



Individual vulnerability to industrial robot adoption increases support for the radical right

Massimo Anelli^{a,b,c,d,1} , Italo Colantone^{a,c,d,1}, and Piero Stanig^{a,b,1,2} 

^aDepartment of Social and Political Sciences, Bocconi University, Milan 20136, Italy; ^bDondena Research Center, Bocconi University, Milan 20136, Italy; ^cBaffi-Carefin Research Center, Bocconi University, Milan 20136, Italy; and ^dCESifo (Center of Economic Studies and ifo Institute), 81679 Munich, Germany

Edited by Robert Franzese, University of Michigan, Ann Arbor, MI, and accepted by the Editorial Board September 20, 2021 (received for review June 30, 2021)

The increasing success of populist and radical-right parties is one of the most remarkable developments in the politics of advanced democracies. We investigate the impact of industrial robot adoption on individual voting behavior in 13 western European countries between 1999 and 2015. We argue for the importance of the distributional consequences triggered by automation, which generates winners and losers also within a given geographic area. Analysis that exploits only cross-regional variation in the incidence of robot adoption might miss important facets of this process. In fact, patterns in individual indicators of economic distress and political dissatisfaction are masked in regional-level analysis, but can be clearly detected by exploiting individual-level variation. We argue that traditional measures of individual exposure to automation based on the current occupation of respondents are potentially contaminated by the consequences of automation itself, due to direct and indirect occupational displacement. We introduce a measure of individual exposure to automation that combines three elements: 1) estimates of occupational probabilities based on employment patterns prevailing in the preautomation historical labor market, 2) occupation-specific automatability scores, and 3) the pace of robot adoption in a given country and year. We find that individuals more exposed to automation tend to display higher support for the radical right. This result is robust to controlling for several other drivers of radical-right support identified by earlier literature: nativism, status threat, cultural traditionalism, and globalization. We also find evidence of significant interplay between automation and these other drivers.

automation | voting | polarization

Populist and radical-right parties and candidates have become increasingly successful in Western democracies over the past decades. A growing body of research has investigated the drivers of this political shift, identifying mainly two broad sets of explanations (see ref. 1 for a review). A first group of studies has emphasized the role of cultural drivers, as related to growing status-threat perceptions and to a rise of nativist and xenophobic attitudes (see, among others, refs. 2 and 3). A second group of studies has stressed the role of structural economic changes, such as globalization and technological progress, seen as drivers of deepening economic and social cleavages that tend to be politically consequential (see, among others, refs. 4–7). Some scholars see cultural and economic explanations as fundamentally alternative and potentially mutually exclusive (e.g., refs. 2 and 5). Others have emphasized their complementarity, and even their interdependence, to the extent that cultural concerns may mediate or moderate the political impact of economic grievances (8) and vice versa (9).

We ask what role automation plays in this political phenomenon, specifically through the robotization of manufacturing, a process that started in the early 1990s and gained momentum in the 2000s. The extant literature on the structural economic drivers of the recent electoral realignments has predominantly focused on globalization and specifically on the shocks generated by surging imports from China (and other

emerging economies) and by offshoring. Yet, automation is arguably the main economic transformation affecting Western democracies after the financial crisis of 2008. In fact, both the “China shock” and offshoring peaked before the crisis, while robotization intensified after 2010 (10, 11). It is true that globalization shocks tend to be persistent in terms of economic and political implications, with long-term effects detected even on recent political events [e.g., Brexit (12) and the election of Trump in 2016 (10)]. At the same time, as noticed by the authors of ref. 7, it is surprising that relatively little attention has been paid to the political implications of the most recent automation wave, despite its well-documented distributional consequences. Indeed, automation leads to skill-biased changes in job opportunities, associated with rising income inequality and the emergence of winners and losers (13–16). Individuals endowed with skills that are complementary to the new technologies benefit from automation, while more substitutable workers lose out. These labor-market developments are likely to be politically consequential. Public opinion seems to be quite concerned with the economic consequences of automation. For instance, according to a survey conducted by the Pew Research Center in 2018, 85% of respondents in Canada, 83% in Italy, and 90% in Japan stated that robots and computers will “definitely” or “probably” do much of the work currently done by humans. According to the same survey, large majorities believed that, as a consequence of automation, there would be fewer jobs (78%

Significance

The success of radical-right parties across western Europe has generated much concern. These parties propose making borders less permeable, oppose ethnic diversity, and often express impatience with the institutions of representative democracy. Part of their recent success has been shown to be driven by structural economic changes, such as globalization, which triggers distributional consequences that, in turn, translate into voting behavior. We ask what are the political consequences of a different structural change: robotization of manufacturing. We propose a measure of individual exposure to automation and show that individuals more vulnerable to negative consequences of automation tend to display more support for the radical right. Automation exposure raises support for the radical left too, but to a significantly lower extent.

Author contributions: M.A., I.C., and P.S. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. R.F. is a guest editor invited by the Editorial Board.

Published under the PNAS license.

¹M.A., I.C., and P.S. contributed equally to this work.

²To whom correspondence may be addressed. Email: piero.stanig@unibocconi.it.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2111611118/-DCSupplemental>.

Published November 19, 2021.

in Canada, 82% in Italy, and 72% in Japan) and the gap in wealth would increase (79% in Canada, 77% in Italy, and 86% in Japan). In two different Pew surveys run in the United States in 2018–2019, 50% of respondents stated that the automation of jobs through new technology in the workplace has mostly hurt American workers, and 75% of respondents believed that the automation of jobs has contributed either “a fair amount” or “a great deal” to economic inequality in the United States.*

This contribution aims at fostering our understanding of the political implications of automation. By so doing, we aim to contribute to the literature on the economic drivers of populism and radical-right support. At the same time, we engage with the other main explanations identified by earlier studies—nativism, status threat, cultural traditionalism, and globalization—and we investigate how automation interacts with these different drivers in affecting electoral outcomes in western Europe.

In this work, we argue and show empirically that studying the effect of automation on voting behavior poses two main challenges:

1. In regional or electoral district-level analysis, the political consequences of automation might be masked to a large extent by the aggregate welfare gains brought about by automation. Regions that adopted more robots in recent decades might be the ones that also experience comparatively more economic dynamism. At the same time, these regions are potentially the most affected by the distributive impacts of technological change. This fact requires researchers to go beyond regional or electoral district-level analyses and study more closely the individual dynamics of economic distress and voting behavior.
2. When performing individual-level analyses, it is necessary to employ individual measures that do not hinge exclusively on the automation exposure of the current occupation. In fact, the current occupation of individuals might be already contaminated by earlier automation dynamics, due to direct or indirect displacement. In particular, some workers may be currently employed in low-automatability service occupations because: 1) They have been directly displaced by robots from an earlier manufacturing job; or 2) as new entrants in the labor market (or dissatisfied, long-term on-the-job searchers), they could not find well-paid and secure “good jobs” in manufacturing, due to indirect displacement caused by automation. Based on their current occupation, these workers would get assigned low automation exposure, but they are arguably canonical automation losers. Moreover, no current occupation is available for unemployed individuals, who could as well be automation losers due to direct or indirect displacement.

