Blowing in the wind: COVID and political support in Brazil*

Rafael Pintro Schmitt[1]

[1]Bocconi University

November 22, 2023

Abstract

I investigate the impact of city-level excess mortality on former Brazilian president Jair Bolsonaro's electoral performance in 2022. First, I provide evidence of a robust relationship between wind and COVID-induced excess mortality. In particular, cities with relatively high wind speeds during COVID-19 waves present lower *overall* excess mortality. Second, I use variations in wind *timing* as an instrumental variable for city-level deaths and present a causal argument that increases in mortality decreased Bolsonaro's vote share in 2022. My most conservative estimates indicate that a one-third reduction in excess mortality during the pandemic would have been enough to secure a win for Bolsonaro. Using data on state governors' reelection campaigns, I do not find evidence that the results are driven by incumbency effects, and instead seem to be idiosyncratic to the ex-president. I also present a novel method for the construction of counterfactuals relying on neural networks, which provides results consistent with those of my main analysis.

*I thank Paolo Pinotti, who supervised my MSc thesis, for constant feedback throughout its development. I also thank seminar participants at CLEAN, the Degree Commission of my thesis, Peter Hull for comments on my neighborhood averages approach, Giulia Comelli for her suggestions and multiple discussions, and my parents for always supporting me.

1 Introduction

The COVID-19 pandemic left more than 700 thousand confirmed deaths in Brazil. Within Latin America, the country had the second-highest official count of deaths per capita, trailing behind Peru.

The particular severity of the pandemic in Brazil has been repeatedly linked to the then-President Jair Bolsonaro's rhetoric and attitudes. Ajzenman et al., 2023, in particular, use mobile location data and an event-study design to show that cities with higher Bolsonaro vote-share in 2018 disproportionately reduced social distancing relative to other cities as a reaction to Bolsonaro's speeches downplaying the extent of the health crisis. Cross-sectionally, cities in which Bolsonaro obtained a higher share of the votes in 2018 incurred more deaths during the pandemic (Cabral et al., 2021). Figueira and Moreno-Louzada, 2021 show that this holds even within the municipality of São Paulo, leveraging variation at the electoral district level.

Jair Bolsonaro took office in 2018. A retired military man and member of Congress (*deputado federal*) from 1991 to 2018, Bolsonaro was a polarizing figure. Early in the pandemic (March 2020), he stated that the disease affected mostly the elderly, that there was no need to close schools, and that due to his "history of athleticism", he was not at risk. The posture of the former president was met with resistance.

On April 15th, the supreme court in Brazil ruled that governors and mayors had freedom in deciding for restrictive measures, such as lockdowns and whether (and which) establishments should be closed[1]. This marked a shift in Brazilian politics and started tensions between local governments' sometimes stringent approaches and the lax federal ones.

Also in April, and similarly Donald Trump in the United States, he started promoting chloroquine as a potential treatment for COVID-19, despite the lack of consensus in the scientific community[2]. There were a total of four health ministers during Bolsonaro's administration, three in the first year of the pandemic alone. The first two were fired due to disagreements over chloroquine and social distancing.

Around this time, when asked about pandemic deaths, his statements were inflammatory: "I am not a gravedigger", he said. "So what? What do you want me to do?".

In defending that the pandemic was an overblown issue, he said that Brazil was a *"país de maricas"*[3], to imply that people were cowards for fearing the pandemic. When Pfizer announced the results of its first successful clinical trials for their vaccine, Bolsonaro said that "if you turn into an alligator [by taking the vaccine], it is your problem".

A scandal broke out in May of 2021 when it became public that the Brazilian government ignored Pfizer's attempts to negotiate a vaccine deal already in August 2020.

This was confirmed in (and fueled) the ongoing parliamentary commission of inquiry - the *CPI da COVID* -, which was installed to investigate irregularities in the government's handling of the pandemic, promotion of and investment of public funds on ineffective medicines, and dismissal of health officials that disagreed with the president. The *CPI* also received a dossier from whistleblowers which indicated the government was linked to irregular clinical trials made by the insurance company *Prevent Senior*. These included the testing and

[1] Note that the federal government still had the power to establish measures, but not to revoke local ones.

[2] Axfors et al., 2021 eventually showed no benefits from chloroquine in the treatment against COVID-19.

[3] Roughly translates to "a country of homosexuals".

administration of the medicines promoted by Bolsonaro, such as chloroquine, azithromycin and ivermectin. The company was also accused of hiding COVID-19 deaths by changing the victims' cause of death.

Given the highly publicized nature of these events and Bolsonaro's insistence on minimizing the pandemic, one can expect them to affect voting behavior. In 2022, Jair Bolsonaro became the first president in the post-dictatorship era to lose a reelection campaign[4]. Luiz Inácio Lula da Silva, from the worker's party (PT), amassed 50.9% of valid votes and took office on January 1st, 2023.

Though it is virtually impossible to reliably estimate a counterfactual election where Bolsonaro's political choices and his handling of the pandemic were different, one may question, given his choices, whether variations in COVID mortality induced people to change their beliefs about and approval of the ex-President.

I will proceed as follows. For the remainder of the introduction, I formalize my research question and briefly review works related to mine. In section (2), I present the data sources, and proceed to establish a robust relationship between wind and excess mortality during the pandemic in section (3). In particular, I show that places with relatively high wind speeds during national COVID waves incurred fewer overall deaths. These results hold exclusively in 2021 (whereas the pandemic started in 2020). I attempt to explain this pattern through changes in the COVID variant profile in Brazil and changes in behavior. Since weather-related instruments are often problematic due to spatial dependencies, I propose a recentering procedure inspired by Borusyak and Hull, forthcoming. I present my main results in (4). I find that every one excess death per thousand caused roughly a 1 to 3% decrease in Bolsonaro's vote share. Furthermore, a one-third reduction in excess mortality during the pandemic would have been enough to flip the election results, securing a win for Bolsonaro. Reassuringly, my instrument is not related to results of previous presidential elections. In (5), I leverage the fact that governors had a substantial role in pandemic policy-making to evaluate whether my results could be explained by incumbency effects. I do not find evidence for this mechanism. Section (6) lists my robustness checks. Finally, (7) introduces a novel method for constructing counterfactuals based on graph attention networks. Results are consistent with the preceding sections. Section (8) concludes.

1.1 Research question and challenges

What was the effect of COVID deaths at the city level on Bolsonaro's electoral performance in the 2022 elections? Formally, my goal is to estimate β in:

$$\Delta_i = \alpha + \beta \text{ Excess mortality}_i + \delta X_i + \varepsilon_i \tag{1}$$

Where Δ_i is the difference in the PT's[5] runoff vote-share from 2018 to 2022, Excess mortality$_i$ is the number of yearly deaths per thousand above a linear extrapolation of the trend of

[4]Fernando Collor was impeached before he had the chance to run for reelection.

[5]The PT has been present in every runoff for the presidential elections since 2002, which allows me to observe each municipality's electoral preferences over time. In 2022 (and 2018), Bolsonaro's vote share is simply 1 minus PT's vote share (in 2018, PT's candidate was Fernando Haddad).

deaths in city i in 2021[6], and X_i is a matrix collecting potentially unobservable/omitted confounders. In appendix (9.1), I plot the population-weighted average excess mortality across Brazilian cities over time.

Even if one attempts to include all relevant (available) controls in equation (1), simply estimating it at the cross-section of Brazilian cities is a naive strategy. Indeed, one can never guarantee that all potential confounders are included - some of them are not available in public or private data, and others are not even measurable.

One such problematic confounder is people's natural predisposition to agree with preventive measures and to be careful during a pandemic. In cities where people assign a high cost to lock-downs, are vaccine-hesitant, and generally did not think COVID-19 was a serious disease, Bolsonaro's statements and opinions were probably well received. If one believes that such pre-disposition positively impacted death rates, a correlation between changes in Bolsonaro's approval and excess mortality would arise even if excess mortality did not have a causal impact on vote shares.

To tackle these identification challenges, I employ an instrumental variables approach, relying on the observation that cities which had relatively more wind during COVID waves incurred in less overall deaths during the pandemic.

1.2 Related literature

Most work covering Bolsonaro's handling of the pandemic focuses on the disparate impact of the crisis on his supporters (Razafindrakoto et al., 2022, Xavier et al., 2022, Ajzenman et al., 2023, Cabral et al., 2021, Mariani et al., 2020) and their different perception of the risk entailed by the situation (Calvo and Dias, 2021). This literature highlights how the president's actions and speeches can be mapped to increases in COVID cases and mortality, but does not link these mechanisms to electoral outcomes.

Partisan differences in reactions to the pandemic are well-known also outside of Brazil. In particular, in the United States, a considerable literature emerged in the wake of the crisis, reporting differences in beliefs between Republicans and Democrats. Allcott et al., 2020 develop a model in which agents react to the pandemic according to their perceptions of risk, and find that partisan differences in mobility behavior cannot be explained exclusively by differences in actual experienced risk. Grossman et al., 2020 show that democratic-leaning counties reacted more strongly to governors' recommendations to stay at home. Barrios and Hochberg, 2020 present similar results showing that higher Trump vote shares predict fewer Google searches about the virus and smaller reductions in mobility.

My work is also related to the literature on retrospective voting, which is concerned with how people evaluate and react to a politician's policies, in particular taking into account their own welfare during a term (see, for instance, Fiorina, 1978, Persson et al., 1997 and Ferejohn, 1986). Healy and Malhotra, 2013 provides a recent review of the topic.

Some have argued, however, that regardless of how a president acts, events can be quickly interpreted through "partisan lenses". A possible formalization for this phenomenon can be found in Szeidl and Szucs, n.d. In their model, agents have an unwarranted prior belief in an

[6]The reason for considering exclusively 2021 will become clear in later sections. I do not use 2020 in constructing the linear extrapolation for 2021, since the pandemic had already started then.

alternative reality where elites conspire against the incumbent politician. They show that this is enough to guarantee the existence an equilibrium in which an incompetent politician can discredit an honest elite[7] informed about the politician's type. I.e., the elite members are unable to convince voters of the politician's incompetence. Thus, even if mismanagement of the pandemic occurred, it is possible that it did not affect voting behavior if voters did not update their beliefs about the former president.

Focusing on the effects of epidemics on election outcomes, both Arroyo Abad and Maurer, 2021 and Gutierrez et al., 2021 find statistically significant negative effects of outbreaks on incumbents' vote-shares during the Spanish Flu (1918) in the United States and the H1N1 outbreak in Mexico (2009), respectively.

However, there might not be an intrinsic incumbent disadvantage in a pandemic. Herrera et al., 2020 document an increase in approval rates of incumbents across 35 different countries at the onset of the crisis, followed by a decrease as cases continued to grow, in particular for countries with lax restrictive measures. Bol et al., 2021 also find that lockdowns had a positive impact on approval rates. These findings were made early on in the pandemic (still in 2020), and it is unlikely that lockdowns kept their popularity as the pandemic went on. Nonetheless, the evidence suggests that governments' handling of the pandemic may have been more important than pure incumbency effects.

As for COVID-19, there has been some work evaluating its impact on the 2020 presidential elections in the US. Bisbee and Honig, 2021 provide evidence that voters favoured more "mainstream" candidates as a response to the uncertainty provoked by the pandemic - in particular opting for Biden over Sanders in the Democratic primaries. Mendoza Aviña and Sevi, 2021 show that both knowing someone who was infected with COVID or who died from it decreased one's probability of supporting Trump (more so for deaths). They argue that these effects were not enough to affect election results. Warshaw et al., 2020 also find that COVID deaths at the local level decreased support for Trump and other Republican candidates.

The work most similar to mine can be found in Baccini et al., 2021. The authors ask whether COVID cases at the county level reduced Donald Trump's vote share in the 2020 elections. They also use an instrumental variable approach, arguing that the share of workers employed in meat-processing factories predicts higher case counts and that this should not affect vote shares except through the latter. Whereas they do control for potential alternative mechanisms and present a reassuring placebo analysis - looking at the impact of the instrument on previous elections -, it is hard to rule out endogeneity of the instrument[8]. They find that Trump lost votes as a result of more COVID cases, and present estimates that a 20% reduction in cases should have been enough to flip the election result and keep Trump in the presidency. I add to their work by considering a different context and a more widely applicable instrument, which is, hopefully, more credibly exogenous. I also provide novel evidence that incumbent effects do not explain the negative effect of excess mortality on Bolsonaro's vote share.

[7] Whether this is an accurate description of the situation at hand is left to the reader.

[8] It could be, for instance, that people working in the meat industry have attitudes towards the pandemic and policies related to it which are different from the rest of the population. Nonetheless, theirs is a useful exercise, and is more convincing than pure correlational analyses, to the extent that it is harder to come up with alternative stories to the one proposed by the authors.

2 Data

I collected data on municipality-level monthly deaths using the microdata from the *Sistema de Informações sobre Mortalidade (SIM)*, the Brazilian Health Ministry's Mortality Information System. For the years from 1996 through 2020, I used the cleaned dataset from Basedosdados, an open-source, non-profit NGO that provides clean data from the raw administrative files. For 2021 and 2022, I collect preliminary data directly from the Health Ministry on death-counts[9].

Also from Basedosdados, I collect (1) population estimates for all Brazilian municipalities until 2021 - the original source is the Brazilian Institute for Geography and Statistics (IBGE) - and (2) results for presidential and state government elections. Since the runoff results for the 2022 presidential election are missing from the data, I supplement them with the *Tribunal Superior Eleitoral* (TSE)'s data[10].

I furthermore collect micro and mesoregion definitions and municipality areas from the IBGE's publicly available data, which also allows me to construct municipality-level characteristics from the 2010 Census microdata. In particular, I make use of the Data Zoom Stata package, which provides processed and merged Census data[11].

Finally, I collect average monthly minimum daily temperatures, precipitation data and, most importantly, mean wind speeds (m/s) at the municipality level using the brclimr R package (Saldanha et al., 2023), which makes use of the TerraClimate dataset (Abatzoglou et al., 2018), a high-resolution global weather dataset, providing monthly data on a $0.04^{\circ} \times 0.04^{\circ}$ spatial resolution grid ($0.01^{\circ} \approx 1.11$km) up to 2021. Brclimr calculates the mean of the measures available in TerraClimate for each Brazilian municipality.

3 Wind-death timing

Ventilation has been consistently put forth as one of the main ways to decrease COVID contagion risk (Wang et al., 2021, Bazant and Bush, 2021, Morawska et al., 2020). Wind can intuitively affect outdoors contagion by dispersing aerosols, but may also have an effect indoors, through wind-induced natural ventilation. Indeed, it has been proposed as a remedy for improving air quality in closed spaces (Bayoumi, 2021), and natural ventilation (NV) is a valuable strategy to reduce the risk of indoors contagion for airborne diseases (Atkinson, 2009; Ghaffari et al., 2022). See section 3.2 of Izadyar and Miller, 2022 for a recent review.

Evidence of the effects of wind speeds on COVID-19 transmission and fatality rates has nonetheless been inconclusive. Some studies indicate a positive association between wind speeds and cases/reproduction rates of the virus (Ali et al., 2021; Habeebullah et al., 2021), and others a negative relationship (Coccia, 2021; Rendana, 2020). Olak et al., 2022 study the Brazilian case and do not find consistent correlations of wind speeds with COVID-19 cases for four Brazilian cities. Moazeni et al., 2023 provide a comprehensive literature review of the impact of weather variables on COVID incidence, highlighting that studies often contradict each other. Most of them look at daily and moving average relationships, sometimes

[9] The data are available at https://svs.aids.gov.br/daent/centrais-de-conteudos/dados-abertos/sim/

[10] Available at https://dadosabertos.tse.jus.br/

[11] Data Zoom is a project by the Economics department of PUC-Rio.

modelling incubation periods to evaluate the effects of weather variations on subsequent registered COVID cases.

