WIDER Working Paper 2024/44

Surviving in the dark: the mortality effects of reducing rolling blackouts

Joshua Budlender*

July 2024

United Nations University World Institute for Development Economics Research
wider.unu.edu

Abstract: South Africa frequently experiences rolling blackouts ('load shedding') due to shortfalls in electricity generation. This is a common problem across the developing world, and yet the developmental impacts of insufficient and unstable electricity supply, and the benefits of mitigating this, are poorly understood. I use the introduction of a unique load shedding reduction policy in parts of South Africa's second-largest city, Cape Town, to investigate the mortality effects of load shedding and its mitigation. To identify these effects, I use a stacked synthetic control design that leverages the episodic nature of load shedding between 2014 and 2019. While the estimates are imprecise, I find robust evidence that the mitigation policy statistically significantly reduces mortality in Cape Town relative to other parts of South Africa experiencing unmitigated load shedding. The incomplete geographic coverage of the mitigation policy entrenches existing inequalities in the city.

Key words: rolling blackouts, mortality, load shedding, synthetic control, Cape Town, South Africa

JEL classification: I15, O13, Q48, C23

Acknowledgements: For comments and suggestions on earlier versions of this work, I thank Arindrajit Dube, Daniele Girardi, Matt Woerman, Nicholas Budlender, Philippe Burger, Richard Kima, and C. Friedrich Kreuser. For sharing data, I thank Pamela Groenewald and Tom Moultrie from the SAMRC, Herman Maritz and Dan Wells from EskomSePush, the City of Cape Town Open Data Portal team, and Adrian Frith. I gratefully acknowledge support from UNU-WIDER as part of the SA-TIED programme.

---

* Southern Africa Labour and Development Research Unit (SALDRU), University of Cape Town (UCT), Cape Town, South Africa, joshua.budlender@uct.ac.za

This study has been prepared within the UNU-WIDER project Southern Africa—Towards Inclusive Economic Development (SA-TIED).

Information and requests: publications@wider.unu.edu

ISSN 1798-7237 ISBN 978-92-9267-506-6

https://doi.org/10.35188/UNU-WIDER/2024/506-6

Typescript prepared by Gary Smith.

United Nations University World Institute for Development Economics Research provides economic analysis and policy advice with the aim of promoting sustainable and equitable development. The Institute began operations in 1985 in Helsinki, Finland, as the first research and training centre of the United Nations University. Today it is a unique blend of think tank, research institute, and UN agency—providing a range of services from policy advice to governments as well as freely available original research.

The Institute is funded through income from an endowment fund with additional contributions to its work programme from Finland and Sweden, as well as earmarked contributions for specific projects from a variety of donors.

Katajanokanlaituri 6 B, 00160 Helsinki, Finland

1 Introduction

Semi-regular power outages are extremely common across the world, in the last 10 years affecting countries representing more than 50% of the global population.[1] What are the implications for development, and what can be done to mitigate negative effects? The question of how and whether electrification 'supercharges' development has attracted significant attention from development economists (Lee et al. 2020a), but this mostly focuses on new electrification *expansions*, and insufficiently considers the impact of unstable and intermittent electricity supply, which is the more prevailing condition across substantial parts of the developing world. South Africa, which frequently experiences these rolling blackouts (called 'load shedding'), provides a useful setting to examine this question due to the introduction of a unique load shedding *reduction* policy in parts of Cape Town, the country's second-largest city. I examine the effects of the policy on premature mortality, a key developmental outcome which is under-investigated in the existing literature.

Leveraging the episodic nature of load shedding in South Africa between 2014 and 2019, I use a stacked synthetic control design to estimate the mortality effects of the load shedding reduction policy in Cape Town. I find that Cape Town experiences statistically significantly lower mortality when load shedding is mitigated. This finding is robust when restricted to the age 65+ population, which is consistent with the prior literature. A variety of robustness and placebo tests support this mortality reduction being directly due to the mitigation policy, rather than differential load shedding responses or artefacts from the synthetic control approach. The estimates are, however, very imprecise: while a robust statistically significant effect is apparent, the 95% confidence intervals are very wide.

After converting load shedding severity into hours without power for the typical household, the point estimates suggests that each (severity-weighted) hour of averted load shedding prevents approximately 16 premature deaths across the city among the age 65+ population, though the 95% confidence intervals include both 3 and 29 deaths averted. The point estimate implies that the mitigation policy prevented approximately 547 premature deaths in Cape Town between 2014 and 2019 for this population (95% confidence intervals of 94 and 1,002). There were 163 severity-weighted hours of load shedding in Cape Town from 2014 to 2019, so this estimate implies that load shedding over this period caused 2,623 deaths among people 65 and older in the city-supplied region (95% confidence intervals of 452 and 4,805). I discuss the conditions under which these mortality magnitudes could be extrapolated beyond the 2019 period, or used to calculate mortality effects of each additional hour of load shedding across the city and nationwide. I show that the policy experiment identifies a particular local average treatment effect which is likely inappropriate to generalize across these dimensions. As part of this discussion, I characterize the regions of Cape Town that receive the mitigation policy, and show that they are older, richer, and (much) whiter than the rest of the city, raising equity considerations and suggesting that the uneven implementation of the policy exacerbates existing inequalities in Cape Town.

That the mortality effects of load shedding and its mitigation are particularly concentrated among the elderly is consistent with the limited existing literature on the mortality effects of electricity rationing from developed countries, which finds that power rationing disproportionately increases mortality among the elderly (Chirakijja et al. 2024; He and Tanaka 2023). This is somewhat unsurprising: age is a major marker of mortality vulnerability, particularly from illness, and one would expect marginal mortality effects to be concentrated among the most vulnerable groups.

What are the mechanisms by which load shedding should increase mortality? Evidence from developed-country contexts suggests an inability to regulate ambient temperate is a potentially significant factor,

[1] Reviewing the rolling blackout record of the 10 most populous developing countries, I count India, China, Indonesia, Pakistan, Nigeria, Bangladesh, Ethiopia, and the Philippines (plus South Africa) as recently experiencing load shedding.

with heat and cold stress increasing mortality especially for the elderly. While South Africa does not experience the temperature extremes of a country such as the United States, large and densely populated parts of the country are at high altitudes and regularly experience lows of around 2 °C (35 °F) in winter, and summer high temperatures can be extreme. This is exacerbated by homes typically being poorly insulated and households frequently relying on electric space heating.[2] But the harmful effects of electricity outages on health are not limited to inability to regulate temperature. Household supplementary oxygen machines or other health equipment such as nebulizers or CPAP machines may stop functioning, particularly during extended outages.[3] In the home, the risk of foodborne disease may increase as food spoils more quickly with inconsistent refrigeration (particularly with cheaper or older refrigerators, which are common), and household sanitation becomes more difficult (Laher et al. 2019). At the municipal level, bulk water treatment and distribution infrastructure relies on electricity to function, and extended outages can limit or stop household water supply even after electricity is restored, as water systems take time to recover.[4] Quality of medical care is also affected, as public hospitals and clinics often have limited, if any, generator capacity and have been known to turn away patients during load shedding.[5] Even in cases where generators are available, elective surgery is often postponed when the hospital is on back-up power (Laher et al. 2019). Cellphone towers sometimes fail when outages are extended, hampering communications and emergency response (e.g. Gavaza 2022; Odendaal 2023). Non-functioning street lights and traffic lights may increase road accidents, while it is plausible that lack of lighting and failure of electronic alarm systems may increase crime (e.g. Buthelezi 2022). Lawson (2022) finds that load shedding causes residential fires, likely due to power surges at reconnection as well as substitution towards combustion-based energy sources in the household.

The mortality data used in this paper lacks detailed cause- or location-of-death information, but does distinguish between natural and unnatural deaths, and it is clear that the mortality results are overwhelmingly but not disproportionately due to effects on natural deaths and supporting causes associated with illness and its treatment, but allowing for increases in non-natural deaths too.

This paper contributes to a number of different literatures. First, while the electricity shortage in South Africa is widely recognized as a national crisis, there is very little credibly causally identified evidence as to its effects. Existing attempts to estimate the economic cost of load shedding in South Africa typically either use an accounting-based ‘cost of unserved energy’ measure (Minnaar et al. 2017; Morema et al. 2019) that imposes very strong assumptions about firms’ production functions, or runtime series regressions (South African Reserve Bank 2019, 2022; Walsh et al. 2021) that do not seem to appropriately account for potentially confounding seasonal and secular trends. In terms of health impacts, Gehringer et al. (2018) find that load shedding increased paediatric hospital admissions, while Lawson (2022) finds that load shedding increased residential fires. As far as I am aware, this is the first paper to leverage the load shedding mitigation policy in Cape Town to implement a design that results in plausibly causal estimates. This also would make this paper the first to evaluate the effects of Cape Town’s mitigation policy using modern causal inference methods, and I believe it is the first paper to document how the incomplete geographic coverage of the policy entrenches racial and income inequalities in the City.

[2] The 2019 General Household Survey shows that 48% of households use mains electricity as the main energy source for space heating. Residential AC use is rare (only 6% of households).

[3] See, for example, Pijoos (2022). Motloung (2023) reports an episode in which a child died after her breathing machine could not function due to load shedding.

[4] See, for example, McCain (2022). Mthethwa (2023b) and SABC (2023) report an episode in which water shortages due to rolling blackouts were understood to have contributed to five farm workers’ heatstroke-related deaths in South Africa’s Northern Cape.

[5] See, for example, Khan and Mkentane (2023) and Monama (2022). Copelyn (2022), Ginindza (2023), JacaNews (2015), and Mthethwa (2023a) relay accounts of patient deaths understood to be caused directly by load shedding.

There is more quasi-experimental international evidence on the effects of electricity rationing, but this is still very limited when it comes to effects on health. Key references for the effects of intermittent electricity supply on firm performance and employment include Fisher-Vanden et al. (2015) on China, Allcott et al. (2016) and Abeberese (2017) on India, Abeberese (2020), Abeberese et al. (2021), and Hardy and McCasland (2021) on Ghana, and Mensah (2024) on a selection of African countries. When it comes to mortality, it seems that the sole study in a developing country is Apenteng et al. (2018), who study the effects of intermittent supply in Ghanaian health facilities and find significant mortality effects.[6] Neidell et al. (2021) and He and Tanaka (2023) study the mortality effects of post-Fukushima electricity rationing in Japan, and Chirakijja et al. (2024) examine the effects of winter heating prices on mortality in the United States; all find substantial mortality effects linked to temperature extremes, and He and Tanaka (2023) and Chirakijja et al. (2024) find that their results are driven by disproportionate mortality among the elderly.

It is striking that the literature on the mortality effects of electricity rationing in developing countries is so limited. There are many reasons why the effects of blackouts in developing countries, which is where they mainly occur, might be quite different to rationing effects in the developed world. Age is a major marker of mortality vulnerability, and developing and developed countries typically have very different demographic profiles. The existing literature also identifies temperature as an important mediator, and climates in developing countries, especially in Africa and Latin America, are often very different to those in the 'global North'. Developed-economy societies may be more reliant on reliable electricity supply and consequently face greater mortality sensitivity when supply is rationed, both because greater wealth means greater use of energy-intensive goods that protect health (e.g. air conditioning), and because the frequency of outages in developing countries may encourage adaptive measures absent in the developed world. Lastly, the supply rationing examined in the developed world is more likely to be rationing on the basis of a price mechanism or moral suasion (this is the case for all of Neidell et al. (2021), He and Tanaka (2023), and Chirakijja et al. (2024)), allowing households and healthcare facilities some choice over when and how to reduce electricity consumption; the incidence of rolling blackouts in the developing world typically allows no such discretion. The South African case therefore would add significantly to the literature, especially given that the design discussed here examines all-cause mortality inside and outside healthcare facilities, and has a clear causal design.

More broadly this paper contributes to a large development literature that examines whether electrification 'supercharges' development (Lee et al. 2020a). However, most of this larger work has focused on electrification *expansions*, particularly in rural areas.[7] Given the reality of widespread rolling blackouts, an important part of the electrification–development question is the effect of insufficient and unstable supply for places already connected to an electricity grid. Compared to electrification expansion policy, intermittent supply is more likely to affect urban populations, vary seasonally, and pose unique challenges for systems designed on the assumption of reliable electricity provision. This raises the question of distinct benefits of load shedding mitigation. As far as I am aware, this is the first paper to study the effects of a load shedding mitigation policy that reduces but does not eradicate the issue of intermittent supply. In this respect it is notable that the Cape Town scheme may be more reproducible than one might imagine: it does not mostly rely on additional generation, but is centred on a pre-existing pumped-storage facility that stores potential energy during periods of low demand when the generation

[6] Indeed, a fairly recent review of the public health literature (Irwin et al. 2020) directly states that Apenteng et al. (2018) is the only article to examine the mortality effects of intermittent electricity supply in low- and middle-income countries. Causal identification relies on facility-specific outages being as good as randomly designed, the plausibility of which the authors do not discuss; they explicitly disclaim causal identification.

[7] See, for example, Dinkelman (2011) on South Africa, Rud (2012), Burlig and Preonas (2016), Aklin et al. (2017), and Thomas et al. (2020) on India, Grogan and Sadanand (2013) on Nicaragua, Lipscomb et al. (2013) on Brazil, Grimm et al. (2017) on Rwanda, and Lee et al. (2020b) on Kenya. Lee et al. (2020a) provide a recent review.

shortfall is lower, and then redistributes to periods of high demand when the generation shortfall is more severe.

In Section 2 I start by explaining how load shedding is implemented in South Africa and the mitigation policy in Cape Town, and then I discuss the different data sources in Section 3. In Section 4 I outline the stacked synthetic control methodology, and present the main synthetic control results in Section 5, where I also discuss various robustness and placebo tests. Section 6 presents reduced form and local average treatment effect estimates for the marginal effect of an hour of load shedding reduction on mortality, and discusses how these results may or may not be extrapolated out-of-sample. It ends by comparing the magnitudes in this paper to those found in the existing literature. I conclude in Section 7 and note some open questions for future research, with a particular focus on the potential inequality effects of the mitigation policy.

2 Load shedding implementation

Electricity in South Africa is overwhelmingly generated and transmitted by the national state-owned power utility, Eskom. Eskom also performs the bulk of electricity distribution, though some municipalities have their own publicly owned distribution companies, which nevertheless source almost all of their power from Eskom. When Eskom experiences a shortfall in electricity generation relative to demand, it has a number of tools to reduce this demand. Broadly, these can be categorized as '(industrial) load *curtailment*' and 'load *shedding*'.[8] Load curtailment is when Eskom instructs large power users (typically large industrial users, with direct Eskom connections) to reduce their energy use, for which they receive compensation at a regulated rate. Load shedding is when the system operator implements scheduled and staggered rolling blackouts across the country, including areas where electricity is distributed by municipalities. This paper focuses on load shedding, both because the bulk of Eskom's 'demand response' programme is implemented via load shedding rather than load curtailment, and also because it is the directly household-facing component.[9]

2.1 National incidence over time

Until recently these rotating outages were clustered and somewhat intermittent, with load shedding in 2007/08, 2014/15, and 2018/19.[10] However, after 2019 load shedding became much more frequent and severe, and in 2022 and 2023 was almost permanently in effect (Figure 1). Out of all severity-weighted load shedding hours between 2014 and October 2023, 92% were implemented after 2019.[11] The dramatic increase in load shedding severity after late 2019 meant that it was experienced as a qualitatively different phenomenon than it was before, and it was accompanied by much more substantial load shedding amelioration efforts by households, the state, and private business. This includes investment in back-up power supplies of various kinds, as well as the state removing barriers associated with private sector energy generation (see Section 6).

[8] I include virtual power station (VPS) in load curtailment. Interruptible load shed (ILS) does not fit neatly into either category but is quantitatively and qualitatively unimportant for the purposes of this paper.

[9] Data on load curtailment is in any case much more limited than that on load shedding.

[10] Though some regional outages were experienced in 2005 and 2006 due to local power station failures and insufficient cross-region transmission capacity, the 2007/08 outages were the first national shortages since the early 1980s.

[11] This weighting method is discussed below.

Figure 1: National post-2008 load shedding hours up to May 2023, by stage

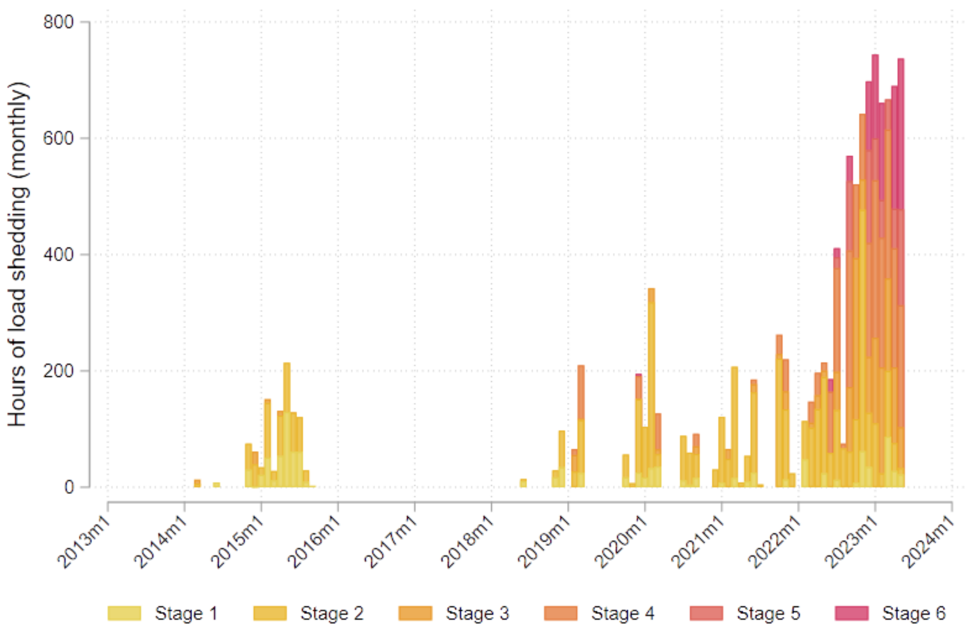

Note: the figure shows national load shedding implementation in South Africa from January 2013 to May 2023. Load shedding is implemented in a staggered fashion such that each stage of load shedding corresponds to an additional 1.5 hours of actual cut-off for the typical household for each 24 hours of load shedding at that stage—see Section 2.
Source: author's compilation based on EskomSePush historic load shedding schedule and the author's collection of Eskom load shedding alert tweets.

For reasons discussed in Section 4, it is difficult to estimate load shedding mortality effects in the post-2019 period. For this reason I restrict my analysis to 2014–19, when load shedding was more episodic. In Section 6 I discuss the conditions under which the results from 2014–19 could be extended to the post-2019 period; in general I am sceptical of the appropriateness of this kind of extrapolation, due to the qualitative distinctness of load shedding across these periods.

While it is less evident when load shedding is almost permanent, in previous years a clear seasonal pattern was evident, with load shedding being more severe and more likely in the summer months. While demand is lower in these months, plant maintenance is higher in preparation for the winter demand peak, and unplanned capacity loss also increases, apparently because some of the coal-fired plants perform worse in hot weather.

2.2 Implementation details

Load shedding is implemented relatively uniformly across the country, with the major exception of the City of Cape Town since 2015, which is discussed in more detail below.[12] However, the general pattern is that load shedding entails electricity outages for 2-, 3-, or 4-hour slots staggered over the day across sub-municipal geographic areas, called 'blocks'. Outages are implemented according to pre-existing schedules, which indicate the times a particular block will be cut-off for each 'stage' of load shedding—examples are provided in Figure A1. The stages reflect the severity of load shedding implementation. Nationally, each stage comprises approximately 1,000 MW being shed from the grid, with higher stages meaning each block will experience more outages in a given period.

[12] There are a few other smaller exceptions, also discussed below.

The general pattern implemented by Eskom and followed by the municipalities is that each 'stage' of load shedding corresponds to the typical block having no electricity for 1.5 out of 24 hours, so that stage 2 means on average 3 hours off per day for a given block, while stage 4 means 6 hours off per day for a given block.[13] While this general pattern is always exactly followed by Eskom and almost always exactly followed by the various municipalities (and, importantly, it is exactly followed in Cape Town), municipal implementation is not *completely* uniform.[14] However, where there are deviations from the national standard, they seem to always result in (usually slightly) *less* load shedding, and therefore do not pose a threat to this study's main findings: if they do cause a bias, it will be in the direction of attenuating the results.[15]

Scheduled outage times for each stage for each block vary by day of the month, and rotate equally across each block. While some blocks may idiosyncratically experience more load shedding than other blocks at times, because higher stages happen to be in effect when their block is scheduled to be affected by those stages, these inequalities will very likely balance out in the long run, especially in recent years when load shedding has become a nearly constant phenomenon. In general, therefore, load shedding is fairly uniformly experienced across the country.

