European Journal of Political Research (2026), 1–12
doi:10.1017/S1475676526101315

RESEARCH NOTE

The (forgotten) atomistic fallacy in political science and its implications for how we interpret elections

Tim Vlandas[1]* and Daphne Halikiopoulou[2]*

[1]University of Oxford, UK and [2]University of York, UK
Corresponding author: Tim Vlandas; Email: tim.vlandas@spi.ox.ac.uk

(Received 2 December 2025; revised 2 February 2026; accepted 13 March 2026)

Abstract

Improvements in the availability, accuracy, and processing of individual-level data have allowed political science literature to address the 'ecological fallacy', whereby inferences are made about individuals based on units of analyses operating at a higher level. Yet there has been limited attention to the risk that individual-level analyses may suffer from the reverse 'atomistic' – or 'individualistic' – fallacy: the erroneous practice of drawing inferences about national-level outcomes based on individual-level analyses. In this research note, we present a mathematical statement and simulations to diagnose and evaluate the extent of this fallacy in the case of voting behaviour. We also illustrate the problem using European Social Survey data on far-right voting. We conclude by identifying three 'perils' of the atomistic fallacy, related to extrapolating conclusions about a party's overall performance from information about an individual's voting propensity. These perils can significantly affect how researchers interpret election results and, in turn, the policy implications of political science research.

Keywords: atomistic/individualistic fallacy; ecological fallacy; voting behaviour; European Social Survey; far-right

Introduction

In the past decades, high-quality, granular data have become readily accessible, while greater computing power and methodological advances have enabled researchers to analyse this data faster and more effectively. Political scientists increasingly rely on individual-level data, paying greater attention to the magnitude of individual-level effects and developing new methods to estimate more internally valid causal effects (see Scheiring, Serrano-Alarcón, Moise et al. 2024). Overall, these developments have led to significant shifts in best practice over time. In particular, voting behaviour scholarship has leveraged better data and new empirical tools to address several methodological issues, most notably the ecological fallacy (Kramer 1983; Erbring 1989; Russo 2017). The latter refers to the risk associated with making inferences about individual behaviour from findings using aggregate-level data, which was problematic in older analyses of voting behaviour relying on aggregate-level voting records. In doing so, newer studies focusing on the objective and attitudinal characteristics that drive political behaviour have significantly advanced our understanding of electoral patterns.

However, methodological issues related to the conflation of different levels of analyses remain and have received renewed attention in recent empirical debates in political science. For example, Marks,

*Both authors contributed equally to this research note and are listed according to a principle of rotation.

Attewell, Hooghe et al. (2023) highlight the importance of the relative weight of social-structural and behavioural factors for party composition. Using the example of the Socialists, they document a 'voter-party paradox': a party's electorate may be composed predominantly of a specific group, but this does not mean that most people in the group actually vote for this party. Indeed, an individual may have a relatively low probability of voting for a specific party, but the group of individuals with this characteristic may make up a large percentage of the party's electorate. Equally, an individual may have a high probability to vote for a party but the group of individuals with this characteristic may make up a small percentage of the party's electorate. A similar problem arises concerning misconceptions about the composition of political parties, as there is often a tendency to overestimate the extent to which party supporters belong to party-stereotypical groups (Ahler and Sood 2018).

Moreover, there is a tendency to conflate changes in voter attributes with changes in voter priorities (Danieli, Gidron, Kikuchi et al. 2022). Dennison and Geddes (2019) argue this has led to the misconception that a rising tide of anti-immigration attitudes has fuelled far-right party success in Europe, whereas they show that the salience of immigration has increased amongst a shrinking number of Europeans with existing anti-immigration sentiments. What matters more to the success of political parties, therefore, is not necessarily an increase in an individual attitude towards an issue, but rather an increase in the salience of this issue, which could be affecting only a small segment of the population.

While this emerging literature discusses somewhat empirically distinct issues, we argue that they capture different dimensions of a common underlying methodological problem, that is, the 'atomistic' – or 'individualistic' – fallacy[1]: the problem associated with (implicitly or explicitly) drawing conclusions about aggregate-level phenomena of interest using individual-level data (Seligson 2002; Subramanian, Jones, Kaddour et al. 2009). In the case of voting behaviour literature where this risk is especially acute, the fallacy entails extrapolating conclusions about what factors affect a party's overall performance from information about an individual's voting propensity. The atomistic fallacy is underpinned by the same logic as the ecological fallacy, namely the erroneous practice of making inferences about one level of analysis by using data from another. Therefore, it can have similarly problematic implications for how we theorise and test key claims about the relevance of different factors. Yet, it has hardly received any attention in the voting behaviour literature.

To remedy this gap, this research note examines the atomistic fallacy and illustrates how overlooking this methodological issue affects our conclusions about electoral outcomes. Our contribution is threefold. First, by presenting a formal statement of the atomistic fallacy, we offer a systematic diagnosis of this problem. This allows researchers to specify the circumstances under which this problem may occur and thus to remedy it. Second, our research note has significant implications for how we assess and theorise the ways in which different factors contribute to a political outcome, for example elections. This can have major implications for our understanding of far-right party success, the Brexit vote, democratic backsliding, and voter turnout. Third, our research can improve the recommendations of policy-relevant research. By showing how interpretations of election results change once the atomistic fallacy is addressed, for example by accounting for the relative size of voter groups within coalitions, our research can help change policy recommendations and party strategies about which constituencies to target and why.