I instead show that the *distribution* of wind throughout the year moderates the impact of national-level COVID waves on city-level deaths. Thus, I present novel evidence and a new method to evaluate the impact of weather variables on the spread and fatality of the COVID-19 pandemic.

Let us start by taking a look at the relationship between excess mortality and average wind speeds at the city-month level in Brazil. I define (monthly) excess mortality as deaths above a linear extrapolation of the month-specific trend of deaths at the city level[12]. This approach is quite flexible and allows for city-specific growth in deaths over time. It also allows me to capture historical heterogeneity in death counts across months[13]. I control for municipality-year and region-month fixed effects, and cluster standard errors at the mesoregion level[14]. To be precise, I estimate the following equation from 2010:

$$\text{Excess mortality}_{i,y,t} = \beta_0 + \sum_{y=2010}^{2021} \beta_y \text{Wind}_{i,y,t} + \gamma_{i,y} + \delta_{r,y,t} + \epsilon_{i,y,t} \tag{2}$$

Where the t, i, y, r subscripts stand, respectively, for the month, city, year and region. The results are presented in panel (a):

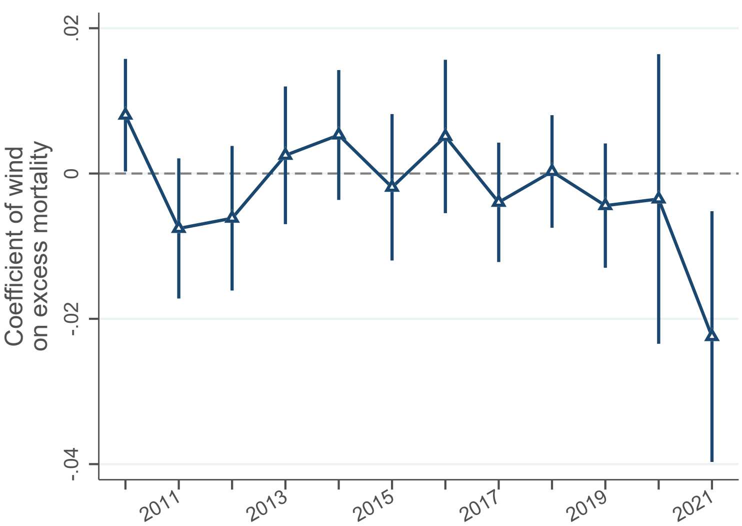

(a) Excess mortality and wind

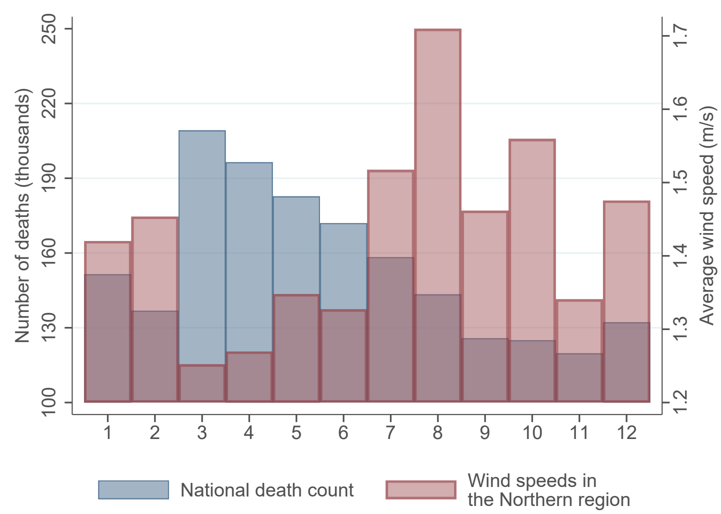

(b) Distribution of wind and deaths throughout the year

Panel (a) shows the β coefficients from equation (2). Note that using excess mortality instead of COVID deaths allows us to make comparisons with pre-COVID time periods. From 2010 to 2020, wind speed is statistically insignificant in predicting excess mortality. The relationship becomes significant only in 2021[15]. One may wonder why the effect appears

[12]For 2021 and 2022 I extrapolate using data up to 2019.

[13]My results are robust to using deviations from a moving average of deaths instead.

[14]There are too few states to cluster by state.

[15]In appendix (9.2), figure (7), I show that adding average monthly precipitation and average daily

exclusively in the last year and not in 2020 when the pandemic started. I will return to this point.

The figure in panel (a) seems promising, but there are a few caveats that deserve mention. Whereas monthly wind speeds negatively impacted monthly excess mortality in 2021, this relationship is lost at the yearly level, as regions that had high yearly average wind speeds had *proportionally less* wind during COVID waves. Intuitively, to get variations in overall mortality, this is a first hint that we should look at the *distributions* of wind and deaths throughout the year.

Furthermore, since both wind and excess mortality are serially correlated variables, estimates of their relationship are sensitive to the selection of different time intervals and may be severely biased. Consider the following illustrative example. Let there be two cities, Alpha and Beta, in periods 1 and 2. Let wind be binary: in period 1, Alpha has wind. Beta doesn't. If people catch COVID in period 1, they cannot catch it in period 2 (short-term immunity).

In period 1, the pandemic hits both cities. 80% of the population is infected in the absence of wind, which reduces contagion by 50%. Alpha thus has 40% of its population sick in period 1, and Beta, 80%. In the first period, wind and contagion are inversely related.

Now, in period 2, since some citizens are immune, 60% of Alpha's population is potentially infected, and only 20% of Beta's. If again 80% of the not-yet-infected population should get the disease, but wind persists (only Alpha has it), then Alpha will have 24% ($0.5 \times 0.8 \times 0.6$) of its population infected in period 2, whereas Beta will have 16%. In period 2, the relationship between wind and contagion flips - even though the causal effect is constant!

Contrary to local mortality, national COVID waves are reasonably not affected by variations in local wind. Therefore, one way to tackle this identification challenge is to consider the relationship between national COVID waves and city-level wind speeds. To build intuition, panel (b) plots the distribution of wind speeds in the Northern region of Brazil and national death counts in 2021. Assume that wind does reduce contagion, as I am arguing. Independently from the *level* of wind in the region, Northern Brazil was unlucky. Its low wind months coincided with the peak of the pandemic in the country. During the wave, wind could have reduced more the number of cases, in absolute terms, than it did when deaths were low. Had the distribution of wind in the North been constant throughout the year, it would have incurred in less overall deaths. I.e., the more wind and excess death co-move, *ceteris paribus*, the fewer excess deaths we can expect for the region in the year.

Formally, for each year and city, I define wind-death timing (WDT) as the covariance between the city's monthly average wind and the leave-one-out share of national yearly deaths for each month.

$$\mathrm{WDT}_{i,y} = \frac{1}{11}\sum_{m=1}^{12}(\mathrm{Wind}_{i,y,m} - \mathrm{Wind}_{i,y})(\text{Death share}_{c,y,m} - \text{Death share}_{c,y})$$

Where subscript y, m denotes monthly averages, y yearly averages, and i and c indicate city- and country-level measures, respectively[16]. I then standardize wind-death timing for each

minimum temperatures does not change the patterns presented here.

[16]Note that local wind speeds should not affect national COVID-19 outcomes, except to the extent that they correlate with national wind levels and that the city influences the country's mortality. Therefore, we

year. Therefore, in all of the following results, coefficients related to this measure should be interpreted as responses to (year-specific) standard deviations in the independent variable[17]. From now on, wind-death timing (WDT) will refer to this standardized measure. Having defined WDT, we can now estimate the following equation:

$$\text{Excess mortality}_{i,y} = \beta_0 + \sum_{y=2010}^{2021} \beta_y \text{WDT}_{i,y} + \delta_{s,y} \times + \varepsilon_{i,y} \tag{3}$$

Where i, y and s stand for city, year and subdivision, such that $\delta_{s,y}$ can be year, state-year, mesoregion-year or microregion-year fixed effects depending on the subdivision represented by s. Errors are clustered at the mesoregion level (137 in Brazil)[18], except for the microregion specification, in which they are clustered at the microregion level (558). How does the relationship between WDT and excess mortality evolve over the years? In figure (1), again and across specifications, we see significant negative effects only in 2021.

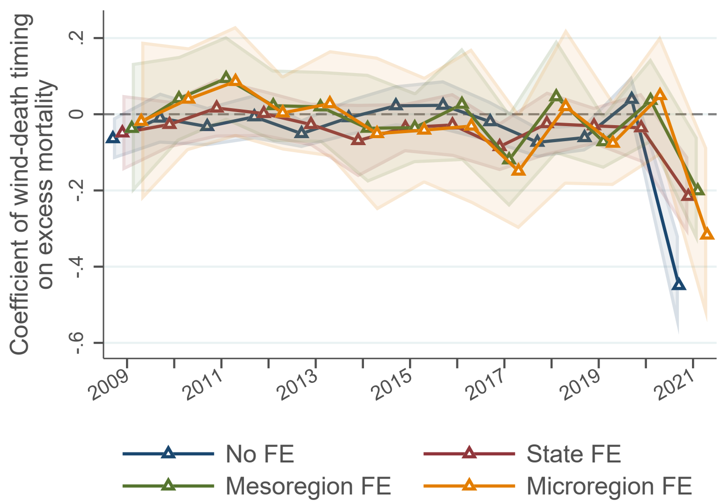

Figure 1: Wind-death timing and excess mortality

Unfortunately, WDT correlates with socioeconomic variables, as shown in appendix (9.3). Therefore, I cannot immediately claim causality, since WDT is possibly correlated to the error term in equation (2). Nonetheless, figure (1) is compelling evidence for its causal effect on *COVID-induced* excess mortality. Indeed, since the group of variables that are significant confounders changes for every level of fixed effects, arguing against a causal effect of WDT would amount to claiming that every group of confounders is a COVID-specific source of

can reasonably expect *variations* in wind-death timing to occur by chance (instead of being determined in equilibrium). Hence the choice of a leave-one-out measure.

[17]In the absence of standardization, results do not meaningfully change. However, since the variance of wind-death covariance changes during the pandemic (excess deaths and its variance increase drastically), standardization allows for clearer visualizations.

[18]Results are robust across different modelling choices, including spatial clustering.

heterogeneity for mortality that goes in the same direction. In appendix (9.2), figure (8), I show that controlling for analogously defined rain-death timing and temperature-death timing does not significantly change the results.

3.1 Why not 2020?

Why are effects null in 2020, when the pandemic had already started? It is hard to make definite statements on this point, but I will try to provide two groups of potential explanations: changes in behavior and changes in the virus itself.

In appendix (9.5), the movement patterns of Brazilians are classified into different locations (Grocery and Pharmacy Stores, Parks, etc.) and plotted over time[19]. All location types, except for residential areas, suffered a sharp decline in attendance early in the pandemic, slowly trending up towards baseline until January 2021, when a smaller decrease took place. In 2021, social distancing (as measured by reductions in attendance) was half as intense as in 2020, despite the larger number of cases and deaths. There are at least three ways, consistent with these data, in which changes in behavior could lead to wind mattering only in 2021. First, it is not clear from Google's mobility data whether people visiting each other's homes counts as time spent in a residential area. As such, the big spike in time spent in residences[20] may mask increased high-risk interactions, and thus poor wind-induced ventilation in social situations in 2020 may dampen the effect of wind speeds on COVID contagion. Second, people learn. Best practices were not widespread at the start of the pandemic, and the status of the virus as airborne was questioned deep into the pandemic, including by the World Health Organization, eventually causing a backlash (Lewis, 2022). Therefore, as time passed and information became widely available, we may expect people to give increasing preference to outdoor meetings and that even indoors, people learned to keep environments ventilated. Third, and relatedly, there may be a selection issue. Whereas average attendance in public spaces did decrease in 2020, it does not mean that *everyone* practiced social distancing to the same extent. If the composition of social interactions in 2021 included more "careful" people - who avoid closed spaces, open windows more, etc. -, then wind should matter more in 2021.

Setting aside changes in behavior, the patterns presented here can also be explained by changes in the virus itself. The turn of the year - from 2020 to 2021 - was accompanied by a quick shift to a new dominant variant of the virus. As shown in appendix (9.6), Gamma, and then Delta variants quickly took over starting from February and August, respectively, whereas the original strand disappeared, for practical purposes. Rowe et al., 2022, in particular, discuss the potential of new variants to present increased airborne transmission. One may therefore hypothesise that changes in the main modes of transmission of the virus may be driving the heterogeneous impact of wind on contagion and subsequent excess mortality.

[19]The data are available online in the form of Google's community mobility reports.

[20]Note that a 20% increase from an already large baseline of time spent at home means larger absolute increases in time relative to other categories.

3.2 Spatial dependencies

As previously mentioned, wind-death timing is not randomly allocated across Brazilian municipalities. This constitutes a hurdle for undoubtedly claiming that WDT reduced COVID excess mortality in Brazil and for the implementation of an instrumental variables estimator. In appendix (9.3), I also present a map with the geographic distribution of WDT in Brazil. There is substantial spatial clustering in the measure.

Spatial dependencies can generate severe omitted variable bias (OVB). Consider an instrument z_i (in my case, WDT), where i is the unit of analysis, which is affected by some sort of spatial dependence. Without loss of generality, we can decompose z_i into $z_i = h(w_i) + g_i$, where $h(w_i)$ is a generic function of the coordinates w_i and g_i are unit-specific zero-mean shocks.

To see why this might be a problem even if we believe z_i does not causally affect the outcome of interest Δ_i except through our main explanatory variable (excess mortality), consider a confounder which also presents spatial dependencies: $c_i = h'(w_i) + g'_i$. Even if c_i is otherwise independent of z_i (i.e. $g \perp\!\!\!\perp g'$), co-movement between h and h', even by chance, can generate powerful correlations at the unit level. For instance, if both confounder and instrument have a north-south gradient, a spurious correlation between z and c will arise.

Whereas a simple north-south gradient would be easy to control for, overlaps in spatial distributions can take complex forms and a curse of dimensionality may follow - and thus these dependencies are not adequately controlled for by latitude and longitude polynomials nor spatial fixed effect estimators[21]. Furthermore, traditional spatial clustering approaches such as Conley, 1999 require the existence of a distance threshold after which $z_i\varepsilon_i$ are uncorrelated, where ε_i is the error term in equation (1). This assumption is particularly problematic for weather variables. See the great discussion in the extended version of Borusyak and Hull, forthcoming (BH).

I thus propose a new approach to dealing with OVB induced by spatial dependencies, and extend the framework of BH. The authors show that if an instrument is a function of an exogenous and an endogenous component, subtracting the expectation of the instrument *given* the endogenous component makes the instrumental variable estimator consistent. I.e. for an instrument $z_i = f_i(w, g)$, where $g \perp\!\!\!\perp \varepsilon \mid w$, the recentered instrument $\tilde{z}_i = z_i - E[f_i(w, g) \mid w]$ identifies β in equation (1). In their framework, the distribution of g is known, and thus $E[f_i(w, g) \mid w]$ can be easily obtained.

In the present case, g is an unknown component of z, and the potentially endogenous component (w) is location itself. Therefore, we need to measure the expectation of the instrument *given coordinates*. A reasonable estimator, in this case, is simply taking neighborhood averages - say, for example, the average WDT for the hundred closest cities. Note that even if neighborhood averages are not an unbiased estimator for $E[f_i(w, g) \mid w]\ \forall\ i$, as long as the city-specific bias is itself exogenous, the BH framework generalizes straightforwardly.