Based on these national standards, I create a 'severity-weighted' measure of load shedding implementation h^w that accounts for the differing severity of different stages. The severity-weighted measure reflects the hours that a typical household would experience actual electricity cut-off, given the number of hours of load shedding h_s being implemented at a given stage s: $h^w = h_s \cdot \frac{1.5 \cdot s}{24}$. Per this measure, 48 hours of load shedding at stage 1 is regarded as equivalent to 24 hours of load shedding at stage 2, for example, because both result in the typical household being without power for 3 hours. This is the preferred measure of load shedding I use throughout this paper.

2.3 Cape Town mitigation policy

The one major exception to this generally uniform experience is the City of Cape Town, which since 2015 has used the Steenbras pumped-storage hydroelectric plant (and some small open cycle gas turbines) to mitigate load shedding in City-supplied areas. While the City's load shedding schedules for a given stage are exactly analogous to the national standard discussed above, they often announce and implement a lower stage of load shedding in the City than is in effect nationally. The general pattern is that load shedding in the City is at the same stage as national load shedding at night, during which time the pumped-storage upper reservoir is replenished, and then lower load shedding levels are implemented during the day if sufficient pumped-storage reserves have been built up.

The Steenbras pumped-storage hydroelectric plant was originally constructed by the City in 1979, and has been in continuous operation since then, though it has only been used to mitigate load shedding since 2015 (Matthews 2015). Prior to that it was used to reduce demand from Eskom during peak periods to save on electricity purchasing costs, via a process called 'peak lopping'.

[13] Because load shedding is implemented in 2-hour slots at a minimum, stage 1 (1.5 hours per 24 hours), for example, would mean three 2-hour slots over four days. The above calculation also does not take into account reconnection time—Eskom and municipal providers typically schedule themselves an additional 30 minutes to reconnect areas after the end of each outage. The actual reconnection time usually varies unpredictably within this 30-minute window.

[14] Specifically, the metropolitan municipality of eThekwini (which includes Durban) received an exemption after devastating floods in 2022, which reduced their load shedding schedule substantially; Nelson Mandela Bay, Buffalo City, and eThekwini (prior to 2022) seem to have implemented slightly less intense schedules, such that rather than 3 hours off per day in stage 2 they respectively implemented roughly 2.5, 2.6, and 2.8 hours per day on average for the typical block; and Polokwane institutes *much* less severe load shedding than the national standard, without any public explanation that I am aware of.

[15] I do in any case exclude the two cases where there are substantial deviations: Polokwane and post-2021 eThekwini.

As is evident from Figure 2, the difference between the Cape Town stage and national stage is not uniform over time. The pumped-storage plant (or particular generating units) sometimes needs to be taken offline for maintenance, and in any case a pumped-storage hydroelectric plant essentially operates as a giant battery and can provide more power when there is off-peak capacity for upper-reservoir replenishing.[16]

Figure 2: Severity-weighted post-2008 load shedding hours, national versus Cape Town

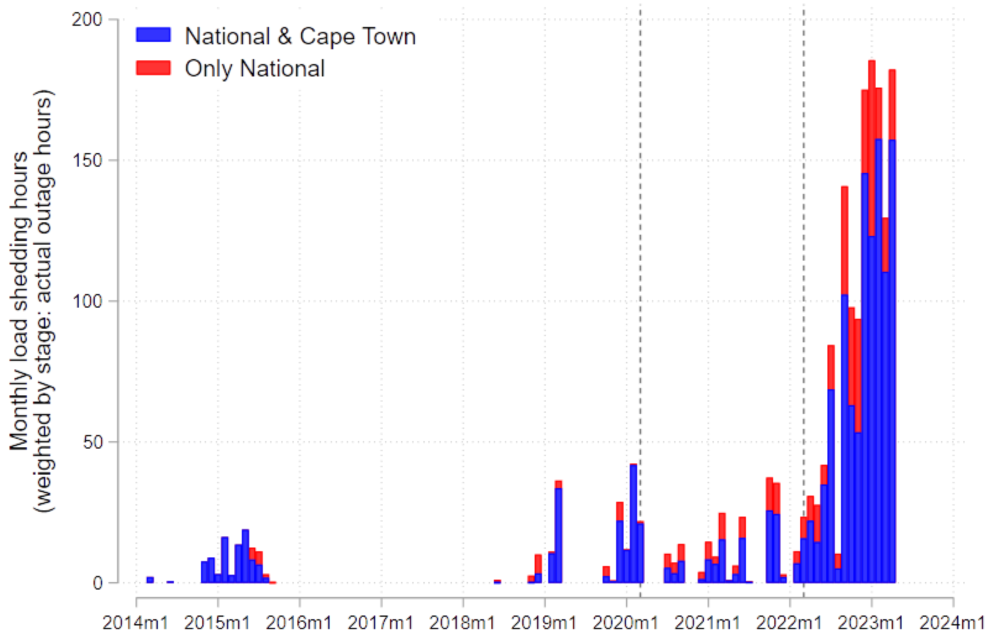

Note: the figure shows severity-weighted load shedding outage time per month for a typical household in the City-supplied region of Cape Town vs the rest of South Africa, from January 2014 to April 2023. Actual load shedding hours are created by weighting each hour of instituted load shedding by the severity of its stage, to match the stage-specific severity outlined in Section 2. Vertical dashed lines show approximate beginning and end of the COVID-19 period in South Africa.
Source: author's compilation based on EskomSePush historic load shedding schedule, the City of Cape Town historic load shedding schedule, and the author's collection of Eskom and City of Cape Town load shedding alert tweets.

The City of Cape Town load shedding mitigation also does not apply to the whole of the metropolitan municipality, but only to City-supplied (rather than Eskom-supplied) areas (Figure A2). Though there are exceptions, in general the Eskom-supplied areas are poorer than City-supplied areas and contain many of the densely populated informal settlements on the city periphery; racial disparities between the areas are particularly stark (see Table 5 in Section 6). The implications of this for the treatment effects estimated in this paper and overall evaluation of the policy are discussed in detail in Sections 6 and 7. I otherwise tend to conflate Cape Town and the City-supplied regions in Cape Town for expositional simplicity.

3 Data

3.1 Load shedding implementation

My main data for national load shedding implementation comes from the developers of a mobile app called 'EskomSePush' (ESP), which collates load shedding schedules for different areas across the country, tracks the current load shedding stage, and provides a personalized schedule for outages in the user's

[16] Or, more accurately, for lower-reservoir emptying; the bottleneck is the capacity of the lower reservoir (Matthews 2015).

chosen area. In response to my queries, they have made publicly available a time series of national load shedding stages (dates and timestamps) going to back to February 2015, when the app was first developed. The series is updated almost daily to include the latest data, and is the main data source researchers and organizations use when tracking or otherwise studying load shedding. This series has no geographical disaggregation, and does not cover the 2007/08, 2014, or January 2015 outages. I hand-built a time series of the 2014 and January 2015 outages based on the Eskom Twitter feed, which ESP has now added to their publicly available data.

I cross-reference the ESP time series with hourly electricity generation and distribution data I received from Eskom, per a PAIA request (South African Freedom of Information Act (FOIA) equivalent), which records generation and 'manual load reduction' (combined load shedding and load curtailment) in megawatt-hours (MWh), going back to January 2006. Because this data is not directly comparable to the Cape Town data below, it is not used in this analysis except to check the ESP data (they match well).

For data on the City of Cape Town's load shedding implementation history, I made a request to their Open Data Portal team, who subsequently publicly uploaded the City's load shedding implementation history back to 2018. This data is at the suburb level and includes geographical disaggregation for the areas they supply. Constructing an aggregate hourly time series comparable to the ESP national data required some minor assumptions and simplifications, because the Cape Town data only shows the applicable load shedding stage at the time of load shedding commencement for each block (and so misses some load shedding step-downs during 2-hour load shedding periods), and also because it seems to potentially include reconnection periods.[17] I therefore assume that the City does not ever implement load shedding at a more severe stage than national, though this affects relatively few records (which are disproportionately at the common step-down time of 5 a.m.), and also round load shedding commencement and termination periods to the nearest hour for both datasets. This data covers load shedding from its recommencement in 2018 up to October 2023. I supplement it with 2015 data I scraped from the City of Cape Town Twitter account, and I know from news sources that there was no load shedding mitigating in Cape Town before this period.

3.2 Health outcomes

I have weekly data on all recorded deaths in South Africa between January 2014 and April 2023, which I requested from the South African Medical Research Council (SAMRC). This data is imputed from the National Population Register (NPR), with adjustments for under-reporting and unregistered individuals. The data is aggregated to preserve individual anonymity. Natural and unnatural deaths are recorded separately, and weekly deaths then provided for each cell, defined by age–sex–geography interactions. Weeks are defined according to US CDC and WHO 'epidemiological weeks'. There are 19 age categories, 2 sex categories, and 52 geography categories. The geography variable is constituted by district councils and metropolitan municipalities, one of which is Cape Town. Collectively these constitute exhaustive and mutually exclusive geography categories for South Africa. For ease of reference I refer to this geographic disaggregation level as simply 'the DC level' henceforth.

In order to construct mortality rates, I merge in Statistics South Africa's (Stats SA) mid-year population estimates (MYPE), which are disaggregated along similar age, sex, and DC categories.

[17] Unlike step-downs, load shedding step-ups are all observed, because they always entail additional areas starting load shedding.

3.3 Predictor variables

My analysis requires a number of DC-level predictor variables for the synthetic control analysis. I merge in data from three distinct sources.

Weather is an important predictor of mortality, so I create weekly population-weighted DC-level weather variables. I download highly granular temperature and precipitation data for South Africa from the Copernicus Climate Data (CDS) Store, which I then aggregate to the weekly DC level. In order to create weather variables that are representative of where people actually live in each DC, I merge in gridded population data from the Gridded Population of the World (GPW) dataset (version 4) provided by NASA's Socioeconomic Data and Applications Center (SEDAC) as well as DC coordinates from Stats SA. I then create population-weighted weather variables for each DC. I create weekly minimum, mean, and maximum temperatures and precipitation variables for each DC.

South African DC-level health data is available from the District Health Barometer (DHB). I extract this data from the 2022/23 release by the Health Systems Trust. Because this data is typically annual, there is no over-time variation in these variables in the pre-period of each (sub-annual) event I study in Section 4. I therefore only extract indicators for which there is data prior to 2014 (when load shedding recommenced) and I extract only the latest available year of data prior to 2014.

For the same reason I only extract 2013 data from Stats SA's annual General Household Survey (2013), except where I use other years' data for the purposes of descriptive statistics. The GHS is a preeminent South African household survey that surveys approximately 26,000 households (94,000 individuals) about various socioeconomic factors. I collapse the microdata to create summary statistics at the DC level. DC codes are not released in the data, but can be constructed from parsing the household identifier variable.[18]

4 Method

4.1 Stacked synthetic control

The synthetic control method (SCM) for estimating treatment effects is widely used in contexts with reasonably long time series but too few cross-sectional units for standard regression analysis. A leading case is analysis of policies or events applicable to aggregated entities, such as countries or their sub-regions. Advantages of the methodology include its transparency and interpretability, excellent theoretical performance when the untreated outcomes are generated per a linear factor model, and inference methods suited for cases with one or a few treated units (Abadie 2021).

Indeed, SCM was originally developed to estimate treatment effects for a single aggregated treated unit (Abadie 2021; Abadie and Gardeazabal 2003; Abadie et al. 2010). The method has, however, subsequently been extended to contexts with multiple treated units (Abadie and L'Hour 2021; Ben-Michael et al. 2022; Cavallo et al. 2013; Dube and Zipperer 2015; Wiltshire 2023). In this paper I implement a version of the 'stacked' method of Wiltshire (2022, 2023), which builds on the similar approach of Dube and Zipperer (2015). Due to widespread use of SCM, in my exposition below I focus on the methodological implications of stacking for multiple treatment events rather than recapitulating the de-

[18] An additional wrinkle is that the 2013 GHS enumeration used geographies defined by the 2001 Census, whereas the rest of my analysis uses 2011 Census geographies. I therefore use a 2001 to 2011 local municipality cross-walk provided by Adrian Frith and aggregate to the DC level.

tails of standard SCM methodology; the reader is referred to Abadie (2021) for a more comprehensive introduction and review of the method.[19]

Stacked episodes for a single treated unit

I observe data for $J+1$ units, $j = 1,2,\ldots,J+1$, and without loss of generality assume that the first unit ($j = 1$) is the treated unit. The 'donor pool' is the set of J untreated units, $j = 2,\ldots,J+1$, which are not only untreated by the policy of interest but also are not affected by spillover effects of the policy.

Data is observed for these units over T periods, $t = 1,2,\ldots,T$. The distinguishing feature of the single-treated-unit stacked design utilized here is that the treated unit switches in and out of treatment multiple times between $t = 1$ and $t = T$. I define an event v as a time-span $T_v < T$ such that $t = T_{0v}+1-\underline{E},\ldots,T_{0v},T_{0v}+1,\ldots,T_{0v}+\overline{E}$, where $t = T_{0v}+1$ is a period in which treatment switches on for unit $j = 1$, which is preceded by at least $\underline{E}$ consecutive 'pre-periods' without treatment (including $t = T_{0v}$), and which is followed by $\overline{E}-1$ post-periods during which time treatment may switch on and off but for simplicity we assume stays on. The common parameters $\underline{E}$ and $\overline{E}$ determine the length of the pre-period and post-period 'windows', respectively.

This structure of distinct 'episodes' in which treatment is 'on' preceded by a sufficiently long pre-period where treatment is 'off' allows me to use (suitably adjusted) techniques from the multiple-treated-unit SCM literature to 'pool' or 'stack' the case studies as distinct events, in order to incorporate as many policy changes as possible and increase statistical power.[20] Essentially I estimate standard single-unit synthetic control estimates for each event, and then aggregate treatment effects across these events.

Defining multiple distinct treatment episodes for a particular treated unit as separate events has been implemented before in the synthetic control literature (e.g. Dube and Zipperer 2015), and is fairly common in the 'stacked' difference-in-differences literature (e.g. Cengiz et al. 2019). The chief drawback of this approach is that it requires an additional assumption about the permanence of dynamic effects of the policy. If effects are sufficiently long-lasting after treatment episodes end, they may contaminate the pre-period of subsequent events. I therefore assume that effects will dissipate after at most T_δ periods, and only use events which are preceded by at least T_δ consecutive periods without treatment since the last time $j = 1$ experienced treatment.

Aggregated treatment effects in event-time

For each unit j in time period t, I observe the outcome Y_{jt}. Restricting attention to observations that fall within identified events, it is useful to recast the calendar time variable t into event-time e_v, where for each event v, $e_v = t - T_{0v}$, so that for all v the immediate period prior to treatment is indexed $e_v = 0$ and the first treatment period is indexed $e_v = 1$. Each outcome Y_{jt} can then be expressed as Y_{je_v}.

My object of interest is the causal effect of the treatment. For each unit j and event-time e_v, define $Y^N_{je_v}$ as the potential outcome in the absence of treatment (**N**on-intervention), and $Y^I_{je_v}$ as the potential outcome under treatment (**I**ntervention). The effect of the policy on unit j in event-time e_v is then

$$\tau_{je_v} = Y^I_{je_v} - Y^N_{je_v} \tag{1}$$

For the treated unit $j = 1$ and the period of interest $e_v > 0$ I necessarily have $Y_{1e_v} = Y^I_{1e_v}$, and so the task of the causal inference exercise is to estimate the counterfactual $Y^N_{1e_v}$ for $e_v > 0$.

[19] My technical exposition below draws substantially from Abadie (2021) and Wiltshire (2022).

[20] Indeed, I easily implement my analysis using the Wiltshire (2022) `allsynth` package.

For each unit j I also observe k predictors of the outcome, which may include linear combinations (in event-time) of pre-intervention values of the outcome Y_{je_v} and other linear combinations of variables that are unaffected by the treatment. While the set of k predictors is common across units and events, the values of the predictors are both unit- and event-specific, so that the $k \times 1$ vector $\mathbf{X}_{jv} = (X_{1jv}, \ldots, X_{kjv})$ contains j's values of these predictors in event v, and it may be that $\mathbf{X}_{jv} \neq \mathbf{X}_{jv'}$. The event-specific $k \times J$ matrix $\mathbf{X}_{0v} = [\mathbf{X}_{2,v}, \ldots, \mathbf{X}_{J+1,v}]$ contains these vectors for the donor pool units for event v.

The synthetic control estimator for $Y^N_{1e_v}$ is a weighted average of the observed outcomes of the J untreated donor pool units Y_{je_v}:

$$\hat{Y}^N_{1e_v} = \sum_{j=2}^{J+1} w^v_j Y_{je_v} \tag{2}$$

The (typically sparse) event-specific weights $\mathbf{W}^v = (w^v_2, \ldots, w^v_{J+1})$ are those that minimize the distance between the values of the k predictors for the treated unit $\mathbf{X}_{1v}$ and the donor pool $\mathbf{X}_{0v}$, for some norm (typically Euclidean), given a set of weights on the k predictors that determine their relative importance.[21] Adding-up and non-negativity constraints are also imposed on the weights $\mathbf{W}^v$, such that $\sum_{j=2}^{J+1} w^v_j = 1$ and $w^v_j \geq 0 \; \forall \; j > 1$, to prevent extrapolation.

The idea is that the estimated weights construct a 'synthetic' control unit as a weighted average of the donor units, by training the weights in the pre-period such that the weighted average of the synthetic control predictors match the treated unit predictors. If the synthetic control is sufficiently similar to the treated unit, divergences of the treated unit from its synthetic counterpart in the post-period reflect the effects of treatment.

Combining Equations 1 and 2, the event-specific estimated treatment effect for the treated unit for event-time $e_v > 0$ is

$$\hat{\tau}_{1e_v} = Y_{1e_v} - \sum_{j=2}^{J+1} w^v_j Y_{je_v} \tag{3}$$

While these event-specific effects may be of interest, I am particularly interested in treatment effects aggregated across events for an average treatment effect on the treated (ATT) for each event-time e.[22] This is naturally achieved by a (possibly weighted) average of the event-specific effects,

$$\hat{\tau}_{1e} = \sum_{v=i}^{V} \gamma_v \hat{\tau}_{1e_v} \tag{4}$$

with weights $\sum_{v=1}^{V} w_v = 1$ and $\gamma_v \geq 0 \; \forall \; v$.

Bias correction

In the main specification I also implement the bias correction procedure of Abadie (2021) as implemented in Wiltshire (2022). This approach, which is similar to the bias correction of Ben-Michael et al. (2021), adjusts the treatment effects for potential bias caused by imperfect pre-treatment fit in the various predictors. Let $\hat{\mu}_{0e_v}(x)$ be a predictor for outcome Y_{je_v} for $X_{jv} = x$, estimated by using the donor pool to regress Y_{e_v} on the set of predictor variables separately for each event. Analogously to Equation 1, the bias-corrected event-specific treatment effects are then calculated as

$$\hat{\tau}^{BC}_{1e_v} = [Y_{1e_v} - \hat{\mu}_{0e_v}(X_{1v})] - \sum_{j=2}^{J+1} w^v_j [Y_{je_v} - \hat{\mu}_{0e_v}(X_{jv})] \tag{5}$$

[21] The reader is referred to Abadie (2021) for discussion of these weights on the k predictors, $\nu^v_1, \ldots, \nu^v_k$.

[22] Note that the $\hat{\tau}_{1e_v}$ are already event-specific ATTs.

and are aggregated across events like in Equation 4. The idea is that the bias correction adjusts for discrepancies between the predictor values of the treated unit and the positively weighted donor units.

Inference

Due to the typically small number of units in the donor pool, and typically even smaller number of treated units, large-sample inferential approaches are generally not appropriate for synthetic control methods. Instead, p-values are usually calculated by permutation inference using 'in-space' placebo tests to create the permutation distribution (Abadie 2021; Wiltshire 2023). While a variety of approaches to permutation inference are possible, here I describe the RMSPE (ratio of the mean squared prediction errors) p-value procedure described in Abadie and L'Hour (2021) as implemented in Wiltshire (2022), which is essentially the inferential procedure of Abadie et al. (2015, 2010) extended to the case of multiple treated units.[23]

Separately for each event v, treatment is permuted across all untreated units $j > 1$, with the treated unit $j = 1$ and the remaining untreated units constituting the donor pool for each j. Event-specific placebo treatment effects ('placebo gaps') $\hat{\tau}_{je_v}$ are then computed for each $j > 1$ and e_v, per Equation 1 (or its bias-corrected equivalent).