We proceed as follows. First, we define and offer a hypothetical illustration of the atomistic fallacy. Second, we present a mathematical statement and generate a series of simulations using hypothetical data to illustrate how generalising results about individual-level behaviour onto the national level can lead to misleading conclusions. Third, we illustrate the problem using European Social Survey (ESS) data on far-right voting. Fourth, we identify three 'perils' associated with the atomistic fallacy, and finally, we conclude by proposing avenues for future research.

[1]Several terms are used in the literature to describe this methodological problem including 'atomistic fallacy', 'individualistic fallacy', and 'fallacy of composition' (Subramanian et al. 2009; Halikiopoulou and Vlandas 2020). In this research note, we predominantly use the term 'atomistic fallacy'.

https://doi.org/10.1017/S1475676526101315 Published online by Cambridge University Press

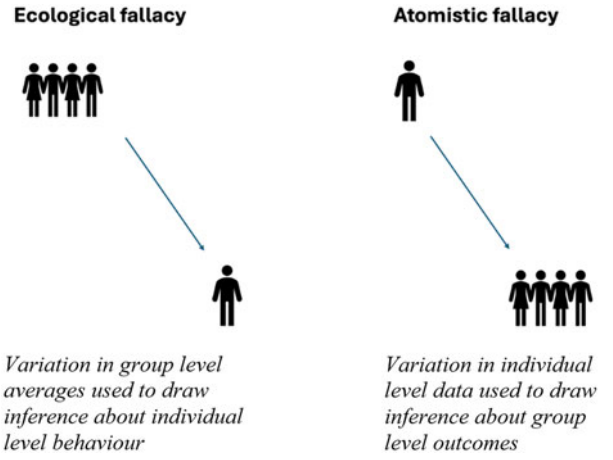

Figure 1. Illustration of the ecological and atomistic fallacies.

A hypothetical illustration of the atomistic fallacy

Both the ecological and atomistic fallacies stem from extrapolating from one level of analysis to another. The ecological fallacy arises from making inferences about individual behaviour from aggregate-level data. By contrast, the atomistic fallacy is the reverse, or in other words, the problem of drawing conclusions about macro-phenomena from findings using individual-level data. Also known as the compositional or individualistic fallacy, it constitutes flawed reasoning for the same reasons as the ecological fallacy (Seligson 2002; Subramanian, Jones, Kaddour et al. 2009): they both emanate from discrepancies between the macro- and micro-levels of analysis resulting in the invalid transfer of results from one level to another (Selvin 1958). The formulation of inferences at one level, based on analyses performed at a different level, is problematic because 'relationships among variables that hold at one level do not necessarily hold at another level of the hierarchy' (Croon and Veldhoven 2007: 45; see also Hox 2010). Indeed, the same data are likely to yield differing results depending on whether it is analysed at the individual or the aggregate level (Robinson 1950; Erbring 1989; Subramanian, Jones, Kaddour et al. 2009).

Numerous studies, as early as the 1950s, have highlighted the perils associated with the ecological fallacy in political science (Robinson 1950; Przeworski and Teune 1970; Kramer 1983), cautioning against the potential biases resulting from this practice. While, however, the flawed reasoning that underpins the ecological fallacy is also present for the atomistic fallacy, the latter is much less discussed, and this tends to also be in other disciplines such as psychology, epidemiology and economics (eg Subramanian, Jones, Kaddour et al. 2009), whereas it is largely overlooked in political science literature, or discussed indirectly without being integrated into a unified framework (for exceptions see: Seligson 2002; Halikiopoulou and Vlandas 2020; Grimmer, Marble and Tanigawa-Lau 2025).

Figure 1 below visualises the reasoning accompanied by the ecological (left panel) and atomistic (right panel) fallacies, respectively. Recall that the ecological fallacy is the problem of using aggregate-level data to draw inferences about individual-level behaviour. An example of this type of reasoning applied to voting behaviour would be the following: if characteristic A at the aggregate level, for instance high unemployment or migration rates, is associated with voting for party X; therefore, people who exhibit characteristic A at the individual level – to stay with the same example, unemployed individuals, or individuals with anti-immigration attitudes – are more likely to vote for party X. This reasoning has often been cautioned against with numerous studies

https://doi.org/10.1017/S1475676526101315 Published online by Cambridge University Press

showing that aggregate-level findings about the relationship between economic conditions and election outcomes do not operate at the individual-voter level (this was debated in the American Political Science Review as early as the 1980s – eg Kramer 1983).

By contrast, the atomistic fallacy is the reverse problem, that is, using variation in individual-level data to draw inferences about group level outcomes. An example of this type of reasoning would be the following: people who exhibit characteristic A, for example unemployed or anti-immigrant individuals, are more likely to vote for party X (individual level); therefore, characteristic A, such as unemployment or migration rates, explains the success of party X at the aggregate level. This is problematic because it is based on the erroneous assumption that the individual-level predictive power of a characteristic implies importance at the national level, whereas, as we show later, this crucially depends on how widespread this characteristic actually is in the country's electorate.