Let us call *net* WDT (NWDT) the difference between a given city's WDT and the average WDT for its 100 closest cities (in terms of distance in kilometres)[22]. In appendix (9.4), I present the geographic distribution of the new measure and its relationship with socioeconomic variables. The F-statistics across specifications are now much smaller. Indeed, for

[21] E.g. a variable may maintain a north-south gradient even within the fixed effect delimited areas.

[22] In appendix (9.8.3) I show the results for different choices of neighborhood size.

microregion fixed effects, we cannot reject the null hypothesis that all the socioeconomic variables included in the model have a 0 coefficient. This indicates that the recentering procedure likely succeeded in reducing OVB. I will use NWDT to investigate the impact of excess mortality at the city-level on Bolsonaro's 2022 vote share, though my results are robust to using WDT instead. Importantly, the patterns of figure (1) do not change when considering NWDT, thus reinforcing the argument that wind caused a reduction in COVID mortality at the city level. Henceforth, wind-death timing refers to NWDT, except when explicitly stated otherwise.

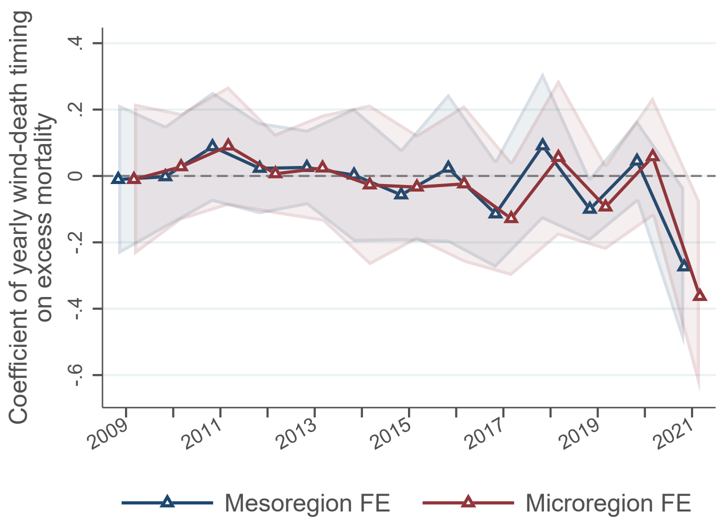

Figure 2: Net wind-death timing and excess mortality

4 Results

Did city-level heterogeneity in COVID mortality drive differences in voting outcomes in the 2022 presidential elections? To answer this question, I use an instrumental variable: net wind-death timing.

The baseline estimates indicate a sizeable effect of excess mortality on the vote shifts towards the worker's party (PT) - and thus away from Bolsonaro. Indeed, a one death per thousand increase in excess mortality is associated with a 2.7% increase in the runoff PT vote-share for mesoregion FE, and 1.3% for microregion FE. Given that at the national level, the pandemic entailed 3 extra deaths per thousand inhabitants in 2020 and 2021 together, the estimates imply an average 3.9% and 8.1% shift for meso and microregion FE, respectively. Effects of this size would have been more than enough to flip the election result, as the difference in vote-share between Lula and Bolsonaro was 1.8%[23,24]. Indeed, using 1% vote-share change

[23]The same is true for the estimates controlling for potential confounders, presented in appendix (9.8.1).

[24]Given that I always refer to valid votes - i.e. not counting absences and null votes -, note that a 0.9%

Table 1: Impact of excess mortality on electoral changes (IV)

	Dep. var.: Difference in PT vote share (2022-2018)	
Excess mortality 2021	0.027***	0.013**
	[0.009, 0.235]	[0.003, 0.065]
Kleibergen-Paap F statistic	5.052	5.570
Number of municipalities	5562	5562
Mesorregion FE	✓	✓
Microrregion FE		✓

Excess mortality in 2021 instrumented by the net covariance between leave-one-out monthly death shares at the national level and city-level monthly average wind. Clustered errors at the level of the FE's for the two specifications. Given the weakness of the instrument in both cases, I present Anderson-Rubin confidence intervals and p-values.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

per death per thousand as a lower bound of the effect, a one-third reduction in excess mortality in 2020 and 2021 would be enough to bring Lula from 50.9% to 49.9% and flip the election result.

In this baseline, I do not include any controls except for subregion fixed effects. In appendix (9.8.1), I present the estimates obtained by including the full set of controls shown in (9.4). The point estimates are slightly smaller but statistically indistinguishable from the ones presented here, though they are significant only at the 10% confidence level. This is mostly due to a widening of the confidence sets.

Also note that, whereas the relationship between wind-death timing and excess mortality in 2021 is quite robust across specifications, the first-stage Kleibergen-Paap (KP) F-statistic of NWDT is below the usual threshold of 10 after controlling for meso and microregion fixed effects. In the just-identified case, KP coincides with the effective statistic from Olea and Pflueger, 2013, which is the appropriate statistic for the case of non-homoskedastic errors (Andrews et al., 2019). As such, I present the Anderson-Rubin p-values and 95% confidence sets, which are consistent under weak instruments.

Due to a lack of statistical power, my main specification does not provide tight confidence intervals, and thus the difference in magnitude between the micro- and mesoregion FE specifications could just reflect noise. However, there is a plausible explanation for why we should expect different levels of fixed effects to imply different estimates. Say that a citizen of city i cares not only about their own city, but also about their microregion and mesoregion. For simplicity, we can represent this with a linear model:

$$\Delta_i = \beta_1 \text{Excess mortality}_i + \beta_2 \text{Excess mortality}_{micro} + \beta_3 \text{Excess mortality}_{meso} + \epsilon_i \quad (4)$$

Where Δ_i is the usual shift in vote-shares for the worker's party and Excess mortality micro and meso are averages for their respective levels, included to capture the effects derived

decrease in PT vote shares would be enough to flip the election, since that would imply a 0.9% increase for Bolsonaro.

from people caring about cities other than their own. Clearly, controlling for mesoregion fixed effects, there is no residual variation in Excess mortality$_{meso}$, and thus β_1 is not biased by the omission of Excess mortality$_{meso}$ in the estimated equation. But it is biased by the omission of Excess mortality$_{micro}$! As we control for progressively more fine-grained FE, the more we take away from our estimates the effects of people caring about their vicinity. If we are generally interested in the effect of the pandemic on votes, this may not be desirable. It boils down to a choice of estimand.

A corollary of this reasoning is that the estimates presented here are likely a lower bound of the total effect of COVID deaths relative to a counterfactual world without the pandemic[25]. Indeed, at the limit, if people cared exclusively about the national situation, then heterogeneity in city-level excess mortality should not imply differences in vote shares between cities. Therefore, if people care about their *local* situation, but also care - separately - about the national situation, then the effect of an extra death relative to the national average does not capture the full effect of the pandemic on votes.

Note that I do not weight observations by population size - doing so, unfortunately, makes the first stage unfeasibly weak[26]. However, as I show in appendix (9.8.9), there does not seem to be a gradient of the point estimates based on population size. As such, the results presented here are likely to be a good approximation of actual election results - which are obviously population-weighed.

4.1 Exogeneity

The instrumental variable estimator presented above requires two main assumptions. First, the instrument must be *valid*. That is, it must be significantly correlated to the independent variable of interest. Although my instrument is weak, my results are significant using weak-instrument-robust statistics.

Secondly, the instrument must be exogenous - that is, the timing of wind must not affect election outcomes *except through* changes in excess mortality. How could NWDT affect vote shares?

At the national level, it could be that regions that have a particular wind distribution have certain geographic characteristics that lead to particular political preferences and reactions to policy choices during the pandemic. Considering variation within fine-grained administrative regions, however, it is unlikely that heterogeneity in the distribution of wind throughout the year affects voting patterns. Nonetheless, as shown in appendix (9.8.7), controlling for the average NWDT *before* 2021 does not change the results.

Exogeneity could also be violated if the timing of wind throughout the year is related to socioeconomic confounders which lead to different voting patterns. Previously, I showed that NWDT seems to be quasi-randomly distributed conditioning on microregion fixed effects. Nonetheless, throughout my analysis, I show the robustness of the results to including the full set of controls I gathered.

I will present three further pieces of evidence, necessary to argue that the exogeneity of the

[25] This rests on an assumption of monotonicity: if excess deaths in a city reduce votes for Bolsonaro, then excess deaths in the country, having an effect of their own, should also reduce votes for Bolsonaro.

[26] This is expected: within meso and microregions, weighting for population removes most of the variation within the subregion, as most of its population will often concentrate in a single city.

instrument is satisfied. These are also generally consistent with the hypothesis that NWDT induced changes exclusively in COVID-specific excess mortality, which in turn affected Bolsonaro's electoral outcome. Each of these three is associated with a slightly different measure.

- **Treatment status**: wind-death timing in 2021 - which I will refer to as treatment - does not affect previous election results nor excess mortality (alternatively: no pre-trends for treated cities).

$$\text{Treatment status}_{\ y} \ = \ \text{NWDT}_{\ 2021}$$

- **Seasonal winds**: The covariance between monthly wind speeds in previous years and monthly death shares *in 2021* does not matter in any other year (to account for, e.g., winds in February influencing elections/deaths).

$$\text{Seasonal winds}_{\ y} \ = \ cov(\text{Wind}_{y,m},\ \text{Death share}_{2021,m})$$

- **Yearly wind-death timing**: The covariance between wind speeds and monthly death shares does not matter in any other year except 2021 (when such covariance does influence overall mortality).

$$\text{Yearly wind-death timing}_{\ y} \ = \ \text{NWDT}_{\ y}$$

Another assumption, which cannot be checked, is that the instrument is not correlated to any omitted confounders which are sources of COVID-19-specific heterogeneity in election outcomes. As NWDT is not correlated with the set of confounders I gathered conditional on microregion FE, and since my results are robust to controlling for such confounders, I find no evidence that this assumption is violated.
What does the relationship of each of these three measures with excess mortality look like? In appendix (9.7), I show that, across specifications, we see the same patterns of figure (2) - the coefficients are mostly insignificant and small in magnitude before 2021. In 2021, across specifications, we find negative and statistically significant effects of NWDT on excess mortality.
In sum, particular seasonal wind distributions throughout the year do not affect excess deaths before 2021. Nor do the selection of treated cities and year-specific wind-death timings. This indicates the link is COVID-specific.
As for the influence of these measures on election results, the pattern is broadly similar. NWDT in 2021 reduces vote shifts towards the worker's party in the 2022 election. Given that it also reduces excess mortality in 2021, under suitable assumptions, excess mortality caused a shift in votes *towards* the workers' party.

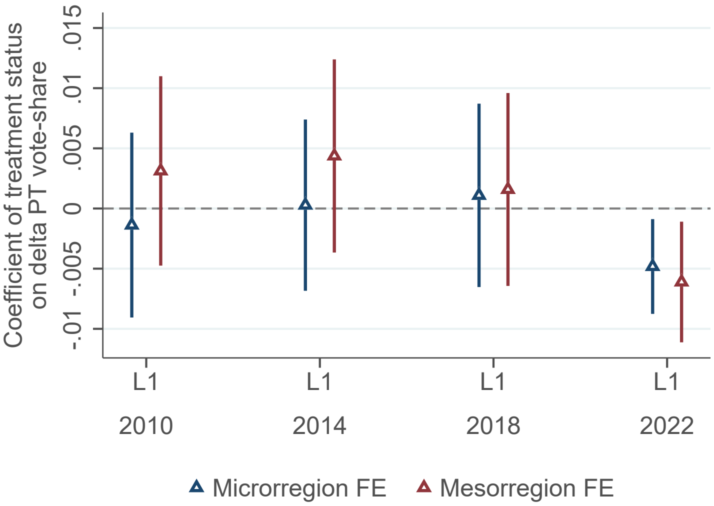

(a) Treatment status

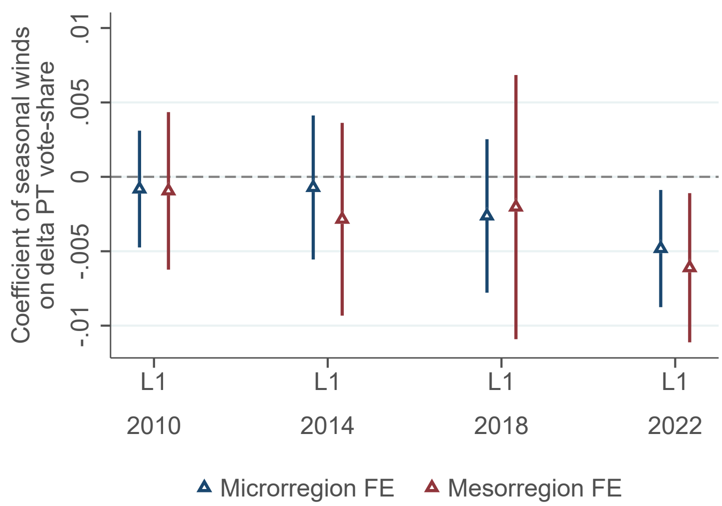

(b) Seasonal winds

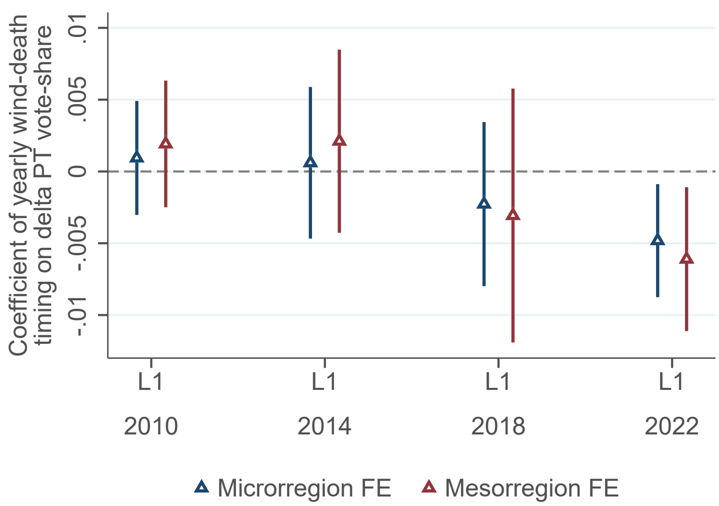

(c) Yearly wind-death timing

5 Incumbent effects

Even if one accepts that COVID-induced excess mortality causally - and negatively - affected Bolsonaro's electoral performance in 2022, voters may simply be unsatisfied with whoever is in power during a period of hardship, especially one so ubiquitous and reaching as the pandemic. Thus, incumbency effects could explain the results presented so far. I will provide evidence that this is not the case.

As mentioned in the introduction, the federal government was not the only source of pandemic-related policy. On April 15th, 2020, the Brazilian Supreme Court ruled that the federal government could not overrule state and municipality decisions to put up additional restrictions on mobility and to close establishments. This did not mean full independence - the federal government could still legislate on the topic. However, the aftermath of this decision was that most lockdowns and closings were decided at the local level, and measures were uncoordinated and heterogeneous. This went as far as independent vaccine purchases by governors.

Given the heterogeneity in policy and local responsibility, this gives us a good set-up to investigate whether, independently from political affiliations and decisions, incumbents' vote shares were negatively affected by excess mortality. That is, did incumbents - both governors and the president - fare relatively worse in cities that suffered from higher excess deaths?

The strategy here is identical to that of the main results. The sole difference is that the outcome is now the change in the incumbent state governor's vote-share in the first-round, to include governors who won or lost the reelection without the need for a runoff[27].