Fig. 1 presents empirical evidence that strongly supports the first argument. In the upper panel, we group individuals based on the exposure to automation of their region of residence[†] and plot the distribution of their individual vulnerability to automation, measured following the method presented later in this article. Two important facts emerge. First, for a given level of regional exposure, there is substantial heterogeneity in individual vulnerability. Second, the distributions across the three groups of individuals are largely overlapping. Research focusing on purely cross-regional variation in exposure to automation might end up

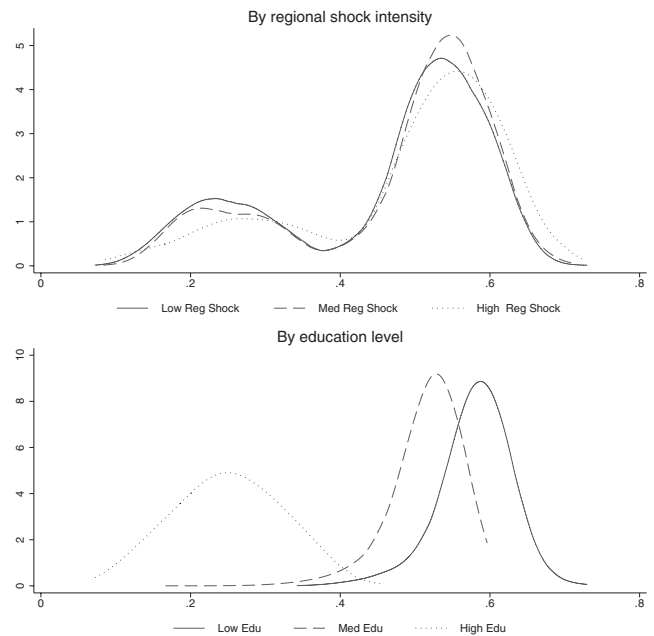


Fig. 1. (Upper) Kernel densities of the individual measure we propose for ESS respondents broken down by tercile of the regional shock of their region of residence, calculated following ref. 13, as described in *SI Appendix, section 1*. Med, medium; Reg, regional. (Lower) Kernel densities of the individual measure for ESS respondents, broken down by level of education, respectively, including university level (High Edu; dotted line), secondary education completed (Med Edu; dashed line), and less than secondary education completed (Low Edu; solid line).

overlooking some of the most politically consequential dynamics. In the bottom panel of Fig. 1, we shed further light on the distributional consequences of robot adoption within regions, by focusing on one driver of the individual vulnerability distribution: education. As one would expect, more educated individuals tend to be less vulnerable to automation. Yet, there is quite some variation in vulnerability within each educational group, resulting in significant overlaps in the distributions, especially for low- and medium-education individuals.

The importance of an individual-level analysis becomes even more evident when empirically exploring the socioeconomic cleavage driven by automation and underlying a potential effect on political behavior. In Fig. 2, we assess the relationship of automation with economic conditions and perceptions, as well as with politically consequential attitudes. We do this both exploiting individual-level variation in exposure to automation within regions and applying the same regional shock to all individuals living in a given region.

There are strong and clearly discernible associations between individual exposure to automation, as we propose to measure it, and self-reported socioeconomic outcomes and attitudes: Individuals with higher robot exposure are less likely to have a permanent contract, display worse perceptions of economic conditions and well-being, and report lower satisfaction with the government and democracy. Conversely, an association between the automation exposure of a region (at a given point in time) and individual conditions and perceptions is detectable only imprecisely and, in most cases, is not statistically discernible. This attests to the importance of leveraging individual variation to better capture the potential effects of automation on political outcomes.

Assessing individual exposure to automation, however, poses the second crucial challenge we highlight: Individual measures based on current occupation are contaminated by past

*Older survey data from the 1990s deliver similar messages. For instance, the 1997 “Work Orientations” module of the International Social Survey Program asked respondents whether they believed that technology would reduce the number of jobs in the future. Across the advanced democracies included in the survey, the public largely agreed that this would be the case. Specifically, the fraction of respondents stating that new technologies would “slightly” or “greatly” reduce the number of jobs ranged from more than 85% in Germany to 48% in Denmark; it was 50.5% in the United States and 75% in Switzerland, the median country in the group.

[†]Regional exposure is measured following ref. 13; see *SI Appendix, section 1* for details.

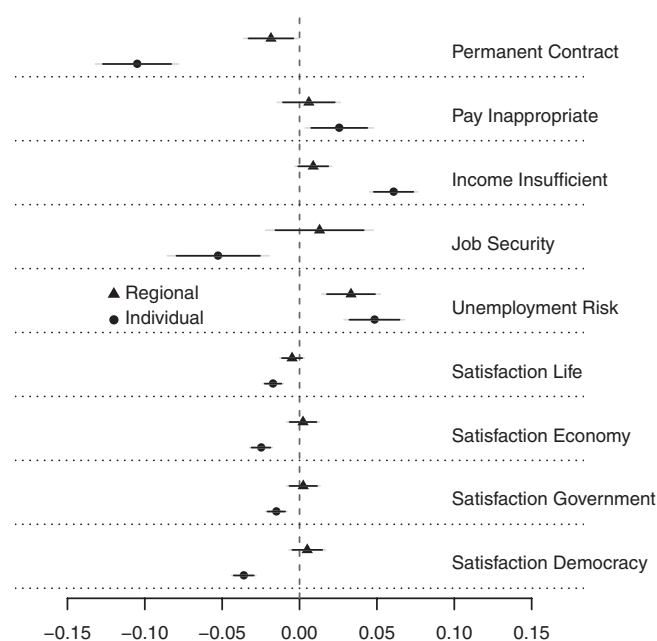


Fig. 2. Coefficients of regressions of individual-level variables from the ESS (standardized on a 0 to 1 scale) on, respectively, the regional exposure of the region of residence of the respondent (triangles) and the individual exposure of the respondent (dots). The bars are 90% and 95% CIs: If the interval crosses the dashed vertical line, the null hypothesis of no relationship cannot be rejected. Fixed effects at the region and country-year level are included; SEs are clustered at the region-year level.

automation dynamics, due to direct or indirect displacement. The role of “direct displacement” is very intuitive. To illustrate, consider an individual who is displaced, due to robot adoption, from a well-paid and stable job in manufacturing. This individual then finds a new job in services, e.g., as a janitor in a fast-food restaurant, at a lower wage and with a temporary contract. If we were to focus on the current (i.e., postdisplacement) occupation to assess the automation exposure of this individual, we would attribute a low score, since the new occupation is not highly automatable. Yet, this hypothetical individual is arguably the canonical case of an automation loser, and a measure based on current occupation would not capture it. Even worse, using current occupation would not allow us to assign a value of automation exposure to workers who are directly displaced and remain unemployed: Hence, we would exclude from the analysis an important group of negatively affected individuals.