We can illustrate this problem with a simple example. Suppose an electorate consists of 110 voters: 10 of these voters exhibit characteristic A, while the remaining 100 exhibit characteristic B. Suppose further that in the last election, 5 out of 10 people with characteristic A voted for party X, so that they have a 50% probability of voting for this party. By contrast, 10 out of 100 with characteristic B voted for party X, so that they have a 10% probability of voting for this party. In this example, a voter with characteristic A is five times more likely than a voter with characteristic B to vote for party X. However, people with characteristic B are more important to the success of party X because of their numerical majority, even though having characteristic A has a stronger effect on individual party support than having characteristic B. In this hypothetical electorate, it may well be that the traditional core voters of party X are voters with characteristic A, but it is nevertheless the voters with characteristic B that allow these parties to expand beyond this core to increase their electoral chances.

Formal statement and simulations

We proceed to illustrate the risks associated with the atomistic fallacy by presenting a formal mathematical statement and simulations. Suppose we know the (true) data generation process leading individuals to vote for party X and that it is a function f(.) of a particular characteristic A (with $A \in [0, 1]$), such that: $f(A) = \alpha + \beta A$; where for example $f(A) = 0.02 + 0.6A$.[2] In other words, individuals who do not exhibit characteristic A have a 2% probability of voting for party X, and as the likelihood of exhibiting characteristic A increases from 0 to 1, this probability reaches a maximum of 62% when $A = 1$. Given that in this hypothetical example, individual voting propensity is entirely a function of characteristic A, the national-level score is then determined by the distribution of this characteristic in the electorate.

Moreover, let us assume that the distribution of characteristic A has the following Beta probability density function[3]: $p(A) = \text{BetaPDF}(A; \gamma, \delta)$, defined for $A \in [0,1]$, where γ and δ are the shape parameters of the distribution of characteristic A with mean $= \gamma/(\gamma + \delta)$. In Figure 2, we show how these two parameters characterise the shape and central tendency of characteristic A for different electorates: an electorate where characteristic A is relatively low and concentrated (Beta(2, 10)); an electorate where characteristic A is widespread (Beta(3, 5)); and an electorate with a high share of individuals who exhibit characteristic A (Beta(3, 5)).

For any given individual-level voting function f(A) and the associated distribution of the underlying characteristic (in this case characteristic A), the expected *national-level* party voting

[2]Although we illustrate the statement with a functional form that is linear and additive, and has only one variable A, our wider point does not depend on the number of variables, linearity, or having additive or multiplicative models.

[3]Our conclusion does not depend on which probability function we choose, since our point is simply to show that the distribution of a characteristic matters to the national level outcome arising from a particular individual-level data generation process.

https://doi.org/10.1017/S1475676526101315 Published online by Cambridge University Press

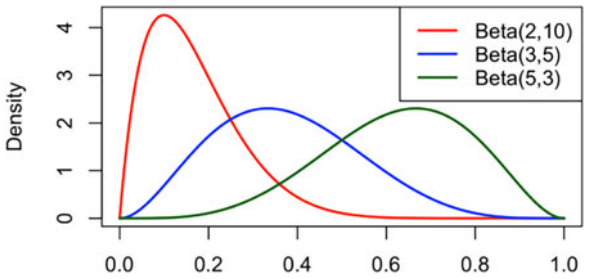

Figure 2. Beta probability distributions of characteristic A.

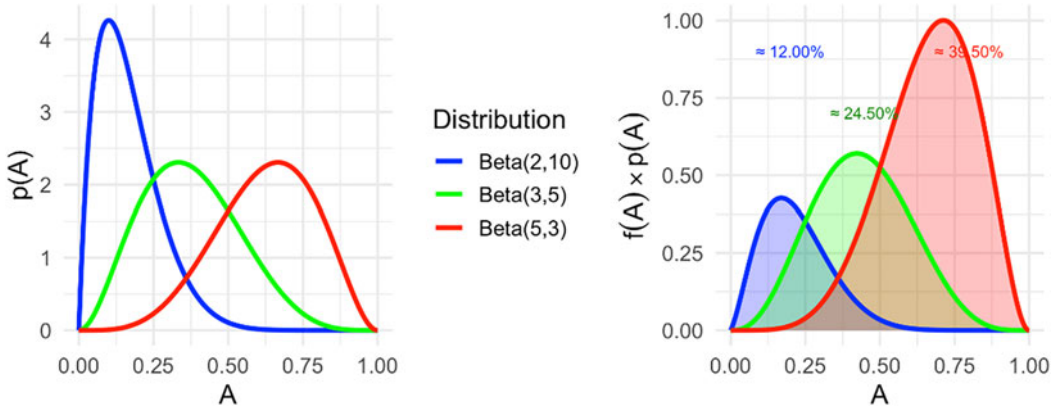

Figure 3. Relationship between the distributions of a characteristic and the overall party vote share.
Note: This figure presents different distributions of characteristic A (left panel) multiplied by the individual voting function (right panel). The values at the top of the distributions (right panel) represent the party's vote share (i.e. the integral of the product between the individual-level voting function f(A) and the BetaPDF function).