The estimated average placebo gaps across events must then be calculated, with each constructed by taking the (possibly weighted) mean of the $\hat{\tau}_{je_v}$ placebo gaps across events, such that exactly one j is chosen from each event, and weighted by the same γ_v as in Equation 4. While the averaging is across v, and the number of events V may be small, given even middling sizes of J the number of placebo gaps becomes very large, $N_G = \prod_{v=1}^{V} J$, and it may be computationally infeasible to estimate all of these gaps. As suggested by Abadie (2021), Wiltshire (2022) randomly samples some subset of these placebo averages, $S < N_G$, to form the permutation distribution.

The actual aggregated treatment effects $\hat{\tau}_{je}$ can then be plotted against the placebo gaps $\hat{\tau}_{se}$ with $s = 1, 2, \ldots, S$, to observe whether the actual treatment effect is in the tails of the (sample) permutation distribution. The RMSPE p-values are then constructed in the standard way, by first calculating for each s the ratio between the mean-squared prediction error (MSPE) up to each post-period $e > 0$ and the average pre-period MSPE, and then using the empirical distribution of these ratios to see whether the RMSPE for the treated unit $j = 1$ is in its tails.[24]

A last step is to construct confidence intervals and standard errors around the treatment effects, to generate estimates of statistical precision, and to allow inference on transformations of the treatment effects. There are a variety of ways to do this, such as the methods used in Dube and Zipperer (2015), Doudchenko and Imbens (2016), Chernozhukov et al. (2018), and Arkhangelsky et al. (2021). In my case I create confidence intervals by inverting the RMSPE test statistics calculated above, in the spirit of Firpo and Possebom (2018). This approach has the advantages of being straightforwardly compatible with the Abadie (2021) and Wiltshire (2022) approach to multiple-event synthetic control, uses a test statistics which adjusts for pre-treatment fit, and by definition will have exactly the same coverage as the p-values calculated as above. The use of an empirical permutation distribution of 150 placebo units and the numerical grid search underlying the test-statistic inversion means the confidence intervals are slightly approximate (and sometimes slightly asymmetric). In particular, one must choose an RMSPE rank to associate with a 5% significance level: I use rank 8 as 8/151 is closest to 0.05.[25]

---

[23] There is now a large literature investigating and extending inference in SCM; the reader is referred to Abadie and L'Hour (2021) and Wiltshire (2023) for discussion and references.

[24] See Abadie (2021) and Wiltshire (2022) for a more detailed exposition.

[25] Additionally, I specify that the grid search steps in 0.01 increments, except for the age ≥ 65 and log of deaths specifications, where reduced-form treatment effects are an order of magnitude smaller and so I step in 0.001 increments.

4.2 Events

The first step in implementing my stacked synthetic control analysis is to define the events in the data. Recall from Section 4.1 that I need to choose pre- and post-event windows $\underline{E}$ and $\overline{E}$, and also specify a time-span T_δ which is the maximum time it may take for dynamic effects of the policy to dissipate after treatment ends.

Cape Town's load shedding mitigation policy started in 2015, and continues today, as is evident from Figure 2. However, a substantial part of this period is not amenable to synthetic control analysis. First, there was no load shedding in 2016 and 2017, trivially precluding analysis of load shedding mitigation effects in this period. More significantly, from March 2020 the COVID-19 pandemic was extremely disruptive in South Africa, causing substantial excess mortality between 2020 and 2022 (Bradshaw et al. 2022). This COVID-19-related mortality was regionally heterogeneous in its timing and severity, and would likely confound attempts to detect the relatively small load shedding mortality effects examined in this paper. In addition, accompanying the pandemic were particularly stringent government lockdowns, which severely restricted economic activity, and therefore electricity demand and load shedding. I therefore judge it inadvisable to attempt to define events which fall between March 2020 and March 2022. Events unfortunately cannot be defined after March 2022 either, however, because by this time load shedding was so severe and constant that one cannot find any pre-period longer than two weeks in which Cape Town did not mitigate relative to national.[26]

Figure 3 shows the weekly incidence of national and Cape Town load shedding for the 2014–15 and 2018–19 epidemiological years, with the time variable defined as epidemiological weeks since 1 January 2014.[27] Up to five events could possibly be defined, assuming at least four pre-periods are required:

a. Nine consecutive weeks of salient differential load shedding starting in epiweek 76 (week starting 7 June 2015). There is no differential load shedding before this date, because the mitigation policy only starts in epiweek 76, but unlike the other events considered here the period preceding did have load shedding.

b. One week of very marginal differential load shedding in epiweek 233 (week starting 10 June 2018), preceded by more than two years of no load shedding.

c. Three consecutive weeks of salient differential load shedding starting in epiweek 256 (week starting 18 November 2018), preceded by 22 consecutive weeks without load shedding.

d. Three non-consecutive weeks with very marginal load shedding mitigation between epiweeks 268 and 273 (six weeks starting 10 February 2019 and ending 23 March 2019), preceded by nine weeks without load shedding.

e. Four non-consecutive weeks with collectively salient differential load shedding mitigation between epiweeks 303 and 311 (nine weeks starting 13 October 2019 and ending 14 December 2019), preceded by 29 weeks without load shedding.

[26] The exclusion of this latest period where load shedding was most intense has implications for the generalizability of this paper's results; this is discussed in Section 6.

[27] The first few months of 2020 before the pandemic struck had virtually no load shedding mitigation in Cape Town. Note that epiweeks are defined on a yearly basis, so that the first and last epiweeks of a given year may contain more or less than seven days, but this will of course be common across all regions in my data.

Figure 3: Weekly severity-weighted load shedding hours, national vs Cape Town, selected years, with events

(a) 2014 and 2015

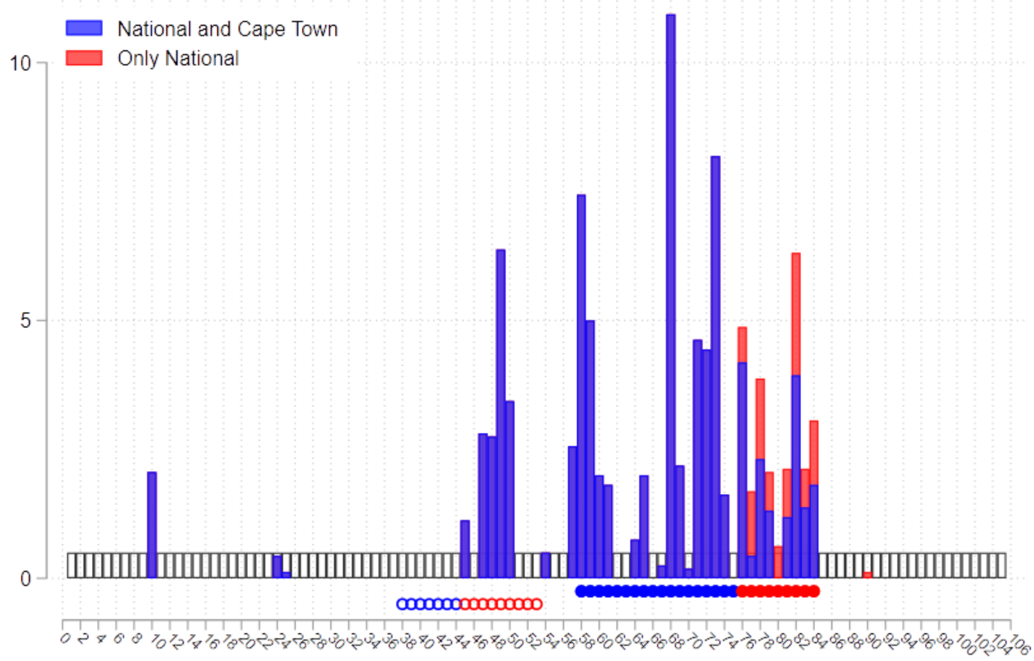

(b) 2018 and 2019

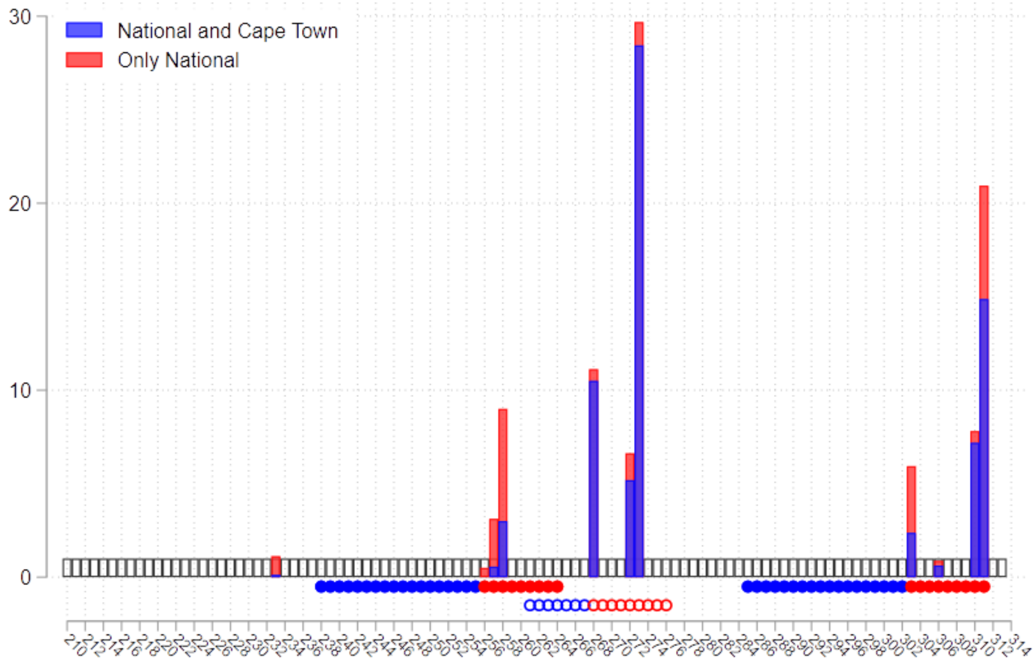

Note: the figures are analogous to Figure 2, but aggregated at the weekly (epiweek) level. The *x*-axis counts epiweeks since 1 January 2014. Filled circles below the figure show event pre-periods (blue) and post-periods (red). Hollow circles do the same for placebo events. Short vertical black bars are shown for each week.

Source: author's compilation.

I exclude the potential events (b) and (d) from my analysis, because the extent of mitigated load shedding is extremely minimal, and event (d) in any case has a much shorter possible pre-period window than the other potential events. Events (a), (c), and (e), in contrast, offer relatively similar amounts of load shedding mitigation (see Table 1) and have long pre-periods. These events are indicated by filled dots in Figure 3. Given that differential load shedding in events (a) and (e) occurs over nine weeks, I use nine weeks for the post-period window, $\overline{E} = 9$. I then use twice that number for the pre-period length, $\underline{E} = 18$. These periods are shown with red and blue dots respectively in Figure 3. This then implies that $T_{\delta} = 4$; that is, I assume that any dynamic effects of mitigated load shedding on mortality do not last more than four weeks after the last period of differential load shedding.[28] I denote potential event (a) as event 1, potential event (c) as event 2, and potential event (e) as event 3. Table 1 shows the extent of load shedding mitigation for each of these three events and the average load shedding per post-period across the events.

Table 1: Severity-weighted load shedding mitigation hours by event and event-time

Event-time	Event 1	Event 2	Event 3	Average
1	0.69	0.5	3.56	1.58
2	1.25	2.56	0	1.27
3	1.56	6	0	2.52
4	0.75	0	0.25	0.33
5	0.63	0	0	0.21
6	0.94	0	0	0.31
7	2.38	0	0	0.79
8	0.75	0	0.63	0.46
9	1.25	0	6.06	2.44
Total	10.19	9.06	10.5	9.92

Note: the table shows the difference in severity-weighted load shedding hours between City-supplied regions of Cape Town and the rest of the country for each post-period of each event. The last column takes a simple average across the events.

Source: author's compilation.

4.3 Main specification

For the main specification I use all-cause deaths per capita as the main outcome variable, though I rescale it to deaths per 10,000 population for ease of interpretation. As I show in Sections 5 and 6, the results are robust to other outcome transformations, such as the log of all-cause deaths. In the synthetic control estimation I also de-mean the outcome variable prior to estimation using the pre-treatment average of Y_{je_v} over the pre-treatment period, which can attenuate bias from time-varying unit-level confounders (Ferman and Pinto 2021). In the baseline specification I estimate bias-corrected treatment effects, but show that the results are very similar when the bias correction is not applied.

For the predictors $\mathbf{X}_{jv}$ I use the following ($k = 12$):

- The outcome variable Y_{je_v} averaged over the pre-period, where I exclude the immediate pre-treatment period $e_v = 0$ from the pre-treatment average to help guard against over-fitting or mean reversion effects in $e_v = 1$.[29]

- The average population-weighted 2-metre temperature and precipitation for each DC over the entire event (from the CDS). While one typically only uses pre-treatment values when constructing synthetic control predictors because of concern that the policy may affect the predictors (Abadie 2021), this issue will clearly not apply to these weather variables.

- The 2013 age structure characteristics from the GHS: separately the proportions of the population which are aged 24 or younger, 25–54, and 54–65.

- The 2013 living condition characteristics from the GHS, to reflect reliance on electricity, municipal amenities, socioeconomic conditions, and house structure type in each DC: the proportion

[28] This seems eminently plausible to me, keeping in mind that the extent of severity-weighted load shedding mitigation in potential event (b), which is five weeks before the start of potential event (c)'s pre-period, is 1 hour over seven days.

[29] Though this issue is more comprehensively dealt with by time placebo robustness tests in Section 5.

of the population which is urban, the proportion using mains electricity for heating, and the proportion of the population living in a government-provided formal house ('RDP' or other state-subsidised dwelling).

- The 2013 and 2012 variables reflecting mortality vulnerability and health services in each DC, from the DHB: 2013 public sector health expenditure per capita, the incidence of adult obesity, and age-adjusted incidence of non-raised blood pressure.[30]

An important decision regarding the main specification concerns the population of interest. Guided by the previous literature, I focus on those aged 65 and older, but show that the results do extend to the general population when very young deaths (age < 5) are excluded. The results for those aged 65 and older are more robust, however, likely due to a greater signal-to-noise ratio (they account for a disproportionate share of the mortality effects), which is another reason to treat them as my main specification, with appropriate limitations regarding the conclusions. Neonatal deaths comprise a substantial share of overall deaths, and based on the poor pre-treatment fit when these deaths are included, it seems that the factor model or other data-generating process underlying my main results does not extend to this quite different type of death, which is not surprising.

5 Results

5.1 Synthetic control weights

How is synthetic Cape Town constructed? Table A1 in the Appendix provides the synthetic control weights for each DC for the main specification, separately by each event and averaged across DCs. The weights seem to make some intuitive sense.

Just over half of synthetic Cape Town is constructed from DCs in the Western Cape, like Cape Town (West Coast, Cape Winelands, Overberg, Eden).[31] Apart from various socioeconomic similarities, these are regions that predominantly experience winter rainfall and dry summers, unlike the rest of the country. The most highly weighted of these Western Cape DCs, the Cape Winelands and Eden, also contain the largest regional urban population centres outside of Cape Town.[32]

Metropolitan municipalities make up nearly 40% of synthetic Cape Town (Buffalo City, Mangaung, eThekwini, Ekurhuleni, and City of Tshwane).[33] The only metropolitan municipalities not represented (other than Cape Town) are Nelson Mandela Bay (mostly constituted by Gqeberha, previously Port Elizabeth) and the City of Johannesburg, but both of these regions are sometimes positively weighted in the alternative specifications. The largest metropolitan municipality weight is associated with Ekuhuleni. Ekurhuleni is South Africa's third largest metropolitan municipality by population (after Johannesburg

---

[30] For the latter two variables, 2012 is the latest pre-period year for which data is available in the Health Systems Trust release.

[31] Common concerns about cross-border spillover effects in regional analysis seem unlikely to apply here. First, the treatment is very strictly restricted to households in the City of Cape Town; there is no mechanism for houses outside the City to have load shedding reduced by the City's policy. Second, as per Figure A2, the regions supplied by the City of Cape Town tend to be more centrally located within the metropolitan municipality anyway; there is close to a border buffer zone already.

[32] For example, Stellenbosch and Paarl, and George, Knysna, Oudtshoorn, and Mossel Bay respectively.

[33] International readers may know these metropolitan municipalities better by their main cities or previous names of, respectively, East London, Bloemfontein, Durban, the East Rand, and Pretoria.

and Cape Town) and is also a major economic centre. It neighbours and is highly integrated with the City of Johannesburg; indeed, it is considered part of 'Greater Johannesburg'.[34]

The remaining 10% of synthetic Cape Town is made up of relatively small weights attached to a few scattered DCs. The DC of Sedibeng (3% average weight) neighbours and is highly integrated with Ekurhuleni and the City of Johannesburg, and its inclusion would seem to follow the same logic as Ekurhuleni's.[35] It is harder to discern the comparative value of the remaining 7%, but these DCs form a very small part of the total synthetic Cape Town in any case.[36]

Figure A3 shows the results if, instead of using the synthetic control methodology with these (sparse) weights, one simply implements a stacked event-study (Cengiz et al. 2019) using the events as defined in the previous section (and otherwise replicates the main specification below). This implicitly weights all DCs equally, and does not seek to construct a specific counterfactual Cape Town, but assumes that the country as a whole provides a good counterfactual. This is a priori unlikely, which is borne out in Figure A3. There are significant pre-period dynamics when comparing Cape Town to the rest of the country—both periods where Cape Town mortality is rising and periods where it falls. This prevents causal interpretation of coefficients in the post-period, and indeed may obscure such effects, and motivates for the synthetic control approach.

5.2 Treatment effects: main specification

The main analysis focuses on the effects of load shedding reduction on mortality (deaths per 10,000 population) among those aged 65 and older. Figure 4(a) shows point estimates for this main specification, where I implement the bias correction procedure of Abadie and L'Hour (2021) as implemented in Wiltshire (2022). While the pre-treatment fit is a bit noisy at the beginning of the pre-period, it soon becomes very stable and close to zero, until there is a substantial negative deviation of Cape Town from synthetic Cape Town (i.e. lower mortality in Cape Town) in the first three post-treatment periods, and a return to baseline from period 5 onwards. Referring to Table 1, this pattern seems to make sense: of the total post-treatment load shedding mitigation hours over the nine post-periods, 54% occur in the first three periods, and the third period has the most mitigation hours. One potential puzzle is the lack of observable effect in period 9, which averaged across the events has similar mitigation magnitude to period 3. A possible explanation is that load shedding mitigation in period 9 comes almost entirely from one event (event 3), so event-specific noise may obscure the effect, and there is little mitigation in the periods before it, so lagged effects are unlikely to contribute to mortality reduction in period 9 (unlike period 3). Magnitudes are not straightforward to interpret, as the intensity of the treatment varies across the post-period and only part of Cape Town actually receives the policy; they are discussed in Section 6.

Figure 4(b) additionally shows results when the classical rather than bias-corrected synthetic control estimator is used. It is apparent that the results are similar, though the pre-treatment fit is a bit noisier and the post-treatment effect somewhat attenuated.

---

[34] Ekurhuleni uses the same telephone dialling code as Johannesburg, the same metropolitan route numbering for its major roads, typically falls under the legal jurisdiction of the Johannesburg High Court, and includes in its boundaries OR Tambo International Airport (previously Johannesburg International Airport), the main airport servicing Johannesburg.

[35] It contains important industrial centres such as Vereeniging and Vanderbijlpark and urban population centres such as Sebokeng and Evaton.

[36] Siyanda is in the Northern Cape and includes the regional population centre of Upington. Xhariep is in the Free State and borders Mangaung (Bloemfontein). Amajuba is in KwaZulu-Natal and includes the regional industrial and population centre of Newcastle.