It is key to recognize that automation not only affects workers initially employed in more automatable occupations (i.e., through direct displacement), but also reduces employment opportunities in such occupations for job seekers: We see this as the “indirect displacement” effect of automation. This is relevant for new entrants in the labor market, who might find themselves unemployed or employed in second-best occupations with low automation intensity. For instance, ref. 14 shows that, in Germany, robot adoption leads to a reduction in the number of job opportunities in manufacturing for new entrants. Indirect displacement applies also to longer-term labor-market insiders who are unsatisfied with their second-best occupation and are searching for good jobs in occupations that—being more automatable—become increasingly hard to find as automation progresses. Both new entrants and long-term on-the-job searchers are thus affected by automation through indirect displacement in terms of reduced job opportunities.

We argue that it is important to develop new approaches that can capture the exposure of individuals to automation in a way that is not contaminated by the consequences of automation

itself. In this work, we propose one strategy, which leverages historical information about the configuration of labor markets in the early 1990s, before the latest spurt of automation. Specifically, rather than using the observed current occupation, we work with predicted probabilities of employment in each occupation. We estimate the parameters of a model of occupation as a function of education, age, gender, and region of residence, using early 1990s European Labor Force Survey (EU-LFS) data (17), separately for each country. We then use such estimates to compute predicted probabilities of working in each occupation for the respondents of the European Social Survey (ESS), interviewed between 2000 and 2016, in times of high-paced robotization. The predicted probabilities are then combined with a measure of automatability for each occupation (18) to obtain a measure of “individual vulnerability to automation.” The latter is multiplied times the pace of robot adoption in a given country and year, to obtain a measure of “individual exposure to automation.” In a nutshell, for a given national pace of robot adoption, our measure assigns higher automation exposure to more vulnerable individuals, i.e., those who would have been more likely, in the preautomation historical labor market, to work in more automatable occupations. Indeed, the idea is to capture the fact that, due to automation, certain job opportunities have been disappearing, potentially leading to unemployment or employment in less desirable occupations.

This measure makes a step forward in the direction of capturing the effects of automation in a way that hinges on the potential vulnerability of each individual to the phenomenon, rather than on the current occupation. We hope our measure will stimulate further methodological reflection and empirical analysis aimed at developing new approaches to measure individual exposure to technological progress and, in general, to structural changes in labor markets. In this respect, in our analysis, we also provide an application of our methodology to offshoring.

In the main analysis, we study the effects of industrial robot adoption on individual voting behavior in 13 western European countries.[‡] We find that, conditional on individual characteristics, on time-invariant regional characteristics (via region fixed effects), and on determinants of voting behavior shared by all individuals in a given country at a given election (via country-year fixed effects), higher individual exposure to automation triggers more electoral support for radical-right parties.

Automation and Politics

From a theoretical perspective, the general idea behind studies linking structural economic transformations to support for populist and radical-right parties is pretty straightforward: Economic processes such as globalization and technological change create aggregate welfare gains, but with winners and losers. Such distributional consequences, through the ensuing perceived economic distress, are then politically consequential. What is not straightforward—and still underdeveloped, according to ref. 7—is the theoretical understanding of the mechanisms through which automation-driven economic distress, in particular, may translate into higher support specifically for radical-right parties. In this respect, we propose several non-mutually-exclusive channels.

First, it is well understood that economic discontent may lead to disaffection with mainstream parties and the political establishment at large. In this context, radical-right forces provide an appealing option for dissatisfied voters, as, in most cases, they can credibly cast themselves as an antiestablishment alternative to mainstream parties (1, 5, 19). Such an antiestablishment advantage can extend to radical-left parties, as we discuss below.

[‡]These are Austria, Belgium, Finland, France, Germany, Italy, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom.

Moving beyond the simple antiestablishment motivation, a second advantage that is specific to radical-right parties may stem from their economic nationalist platforms. These are particularly alluring in the wake of structural transformations of the economy, as they offer a very generic promise of protection. This crucially involves the broad idea of “taking back control” of the country from global impersonal forces—such as those behind globalization and technological change (20)—which are perceived to be detrimental to economic prosperity by distressed constituencies. In reality, it might be very hard for people to disentangle precisely the causes of perceived economic hardship; these are difficult to identify even for researchers. Different causes of economic distress may be conflated; for instance, there is evidence that the same type of workers may be more exposed to both international competition and automation (21). Moreover, situations of economic hardship may be misattributed; in particular, there is evidence that people may misattribute to immigration and international trade economic distress that is actually due to automation (22, 23). This phenomenon may be reinforced by political entrepreneurs who find it easier to provide responses to immigration and trade rather than to automation (24, 25). In addition, even when explicitly primed about hypothetical layoffs being caused by automation in experimental contexts, people still demand more trade protectionism as a policy response (26). All this may fuel support for radical-right parties, which are strongly protectionist, antiimmigrant, and in favor of a broad concept of workers’ protection (6).

Third, and relatedly, besides the economic platforms, the idea of taking back control is often combined with the defense of a traditional way of life that supposedly characterized the nation before globalization and immigrants—but also computers and robots—had a disruptive impact on society. Evidence shows that nostalgia for a mythical (recent) past plays a significant role in radical-right support (27–29). The nostalgic rhetoric typically involves an emphasis on traditional family structure, with a strong role for the male head of household empowered by a well-paid and stable job (30, 31). Intuitively, the nostalgic appeal to a bygone era is particularly attractive for individuals whose relative standing in society is threatened or declining, in line with recent findings by refs. 32–34. As the authors of ref. 24 effectively summarize, “rightwing populist parties’ promises to turn back the clock seem to strike a chord with routine workers’ fears of social regression.”

Fourth, a growing literature shows that structural changes such as automation have not only economic, but also social and psychological, consequences. There is accumulating evidence that economic vulnerability and inequality lead to increased authoritarianism and nativism (e.g., refs. 35 and 36). Recent work in psychology directly documents a robust association between concerns for automation and nativism; in addition, experimental subjects propose to shed more immigrant workers when layoffs are motivated by automation, as opposed to generic “company restructuring” (37). These types of reactions naturally push voters toward radical-right parties, whose platforms feature strong authoritarian and nativist elements (38, 39).

The main focus of this paper is on radical-right parties. Yet, one could expect automation-driven distress to determine higher support for radical-left parties too. As a matter of fact, these parties could also benefit from an antiestablishment advantage. Moreover, one could argue that preredistribution platforms—typical of the extreme left—would be appealing to distressed constituencies. However, as reviewed in ref. 7, the evidence on whether exposure to automation increases demand for redistribution is actually mixed. Moreover, the automation-driven authoritarian and nativist shift in attitudes is unlikely to tilt voters toward radical-left parties, which are traditionally more libertarian and solidaristic (20, 40). Whether automation exposure raises support for the radical left is ultimately an empirical question,

which we address in our analysis. As a preview, we find evidence of a positive effect, although of a substantially smaller magnitude compared to the radical right.

