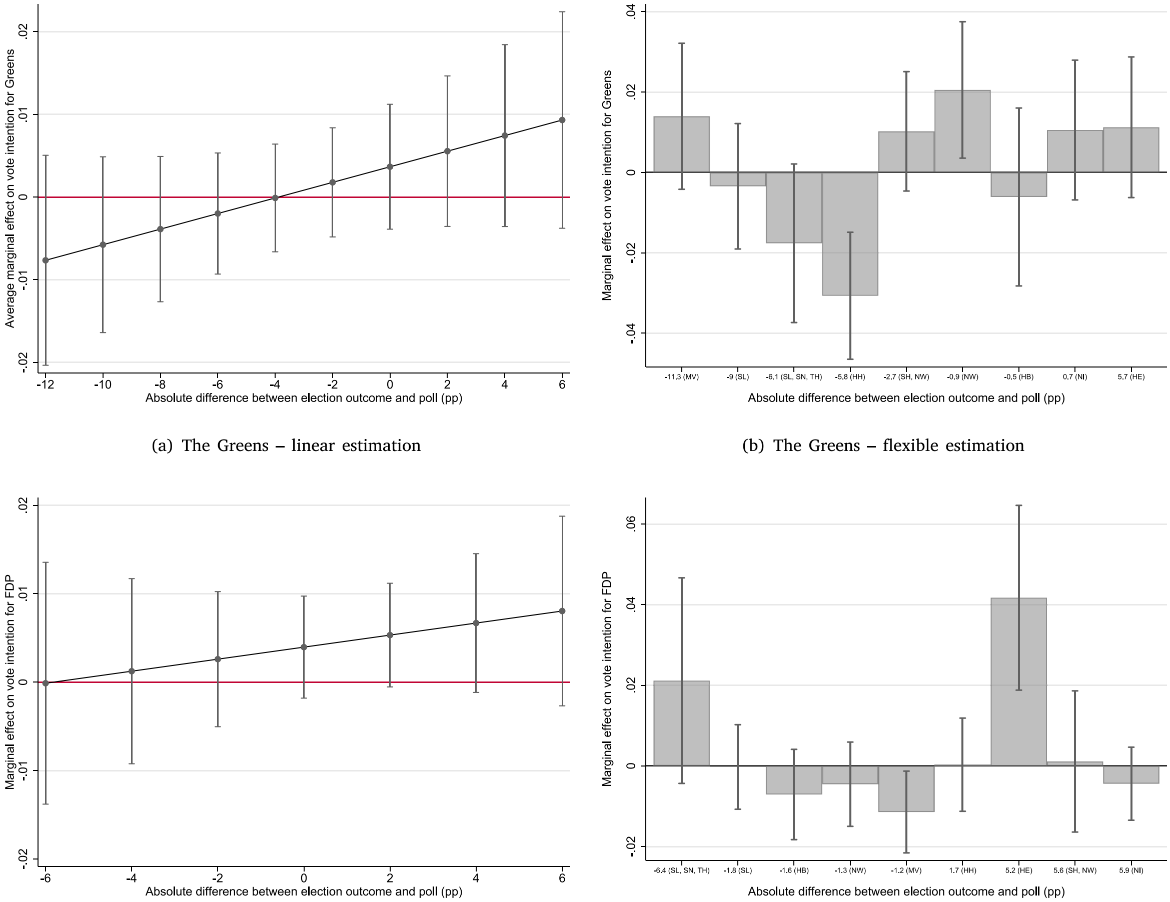

(a) The Greens – linear estimation

(b) The Greens – flexible estimation

(c) Free Democratic Party (FDP) – linear estimation

(d) Free Democratic Party (FDP) – flexible estimation

Fig. 5. Shock intensity and vote intentions for non-populist parties: Conditional treatment effects.
Note: Dependent variable: (a) and (b) vote intentions for The Greens (Bündnis 90/Die Grünen), (c) and (d) vote intentions for the Free Democratic Party (FDP). The graphs in (a) and (c) display conditional effects of *post-election* at different values of *shock intensity*. The conditional effects are calculated based on the specification of columns (1) and (2) Table III.7 in the online supplementary material. *Shock intensity* measures the absolute difference (in percentage points) between the election outcome and the most recent federal-level pre-election poll for the respective party. For election windows with more than one election, we consider the largest shock in absolute terms. All other covariates are held constant at their means. Figures (b) and (d) show conditional effects of *post-election* for each election window. The conditional effects are calculated by estimating the specification of model (5) in Tables III.2 and III.3, respectively, for the pooled sample, including a full set of election-window dummies and their interaction with *post-election* (Hamburg is the omitted baseline category). All other covariates are held constant at their means. Sample period: 2009 to 2013. Caps indicate 95% confidence intervals. Red horizontal line marks a marginal effect of 0.

(along with bandwagon motives).[26] In line with our theoretical predictions, positive election information shocks seem to alleviate reputational concerns with regard to AfD vote intentions. While we cannot ultimately rule out the bandwagon channel in our setting, these results make us confident that reputational concerns are the dominant driving force explaining the specific contagion effect when it comes to populist party support.

4.3.2. Strategic motivations

In addition to the proposed channels, we discuss one alternative mechanism, namely *strategic* motivations to change one's vote intention upon the observation of an election information shock (Feddersen and Pesendorfer, 1997; Franklin et al., 1994).

It has been established that individuals may use their reported vote intentions strategically in order to affect other voters and/or party programs in later elections (Piketty, 2000; Castanheira, 2003). Transferring this idea to our context, it is conceivable that respondents who would not actually vote for a controversial party like the AfD may still report a respective vote intention in the *Politbarometer* in order to communicate their discontent with the ruling parties or to promote political change. In Germany, this type

[26] In appendix Table C.6, we show that these results are robust to a number of alternative specifications.

Table 5
Election information shocks and vote intentions: Mechanisms.

DV: AfD vote intention	Sample restriction: *AfD vote in 2013* = 0		Sample restriction: *Preferred party* = *CDU*	
	Positive shocks (1)	Negative shocks (2)	Positive shocks (3)	Negative shocks (4)
Post-election	0.0113	−0.0016	0.0193	−0.0057
	[0.00353]	[0.00275]	[0.00595]	[0.00507]
Cluster-robust SE, *p*-value	0.002***	0.564	0.002***	0.267
WCRB by state of residence, *p*-value	0.007***	0.596	0.003***	0.419
WCRB by survey round, *p*-value	0.105	0.724	0.025**	0.516
Post-election × AfD preferred party	0.0266	0.0633		
	[0.05534]	[0.05802]		
Cluster-robust SE, *p*-value	0.632	0.278		
WCRB by state of residence, *p*-value	0.558	0.146		
WCRB by survey round, *p*-value	0.724	0.554		
Post-election × CDU vote share			0.0083	0.0052
			[0.00545]	[0.00670]
Cluster-robust SE, *p*-value			0.131	0.441
WCRB by state of residence, *p*-value			0.280	0.529
WCRB by survey round, *p*-value			0.088*	0.167
Adj. R^2	0.41	0.42	0.15	0.07
Observations	8,809	7,137	3,809	2,734
State-of-residence FE	Y	Y	Y	Y
Election-window FE	Y	Y	Y	Y
Demographics	Y	Y	Y	Y
Socio-economics	Y	Y	Y	Y
Political attitudes	Y	Y	Y	Y
State-level controls	Y	Y	Y	Y
Entropy balancing	Y	Y	Y	Y

Note: Table displays estimates for OLS regressions. Standard errors (SE) clustered by *state of residence* × *survey round* are reported in square brackets. The regressions in columns (1) and (2) are restricted to respondents who report to have voted for a different party than the AfD in the 2013 federal election. *AfD preferred party*= 1 if the respondent ranks the AfD first among all parties; 0 if the respondent ranks a different party first; missing if no party rank is reported. The regressions in columns (3) and (4) are restricted to respondents who rank the CDU/CSU first among all parties. *CDU vote share* measures the vote share that the CDU obtained in the respective state-level election. The variable is standardized with a mean of 0 and a standard deviation of 1. Control variables as in Table 2. Matching weights from entropy balancing based on the demographic and socio-economic covariates as well as the state of residence for each sample restriction. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

of protest behavior has been observed for many supporters of the conservative party (CDU)[27] during our sample period (Dostal, 2017; Wurthmann et al., 2020). Following the refugee crisis in 2015, many supporters of the CDU, which was the leader of the federal government at that time, were disappointed with the liberal immigration policies. Since the AfD was the only political platform openly promoting strict anti-immigration policies, frustrated conservatives thus had an incentive to report an AfD vote intention in order to voice their protest and encourage the CDU to adapt some AfD policies.

Linking this to our model of election information shocks, a larger-than expected AfD vote share in a state election improves the party's image as a serious electoral competitor. In line with the bandwagon channel, disappointed conservative supporters may feel encouraged to report an AfD vote intention. Yet, respondents with a genuine CDU preference do not actually want to see the AfD win. Thus, they only report this strategic vote intention for the AfD upon observation of a large winning margin for the CDU. Following Piketty (2000), self-reported AfD vote intentions are then less likely to translate into an office win for the AfD by motivating others to actually cast an AfD vote.

