Equity and Adaptation to Wildfire Risk: Evidence from California Public Safety Power Shutoffs

Max Vilgalys*

April 17, 2023

Abstract

In the past decade, California's investor-owned electric utilities have begun implementing Public Safety Power Shutoffs (PSPS) as part of their effort to adapt to increasing risk of catastrophic wildfires. I examine the extent that these decisions are correlated with two measures of community vulnerability: health risk factors and socioeconomic status (SES). I first construct a dataset linking weather, vulnerability indices, and PSPS decisions for electric circuits in California's three largest investor-owned utilities. I show that PSPS is used more frequently in circuits with lower average SES among two of California's major utilities, and circuits with higher average health risk in one of the major utilities. To focus on utilities' decisions, rather than other sources of inequality that may place vulnerable communities in areas with higher wildfire risk, I repeat this analysis after controlling for population and weather variation. The results are qualitatively similar. I then examine two key factors in firms' shutoff decisions: ignitions along power lines and the magnitude of service interruption. After controlling for weather variation, I find that ignitions are more frequent in low-SES circuits and in lower health risk circuits for one utility. When observations are aggregated across all utilities in the sample, the probability of ignition and the magnitude of PSPS disruption are higher in low-SES circuits, and lower in circuits with high health risks.

1 Introduction

In the last decade, electric utilities in California have been forced to adapt to increasing risk of catastrophic wildfire. Climate change, forest management practices, and shifting wildland-urban interface have contributed to the most severe wildfire seasons in California's history. Electric utility infrastructure has sparked some of the costliest wildfires. Under California law, utilities are financially responsible for these damages. Utilities have already faced billions in dollars in fines, driving Pacific Gas and Electric (PG&E) to declare bankruptcy in 2019. To make their electric lines safer, utilities invest in managing vegetation, upgrading infrastructure, and moving lines underground. However, these improvements are relatively slow and wildfire risk can require utilities to respond quickly. In these cases, utilities are sometimes forced to de-energize power lines to avoid sparking wildfires.

*Eastern Research Group, Inc. max.vilgalys@erg.com

This paper focuses on this last-ditch effort to prevent wildfire, the Public Safety Power Shutoff (PSPS). In a PSPS, a utility preemptively de-energizes lines that are likely to spark large wildfires. Sunce first PSPS in 2013, over 5,000 circuits (small segments of the electric grid) have been de-energized and over 1 million customers impacted (Hill et al., 2020). I focus on shutoffs from 2013-2021 by the three largest investor-owned utilities in California: Southern California Edison (SCE), Pacific Gas and Electric (PG&E), and San Diego Gas and Electric (SDG&E). Over 99% of shutoffs during this time period were conducted by one of these utilities.

PSPS events are subject to strict regulation, which makes them a useful data source to examine the equitability of utilities' adaptation to wildfire risk overall. Because shutoffs are costly to impacted communities, utilities are forced to justify each shutoff decision. Utilities must demonstrate that they carefully weigh the costs and benefits of each de-energization event (CPUC, 2019), and publish reports on their decision-making process (SCE, 2021; PG&E, 2021; SDG&E, 2021a). Records of each PSPS event are then disclosed to the public. This provides a rare window into utilities' adaptation to wildfire risk, as information on other investments are classified to avoid risks to critical national infrastructure.

The main hypothesis this paper evaluates is whether PSPS events have been equitably distributed. While PSPS is necessary in the short run to adapt to rising wildfire risk, shutoffs could exacerbate inequalities if disadvantaged communities receive fewer benefits or bear more costs. I examine two dimensions of inequality that increase vulnerability during electricity outages: socioceconomic status (SES) and health risk. Low-SES communities may have limited resources to adapt to electricity failures, and those with health risks may experience complications from wildfire smoke or electricity outages. I use definitions of health risk and SES from CalEnviroScreen (August et al., 2021). I focus on the costs of PSPS events, as evaluating the benefit of PSPS depends on predicting the size and damages of potential wildfires. This is a notoriously challenging problem, even given modern machine learning techniques (Taylor et al., 2013; Xi et al., 2019; Jain et al., 2020). Many Californians benefit from reduced wildfire risk, and rural populations or those with health risk factors may benefit most from reducing wildfire smoke (D'Evelyn et al., 2022).