Table 2: Impact of excess mortality on state government electoral changes (IV)

	Dep. var.: Difference in incumbent vote share (2022-2018)	
Panel I: effect for governors running for re-election		
Excess mortality 2021	-0.012	-0.010
	[-0.091, 0.093]	[-0.083, 0.081]
Kleibergen-Paap F statistic	5.901	5.781
Number of municipalities	4243	4243
Panel II: excluding Bolsonaro-aligned governors		
Excess mortality 2021	0.010	0.006
	[-0.047, 0.412]	[-0.065, 0.191]
Kleibergen-Paap F statistic	5.213	5.017
Number of municipalities	3774	3774
Mesorregion FE	✓	✓
Microrregion FE		✓

Excess mortality in 2021 instrumented by the net covariance between leave-one-out national monthly death shares and city-level average monthly wind speeds. Clustered errors at the level of the FE's. Given the weakness of the instrument, I present Anderson-Rubin confidence intervals and p-values. Bolsonaro-aligned governors are defined as those from Republicanos and PP. Vote-share measured in the first round.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Governors can run for a maximum of two consecutive terms in Brazil. In the 2022 general elections, a total of 20 candidates ran for reelection[28]. A whopping 18 were successful. I retrieved city-level vote-shares for governors running for reelection in 19 out of the 20,

[27] In Brazil, governors need 50% of the votes in the first round to win right away.
[28] Only 7 disputes did not have incumbents running.

amounting to a total of 4243 municipalities[29]. Note that restricting the presidential analysis to these municipalities does not meaningfully change the results.
In panel I, I present the estimates for all governors running for reelection. The point estimates in both meso and microregion specifications are slightly negative - though not far off from the point estimates we obtained for the presidential runoff. However, if we restrict the analysis to non Bolsonaro-aligned governors (PP and Republicanos), the estimates increase and flip in sign. That is, the effect for incumbents whose parties did not support Bolsonaro before the first round of the elections is if anything, slightly positive[30].
The significant effects in the presidential analysis do not seem to stem from an unconditional incumbent disadvantage during a crisis. This reinforces the thesis that Bolsonaro, due to his approach towards the pandemic and denialism, and not some incumbent curse, received a backlash from those most affected by the pandemic.

6 Robustness tests

I conducted a series of tests to validate the robustness of my results. First, as previously mentioned, adding the full set of controls to my main specification does not significantly change coefficients, though the magnitude of the point estimates is slightly smaller - 0.019 and 0.011 for meso and microregion definitions, respectively. They remain significant at the 10% level, and the changes in significance are likely due to a widening of the confidence sets following a decrease in the first stage's F-test. The table with these results can be found in appendix (9.8.1). In appendix (9.8.7), I also show that controlling for a moving average of the instrument does not affect my estimates.
Second, in appendix (9.8.2), I show that using WDT instead of the *net* wind-death timing does now change my conclusions. If anything, the point estimates are larger - conditional on the full set of controls, they are 0.038 and 0.011, and slightly greater unconditionally. It is reassuring that using WDT provides similar results - it doesn't seem that any bias stemming from spatial correlation in the instrument significantly affects the estimates.
Third, table (7) and figure (11) in appendix (9.8.3) highlight the robustness of the results to varying the neighborhood size definition. In the table, I present the main result (the simple IV, with and without controls), for neighborhoods of size 20, 40, 80, 100 (baseline) and 150. As can be expected, smaller neighborhoods yield lower magnitudes for the point estimates. The reasons are likely threefold: (1) wind-death timing net of small neighborhood averages is potentially associated with proportionally larger measurement error, leading to estimates biased towards 0; (2) COVID contagion spillovers to nearby cities may make local differentials in wind have an attenuated effect on mortality; (3) the mechanism enshrined in equation (4). Nonetheless, unconditional estimates are significant for all but neighborhoods of size 20, and estimates are almost identical for neighborhoods from sizes 80 to 150.
In figure (11) I plot the coefficients of an event study specification, for each neighborhood size and for WDT (i.e. without netting out neighborhood averages). Each shade of blue represents a different equation, and the lighter colors represent progressively smaller neigh-

[29] The data from Basedosdados does not include information for the state of Mato Grosso in 2018.
[30] Note that this is a very lenient definition of Bolsonaro ally. Some politicians such as Zema, governor from Minas Gerais, declared their support after the first round.

borhoods. Note that these equations track a constant selection of cities into treatment across time. Clearly, using WDT alone entails an issue - the cities that had higher wind-death timing in 2021 were decreasing their PT vote share in every election by a greater amount than in 2018. But selection into treatment should not affect previous elections! This is suggestive of some omitted confounder, though if we were to extrapolate the pre-trend (which is slightly increasing), the estimates would be biased *towards* 0, and my main results would still hold. However, subtracting neighborhood averages from WDT seems to eliminate the issue, and progressively more so for smaller neighborhoods. The point estimates for a neighborhood of size 20 are pretty much constant at 0 from 2008 to 2018, and the pattern is similar for other neighborhood definitions up to size 100. Thus, if neighborhood size choice can be thought of as a trade-off between power and less confounding - as captured by flatter pre-trends -, my baseline specification is a good choice.

Fourth, I test whether having COVID death and case rates as the main independent variable yields results similar to those of the baseline specification. Here, unfortunately, NWDT is too weak an instrument to obtain reliable estimates. However, WDT provides a good first stage, and the point estimates using either NWDT or WDT are very similar. This, in conjunction with previous evidence that WDT, conditional on the full set of controls, does not entail significant bias, validates its usage here. Results are indeed encouraging: not only are estimates significant, but adding controls leaves point estimates practically unchanged! Note, however, that these estimates are quite large: 0.086 and 0.033 for meso and microregion fixed effects, respectively. That is 3 to 4 times larger than the estimates found when using excess mortality. Whereas this could reflect the usage of a different instrument and a higher impact of COVID deaths relative to excess mortality in general, there is another good reason why this might happen. Since the IV estimate is a local average treatment effect (LATE), we have to be careful about who are the compliers in the present case. Compliance is likely not a big issue when it comes to excess deaths - people probably can't decide whether to die or not -, but it takes a front seat when it comes to characterizing COVID deaths. Testing can play a role here since the municipalities most responsive to WDT in terms of COVID death rates are likely the ones which test more. Also, the cities that test more are likely to be more careful towards COVID and react more against Bolsonaro in case of a higher death toll. This could explain larger point estimates.

Fifth, I check that the results for the presidential election hold in the subsection of states for which the incumbent governor was running for reelection. I do this to ascertain that the null findings for incumbent governors do not stem from the particular subset of cities with incumbents running, which could drive the results. Indeed, point estimates for the main specification in this group of cities are basically unchanged and more precisely estimated.

Sixth, I check that using an alternative definition of excess mortality does not change the results. Estimates do not substantially change when considering excess mortality in 2021 to be defined as deaths above the average from 2017 to 2019 (I exclude 2020 since the pandemic had already started then) per thousand inhabitants.

Seventh, I present the results for a variation in the dependent variable: using the first round instead of runoff results. Estimates are quite similar (0.032 and 0.019 for meso and microregion), though a bit larger in magnitude. This is consistent with the marginal voter changing their mind enough not to vote for Bolsonaro in the first round but not sufficiently to support the worker's party - which gathers staunch opposition from a substantial part of

the population in Brazil.
Eighth, I show that there is no evidence for effect heterogeneity by population size in appendix (9.8.9). This suggests that my results are representative of population-weighed estimates[31]. Note that I use WDT instead of NWDT to provide more power to the heterogeneity analysis.

7 Graph neural network counterfactuals

The main issue I encountered when trying to use (non-net) wind-death timing as a treatment, as shown in figure (11), was that cities selected into treatment - i.e. WDT_{2021} - presented differential trends in prior elections. In previous subsections, I tackled this issue using NWDT, which reduces the scope for spatial clustering and omitted variable bias (as measured by correlation with observable potential confounders). In this section, I show that graph neural networks can provide credible counterfactuals even if there is selection into treatment.
Let us consider, for the purposes of this analysis, that cities with WDT above the median in their state are "treated", and assign them to "control" otherwise. How can we get an appropriate counterfactual if the treated and control groups are substantially different and their voting patterns do not evolve in parallel?
The two groups still carry information about each other. Traditionally, one would perform some kind of matching (e.g. using propensity scores or synthetic controls), so that treated and control units are paired according to some underlying measure of similarity. Though at its core my approach does match treated and (a function of the) control units, I try to tackle the issue as a prediction problem: using only the control units, I attempt to construct a model that approximates the treated units' outcomes as well as possible.
To do so, I use a graph attention network (GAT). Explaining GATs from scratch is outside the scope of this project, but I will try to provide some intuition as I go. The technicalities of the machine learning algorithm are not necessary to understand the application presented here[32,33]. I start with an overview of how counterfactual prediction fits within a potential outcomes framework.

7.1 Potential outcomes

In this section, I propose a method to evaluate treatment effects based on a neural network's (lack of) accuracy post-treatment. This comes down to recasting causal inference as a prediction problem. I refer to Chernozhukov et al., 2021 for a summary of the related literature and a general framework that includes the method presented here. My discussion will be inspired by their conformal inference test[34].

[31] As mentioned above, the instrument becomes unfeasibly weak if one applies population weights and controls for fine-grained fixed effects.

[32] The Jupyter Notebook containing the code is available at https://github.com/RafaelPintroSchmitt/neuralnet-counterfactuals.

[33] The definition of the attention layer can be found at https://pytorch-geometric.readthedocs.io/en/latest/generated/torch_geometric.nn.conv.GATConv.html

[34] Their test relies on permutations across the time dimension. Since I have very few (4) years of observation and a single post-treatment period, the test becomes trivial.

Following the potential outcomes framework of Neyman, 1923 and Rubin, 1974, and borrowing the notation from Chernozhukov et al., 2021, let $t \in \{1, 2, ..., T\}$ be a generic time period in a sequence of length T. I consider $\{Y_{i,t}^I\}_{t=1}^T$ to be a sequence of outcomes for a unit of observation i followed over time, under some intervention at time T_0. Let θ_t be a scalar capturing the effect of the intervention, with $\theta_t = 0$ for $t < T_0$. $\{Y_{i,t}^N\}_{t=1}^T$ denotes $Y_{i,t}^I - \theta_t$ for each t, that is, it represents a counterfactual world where the intervention did not take place.

Finally, let $\{P_{i,t}^N\}$ be a sequence of mean-unbiased predictors or proxies for $Y_{i,t}^N$, such that $Y_{i,t}^N = P_{i,t}^N + u_t$, where $E(u_t) = 0 \; \forall \; t \in \{1, ..., T\}$. The potential outcomes can thus be written as:

$$\begin{aligned} Y_{i,t}^N &= P_{i,t}^N + u_t \\ Y_{i,t}^I &= P_{i,t}^N + \theta_t + u_t \end{aligned} \tag{5}$$

Under suitable assumptions, having $P_{i,t}^N$ allows us to test hypotheses on θ_t. One such hypothesis is whether $\theta_t = 0 \; \forall \; t \geq T_0$. If we reject it, we have evidence that the intervention had *some* effect.

My contribution will be to provide a method to construct such counterfactuals (i.e. to get a plausible sequence $\{P_{i,t}^N\}$), using a graph neural network architecture. If such a method can extrapolate from the training sample, i.e. it can predict the dependent variable accurately for t outside the periods used for training, then it provides a good candidate for $P_{i,t}^N$.

7.2 Defining the neural network

The first step is setting up the data structure which will be fed into the model. I construct a graph - also called a network - for each year. Each city is connected to its 10 closest neighbors in terms of geographical distance, which is itself used as an edge attribute (a weight used by the neural network). I include the covariates listed in appendix (9.4) as node attributes, together with the year-percentile of the difference in PT vote-share relative to the previous election, which is the outcome of interest[35], and treatment status.

In machine learning jargon, predicting the outcome of interest for each city is a node-level classification task. As such, it is odd to include the percentile of the difference in vote shares as a covariate, since - usually - the network could just learn to use the outcome to predict the outcome. This would not be very useful. My reasons will become clear.

The objective of the model is simple: given a sample of Brazilian cities' changes in PT vote share, it should predict the outcome for all other cities. If it learns to do so, we can hope that by inputting the control cities' information, the model will be able to predict the treated cities' outcomes. Note that the model learns a generic instruction: given any set of cities, predict the rest. But by learning to do so, it becomes well-suited for our objective, which is creating a counterfactual for the treated group.

I split the sample into training and test sets. The model's parameters are obtained from one set of cities, and my results are derived by applying the model to a different set.

Now I can explain the neural network's architecture - i.e. how it learns. I start with an

[35] This measure ranges from 1 to 100 where 1 is the smallest difference (more negative) and 100 is the highest.

overview for the reader unfamiliar with neural networks and then go into the specifics. At the first epoch (or training round), I feed the model with the 2010 network. Each layer of the network is basically a set of instructions to receive, transform and transmit the vector/matrix received from the previous layer. Then, the model applies an optimizer step, more precisely a variant of gradient descent called Adam[36]. Intuitively, this optimizer step consists of changing the parameters of every layer a little so that the training loss[37] is smaller. I use a cross-entropy loss function. In the next training round, I use the 2014 network, and then the 2018 one. The cycle resumes, from 2010 to 2018[38]. It may seem that I have exhausted my data already at the third training round, but the key here is that each year's dataset has many *subsets* of cities.

The actual architecture and structure of the layers is quite simple. First, and very importantly, I apply a dropout layer, which consists of dropping, at random, all the information for a given percentage of cities. Note that the model will eventually reduce *overall* losses. By using a dropout layer right away, the model will have to learn to predict the outcomes for the omitted cities using information coming only from the non-dropped ones. This also explains why I kept the outcome as one of the covariates: I want the network to *use the outcomes* of the non-dropped cities to predict those of the omitted ones.

The second layer is then a graph attentional operator from Veličković et al., 2018. In simple terms, this is a layer that performs a weighted sum of the covariates of the node's neighbors and its own. The weights for each neighbor are learned - and can depend on similarity across any covariate[39]. Then, I apply a ReLU layer, which is a piece-wise linear function, and another graph attentional operator. Finally, I apply a linear layer.

Summarizing:

- Architecture:
    1. A dropout layer (with varying dropout rates).
    2. Two attention layers, with a ReLU in between.
    3. A linear layer.
- Training procedure:
    1. Perform a forward pass and take an optimizer step for each year in the training data (2010-2018).
    2. 100 training rounds for each year - 300 forward passes and optimization steps.

I trained the model, 300 forward passes and optimization steps, a hundred times, obtaining different parametrizations for each[40]. I call each of the hundred iterations a *training procedure*.

[36]Kingma and Ba, 2017

[37]Which is a function summarizing how wrong the model is.

[38]Note that I never include 2022 in the training process so that treatment outcomes - the effect of WDT on excess mortality - do not affect training.

[39]In this sense, this is not too different from a synthetic control approach.

[40]This is expected, since the dropout layer randomly selects cities to be dropped, causing different training paths. I also use different random seeds for each training procedure.