We test this mechanism by focusing on the sub-sample of respondents who report that they generally prefer the conservative party according to the party ranking variable introduced above. We then interact our treatment variable with the (standardized) CDU vote share in the respective state election. If our results are – at least partly – driven by strategic motivations of conservative supporters, we would expect to find a larger treatment effect after positive AfD support shocks when the winning margin of the CDU is high (and a smaller effect after negative shocks). Yet, as shown in columns (3) and (4) of Table 5, we find no robust evidence of such a mechanism.[28] Strategic motivations may still play a role, yet, the results support the notion that the observed contagion effect in populist support cannot solely be explained by this alternative mechanism.

[27] In Germany, the conservative platform is a party union comprising the CDU (Christian Democratic Union) and the CSU (Christian Social Union, active only in Bavaria). In the following, we refer to the entire platform as the CDU or the *conservative party*.

[28] In appendix Fig. C.3, we also display the conditional treatment effects at different values of *CDU vote share* and show that there is no significant association between CDU vote shares and vote intentions.

Overall, the results in this section suggest that shocks in aggregate party support have different effects on vote intentions for populist and non-populist parties, and that reputational concerns are the most likely mechanism to explain the dynamics of populist vote intentions that we observe in the baseline analysis.

5. Conclusion

Is populist voting contagious? In order to explain the dynamic rise of populist and nationalist movements around the globe, it appears important to understand the role of social contagion in political behavior. In this paper, we examine whether unexpected shifts in observed social support for a right-wing populist party encourage individuals to report a respective vote intention in survey interviews.

We apply a quasi-experimental event-study design for Germany, where the right-wing populist AfD has registered considerable support among the electorate since it was founded in 2013. We presuppose that voters are hesitant to openly support this controversial movement in an attempt to avoid reputational losses among others. However, the observation of a higher aggregate support level may serve as a signal of the party's improved social standing and thereby increase the individual propensity to openly report a respective preference. We use the staggered state elections in the German federal system as a source of new information about general AfD preferences and define *election information shocks* as the deviation of AfD state election outcomes from previously known opinion polls at the federal level. Using repeated cross-sectional data from the German *Politbarometer* survey, we then test if the average likelihood that an individual will report an AfD vote intention differs between individuals interviewed right after a state election and individuals interviewed right before the election, who are, therefore, unaffected by the information shock.

Our empirical results provide systematic evidence that information shocks associated with larger-than-expected AfD vote shares in state elections raise subsequently reported AfD vote intentions in *other* states by up to 2.7 percentage points. Our results are consistent with other measures of individual-level AfD support and are robust to accounting for different levels of within-cluster correlation of model errors. Applying entropy balancing and conducting placebo tests, we find convincing support for a causal effect of election information shocks on individual vote intentions. Our results on the heterogeneity of the effects confirm that the contagion effect increases with cultural-spatial proximity and is particularly strong when media consumption is high. Investigating the underlying mechanisms, our results do not seem to be driven by bandwagoning or general time trends. Instead, we find evidence that higher aggregate AfD support reduces the reputational concerns associated with populist voting.

Our findings provide clear and quantitatively relevant evidence that social contagion plays an important role in shaping populist attitudes. Given the reputational benefits from political behavior, the interaction between aggregate and individual-level support can lead to a contagious dynamic of political advocacy for right-wing populist ideas. As shown both theoretically and empirically, the existence of reputational concerns actually reinforces the impact of shocks in perceived aggregate support for such platforms compared to established parties. In fact, because our sample consisted of only very small negative shocks, our results provide some tentative evidence of a reinforcing upward trend of populist support. In this regard, our findings bear some important implications for real-world politics. While we remain silent about the normative assessment of how populist contagion potentially changes the political process, we acknowledge that the authoritarian elements inherent in right-wing populist platforms pose risks to democratic functionality. Our results indicate that small shocks may be sufficient to change the public discussion culture to such an extent that platforms rejecting some basic democratic rights become substantially more acceptable. Policymakers are thus challenged to address the concerns of voters who are willing to turn towards populist movements and integrate them in the democratic process. Furthermore, political education is called to foster the process of individual, independent and critical opinion forming. This is especially true in these times of online information flows which may, as our results show, boost social contagion in political behavior.

Ultimately, the empirical results presented here are related to the social desirability bias in survey research. However, our data does not allow us to draw conclusions regarding changes in actual voting behavior in response to observable shifts in aggregate populist advocacy. Similarly, our quasi-experimental set-up limits conclusions about the persistence of social contagion effects in populist voting. Nevertheless, the results provide meaningful evidence of a contagion effect on individual populist preferences that helps to assess the dynamics of populist rises in other political systems where the lack of a suitable empirical set-up does not allow the development to be quantified.

Declaration of competing interest

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

Data availability

Data will be made available on request.

Acknowledgments

The authors wish to thank two anonymous reviewers, Toke Aidt, Thomas Apolte, Niklas Potrafke, Alois Stutzer, Celeste Brennecka and Julie Davies for excellent comments. The paper benefited from helpful discussions at the 2nd International Conference on The Political Economy of Democracy and Dictatorship in Münster, the 2019 Meeting of the European Public Choice Society in Jerusalem, the 2019 Annual Congress of the European Economic Association in Manchester, the 2019 Annual Meeting of the German Economic Association in Leipzig as well as the 13th CESifo Workshop on Political Economy in Dresden. Lennard Maßmann, Lukas Möller and Camilla Nienhaus provided excellent research assistance.

Appendix A. Variables and summary statistics

Table A.1
Summary statistics.

Variable	Obs	Mean	Std. Dev.	Min	Max
Dependent variable:					
AfD vote intention	20,861	0.058	0.234	0	1
Independent variables:					
Election information shock	20,861	0.498	0.5	0	1
Shock intensity	20,861	3.100	6.106	−2.8	14.3
Gender	20,728	0.473	0.499	0	1
Age	20,728	7.618	2.209	1	10
Age squared	20,728	62.916	28.865	1	100
Married	20,728	0.597	0.490	0	1
Low education	20,728	0.183	0.387	0	1
Medium-level education	20,728	0.361	0.480	0	1
In school	20,728	0.005	0.068	0	1
Part-time employment	20,728	0.122	0.327	0	1
Marginally employed	20,728	0.001	0.039	0	1
Unemployed	20,728	0.02	0.142	0	1
In vocational training	20,728	0.027	0.163	0	1
Retired	20,728	0.345	0.475	0	1
Other employment status	20,728	0.032	0.177	0	1
Voted for AfD in the last federal election?	17,821	0.022	0.148	0	1
Self-positioning on left–right scale	17,821	5.445	1.894	1	11
Satisfaction with current government	17,821	7.052	2.489	1	11
Immigration as most important issue?	17,821	0.536	0.499	0	1
Household income p.c.	17,821	1718.349	166.226	1434.833	2022.146
Share of foreigners	17,821	8.706	4.367	2.2	16.7
Crimes p.c.	17,821	0.078	0.024	0.050	0.162
Unemployment rate	17,821	7.014	2.280	3.2	12.3

Appendix B. Case selection and sample

In this section, we provide further information on perceptions of the AfD that supports our claim that German state elections can be understood as election information shocks. We show that these shocks split the sample of survey respondents evenly into a treatment and a control group and that there are no structural differences in the distribution of personal characteristics of the respondents across these groups.

Fig. B.1 shows how respondents in the German Longitudinal Election Study rated the AfD on a left–right-scale from 1 to 11. Fig. B.2 displays deviations of the realized AfD vote shares from earlier polls for each state election. Fig. B.2(b) uses as poll data the most recent forecast for the specific election in question which is based on survey interviews conducted in the state of the election (these surveys are not part of our sample). Normally, the last poll based on these interviews is published three days before the election. In contrast, Fig. B.2(a) uses as poll data the most recent estimate of the current AfD vote share at the federal level. This estimate is based on the regular *Politbarometer* surveys which we use in our analysis. Fig. B.3 shows relative frequencies of Google searches of the respective state name in the four weeks around an election. Together, these figures show that (i) during our sample period public perceptions of the AfD as a nationalist far-right party clearly increased; (ii) the party has realized vote shares in German state elections that substantially deviated from pre-election polls both at the state and at the federal level; and (iii) these deviations were unanticipated due to the low interest in state elections before the election day.

Table B.1 displays covariate imbalance tests for the full sample. For each covariate, the reported coefficient reports the estimated average difference of this variable in the treatment and control group. Imbalance statistics for each election window and after entropy balancing can be found in Section I of the online supplementary material.