share V – which is not to be confused with the *individual-level* voting function f(A) – is calculated as the integral of the product of the individual-level function f(A) and the probability distribution function p(A):

$$V = \int_0^1 f(A) \cdot p(A) dA$$

Solving for our function f(A), we get: $E[\int(\alpha + \beta A) \cdot \text{Beta}(A; \gamma, \delta) dA)] = E[\alpha \int \text{Beta}(A; \gamma, \delta) dA + \beta \int A \text{Beta}(A; \gamma, \delta) dA] = \alpha + \beta[\gamma/(\delta + \gamma)]$. Thus, for $f(A) = 0.02 + 0.6A$ and BetaPDF(2, 10), the resulting national-level vote share $V = 0.02 + 0.6[2/(10 + 2)] = 0.02 + 0.6(1/6) = 0.12$. Figure 3 shows the national-level results for three hypothetical probability density functions: BetaPDF(2,10); BetaPDF(3,5); and BetaPDF(5,3). Each distribution shown in the left panel is multiplied by the function f(A) resulting in a new national-level vote share distribution shown in the right panel, which is then integrated to get the national-level vote share, shown in the numbers placed above the curves on right-hand side. The figure clearly illustrates the large effects that changes in the distribution of a characteristic can have on the national-level voting share (rising from 12% to 40%) even when holding the magnitude of the effect of that characteristic constant.

How large is the problem likely to be for plausible parameters capturing the effect size of an individual-level characteristic and its national-level distribution? Assuming only one factor determines individual votes, the national party vote share is the integral of the individual voting

https://doi.org/10.1017/S1475676526101315 Published online by Cambridge University Press

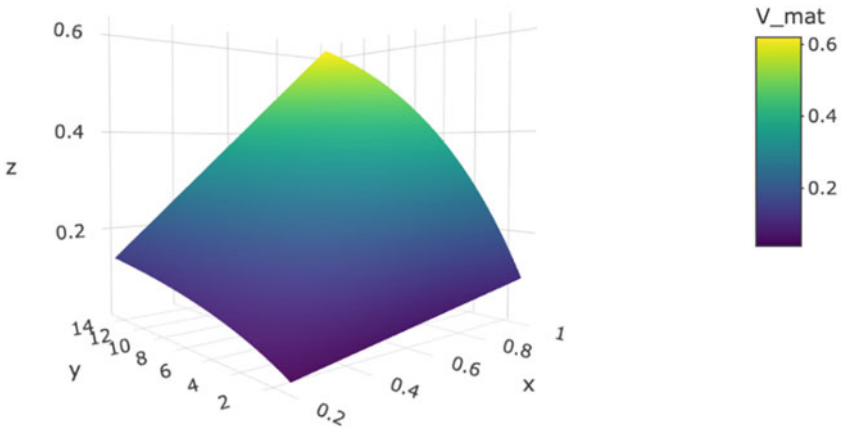

Figure 4. Simulations of election results.

function multiplied by the probability distribution of the characteristic in question. Formally, this can be represented by the following equation:

$$V(A) = \int_0^1 (\alpha + \beta A)\text{Beta}(A; \gamma, \delta)dA = \alpha + \beta\frac{\gamma}{\gamma + \delta}$$

where α is the intercept of individual voting function, β is the coefficient capturing the effect size, while **γ** and **δ** determine the shape and central tendency of the distribution. Holding the intercept $\alpha = 0.2$ and one of the distribution parameters $\delta = 10$, Figure 4 shows how the national vote share (shown on z-axis in the figure) is a function of two factors: (1) β, which in this example varies between 0.2 and 1 (shown on the x-axis in the figure); and (2) the distribution of A, which varies depending on the value of γ, set between 1 and 15 (shown on the y-axis in the figure), that is, we progressively increase the mean value of characteristic A in the population – or in other words its prevalence. The figure makes clear that the size of the individual-level effect of a characteristic only matters to the national-level performance outcome if a sufficiently large share of the electorate has that characteristic. As a result, an increase in the prevalence of characteristic A over time could lead to higher national-level vote share without any change in coefficient, and conversely, an increase in the effect of A may (or may not) lead to higher national scores, depending on how the distribution of the characteristic has evolved over time.

It is common practice among political scientists to compare the relative importance of two different factors. Here again, there has been a (welcome) move over time from a focus on statistical significance to one on magnitude (and causal inference). Our formal statement of the atomistic fallacy can be straightforwardly generalised to a scenario of two individual-level factors A and B, where individual-level effects are captured by f(A) and f(B). For example, one could assume that $f(A) = 0.1 + 0.4A$ and $f(B) = 0.1 + 0.3B$; such that the effect of B is 25% less than the effect of A.[4] Suppose further that each characteristic has a distribution characterised by its respective BetaPDF. Combining both the functions linking a characteristic to party voting and the distribution of the characteristics, we can calculate the independent contribution of A and B on the overall national vote share V in the following way:

$$V = V_A + V_B$$

[4]For simplicity, two separate individual-level functions are assumed with completely independent distributions. By contrast, in many empirical applications, coefficients and distributions are often jointly determined. Although our assumption is to some extent implausible in reality, it facilitates the exposition. Relaxing this assumption does not change our conclusions.