Figure 4: Treatment effects, deaths per 10,000 population, age ≥ 65
(a) Bias-corrected

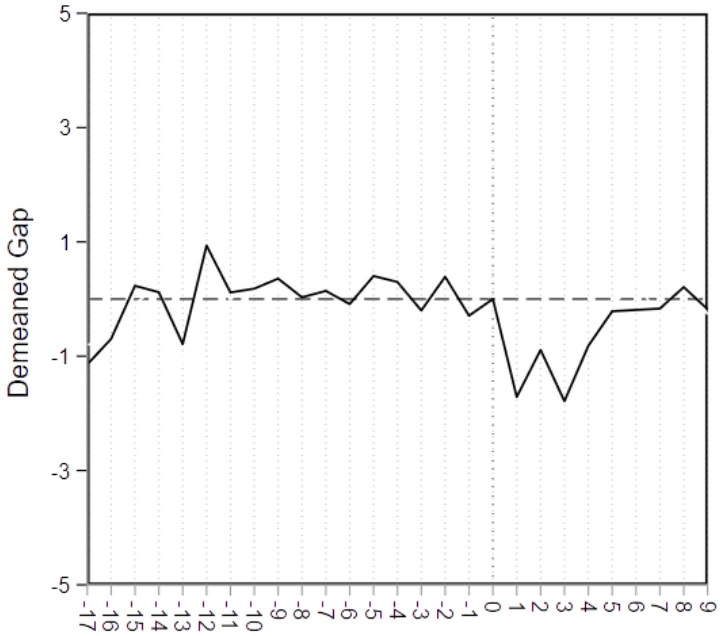

(b) Bias-corrected and classic

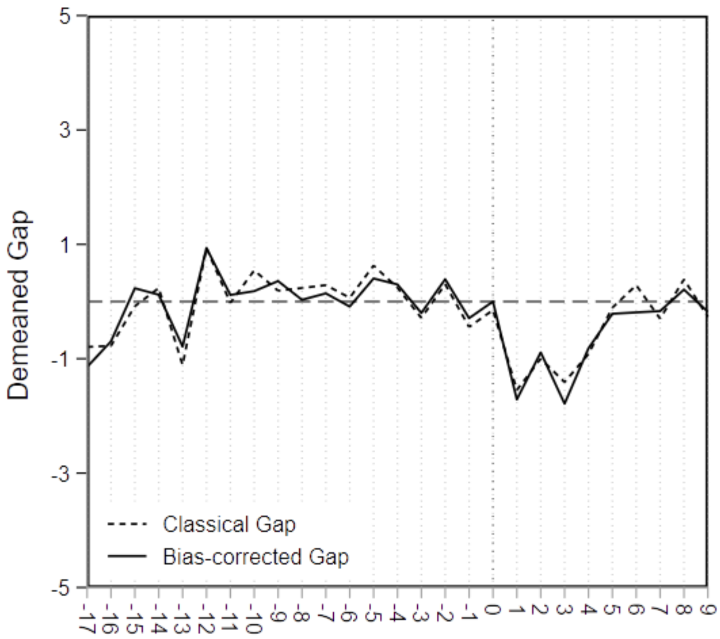

Note: figure shows synthetic control gaps between Cape Town and synthetic Cape Town for the main specification, averaged across the three events. The outcome is deaths per 10,000 population and the sample is restricted to those aged 65 or older. Panel (a) shows bias-corrected gaps while Panel (b) shows bias-corrected (solid line) and classic (dashed-line) gaps. Treatment begins in period 1, the period immediately after the dashed vertical line.

Source: author's compilation.

Table 2 provides point estimates and p-values for the bias-corrected and classical treatment effects. The p-values come from the 'in-space' placebo permutation test described in Section 4.1; these placebo gaps are shown in Figure 5. The treatment effects are all statistically significant at conventional levels, except for the classical specification for post-periods 7, 8, and 9, where the treatment effect magnitude is in any case very small, as discussed above.[37] It is apparent from visual inspection of Figure 5 that the treatment effect in the first three periods is in the tails of the permutation distribution of the in-space placebo effects, at least for the main bias-corrected specification, even before accounting for the superior pre-treatment fit as in the p-value calculation.

Table 2: Treatment effects, deaths per 10,000 population, age ≥ 65

	Bias-corrected		Classical	
Event-time	Gap	*p*-value	Gap	*p*-value
1	–1.712	0.0132	–1.548	0.0265
2	–0.890	0.0066	–1.000	0.0331
3	–1.786	0.0066	–1.410	0.0132
4	–0.819	0.0066	–0.917	0.0132
5	–0.214	0.0066	–0.108	0.0331
6	–0.190	0.0066	0.286	0.0397
7	–0.167	0.0132	–0.299	0.0596
8	0.209	0.0199	0.388	0.0861
9	–0.178	0.0265	–0.270	0.1060

Note: the table shows (reduced-form) synthetic control treatment effects for the main specification, for each post-period averaged across the three events. The *p*-values come from in-place permutation tests. Bias-corrected treatment effects are shown in the left super-column while classic treatment effects are shown in the right super-column.

Source: author's compilation.

[37] It is worth highlighting that the p-values shown in Table 2 for some time period t are calculated using the post-period MSPE from period 1 *up to* period t, not the deviation of period t exclusively. This is the standard approach in the literature (Abadie 2021), and guards against one-period fluctuations in treated or placebo units being over-interpreted (Firpo and Possebom 2018).

Figure 5: Treatment effects with placebos, deaths per 10,000 population, age ≥ 65

(a) Bias-corrected

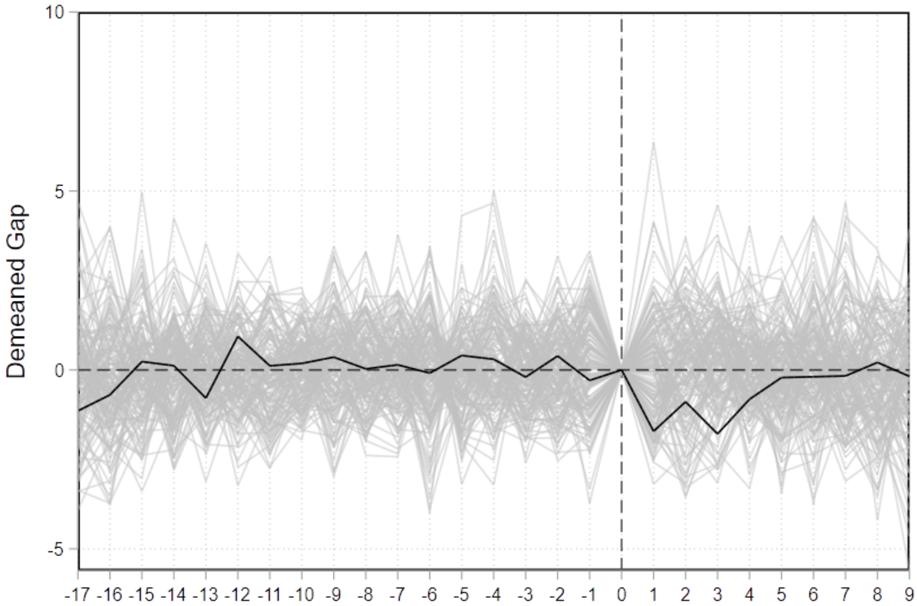

(b) Classic

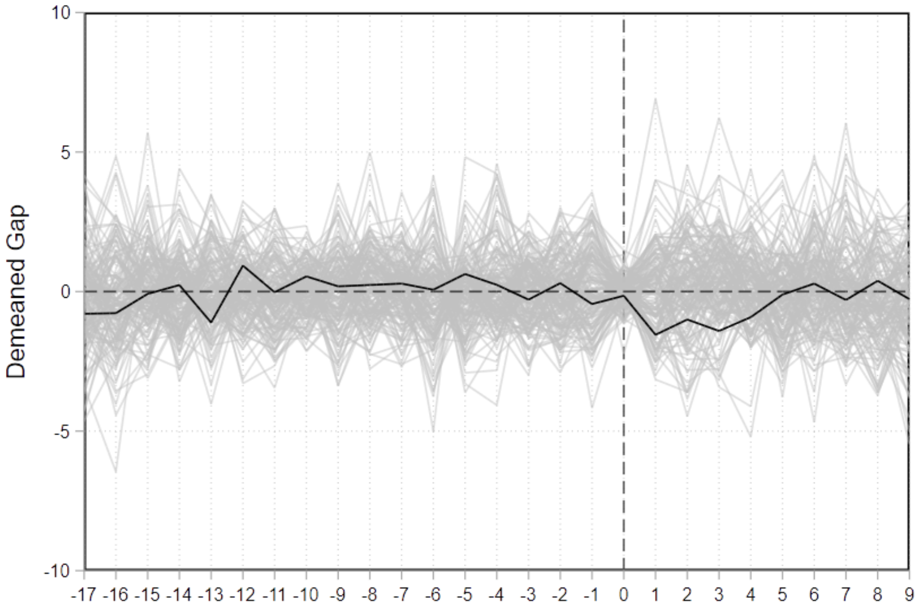

Note: the figure shows synthetic control gaps between Cape Town and synthetic Cape Town for the main specification, averaged across the three events (solid dark line), with placebo gaps from in-place permutation tests (light grey lines). The outcome is deaths per 10,000 population and the sample is restricted to those aged 65 or older. Panel (a) shows bias-corrected gaps while Panel (b) shows classic gaps. Treatment begins in period 1, the period immediately after the dashed vertical line.

Source: author's compilation.

5.3 Treatment effects: other age groups

The results above restrict attention to those 65 or older, in line with much of the literature. What if the full population is considered? Figure 6 shows bias-corrected gaps when no age restriction is imposed (Panel (a)) and when the population is restricted to those aged 5 or older (Panel (b)). The figure for all ages is clearly far less persuasive than when looking only at those aged 65 and older. The post-treatment gaps, while suggestive of similar results, are much more similar in size to the noise in the pre-treatment fit, especially just before treatment commences. Indeed, Table A2 shows that while some treatment effects are statistically significant at the 10% level, none are significant at the 5% level. These full population point estimates are also not robust to transformation of the outcome variable (e.g. log deaths)—and the pre-treatment fit becomes very noisy—unlike the case with the main specification.

Figure 6: Bias-corrected treatment effects, deaths per 10,000 population, various ages

(a) All ages

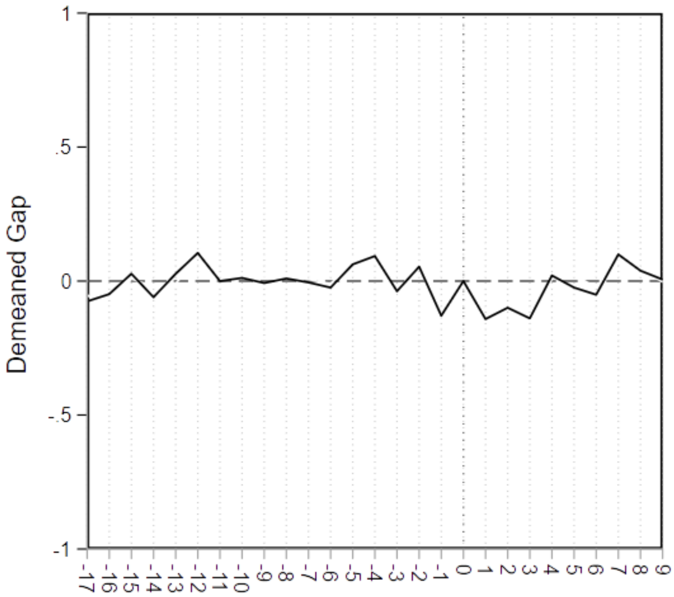

(b) Age ≥ 5

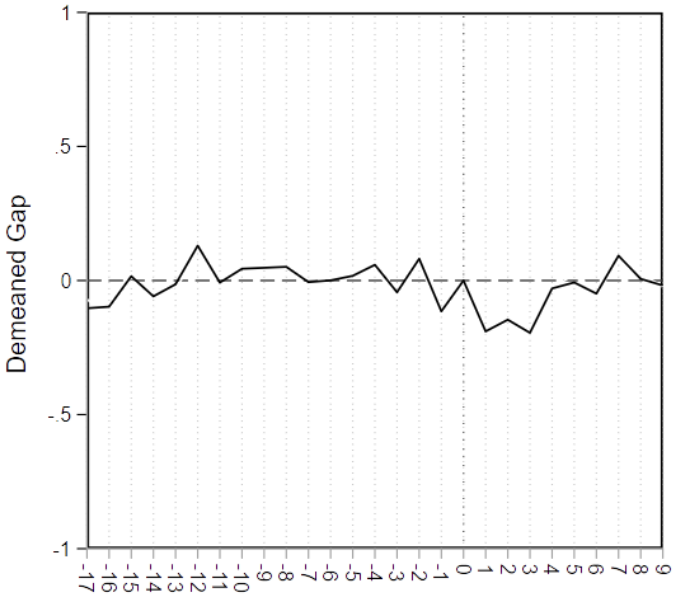

Note: the figure shows synthetic control gaps between Cape Town and synthetic Cape Town for various age groups, averaged across the three events. The outcome is deaths per 10,000 population and bias-corrected treatment effects are shown. Panel (a) shows results when no age restriction is applied, while Panel (b) shows results when the sample is restricted to those aged 5 or older. Treatment begins in period 1, the period immediately after the dashed vertical line.

Source: author's compilation.

Figure 6(b) shows that the results become much more similar to our main specification once one excludes the deaths of those younger than 5. Per Table A2, these age 5 and older results are also statistically significant at the 5% level until period 6. As discussed in Section 4.3, a plausible explanation is that neonatal deaths (which comprise a large share of overall deaths) are driven by a factor model or other data-generating process which is distinct from general-population deaths, and so would require a different synthetic control specification to generate a plausible synthetic Cape Town.[38]

While I do use and discuss the results for the 5 years and older specification in Sections 1 and 6, I retain focus on the 65+ group as my main specification. This focus is not only more typical in the literature, but results are more robust across the various robustness and placebo tests I conduct below. As is outlined in Section 6, ages 65+ in any case make up a preponderance of the treatment effect, disproportionate to their share of total deaths.

5.4 Robustness and falsification tests

Robustness

Section 5.2 shows that the main results are robust to using the bias-corrected or classic synthetic control estimator, and Section 5.3 shows that there is some reasonable robustness to different age groups. The results are also robust to different outcome transformations—Figure A5 shows the main specification when the outcome used is the natural log of deaths, rather than (levels of) deaths per capita, and per Section 6, the implied mortality magnitudes are similar too. As is discussed further in Section 6, the results are also robust (though appropriately slightly quantitatively smaller) to using natural rather than total deaths (Figure A6 and Table A3), which is consistent with causes of death associated with power rationing in the prior literature.

Backdated in-time placebo falsification test

Perhaps more important for a synthetic control analysis is an 'in-time placebo' falsification test, also called 'backdating' (Abadie and L'Hour 2021). It is possible for post-treatment divergences in a synthetic control analysis to be driven by over-fitting in the pre-treatment period. In this case the synthetic control weights reflect 'irreducible error' rather than fundamentals, and the treated unit diverges from the synthetic unit immediately after the training period not because of genuine treatment effects, but simply because the weights provide a poor match outside of the training sample (the pre-period). A common falsification test for this problem is to re-implement the synthetic control analysis, but to artificially backdate the treatment a few periods from when it actually begins, to ensure that (1) the synthetic unit does not suddenly diverge from the treated unit after the artificially backdated treatment but before the true treatment, and (2) the treated unit diverges from the synthetic unit at the true treatment period in a fashion similar to when the treatment period is not backdated.

Concern that our post-treatment effects may be driven by over-fitting is probably somewhat mitigated by the pattern of post-treatment effects: as discussed above, the vast majority of the estimated mortality reduction effect is in the first three periods, when load shedding mitigation is highest, and then the estimated treatment effects gradually return to zero as load shedding mitigation drops. It would seem a surprising coincidence for spurious treatment effects based on over-fitting to match exactly this pattern in treatment intensity. However, one might remain concerned by the size of the gap especially in the very first treatment period, when divergence due to over-fitting can be particularly salient due to issues such as mean reversion.

[38] While I do separately observe deaths for those aged under one or two years old in the SAMRC data, the more aggregated nature of the Stats SA MYPE population data prevents me from excluding just these deaths and confirming that it is primarily neonatal deaths that drive these results.

Figure 7 shows an in-time placebo falsification test when treatment is artificially backdated by two periods, so that it starts at period –1 rather than period 1.[39] The left-most vertical line is the last pre-period for the artificial treatment, and the right-most vertical line is the last pre-period for the true treatment, which as always starts at $e = 1$. It is apparent from the figure that the synthetic control procedure does not generate some kind of immediate artificial post-treatment divergence; period –1 is almost exactly 0 (0.018), and the small divergence in period 0 is also small (–0.551) when considered in comparison to the deviations in the pre-treatment fit. Indeed, both estimates are not statistically distinguishable from zero, with p-values of 0.9735 and 0.8344 respectively. The same pattern of genuinely post-period treatment effects then begins at period 1, as in the main specification.

Figure 7: Bias-corrected treatment effects with backdated time placebo, deaths per 10,000 population, age ≥ 65

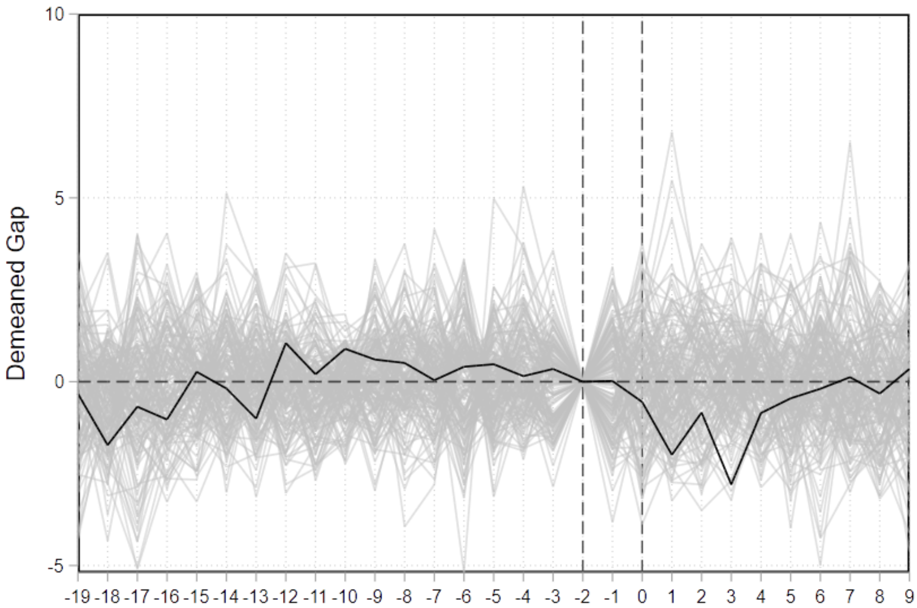

Note: the figure shows synthetic control gaps between Cape Town and synthetic Cape Town (solid dark line) for the main specification but where treatment is artificially backdated by two periods to start in period –1, averaged across the three events. Equivalently specified placebo gaps from in-place permutation tests are also shown (light grey lines). The outcome is deaths per 10,000 population, the sample is restricted to those aged 65 or older, and bias-corrected estimates are shown. The left-most dashed vertical line shows the last pre-period before the artificially backdated treatment starts in period –1. The right-most dashed vertical line shows the last pre-period before the true treatment starts in period 1.

Source: author's compilation.

Alternative event (common load shedding treatment) in-time placebo falsification test

While one may accept the evidence above as showing that load shedding causes fewer deaths in Cape Town than elsewhere, one may still question whether it is specifically load shedding mitigation policy in Cape Town that causes this. Given that in two out of the three events the start of load shedding mitigation coincides with the recommencement of load shedding nationally (see Section 4.2), with the pre-treatment period not having any load shedding, could the results instead be driven by heterogeneous responses to commonly experienced load shedding? This seems unlikely a priori: the first event does

[39] In order to retain the same pre-period length ($\underline{E} = 18$) as in the baseline case, I extend the pre-period to include $e = -19$. This means adjusting the assumption on the maximum period of post-treatment dynamic effects from $T_\delta = 4$ to $T_\delta = 2$ (see Section 4.2).

have common load shedding in the pre-period, and the pattern of post-treatment effects does follow the intensity of load shedding mitigation, not the intensity of load shedding itself.[40]

But this can be tested directly by creating events where the treatment is common load shedding, at times when Cape Town either didn't mitigate or such mitigation was negligible. Unfortunately, in the period of the data only two such events can be constructed: the recommencement of load shedding at the end of 2014, and potential event (d) discussed in Section 4.2, where very marginal mitigation was undertaken by Cape Town at the beginning of 2019. In the former case, this event was before the City of Cape Town implemented its mitigation scheme in 2015. In the latter case, the City did not significantly mitigate load shedding because the Steenbras plant was undergoing regular scheduled maintenance.[41] These events are both shown with hollow circles in Figure 3. The use of only two events means that this falsification test does not have the power of the three-event main specification, but in results not reported here I do find that separately implementing the main analysis for two-by-two combinations of events 1, 2, and 3 does yield statistically significant and qualitatively similar results as in the main analysis, even though pre-treatment and post-treatment gaps are much noisier. A bigger issue is that, as discussed in Section 4.2, potential event (d) only has 9 potential pre-periods, rather than 18, and this must be further restricted to allow for dynamic effects from event 2 (end of 2018) to dissipate. In order to maximize pre-treatment period length and allow for a decent matching period, I assume dynamic effects dissipate after at most two periods (so $T_\delta = 2$ rather than $T_\delta = 4$, and I have seven pre-treatment periods), and I include period $e = 0$ when constructing the average of pre-treatment outcomes as a predictor.