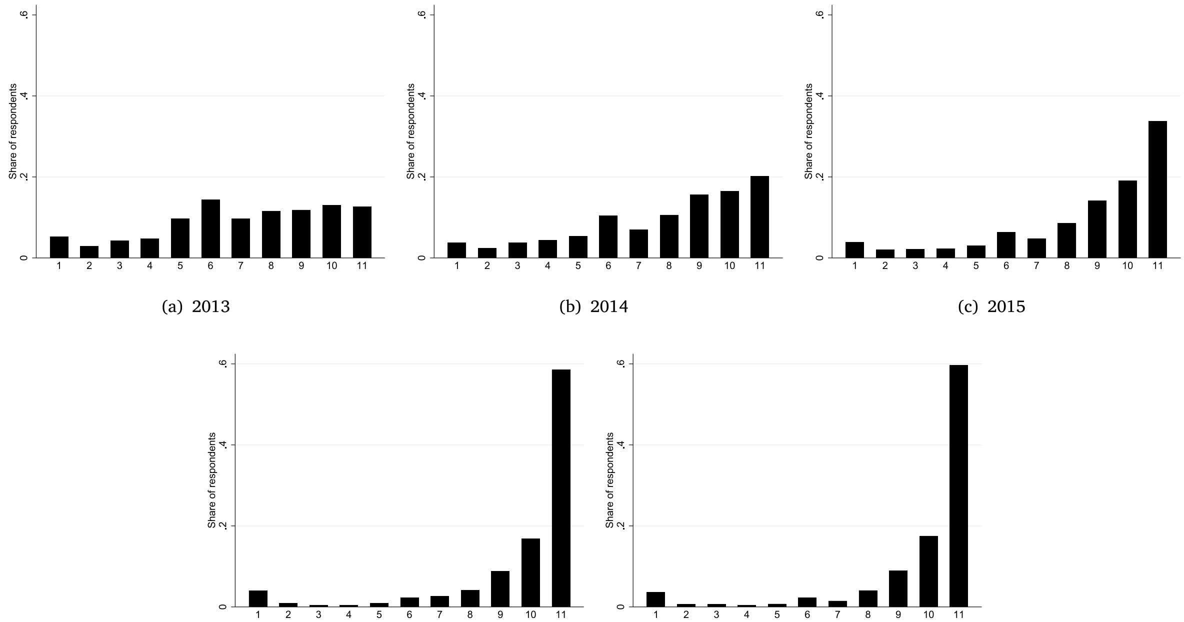

(a) 2013 (b) 2014 (c) 2015

(d) 2016 (e) 2017

Fig. B.1. AfD ratings on a 1- to −11 left–right-scale in the German Longitudinal Election Study.
Note: Figure displays the share of respondents assigning the respective rating to the AfD. Scale runs from 1 = left to 11 = right.

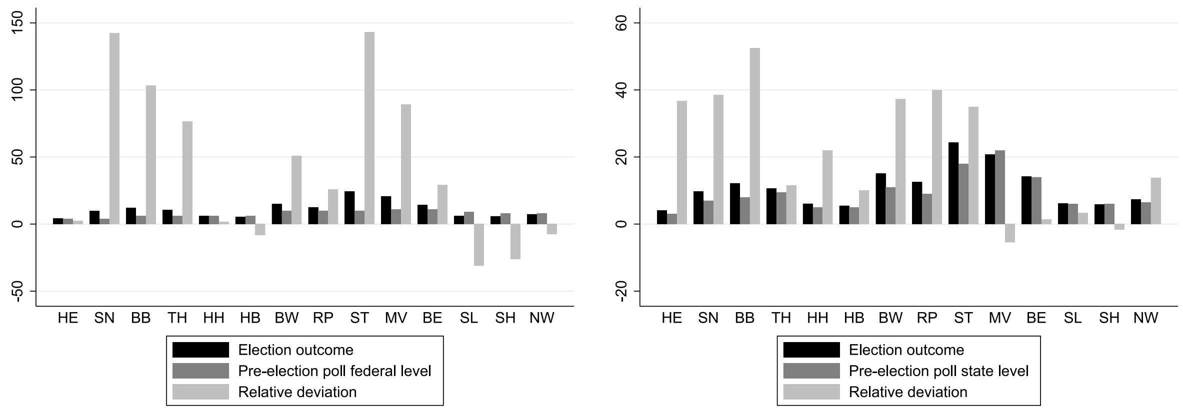

(a) Federal-level polls (b) State-level polls

Fig. B.2. Vote shares for the AfD in German state elections and pre-election polls.
Note: Results for state elections in chronological order from 2013 to 2017. *Relative deviation* measures the percentage deviation of the realized vote share from the pre-election poll. *Pre-election poll federal level* is the current estimated AfD share at the federal level (as published by the most recent general poll before the state election). *Pre-election poll state level* reports the most recent pre-election poll for the respective state election (published ca. three days before the election). HE = Hesse, SN = Saxony, BB = Brandenburg , TH = Thuringia , HH = Hamburg, HB = Bremen, BW = Baden-Wuerttemberg, RP = Rhineland-Palatinate, ST = Saxony-Anhalt, MV = Mecklenburg-Hither Pomerania, BE = Berlin, SL = Saarland, SH = Schleswig-Holstein, NW = North Rhine-Westphalia.

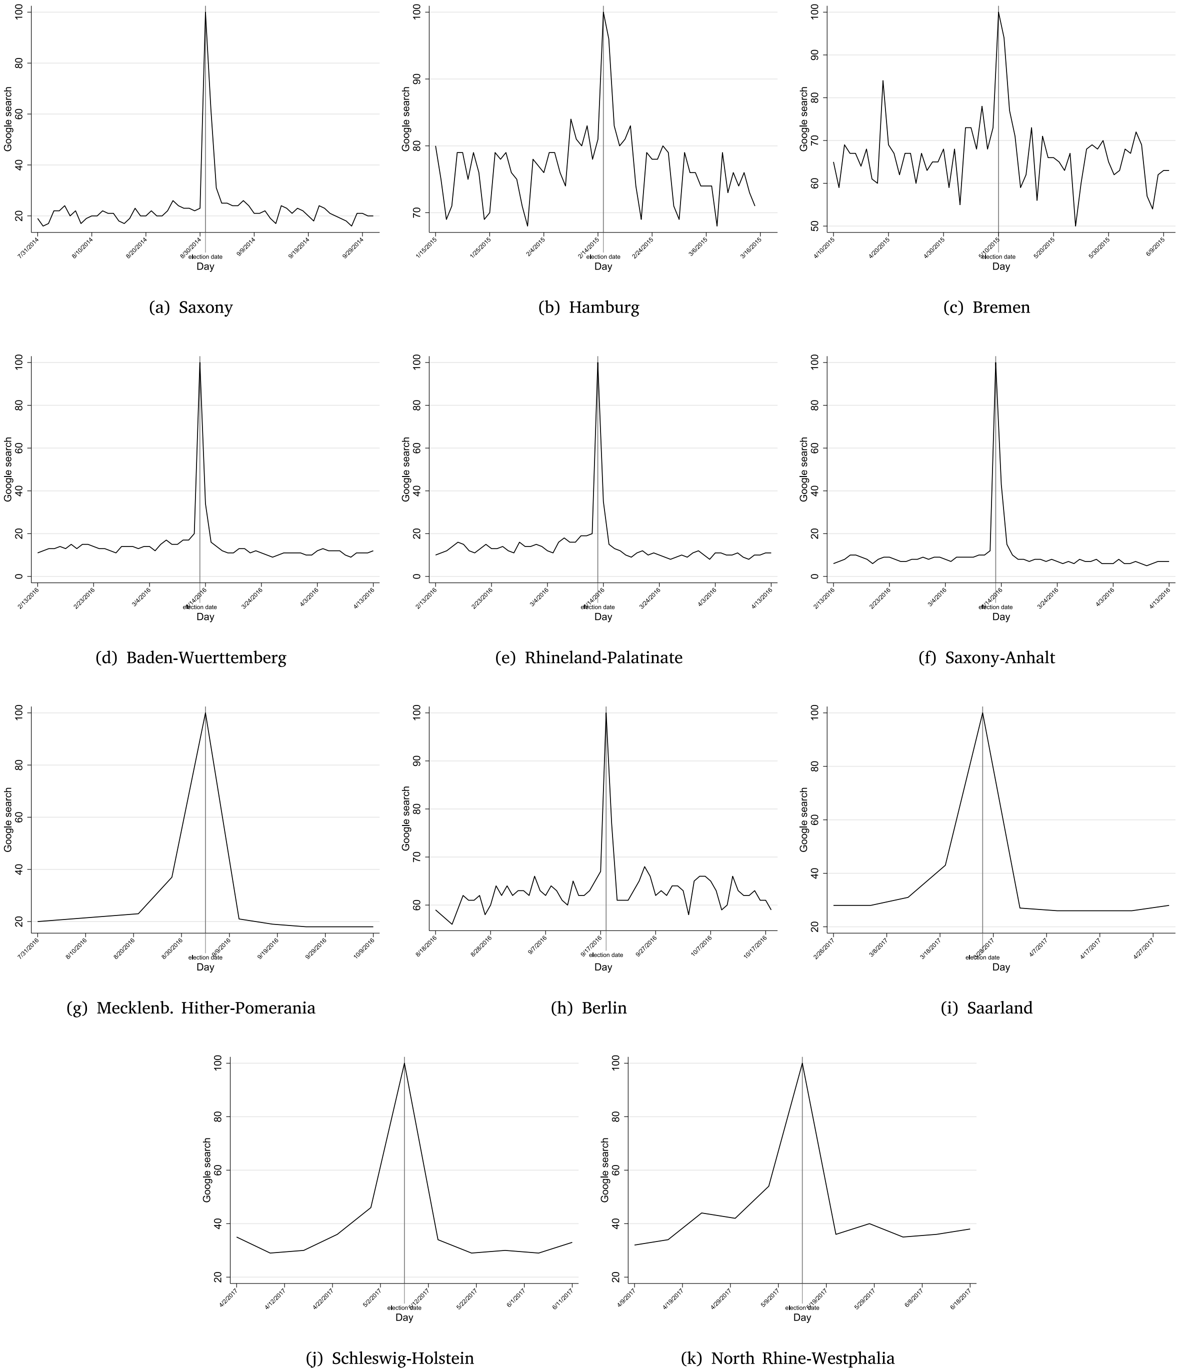

(a) Saxony (b) Hamburg (c) Bremen

(d) Baden-Wuerttemberg (e) Rhineland-Palatinate (f) Saxony-Anhalt

(g) Mecklenb. Hither-Pomerania (h) Berlin (i) Saarland

(j) Schleswig-Holstein (k) North Rhine-Westphalia

Fig. B.3. Relative frequencies of Google searches for a state name around the respective state election.