7.3 Performance

The model performs reasonably well. Be reminded that the outcome of interest is the percentile of the difference in worker's party vote share for each year. Also, let us define $R^2_{GAT} = 1 - \frac{MSE_M}{MSE_R}$, where MSE_M is the mean-squared error of the model and MSE_R is the mean-squared error under random guesses.
The testing is simple: I take the set of cities in the test set, which were not used in training, select half at random and make their outcome variable equal to zero. Then, I input the resulting dataset into the model and evaluate it on its ability to guess the dropped cities' outcomes correctly.
Across training procedures, on average, $R^2_{GAT} \approx 0.8$ in the test set, and its accuracy - getting the percentile exactly right - is 2 to 3 times higher than random guesses. The performance on training and test sets is similar, which indicates overfitting is likely not an issue, and the model is on average unbiased from 2010 to 2018. I will come back to this last point.
A fairly representative training procedure yields model predictions such as this:

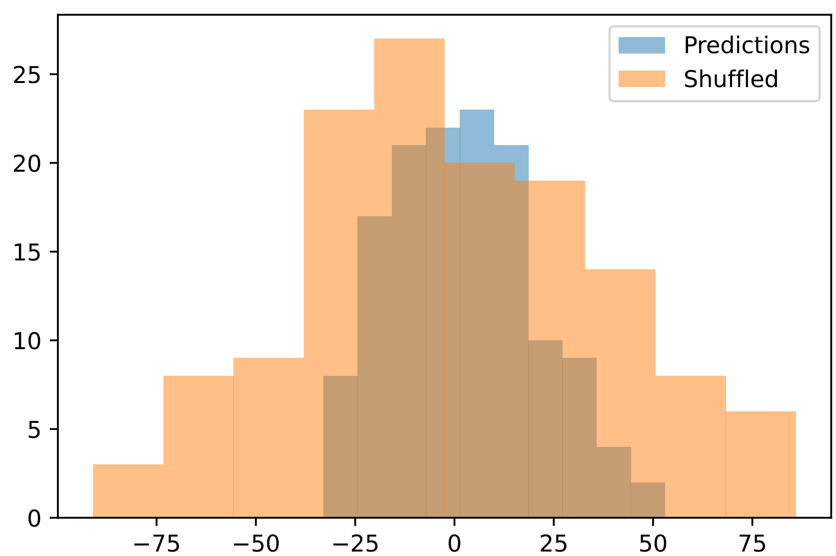

Figure 3: Perfomance of the model

Where the x-axis is the error in the prediction (e.g. a prediction of 25 for a true value of 10 would yield an error of 15). In blue, we can see the model's errors, and in orange, the errors stemming from guessing using a random shuffle of the percentiles (random guesses). Clearly, the neural network learned to predict a city's voting outcomes using the information of the cities in its neighborhood and its covariates.

7.4 Treatment effect evaluation

We can finally turn to the evaluation of treatment effects using the neural network's prediction as counterfactuals. Be reminded that I define as "treated" those cities with wind-death

covariance above the median in their state.
For the years 2010 through 2022, I omit the outcomes of the treated cities and feed the data into the network. I get predictions for each treated city which are based only on the control cities (and covariates of the treated cities). For each year, I record the bias of the predictions (average error across cities) and repeat the training and evaluation 100 times. I perform the same exercise randomly selecting "treatment" units at every training procedure and year, as a placebo test.
The results are plotted in the following picture:

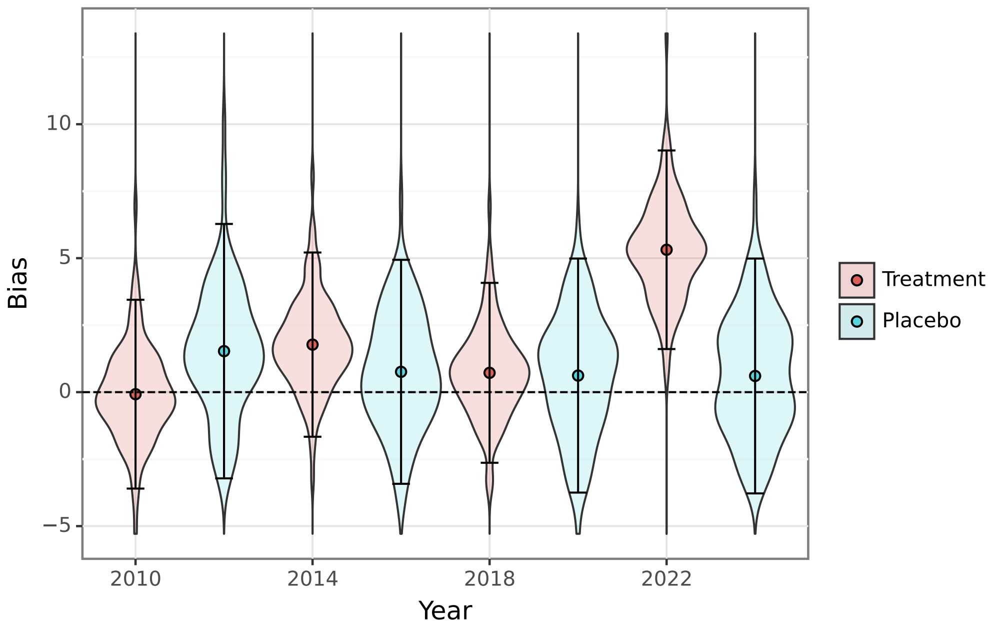

Figure 4: Bias distributions over time, training on 2010-2018

The volume around the vertical lines represents the distribution of the bias obtained for the hundred training procedures. The error bars represent the 95% confidence intervals for the bias, calculated from the observed distribution[41].
The model is roughly unbiased before treatment for the treated, and unbiased before *and after* treatment for the placebo. Using the notation introduced at the beginning of this section, the neural network's predictions are a good candidate for a sequence $\{P_{i,t}^N\}$ of mean-unbiased predictors or proxies for $Y_{i,t}^N$.
The bias is positive for 2022. This means that the predictions of the model - our counterfactual - *overshot* the worker's party vote share in treated cities[42]. From (5), we get that θ_{2022}, the treatment effect, should be negative. Knowing that WDT negatively impacted excess mortality in 2021, we can revert back to the IV reasoning of previous sections to argue that excess mortality caused a decrease in Bolsonaro's vote share. Note, however, that I tackled

[41] I.e. the space between the limits of the confidence intervals contains 95 training procedures out of the 100.
[42] The model gets the treated group's percentile wrong by about 5 positions up, on average.

the issue of selection of cities into treatment[43], particularly when using WDT instead of NWDT, by constructing a credible counterfactual.
From the error bars for the placebo, we can see that, across the training procedures, a random selection of cities into treatment would generate results as extreme as the ones observed for the treatment group less than 5% of the time.
Furthermore, in the spirit of the conformal test proposed by Chernozhukov et al., 2021, I provide some back-of-the-envelope calculations on how likely it would be to find results as extreme as the ones presented here if there were no effect to be found. Let $\{0, 0, 0, 1\}$ represent the current result, indicating that I do not find effects for the first three years and do find an effect for the fourth. This configuration has a 25% chance of occurring under the set of all permutations of the observed results, which can be interpreted as a sort of p-value.
Taking the placebo in conjunction with the permutations test, it is unlikely that the results here do not reflect a treated-group-specific effect in 2022. However, the same caveat from before applies: it could be that the treated group is different from the control across some dimension that moderates the pandemic's impact on the elections, and the neural network did not have the opportunity to learn this pattern.
The main patterns do not change if I restrict training to 2010 and 2014:

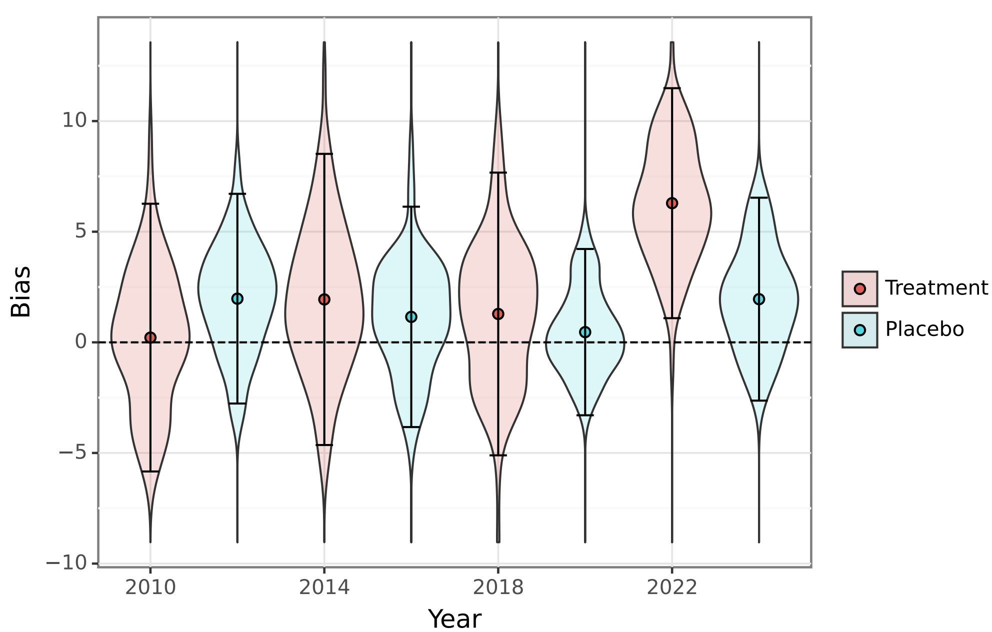

Figure 5: Bias distributions over time, training on 2010-2014

Finally, note that this is also an exercise in model selection: given the architecture, the results here are an average of many parametrizations stemming from different random seeds. As such, I provide evidence for the robustness of my results with respect to control group selection and counterfactual generation procedures.

[43]One would be right to argue that I should have also trained the model to predict excess mortality to make this argument. Since pre-trends were never an issue for excess mortality, I omit this exercise.

8 Discussion and conclusion

This paper has examined the causal effect of city-level variation in excess mortality in 2021 on Bolsonaro's performance in the 2022 elections. Whereas most of the past literature has focused on the consequences of a leader's speeches and attitudes in driving changes in behavior during the COVID-19 pandemic, I shed light on how they translate into electoral accountability, by exploiting the nuanced relationship between meteorological factors, pandemic-driven mortality, and political outcomes, offering new insights for future work.

First, I establish a connection between wind speeds and COVID-induced excess mortality. In particular, variations in the timing of wind throughout the year - beyond the mere *level* of wind - are key in moderating the impact of national pandemic waves on local outcomes. Wind serves as a protection, and having more of it during critical periods translates into fewer overall deaths. Importantly, I operationalize this idea by measuring, for each city in Brazil, the covariance between local monthly average wind speeds and the national share of yearly deaths in a given month. This allows me to implement an instrumental variable strategy to examine the causal link between city-level mortality and Bolsonaro's vote share in 2022.

This study underscores the importance of public health outcomes in driving election results, and in particular how crisis, politicians' reactions and accountability intersect. My most conservative estimates indicate that a one-third reduction in excess mortality during the pandemic would have swayed the election result in favor of Bolsonaro. The results stand up to rigorous scrutiny, encompassing alternative specifications and robustness checks. Furthermore, I address and dismiss the notion that these findings are driven solely by incumbency effects, highlighting the unique nature of the relationship between excess mortality and the former president's electoral outcomes, which is suggestive that his attitudes and public persona took a role in shaping the voters' reaction.

I also provide a novel method for constructing counterfactuals, relying on neural networks, more specifically using a graph attention network architecture. I show how a simple model is capable of producing reliable estimates, and may thus be of use in other applications for which treated and control units are numerous but not immediately comparable. I see particular promise in using this method under a broader difference-in-differences framework when parallel trends are not satisfied.

It would also be interesting to evaluate whether the interplay between wind and excess deaths can be replicated in other parts of the world. If so, the strategy employed here could be extremely useful in examining the consequences of the pandemic at the city level across a wide variety of outcomes. Nonetheless, the work presented here makes the Brazilian case particularly promising for future studies.

In summary, this paper advances our comprehension of the connections between leaders' actions, their immediate impact and eventual electoral consequences. Partisan lenses may play a role in the voters' interpretation of events, but the evidence presented here highlights that politician's policy choices remain subject to democratic accountability.

References

Abatzoglou, J. T., Dobrowski, S. Z., Parks, S. A., & Hegewisch, K. C. (2018). TerraClimate, a high-resolution global dataset of monthly climate and climatic water balance from 1958–2015 [Number: 1 Publisher: Nature Publishing Group]. *Scientific Data*, *5*(1), 170191. https://doi.org/10.1038/sdata.2017.191

Ajzenman, N., Cavalcanti, T., & Da Mata, D. (2023). More than words: Leaders' speech and risky behavior during a pandemic. *American Economic Journal: Economic Policy*, *15*(3), 351–71. https://doi.org/10.1257/pol.20210284

Ali, Q., Raza, A., Saghir, S., & Khan, M. T. I. (2021). Impact of wind speed and air pollution on COVID-19 transmission in pakistan. *International Journal of Environmental Science and Technology*, *18*(5), 1287–1298. https://doi.org/10.1007/s13762-021-03219-z

Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., & Yang, D. (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics*, *191*, 104254. https://doi.org/10.1016/j.jpubeco.2020.104254

Andrews, I., Stock, J. H., & Sun, L. (2019). Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics*, *11*(1), 727–753. https://doi.org/10.1146/annurev-economics-080218-025643

Arroyo Abad, L., & Maurer, N. (2021, January 1). Do pandemics shape elections? retrospective voting in the 1918 spanish flu pandemic in the united states. Retrieved August 15, 2023, from https://papers.ssrn.com/abstract=3783893

Atkinson, J. (2009). *Natural ventilation for infection control in health-care settings* [Google-Books-ID: HID6NpLrwOsC]. World Health Organization.

Axfors, C., Schmitt, A. M., Janiaud, P., van't Hooft, J., Abd-Elsalam, S., Abdo, E. F., Abella, B. S., Akram, J., Amaravadi, R. K., Angus, D. C., Arabi, Y. M., Azhar, S., Baden, L. R., Baker, A. W., Belkhir, L., Benfield, T., Berrevoets, M. A. H., Chen, C.-P., Chen, T.-C., … Hemkens, L. G. (2021). Mortality outcomes with hydroxychloroquine and chloroquine in COVID-19 from an international collaborative meta-analysis of randomized trials [Number: 1 Publisher: Nature Publishing Group]. *Nature Communications*, *12*(1), 2349. https://doi.org/10.1038/s41467-021-22446-z

Baccini, L., Brodeur, A., & Weymouth, S. (2021). The COVID-19 pandemic and the 2020 US presidential election. *Journal of Population Economics*. https://doi.org/https://doi.org/10.1007/s00148-020-00820-3

Barrios, J. M., & Hochberg, Y. V. (2020). Risk perception through the lens of politics in the time of the COVID-19 pandemic. *National Bureau of Economic Research*. https://doi.org/https://doi.org/10.3386/w27008

Bayoumi, M. (2021). Improving indoor air quality in classrooms via wind-induced natural ventilation [Publisher: Hindawi]. *Modelling and Simulation in Engineering*, *2021*, e6668031. https://doi.org/10.1155/2021/6668031

Bazant, M. Z., & Bush, J. W. M. (2021). A guideline to limit indoor airborne transmission of COVID-19 [Publisher: Proceedings of the National Academy of Sciences]. *Proceedings of the National Academy of Sciences*, *118*(17), e2018995118. https://doi.org/10.1073/pnas.2018995118

Bisbee, J., & Honig, D. (2021). Flight to safety: COVID-induced changes in the intensity of status quo preference and voting behavior. *American Political Science Review*. https://doi.org/https://doi.org/10.1017/s0003055421000691

Bol, D., Giani, M., Blais, A., & Loewen, P. J. (2021). The effect of COVID-19 lockdowns on political support: Some good news for democracy? *European Journal of Political Research*, *60*(2), 497–505. https://doi.org/10.1111/1475-6765.12401

Borusyak, K., & Hull, P. (forthcoming). *Econometrica*.