Correlation between PSPS decisions and vulnerability could be explained by differences in utilities' investments or by various weather, vegetation, and development conditions that impact wildfire risk. It's important to identify the source of any inequality, to help determine the appropriate response. While there are not publicly available data on utilities' infrastructure investments, it is possible to consider the role of some key factors such as population and weather. I therefore compare circuits that differ in average SES or health risk, with and without population and weather controls. Without population or weather controls, I can assess whether vulnerable populations experience different rates of shutoffs. With these controls, I can assess whether this difference is explained by key observable factors outside the utility's control. To conduct this analysis, I first construct a dataset linking vulnerability indices, weather, population, and PSPS records from 2013-2021.

This hypothesis relates to several literatures. First is a literature studying the environmental justice of wildfire risk. In early work to explore this topic, Niemi and Lee (2001) describe how poverty can increase wildfire incidence and damages and Ojerio (2008) shows that federal wildfire preparedness grants are concentrated in higher-SES communities. One strand of this literature focuses on comparing populations that live in high wildfire risk regions. Wigtil et al. (2016) document that places with higher wildfire potential generally have lower social vulnerability to wildfire risk. Wibbenmeyer and Robertson (2022) find higher average property value, older residents, and more white residents in places with high wildfire potential. Another strand focuses on the impacts and responses of wildfires. D'Evelyn et al. (2022) argue that the health effects of wildfire smoke disproportionately impact populations with limited adaptive capacity. Anderson, A. Plantinga, Wibbenmeyer, et al. (2020) study inequality in firefighting responses, and document preferential treatment to higher SES communities following salient wildfire events. A. J. Plantinga, Walsh, and Wibbenmeyer (2022) study the historical spread of fires and find that firefighting efforts prioritize high-value properties.

Within this literature, several recent studies have examined PSPS as a tool to combat wildfire risk. Guliasi (2021) gives an analysis of the political economy and history of the PSPS. Hill et al. (2020) examines potential health costs from PSPS, and Wong-Parodi (2020) surveys impacted California residents about attitudes towards PSPS events. Rhodes, Ntaimo, and Roald (2020) studies the PSPS as an optimization problem, and suggests improvements to current decision processes using a test case. This paper is the first, to my knowledge, to empirically study the equity of these shutoff decisions.

This hypothesis is also related to literature on measuring equity in adaptation to climate change. Among environmental advocates, there has long been a call to focus on equity in climate change adaptation (Smit and Pilifosova, 2003; Thomas and Twyman, 2005). In their report, IPCC (2022) identifies several settings where inequality and poverty have set "soft limits" on the ability of groups to adapt to climate change. Coggins et al. (2021) conducted a review of literature on equity in climate change adaptation and highlighted several examples of work assessing the equity of climate adaptation. Sheller and Leon (2016) use interviews to study how historical inequalities between Haiti and the Dominican Republic impacted government responses to similar environmental crises, and Satyal, Byskov, and Hyams (2021) use environmental justice theory to examine how systemic injustices facing an indigenous group in Uganda undermine adaptation planning. However, Coggins et al. (2021) ultimately conclude that more work is needed in this area, especially in empirical assessment of equity and justice. This paper addresses this gap by providing an empirical assessment of equity in PSPS decisions.

As a secondary hypothesis, I evaluate the extent that utilities' calculations of two factors in the shutoff decision contribute to any inequitable distribution of shutoff decisions. These two primary factors are the probability of ignition and the cost of an outage to the impacted communities. These factors are identified based on guidelines in utilities' published Wildfire Mitigation Plans (SCE, 2021; PG&E, 2021; SDG&E, 2021a).

The basic elements of these plans are similar across utilities, although each utility uses proprietary data and modeling resources. For each circuit, the utility weighs the expected value of damages from a wildfire against the expected cost of the PSPS outage to impacted communities. Teams of meteorologists, fire scientists, and data scientists predict regions where ignitions are likely to spread to large fires. They use data of line conditions from public weather reports, service crews, and private weather stations to identify lines that could spark wildfires. These experts form predictions using public and private information, relying on machine learning models and extensive simulations of wildfire behavior. If their predictions find that the likely costs of wildfire exceed the costs of shutting off power, they notify residents and de-energize the circuit. Power remains off until weather conditions are less severe and the utility inspects affected circuits for any debris or damage.