There is limited evidence, thus far, on the consequences of the most recent spurts of technological change on political behavior. A handful of recent studies have taken a regional approach. Studying the 2016 US presidential election, ref. 41 shows that voters in regions more affected by robotization in manufacturing were more supportive of the Republican candidate, Donald Trump, who was running on a nationalist platform akin to those of the European radical right. Similar evidence is provided by refs. 42 and 43 for radical-right parties in Europe. At the individual level, ref. 44 finds that computerization winners in the United Kingdom become more likely to vote Conservative and less likely to vote Labor. Closer to our work, refs. 43, 45, and 46 show that individual workers currently employed in occupations at higher risk of automation are more prone to vote for radical-right parties across several European countries. Yet, using the current occupation to assess automation exposure is potentially problematic, due to direct and indirect displacement caused by automation.

Two studies have started to shed light on this issue using panel data at the individual level. Focusing on the United Kingdom, ref. 16 finds that only 64% of workers initially employed in routine-intensive (and thus more automatable) occupations survive in those types of jobs over the sample period (1991–2015). Similar evidence is provided by ref. 34 across the United Kingdom, Germany, and Switzerland. Interestingly, ref. 34 provides evidence of potentially different political responses for workers initially employed in routine-intensive occupations, depending on whether they are actually displaced or not. Specifically, those who cling to their jobs are more likely to vote for populist-right parties, while displaced workers who remain unemployed may, rather, display a higher probability of supporting left parties. Both refs. 16 and 34 reckon that the bulk of compositional changes in the labor market—by which more automatable occupations shrink over time—are driven by labor-market entrants, who are less likely to find jobs in such occupations. The approach that we propose for assessing individual exposure to automation allows us to take into account such indirect displacement of labor-market entrants, which is likely to be politically consequential. This is not captured by the design in ref. 34, which focuses on workers initially employed in routine occupations. The same applies to long-term unemployed individuals or to the indirect displacement of workers initially employed in nonroutine occupations, who potentially cannot find the desired good job in routine-intensive activities due to automation. By taking into account both direct and indirect displacement caused by automation, our approach allows us to assess in a very comprehensive way automation-driven displacement, as compared to expected/desired occupational outcomes. More in general, we provide a methodological tool that can be applied to any cross-sectional database in order to measure exposure to automation—and potentially to other structural changes—in a way that is not contaminated by the consequences of the structural change that is being studied.

Materials and Methods

We rely on the first seven waves of the ESS (47). We exploit information on the party voted for by each individual in the last election before the interview date. Elections span the period 1999–2015. Our main focus is on a dummy equal to one if the chosen party is categorized as a radical-right party. We classify radical-right parties following the conventional wisdom in the literature, along the same lines as ref. 6.⁵

⁵See *SI Appendix, section 2* for the full list of radical-right parties.

The baseline specification we estimate has the general form:

$$\text{Vote Choice}_{icrt} = \alpha_{ct} + \alpha_r + \beta_1 \text{Individual Exposure}_{it} + \beta_2 X_{it} + \varepsilon_{icrt}, \quad [1]$$

where i indexes individuals, c countries, r regions, t election years, and ε_{icrt} is an error term. Individual Exposure_{it} is the individual exposure to automation, computed as outlined in Eq. 3, over 2 y prior to the election. X_{it} is a vector of individual-level observable characteristics, including, in all models, age, gender, and years of education. α_r are region fixed effects, while α_{ct} are country-year fixed effects. Country-year fixed effects are equivalent to election fixed effects. They provide a parsimonious—albeit crude—way to account for factors that are specific to a given country at the time of a given election. Examples of such factors include political supply-side characteristics, such as the quality of the incumbents, as well as other contextual elements, such as the general economic and political climate. SEs are clustered at the region-year level. The coefficients refer to standardized individual exposure, so they can be interpreted as the effect of a one-SD change.¹¹

In order to account for other determinants of radical-right support, as identified by earlier literature, we also estimate models where we augment Eq. 1 with a number of controls for such other factors, namely: cultural traditionalism, nativism, status threat, China shock, and offshoring. In the most complete specifications, we also include the interaction between individual exposure to automation and each “Other” factor, as follows:

$$\text{Vote Choice}_{icrt} = \alpha_{ct} + \alpha_r + \beta_1 \text{Individual Exposure}_{it} + \beta_2 X_{it} + \beta_3 \text{Other}_{it} + \beta_4 \text{Individual Exposure}_{it} \cdot \text{Other}_{it} + \varepsilon_{icrt}. \quad [2]$$

We construct our individual measure of automation exposure based on a vector of predicted probabilities for each individual to be employed in each occupation. Crucially, these probabilities are computed based on individual characteristics and on the preautomation, historical composition of employment at the occupation level. The individual vulnerability to automation is then obtained as the scalar product between this vector of probabilities and a vector of automatability scores of the occupations. In other words, the vulnerability score for individual i is a weighted average of the automatability scores of each occupation, where the weights are the predicted probabilities of employment of individual i in each occupation. To obtain the individual exposure to automation at the time of a given election, the vulnerability score is further interacted with the pace of robot adoption in the specific country and election year.

Specifically, we compute:

$$\text{Exp}_{icrt} = \underbrace{\left[\sum_j \hat{Pr}(o_i = j | \text{age, gender, edu, } r) * \theta_j \right]}_{\text{Individual Vulnerability}} * \Delta R_{ct}, \quad [3]$$

where $\hat{Pr}(o_i = j | \text{age, gender, edu, } r)$ is individual i 's probability of working in each occupation j , predicted based on age, gender, educational attainment, and region dummies, using historical employment data from the EU-LFS. The score θ_j is the automatability of occupation j . Summing the product of $\hat{Pr}(o_i = j | \text{age, gender, edu, } r)$ times θ_j over all occupations, we obtain a value of individual vulnerability to automation. The individual exposure to robot adoption in year t is then obtained by multiplying this individual vulnerability by ΔR_{ct} , which is the national percentage change in total operational robots between year $t - 1$ and $t - 3$, in the country c where individual i resides, i.e., $\frac{R_{c(t)}^{t-1} - R_{c(t)}^{t-3}}{R_{c(t)}^{t-3}}$.¹² Data on robot adoption are sourced

from the International Federation of Robotics (IFR) (48).¹³

Intuitively, for a given national pace of robot adoption, our measure of individual exposure assigns higher scores to individuals that would have been more likely—in the preautomation historical labor market—to work in occupations whose automatability is higher. To illustrate, consider a hypothetical 20-y-old man without a college degree, observed in the year 2015,

who lives in a region where many individuals with a similar demographic profile used to work in high-automatability occupations at the beginning of the 1990s. His vector of employment probabilities, based on the historical estimates, would contain relatively high values corresponding to more automatable occupations. This man would then receive a relatively high score of exposure to automation in 2015, even though he is currently unemployed or employed in an occupation that is not highly automatable, e.g., delivery service. This approach allows us to capture the fact that automation has reduced the availability of certain job opportunities, potentially leading to unemployment or employment in second-best occupations. In *SI Appendix, section 3*, we provide further details and evidence on how our exercise captures and reflects the changes in labor markets and occupational opportunities that have been taking place over the past three decades.