Table B.1
Balance in covariates between treatment and control groups.

Covariate	Estimates	Observations	R-squared
State of residence	−0.0014 (0.06397)	20,861	0.000
Gender	0.0049 (0.00694)	20,728	0.000
Age	−0.0036 (0.03069)	20,728	0.000
Age squared	−0.0078 (0.40099)	20,728	0.000
Married	−0.0040 (0.00681)	20,728	0.000
Low education	0.0065 (0.00537)	20,728	0.000
Medium-level education	−0.0040 (0.00667)	20,728	0.000
In school	−0.0010 (0.00095)	20,728	0.000
Part-time employed	−0.0010 (0.00454)	20,728	0.000
Marginally employed	0.0007 (0.00054)	20,728	0.000
Unemployed	0.0002 (0.00197)	20,728	0.000
In vocational training	0.0021 (0.00227)	20,728	0.000
Retired	0.0033 (0.00660)	20,728	0.000
Other employment status	0.0024 (0.00246)	20,728	0.000
Voted for AfD in the last federal election?	0.0021 (0.00221)	17,821	0.000
Self-positioning on left–right-scale	0.0070 (0.02838)	17,821	0.000
Satisfaction with current government	0.0356 (0.03729)	17,821	0.000
Immigration as most important issue?	−0.0090 (0.00747)	17,821	0.000

Notes: Coefficients for 18 OLS regressions of a covariate on *post-election*. *post-election* takes the value 1 if the respondent was interviewed in the first survey round after an election (treated), and takes the value 0 if the respondent was interviewed in the last survey round before the election (control). ***p < 0.01, **p < 0.05, *p <0.1.

Appendix C. Additional results

C.1. Additional figures

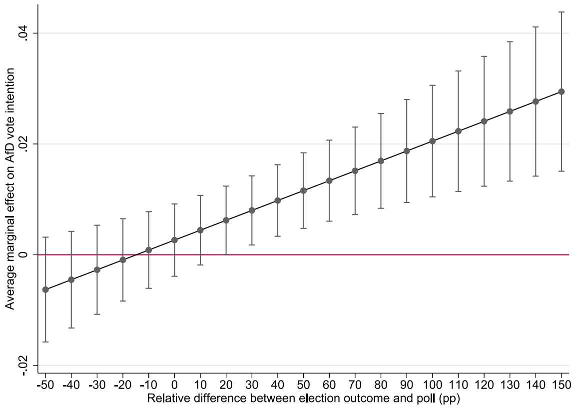

(a) Relative difference from federal poll.

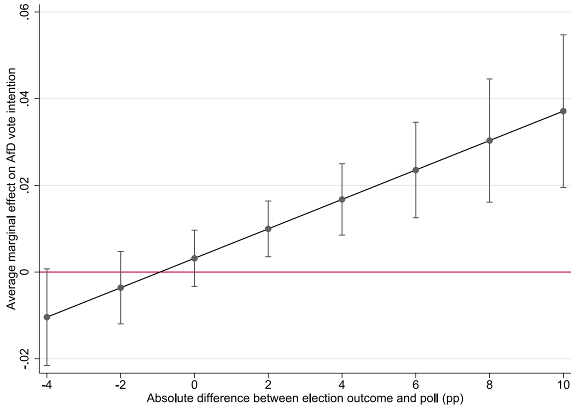

(b) Mean difference from federal poll.

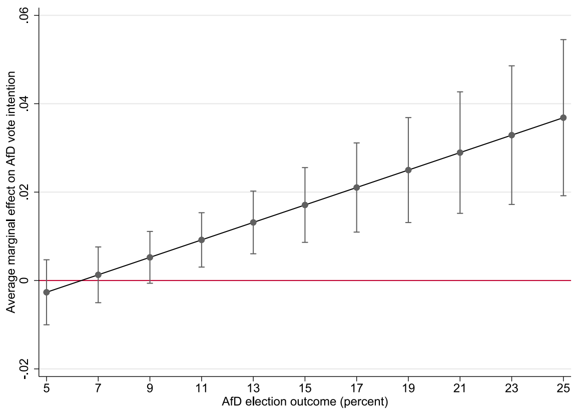

(c) Absolute AfD vote share.

Fig. C.1. Alternative definition of shock intensity.
Note: Graphs show the conditional effects of *post-election* at different values of *shock intensity*. The conditional effects are calculated based on the specifications of columns (2), (4) and (6) in appendix Table C.1. In Fig. C.1(a), *shock intensity* measures the *relative* difference (as a percentage) between the election outcome and the federal-level pre-election poll. For election windows with more than one election, we consider the largest shock in relative terms. In Fig. C.1(b), *shock intensity* measures the mean absolute difference between election outcomes and federal polls for election windows with more than one election. In Fig. C.1(c), *shock intensity* measures the AfD vote share (as a percentage) in the respective election. For election windows with more than one election, we consider the largest vote share that the AfD obtained. All other covariates are held constant at their means. Caps indicate displays 95% confidence intervals. Red horizontal line marks a marginal effect of 0.

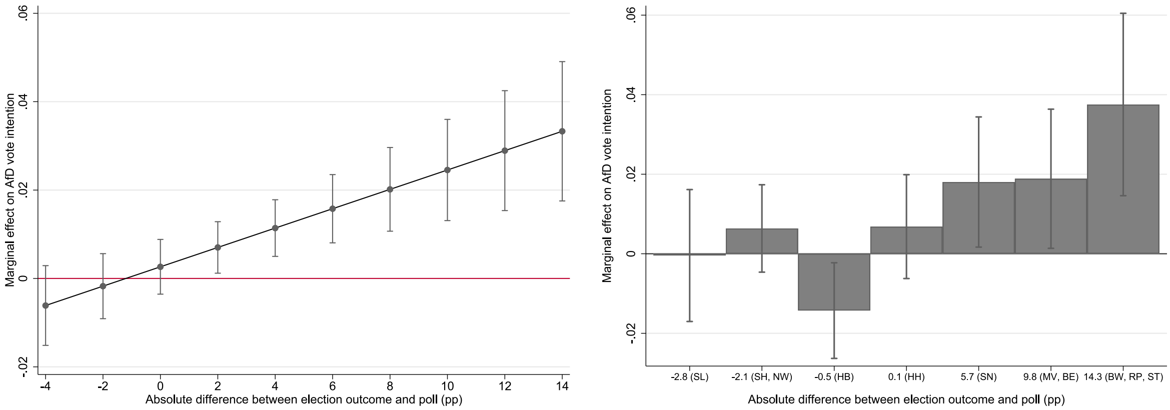

(a) Linear estimation

(b) Flexible estimation

Fig. C.2. Shock intensity and vote intentions: Conditional treatment effects, full sample including all survey respondents.
Note: Fig. C.2(a): Conditional effects of *post-election* at different values of *shock intensity*. The conditional effects are calculated based on the specification of column (1) in appendix Table C.1 but adding respondents from election states. *shock intensity* measures the absolute difference (in percentage points) between the election outcome and the federal-level pre-election poll. For election windows with more than one election, we consider the largest shock in absolute terms. All other covariates are held constant at their means. Caps indicate displays 95% confidence intervals. Red horizontal line marks a marginal effect of 0. Fig. C.2(b): Conditional effects of *post-election* for each election window. The conditional effects are calculated from estimating the specification of column (5) in Table 2 for the pooled sample of all respondents, including a full set of election-window dummies and their interaction with *post-election* (Hamburg is the omitted baseline category). All other covariates are held constant at their means. Caps indicate 95% confidence intervals.

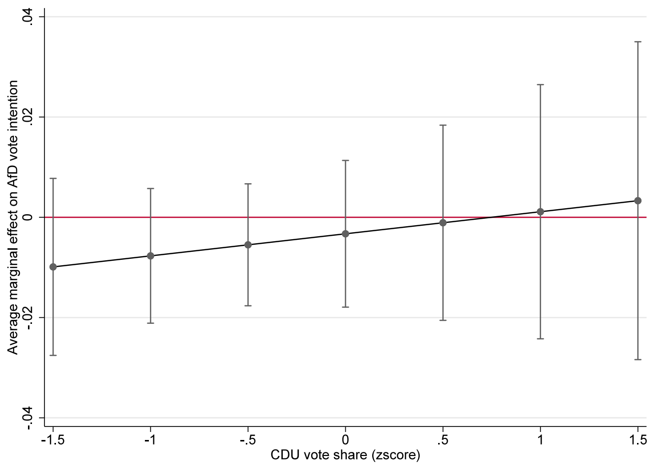

Fig. C.3. Conditional treatment effects at different values of CDU vote share.
Note: Conditional effects of *post-election* at different values of *CDU vote share*. The conditional effects are calculated based on estimating the specifications of columns (3) and (4) in Table 5 for the pooled sample. *CDU vote share* measures the vote share that the CDU obtained in the respective state-level election. The variable is standardized with a mean of 0 and a standard deviation of 1. All other covariates are held constant at their means. Caps indicate 95% confidence intervals. Red horizontal line marks a marginal effect of 0.