Publicly available data allow me to evaluate the association between vulnerability and both the probability of ignition and the the cost of an outage to the impacted communities. Each factor may positively or negatively associated with vulnerability. For example, the probability of ignition may be higher in more vulnerable communities if utilities underinvest in infrastructure, or may be lower if utilities act to minimize impacts to vulnerable populations. The cost of an outage are likely higher for vulnerable populations (given health risks and lower access to backup power), but the way that utilities calculate the cost of an outage could place less weight on these populations.

To find the probability of ignition, I use logistic regression with records of ignitions along circuits from 2013-2021. To find the cost of an outage to the community, I select two proxies based on wildfire management plans and post-event reports: the size of interruption (in customer minutes interrupted) and the number of customers impacted. In their regulatory filings, utilities state that they model the cost of declaring PSPS as linear in the expected size of interruption (SCE, 2021; PG&E, 2021; SDG&E, 2021a). However, the calculations I find from post-event reports indicate that cost is linear in number of customers (Valdberg, Tozer, and Kilberg, 2021). Without further insight into how utilities make decisions, I report results using both proxies.

This hypothesis is related to a literature on identifying bias in decision making, specifically in cases where agents make decisions relying on complex algorithms. There is a broad literature on studying discrimination in decision-making, dating back to at least Becker (1957). Lang and Kahn-Lang Spitzer (2020) and Mehrabi et al. (2021) provide reviews of economics and machine learning literature, respectively, on identifying bias in decision making. Recent examples examining bias in human decisions include an analysis of racial bias in healthcare decision rules (Obermeyer et al., 2019) and in pretrial appearance risk (Rambachan, 2021). Examples examining bias in algorithms include facial recognition software (Buolamwini and Gebru, 2018) and predicting risks from medical records data (Gianfrancesco et al., 2018; Parikh, Teeple, and Navathe, 2019). Like these studies, I examine decisions and look for evidence of unequal treatment after controlling for relevant, exogenous variation. This setting, where agents make algorithm-supported decisions, is less well-studied by the existing literature.

The remainder of this paper is structured as follows. Section 2 describes how the dataset is constructed, and provides summary statistics for several sources of that dataset. Section 3 provides background on modeling choices used to investigate the hypotheses, including institutional background on utilties' decisions. Section 4 gives the results of my analysis, and discusses their interpretation. Section 5 concludes.

2 Data

My analysis relies on a dataset with records of weather variation, vulnerability indices, ignitions along electric circuits, and shutoff decisions from 2013-2021. The unit of analysis is the electric circuit, a small unit of the electricity distribution network. PSPS decisions are generally made and recorded at this level. To construct this dataset, I merge administrative records of shutoffs and ignitions, gridded daily weather observations, and Census tract-level data of socioeconomic status and health vulnerability.

2.1 PSPS Events

Filings from firms to the CPUC provide a complete record of PSPS events. Firms are required to report statistics after each shutoff, so this dataset represents the universe of shutoffs between October 2013 and December 2021. The CPUC summarizes these reports and publishes a record of each shutoff. Each record includes the circuit targeted, the date and time of the shutoff, the duration of the outage, the number of customers impacted, and information on what types of customers are impacted. Table 1 summarizes these filings by year and firm.

In order to link these with other geospatial records, I use integrated capacity analysis (ICA) maps from each electric utility. ICA maps are circuit-level maps of the distribution infrastructure, although some circuit segments are not published due to privacy concerns. I am able to match over 98% of PSPS records to their corresponding geographic file. The ICA maps include 5,411 circuits; there are PSPS events recorded on 20.3% of these circuits. Figure 1a shows the location of these circuits.

PSPS events are generally reported at the circuit level. In some cases, a firm can conduct a sub-circuit level outage. As the circuit level is the most specific level in the ICA maps, I am unable to match sub-circuit level outages to geographic information or other datasets. I sum these outages to the circuit level to merge with the other datasets.