In the baseline analysis, the occupation-level automatability score, θ_j , is sourced from ref. 18. It is a computerization probability based on a combination of expert data and detailed task content. In *SI Appendix, section 5*, we probe the robustness of our results to using an alternative score, based on the perceived threat of automation of workers employed in each occupation, as self-reported by respondents of the International Social Survey Program in 1997.

When regressing electoral outcomes on exposure to robots, endogeneity issues linked to the change in total operational robots could arise. First, robot adoption tends to be procyclical: Firms install more robots during periods of stronger economic growth. If economic cycles were associated with support for political parties, the ordinary least squares (OLS) estimates of the impact of robots on voting would be biased. In particular, if voters in good times tend to support more mainstream parties, rather than radical-right parties, we would expect a downward bias in the estimates. Second, more robots might be installed in regions with stronger employment-protection legislation, which makes labor relatively more costly. Given that employment legislation in Europe is determined at the national level, we reduce this concern by including country-year fixed effects in our regressions. The pace of robot adoption in a region may also be influenced by the local strength of labor unions. To the extent that unionization is systematically associated with performance of political parties, we would have a confounding factor biasing the estimates.

To address these endogeneity concerns, similarly to ref. 13, we instrument robot adoption in each country with robot adoption in other countries. In practice, the instrument is based on the same individual vulnerability component, but this time interacted with the average percentage change in the stock of operational robots across all other sample countries (i.e., excluding c), between year $t - 1$ and $t - 3$. The instrument is meant to exploit trajectories in automation driven by technological innovations shared across countries. Its validity hinges on the fact that the adoption of robots in other countries is plausibly exogenous to domestic political dynamics. The exclusion restriction could be violated in the presence of demand shifts correlated across countries. To address this concern, in *SI Appendix, section 5B*, we probe the robustness of our results to alternative instruments. We consider the more distant—and less cyclically correlated—non-European advanced economies for which IFR data are available (i.e., countries of North America, Japan, and South Korea). We also employ four variables meant to directly capture technological shifts in robots and computing: the average unit price of industrial robots sold in the United States (IFR), a producer index of computer prices (Federal Reserve Economic Data, Federal Reserve Bank of St. Louis) (49), single-thread performance, and number of transistors per microprocessor (50). Results, presented in *SI Appendix, Table S5*, are robust to these alternative instrument specifications.

The Effect of Individual Exposure to Automation on Voting Behavior

The top row of Fig. 3 reports the baseline estimates of Eq. 1. Both the OLS and the instrumental variable (IV) coefficients on individual exposure to robots are positive and significant. The IV estimate is somewhat higher than the OLS one, consistent with the well-known procyclicality of robot adoption. The F -statistic, reported in *SI Appendix, Table S2*, is comfortably high. These results suggest that individuals more exposed to automation are more likely to support radical-right parties. In terms of magnitude, the IV coefficient can be directly read as the effect of a one-SD increase in automation exposure. This leads to a 2.8-percentage-point increase in the probability of voting for a radical-right party. This effect is substantively important,

¹¹We estimate linear probability models in the baseline analysis, since our specifications include multiway fixed effects and are estimated using IVs. Nevertheless, in *SI Appendix, Table S6*, we show that our main finding is robust to using IV-probit.

¹²See *SI Appendix, section 4* for estimations based on an alternative interpretation of the multiplicative nature of individual exposure.

¹³See *SI Appendix, section 3* for an assessment of the fit of the occupational models and for further methodological details on the construction of the individual exposure measure.

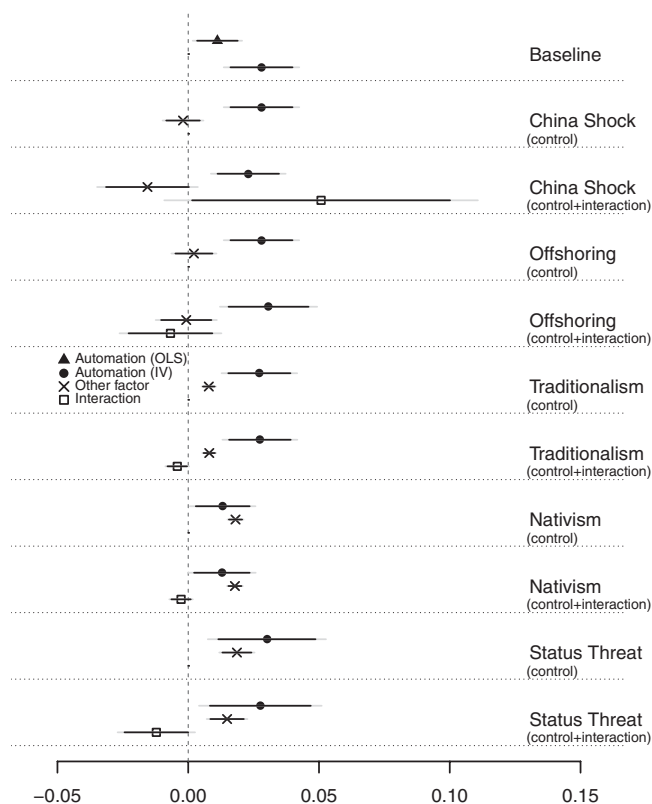


Fig. 3. Regression estimates of the impact of a one-SD increase in individual-level robot exposure on voting for a radical-right party. For each potential driver, we report two specifications: one that controls for the other factor (labeled “control”) and one that also interacts robot exposure and the other factor (labeled “control + interaction”). The dots represent the IV estimate of the linear term of automation exposure, the x marks the linear term of the variable listed, and the open squares represent the interaction between automation and the variable listed. The bars correspond to 90% and 95% CIs: If the interval crosses the dashed vertical line, the null hypothesis of no relationship cannot be rejected. Fixed effects at the region level and the country-year level are included; SEs are clustered at the region-year level. The *F*-statistics for the excluded instrument in the first-stage regressions are reported in *SI Appendix, Table S2*.

considering that the baseline probability of voting for a radical-right party in the sample is about 4.8%.^{**}

In *SI Appendix, section 5*, we present an extensive set of robustness specifications for our baseline analysis. Notably, results are robust to alternative estimations of our individual exposure measure (*SI Appendix, section 5A*); alternative instruments (*SI Appendix, section 5B*); propagation of uncertainty from the occupation models to the voting models (*SI Appendix, section 5C*); and the exclusion of automotive workers and the inclusion of controls for the current occupation (*SI Appendix, section 5D*). In particular, in one robustness check, we augment the specification with a measure of automation exposure based on the current occupation. This is found to be positively related to radical-right support. Yet, the coefficient on our main individual exposure variable remains stable in size and statistical significance, suggesting that our measure is capturing a different (and plausibly not contaminated) source of variation compared to current employment status.