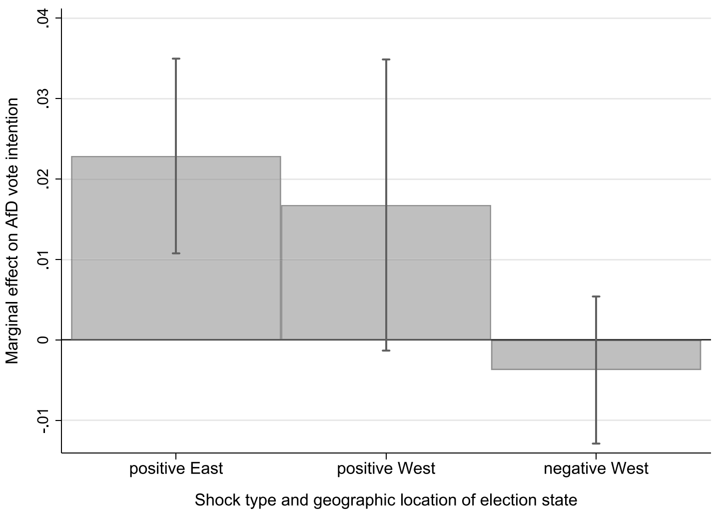

(a) Respondent lives in former western region

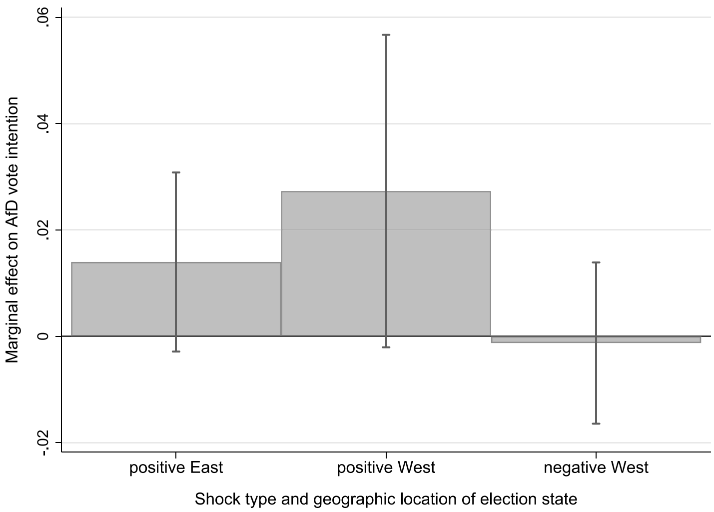

(b) Respondent lives in former eastern region

Fig. C.4. Election information shocks and vote intentions by geographic groups.
Note: Coefficient estimates from linear regressions for different geographic groups. The coefficient estimates are based on separate regressions for (i) positive shocks in former East German states, (ii) positive shocks in former West German states and (iii) negative shocks in former West German states (our sample contains no negative shocks in the eastern part). Figure (a) shows results for the subsample of respondents that live in former West German states. Figure (b) shows results for the subsample of respondents that live in former East German states (including Berlin). Elections in *West* refer to the election windows HH, HB, SL and (SH, NW). Elections in *East* refer to the election windows SN and (MV, BE). Caps indicate displays 95% confidence intervals. Horizontal line marks a marginal effect of 0.

C.2. Additional tables

Table C.1
Alternative definitions of shock intensity: Coefficient estimates.

DV: AfD vote intention	Federal poll				State poll	
	Absolute difference	Relative difference (percent)	Relative difference (std. dev.)	Mean difference	Absolute difference	AfD election outcome
	(1)	(2)	(3)	(4)	(5)	(6)
Post-election	0.0022	0.0026	0.0015	0.0032	0.0044	−0.0125
	[0.00324]	[0.00331]	[0.00331]	[0.00328]	[0.00403]	[0.00576]
Cluster-robust SE, *p*-value	0.499	0.426	0.662	0.334	0.277	0.031**
WCRB by state of residence, *p*-value	0.602	0.563	0.731	0.485	0.386	0.069*
WCRB by survey round, *p*-value	0.679	0.627	0.793	0.575	0.841	0.035**
Post-election × shock intensity	0.0023	0.0002	0.0055	0.0034	0.0038	0.0020
	[0.00064]	[0.00005]	[0.00153]	[0.00093]	[0.00204]	[0.00054]
Cluster-robust SE, *p*-value	0.000***	0.001***	0.000***	0.000***	0.063*	0.000***
WCRB by state of residence, *p*-value	0.010**	0.001***	0.011**	0.002***	0.116	0.017**
WCRB by survey round, *p*-value	0.005***	0.004***	0.007***	0.003***	0.528	0.014**
Adj. R^2	0.231	0.230	0.231	0.230	0.230	0.231
Observations	17,821	17,821	17,821	17,821	17,821	17,821
State-of-residence FE	Y	Y	Y	Y	Y	Y
Election-window FE	Y	Y	Y	Y	Y	Y
Demographics	Y	Y	Y	Y	Y	Y
Socio-economics	Y	Y	Y	Y	Y	Y
Political attitudes	Y	Y	Y	Y	Y	Y
State-level controls	Y	Y	Y	Y	Y	Y
Entropy balancing	Y	Y	Y	Y	Y	Y

Note: Table displays estimates for OLS regressions. Standard errors (SE) clustered by *state of residence* × *survey round* are reported in square brackets. In column (1), *shock intensity* measures the absolute difference between the most recent federal-level poll and the AfD election outcome. In columns (2) and (3), *shock intensity* measures the relative difference between the most recent federal-level poll and the AfD election outcome, expressed in percent and standard deviations of the poll, respectively. In column (4), *shock intensity* is measured analogous to column (1), but taking the mean absolute difference for windows with more than one election. In column (5), *shock intensity* measures the absolute difference between the most recent state-level poll and the AfD election outcome. In column (6), *shock intensity* measures the AfD election outcome in percentage terms. For windows with more than one election, the largest vote share that the AfD received in these elections is included. Control variables as in Table 2. Matching weights from entropy balancing based on the demographic and socio-economic covariates as well as the state of residence. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table C.2
Shock intensity and vote intentions: Marginal effects for flexible estimation.

Election window:	SL	SH, NW	HB	HH	SN	MV, BE	BW, RP, ST
Shock intensity (pp):	−2.8	−2.1	−0.5	0.1	5.7	9.8	14.3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Marginal effect for *Post-election*	−0.0003	0.0057	−0.0144**	0.0072	0.0135*	0.025***	0.0367***
	[0.00812]	[0.00697]	[0.0065]	[0.00647]	[0.00708]	[0.00864]	[0.01275]
Adj. R^2: 0.231							
Observations: 17,821							

Note: Table displays estimates for an OLS regression for the pooled baseline sample of positive and negative shocks, including a full set of election-window dummies and their interaction with *post-election* (Hamburg is the omitted baseline category). In each column, the marginal effect of *post-election* for the respective election window is reported (see Fig. 2(b)). Standard errors clustered by *state of residence* × *survey round* in square brackets. Included controls: state-of-residence FE, election-window FE, demographic controls, socio-economic controls, political attitudes, state-level controls. Matching weights from entropy balancing have been used. *shock intensity* reports the absolute difference (in percentage points) between the election outcome and the federal poll. ***p < 0.01, **p <0.05, *p <0.1.

Table C.3
Results for full sample including all survey respondents.

DV: AfD vote intention	Placebo treatment I		Placebo treatment II		Sample restriction: *Preferred party = CDU*	
	Positive shocks	Negative shocks	Positive shocks	Negative shocks	Positive shocks	Negative shocks
	(1)	(2)	(3)	(4)	(5)	(6)
Counterfactual post-election	−0.0037	−0.0068	−0.0020	−0.0001		
	[0.00330]	[0.00410]	[0.00452]	[0.00340]		
Cluster-robust SE, p-value	0.262	0.100	0.660	0.980		
WCRB by state of residence, p-value	0.278	0.221	0.659	0.984		
WCRB by survey round, p-value	0.421	0.195	0.897	0.923		
Post-election					0.0209	−0.0048
					[0.00587]	[0.00477]
Cluster-robust SE, p-value					0.001***	0.314
WCRB by state of residence, p-value					0.003***	0.451
WCRB by survey round, p-value					0.005***	0.565
Post-election × CDU vote share					0.0105	0.0054
					[0.00535]	[0.00670]
Cluster-robust SE, p-value					0.051*	0.425
WCRB by state of residence, p-value					0.158	0.509
WCRB by survey round, p-value					0.025**	0.191
Adj. R^2	0.252	0.207	0.243	0.205	0.141	0.069
Observations	11,800	8,464	11,562	8,445	4,426	2,910
State-of-residence FE	Y	Y	Y	Y	Y	Y
Election-window FE	Y	Y	Y	Y	Y	Y
Demographics	Y	Y	Y	Y	Y	Y
Socio-economics	Y	Y	Y	Y	Y	Y
Political attitudes	Y	Y	Y	Y	Y	Y
State-level controls	Y	Y	Y	Y	Y	Y
Entropy balancing	Y	Y	Y	Y	Y	Y

Note: Table displays estimates for OLS regressions. Sample includes all survey respondents, both in election and non-election states. Standard errors (SE) clustered by *state of residence* × *survey round* are reported in square brackets. For *Placebo treatment I*, we use the two most recent survey rounds *before* the actual election. For *Placebo treatment II*, we use the two most recent survey rounds *after* the actual election. *Counterfactual* election windows for sample with positive shocks: SN, HH, (BW, RP, ST), (MV, BE). *Counterfactual* election windows for sample with negative shocks: HB, (SH, NW), SL. The regressions in columns (5) and (6) are restricted to respondents who rank the CDU/CSU first among all parties. *CDU vote share* measures the vote share that the party obtained in the respective state-level election. The variable is standardized with a mean of 0 and a standard deviation of 1. Control variables as in Table 2. State-level control variables have been adapted to the corresponding time periods when necessary. Matching weights from entropy balancing are based on the demographic and socio-economic covariates as well as the state of residence. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Table C.4
Additional results for *party affinity*.