2.2 Fire Ignitions

To study the risk of igniting a fire, I use administrative records of fires ignited along utility lines from 2014 through 2021. Per CPUC guidelines, firms must report all fires to their knowledge larger than one meter (CPUC, 2014). This dataset includes 4,550 ignitions from the three firms I study. These filings are required to include the ignition location, but not the corresponding circuit segment. To match these to the circuit records, I find the closest circuit segment from the ICA maps to the ignition location. Figure 1b shows the location of these circuits.

2.3 Vulnerability

I use data from CalEnviroScreen to measure population vulnerability (August et al., 2021). The authors construct a Census tract-level database of health risk factors and socioeconomic status (SES) indicators. This database is primarily intended to assess environmental and energy justice in the state of California. I use these indexes, as well as tract-level population, in my analysis.

I construct an SES index and a health risk index to summarize health risk factors and SES

Utility	year	2013	2014	2017	2018	2019	2020	2021
SCE	Customers	–	–	–	–	196,879	235,879	117,690
	Million CMI	–	–	–	–	353	280	372
	# PSPS Events	–	–	–	–	246	1,501	122
PG&E	Customers	–	–	–	47,324	1,987,783	645,859	79,630
	Million CMI	–	–	–	89.8	6,670	1,560	174
	# PSPS Events	–	–	–	32	1,458	670	219
SDG&E	Customers	179	884	17,111	21,036	45,337	93,058	–
	Million CMI	0.0797	0.665	40.5	65.4	78.2	165	–
	# PSPS Events	3	6	51	38	218	110	–

Table 1: Number of PSPS events by firm, by year, and the number of customers impacted. CMI is Customer Minutes Impacted, the product of the minutes of shutoff and number of customers per circuit. Note that number of customers impacted is the sum of customer shutoffs experiences, but not the unique number of customers impacted.

indicators. Each index ranges from 0-100, with 100 being the most vulnerable and 0 being the least. The indices are constructed as the average of ranks of several factors, as in August et al. (2021). For socioeconomic vulnerability, this includes rate of high school non-attainment, rent-burdened low-income households, limited English proficiency, living below twice the federal poverty line, and share unemployed. For the health risk index, this includes asthma incidence, cardiovascular disease incidence, and rate of low birth weight infants.

To match these records to circuits, I take the average of values from each census tract that contains a given circuit segment. I weight these averages by the length of the circuit in each census tract. I am able to match records for 5,000 out of 5,411 circuits, and for 1,071 of the 1,103 circuits with a PSPS event.

Figure 2 plots these scores per circuit against the total number of PSPS events (among circuits with at least one event), the total number of recorded ignitions (among circuits with at least one ignition), and the total customer minutes interrupted (among circuits with at least one event). Each plot also includes the best-fitting line to these observations, to help summarize the trend among these scatter plots.

2.4 Weather data

For weather observations, I use the GridMET weather dataset from Abatzoglou (2013) and an archive of areas with a red flag warning. GridMET was designed to support applications in modeling wildfire risk, and includes a rich set of relevant weather variables. GridMET includes primary variables, constructed via satellite- and geography-guided interpolation from weather stations, and variables derived from these primary observations.

Primary variables are specific humidity, precipitation, minimum relative humidity, maximum relative humidity, surface downwelling shortwave flux in air (a measure of solar radiation), minimum air temperature, maximum air temperature, wind speed, and wind direction. Derived variables are expected to be relevant for predicting wildfire risk: burning index, energy release component,