https://doi.org/10.1017/S1475676526101315 Published online by Cambridge University Press

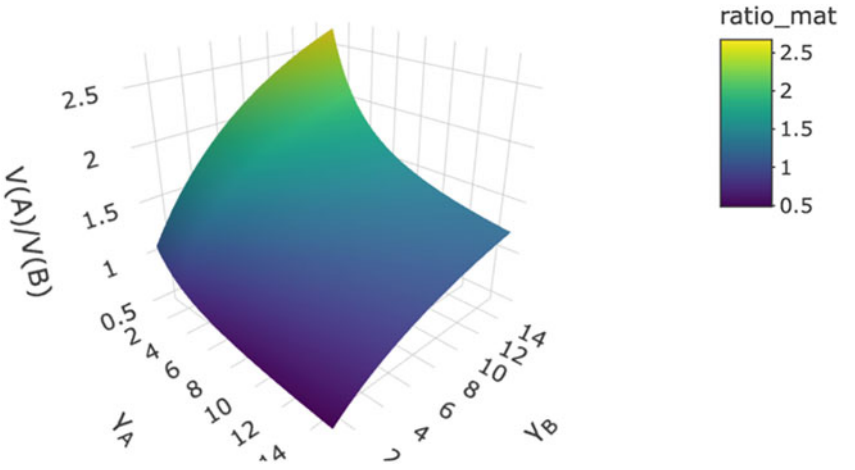

Figure 5. Contribution of each factor to the national party vote.

where:

$$V_A = \int_0^1 f(A) \cdot p(A) dA$$

$$V_B = \int_0^1 f(B) \cdot p(B) dB$$

Next, we calculate the ratio of the contribution of A to B: where this ratio is superior to 1, the factor A contributes more to the overall national party vote; and where it is under 1, then it is factor B that matters relatively more. Supposing that the distributions of each characteristic are captured by BetaPDF(A, $\boldsymbol{\gamma}_A$, $\delta = 10$) and BetaPDF(B, $\boldsymbol{\gamma}_B$, $\delta = 10$), then the ratio of the contribution of A and B to total vote share can be calculated as:

$$\frac{V(A)}{V(B)} = \frac{\int_0^1 (0.1 + 0.4A)\,\mathrm{Beta}(A; \gamma_A, 10)\, dA}{\int_0^1 (0.1 + 0.3B)\,\mathrm{Beta}(B; \gamma_B, 10)\, dB}$$

Figure 5 shows that even with an effect size that is 25% larger, the contribution of A to the national vote is sometimes lower than B (captured by a ratio < 1) depending on its underlying distribution relative to the underlying distribution of B. Thus, comparing the magnitudes of two coefficients does not automatically allow for drawing valid conclusions about which factor matters most to national party success. This suggests that political science researchers interested in assessing the relative macro-level importance of two factors on the basis of analyses of micro-level results must be cautious to also check their distribution and to then calculate the appropriate combined effects of each. In the next section, we illustrate that this is not merely a theoretical possibility but also a problem that can materialise when using survey data.

An empirical example of the atomistic fallacy: interpreting far-right party success

Our mathematical statement and simulations demonstrate how the atomistic fallacy in political science can result in the erroneous conclusion that individual-level predictive power implies substantive importance. This problem can be empirically illustrated in the case of far-right party success. We focus on electoral support for the far-right for two reasons. First, far-right parties have significantly increased their electoral presence in recent years in Europe and beyond, also often participating in government (de Jonge et al. 2025). This has led to a rapid proliferation of studies on the far-right in the past few years (Rooduijn, Pirro, Halikiopoulou et al. 2024). Providing

https://doi.org/10.1017/S1475676526101315 Published online by Cambridge University Press

concrete data, Arzheimer's (2024) extreme right bibliography records a total of 1511 titles, with 222 new titles in 2024 alone – the biggest update to this bibliography and more than twice the size of the 2023 update which increased the total number of titles by some 18%. Second, the rise of the far-right has significant implications for democratic politics. Scholars agree that the rise of far-right parties and their entrenchment is associated with a broad trend towards democratic backsliding (Bermeo 2016) and constitutes a threat to democracy (Mudde 2019). As such, far-right research is often policy-relevant and can have a significant impact on party strategy design and formulation.

Despite this substantive importance, however, the risks associated with the atomistic fallacy remain largely overlooked. The problem is particularly acute within the quantitative electoral behaviour tradition in political science, which tends to focus more on micro-level processes and less on structural and macro-level explanations. It also has significant public-facing implications as overlooking the atomistic fallacy can lead to misinterpretations about far-right party success and flawed policy recommendations about how to address it. Much of the research on the far-right is on voter attitudes and uses individual-level data. For example, according to Rooduijn, Pirro Andrea, Halikiopoulou et al. (2024), a large percentage of items that use the PopuList classification do so in conjunction with other datasets such as the ESS to identify the individual-level drivers of far-right party success.

While this has significantly advanced our understanding of voting patterns and has helped address the risks associated with the ecological fallacy, it runs the reverse risk of drawing implicit or explicit conclusions about the drivers of the phenomenon at the national level from individual-level data. This is problematic because, as we argue in the next paragraphs, the most important factor shaping individual vote choices does not automatically tell us what 'explains' the party's overall performance in elections.