DV: Party affinity	Baseline sample	
	Positive shocks (1)	Negative shocks (2)
Post-election	0.0115	−0.0011
	[0.00196]	[0.00328]
Cluster-robust SE, p-value	0.000***	0.743
WCRB by state of residence, p-value	0.002***	0.723
WCRB by survey round, p-value	0.110	0.761
Adj. R^2	0.098	0.108
Observations	4,987	4,926
State-of-residence FE	Y	Y
Election-window FE	Y	Y
Demographics	Y	Y
Socio-economics	Y	Y
Political attitudes	Y	Y
State-level controls	Y	Y
Entropy balancing	Y	Y

Note: Table displays estimates for OLS regressions. Standard errors (SE) clustered by *state of residence* × *survey round* are reported in square brackets. *party affinity*= 1 if the respondents states that, of all the parties, they feel most closely connected to the AfD; 0 if the respondent states a different party affinity; missing if the respondent states no party affinity. The underlying item in the *Politbarometer* lists AfD as a distinct only after 2015; therefore, the regressions only include elections in 2016 and 2017. Election windows for the samples of positive shocks: SN, HH, (BW, RP, ST), (MV, BE). Election windows for the samples of negative shocks: HB, (SH, NW), SL. We include only those respondents living in states other than the respective election state(s) for each election window. Control variables as in Table 2. Matching weights from entropy balancing based on the demographic and socio-economic covariates as well as the state of residence. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C.5
Heterogeneous effects of election information shocks.

Interacted variable (X =)	Local support base	AfD largest party	Post programmatic shift		Media consumption	Distance	
DV: AfD vote intention	Pooled sample (1)	Pooled sample (2)	Positive shocks (3)	Negative shocks (4)	Pooled sample (5)	Positive shocks (6)	Negative shocks (7)
Post-election	−0.0023	0.0013	0.0096	−0.0135	−0.0050	0.0247	−0.0085
	[0.00375]	[0.00327]	[0.00464]	[0.00545]	[0.00713]	[0.00993]	[0.00828]
Cluster-robust SE, p-value	0.540	0.683	0.041**	0.015**	0.487	0.014**	0.306
WCRB by state of residence, p-value	0.616	0.730	0.086*	0.043**	0.608	0.053*	0.504
WCRB by survey round, p-value	0.765	0.835	0.066*	0.336	0.860	0.065*	0.241
Post-election × X	0.0019	0.0293	0.0215	0.0158	0.00821	−0.00001	0.00002
	[0.00053]	[0.00833]	[0.00845]	[0.00746]	[0.00424]	[0.00003]	[0.00003]
Cluster-robust SE, p-value	0.000***	0.001***	0.012**	0.037**	0.054*	0.592	0.547
WCRB by state of residence, p-value	0.016**	0.023**	0.182	0.019**	0.240	0.652	0.687
WCRB by survey round, p-value	0.021**	0.027**	0.005***	0.109	0.395	0.644	0.019**
Adj. R^2	0.231	0.231	0.254	0.198	0.230	0.253	0.197
Observations	17,821	17,821	9,909	7,912	17,821	9,909	7,912
State-of-residence FE	Y	Y	Y	Y	Y	Y	Y
Election-window FE	Y	Y	Y	Y	Y	Y	Y
Demographics	Y	Y	Y	Y	Y	Y	Y
Socio-economics	Y	Y	Y	Y	Y	Y	Y
Political attitudes	Y	Y	Y	Y	Y	Y	Y
State-level controls	Y	Y	Y	Y	Y	Y	Y
Entropy balancing	Y	Y	Y	Y	Y	Y	Y

Note: Table displays estimates for OLS regressions. Standard errors (SE) clustered by *state of residence* × *survey round* are reported in square brackets. *local support base* measures the difference (in percentage points) between the AfD vote share in state election e and the vote share that the AfD obtained in the state of residence of respondent i in the 2013 federal election. For state election windows with more than one election, we consider the largest vote share. *AfD largest party* is 1 is the AfD turned out to be one of the three largest parties in the respective state election; 0 otherwise. For election windows with more than one election, the variable is 1 if the AfD is among the three largest parties in at least one election included in that window. *Post-shift* is 1 for all elections held after the AfD's program shift in summer 2015 and 0 for elections held prior to these events. *Distance* measures the shortest distance between the capitals of the election state and the state of residence of the respondent. For election windows with more than one election, we consider the distance to the closest capital of the respondent's state of residence. Control variables as in Table 2. Matching weights from entropy balancing based on the demographic and socio-economic covariates as well as the state of residence for each sample restriction separately. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C.6
Mechanisms: Robustness.

DV: AfD vote intention	No sample restriction		Sample restriction: *AfD vote in 2013*= 0		Sample restriction: *Preferred party*= *CDU*	
	Positive shocks (1)	Negative shocks (2)	Positive shocks (3)	Negative shocks (4)	Positive shocks (5)	Negative shocks (6)
Post-election	0.0122	−0.0026	0.0280	0.0033	0.0169	−0.0047
	[0.00386]	[0.00308]	[0.00717]	[0.00420]	[0.00627]	[0.00561]
Cluster-robust SE, *p*-value	0.002***	0.396	0.000***	0.436	0.008***	0.401
WCRB by state of residence, *p*-value	0.005***	0.470	0.008***	0.616	0.013**	0.561
WCRB by survey round, *p*-value	0.042**	0.569	0.139	0.145	0.039**	0.517
Post-election × AfD preferred party	−0.0067	0.0444	−0.0175	0.1178		
	[0.04025]	[0.04477]	[0.10394]	[0.07593]		
Cluster-robust SE, *p*-value	0.868	0.324	0.867	0.126		
WCRB by state of residence, *p*-value	0.882	0.247	0.781	0.040**		
WCRB by survey round, *p*-value	0.871	0.556	0.776	0.055*		
Post-election × federal CDU vote share					−0.0026	−0.0043
					[0.00627]	[0.00557]
Cluster-robust SE, *p*-value					0.680	0.445
WCRB by state of residence, *p*-value					0.795	0.532
WCRB by survey round, *p*-value					0.674	0.167
Adj. R^2	0.469	0.471	0.199	0.276	0.149	0.073
Observations	9,019	7,278	4,902	4,839	3,809	2,734
State-of-residence FE	Y	Y	Y	Y	Y	Y
Election-window FE	Y	Y	Y	Y	Y	Y
Demographics	Y	Y	Y	Y	Y	Y
Socio-economics	Y	Y	Y	Y	Y	Y
Political attitudes	Y	Y	Y	Y	Y	Y
State-level controls	Y	Y	Y	Y	Y	Y
Entropy balancing	Y	Y	Y	Y	Y	Y

Note: Table displays estimates for OLS regressions. Standard errors (SE) clustered by *state of residence* × *survey round* are reported in square brackets. In columns (1) and (2), we do not restrict the subsample on the reported vote decision in the 2013 federal election. In columns (3) and (4), support for the AfD is measured using the variable *party affinity* from the *Politbarometer*. We construct two indicator variables that equal 1 if the respondent states that, of all the parties, they feel most closely connected to the AfD; 0 if the respondent states a different party affinity; missing if the respondent states no party affinity. In columns (5) and (6), *federal CDU vote share* measures the vote share that the CDU/CSU obtains at the *federal* level according to the most recent federal-level poll (*Sonntagsfrage*). The variable is standardized with a mean of 0 and a standard deviation of 1. Control variables as in Table 2. Matching weights from entropy balancing are based on the demographic and socio-economic covariates as well as the state of residence for each sample restriction. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Appendix D. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.ejpoleco.2021.102098.

References

Acharya, A., Blackwell, M., Sen, M., 2016. Explaining causal findings without bias: Detecting and assessing direct effects. Am. Political Sci. Rev. 110 (3), 512–529. https://doi.org/10.1017/s0003055416000216.

Aidt, T., Leon, G., Satchell, M., 2021. The social dynamics of collective action: Evidence from the diffusion of the swing riots, 1830-31. J. Politics https://doi.org/10.1086/714784, (forthcoming).