Cabral, S., Ito, N., & Pongeluppe, L. (2021, April 28). The disastrous effects of leaders in denial: Evidence from the COVID-19 crisis in brazil. https://doi.org/10.2139/ssrn.3836147

Calvo, E. J., & Dias, J. (2021). Will i get COVID-19? partisanship, social media frames, and perceptions of health risk in brazil. *Latin American Politics and Society*. https://doi.org/https://doi.org/10.1017/lap.2020.30

Chernozhukov, V., Wüthrich, K., & Zhu, Y. (2021). An exact and robust conformal inference method for counterfactual and synthetic controls [Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/01621459.2021.1920957]. *Journal of the American Statistical Association*, *116*(536), 1849–1864. https://doi.org/10.1080/01621459.2021.1920957

Coccia, M. (2021). The effects of atmospheric stability with low wind speed and of air pollution on the accelerated transmission dynamics of COVID-19. *International Journal of Environmental Studies*, *78*(1), 1–27. https://doi.org/10.1080/00207233.2020.1802937

Conley, T. G. (1999). Gmm estimation with cross sectional dependence. *Journal of econometrics*, *92*(1), 1–45.

Ferejohn, J. (1986). Incumbent performance and electoral control [Publisher: Springer]. *Public Choice*, *50*(1), 5–25. Retrieved August 15, 2023, from https://www.jstor.org/stable/30024650

Figueira, G., & Moreno-Louzada, L. (2021). Influência de messias? relação intramunicipal entre preferências políticas e mortes em uma pandemia (messias' influence? intramunicipal relationship between political preferences and deaths in a pandemic). *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3849383

Fiorina, M. P. (1978). Economic retrospective voting in american national elections: A microanalysis [Publisher: [Midwest Political Science Association, Wiley]]. *American Journal of Political Science*, *22*(2), 426–443. https://doi.org/10.2307/2110623

Ghaffari, H. R., Farshidi, H., Alipour, V., Dindarloo, K., Azad, M. H., Jamalidoust, M., Madani, A., Aghamolaei, T., Hashemi, Y., Fazlzadeh, M., & Fakhri, Y. (2022). Detection of SARS-CoV-2 in the indoor air of intensive care unit (ICU) for severe COVID-19 patients and its surroundings: Considering the role of environmental conditions. *Environmental Science and Pollution Research*, *29*(57), 85612–85618. https://doi.org/10.1007/s11356-021-16010-x

Grossman, G., Kim, S., Rexer, J. M., & Thirumurthy, H. (2020). Political partisanship influences behavioral responses to governors' recommendations for COVID-19 prevention in the united states. *Proceedings of the National Academy of Sciences*, *117*(39), 24144–24153. https://doi.org/10.1073/pnas.2007835117

Gutierrez, E., Meriläinen, J., & Rubli, A. (2021, November 1). Electoral repercussions of a pandemic: Evidence from the 2009 h1n1 outbreak. https://doi.org/10.2139/ssrn.3667065

Habeebullah, T. M., Abd El-Rahim, I. H. A., & Morsy, E. A. (2021). Impact of outdoor and indoor meteorological conditions on the COVID-19 transmission in the western region of saudi arabia. *Journal of Environmental Management*, *288*, 112392. https://doi.org/10.1016/j.jenvman.2021.112392

Healy, A., & Malhotra, N. (2013). Retrospective voting reconsidered. *Annual Review of Political Science*, *16*(1), 285–306. https://doi.org/10.1146/annurev-polisci-032211-212920

Herrera, H., Ordoñez, G., Konradt, M., & Trebesch, C. (2020, September 10). Corona politics: The cost of mismanaging pandemics. https://doi.org/10.2139/ssrn.3690490

Izadyar, N., & Miller, W. (2022). Ventilation strategies and design impacts on indoor airborne transmission: A review. *Building and Environment*, *218*, 109158.

Kingma, D. P., & Ba, J. (2017, January 29). Adam: A method for stochastic optimization. https://doi.org/10.48550/arXiv.1412.6980

Lewis, D. (2022). Why the WHO took two years to say COVID is airborne [Bandiera_abtest: a Cg_type: News Feature Number: 7904 Publisher: Nature Publishing Group Subject_term: Diseases, SARS-CoV-2, Infection, Public health]. *Nature*, *604*(7904), 26–31. https://doi.org/10.1038/d41586-022-00925-7

Mariani, L. A., Gagete-Miranda, J., & Rettl, P. (2020). Words can hurt: How political communication can change the pace of an epidemic. *OSF Preprints*. https://doi.org/https://doi.org/10.31219/osf.io/ps2wx

Mendoza Aviña, M., & Sevi, S. (2021). Did exposure to COVID-19 affect vote choice in the 2020 presidential election? [Publisher: SAGE Publications Ltd]. *Research & Politics*, *8*(3), 20531680211041505. https://doi.org/10.1177/20531680211041505

Moazeni, M., Rahimi, M., & Ebrahimi, A. (2023). What are the effects of climate variables on COVID-19 pandemic? a systematic review and current update. *Advanced Biomedical Research*, *12*, 33. https://doi.org/10.4103/abr.abr_145_21

Morawska, L., Tang, J. W., Bahnfleth, W., Bluyssen, P. M., Boerstra, A., Buonanno, G., Cao, J., Dancer, S., Floto, A., Franchimon, F., Haworth, C., Hogeling, J., Isaxon, C., Jimenez, J. L., Kurnitski, J., Li, Y., Loomans, M., Marks, G., Marr, L. C., ... Yao, M. (2020). How can airborne transmission of COVID-19 indoors be minimised? *Environment International*, *142*, 105832. https://doi.org/10.1016/j.envint.2020.105832

Neyman, J. (1923). On the application of probability theory to agricultural experiments. essay on principles. *Ann. Agricultural Sciences*, 1–51.

Olak, A. S., Santos, W. S., Susuki, A. M., Pott-Junior, H., V. Skalny, A., Tinkov, A. A., Aschner, M., Pinese, J. P. P., Urbano, M. R., & Paoliello, M. M. B. (2022). Meteorological parameters and cases of COVID-19 in brazilian cities: An observational study [Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/15287394.2021.1969304]. *Journal of Toxicology and Environmental Health, Part A*, *85*(1), 14–28. https://doi.org/10.1080/15287394.2021.1969304

Olea, J. L. M., & Pflueger, C. (2013). A robust test for weak instruments [Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00401706.2013.806694]. *Journal of Business & Economic Statistics*, *31*(3), 358–369. https://doi.org/10.1080/00401706.2013.806694

Persson, T., Roland, G., & Tabellini, G. (1997). Separation of powers and political accountability*. *The Quarterly Journal of Economics*, *112*(4), 1163–1202. https://doi.org/10.1162/003355300555457

Razafindrakoto, M., Roubaud, F., Saboia, J., Castilho, M. R., & Pero, V. (2022). Municípios in the time of covid-19 in brazil: Socioeconomic vulnerabilities, transmission factors and public policies. *The European Journal of Development Research*, *34*(6), 2730–2758. https://doi.org/10.1057/s41287-021-00487-w

Rendana, M. (2020). Impact of the wind conditions on COVID-19 pandemic: A new insight for direction of the spread of the virus. *Urban Climate*, *34*, 100680. https://doi.org/10.1016/j.uclim.2020.100680

Rowe, B. R., Canosa, A., Meslem, A., & Rowe, F. (2022). Increased airborne transmission of COVID-19 with new variants, implications for health policies. *Building and Environment*, *219*, 109132. https://doi.org/10.1016/j.buildenv.2022.109132

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, *66*(5), 688.

Saldanha, R., Akbarinia, R., Valduriez, P., Pedroso, M., Ribeiro, V., Cardoso, C., Pena, E., & Porto, F. (2023). *Brclimr: Fetch zonal statistics of weather indicators for brazilian municipalities* [https://rfsaldanha.github.io/brclimr/, https://github.com/rfsaldanha/brclimr].

Szeidl, A., & Szucs, F. (n.d.). The political economy of alternative realities.

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018, February 4). Graph attention networks. https://doi.org/10.48550/arXiv.1710.10903

Wang, C. C., Prather, K. A., Sznitman, J., Jimenez, J. L., Lakdawala, S. S., Tufekci, Z., & Marr, L. C. (2021). Airborne transmission of respiratory viruses [Publisher: American Association for the Advancement of Science]. *Science*, *373*(6558), eabd9149. https://doi.org/10.1126/science.abd9149

Warshaw, C., Vavreck, L., & Baxter-King, R. (2020). Fatalities from COVID-19 are reducing americans' support for republicans at every level of federal office [Publisher: American Association for the Advancement of Science]. *Science Advances*, *6*(44), eabd8564. https://doi.org/10.1126/sciadv.abd8564

Xavier, D. R., Lima E Silva, E., Lara, F. A., E Silva, G. R. R., Oliveira, M. F., Gurgel, H., & Barcellos, C. (2022). Involvement of political and socio-economic factors in the spatial and temporal dynamics of COVID-19 outcomes in brazil: A population-based study. *Lancet Regional Health. Americas*, *10*, 100221. https://doi.org/10.1016/j.lana.2022.100221

9 Appendix

9.1 Excess mortality over time

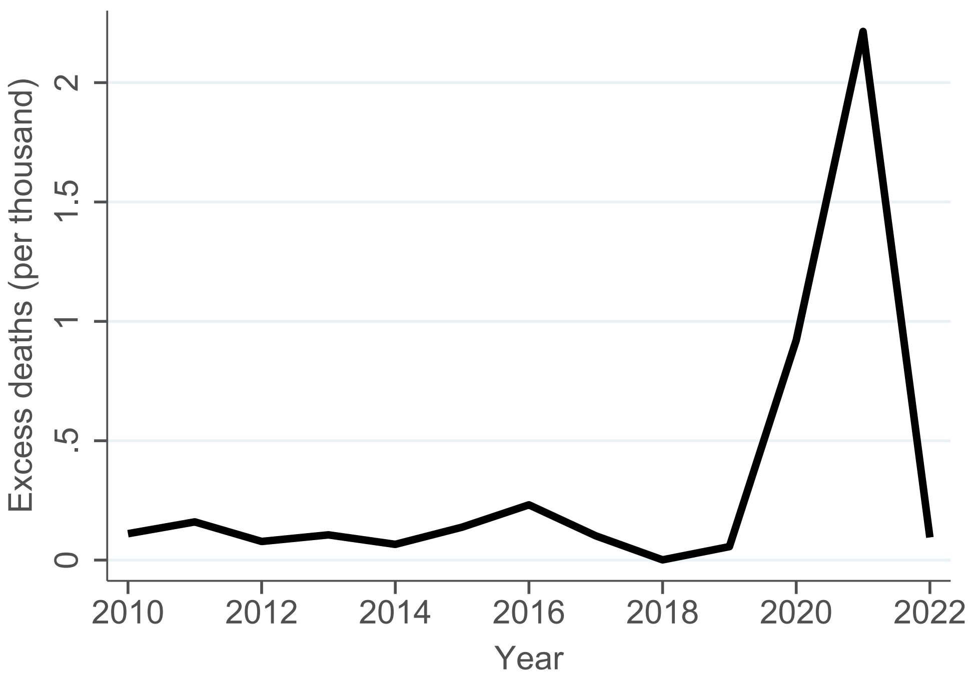

Figure 6: Excess mortality over time

9.2 Other meteorological factors

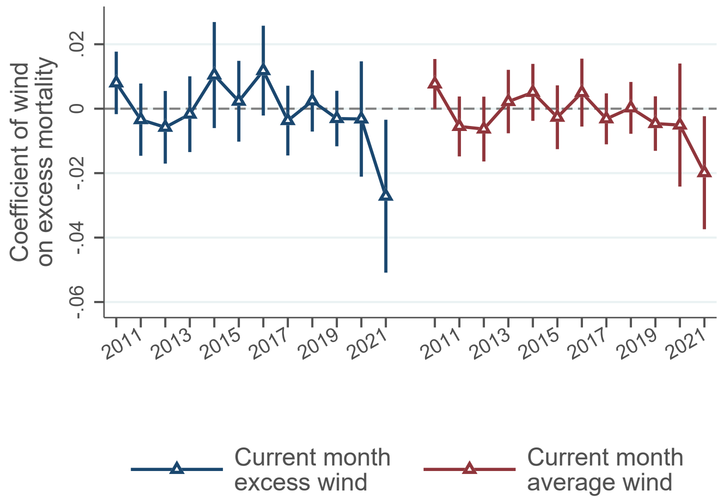

Figure 7: Excess mortality and wind, controlling for precipitation and temperature. Excess wind defined as wind speeds above a 12-year month-specific moving average.

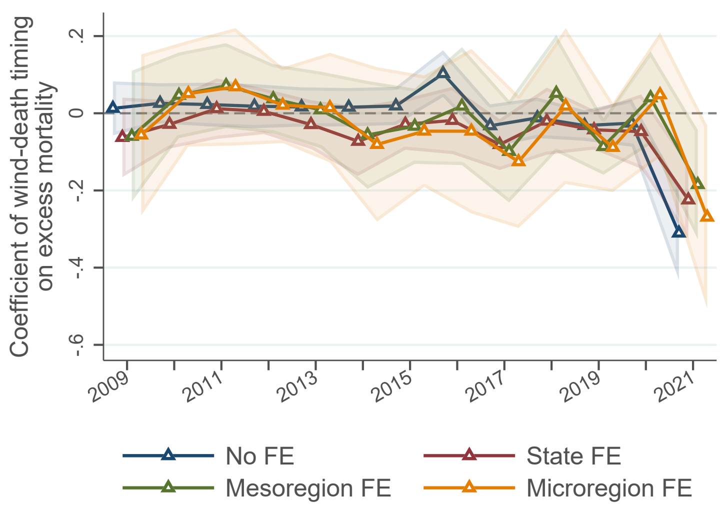

Figure 8: Wind-death timing and excess mortality, controlling for rain- and temperature-death timing.