A good example of this conundrum may be observed in debates about the role of immigration attitudes, which have received significant attention in the literature (eg Ivarsflaten 2008), both with regard to the formation of these attitudes themselves, which can be partly linked to economic hardship, and their role in driving far-right party success. According to much of this literature, anti-immigration attitudes are mostly about cultural concerns, and far-right party success, in turn, depends largely on mobilising these cultural grievances over immigration (Malhotra, Margalit and Mo 2013). In these accounts, far-right party success can be largely attributed to cultural concerns over immigration.

However, recent work has indirectly touched upon how discrepancies between individual and national-level dynamics, as well as overlooking group size, may be affecting our conclusions about the role of immigration attitudes. For example, Dennison and Geddes (2019) argue that it is erroneous to assume that the rise of the far-right is driven by an overall increase in anti-immigration sentiment. Instead, they contend that aspects of immigration in the last decade have activated pre-existing opposition to immigration amongst a shrinking segment of the population. Danieli, Gidron, Kikuchi et al. (2022) also point out that it is not changes in voter attributes that account for the recent success of far-right parties, but rather voters' changing priorities. In other research, we compare the effects of cultural and economic concerns over immigration on far-right party support (Halikiopoulou and Vlandas 2020). While we find that both types of attitudes drive far-right party support, we also show that individuals with cultural concerns over immigration do not necessarily matter more for far-right party success at the national level in substantive terms.

Similar debates concern the role of other factors in driving far-right party success, for instance downward social mobility (Gugushvili, Halikiopoulou, and Vlandas 2025): despite their propensity to support the far-right, downwardly mobile people constitute a small proportion of the far-right electorate. Thus, this social group is not one of the main success drivers of the far-right and should rather be seen as part of the often larger set of voters with economic grievances that make it possible for far-right parties to broaden their electoral base beyond their core supporters.

https://doi.org/10.1017/S1475676526101315 Published online by Cambridge University Press

In the online Appendix Section B, we report results from analyses using data from the ESS which has been used extensively by previous research on the far-right (Rooduijn, Pirro, Halikiopoulou et al. 2024). We combine eleven waves of the European Social Survey in 25 European countries and focus on the effects of age, gender, and immigration concerns. Our purpose is not to carry out a systematic empirical analysis of all the drivers of far-right party voting, but instead to illustrate how the problem plays out in a concrete empirical case. Given that the core logic of the atomistic fallacy does not hinge on linearity, additive effects, or the absence of interactions, our conclusions are not dependent on the estimation strategy or model specification. Therefore, we do not carry out checks with such alternative specifications. Results from all three sets of analyses suggest that the magnitude of the individual-level coefficient does not capture how strongly each variable (age, gender, and immigration attitudes) predicts who votes for a far-right party. While older male individuals with strong immigration concerns are more likely to vote for a far-right party, it would be erroneous to conclude that such individuals account for the success of these parties, or that the magnitude of the coefficient in different waves captures the relative importance of this variable in the national-level far-right party score.

Discussion: the three perils of the atomistic fallacy

Our analyses highlight three related but distinct 'perils' associated with the atomistic fallacy. First, it is problematic to use the magnitude of an individual-level coefficient to infer its substantive importance for national-level dynamics. As we show in Section B of our Appendix, it is not because being an older male makes one more likely to vote for the far-right that older male individuals account for these parties' success. Similarly, it is not because holding strong cultural concerns over immigration makes one more likely to vote for the far-right, that cultural concerns over immigration 'explain' far-right party national level success. This illustrates intuitively the main point of our formal statement and simulations that the effect of a variable and its distribution must both be considered to draw any conclusions about its national-level importance.

The second peril concerns comparing magnitudes of individual-level coefficients of different independent variables to infer their relative substantive importance for national-level dynamics. Take the example of far-right voters with cultural and economic concerns over immigration. It is not because individuals with cultural concerns over immigration are more likely to vote for the far-right than individuals with economic concerns over immigration, that the former group matters 'more' for far-right party success than the latter (cf. Halikiopoulou and Vlandas 2020). Beyond these distributional considerations arising from the atomistic fallacy, comparisons of coefficient magnitudes can be further complicated by analytically separate issues of causal ordering and mediation. For example, a growing body of work suggests that cultural attitudes may function as mediators of underlying material grievances rather than as independent drivers of political behaviour (Halikiopoulou and Vlandas 2020). In other words, the distribution of one factor (B in the formal statement provided earlier) could itself be determined by another factor (A in our previous formal statement) even if B has a larger effect. This implies that associations between cultural attitudes and far-right support cannot be straightforwardly interpreted as evidence of greater substantive importance, even when aggregation issues are sufficiently considered to avoid the atomistic fallacy. While our framework does not offer a direct solution for these issues, it highlights that inferences based on individual-level coefficients can be further complicated when multiple, potentially mediated mechanisms are at play.