Algan, Y., Guriev, S., Papaioannou, E., Passari, E., 2017. The European trust crisis and the rise of populism. Brook. Papers Econ. Activity 2017 (2), 309–400. https://doi.org/10.1353/eca.2017.0015.

Anderson, C.J., Ward, D., 1996. Barometer elections in comparative perspective. Elect. Stud. 15 (4), 447–460. https://doi.org/10.1016/0261-3794(95)00056-9.

Arbatli, C.E., Gomtsyan, D., 2019. Voting retrospectively: Critical junctures and party identification. Eur. Econ. Rev. 119, 356–390. https://doi.org/10.1016/j.euroecorev.2019.08.005.

Arzheimer, K., 2015. The AfD: Finally a successful right-wing populist eurosceptic party for Germany?. West Eur. Politics 38 (3), 535–556. https://doi.org/10.1080/01402382.2015.1004230.

Bakker, R., Hooghe, L., Jolly, S., Marks, G., Polk, J., Rovny, J., Steenbergen, M., Vachudova, M.A., 2019. 2019 Chapel Hill Expert Survey. Version 2019.1.. Available at: chesdata.eu.

Barbera, S., Jackson, M.O., 2020. A model of protests, revolution, and information. Q. J. Political Sci. 15 (3), 297–335. https://doi.org/10.1561/100.00019002.

Barnfield, M., 2020. Think twice before jumping on the Bandwagon: Clarifying concepts in research on the Bandwagon effect. Political Stud. Rev. 18 (4), 553–574. https://doi.org/10.1177/1478929919870691.

Baudains, P., Johnson, S.D., Braithwaite, A.M., 2013. Geographic patterns of diffusion in the 2011 London riots. Appl. Geogr. 45, 211–219. https://doi.org/10.1016/j.apgeog.2013.09.010.

Benesch, C., Loretz, S., Stadelmann, D., Thomas, T., 2019. Media coverage and immigration worries: Econometric evidence. J. Econ. Behav. Organ. 160, 52–67. https://doi.org/10.1016/j.jebo.2019.02.011.

Berbuir, N., Lewandowsky, M., Siri, J., 2015. The AfD and its sympathisers: Finally a right-wing populist movement in Germany?. Ger. Politics 24 (2), 154–178. https://doi.org/10.1080/09644008.2014.982546.

Bergmann, K., Diermeier, M., 2017. Die AfD: Eine unterschätzte Partei. Soziale Erwünschtheit als Erklärung für fehlerhafte Prognosen. IW-Reports 7/2017.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. J. Political Econ. 100 (5), 992–1026. https://doi.org/10.1086/261849.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1998. Learning from the behavior of others: Conformity, fads, and informational cascades. J. Econ. Perspect. 12 (3), 151–170. https://doi.org/10.1257/jep.12.3.151.

Boomgarden, H.G., Vliegenthart, R., 2009. How news content influences anti-immigration attitudes: Germany, 1993-2005. Eur. J. Political Res. 48, 516–542. https://doi.org/10.1111/j.1475-6765.2009.01831.x.

Bos, L., van der Brug, W., de Vreese, C., 2011. How the media shape perceptions of right-wing populist leaders. Political Commun. 28 (2), 182–206. https://doi.org/10.1080/10584609.2011.564605.

Brown, R., Montalva, V., Thomas, D., Velásquez, A., 2019. Impact of violent crime on risk aversion: Evidence from the Mexican drug war. Rev. Econ. Stat. 101 (5), 892–904. https://doi.org/10.1162/rest_a_00788.

Bursztyn, L., Egorov, G., Fiorin, S., 2020. From extreme to mainstream: The erosion of social norms. Amer. Econ. Rev. 110 (11), 3522–3548. https://doi.org/10.1257/aer.20171175.

Cameron, A.C., Gelbach, J.B., Miller, D.L., 2008. Bootstrap-based improvements for inference with clustered errors. Rev. Econ. Stat. 90 (3), 414–427. https://doi.org/10.1162/rest.90.3.414.

Cameron, A.C., Miller, D.M., 2015. A practitioner's guide to cluster-robust inference. J. Hum. Resour. 50 (2), 317–372. https://doi.org/10.3368/jhr.50.2.317.

Cantoni, D., Hagemeister, F., Westcott, M., 2019. Persistence and activation of right-wing political ideology. Rationality Compet. Discuss. Pap. 143.

Cao, H.H., Han, B., Hirshleifer, D., 2011. Taking the road less traveled by: Does conversation eradicate pernicious cascades?. J. Econ. Theory 146, 1418–1436. https://doi.org/10.1016/j.jet.2011.03.010.

Caprettini, B., Voth, H.-J., 2020. Rage against the machines: Labor-saving technology and unrest in industrializing England. Am. Econ. Rev. Insights 2 (3), 305–320. https://doi.org/10.1257/aeri.20190385.

Carpena, F., Zia, B., 2020. The causal mechanism of financial education: Evidence from mediation analysis. J. Econ. Behav. Organ. 177, 143–184. https://doi.org/10.1016/j.jebo.2020.05.001.

Castanheira, M., 2003. Why vote for losers? J. Eur. Econom. Assoc. 1 (5), 1207–1238. https://doi.org/10.1162/154247603770383433.

Coate, S., Conlin, M., 2004. A group rule-utilitarian approach to voter turnout: Theory and evidence. Amer. Econ. Rev. 94 (5), 1476–1504. https://doi.org/10.1257/0002828043052231.

Dahlgaard, J.O., Hansen, J.H., Hansen, K.M., Larsen, M.V., 2016. How are voters influenced by opinion polls? The effect of polls on voting behavior and party sympathy. World Political Sci. 12 (2), 283–300. https://doi.org/10.1515/wps-2016-0012.

DellaVigna, S., Kaplan, E., 2007. The Fox News effect: Media bias and voting. Q. J. Econ. 122 (3), 1187–1234. https://doi.org/10.1162/qjec.122.3.1187.

DellaVigna, S., List, J.A., Malmendier, U., 2012. Testing for altruism and social pressure in charitable giving. Q. J. Econ. 127 (1), 1–56. https://doi.org/10.1093/qje/qjr050.

Depetris-Chauvin, E., Durante, R., Campante, F., 2020. Building nations through shared experiences: Evidence from African football. Amer. Econ. Rev. 110 (5), 1572–1602. https://doi.org/10.1257/aer.20180805.

Depetris-Chauvin, E., Özak, Ö., 2020. The origins of the division of labor in pre-industrial times. J. Econ. Growth 25 (3), 297–340. https://doi.org/10.1007/s10887-020-09179-2.

Dewenter, R., Linder, M., Thomas, T., 2019. Can media drive the electorate? The impact of media coverage on voting intentions. Eur. J. Political Econ. 58, 245–261. https://doi.org/10.1016/j.ejpoleco.2018.12.003.

Dostal, J.M., 2017. The German federal election of 2017: How the wedge issue of refugees and migration took the shine off Chancellor Merkel and transformed the party system. Political Q. 88 (4), 589–602. https://doi.org/10.1111/1467-923x.12445.

Downs, A., 1957. An economic theory of democracy. Harper & Brothers.

Durante, R., Knight, B., 2012. Partisan control, media bias, and viewer responses: Evidence from Berlusconi's Italy. J. Eur. Econom. Assoc. 10 (3), 451–481. https://doi.org/10.1111/j.1542-4774.2011.01060.x.

Durante, R., Pinotti, P., Tesei, A., 2019. The political legacy of entertainment TV. Amer. Econ. Rev. 109 (7), 2497–2530. https://doi.org/10.1257/aer.20150958.

Faas, T., Mackenrodt, C., Schmitt-Beck, R., 2008. Polls that mattered: Effects of media polls on voters' coalition expectations and party preferences in the 2005 German parliamentary election. Int. J. Public Opin. Res. 20 (3), 299–325. https://doi.org/10.1093/ijpor/edn034.

Feddersen, T., Pesendorfer, W., 1997. Voting behavior and information aggregation in elections with private information. Econometrica 65 (5), 1029. https://doi.org/10.2307/2171878.

Forschungsgruppe Wahlen, 2019. Politbarometer 1977-2017 (partial cumulation). ZA2391. Version 10.0.0. GESIS Data Archive Cologne. https://doi.org/10.4232/1.13243.

Franklin, M., Niemi, R., Whitten, G., 1994. The two faces of tactical voting. Br. J. Political Sci. 24 (4), 549–557. https://doi.org/10.1017/S0007123400007006.

Friess, H.-J., Neu, V., 2018. Die Heterogenität des Protests. Eine Untersuchung von Einstellungen potenzieller Protestwähler und politischem Protest im Internet. Available at: https://www.kas.de/documents/252038/253252/7_dokument_dok_pdf_53242_1.pdf/a75505bc-46d4-ce63-e141-2f1d01ec788b?version=1.0&t=1539647113548.

Gehring, K., Schneider, S.A., 2018. Towards the greater good? EU commissioners' nationality and budget allocation in the European Union. Am. Econ. J. 10 (1), 214–239.

Giani, M., Méon, P.-G., 2021. Global racist contagion following Donald Trump's election. Br. J. Political Sci. 51 (3), 1332–1339. https://doi.org/10.1017/S0007123419000449.