9.3 WDT: geography and relationship with confounders

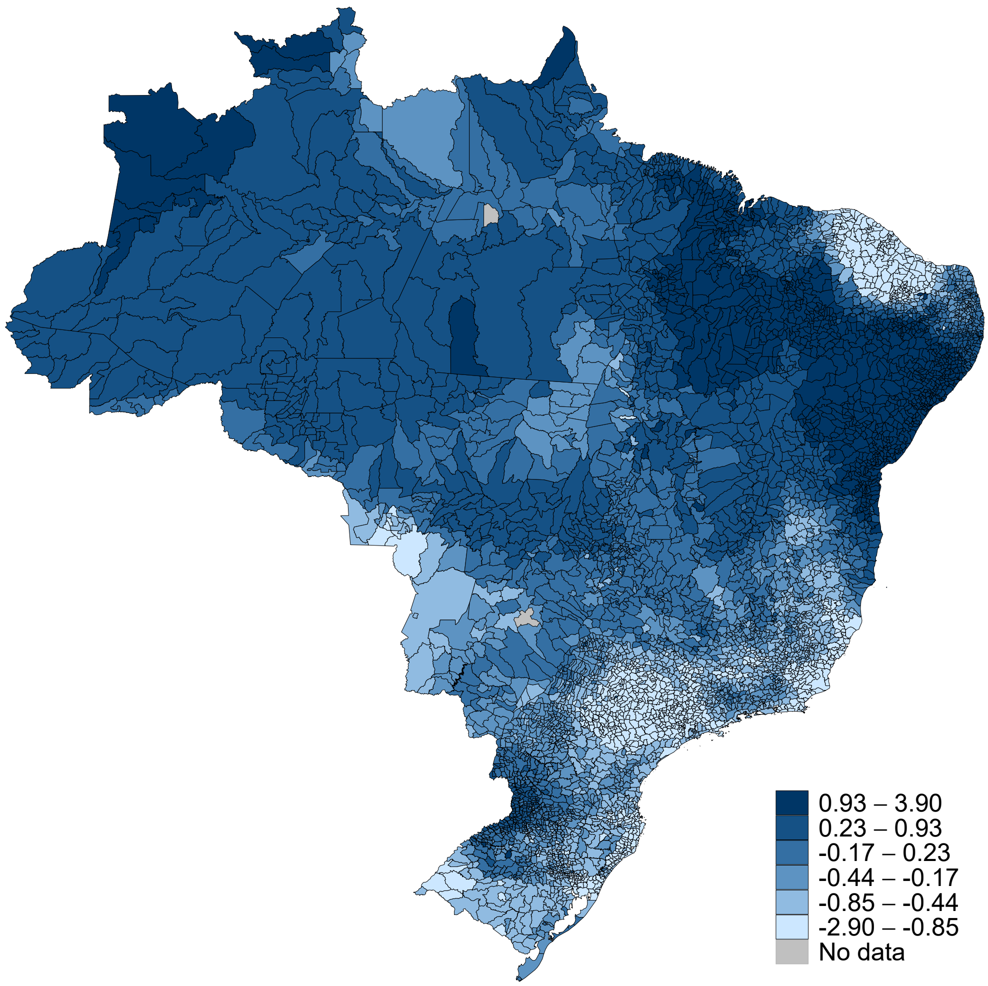

Figure 9: Geographic distribution of Wind-Death Timing

Table 3: Relationship of WDT with potential confounders

	Dep. var.: Wind-death timing			
PT vote share previous election	1.743***	1.231***	0.366***	0.038
	(0.413)	(0.297)	(0.122)	(0.080)
Evangelical share	-0.443	-0.330	0.063	-0.090
	(0.357)	(0.228)	(0.116)	(0.077)
Share of population living in an urban area	-0.320	-0.245**	-0.115*	-0.064*
	(0.194)	(0.099)	(0.064)	(0.034)
Share of population that owns a radio	0.249	0.761***	0.407*	0.100
	(0.396)	(0.268)	(0.215)	(0.107)
Share of population that owns a TV	-0.260	0.415	0.109	-0.035
	(0.590)	(0.291)	(0.188)	(0.161)
Average age of the population	-0.010	-0.007	-0.008	-0.010**
	(0.020)	(0.012)	(0.007)	(0.004)
Literacy rate	-0.898	0.160	0.039	-0.194
	(1.055)	(0.641)	(0.292)	(0.194)
Average family income (thousands of Reais)	0.299***	0.052	-0.050*	-0.015
	(0.084)	(0.049)	(0.029)	(0.016)
Share white	-0.706	-0.477*	-0.415**	-0.013
	(0.474)	(0.281)	(0.160)	(0.098)
Share born in municipality	-0.760***	-0.527***	-0.118	-0.037
	(0.205)	(0.172)	(0.114)	(0.054)
Hours worked (main job)	-0.030***	-0.009	0.001	0.001
	(0.010)	(0.005)	(0.003)	(0.002)
High school degree	0.363	0.772	0.576*	0.512***
	(1.075)	(0.559)	(0.302)	(0.193)
Further education	-4.023	2.515**	1.819**	0.252
	(3.082)	(1.268)	(0.841)	(0.414)
Average wind velocity	-0.136	-0.048	-0.096	-0.088
	(0.144)	(0.103)	(0.070)	(0.071)
Log pop. density	-0.058	-0.078***	-0.025	-0.013
	(0.040)	(0.029)	(0.018)	(0.009)
Log population	0.044	-0.003	-0.012	-0.009
	(0.050)	(0.022)	(0.013)	(0.008)
F-statistic	20.061	6.326	3.999	2.293
F p-value	0.000	0.000	0.000	0.003
State FE		✓	✓	✓
Mesorregion FE			✓	✓
Microrregion FE				✓

Errors clustered at the mesoregion level, except for the microregion FE specification, where errors are clustered at the level of the FE.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

9.4 NWDT: geography and relationship with confounders

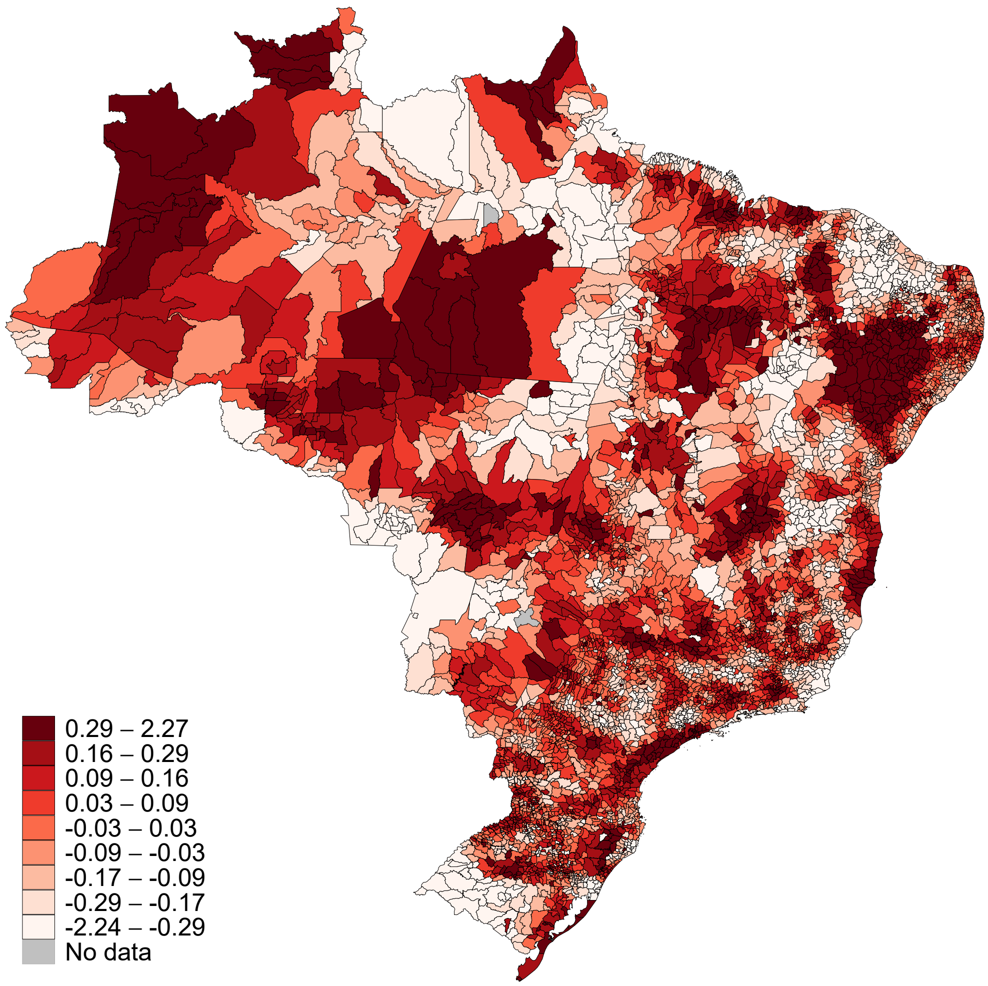

Figure 10: Geographic distribution of Net Wind-Death Timing

Table 4: Relationship of NWDT with potential confounders

	Dep. var.: Net wind-death timing			
PT vote share previous election	0.191*	0.094	0.008	-0.068
	(0.113)	(0.094)	(0.076)	(0.069)
Evangelical share	-0.030	-0.068	0.088	-0.024
	(0.094)	(0.088)	(0.083)	(0.066)
Share of population living in an urban area	-0.067	-0.115***	-0.077**	-0.057*
	(0.048)	(0.043)	(0.034)	(0.029)
Share of population that owns a radio	0.259**	0.117	0.121	0.065
	(0.111)	(0.108)	(0.096)	(0.090)
Share of population that owns a TV	0.103	0.056	-0.051	0.026
	(0.228)	(0.212)	(0.151)	(0.135)
Average age of the population	-0.011**	-0.012***	-0.008**	-0.009**
	(0.005)	(0.004)	(0.004)	(0.004)
Literacy rate	-0.132	0.017	0.048	-0.135
	(0.189)	(0.249)	(0.199)	(0.174)
Average family income (thousands of Reais)	-0.001	-0.004	-0.025	-0.017
	(0.022)	(0.022)	(0.020)	(0.015)
Share white	0.106	-0.002	-0.041	0.035
	(0.084)	(0.100)	(0.085)	(0.083)
Share born in municipality	-0.077	-0.092	-0.046	0.001
	(0.066)	(0.071)	(0.067)	(0.050)
Hours worked (main job)	0.002	0.003	0.003*	0.002
	(0.002)	(0.002)	(0.002)	(0.002)
High school degree	0.115	0.234	0.302	0.344**
	(0.244)	(0.227)	(0.187)	(0.154)
Further education	1.182**	1.119**	1.000**	0.261
	(0.542)	(0.507)	(0.442)	(0.360)
Average wind velocity	-0.073***	-0.079**	-0.031	-0.044
	(0.027)	(0.037)	(0.045)	(0.054)
Log pop. density	-0.006	-0.020*	-0.021**	-0.008
	(0.009)	(0.011)	(0.010)	(0.007)
Log population	-0.007	-0.001	-0.006	-0.009
	(0.012)	(0.010)	(0.008)	(0.006)
F-statistic	2.160	2.455	2.716	1.321
F p-value	0.009	0.003	0.001	0.179
State FE		✓	✓	✓
Mesorregion FE			✓	✓
Microrregion FE				✓

Errors clustered at the mesoregion level, except for the microregion FE specification, where errors are clustered at the level of the FE.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

9.5 Social distancing by locations

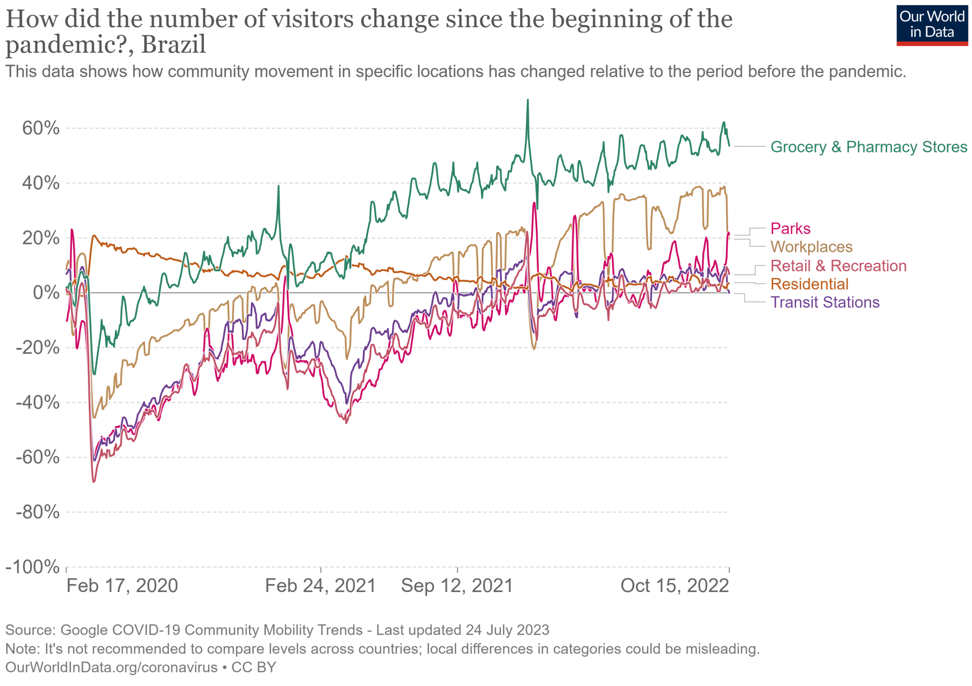

9.6 Variants over time

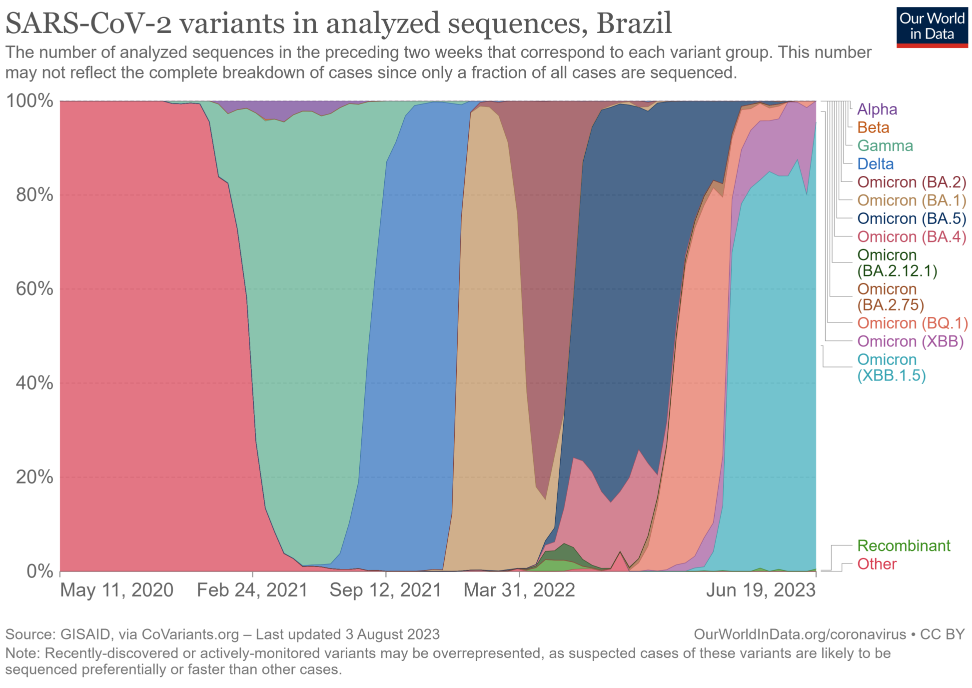

9.7 Three measures, same pattern: excess mortality

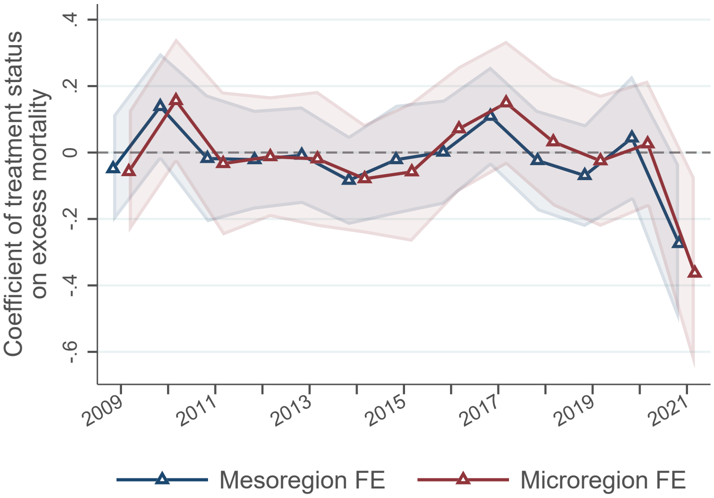

(a) Treatment status

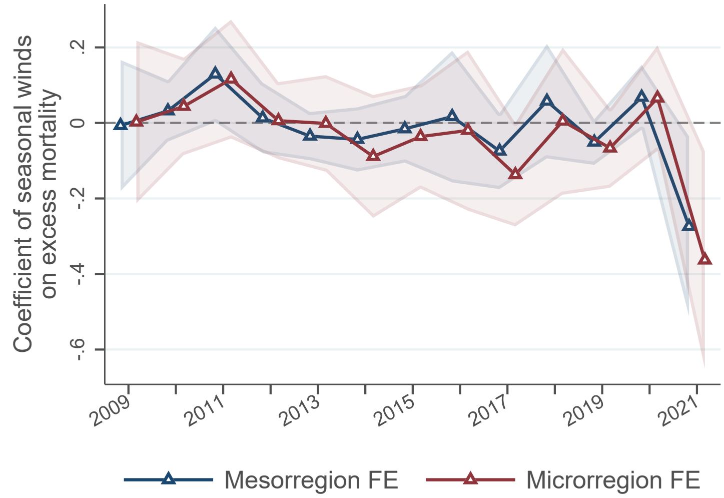

(b) Seasonal winds

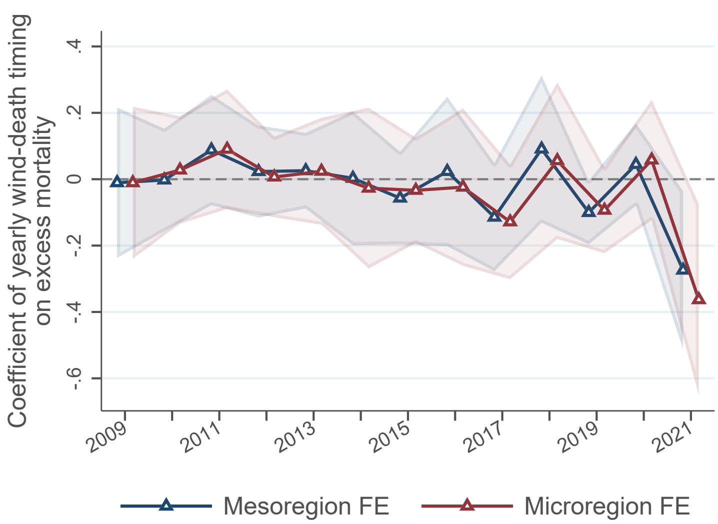

(c) Yearly wind-death timing

9.8 Robustness tests

9.8.1 Adding controls

Table 5: Impact of excess mortality on electoral changes (IV), including controls