Third, one cannot automatically infer from differences in the magnitude of an individual-level coefficient (over time/across countries) their relative importance for national level over time dynamics. For example, as we showed in the case of older male respondents, the changing magnitude of this coefficient across different waves does not necessarily help us make sense of the variation in far-right party success at the national level over time. Conversely, the increase in the magnitude of the coefficient capturing the effects of anti-immigration attitudes over time does not

https://doi.org/10.1017/S1475676526101315 Published online by Cambridge University Press

imply that anti-immigration attitudes 'explain' the rise of the far-right over time. In fact, the predictive power of these attitudes (in correctly predicting whether an individual votes or not for the far-right in a given electorate) has fallen over time (Section B in Appendix).

This third peril highlighted here applies to the specific cases where dynamic claims about aggregate change over time are inferred directly from over time changes in individual-level coefficients. It should not be conflated with the problem of inferring who drives change over time from static information about who supports a party at a given moment. Arguing that anti-immigrant individuals are more likely to vote for far-right parties is not the same as arguing that anti-immigrant individuals account for the *increase* in support for these parties. This is a distinct within-level inferential error that could be problematic even in the absence of the atomistic fallacy where the micro-macro leap is the issue.

Conclusion

Our research incorporates several methodological concerns raised in recent literature into an integrated framework. For instance, Marks, Attewell, Hooghe et al. (2023) have pointed out that an individual's probability to vote for a specific party tells us little about the makeup of this party's electorate. Often, individuals most likely to vote for a party constitute a small part of its (core) electorate (Halikiopoulou and Vlandas 2020). As a result, tipping a party to electoral victory may be less about changes in political preferences overall, but more about an increase in the salience of an issue among a 'loud minority' (Dennison and Geddes 2019). By offering a formal diagnosis of the problem and a mathematical statement as well as simulations of its likely prevalence, this research note allows researchers to identify and assess whether the atomistic fallacy is likely to alter the conclusions of their research.

We suggest several potential ways forward to minimise the atomistic fallacy. First, researchers should pay much more attention to the structure of their dataset and descriptive statistics, most notably concerning the distribution of the independent variables they are interested in. In the case of voting behaviour research, the size and composition of different parts of the electorate appear crucial in understanding how a given individual-level effect 'scales up' to a national-level outcome. With the (welcome) move to more sophisticated estimation strategies, the structure of the underlying data is too often neglected. Second, researchers could implement disaggregation techniques. In this research note, we show that the national-level outcome is the integral of the product of a characteristic and its distribution. This approach could be adopted to combine regression coefficients of key voter characteristics with their underlying distribution to decompose their national-level implications. Third, simulations can help us capture the macro implications of micro-level findings. Researchers could vary the effect size and distribution of characteristics to simulate how counterfactual scenarios would lead to different national-level outcomes.

Moreover, novel methodological tools could further help disaggregate the different levels of analysis. For example, Marks, Attewell, Hooghe et al. (2023) analysis offers a new way of measuring social structuration that allows researchers to assess the extent to which a political party is socially rooted in a particular constituency by estimating the overrepresentation or underrepresentation of a social group in a party relative to its size in the society. In addition, Grimmer, Marble and Tanigawa-Lau's (2025) analysis of US election outcomes provides a set of statistical tools enabling researchers to measure the contribution of voting blocs to a candidate's vote total and to assess how their contributions change from one election to another. Future research should use these approaches to better gauge how the effect and distribution of individual-level variables scale up to the national level.

In sum, one key takeaway of our research note is that voting behaviour scholarship should pay more attention to discrepancies arising from extrapolating from the individual to the aggregate levels. Another concerns the public-facing implications of the atomistic fallacy and the ways in which research findings are communicated beyond academia. A reductionist reading of a

https://doi.org/10.1017/S1475676526101315 Published online by Cambridge University Press

correlation between variables can obscure the public debate and lead to flawed policy recommendations. To avoid misleading conclusions about how support from different voter groups translates into electoral success, researchers should interpret their findings in ways that account for the composition of electorates and the distribution of various characteristics. At a time of widespread far-right party success, and democratic decline more generally, it is crucial not only to correctly diagnose what drives these political phenomena but also to communicate research findings clearly beyond academia.

Supplementary material. To view supplementary material for this article, please visit https://doi.org/10.1017/S1475676526101315.

Data availability statement. Data and added supplementary material are available in the online Appendix.

Acknowledgements. We are grateful for the feedback we received during presentations at the CES 2023, CES 2024, and IPSA 2023 conferences. For discussions, suggestions, and comments on earlier versions, we would like to thank Tarik Abou Chadi, Cyril Benoit, Diane Bolet, Björn Bremer, Marius Busemeyer, Tom Chevalier, Juliana Chueri, Elias Dinas, Thalia Gerszo, Juho Härkönen, Peter Hall, Anton Hemerijck, Simon Hix, Alison Johnston, Filip Kostelka, Margaryta Klymak, Herbert Kitschelt, Paul Marx, Reto Mitteregger, Sascha Riaz, David Rueda, Alexander Sahn, Waltraud Schelkle, Sarah Shair-Rosenfield, Stuart Turnbull-Dugarte, Vicente Valentim, Herman Van De Werfhorst, Clara Weissenfels, and Delia Zollinger. Finally, we thank the anonymous reviewers of the EJPR for their constructive comments and the Editors for the smooth review process.

Financial support. N.A.

Competing interests. N.A.

Ethical standards. N.A.