^{**}To provide a more concrete sense of the kind of regularities that these estimates capture, *SI Appendix, Fig. S4* plots the data from several districts that are illustrative of the general pattern we detect.

Finally, to characterize more broadly the impact of automation on voting, we estimate our baseline regression using three alternative outcome variables. These are dummies denoting whether the party voted for by the respondent belongs to one of the following party families: Radical Left, Mainstream Left, and Mainstream Right.^{††} In *SI Appendix, section 5E*, we show that parties belonging to different families are systematically different along three ideological dimensions that are key in our context, in light of the theoretical discussion: Net Autarky, Nationalism, and Economic Conservatism. These are measured for each party, in each election, using Manifesto Project data (51), as in ref. 6. Conditional on country-year fixed effects, radical-right parties are the most protectionist and isolationist, as denoted by higher levels of the net autarky index. They are followed in the ranking by radical-left parties and mainstream-right tied with mainstream-left parties. Radical-right parties also display the highest levels of the nationalism score, denoting more traditionalist, authoritarian, and nationalist positions. They are followed in this case by mainstream-right, mainstream-left, and radical-left parties. The ranking in terms of economic conservatism sees radical-right parties tied with mainstream-right parties, followed by mainstream-left and radical-left parties, which thus emerge as the most proredistribution.

In line with the theoretical expectations, higher automation exposure raises support for radical-left parties. Yet, the estimated effect is significantly smaller—about one-third—than the one estimated for the radical right (0.01 vs. 0.028).^{‡‡} This is consistent with radical-left parties offering relatively high levels of protectionism, similarly to the radical right, but not high levels of authoritarian nationalism. If automation exposure increases both protectionist and authoritarian nationalist demands—as per our theoretical discussion—radical-right parties then enjoy an advantage over the radical left. This advantage does not seem to be compensated by the fact that radical-left parties are more proredistribution, in line with the mixed evidence reviewed in ref. 7 on the effect of automation exposure on demand for redistribution. Support for mainstream-right parties is significantly reduced by automation exposure (−0.037), while for the mainstream left, the estimated coefficient is negative, but not statistically significant (see *SI Appendix, Table S7* for full results). Overall, automation seems to determine an increase in political polarization, with rising support for extremist forces at both ends of the political spectrum and diminishing support for mainstream parties. This evidence is in line with earlier findings on the effects of globalization threats in the United States (10). In Europe, on the other hand, globalization has been documented to tilt voters toward the radical right, but not the radical left (6).

An interesting question to ask is whether economic distress driven by automation vs. globalization determines different effects in terms of policy preferences. According to ref. 7, potential differences may arise, in theory, due to several reasons, including the more gradual and sustained nature of automation processes over time, as compared to international trade, and the absence of a clearly identifiable outgroup to mobilize against in response to automation. Based on a survey conducted in Spain, ref. 52 finds that workers more exposed to automation are not more likely to demand compensation policies, while they are more likely to support policies that would slow down the pace of technological innovation. The authors of ref. 26 perform a survey experiment, in which US respondents are primed about the causes of

^{††}Radical-left parties are identified in the same way as radical-right parties, based on the general consensus in the literature. For mainstream parties, we follow the Comparative Manifesto classification (51). See *SI Appendix, section 2* for more details.

^{‡‡}The difference between the two coefficients is statistically significant. This is corroborated by the joint estimation of the two equations in a seemingly unrelated regression model, whereby the hypothesis of equality of the two effects is rejected at the 1% level.

hypothetical job losses (e.g., trade/offshoring vs. automation). In line with evidence in ref. 52, there is no strong evidence of substantially different effects in terms of preferences for compensatory policies targeted to affected workers. Interestingly, individuals seem to demand more trade protectionism not only in the case of job losses due to trade/offshoring, but also in the case of automation-driven layoffs (26). This is, indeed, one channel through which exposure to automation may raise support for radical-right parties, which tend to be strongly protectionist (6).

In reality, globalization and automation dynamics may actually be closely intertwined, and their economic and political effects may be mutually reinforcing. For instance, consider a region that has suffered job losses in manufacturing due to import competition from China. Surviving firms in that region may try to react and become more competitive through automation, which poses an additional threat to employment. On the other hand, globalization and automation threats may also be seen as substitutes when considering offshoring dynamics. In fact, there is evidence that automation has been facilitating the reshoring of manufacturing activities back to high-wage contexts, such as the United States and Europe (11). Reshoring may create new job opportunities, although primarily for workers whose skills are complementary to robots. This speaks again to the importance of individual-level measures of vulnerability to automation.

Against this backdrop, in rows 2–5 of Fig. 3, we show results from augmented specifications, where we take into account two different dimensions of globalization: Chinese imports and offshoring. First, in row 2, we include as a control the exposure to Chinese imports in the region of residence, computed over 2 y prior to each election, as in ref. 6. Then, in row 3, we also include the interaction between the China shock and our measure of individual exposure to automation. That is, we estimate Eq. 2 using exposure to Chinese imports as the “Other” factor. In both cases, the coefficient on automation exposure remains in line with the baseline IV estimate. Moreover, the coefficient on the interaction term is positive and marginally significant, suggesting that there might indeed be some reinforcing effect between regional exposure to Chinese imports and individual exposure to automation.## In row 4, we include as a control the individual exposure to offshoring. This is computed in a similar way as exposure to automation, thus providing another application of our methodology. Specifically, compared to Eq. 3: 1) We employ the same occupational probabilities; 2) θ_j is the offshorability index by ref. 53^{##}; and 3) the percentage change in national robots (ΔR_{ct}) is replaced by the percentage change in national imports. The inclusion of this individual control leaves our main estimate on automation exposure essentially unaffected.### In row 5, we also include the interaction term between exposure to offshoring and automation. The estimated coefficient is negative, though imprecisely estimated, broadly in line with the idea of some substitutability between automation and offshoring threats at the individual level.

Culture, Status, and Radical-Right Support

Thus far, our analysis has focused on structural economic drivers of radical-right support. Yet, the literature has provided abundant evidence on the role of cultural drivers, too (see, among others, refs. 2 and 3). In rows 6–11 of Fig. 3, we engage with such factors. Specifically, we employ three individual-level variables

based on ESS items: cultural traditionalism, nativism, and status threat. Cultural traditionalism is proxied as in ref. 27, through the agreement with the statement that gays and lesbians should be free to live life as they wish. While domain-specific, this item is related to one of the main noneconomic components of the “cosmopolitan values” package. Nativism is measured with the average of two items, on whether immigrants are good or bad for the national economy and cultural life. Status threat is measured with the first factor from factor analysis of three items: 1) whether the respondents find it hard to be hopeful about the future of the world; 2) whether for most people life is getting worse; and 3) whether the respondents feel they are treated with respect. Higher values of the three cultural variables correspond to more traditionalism, more nativism, and stronger perceived status threat, respectively. Each variable is first included only as a control and then interacted with automation exposure, as per Eq. 2.