	Dep. var.: Difference in PT vote share (2022-2018)	
Excess mortality 2021	0.019*	0.011*
	[0.001, .]	[0.001, 0.083]
Kleibergen-Paap F statistic	3.882	4.816
Number of municipalities	5560	5560
Mesorregion FE	✓	✓
Microrregion FE		✓
Full set of controls	✓	✓

Excess mortality in 2021 instrumented by the net covariance between leave-one-out monthly death shares at the national level and city-level monthly average wind. Clustered errors at the level of the FE's for the two specifications. Given the weakness of the instrument in the latter cases, I present Anderson-Rubin confidence intervals and p-values. For each specification, I include the full set of controls as per appendix (9.4). The dot in the mesoregion FE specification's confidence set indicates that the null of the coefficient being equal to 0 for some positive number cannot be rejected for all positive numbers.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

9.8.2 Using WDT instead of NWDT

Table 6: Impact of excess mortality on PT vote share (IV)

	Dep. var.: Difference in PT vote share (2022-2018)			
Excess mortality 2021	0.050***	0.017***	0.038***	0.011*
	[0.029, 0.135]	[0.007, 0.063]	[0.017, 0.187]	[0.001, 0.073]
Kleibergen-Paap F statistic	7.9	6.71	5.687	5.088
Number of municipalities	5562	5562	5560	5560
Mesorregion FE	✓	✓	✓	✓
Microrregion FE		✓		✓
Full set of controls			✓	✓

Excess mortality in 2021 instrumented by the raw (not netted from neighborhood averages) covariance between leave-one-out monthly death shares at the national level and city-level monthly average wind. Clustered errors at the level of the FE's for all the specifications. Given the weakness of the instrument, I present Anderson-Rubin confidence intervals and p-values.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

9.8.3 Varying neighborhood sizes

Table 7: Impact of excess mortality on electoral changes (IV), varying neighborhood size

	Difference in PT vote share (2022-2018)			
Panel I: neighborhood size 20				
Excess mortality 2021	0.011	0.007	0.008	0.007
	[-0.005, 0.061]	[-0.003, 0.033]	[-0.009, 0.089]	[-0.003, 0.037]
Kleibergen-Paap F statistic	6.0	6.20	5.074	5.975
Panel II: neighborhood size 40				
Excess mortality 2021	0.017*	0.010*	0.012	0.009
	[0.000, .]	[-0.001, 0.053]	[-0.006, .]	[-0.003, 0.063]
Kleibergen-Paap F statistic	5.5	5.55	4.428	5.037
Panel III: neighborhood size 80				
Excess mortality 2021	0.025**	0.013**	0.017*	0.011*
	[0.006, .]	[0.001, 0.066]	[-0.001, .]	[-0.001, 0.084]
Kleibergen-Paap F statistic	5.1	5.53	3.934	4.830
Panel IV: neighborhood size 100				
Excess mortality 2021	0.027***	0.013**	0.019*	0.011*
	[0.007, .]	[0.001, 0.066]	[0.000, .]	[-0.001, 0.084]
Kleibergen-Paap F statistic	5.1	5.57	3.882	4.816
Panel V: neighborhood size 150				
Excess mortality 2021	0.031***	0.014**	0.023**	0.012*
	[0.010, .]	[0.001, 0.081]	[0.002, .]	[-0.001, .]
Kleibergen-Paap F statistic	4.7	5.19	3.577	4.433
Mesorregion FE	✓	✓	✓	✓
Microrregion FE		✓		✓
Full set of controls			✓	✓
Number of municipalities	5562	5562	5560	5560

Excess mortality in 2021 instrumented by the covariance between leave-one-out monthly death shares at the national level and city-level monthly average wind, net of neighbordhood averages of varying size. Clustered errors at the level of the FE's for all the specifications. Given the weakness of the instrument, I present Anderson-Rubin confidence intervals and p-values. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

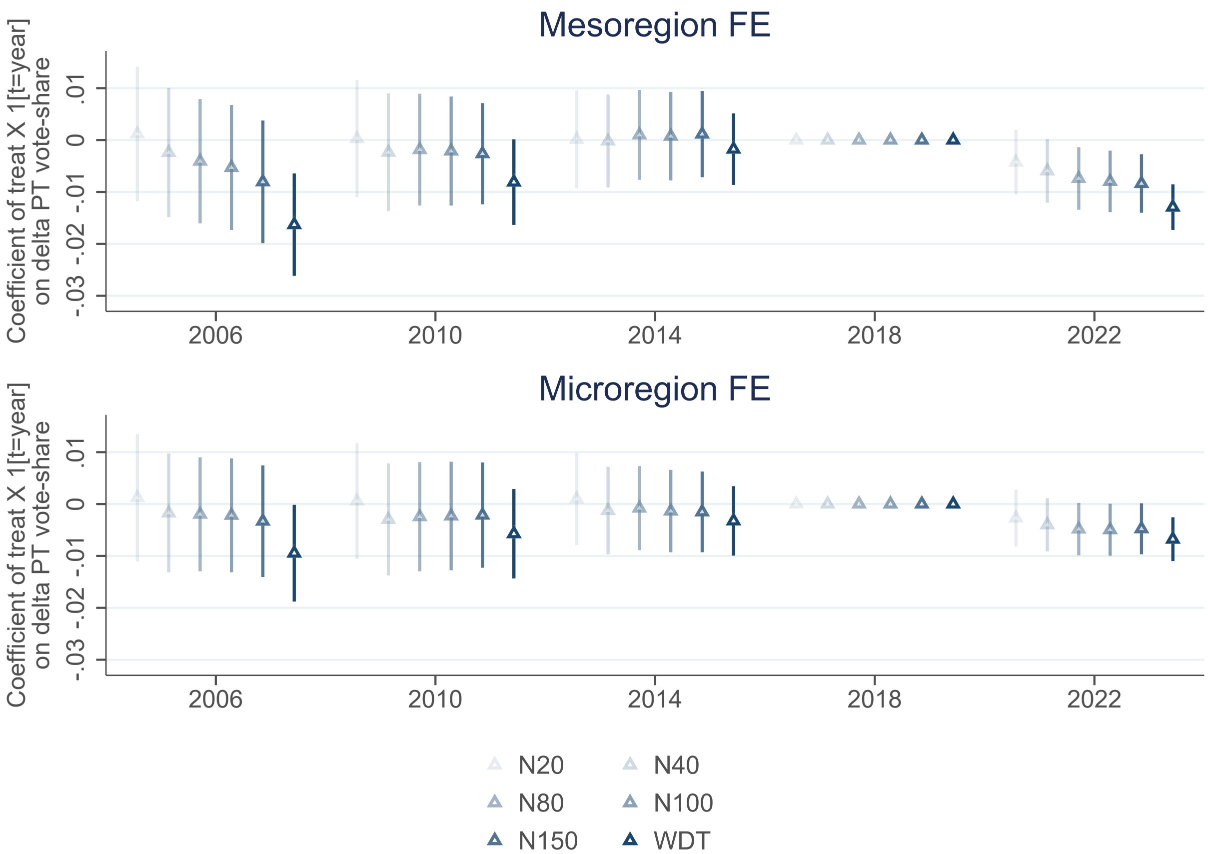

Figure 11: Coefficient of interest for various neighborhood size definitions

9.8.4 COVID case and death rates as main independent variable

Table 8: Impact of COVID case and death rates on electoral changes (IV)

	Dep. var.: Difference in PT vote share (2022-2018)			
Panel I: COVID death rate as main independent variable				
Covid death rate 2021	0.089***	0.036***	0.086***	0.033*
	[0.049, 0.260]	[0.013, 0.101]	[0.033, .]	[-0.001, .]
Kleibergen-Paap F statistic	7.4	7.45	3.300	3.484
Number of municipalities	5562	5562	5560	5560
Panel II: COVID case rate as main independent variable				
Covid case rate 2021	0.002***	0.001***	0.002***	0.001
	[0.001, 0.005]	[0.001, 0.022]	[0.001, 0.015]	[., .]
Kleibergen-Paap F statistic	10.6	3.75	5.102	2.464
Number of municipalities	5562	5562	5560	5560
Mesorregion FE	✓	✓	✓	✓
Microrregion FE		✓		✓
Full set of controls			✓	✓

COVID death and case rates in 2021 - defined as city deaths and cases per thousand - instrumented by the (non-net) covariance between leave-one-out monthly death shares at the national level and city-level monthly average wind. Clustered errors at the level of the FE's for all the specifications. Given the weakness of the instrument, I present Anderson-Rubin confidence intervals and p-values.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

9.8.5 Presidential runoff results in the subsection of states in which incumbent ran for re-election

Table 9: Impact of excess mortality on worker's party vote-share (IV)

	Dep. var.: Difference in PT vote share (2022-2018)	
Panel I: effect in states with governors running for re-election		
Excess mortality 2021	0.026***	0.012**
	[0.009, 0.117]	[0.003, 0.051]
Kleibergen-Paap F statistic	6.549	6.172
Number of municipalities	4384	4384
Panel II: excluding Bolsonaro-aligned governors		
Excess mortality 2021	0.031***	0.017***
	[0.013, 0.193]	[0.005, 0.099]
Kleibergen-Paap F statistic	5.690	5.016
Number of municipalities	3915	3915
Mesorregion FE	✓	✓
Microrregion FE		✓

Excess mortality in 2021 instrumented by the net covariance between leave-one-out monthly death shares at the state level and excess wind. Clustered errors at the level of the FE's. Given the weakness of the instrument, I present Anderson-Rubin confidence intervals and p-values. Bolsonaro-aligned governors are defined as those from Republicanos and PP.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

9.8.6 Alternative measure of excess deaths

Table 10: Impact of excess mortality on electoral changes (IV)

	Dep. var.: Difference in PT vote share (2022-2018)	
Excess mortality 2021	0.025***	0.012**
	[0.007, 0.121]	[0.003, 0.043]
Kleibergen-Paap F statistic	6.352	7.126
Number of municipalities	5562	5562
Mesorregion FE	✓	✓
Microrregion FE		✓

Excess mortality in 2021 is defined here as deaths above the three-year average number of deaths in each city from 2017 to 2019 (I exclude 2020 since the pandemic had already started then). I instrument it with the net covariance between leave-one-out monthly death shares at the national level and city-level monthly average wind. Clustered errors at the level of the FE's for the two specifications. Given the weakness of the instrument, I present Anderson-Rubin confidence intervals and p-values.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

9.8.7 Using first round results

Table 11: Impact of excess mortality on electoral changes (IV)

	Dep. var.: Difference in PT first round vote share (2022-2018)			
Excess mortality 2021	0.039**	0.023**	0.032*	0.019*
	[0.005, 0.390]	[0.005, 0.131]	[-0.003, .]	[0.001, 0.177]
Kleibergen-Paap F statistic	5.1	5.57	3.882	4.816
Number of municipalities	5562	5562	5560	5560
Mesorregion FE	✓	✓	✓	✓
Microrregion FE		✓		✓
Full set of controls			✓	✓

Excess mortality in 2021 instrumented by the net covariance between leave-one-out monthly death shares at the national level and city-level monthly average wind. Clustered errors at the level of the FE's for all the specifications. Given the weakness of the instrument, I present Anderson-Rubin confidence intervals and p-values.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

9.8.8 Controlling for the 5-year moving average of NWDT

Table 12: Impact of excess mortality on electoral changes (IV)

	Dep. var.: Difference in PT vote share (2022-2018)			
Excess mortality 2021	0.026***	0.013**	0.018*	0.011*
	[0.009, 0.135]	[0.003, 0.061]	[0.001, .]	[-0.001, 0.089]
Kleibergen-Paap F statistic	5.9	5.84	4.080	4.647
Number of municipalities	5562	5562	5560	5560
Mesorregion FE	✓	✓	✓	✓
Microrregion FE		✓		✓
Full set of controls			✓	✓
5-year moving average of NWDT	✓	✓	✓	✓

Excess mortality in 2021 instrumented by the net covariance between leave-one-out monthly death shares at the national level and city-level monthly average wind. Clustered errors at the level of the FE's for all the specifications. Given the weakness of the instrument, I present Anderson-Rubin confidence intervals and p-values.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

9.8.9 Main results by population quintile

Table 13: Impact of excess mortality on electoral changes (IV), by population quintile

	Dep. var.: Difference in PT vote share (2022-2018)			
Excess mortality 2021 $\times PopQuint_1$	0.049***	0.015*	0.040**	0.012
	(0.018)	(0.008)	(0.018)	(0.008)
Excess mortality 2021 $\times PopQuint_2$	0.051**	0.014	0.040*	0.010
	(0.021)	(0.009)	(0.021)	(0.009)
Excess mortality 2021 $\times PopQuint_3$	0.063***	0.019*	0.049*	0.013
	(0.024)	(0.011)	(0.025)	(0.011)
Excess mortality 2021 $\times PopQuint_4$	0.055***	0.017*	0.043**	0.011
	(0.020)	(0.009)	(0.022)	(0.009)
Excess mortality 2021 $\times PopQuint_5$	0.055***	0.016*	0.043**	0.011
	(0.020)	(0.009)	(0.021)	(0.009)
Anderson-Rubin p-values	.0004	.0731	.0171	.3812
Number of municipalities	5562	5562	5560	5560
Mesorregion FE	✓	✓	✓	✓
Microrregion FE		✓		✓
Full set of controls			✓	✓
Population quintile FE	✓	✓	✓	✓

The interaction between Excess mortality in 2021 and population quintiles is instrumented by the (non-net) covariance between leave-one-out monthly death shares at the national level and city-level monthly average wind, interacted with population quintiles. Clustered errors at the level of the FE's for all the specifications. Given the weakness of the instrument, I present Anderson-Rubin p-values testing the hypothesis that the instrumented terms' coefficients are jointly equal to 0. $PopQuint_1$ has a mean population of roughly 3000 and $PopQuint_5$, 150000. P-values for each individual interaction should be interpreted with caution as they are not robust to weak instrumentation.

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.