While the pre-treatment fit is unsurprisingly somewhat noisy, given fewer events and a truncated pre-period, the main result that comes out of this analysis (Figure 8) is that there is no hint of lower mortality sensitivity in Cape Town to common load shedding. The post-treatment point estimates are in fact slightly positive, though as one would expect they are not statistically significant at the 10% level. Figure A7 shows the results if one only uses the 2014 event by itself, which allows a full pre-period but is underpowered because there is no event stacking. It also shows no detectable effect and indeed no suggestion of any kind of effect.

[40] Unlike the case with load shedding mitigation, where the bulk of the treatment is in the first three periods (54%, see Table 1), when looking at load shedding implementation, 39% is in the first three periods and 54% is in the last three periods (see Figure 3).

[41] This is attested to by contemporaneous social media posts by the City: see City of Cape Town (2019a), City of Cape Town (2019b), and City of Cape Town (2019c).

Figure 8: Bias-corrected treatment effects with alternative event-time placebo, deaths per 10,000 population, age ≥ 65

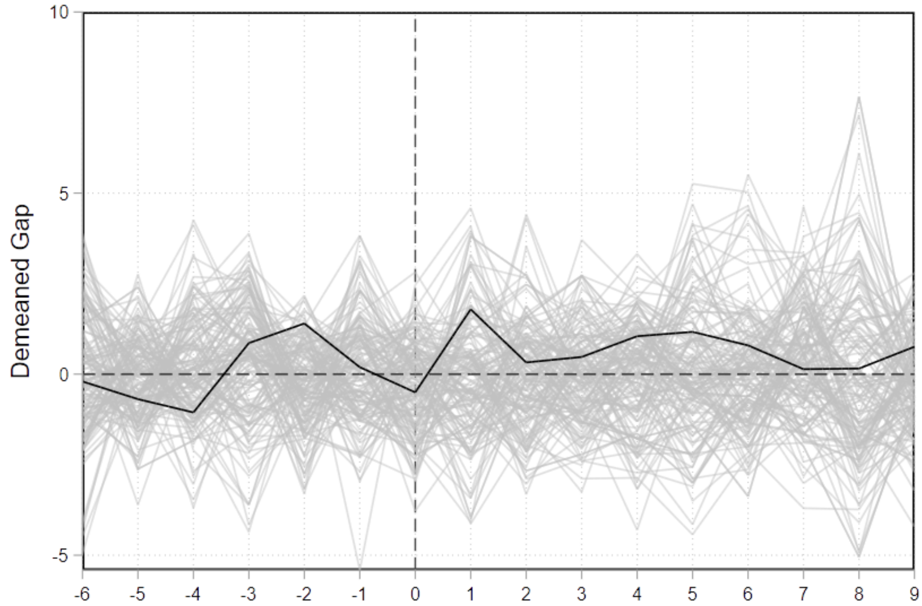

Note: the figure shows synthetic control gaps between Cape Town and synthetic Cape Town (solid dark line) for an in-time placebo test in 2014, where common load shedding starts in period 1 but there is no differential load shedding between Cape Town and national. Placebo gaps from in-place permutation tests are also shown (light grey lines). The outcome is deaths per 10,000 population, the sample is restricted to those aged 65 or older, and bias-corrected estimates are shown. The placebo treatment begins in period 1, the period immediately after the dashed vertical line.

Source: author's compilation.

6 Discussion

While the previous section discussed the plausibility of the synthetic control methodology in this context and the statistical significance of the results, in this section I focus on interpreting effect magnitudes.

6.1 Adjusting for partial and continuous treatment

Per the 2019 GHS, only 51% of the Cape Town population receives their mains electricity from the City itself, rather than from Eskom, and so only approximately this proportion of the population receives the City's load shedding mitigation.[42] The treatment effects estimated throughout this paper will therefore be attenuated if interpreted as the effects of the policy on those who actually receive treatment, because they are estimated by examining city-wide outcomes when only about half of the city receives treatment. One can, however, recover an estimate of the (local average) treatment effect on the treated, and further can extrapolate to the city at large if one is willing to make a (probably strong) assumption of intra-city homogeneity of treatment effects.

[42] There are a few geographically small areas which are City-supplied but load shed by Eskom; they cannot be identified in the GHS data so are ignored. The GHS data on electricity supplier is also based on self-reporting and not designed for intra-city analysis, so should be regarded as approximate in any case. Figure A9 compares GHS estimates to 2011 Census estimates of the City-supplied population proportion. These results suggest that while there is some noise or measurement error, the GHS and Census results match reasonably well in 2011 and 2012, supporting the approximate reliability of the GHS for this exercise. The 2011 Census estimates are not used for the main analysis because the GHS estimates imply that the City-supplied proportion steadily shrinks over time. This is likely due to disproportionate population growth in the city periphery, which has been established in prior literature (Scheba et al. 2021).

A useful framing is to consider residence in Cape Town as an instrument for the actual treatment, which is having electricity supplied directly by the City of Cape Town municipality. The results thus far presented are then interpretable as 'reduced-form' effects in an instrumental variables setting. The 'first stage', the effect of the instrument on treatment probability, is simply the proportion of Cape Town which is supplied by the City directly. Dividing the reduced form by the first stage—the Wald estimator—then yields the local average treatment effect (LATE) of the policy on those in the City's direct supply region, which is also the average treatment effect on the treated (ATT) (Angrist and Pischke 2009).[43]

The other aspect of partial treatment that needs to be considered is that—contrary to typical synthetic control settings—the treatment is not constantly 'on' in the post-period, and indeed varies in intensity over the post-period. While the synthetic control results provide estimates of the effect of these 'load shedding mitigation periods', of greater policy interest is the mortality-reducing effect of one less hour of (severity-weighted) load shedding. In order to calculate an approximate cost–benefit quantity of 'average total effect per unit of treatment' of this nature as described in De Chaisemartin and d'Haultfoeuille (2024), I follow their logic and divide the sum of cumulative post-period treatment effects (across events and periods) by the sum of cumulative post-period treatments (hours of load shedding mitigation, also across events and periods), to get an 'average mortality reduction per severity-weighted hour of load shedding'.[44]

These 'dosage responses' can be calculated directly from Table 1 and the relevant table of treatment effects. For example, using Table 2, with results from the main specification, the cumulative post-period treatment effect (averaged across events) is -5.747 deaths per 10,000 population, and per Table 2 the cumulative post-period treatment (averaged across events) is 9.92 mitigated severity-weighted load shedding hours.[45] The implied (reduced-form) treatment effect (those aged 65+) is therefore $5.747/9.92 = 0.579$ deaths per 10,000 population reduced per hour of mitigated load shedding.

6.2 Effects per hour of mitigated load shedding

For the specifications discussed in Section 5, Panel (a) of Table 3 presents these cumulative (over the whole post-period) reduced-form treatment effects, with their associated standard errors and 95% confidence intervals from the Firpo and Possebom (2018) RMSPE test-statistic inversion procedure.[46] Panel (b) then shows the implied reduced-form treatment effects per (severity-weighted) hour of reduced load shedding by dividing through by 9.92, as above. It is immediately evident that the aggregate treatment effects are quite imprecisely estimated when evaluated over the entire nine periods of the post-period. While the treatment effects of the main specification (bias-corrected, deaths per 10,000, age ≥ 65) are statistically significantly different from zero (as is also evident from Table 2), the 95% confidence intervals are very wide. The same applies to the natural deaths and log of deaths specifications. For the classic gaps and age ≥ 5 specifications, the aggregate post-period effects are not statistically significantly different from zero (as is also apparent from Tables 2 and A2).

[43] This is because those in the City's direct supply region are the 'compliers', and there is no always-taker group. The LATE monotonicity assumption clearly holds. The synthetic control methodology and results motivate for the conditional independence and exclusion restriction assumptions to also be satisfied.

[44] Fuller integration of the De Chaisemartin and d'Haultfoeuille (2024) methodology into this paper and its synthetic control setting is the subject of ongoing research.

[45] Using these averages across events makes no difference compared to summing across events; just multiply each quantity by 3 to get the summed totals.

[46] These confidence intervals can also be shown graphically, for average post-period effects, as for the main specification in Figure A8.

Table 3: Reduced-form treatment effects

Panel(a)			Cumulative post-period effect	
Outcome	Gap	Age	Estimate	95% CI
Deaths per 10,000 pop.	BC	≥65	–5.75 (2.43)	(–10.53, –0.99)
Nat. deaths per 10,000 pop.	BC	≥65	–4.81 (2.18)	(–9.09, –0.54)
Deaths per 10,000 pop.	BC	≥5	–0.54 (0.36)	(–1.23, 0.16)
Deaths per 10,000 pop.	Classic	≥65	–4.88 (3.21)	(–11.16, 1.44)
Log of deaths	BC	≥65	–0.56 (0.26)	(–1.07, –0.05)
Panel(b)			Effect per severity-weighted hr	
Outcome	Gap	Age	Estimate	95% CI
Deaths per 10,000 pop.	BC	≥65	–0.579 (0.245)	(–1.061, –0.1)
Nat. deaths per 10,000 pop.	BC	≥65	–0.485 (0.22)	(–0.916, –0.054)
Deaths per 10,000 pop.	BC	≥5	–0.054 (0.036)	(–0.124, 0.016)
Deaths per 10,000 pop.	Classic	≥65	–0.492 (0.324)	(–1.125, 0.145)
Log of deaths	BC	≥65	–0.057 (0.026)	(–0.108, –0.005)

Note: the table shows the reduced-form treatment effects of load shedding mitigation on mortality, cumulatively over the post-period (Panel (a)) and per severity-weighted hour of load shedding mitigation (Panel (b)). 'BC' refers to bias-corrected and 'Classic' to classic synthetic control gaps. The effect per severity-weighted hour is calculated by dividing the cumulative post-period effect (which is an average across the events) by the cumulative mitigation in the post-period (averaged across events) of 9.92 severity-weighted hours. The 95% confidence intervals are calculated using the RMSPE test-statistic inversion method of Firpo and Possebom (2018); the use of numerical methods and inference from a permutation distribution of 150 placebo units means confidence intervals are not always exactly symmetrical. The implied standard errors are shown in parentheses.

Source: author's calculations.

What do these magnitudes mean in terms of actual premature deaths averted by the policy? In order to estimate a treated effect on the treated, the reduced-form treatment effect needs to be divided by the first stage, as discussed above. Per the 2019 GHS, 72.9% (0.0355) of Cape Town's age 65+ population falls in the City-supplied rather than Eskom-supplied regions.[47] Using the reduced-form estimate above, the Wald estimate for the LATE and ATT on those aged 65 and older and actually in City-supplied regions is therefore $0.579/0.729 = 0.794$ deaths per 10,000 population reduced per hour of mitigated load shedding. Combining the average population age 65+ in the City of Cape Town over the three event periods (277,736) from the MYPE with this City-supplied vs Eskom-supplied proportion from the GHS (72.9%), I estimate the population of age 65+ in the City-supplied regions to be 202,470. Finally, combining this population number and adjusting for the scale of the outcome variable, I find $\frac{0.579}{0.729} \times \frac{202{,}470}{10{,}000} = 16.08$ premature deaths averted per (severity-weighted) hour of mitigated load shedding.

It is worth pointing out that the magnitude expressed in these terms (number of premature deaths averted per hour of mitigated load shedding) will be exactly the same for the LATE and the reduced-form treatment effect, because the (smaller) reduced-form effect applies to a proportionally larger population.[48] This is a desirable result given that the City-supplied vs Eskom-supplied proportionate population split is a highly approximate estimate (the GHS is not designed for intra-city analysis, see footnote 42), but it is worth reiterating that it comes from two different kinds of treatment effects. Using the reduced-form effects, one can conclude about the effect of the policy on deaths across Cape Town, keeping in mind that the policy only affected some households. One uses the LATE from the Wald estimator to discuss effects of the policy directly on the treated.

Table 4 shows the LATEs and implied deaths averted per hour of severity-weighted load shedding, with standard errors. It is immediately apparent again that the estimates are imprecise; the standard errors are large relative to the estimated treatment effects even where the treatment effects are statistically signif-

[47] Standard error shown in parentheses is calculated accounting for the complex survey design of the GHS. Per the same survey and calculation, 52.1% (0.0294) of the age 5+ population is in City-supplied rather than Eskom-supplied regions.

[48] i.e. $\frac{0.579}{0.729} \times \frac{277{,}736 \cdot 0.729}{10{,}000} = 0.579 \times \frac{277{,}736}{10{,}000}$.

icantly different from zero (the baseline, natural deaths, and log of deaths specifications). Comparison of the different results in Table 4 must keep in mind that the wide confidence intervals mean that they are not statistically distinguishable from each other. Taking the point estimate for the age 5+ category at face value, it seems that the vast majority of premature deaths averted come from the age 65+ category (75%), despite deaths among this latter group only constituting 42% of deaths in the 5+ category in Cape Town in the event pre-periods. This heightened sensitivity of elderly mortality to electricity rationing is consistent with the prior literature from other contexts. While deaths averted in the 65+ category are almost all natural deaths (84%), this needs to be understood in the context of almost all deaths among the 65+ group (97%) in Cape Town during the event pre-periods being natural rather than unnatural deaths, as well as the imprecision of the estimates. This is some tentative evidence that load shedding mitigation also reduces unnatural deaths, but unnatural deaths are too rare and noisy to be amenable to direct analysis of their effects via this paper's SCM; the pre-treatment fit is simply too poor to make substantive conclusions. The last two rows show that the main result is relatively robust to an alternative outcome transformation and estimation method.

Table 4: Treatment effects per (severity-weighted) hour of mitigated load shedding

Specification			Treatment effect	
Outcome	Gap	Age	LATE	Deaths
Deaths per 10,000	BC	$\geq$65	–0.795 (0.3387)	–16.1 (6.81)
Nat. deaths per 10,000 pop.	BC	$\geq$65	–0.666 (0.3033)	–13.5 (6.11)
Deaths per 10,000 pop.	BC	$\geq$5	–0.104 (0.0691)	–21.6 (14.36)
Deaths per 10,000 pop.	Classic	$\geq$65	–0.674 (0.4456)	–13.7 (9)
Log of deaths	BC	$\geq$65	–0.078 (0.0361)	–14.9 (6.87)

Note: the table shows the implied treatment effects of load shedding mitigation on mortality per severity-weighted hour of load shedding mitigation. 'BC' refers to bias-corrected and 'Classic' to classic synthetic control gaps. Local average treatment effects (LATE) refer to effects on the actually treated City-supplied households, per the Wald estimator, using the GHS to determine the proportion of the relative age group which is in a City-supplied region. As discussed in Section 6, the reduced-form and LATE estimates imply the same magnitudes in deaths averted per severity-weighted hour, which is shown in the last column. This may not apply when the outcome is log of deaths (it depends on whether deaths are distributed proportionally to population); for the log of deaths specification, the reduced-form estimate (Table 3) is used. Standard errors are shown in parentheses from Firpo and Possebom (2018) RMSPE test-statistic inversion. The LATE estimates use the Delta method to account for the standard errors around the estimates of the municipal/Eskom population split from the GHS.

Source: author's calculations.

Given that the City's policy reduced load shedding by 34 severity-weighted hours between 2014 and the end of 2019, the main point estimate implies that over this period the policy averted 547 premature deaths among those aged 65 and older, though averting 94 and 1,002 premature deaths are within the 95% confidence intervals.

6.3 Extrapolating beyond City-supplied regions and after 2019

While the results above concern how load shedding mitigation *reduced* mortality, it seems natural to also use them to estimate how much unmitigated load shedding *increased* mortality. For the effect of load shedding on the City-supplied regions, this is relatively straightforward: one can just use the LATE and implied deaths averted per severity-weighted hour of mitigation to conclude that each additional hour of load shedding causes 16.1 premature deaths among the age 65+ group in the City-supplied region. While this requires symmetry and linearity assumptions, it seems likely to be a reasonable approximation. There were 163 severity-weighted hours of load shedding in Cape Town from 2014 to 2019, so this estimate implies that load shedding over this period caused 2,623 deaths among the age 65+ group in the City-supplied region (95% confidence intervals of 452 and 4,805).

If the objective is to understand the effect on deaths across the City of Cape Town, or indeed across the country as a whole, the exercise becomes much more fraught, as the LATE must be extrapolated to these regions. This is fraught because the City-supplied region to which the LATE applies is very different to the rest of Cape Town or indeed the rest of the country, as Table 5 shows, and it is very likely that treatment effects will be heterogeneous across these regions. In particular, the City-supplied region is much richer and older.[49] The implication for treatment effect magnitudes are a priori ambiguous. On the one hand, greater age and higher use of electrical appliances will likely mean that the City-supplied region is atypically sensitive to load shedding. On the other hand, higher income means more possibilities to ameliorate load shedding, and better healthcare (illustrated by much higher rates of private medical insurance).

Table 5: Demographic characteristics of electricity supply regions

	CPT direct	CPT Eskom	Rest of country
African	21%	60%	86%
Coloured	52%	33%	6%
White	25%	6%	6%
Mean age	34	29	28
Urban	100%	100%	62%
RDP	16%	42%	21%
Mean household income pc	4,254	2,707	2,312
Household grant receipt	39%	55%	63%
Medical insurance	42%	16%	16%
Employed	39%	37%	29%
Free basic electricity	22%	52%	22%
Electric heating	51%	39%	49%
Piped water or shared tap	99%	99%	86%

Note: the table shows characteristics of three different regions: City-supplied households in Cape Town (CPT direct), Eskom-supplied households in Cape Town (CPT Eskom), and the rest of South Africa. Statistics are calculated using the 2019 GHS and weighted using the household weight. In all cases statistics are at the individual level—for example, 'Household grant receipt' means the proportion of individuals who live in a household which receives at least one grant. 'African', 'Coloured', and 'White' follow the racial classifications of the South African census. 'RDP' means the respondent lives in an RDP or state-subsidised dwelling. 'Mean household income pc' is approximate monthly mean household income per capita in rand, derived from the GHS total household income variable, which is severely top-censored. 'Free basic electricity' means the respondent's household receives the free basic electricity subsidy. Piped water or shared tap refers to drinking water.

Source: author's compilation.

If the LATE is extended to include the Eskom-supplied region of Cape Town, with its 65+ population of 75,266, so a total Cape Town 65+ population of 277,736, this implies that each additional hour of severity-weighted load shedding increases premature deaths for this age group across the city by 22.08 (95% confidence intervals of approximately 3.64 and 40.52). That this is not substantially larger than the effect just in the City-supplied regions reflects that a large majority of this population lives in City-supplied regions.

While an equivalent exercise can be conducted for the age 5+ population in the City, or for any age group for the country at large, I judge this to be too speculative. In the former case, one would expect the age differences between the LATE population and Eskom-supplied region to be particularly salient, while in the latter case Table 5 shows that the rest of the country is just very substantially different from the City-supplied region in Cape Town across multiple relevant dimensions.[50] In general, extrapolating the 65+ LATE to the rest of the city is more defensible because the LATE population forms a large majority of the city's population that is 65 and older, meaning there is less extrapolation. This is not the case for the age 5+ results (which are so imprecise as to not be statistically significant in any case), and certainly not for extrapolation to the rest of the country.

A last extrapolation that might be considered is applying the 'deaths per hour' result not just to the load shedding (mitigation) of 2014–19, but also to the 2020–23 period when load shedding (and its

[49] This is likely driven to a significant degree by enduring apartheid legacies and its substantial over-representation of whites. See Section 7.

[50] In particular, the high urbanity of the City compared to the rest of the country is likely be important. Indeed, for Ghana Apenteng et al. (2018) find a lower correlation between electricity shortages and mortality in rural areas than in urban areas.

mitigation) was substantially more intense (see Figure 1).[51] Between 2014 and October 2023, 2,576 and 1,955 severity-weighted hours of load shedding were implemented nationally and in the City-supplied region of Cape Town, respectively, while the respective figures between 2014 and the end of 2019 are 197 and 163 severity-weighted hours. It is clear that by focusing on 2014 to 2019, this paper's results examine a relatively small part of load shedding implementation in South Africa.