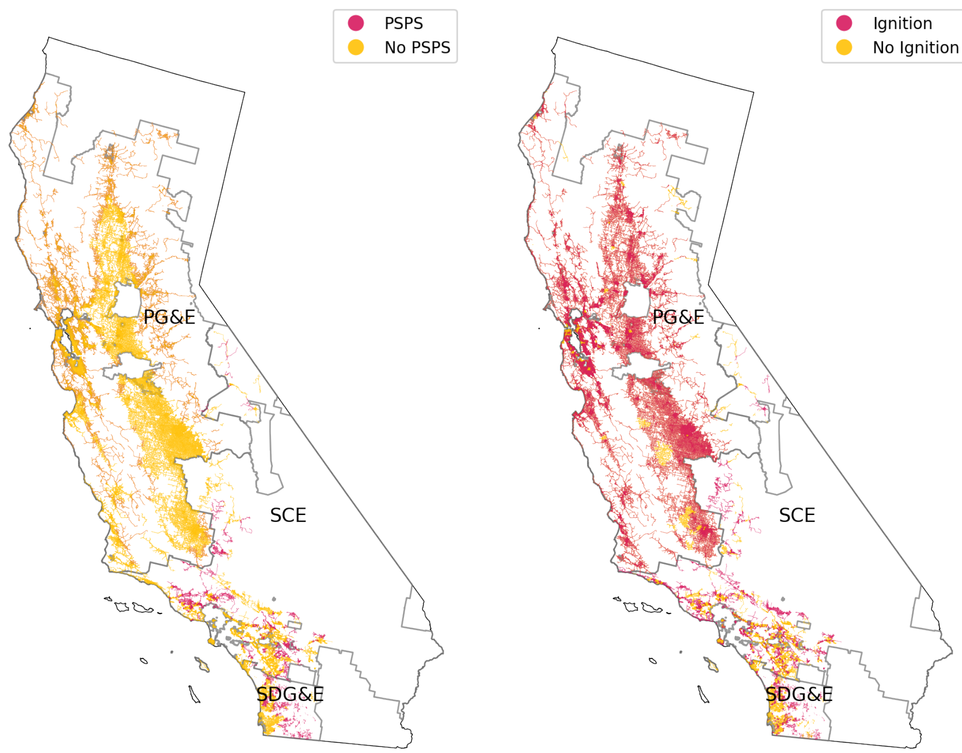

(a) PSPS Events

(b) Ignitions

Figure 1: Map of which circuits had PSPS events (a) and ignitions (b) from my data. State borders and boundaries between electric utilities are shown.

potential evapotranspiration, reference evapotranspiration, dead fuel moisture at 100 hours, and dead fuel moisture at 1000 hours. See Abatzoglou (2013) for more details on the development of this dataset. Each variable is reported daily at a 4 km resolution across the United States. To merge GridMET records with my dataset, I take the average of weather records in each grid cell a circuit passes through, weighted by the length of the circuit in each grid cell.

In addition to weather variation, firms use Red Flag Warnings from the National Weather Service to make PSPS decisions. Over 98% of shutoffs occur during a Red Flag Warning, a period when the National Weather Service has identified weather conditions that could sustain catastrophic fires. I include an indicator of whether a Red Flag Warning was in effect in any part of a circuit by merging a historical archive of Red Flag Warning shapefiles.[1]

Summary statistics of the weather data are given in Appendix A. The trends in weather data show that conditions are becoming drier, hotter, and more conducive to wildfire, consistent with a literature studying the consequences of climate change for California (Westerling and Bryant, 2008). Relative to the full sample, Red Flag Warnings are drier, hotter, and more elevated, and PSPS events occur in windier and drier conditions and at higher elevation.

[1] From `https://mesonet.agron.iastate.edu/info/datasets/vtec.html`, accessed 9 December 2021.