Giuliano, P., Wacziarg, R., 2020. Who voted for Trump? Populism and social capital. In: NBER Working Papers, 27651, https://doi.org/10.3386/w27651.

Glaser, W.A., 1959. The family and voting turnout. Public Opin. Q. 23 (4), 563–570.

Granovetter, M., 1978. Threshold models of collective behaviour. Am. J. Sociol. 83 (6), 1420–1443. https://doi.org/10.1086/226707.

Guriev, S., 2018. Economic drivers of populism. AEA Pap. Proc. 108, 200–203. https://doi.org/10.1257/pandp.20181123.

Guriev, S., Papaioannou, E., 2020. The political economy of populism. J. Econ. Lit. (forthcoming).

Hainmueller, J., 2012. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. Political Anal. 20, 25–46. https://doi.org/10.1093/pan/mpr025.

Hainmueller, J., Hangartner, D., 2013. Who gets a Swiss passport? A natural experiment in immigrant discrimination. Am. Political Sci. Rev. 107 (1), 159–187. https://doi.org/10.1017/s0003055412000494.

Huang, J., Low, C., 2017. Trumping norms: Lab evidence on aggressive communication before and after the 2016 US presidential election. Amer. Econ. Rev. 107 (5), 120–124. https://doi.org/10.1257/aer.p20171016.

Huckfeld, R., Sprague, J., 1991. Discussant effects on vote choice: Intimacy, structure, and independence. J. Politics 53 (1), 122–158. https://doi.org/10.2307/2131724.

Imai, K., Keele, L., Yamamoto, T., 2010. Identification, inference and sensitivity analysis for causal mediation effects. Statist. Sci. 25 (1), 51–71. https://doi.org/10.1214/10-sts321.

Jeffery, C., Hough, D., 2001. The electoral cycle and multi-level voting in Germany. Ger. Politics 10 (2), 73–97. https://doi.org/10.1080/772713264.

Jeffery, C., Hough, D., 2003. Regional elections in multi-level systems. Eur. Urban Reg. Stud. 10 (2), 199–212. https://doi.org/10.1080/772713264.

Khalil, U., Mookerjee, S., Tierney, R., 2019. Social interactions in voting behavior: Evidence from India. J. Econ. Behav. Organ. 163, 158–171. https://doi.org/10.1016/j.jebo.2019.04.024.

Klor, E.F., Winter, E., 2018. On public opinion polls and voters' turnout. J. Public Econ. Theory 20, 239–256. https://doi.org/10.2139/ssrn.895946.

Kuran, T., 1987. Chameleon voters and public choice. Public Choice 53, 53–78. https://doi.org/10.1007/BF00115654.

Kuran, T., 1989. Sparks and prairie fires: A theory of unanticipated political revoultion. Public Choice 61, 41–74. https://doi.org/10.1007/BF00116762.

Lubbers, M., Gijsberts, M., Scheepers, P., 2002. Extreme right-wing voting in Western Europe. Eur. J. Political Res. 41 (3), 345–378. https://doi.org/10.1111/1475-6765.00015.

MacKinnon, J.G., 2019. How cluster-robust inference is changing applied econometrics. Can. J. Econ. 52 (3), 1–31. https://doi.org/10.1111/caje.12388.

MacKinnon, J.G., Webb, M.D., 2017. Wild bootstrap inference for wildly different cluster sizes. J. Appl. Econometrics 32, 233–254. https://doi.org/10.1002/jae.2508.

MacKinnon, J.G., Webb, M.D., 2018. The wild bootstrap for few (treated) clusters. Econom. J. 21, 114–135. https://doi.org/10.1111/ectj.12107.

Manski, C.F., 1993. Identification of endogenous social effects: The reflection problem. Rev. Econom. Stud. 60 (3), 531–542. https://doi.org/10.2307/2298123.

Manski, C.F., 2000. Economic analysis of social interaction. J. Econ. Perspect. 14 (3), 115–136. https://doi.org/10.1257/jep.14.3.115.

Mikulaschek, C., Pant, S., Tesfaye, B., 2020. Winning hearts and minds in civil wars: Governance, leadership change, and support for violent groups in Iraq. Am. J. Political Sci. 64 (4), 773–790. https://doi.org/10.1111/ajps.12527.

Mitra, A., Mitra, S., Mukherji, A., 2017. Cash for votes: Evidence from India. In: University of Kent, School of Economics Discussion Papers, 17/11.

Morton, R.B., Muller, D., Page, L., Torgler, B., 2015. Exit polls, turnout, and bandwagon voting: Evidence from a natural experiment. Eur. Econ. Rev. 77, 65–81. https://doi.org/10.1016/j.euroecorev.2015.03.012.

Moya, A., Carter, M.R., 2019. Violence and the formation of hopelessness: Evidence from internally displaced persons in Colombia. World Dev. 113, 100–115. https://doi.org/10.1016/j.worlddev.2018.08.015.

Muñoz, J., Falcó-Gimeno, A., Hernández, E., 2019. Unexpected event during survey design: Promise and pitfalls for causal inference. Political Anal. 28 (2), 186–206. https://doi.org/10.1017/pan.2019.27.

Mudde, C., 2004. The populist Zeitgeist. Gov. Oppos. 39, 541–563. https://doi.org/10.1111/j.1477-7053.2004.00135.x.

Müller, K., Schwarz, C., 2020. From hashtag to hate crime: Twitter and anti-minority sentiment. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3149103.

Nickerson, D.W., 2008. Is voting contagious? Evidence from two field experiments. Am. Political Sci. Rev. 102 (1), 49–56. https://doi.org/10.1017/S0003055408080039.

Niedermayer, O., Hofrichter, J., 2016. DIe Wählerschaft der AfD: Wer ist sie, woher kommt sie und wie weit rechts steht sie?. Zeitschrift für Parlamentsfragen 47 (2), 267–285. https://doi.org/10.5771/0340-1758-2016-2-267.

Noury, A., Roland, G., 2020. Identity politics and populism in Europe. Annu. Rev. Political Sci. 23 (1), 421–439. https://doi.org/10.1146/annurev-polisci-050718-033542.

Piketty, T., 2000. Voting as communicating. Rev. Econom. Stud. 67 (1), 169–191. https://doi.org/10.1111/1467-937x.00126.

Riker, W.H., Ordeshook, P.C., 1968. A theory of the calculus of voting. Am. Political Sci. Rev. 62 (1), 25–42. https://doi.org/10.2307/1953324.

Rodrik, D., 2018. Populism and the economics of globalization. J. Int. Business Policy 1, 12–33. https://doi.org/10.1057/s42214-018-0001-4.

Roodman, D., MacKinnon, J.G., Nielsen, M.Ø., Webb, M.D., 2019. Fast and wild: Bootstrap inference in Stata using boottest. Stata J. 19 (1), 4–60. https://doi.org/10.1177/1536867X19830877.

Rydgren, J., 2005. Is extreme right-wing populism contagious? Explaining the emergence of a new party family. Eur. J. Political Res. 44, 413–437. https://doi.org/10.1111/j.1475-6765.2005.00233.x.

Salzborn, S., 2016. Renaissance of the New Right in Germany? A discussion of New Right elements in German right-wing extremism today. Ger. Politics Soc. 34 (2), https://doi.org/10.3167/gps.2016.340203.

Schaub, M., Morisi, D., 2020. Voter mobilisation in the echo chamber: Broadband internet and the rise of populism in Europe. Eur. J. Political Res. 59 (4), 752–773. https://doi.org/10.1111/1475-6765.12373.

Schmitt-Beck, R., 2016. The 'Alternative für Deutschland in the electorate': Between single-issue and right-wing populist party. Ger. Politics 26 (1), 124–148. https://doi.org/10.1080/09644008.2016.1184650.

Straits, B.C., 1990. The social context of voter turnout. Public Opin. Q. 54 (1), 64–73. https://doi.org/10.1086/269184.

van der Brug, W., Fennema, M., Tillie, J., 2000. Anti-immigrant parties in Europe: Ideological or protest vote? Eur. J. Political Res. 37 (1), 77–102. https://doi.org/10.1111/1475-6765.00505.

van der Meer, T.W.G., Hakhverdian, A., Aaldering, L., 2016. Off the fence, onto the bandwagon? A large-scale survey experiment on effect of real-life poll outcomes on subsequent vote intentions. Int. J. Public Opin. Res. 28 (1), 46–72. https://doi.org/10.1093/ijpor/edu041.

Volkens, A., Burst, T., Krause, W., Lehmann, P., Matthieß, T., Merz, N., Regel, S., Weßels, B., Zehnter, L., 2020. The Manifesto Data Collection. Manifesto Project (MRG/CMP/MARPOR). Version 2020a.. http://dx.doi.org/10.25522/manifesto.mpds.2020a.

Webb, M.D., 2013. Reworking wild bootstrap based inference for clustered errors.. Available at: https://www.econstor.eu/handle/10419/97480.

Wurthmann, L.C., Marschall, S., Triga, V., Manavopoulos, V., 2020. Many losers – one winner? An examination of vote switching to the AfD in the 2017 German federal election using VAA data. Party Politics https://doi.org/10.1177/1354068820914959, (forthcoming).

Zeitzoff, T., 2017. How social media is changing conflict. J. Confl. Resolut. 61 (9), 1970–1991. https://doi.org/10.1177/0022002717721392.