And yet it is precisely the extraordinary increase in load shedding from 2020, and especially from 2022, which leads me to judge an extrapolation of this nature to be ill-advised.[52] Load shedding was experienced as a qualitatively different phenomena during this latter time period, which entailed a 'not-so-quiet revolution' (Creamer 2024a) in South Africa's energy landscape as households and businesses responded to government failure by taking energy supply into their own hands where they could. Consumers started to invest heavily in load shedding ameliorating mechanisms, as is reflected by retailer reports of substantial increases in demand and new product lines for household batteries and inverters, LPG gas cylinders, uninterrupted power supplies (UPS), and generators, which they directly attribute to increased load shedding.[53] Telecoms companies substantially expanded investments in back-up power to keep networks operating, while imports of solar panels, lithium-ion batteries, and inverters spiked dramatically from 2022.[54] In the public sector, starting from 2022 the government began a programme of exempting large state hospitals from load shedding (Mnyandu 2023; Phaahla 2023). This reflected a general recognition in the state that load shedding had reached unprecedented crisis levels, and it prompted a variety of reforms, such as dramatically increasing the capacity limit on private distributed generation in August 2021, before abolishing the limit entirely in December 2022 (Creamer 2023), declaration of a National State of Disaster in February 2023 (Mnukwa 2023), and the creation of a new cabinet post of Minister of Electricity in May 2023 (Magwenya 2023).

These particular interventions were unevenly experienced and far from all-encompassing, but reflect a qualitative shift in how load shedding was experienced and responded to, and thus make extrapolation of 2015–19 effects to this period probably ill-advised. This does of course mean that the mortality effects estimated here fall far short of a full accounting of the costs of load shedding or the benefits of its mitigation.

6.4 Comparing magnitudes to existing estimates

Per the SAMRC mortality data, approximately 50,373 people age 65 and older in Cape Town died in 2014, 2015, 2018, and 2019. Extending the 65+ LATE across Cape Town, and multiplying by 163 severity-weighted load shedding hours in this period, the point estimate from this paper implies that 7% of these deaths in Cape Town in these years were linked to load shedding. A priori, this effect seems improbably large. And indeed, the imprecision of the estimated treatment effects make this quite possible; at the 95% confidence interval we cannot reject as few as 1% of these deaths being caused by load shedding. The main contribution of this paper is to show that load shedding mitigation reduces mortality, and conversely that load shedding itself increases mortality; there should not be too much emphasis placed on the specific point estimates.

But how do these point estimates compare to those in the existing literature? There are no directly comparable estimates. First, as detailed above, we would expect varying treatment effects depending on the local demographic and social conditions, and the existing literature all studies contexts outside

[51] The 2007/08 load shedding period cannot be included because we lack hourly data on load shedding implementation.

[52] An extrapolation to the 2020–22 period would also need to consider how load shedding treatment effects may have interacted with the effects of COVID-19 and it's associated lockdowns, which is beyond my scope.

[53] See, for example, Mashego (2022), Mdaka (2023), Patrick (2022), Payi (2023), and Wilson (2023).

[54] See, for example, Gavaza (2022), Mahlaka (2024), and Odendaal (2023) on telecoms, and Creamer (2024b) on imports.

South Africa. More fundamentally, none of the existing literature examines effects on all-cause mortality as we do here, and indeed most of it does not concern power cuts, but power rationing via prices or moral suasion. Nonetheless, it is useful to discuss the existing estimates when considering this paper's magnitudes, mostly because they are also surprisingly large.

Apenteng et al. (2018), the only study that looks at the effects of rolling blackouts rather than rationing, finds an astonishingly large effect; that for each day the power was out for over 2 hours, in-facility mortality at healthcare facilities in Ghana (excluding referral hospitals) increased by 43%.

Neidell et al. (2021), He and Tanaka (2023), and Chirakijja et al. (2024) examine the mortality effects of electricity rationing, but only examine how rationing affects mortality through the channel of temperature regulation. Neidell et al. (2021) and He and Tanaka (2023) examine how electricity rationing in post-Fukushima Japan affects the temperature–mortality relationship, while Chirakijja et al. (2024) examine how variation in heating prices (specifically domestic natural gas) affects mortality. Neidell et al. (2021) find that 19% of cold-related deaths in Japan during their study period are associated with the price increase due to the nuclear power shutdown, while He and Tanaka (2023) find that energy-saving campaigns caused 7,710 deaths per year in Japan among those aged 65+, which accounts for 0.7% of total national mortality. Chirakijja et al. (2024) find that the 42% relative decline in the price of natural gas in the United States between 2005 and 2010 caused a 1.7% decrease in the winter mortality rate for households using natural gas.

While these estimates are not comparable to my all-cause mortality estimates, it is notable that they are quite substantial. The Apenteng et al. (2018) effect is self-evidently very big. But the Neidell et al. (2021), He and Tanaka (2023), and Chirakijja et al. (2024) estimates are also surprisingly large when one considers that they are mortality effects via one channel (temperature regulation) in a context in which consumers have some discretion over their cut-offs, but they still have substantial national impacts. *Ceteris paribus*, one would expect greater impacts when examining all-cause mortality, and when household discretion over electricity cut-offs is more limited. There are of course other countervailing factors, such as the older age structure and greater temperature extremes in these countries.

7 Conclusion

This paper presents causal evidence that load shedding mitigation in Cape Town reduces mortality, particularly among the elderly. While the magnitudes are imprecisely estimated, and depend on the characteristics of the treated region and time period, the results accord with the small existing literature on this topic, which finds that electricity rationing can cause premature deaths. The main results, looking at the age 65+ population in the City-supplied regions, imply that the mitigation policy prevented approximately 547 premature deaths in this population in Cape Town between 2014 and 2019 (95% confidence intervals of 94 and 1,002). The estimate further implies that load shedding over this period caused 2,623 deaths in this population in the City-supplied region (95% confidence intervals of 452 and 4,805). A few additional points for consideration should be kept in mind when interpreting the results, and suggest areas for future research.

First, I use the term 'premature deaths' throughout this paper advisedly. It is quite plausible that the mortality effects of load shedding are concentrated among those who are already highly vulnerable and unlikely to live much longer even in its absence. It is possible that the effects are mainly 'timing' or so-called 'harvesting' effects, such that mortality is pushed ahead in time, but causes few deaths where the counterfactual is much longer life. We do not observe a mortality *increase* in Cape Town in the latter part of the post-period, which is what we would expect to see if the results were completely driven by harvesting effects, but the imprecision of the estimates means this cannot be ruled out. An investigation that accounts for 'life years lost' would provide some measure of the extent to which this drives results.

Second, the local average treatment effects estimated here, while useful, clearly do not answer the larger question of the total mortality effects of load shedding in South Africa. This would require a method that can estimate effects during recent years, and estimate effects for the country at large. If load shedding were implemented at random, a simple regression of national mortality on load shedding severity might provide useful estimates, but load shedding has seasonal and secular trends which would confound this analysis. Future research that could answer this broader question using a plausibly causal design would be very valuable.

Lastly, while the Cape Town load shedding mitigation policy is undoubtedly a positive development, its implications for inequality in the City deserves greater attention. The racial and income disparities between areas that are City-supplied versus Eskom-supplied in Table 5 are incredibly stark, if unsurprising for someone familiar with Cape Town's geography who has also seen Figure A2. Figure A9 shows that per the 2019 GHS, 81% of whites in Cape Town receive load shedding mitigation while only 27% of black Africans in the City receive such mitigation. Granular mortality data with greater demographic information such as race and socioeconomic status would be useful for understanding the inequality implications of the policy, as well as for extrapolating the results beyond Cape Town's City-supplied regions.

In this last respect, the policy reflects a not atypical dilemma in the South African context. The Steenbras pumped-storage hydroelectric power station and (presumably) the extent of the City-supplied electricity network both preceded post-apartheid government in the City of Cape Town. It is to the City's credit that it took advantage of this existing well-developed infrastructure to ameliorate load shedding for some of its residents. And the City has, since 2022, initiated efforts to expand its supply remit to include the current Eskom-supplied regions.[55] But it is unsurprising that the present-day use of this historical infrastructure disproportionately benefits a whiter and richer population, given that this is the population it was most likely originally designed to prioritize.

References

Abadie, A. (2021). 'Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects'. *Journal of Economic Literature*, 59(2): 391–425. https://doi.org/10.1257/jel.20191450

Abadie, A., and J. Gardeazabal (2003). 'The Economic Costs of Conflict: A Case Study of the Basque Country'. *American Economic Review*, 93(1): 113–32. https://doi.org/10.1257/000282803321455188

Abadie, A., and J. L'Hour (2021). 'A Penalized Synthetic Control Estimator for Disaggregated Data'. *Journal of the American Statistical Association*, 116(536): 1817–34. https://doi.org/10.1080/01621459.2021.1971535

Abadie, A., A. Diamond, and J. Hainmueller (2015). 'Comparative Politics and the Synthetic Control Method'. *American Journal of Political Science*, 59(2): 495–510. https://doi.org/10.1111/ajps.12116

Abadie, A., A. Diamond, and J. Hainmueller (2010). 'Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program'. *Journal of the American Statistical Association*, 105(490): 493–505. https://doi.org/10.1198/jasa.2009.ap08746

Abeberese, A.B. (2017). 'Electricity Cost and Firm Performance: Evidence from India'. *Review of Economics and Statistics*, 99(5): 839–52. https://doi.org/10.1162/REST_a_00641

[55] See Githahu (2023) and Naidoo (2023). Eskom recently rejected the City's request, reportedly because it seeks to manage the process as part of its broader ongoing unbundling process (Omarjee 2024).

Abeberese, A.B. (2020). 'The Effect of Electricity Shortages on Firm Investment: Evidence from Ghana'. *Journal of African Economies*, 29(1): 46–62.

Abeberese, A.B., C.G. Ackah, and P.O. Asuming (2021). 'Productivity Losses and Firm Responses to Electricity Shortages: Evidence from Ghana'. *World Bank Economic Review*, 35(1): 1–18. https://doi.org/10.1093/wber/lhz027

Aklin, M., P. Bayer, S. Harish, and J. Urpelainen (2017). 'Does Basic Energy Access Generate Socioeconomic Benefits? A Field Experiment with Off-Grid Solar Power in India'. *Science Advances*, 3(5): e1602153. https://doi.org/10.1126/sciadv.1602153

Allcott, H., A. Collard-Wexler, and S.D. O'Connell (2016). 'How Do Electricity Shortages Affect Industry? Evidence from India'. *American Economic Review*, 106(3): 587–624. https://doi.org/10.1257/aer.20140389

Angrist, J.D., and J.-S. Pischke (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press. https://doi.org/10.1515/9781400829828

Apenteng, B.A., S.T. Opoku, D. Ansong, E.A. Akowuah, and E. Afriyie-Gyawu (2018). 'The Effect of Power Outages on In-Facility Mortality in Healthcare Facilities: Evidence from Ghana'. *Global Public Health*, 13(5): 545–55. https://doi.org/10.1080/17441692.2016.1217031

Arkhangelsky, D., S. Athey, D.A. Hirshberg, G.W. Imbens, and S. Wager (2021). 'Synthetic Difference-in-Differences'. *American Economic Review*, 111(12): 4088–118. https://doi.org/10.1257/aer.20190159

Ben-Michael, E., A. Feller, and J. Rothstein (2021). 'The Augmented Synthetic Control Method'. *Journal of the American Statistical Association*, 116(536): 1789–803.

Ben-Michael, E., A. Feller, and J. Rothstein (2022). 'Synthetic Controls with Staggered Adoption'. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(2): 351–81. https://doi.org/10.1111/rssb.12448

Bradshaw, D., R. Dorrington, R. Laubscher, P. Groenewald, and T. Moultrie (2022). 'COVID-19 and All-Cause Mortality in South Africa: The Hidden Deaths in the First Four Waves'. *South African Journal of Science*, 118(5–6): 1–7. https://doi.org/10.17159/sajs.2022/13300

Burlig, F., and L. Preonas (2016). 'Out of the Darkness and into the Light? Development Effects of Rural Electrification'. Energy Institute Working Paper 268. Berkeley, CA: Energy Institute at Haas.

Buthelezi, L. (2022). 'Load Shedding Insurance Claims Surge Up to 60% as Burglars Take Advantage of Blackouts'. *News24 Business*, 21 September. Available at: https://www.news24.com/fin24/Companies/load-shedding-insurance-claims-surge-up-to-60-as-burglars-take-advantage-of-blackouts-20220921 (accessed 7 April 2024).

Cavallo, E., S. Galiani, I. Noy, and J. Pantano (2013). 'Catastrophic Natural Disasters and Economic Growth'. *Review of Economics and Statistics*, 95(5): 1549–61. https://doi.org/10.1162/REST_a_00413

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). 'The Effect of Minimum Wages on Low-Wage Jobs'. *Quarterly Journal of Economics*, 134(3): 1405–54. https://doi.org/10.1093/qje/qjz014

Chernozhukov, V., K. Wuthrich, and Y. Zhu (2018). 'A *t*-Test for Synthetic Controls'. arXiv preprint arXiv:1812.10820.

Chirakijja, J., S. Jayachandran, and P. Ong (2024). 'The Mortality Effects of Winter Heating Prices'. *The Economic Journal*, 134(657): 402–17. https://doi.org/10.1093/ej/uead072

City of Cape Town (2019a). Facebook post from City of Cape Town (CityofCT). 11 February. Available at: https://www.facebook.com/CityofCT/posts/2067126506657591 (accessed 30 May 2024).

City of Cape Town (2019b). Tweet from City of Cape Town (@CityofCT). 14 March. Available at: https://x.com/CityofCT/status/1106198674500186114 (accessed 30 May 2024).

City of Cape Town (2019c). Tweet from City of Cape Town (@CityofCT). 18 March. Available at: https://x.com/CityofCT/status/1107575500434784256 (accessed 30 May 2024).

Copelyn, J. (2022). 'Torchlight Tragedy: Inside a Joburg Emergency Room During Load-Shedding'. *Mail & Guardian*, 22 September. Available at: https://mg.co.za/top-six/2022-09-22-torchlight-tragedy-inside-a-joburg-emergency-room-during-load-shedding/ (accessed 7 April 2024).

Creamer, T. (2023). 'Power of Reform'. *Engineering News*, 9 June. Available at: https://www.engineeringnews.co.za/print-version/power-of-reform-2023-06-09 (accessed 7 April 2024).

Creamer, T. (2024a). 'Not so Quiet Revolution'. *Engineering News*, 15 March. Available at: https://www.engineeringnews.co.za/article/not-so-quiet-revolution-2024-03-15 (accessed 7 April 2024).

Creamer, T. (2024b). 'Solar, Battery, Inverter Imports Surged to R70bn in 2023 as Wind Turbines Recovered from Two-Year Lull'. *Engineering News*, 4 March. Available at: https://www.engineeringnews.co.za/article/solar-battery-inverter-imports-surged-to-r70bn-in-2023-as-wind-turbines-recovered-from-two-year-lull-2024-03-04 (accessed 7 April 2024).

De Chaisemartin, C., and X. d'Haultfoeuille (2024). 'Difference-in-Differences Estimators of Intertemporal Treatment Effects'. *Review of Economics and Statistics*. https://doi.org/10.1162/rest_a_01414

Dinkelman, T. (2011). 'The Effects of Rural Electrification on Employment: New Evidence from South Africa'. *American Economic Review*, 101(7): 3078–108. https://doi.org/10.1257/aer.101.7.3078

Doudchenko, N., and G.W. Imbens (2016). 'Balancing, Regression, Difference-in-Differences and Synthetic Control Methods: A Synthesis'. Working Paper 22791. Cambridge, MA: NBER. https://doi.org/10.3386/w22791

Dube, A., and B. Zipperer (2015). 'Pooling Multiple Case Studies Using Synthetic Controls: An Application to Minimum Wage Policies'. Discussion Paper 8944. Bonn: IZA. https://doi.org/10.2139/ssrn.2589786

Ferman, B., and C. Pinto (2021). 'Synthetic Controls with Imperfect Pretreatment Fit'. *Quantitative Economics*, 12(4): 1197–221. https://doi.org/10.3982/QE1596

Firpo, S., and V. Possebom (2018). 'Synthetic Control Method: Inference, Sensitivity Analysis and Confidence Sets'. *Journal of Causal Inference*, 6(2): 20160026. https://doi.org/10.1515/jci-2016-0026

Fisher-Vanden, K., E.T. Mansur, and Q.J. Wang (2015). 'Electricity Shortages and Firm Productivity: Evidence from China's Industrial Firms'. *Journal of Development Economics*, 114: 172–88. https://doi.org/10.1016/j.jdeveco.2015.01.002

Gavaza, M. (2022). 'MTN Sounds Alarm on Network Crisis Due to Load-Shedding'. *Business Day*, 16 December. Available at: https://www.businesslive.co.za/bd/companies/telecoms-and-technology/2022-12-16-mtn-sounds-alarm-on-network-crisis-due-to-load-shedding/ (accessed 7 April 2024).

Gehringer, C., H. Rode, and M. Schomaker (2018). 'The Effect of Electrical Load Shedding on Pediatric Hospital Admissions in South Africa'. *Epidemiology*, 29(6): 841. https://doi.org/10.1097/EDE.0000000000000905

Ginindza, B. (2023). 'Clinic Load Shedding Blamed for Death!' *Daily Sun*, 6 March. Available at: https://www.snl24.com/dailysun/news/clinic-load-shedding-blamed-for-death-20230306 (accessed 7 April 2024).

Githahu, M. (2023). 'City, Eskom in Talks to Protect All of Cape Town from Some Load Shedding Stages'. *Cape Argus*, 8 February. Available at: https://www.iol.co.za/capeargus/news/city-eskom-in-talks-to-protect-all-of-cape-town-from-some-load-shedding-stages-066308a0-e351-4bc2-b59f-72637c535847 (accessed 5 May 2024).

Grimm, M., A. Munyehirwe, J. Peters, and M. Sievert (2017). 'A First Step Up the Energy Ladder? Low Cost Solar Kits and Household's Welfare in Rural Rwanda'. *World Bank Economic Review*, 31(3): 631–49.

Grogan, L., and A. Sadanand (2013). 'Rural Electrification and Employment in Poor Countries: Evidence from Nicaragua'. *World Development*, 43: 252–65. https://doi.org/10.1016/j.worlddev.2012.09.002

Hardy, M., and J. McCasland (2021). 'Lights Off, Lights On: The Effects of Electricity Shortages on Small Firms'. *World Bank Economic Review*, 35(1): 19–33. https://doi.org/10.1093/wber/lhz028

He, G., and T. Tanaka (2023). 'Energy Saving May Kill: Evidence from the Fukushima Nuclear Accident'. *American Economic Journal: Applied Economics*, 15(2): 377–414. https://doi.org/10.1257/app.20200505

Irwin, B.R., K. Hoxha, and K.A. Grépin (2020). 'Conceptualising the Effect of Access to Electricity on Health in Low- and Middle-Income Countries: A Systematic Review'. *Global Public Health*, 15(3): 452–73. https://doi.org/10.1080/17441692.2019.1695873

JacaNews (2015). 'Investigation Confirms Patient's Death Due to Load Shedding'. *Jacaranda FM*, 22 September. Available at: https://www.jacarandafm.com/news/news/investigation-confirms-patients-death-due-load-shedding/ (accessed 7 April 2024).

Khan, T., and L. Mkentane (2023). 'How Power Cuts Put Lives and Health at Risk'. *Business Day*, 18 January. Available at: https://www.businesslive.co.za/bd/national/2023-01-18-how-power-cuts-put-lives-and-health-at-risk/ (accessed 7 April 2024).

Laher, A., B. Van Aardt, A. Craythorne, M. Van Welie, D. Malinga, and S. Madi (2019). '"Getting Out of the Dark": Implications of Load Shedding on Healthcare in South Africa and Strategies to Enhance Preparedness'. *South African Medical Journal*, 109(12): 899–901. https://doi.org/10.7196/SAMJ.2019.v109i12.14322

Lawson, K. (2022). 'Electricity Outages and Residential Fires: Evidence from Cape Town, South Africa'. *South African Journal of Economics*, 90(4): 469–85. https://doi.org/10.1111/saje.12329

Lee, K., E. Miguel, and C. Wolfram (2020a). 'Does Household Electrification Supercharge Economic Development?' *Journal of Economic Perspectives*, 34(1): 122–44. https://doi.org/10.1257/jep.34.1.122

Lee, K., E. Miguel, and C. Wolfram (2020b). 'Experimental Evidence on the Economics of Rural Electrification'. *Journal of Political Economy*, 128(4): 1523–65. https://doi.org/10.1086/705417

Lipscomb, M., A.M. Mobarak, and T. Barham (2013). 'Development Effects of Electrification: Evidence from the Topographic Placement of Hydropower Plants in Brazil'. *American Economic Journal: Applied Economics*, 5(2): 200–31. https://doi.org/10.1257/app.5.2.200

Magwenya, V. (2023). 'President Cyril Ramaphosa Sets Out Role and Responsibilities of Minister of Electricity. Issued by: The Presidency'. Media Statement, 26 May. Available at: https://www.gov.za/speeches/president-cyril-ramaphosa-sets-out-role-and-responsibilities (accessed 8 April 2024).