In line with earlier studies, cultural traditionalism, nativism, and status threat are positively and significantly related to radical-right support. At the same time, the coefficient on individual exposure to automation remains positive and precisely estimated across the board. This evidence is in line with the view that cultural and economic drivers are not necessarily alternative explanations for the success of the radical right.

One important question worth asking is to what extent automation and cultural drivers may be complements or substitutes. In this respect, the interaction terms are all negative, and marginally significant in the case of cultural traditionalism and status threat. This provides mild evidence of a substitution effect: Individuals with higher probability to choose the radical right for cultural reasons are less tilted by automation, while the effect of automation exposure is stronger for individuals who would be less attracted by the radical right in virtue of cultural traits.

The evidence just discussed points to the fact that cultural attitudes may moderate the impact of automation. However, they may also work as channels through which the effect of automation exposure on voting unfolds. Fig. 4 presents suggestive evidence in this direction. Specifically, coefficients refer to the effect of automation exposure on cultural traditionalism, nativism, and status threat, estimated as per Eq. 1. The three coefficients are positive and statistically significant, suggesting that economic hardship induced by automation may indeed tilt individuals’ attitudes in a way that pushes them toward radical-right platforms. This evidence is in line with earlier results showing that cultural drivers may be at least partially posttreatment with respect to economic shocks (e.g., refs. 32, 34, and 35).

All in all, two main messages emerge from this section. First, we are able to identify the role of automation while controlling for the role of cultural factors. The automation effect is




Fig. 4. Coefficients of IV regressions showing the impact of a one-SD increase in individual-level robot exposure on three potential channels linking automation and radical-right vote: cultural traditionalism (measured with an item regarding rights for gays and lesbians), an index of nativism, and an index of perceived status threat. The bars correspond to 90% and 95% CIs: If the interval crosses the dashed vertical line, the null hypothesis of no relationship cannot be rejected. Fixed effects at the region level and the country-year level are included; SEs are clustered at the region-year level.

Full results from all regressions reported in Fig. 3 are displayed in *SI Appendix, Table S2*.

Results are robust to using all the alternative offshorability indexes made available by ref. 54.

Importantly, if we run a falsification test, in which we multiply the vulnerability to offshoring times ΔR_{ct} , we do not get a significant result, suggesting that our measure of individual robot exposure is indeed capturing a dimension of economic distress that is specifically related to automation.

substantively stable when accounting for cultural factors compared to the baseline evidence, notwithstanding the inclusion of potential posttreatment controls. Second, we corroborate that cultural factors are relevant drivers of radical-right support, and we document that they interact with automation in different ways. The aim of this paper is to shed light on one specific economic factor, i.e., automation. Yet, thinking in broader terms, a research program aimed at thoroughly understanding extremist political dynamics definitely requires considering both economic and cultural drivers.

Conclusion

We study the effects of robot adoption on voting behavior in western Europe. We find that higher exposure to automation increases support for radical-right parties. We argue that an individual-level analysis of vulnerability to automation is required, given the prominent role played by the distributional effects of automation unfolding within geographic areas. We also argue that measures of automation exposure based on an individual's current occupation, as used in previous studies, are

potentially problematic, due to direct and indirect displacement induced by automation. We then propose an approach that combines individual observable features with historical labor-market data. Our paper provides further evidence on the material drivers behind the increasing support for the radical right. At the same time, it takes into account the role of cultural factors and shows evidence of their interplay with automation in explaining the political realignment witnessed by advanced Western democracies.

Data Availability. Individual-level data and replication code for all the results have been deposited in the Harvard Dataverse (<https://doi.org/10.7910/DVN/TF470>).

ACKNOWLEDGMENTS. We thank the editor (Robert Franzese) and five anonymous referees for insightful comments. We thank Paolo Agnolin, Jakob Wall, and Simon Tobias Weinmann for excellent research assistance. We thank Catherine de Vries, Livio Di Leonardo, Andrea Fracasso, Jeff Frieden, Aina Gallego, Peter Hall, Dominik Hangartner, Thomas Kurer, Arlo Poletti, Dani Rodrik, and seminar participants at various venues for useful comments. This work was supported by the European Union Horizon 2020 research and innovation programme under Grant Agreement 822390 Raising EU Productivity: Lessons from Improved Micro Data.