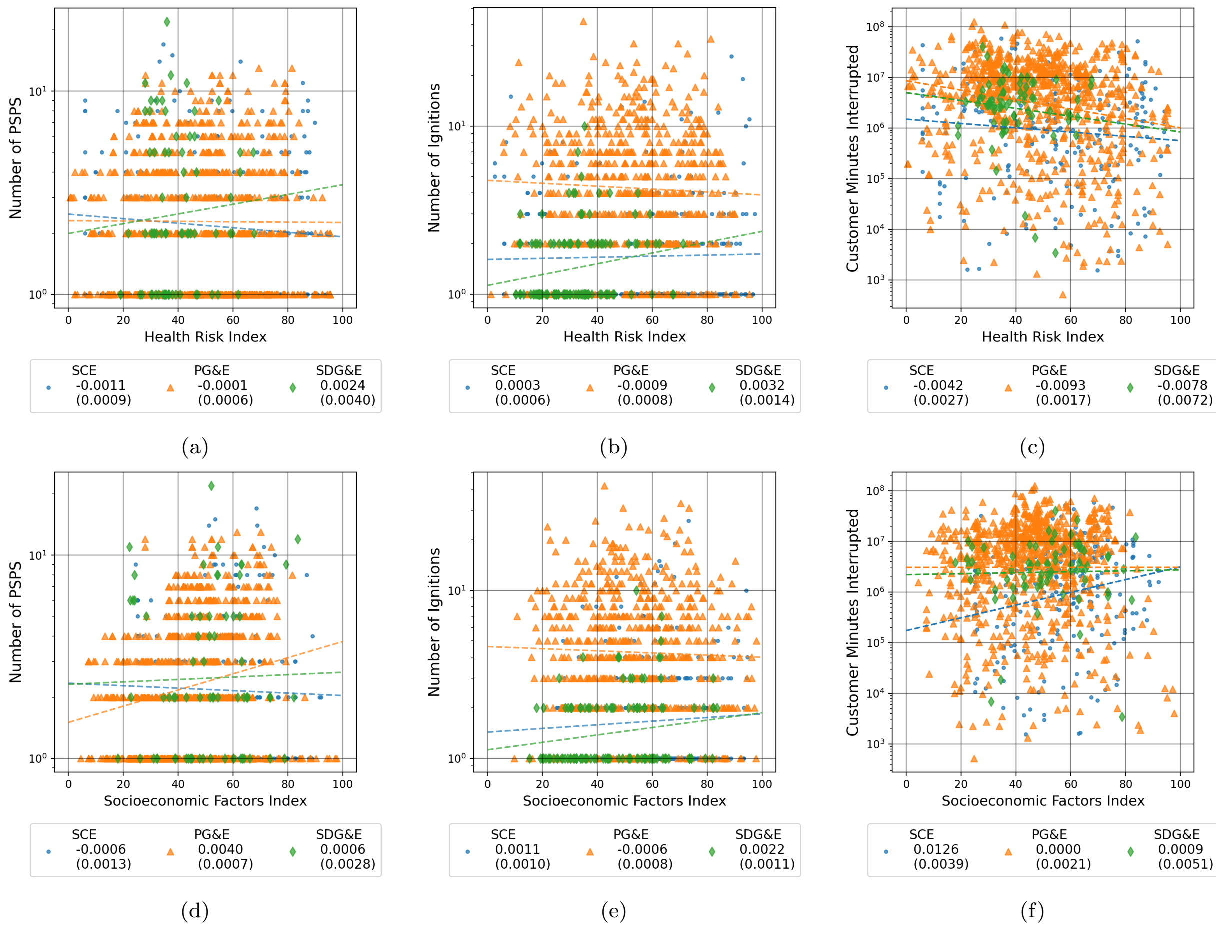

Figure 2: Scatter plots showing vulnerability indices and various outcomes, for circuits with nonzero values of PSPS events or ignitions. Each plot includes a best-fitting line for the observations. The coefficient on the index and and the standard error (in parentheses) of each line are reported in the legend.

3 Methods

In this section, I describe regression models based on utilities' wildfire management plans that allow me to measure the association between vulnerability indices and factors in utilities' decisions. I estimate separate regressions for three components: PSPS shutoffs, ignitions along lines, and the cost of an outage to the impacted community.

Per their filings to the CPUC, utilities initiate a PSPS if the expected degree of damages (that is, the product of expected damages conditional on ignition and the probability of ignition) exceeds the cost of failing to provide power.[2] The shutoff decision is a binary choice model, where the utility weighs the expected damages from a wildfire ("Wildfire Risk") against the utility's cost of failing to provide power ("PSPS Risk"). Utilities use separate prediction problems for probability of an ignition and size of fire conditional on ignition (PG&E, 2021; SCE, 2021; SDG&E, 2021a). This approach is common in both classical statistical (Xi et al., 2019) and machine learning (Jain

[2]Legislation requires that utilities making shutoff decisions must quantify benefits and risks of de-energization events, and document "how the power disruptions to customers, residents, and the general public is weighed against the benefits of a proactive de-energization."

et al., 2020) approaches to predicting wildfire size.