Mahlaka, R. (2024). 'MTN Pours Billions More into Network to Boost Resilience to Eskom's Rolling Blackouts'. *Daily Maverick*, 25 March. Available at: https://www.dailymaverick.co.za/article/2024-03-25-mtn-pours-billions-more-into-network-to-boost-resilience-to-eskoms-rolling-blackouts/ (accessed 7 April 2024).

Mashego, P. (2022). 'Game Sees 27

Matthews, S. (2015). 'Water and Energy: Hydropower—Keeping the Lights on in the Mother City'. *Water Wheel*, 14(6): 22–25.

McCain, N. (2022). 'Calls for Reduced Load Shedding at Water Plants Amid Fears of Water Cuts, Sewage Overflows'. *News24*, 26 September. Available at: https://www.news24.com/news24/southafrica/news/calls-for-reduced-load-shedding-at-water-plants-amid-fears-of-water-cuts-sewage-overflows-20220926 (accessed 7 April 2024).

Mdaka, Y. (2023). 'Load Shedding: "UPS" the Third Most-Searched Term on Takealot Last Year'. *News24*, 24 March. Available at: https://www.news24.com/news24/tech-and-trends/load-shedding-ups-the-third-most-searched-term-on-takealot-last-year-20230324 (accessed 7 April 2024).

Mensah, J.T. (2024). 'Jobs! Electricity Shortages and Unemployment in Africa'. *Journal of Development Economics*, 167: 103231. https://doi.org/10.1016/j.jdeveco.2023.103231

Minnaar, U., W. Visser, and J. Crafford (2017). 'An Economic Model for the Cost of Electricity Service Interruption in South Africa'. *Utilities Policy*, 48: 41–50. https://doi.org/10.1016/j.jup.2017.08.010

Mnukwa, N. (2023). 'Government Issues Disaster Management Regulations Pertaining to the Impact of Severe Electricity Supply Constraints. Issued by: Government Communications'. Media Statement, 28 February. Available at: https://www.gov.za/news/media-statements/government-issues-disaster-management-regulations-pertaining-impact-severe (accessed 8 April 2024).

Mnyandu, E. (2023). 'Government Reiterates Its Determination to End Load Shedding. Issued by: Department of Public Enterprises'. Media Statement, 29 May. Available at: https://www.gov.za/news/media-statements/government-reiterates-its-determination-end-load-shedding-29-may-2023 (accessed 8 April 2024).

Monama, T. (2022). 'Calls Grow for Load Shedding Exemptions for Healthcare Facilities'. *News24*, 22 September. Available at: https://www.news24.com/news24/health/calls-grow-for-load-shedding-exemptions-for-healthcare-facilities-20220922 (accessed 7 April 2024).

Morema, K., J. Rakgalakane, T. Alton, and P. Mjandana (2019). 'A Novel Supply-Side Approach to Estimating the Impact of Load Shedding on GDP Growth in South Africa'. SARB Occasional Bulletin of Economic Notes 1902. Pretoria: South African Reserve Bank.

Motloung, T. (2023). 'Joburg Girl, 3, Dies after Breathing Machine Reportedly Could Not Function Due to Load Shedding'. *News24*, 26 May. Available at: https://www.news24.com/news24/southafrica/news/joburg-girl-3-dies-after-breathing-machine-reportedly-could-not-function-due-to-load-shedding-20230526 (accessed 28 May 2024).

Mthethwa, C. (2023a). 'Patient Dies in Front of Limpopo Clinic after Security Guard Denied Him Access Due to Load Shedding'. *News24*, 26 November. Available at: https://www.news24.com/news24/southafrica/news/patient-dies-in-front-of-limpopo-clinic-after-security-guard-denied-him-access-due-to-load-shedding-20231126 (accessed 7 April 2024).

Mthethwa, C. (2023b). '"We Were Struggling to Breathe": Five Farm Workers Died of Heat Stroke in Sweltering Northern Cape Heat. *News24*, 22 January. Available at: https://www.news24.com/news24/southafrica/news/we-were-struggling-to-breathe-five-farm-workers-died-of-heat-stroke-in-sweltering-northern-cape-20230122 (accessed 7 April 2024).

Naidoo, P. (2023). 'Cape Town Wants to Give Eskom the Hoof'. *Tech Central*, 25 May. Available at: https://techcentral.co.za/cape-town-wants-to-give-eskom-the-hoof/226102/ (accessed 5 May 2024).

Neidell, M., S. Uchida, and M. Veronesi (2021). 'The Unintended Effects from Halting Nuclear Power Production: Evidence from Fukushima Daiichi Accident'. *Journal of Health Economics*, 79: 102507. https://doi.org/10.1016/j.jhealeco.2021.102507

Odendaal, N. (2023). 'Telcos Feeling Pressure of Sustained Levels of Loadshedding'. *Engineering News*, 24 February. Available at: https://www.engineeringnews.co.za/article/telcos-feeling-pressure-of-sustained-levels-of-loadshedding-2023-02-24 (accessed 7 April 2024).

Omarjee, L. (2024). 'Eskom Says It Halted Cape Town Supply Takeover to Avoid "Haphazard Approach" During Unbundling'. *News24*, 15 May. Available at: https://www.news24.com/fin24/climate_future/energy/eskom-says-it-halted-cape-town-supply-takeover-to-avoid-haphazard-approach-during-unbundling-20240515 (accessed 28 May 2024).

Patrick, A. (2022). 'Customers Rushing to Buy Inverters as Load-Shedding Crisis Worsens'. *Sowetan Live*, 4 July. Available at: https://www.sowetanlive.co.za/news/south-africa/2022-07-04-customers-rushing-to-buy-inverters-as-load-shedding-crisis-worsens/ (accessed 7 April 2024).

Payi, B. (2023). 'Gas Demand Rises in Wake of Load Shedding'. *Weekend Argus*, 5 February. Available at: https://www.iol.co.za/weekend-argus/news/gas-demand-rises-in-wake-of-load-shedding-4bd0c8b4-df32-42b6-9ca8-5010d98dd312 (accessed 7 April 2024).

Phaahla, J. (2023). 'Minister Joe Phaahla on the Impact of Loadshedding on the Provision of Healthcare Services and Intervention Measures. Issued by: Department of Health'. Media Statement, 30 September. Available at: https://www.gov.za/news/media-statements/minister-joe-phaahla-impact-loadshedding-provision-healthcare-services-and (accessed 8 April 2024).

Pijoos, I. (2022). '"My Health Is in the Hands of Someone Else": Joburg Resident Amid Ongoing Power Cuts'. *News24*, 21 September. Available at: https://www.news24.com/news24/SouthAfrica/News/my-health-is-in-the-hands-of-someone-else-joburg-resident-amid-ongoing-power-cuts-20220921 (accessed 7 April 2024).

Rud, J.P. (2012). 'Electricity Provision and Industrial Development: Evidence from India'. *Journal of Development Economics*, 97(2): 352–67. https://doi.org/10.1016/j.jdeveco.2011.06.010

SABC (2023). 'Seven Die of Heatstroke in the Northern Cape'. *SABC News*, 21 January. Available at: https://www.youtube.com/watch?v=ZINsYehdA8M (accessed 7 April 2024).

Scheba, A., I. Turok, and J. Visagie (2021). 'Inequality and Urban Density: Socio-Economic Drivers of Uneven Densification in Cape Town'. *Environment and Urbanization ASIA*, 12(1 suppl.): S107–26. https://doi.org/10.1177/0975425321998026

South African Reserve Bank (2019). *Quarterly Bulletin, September 2019*. Pretoria: South African Reserve Bank.

South African Reserve Bank (2022). *Quarterly Bulletin, December 2022*. Pretoria: South African Reserve Bank.

Thomas, D.R., S. Harish, R. Kennedy, and J. Urpelainen (2020). 'The Effects of Rural Electrification in India: An Instrumental Variable Approach at the Household Level'. *Journal of Development Economics*, 146: 102520. https://doi.org/10.1016/j.jdeveco.2020.102520

Walsh, K., R. Theron, and C. Reeders (2021). 'Estimating the Economic Cost of Load Shedding in South Africa'. Technical Report. Mimeo.

Wilson, N. (2023). 'Big Rush for Inverters, Generators and Gas Stoves, as Retailers Warn of Erratic Supplies'. *News24 Business*, 26 January. Available at: https://www.news24.com/fin24/companies/big-rush-for-inverters-generators-and-gas-stoves-as-retailers-warn-of-erratic-supplies-20230126 (accessed 7 April 2024).

Wiltshire, J.C. (2022). 'Allsynth: (Stacked) Synthetic Control Bias-Correction Utilities for Stata'. Working Paper. Mimeo.

Wiltshire, J.C. (2023). 'Walmart Supercenters and Monopsony Power: How a Large, Low-Wage Employer Impacts Local Labor Markets'. Working Paper. Mimeo.

Appendix A: Additional figures and tables

Figure A1: Examples of load shedding schedules

(a) Westdene, 1 August 2023

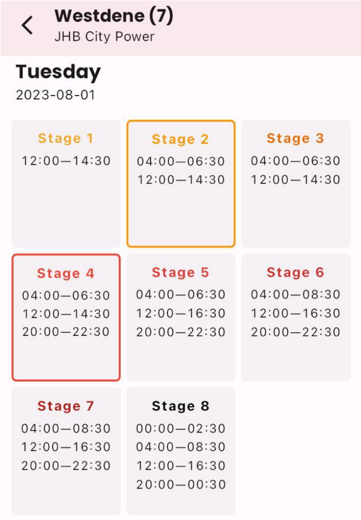

(b) Sophiatown, 1 August 2023

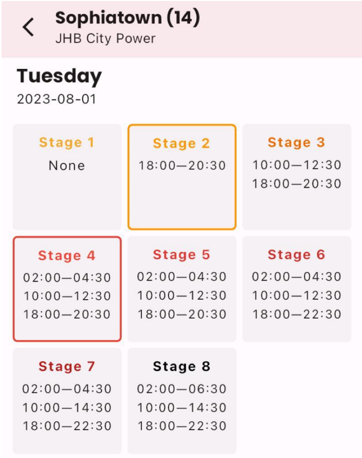

(c) Westdene, 2 August 2023

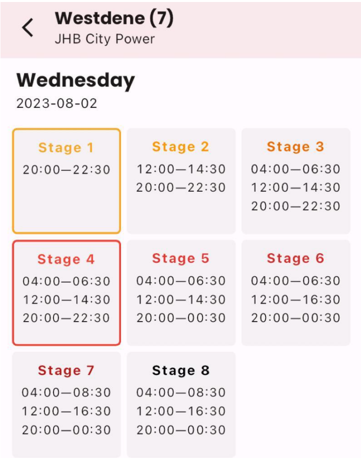

(d) Sophiatown, 2 August 2023

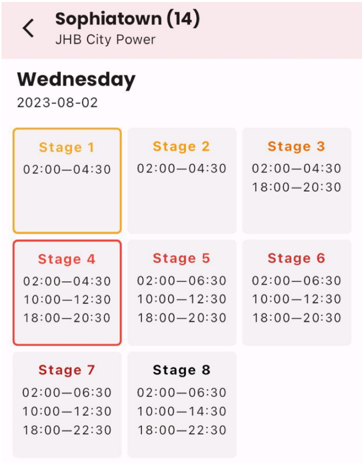

Note: examples of load shedding schedules from the EskomSePush (ESP) cellphone app, for adjacent neighbourhoods in the City of Johannesburg: Westdene and Sophiatown. Westdene falls under Zone 7 and Sophiatown under Zone 14 of Johannesburg City Power's broader (non-contiguous) load shedding 'blocks'. The highlighted sections refer to those stages/outages being relevant for the user on the particular day: on 1 August 2023, stage 2 was in force 05:00–16:00 and stage 4 00:00–05:00 and 16:00–24:00; on 2 August 2023, stage 1 was in force 05:00–16:00 and stage 4 00:00–05:00 and 16:00–24:00.

Source: author's screenshots from EskomSePush.

Figure A2: City of Cape Town-supplied vs Eskom-supplied regions in Cape Town

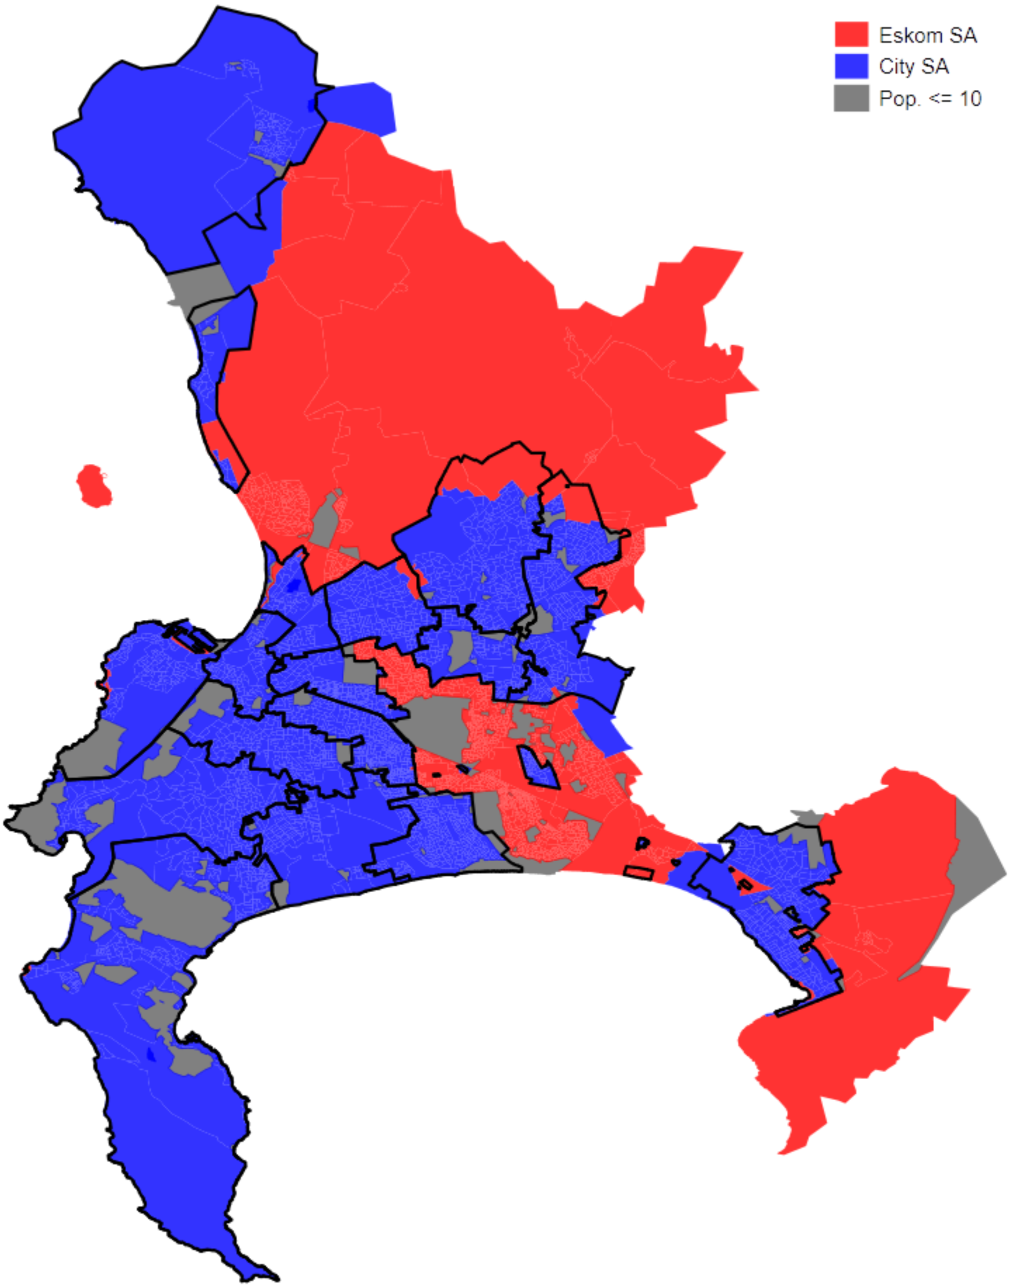

Note: map showing City of Cape Town Small Areas (SAs) supplied by the City (blue) versus directly by Eskom (red). Grey areas have a population of less than 11 and data is suppressed by Stats SA to preserve anonymity. Thick black outlines show the borders of City-supplied load shedding blocks. SAs are assigned to City- or Eskom-supply according to whether their centroid falls within a City-supplied load shedding block. A (very) few SAs are incorrectly assigned along borders due to imprecision when merging the SA data from Stats SA with the load shedding block data from the City of Cape Town.

Source: author's compilation.

Figure A3: Stacked event-study, deaths per 10,000 population, age $\geq$65

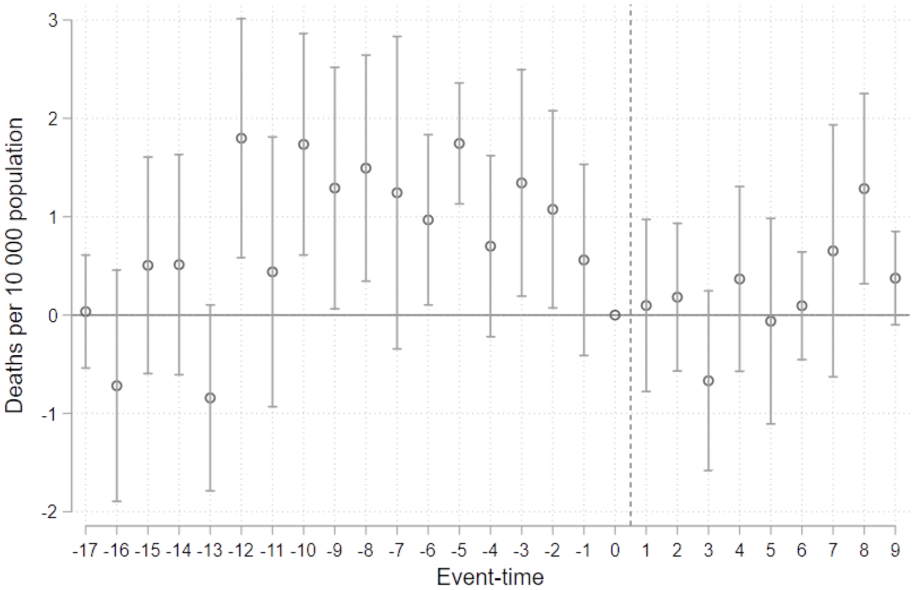

Note: the figure shows a stacked event-study where Cape Town is the treated unit. The outcome is deaths per 10,000 population and the sample is restricted to those aged 65 or older. The treatment begins in period 1, the period immediately after the dashed vertical line. Standard errors are clustered at the event–DC level.

Source: author's compilation.