References

Ahler, D.J. and Sood, G. (2018). 'The parties in our heads: misperceptions about party composition and their consequences'. *Journal of Politics*, **80**(3), 964–981.

Arzheimer, K. (2024). 222 new pieces of research on the radical right, available at https://www.kai-arzheimer.com/222-new-pieces-of-research-on-the-radical-right/, accessed 7 February 2025.

Bermeo, N. (2016). 'On democratic backsliding'. *Journal of Democracy*, **27**(1), 5–19.

Croon, M.A. and van Veldhoven, M.J. (2007). 'Predicting group-level outcome variables from variables measured at the individual level: a latent variable multilevel model'. *Psychological Methods*, **12**(1), 45–57.

Danieli, O., Gidron, N., Kikuchi, S. and Levy, R.E. (2022). Decomposing the rise of the populist radical right, available at https://doi.org/10.2139/ssrn.4255937, accessed 19 May 2026.

de Jonge, L., Georgiadou, V., Halikiopoulou, D., Kaltwasser, C. R. and Tanscheit, T. (2025). 'Is the far right a global phenomenon? Comparing Europe and Latin America: A scholarly exchange'. *Nations and Nationalism*, **31**(1), 7–24. https://doi.org/10.1111/nana.13074

Dennison, J. and Geddes, A. (2019). 'A rising tide? The salience of immigration and the rise of anti-immigration political parties in western Europe'. *The Political Quarterly*, **90**, 107–116.

Erbring, L. (1989). 'Individuals writ large: an epilogue on the ecological fallacy'. *Political Analysis*, **1**, 235–269.

Grimmer, J., Marble, W. and Tanigawa-Lau, C. (2025). 'Measuring the contribution of voting blocs to election outcomes.' *Journal of Politics*, **87**(3), 905–921.

Gugushvili, A., Halikiopoulou, D. and Vlandas, T. (2025). 'Downward class mobility and far-right party support in Western Europe', *Political Behavior*, **48**(1), 349–380.

Halikiopoulou, D. and Vlandas, T. (2020). 'When economic and cultural interests align: the anti-immigration voter coalitions driving far right party success in Europe'. *European Political Science Review*, **12**(4), 427–448.

Hox, J. (2010). *Multilevel Analysis: Techniques and Application*. Routledge.

Ivarsflaten, E. (2008). 'What unites right-wing populists in Western Europe? Re-examining grievance mobilization models in seven successful cases'. *Comparative Political Studies*, **41**(1), 3–23.

Kramer Gerald, H. (1983). 'The ecological fallacy revisited: aggregate- versus individual-level findings on economics and elections, and sociotropic voting'. *American Political Science Review*, **77**(1), 92–111.

Malhotra, N., Margalit, Y. and Mo, C.H. (2013). 'Economic explanations for opposition to immigration: distinguishing between prevalence and conditional impact'. *American Journal of Political Science*, **57**, 391–410.

Marks, G., Attewell, D., Hooghe, L., Rovny, J. and Steenbergen, M. (2023). 'The social bases of political parties: a new measure and survey'. *British Journal of Political Science*, **53**(1): 249–260.

Mudde, C. (2019). *The Far Right Today*. Polity Press.

Przeworski, A. and Teune, H. (1970). *The Logic of Comparative Social Inquiry*. New Interscience. https://doi.org/10.2307/1958372

Robinson, W.S. (1950). 'Ecological correlations and the behavior of individuals'. *American Sociological Review*, **15**(3), 351–357.

Rooduijn, M., Pirro Andrea, L.P., Halikiopoulou, D., Froio, C., Van Kessel, S., De Lange, S.L., Mudde, C. and Taggart, P. (2024). 'The populist: a database of populist, far left, and far right parties using expert-informed qualitative comparative classification.' *British Journal of Political Science*, **54**(3), 969–978. https://doi.org/10.1017/S0007123423000431

Russo, L. (2017). 'The use of aggregate data in the study of voting behaviour: ecological inference, ecological fallacy and other applications', in J. Fisher, E. Fieldhouse, M. Franklin, R. Gibson, M. Cantijoch and C. Wlezien (Eds.), *The Routledge Handbook of Elections, Voting Behavior and Public Opinion* (pp. 485–495). Routledge.

Scheiring, G., Serrano-Alarcón, M., Moise, A., McNamara, C. and Stuckler, D. (2024). 'The populist backlash against globalization: a meta-analysis of the causal evidence'. *British Journal of Political Science*, **54**(3), 892–916. https://doi.org/10.1017/S0007123424000024

Seligson, M.A. (2002). 'The renaissance of political culture or the renaissance of the ecological fallacy?' *Comparative Politics*, **34**(3), 273–292.

Selvin, H.C. (1958). 'Durkheim's suicide and problems of empirical research'. *American Journal of Sociology*, **63**(6), 607–619.

Subramanian, S.V., Jones, K., Kaddour, A. and Krieger, N. (2009). 'Revisiting Robinson: the perils of individualistic and ecologic fallacy'. *International Journal of Epidemiology*, **38**(2), 342–360.

Cite this article: Vlandas T and Halikiopoulou D (2026). The (forgotten) atomistic fallacy in political science and its implications for how we interpret elections. *European Journal of Political Research*, 1–12. https://doi.org/10.1017/S1475676526101315