While I have data (or proxies) to determine the probability of shutoffs, the probability of ignition, and the cost of failing to provide power, I am not able to estimate the damages from a fire conditional on ignition. Damages from a wildfire are a function of wildfire size and the features of land damaged by the wildfire. Both elements may be related to vulnerability factors, and the association could be positive or negative. For example, lower-SES areas have lower property values (decreasing the estimated damages from a fire) and may have lower firefighting capacity (increasing the size of a potential wildfire). Predicting fire size is a notoriously challenging problem, even given modern machine learning techniques (Taylor et al., 2013; Xi et al., 2019; Jain et al., 2020). Utilities use proprietary software to make these fire size predictions, with relatively high levels of accuracy. In Appendix B, I document an attempt to predict fire size using linear regression and random forest regressions. With publicly available data, I am unable to provide informative bounds on the degree of fire size.

I express all regressions as a generalized linear model, to explain choices of fixed effects and explanatory variables that apply to each regression. For a circuit i in utility service area U in time period t, X_i is the set of vulnerability indices, Z_{it} are additional controls, and $\alpha_{U,t}$ is a utility-by-time period fixed effect. y_{it} is the outcome variable, which varies depending on the regression. The additional controls vary by model specification, and may include weather variation, elevation, fire risk scores, and population. Then the general model is:

$$g(y_{it}) = \alpha_{U,t} + \beta X_i + \gamma_{U,t} Z_{it} \tag{1}$$

Where $g(\cdot)$ is the link function that determines the form of regression in a generalized linear model. For linear regression (for cost of an outage) this is the identity function, and for logistic regression (for shutoff decisions and ignition) this is the logit link function. Note that there is a separate coefficient $\gamma_{U,t}$ for each utility, for each time period. This choice captures the fact that each utility develops their own shutoff decision rules each year.

This regression does not include interactions between terms, implying the assumption that components are additively separable. This restriction rules out associations where the level of one explanatory factor changes the effect of another explanatory factor on the outcome variable.

The main coefficient of interest is β, which captures the association between the outcome variable and the vulnerability indices. Vulnerability indices are the SES index and health risk index, as described in Section 2.3. I estimate two specifications of the outcome variables: one with utility-level β coefficients, and one where I report aggregate estimates across all utilities.

I consider four different sets of variables for Z_{it}: no controls, only population, population with primary weather variables, and population with all weather variables. Primary weather variables are those that are directly observed, such as temperature and humidity. All weather variables also includes variables that were derived from other variables to represent wildfire hazard, such as burning index or energy release component. In Section 2.4, I list each of the eleven primary weather variables and the six derived weather variables. Appendix A gives summary statistics.

I choose not to include vegetation in my set of explanatory variables. By pruning tree limbs or other vegetation along lines, utilities can influence the level of vegetation. Vegetation is therefore considered a bad control, and should be omitted even though it may influence the probability of ignition or shutoff (Cinelli, Forney, and Pearl, 2022).

In the remainder of this section, I describe the model specifications for each outcome variable I study.

3.1 PSPS Probability

To estimate association between PSPS shutoffs and vulnerability, I use logistic regression with PSPS shutoffs as the outcome variable. I use a subset of data during red flag warnings from October 2013 (the month of the first PSPS event) onward. I limit the sample to red flag warnings because these are widely used indicators of fire hazard. All utilities mention using red flag warnings as part of their process for determining shutoffs (PG&E, 2021; SCE, 2021; SDG&E, 2021a). Over 90% of PSPS events are declared during a red flag warning.

In Appendix D, I show results from a robustness exercise using the full sample. Generally, the estimates are similar to those using only the subset with red flag warnings. Magnitudes of most estimates are closer to zero.

3.2 Ignition Probability

To estimate the probability of ignition, I use logistic regression of ignitions along power lines. I do so using a subset of data from years where utilities do not use PSPS. Table 1 shows the years with PSPS observations. I use data from all firms in 2015 and 2016 and from a subset of firms in 2014, 2017, 2018, and 2021. In years where utilities are able to use PSPS, the researcher does not observe whether an ignition would have occurred without PSPS. Because the utility can choose when to use PSPS and prevent the researcher from observing a potential ignition, it is only possible to partially identify regression functions (Khan and Tamer, 2009). Subsets from years when firms do not use PSPS are not subject to this censoring concern. I assume that the relationship between weather variation and ignition probability is consistent between years when utilities do and do not use PSPS; this could be violated if dry vegetation accumulates and fire risk increases over time, or utilities choose to use other wildfire management strategies in years without PSPS.