Figure A4: Bias-corrected treatment effects with placebos, deaths per 10,000 population, various ages

(a) All ages

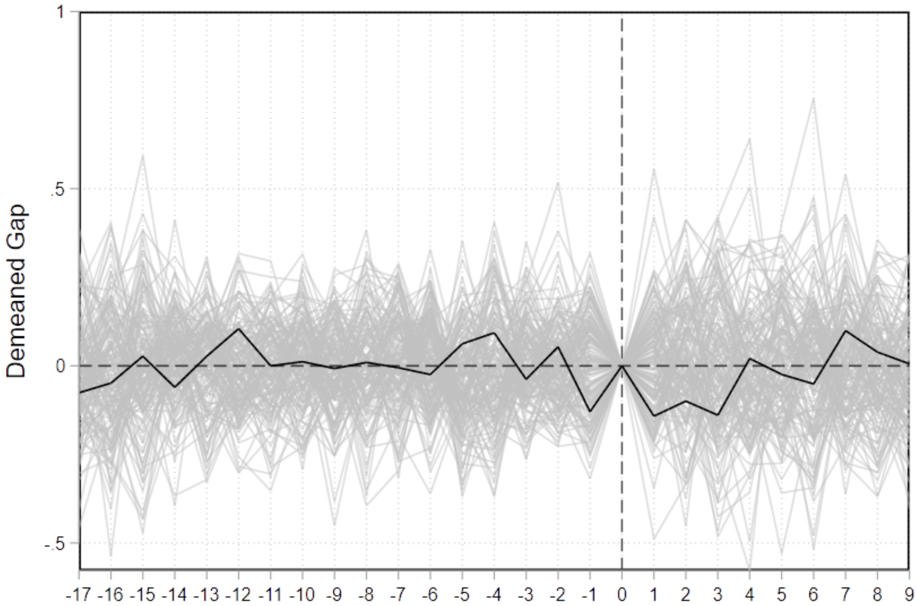

(b) Age $\geq$5

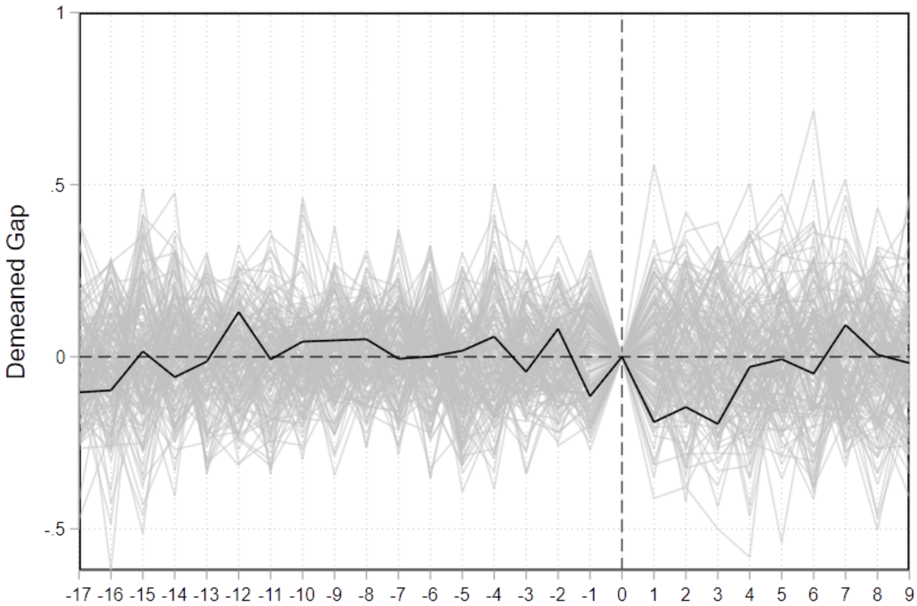

Note: the figure shows synthetic control gaps between Cape Town and synthetic Cape Town for various age groups, averaged across the three events (dark lines), with placebo gaps from in-place permutation tests (light grey lines). The outcome is deaths per 10,000 population and bias-corrected treatment effects are shown. Panel (a) shows results when no age restriction is applied, while Panel (b) shows results when the sample is restricted to those aged 5 and older. Treatment begins in period 1, the period immediately after the dashed vertical line.

Source: author's compilation.

Figure A5: Bias-corrected treatment effects, natural log of deaths, age $\geq$65

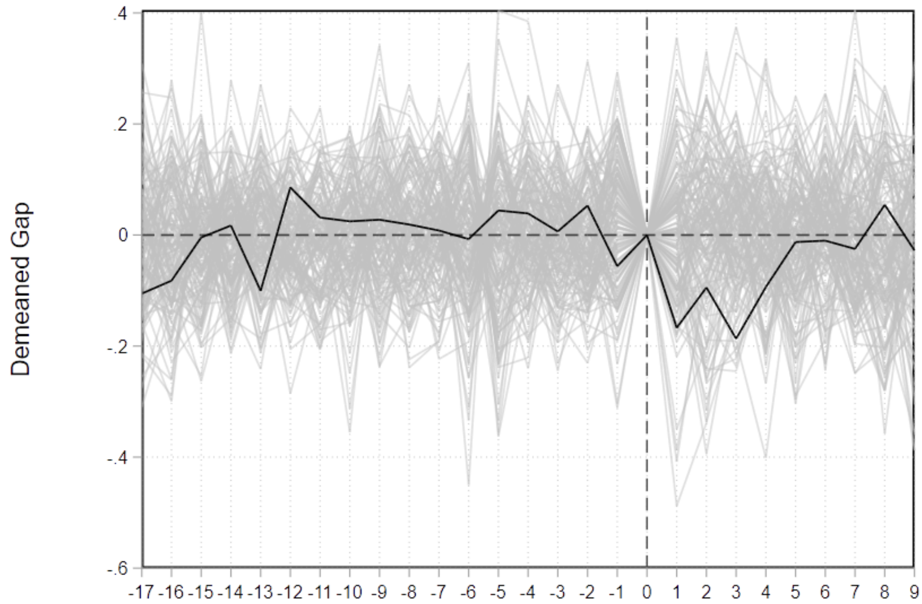

Note: the figure shows synthetic control gaps between Cape Town and synthetic Cape Town for those aged 65 and older, averaged across the three events (dark lines), with placebo gaps from in-place permutation tests (light grey lines). The outcome is the natural logarithm of deaths and bias-corrected treatment effects are shown. Treatment begins in period 1, the period immediately after the dashed vertical line.

Source: author's compilation.

Figure A6: Bias-corrected treatment effects, natural deaths per 10,000 population, age $\geq$65

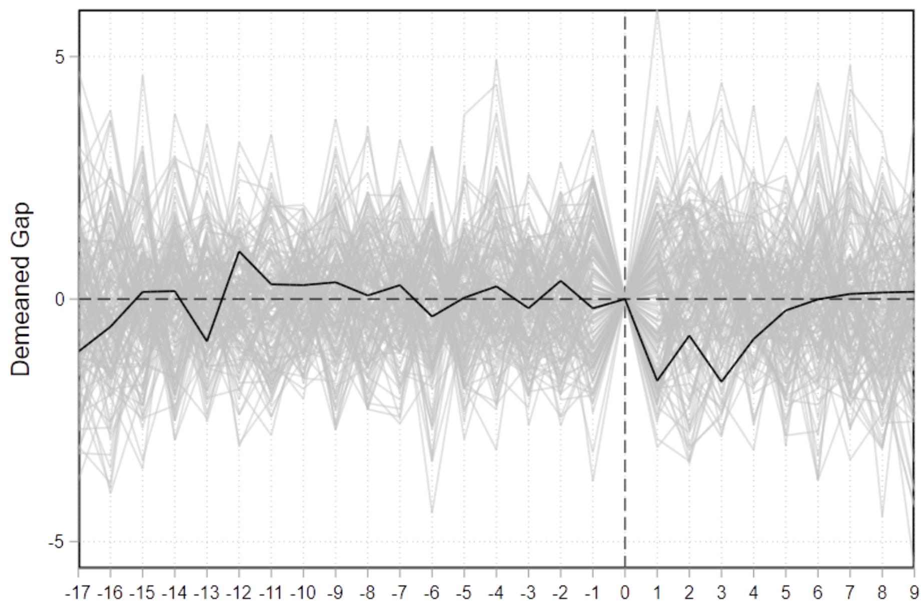

Note: the figure shows synthetic control gaps between Cape Town and synthetic Cape Town for those aged 65 and older, averaged across the three events (dark lines), with placebo gaps from in-place permutation tests (light grey lines). The outcome is natural deaths per 10,000 population and bias-corrected treatment effects are shown. Treatment begins in period 1, the period immediately after the dashed vertical line.

Source: author's compilation.

Figure A7: Bias-corrected treatment effects with 2014 event-time placebo, deaths per 10,000 population, age $\geq$65

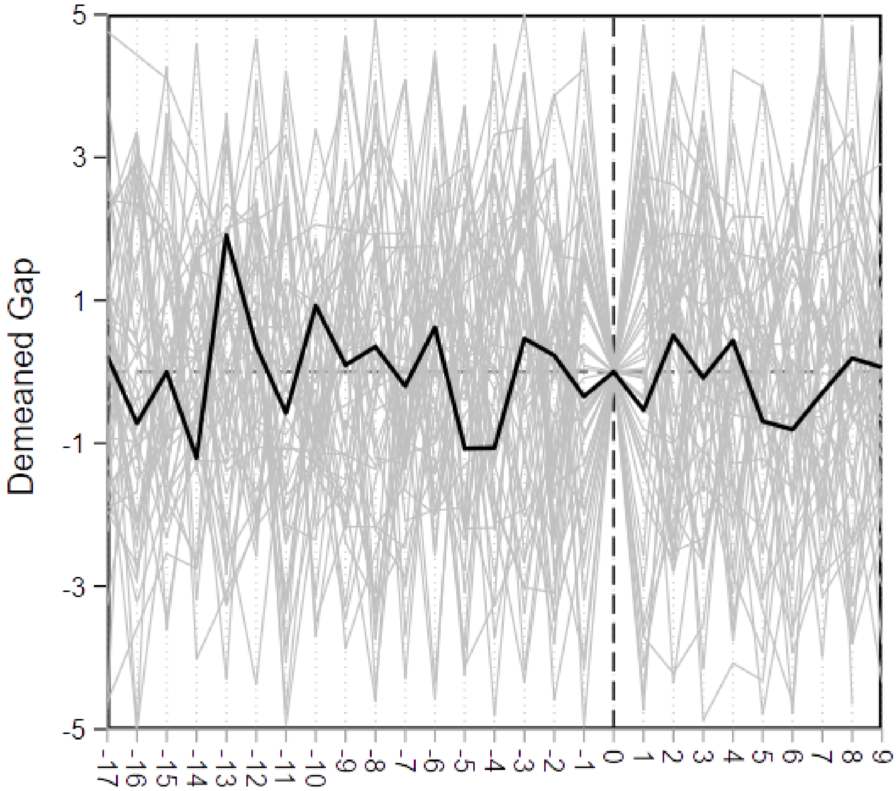

Note: the figure shows synthetic control gaps between Cape Town and synthetic Cape Town (solid dark line) for an in-time placebo test where common load shedding starts in period 1 but there is practically no differential load shedding between Cape Town and national. Treatment effects are averaged across the two events that meet these criteria. Equivalently specified placebo gaps from in-place permutation tests are also shown (light grey lines). The outcome is deaths per 10,000 population, the sample is restricted to those aged 65 or older, and bias-corrected estimates are shown. The placebo treatment begins in period 1, the period immediately after the dashed vertical line.

Source: author's compilation.

Figure A8: Bias-corrected treatment effects with constant post-period confidence interval, deaths per 10,000 population, age $\geq$65

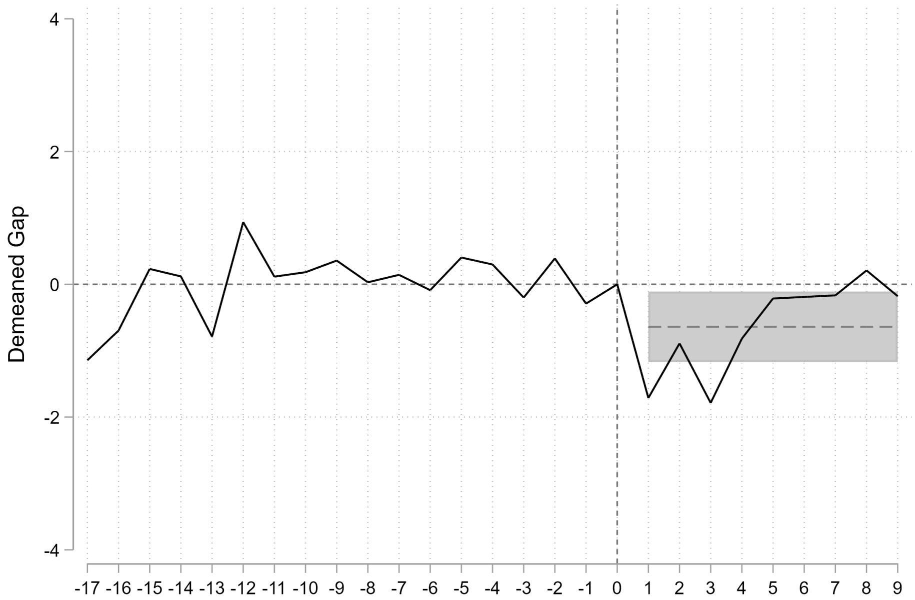

Note: the figure shows the synthetic control gap between Cape Town and synthetic Cape Town for those aged 65 and older, averaged across the three events (solid dark line), with average post-period effect (dashed grey line) and associated 95% confidence interval (light grey area). The outcome is deaths per 10,000 population and bias-corrected treatment effects are shown. Treatment begins in period 1, the period immediately after the dashed vertical line.

Source: author's compilation.

Figure A9: Proportion of population living in city-supplied (vs Eskom-supplied) regions in Cape Town

(a) Full population

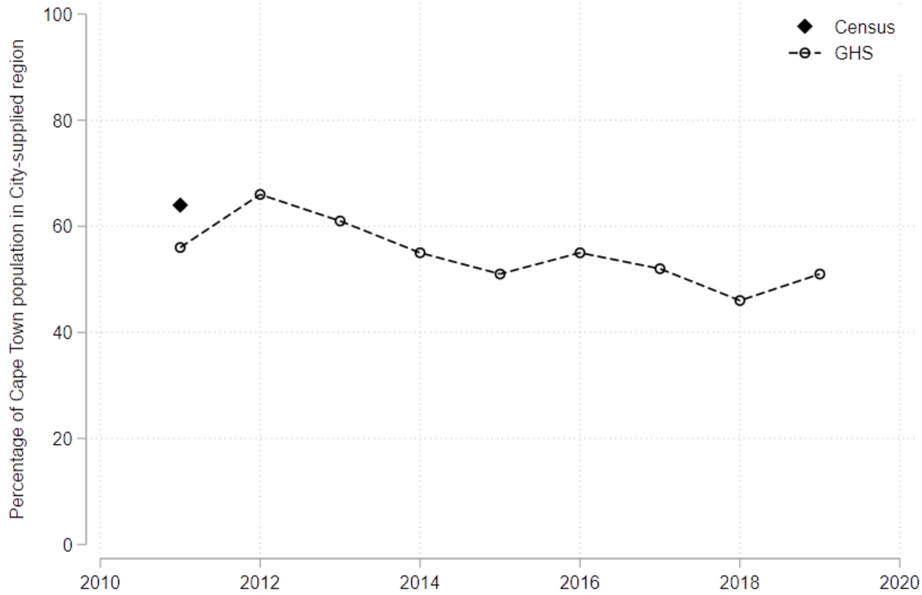

(b) By racial group

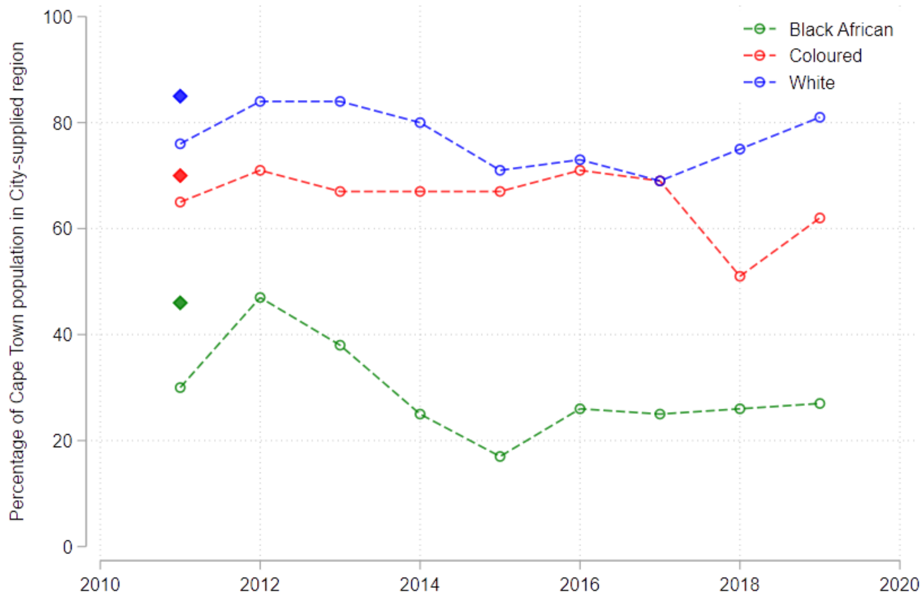

Note: the figure shows the proportion of the population living in City-supplied (vs Eskom-supplied) regions in Cape Town, using different data sources. Estimates from the 2011 Census are shown with solid diamonds, while estimates from the yearly GHS are shown with connected hollow circles. The census estimates are derived by matching census Small Areas to geographic coordinates of City of Cape Town load shedding blocks. The GHS estimates are derived from the self-reported 'Supplier of electricity' question in the GHS survey, and use the GHS household weights. Panel (a) shows the proportion of the total Cape Town population supplied by the City. Panel (b) shows the same proportion but separately by Stats SA racial groups. The 'Indian/Asian' group is not shown due to very small sample sizes.

Source: author's compilation.

Table A1: Synthetic control weights, main specification

DC/Metro municipality	Event 1	Event 2	Event 3	Average
West Coast	0.034	0	0.129	0.054
Cape Winelands	0	0.242	0.319	0.187
Overberg	0	0.026	0.084	0.037
Eden	0.358	0.361	0	0.24
Central Karoo	0	0	0	0
Cacadu	0	0	0	0
Amathole	0	0	0	0
Chris Hani	0	0	0	0
Joe Gqabi	0	0	0	0
O.R.Tambo	0	0	0	0
Alfred Nzo	0	0	0	0
Buffalo City	0.125	0.114	0	0.08
Nelson Mandela Bay	0	0	0	0
Namakwa	0	0	0	0
Pixley ka Seme	0	0	0	0
Siyanda	0.091	0	0	0.03
Frances Baard	0	0	0	0
John Taolo Gaetsewe	0	0	0	0
Xhariep	0.017	0.045	0.033	0.032
Lejweleputswa	0	0	0	0
Thabo Mofutsanyane	0	0	0	0
Fezile Dabi	0	0	0	0
Mangaung	0	0	0.107	0.036
Ugu	0	0	0	0
uMgungundlovu	0	0	0	0
Uthukela	0	0	0	0
Umkhanyakude	0	0	0	0
Uthungulu	0	0	0	0
Sisonke	0	0	0	0
Umzinyathi	0	0	0	0
Amajuba	0	0.009	0	0.003
Zululand	0	0	0	0
iLembe	0	0	0	0
eThekwini	0.002	0	0	0.001
Bojanala	0	0	0	0
Ngaka Modiri Molema	0	0	0	0
Dr Ruth Segomotsi Mompati	0	0	0	0
Dr Kenneth Kaunda	0	0	0	0
Sedibeng	0	0	0.09	0.03
West Rand	0	0	0	0
Ekurhuleni	0.357	0.203	0.128	0.229
City of Johannesburg	0	0	0	0
City of Tshwane	0.016	0	0.11	0.042
Gert Sibande	0	0	0	0
Nkangala	0	0	0	0
Ehlanzeni	0	0	0	0
Mopani	0	0	0	0
Vhembe	0	0	0	0
Waterberg	0	0	0	0
Greater Sekhukhune	0	0	0	0

Note: synthetic control weights when outcome is deaths per capita and age $\geq$65.
Source: author's compilation.

Table A2: Bias-corrected treatment effects, deaths per 10,000 population

	All ages		Age $\geq$5	
Event-time	Gap	*p*-value	Gap	*p*-value
1	–0.142	0.0530	–0.190	0.0199
2	–0.100	0.0596	–0.146	0.0265
3	–0.139	0.0530	–0.195	0.0132
4	0.021	0.0861	–0.030	0.0199
5	–0.024	0.1457	–0.007	0.0265
6	–0.051	0.2119	–0.049	0.0530
7	0.099	0.1854	0.092	0.0530
8	0.039	0.2185	0.006	0.0795
9	0.006	0.2583	–0.018	0.1192

Note: the table shows (reduced-form) bias-corrected synthetic control treatment effects separately by age group for each post-period averaged across the three events. The *p*-values come from in-place permutation tests. Results for the population as a whole are shown in the left super-column while results for those aged 5 and older are shown in the right super-column.

Source: author's compilation.

Table A3: Bias-corrected treatment effects, natural deaths per 10,000 population

	Natural deaths	
Event-time	Gap	*p*-value
1	–1.685	0.0132
2	–0.752	0.0132
3	–1.706	0.0132
4	–0.823	0.0132
5	–0.234	0.0132
6	–0.006	0.0132
7	0.106	0.0132
8	0.135	0.0397
9	0.152	0.0397

Note: the table shows (reduced-form) bias-corrected synthetic control treatment effects when the outcome is natural deaths per 10,000 population, for each post-period averaged across the three events. The *p*-values come from in-place permutation tests. The sample is restricted to those aged 65 and older.

Source: author's compilation.