- S. Guriev, E. Papaioannou, The political economy of populism. *J. Econ. Lit.*, in press.
- D. C. Mutz, Status threat, not economic hardship, explains the 2016 presidential vote. *Proc. Natl. Acad. Sci. U.S.A.* **115**, E4330–E4339 (2018).
- P. Norris, R. Inglehart, *Cultural Backlash: Trump, Brexit, and Authoritarian Populism* (Cambridge University Press, Cambridge, UK, 2019).
- D. Swank, H. G. Betz, Globalization, the welfare state and right-wing populism in Western Europe. *Socioecon. Rev.* **1**, 215–245 (2003).
- L. Guiso, H. Herrera, M. Morelli, T. Sonno, "Populism: Demand and supply" (Tech. Rep. 11871, Centre for Economic Policy Research, London, 2017).
- I. Colantone, P. Stanig, The trade origins of economic nationalism: Import competition and voting behavior in Western Europe. *Am. J. Pol. Sci.* **62**, 936–953 (2018).
- A. Gallego, T. Kurer, Politics in the age of automation. *Annu. Rev. Polit. Sci.*, in press.
- R. J. Franzese, "The comparative and international political economy of anti-globalization populism" in *Research Encyclopedia of Politics*, W. R. Thompson et al., Eds. (Oxford University Press, Oxford, UK, 2019).
- Y. Margalit, Economic insecurity and the causes of populism, reconsidered. *J. Econ. Perspect.* **33**, 152–170 (2019).
- D. Autor, D. Dorn, G. Hanson, K. Majlesi, Importing political polarization? The electoral consequences of rising trade exposure. *Am. Econ. Rev.* **110**, 3139–3183 (2020).
- A. Bonfiglioli, R. Crinó, G. Gancia, I. Papadakis, "Robots, offshoring and welfare" in *Robots and AI: A New Economic Era*, L. Y. Ing, G. M. Grossman, Eds. (Routledge, Oxon, UK, 2021).
- I. Colantone, P. Stanig, Global competition and Brexit. *Am. Polit. Sci. Rev.* **112**, 128–151 (2018).
- D. Acemoglu, P. Restrepo, Robots and jobs: Evidence from US labor markets. *J. Polit. Econ.* **128**, 2188–2244 (2020).
- W. Dauth, S. Findeisen, J. Suedekum, N. Woessner, The adjustment of labor markets to robots. *J. Eur. Econ. Assoc.*, 10.1093/jeaa/fvab012 (2021).
- G. Graetz, G. Michaels, Robots at work. *Rev. Econ. Stat.* **100**, 753–768 (2018).
- T. Kurer, A. Gallego, Distributional consequences of technological change: Worker-level evidence. *Res. Polit.* **6**, 2053168018822142 (2019).
- Eurostat, European Labor Force Survey. <https://ec.europa.eu/eurostat/web/micro-data/european-union-labour-force-survey>. Accessed 12 November 2021.
- C. B. Frey, M. A. Osborne, The future of employment: How susceptible are jobs to computerisation? *Technol. Forecast. Soc. Change* **114**, 254–280 (2017).
- Y. Algan, S. Guriev, E. Papaioannou, E. Passari, The European trust crisis and the rise of populism. *Brookings Pap. Econ. Act.* **2017**, 309–400 (2017).
- I. Colantone, P. Stanig, The surge of economic nationalism in western Europe. *J. Econ. Perspect.* **33**, 128–151 (2019).
- E. Owen, N. P. Johnston, Occupation and the political economy of trade: Job routineness, offshorability, and protectionist sentiment. *Int. Organ.* **71**, 665–699 (2017).
- A. Kaihovaara, Z. J. Im, Jobs at risk? Task routineness, offshorability, and attitudes toward immigration. *Eur. Polit. Sci. Rev.* **12**, 327–345 (2020).
- D. C. Mutz, (Mis)attributing the causes of American job loss: The consequences of getting it wrong. *Public Opin. Q.* **85**, 101–122 (2021).
- T. Kurer, B. Palier, Shrinking and shouting: The political revolt of the declining middle in times of employment polarization. *Res. Polit.* **6**, 1–6 (2019).
- D. Rodrik, Why does globalization fuel populism? Economics, culture, and the rise of right-wing populism. *Annu. Rev. Econ.* **13**, 133–170 (2021).
- R. Di Tella, D. Rodrik, Labour market shocks and the demand for trade protection: Evidence from online surveys. *Econ. J. (Lond.)* **130**, 1008–1030 (2020).
- S. Bornschieer, H. Kriesi, "The populist right, the working class, and the changing face of class politics" in *Class Politics and the Radical Right*, J. Rydgren, Ed. (Routledge, London, 2012), pp. 1–23.
- J. Gest, T. Remy, J. Mayer, Roots of the radical right: Nostalgic deprivation in the United States and Britain. *Comp. Polit. Stud.* **51**, 1694–1719 (2018).
- E. Steenvoorden, E. Hartevelde, The appeal of nostalgia: The influence of societal pessimism on support for populist radical right parties. *West Eur. Polit.* **41**, 28–52 (2018).
- T. Akkerman, Gender and the radical right in western Europe: A comparative analysis of policy agendas. *Patterns Prejudice* **49**, 37–60 (2015).
- N. Spierings, A. Zaslove, Gendering the vote for populist radical-right parties. *Patterns Prejudice* **49**, 135–162 (2015).
- N. Gidron, P. A. Hall, The politics of social status: Economic and cultural roots of the populist right. *Br. J. Sociol.* **68** (suppl. 1), S57–S84 (2017).
- B. Burgoon, S. van Noort, M. Rooduijn, G. Underhill, Positional deprivation and support for radical right and radical left parties. *Econ. Policy* **34**, 49–93 (2019).
- T. Kurer, The declining middle: Occupational change, social status, and the populist right. *Comp. Polit. Stud.* **53**, 1798–1835 (2020).
- C. Ballard-Rosa, M. Malik, S. Rickard, K. Scheve, The economic origins of authoritarian values: Evidence from local trade shocks in the United Kingdom. *Comp. Polit. Stud.*, 10.1177/00104140211024296 (2021).
- N. Gennaioli, G. Tabellini, "Identity, beliefs, and political conflict" (Tech. Rep. 7707, CESifo, Munich, 2019).
- M. Gamez-Djokic, A. Waytz, Concerns about automation and negative sentiment toward immigration. *Psychol. Sci.* **31**, 987–1000 (2020).
- K. Arzheimer, Contextual factors and the extreme right vote in Western Europe, 1980–2002. *Am. J. Pol. Sci.* **53**, 259–275 (2009).
- M. Golder, Explaining variation in the success of extreme right parties in Western Europe. *Comp. Polit. Stud.* **36**, 432–466 (2003).
- L. Hooghe, G. Marks, Cleavage theory meets Europe's crises: Lipset, Rokkan, and the transnational cleavage. *J. Eur. Public Policy* **25**, 109–135 (2018).
- C. B. Frey, T. Berger, C. Chen, Political machinery: Did robots swing the 2016 US presidential election? *Oxf. Rev. Econ. Policy* **34**, 418–442 (2018).
- M. Caselli, A. Fracasso, S. Traverso, Globalization, robotization, and electoral outcomes: Evidence from spatial regressions for Italy. *J. Reg. Sci.* **61**, 86–111 (2021).
- H. V. Milner, Voting for populism in Europe: Globalization, technological change, and the extreme right. *Comp. Polit. Stud.*, 10.1177/0010414021997175 (2021).
- A. Gallego, T. Kurer, N. Schöll, Neither left-behind nor superstar: Ordinary winners of digitalization at the ballot box. *J. Polit.*, 10.1086/714920.
- Z. J. Im, N. Mayer, B. Palier, J. Rovny, The 'losers of automation': A reservoir of votes for the radical right? *Res. Polit.* **6**, 1–7 (2019).
- J. Gingrich, Did state responses to automation matter for voters? *Res. Polit.* **6**, 2053168019832745 (2019).
- Norwegian Centre for Research Data, European Social Survey cumulative file, ESS 1-9. Data archive and distributor for ESS data for ESS ERIC Edition 1.0. (2020). <https://doi.org/10.21338/NSD-ESS-CUMULATIVE>. Accessed 23 March 2021.
- International Federation of Robotics, World Robotics: Industrial Robots and Service Robots. <https://ifr.org/worldrobotics/>. Accessed 14 February 2018.
- US Bureau of Labor Statistics, Producer Price Index by Industry: Computer and Electronic Product Manufacturing [PCU334334], retrieved from FRED, Federal Reserve Bank of St. Louis. <https://fred.stlouisfed.org/series/PCU334334>. Accessed 1 December 2019.
- K. Rupp, Data from "42 Years of microprocessor trend data." GitHub. <https://github.com/karlrupp/microprocessor-trend-data>. Accessed 12 November 2021.
- A. Volkens et al., The manifesto data collection. Manifesto Project (Manifesto Research Group / Comparative Manifestos Project / Manifesto Research on Political Representation [MRG/CMP/MARPOR]). Version 2018a. WZB, Berlin (2018). <https://manifesto-project.wzb.eu/>. Accessed 12 November 2021.
- A. Gallego, A. Kuo, D. Manzano, J. Fernandez-Albertos, Technological risk and policy preferences. *Comp. Polit. Stud.*, 10.1177/00104140211024290 (2021).
- A. S. Blinder, A. B. Krueger, Alternative measures of offshorability: A survey approach. *J. Labor Econ.* **31**, S97–S128 (2013).
- M. Goos, A. Manning, A. Salomons, Explaining job polarization: Routine-biased technological change and offshoring. *Am. Econ. Rev.* **104**, 2509–2526 (2014).