In Appendix E, I show results from a robustness exercise and estimate ignition probabilities using the full sample. I include results using the full time period, with two different assumptions on the observed ignitions and shutoffs: that, absent a shutoff, each circuit with a PSPS event would either have an ignition or no ignition. Overall, these results suggest that my conclusions in the main analysis are robust to including ignition data from the full sample.

3.3 Cost of PSPS

Due to ambiguity between various documents, I both the number of customers impacted and the customer minutes interrupted (CMI) as proxies to the firm's cost of PSPS. These proxies capture major sources of variance in the firm's expected costs of PSPS. While neither is a perfect approximation, they provide a reasonable estimate of how vulnerability indices influence the firm's cost of PSPS. I estimate linear regression of the log of each outcome variable. I use logs because firms state that their costs are multiplicative in CMI or customers impacted.

	(1)	(2)	(3)	(4)	(5)	(6)
Sens Index	53.28	50.01	50.16	48.21	39.65	39.19
	(1.610)	(1.160)	(1.010)	(0.439)	(1.130)	(0.768)
SES Index	54.18	57.34	49.01	48.65	50.35	51.91
	(1.149)	(0.712)	(0.758)	(0.349)	(2.018)	(1.193)
Observations	223	469	451	1885	81	181
Utility	SCE	SCE	PGE	PGE	SDGE	SDGE
$\geq$ 24 Hours		X		X		X

mean coefficients; se in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Summary of vulnerability indices (SES index and health risk index), by whether the observed outage exceeds 24 hours. Standard error of the mean is in parentheses.

In firms' Wildfire Mitigation Plans, they provide formulas to calculate the cost of a PSPS as a simple function of customer minutes interrupted (CMI) and the total number of customers interrupted (SCE 2021, p. 61; PG&E 2021, p. 52; SDG&E 2021a, p. 26). Firms also incorporate the safety cost and financial cost of PSPS, as well as a reliability score. This safety cost is calculated as a constant factor multiplied by CMI, and the financial cost scales with the cost of shutoff (SCE, 2022; PG&E, 2020; SDG&E, 2021b). PG&E incorporates a scaling function if the safety, reliability, or financial costs of PSPS in a circuit exceed 10% of the largest recorded wildfire damages; I neglect this nonlinearity for simplicity. SDG&E plans to incorporate the health sensitivity of subpopulations, but I do not observe decisions made with these rules (SDG&E, 2021a, p. 30). In 2021, SCE began weighting some components of its cost function by the number of vulnerable customers per line; I do not have access to their conversion formula and do not attempt to model this improvement. By taking the regression with logs of CMI or number of customers impacted, this conversion factor is absorbed into the constant during linear regression and does not impact my estimates.

SCE is the only firm to specify how they form ex-ante predictions of the CMI. In their post-event reports, SCE calculates their CMI as a constant number of minutes multiplied by number of customers impacted, effectively making the cost of a shutoff a function of function solely of the number of customers (Valdberg, Tozer, and Kilberg, 2021, p. 16). No other firms publish their ex-ante PSPS cost calculations. I assume that they either use a constant factor, or the expected CMI per outage based on the empirical duration of PSPS outages.

If firms use a constant outage duration to estimate costs, this approximation may systematically undervalue the cost to low SES or high health risk communities. The number of customers impacted and CMI for a given outage are stochastic; depending on weather conditions, firms may be able to de-energize a smaller section of the circuit or be forced to prolong the outage. To inspect this, I compare the average health risk and SES indices for circuits with PSPS outages above and below 24 hours. Table 2 shows these summary statistics. For SCE, outages over 24 hours occur in circuits with significantly higher SES index (indicating lower-SES circuits), and significantly lower health risk index. This shows that SCE's stated decision systematically undervalues the cost of an outage to low-SES populations.

4 Results

I evaluate the association between vulnerability factors and the rates of PSPS shutoff by using logistic regression. I then evaluate the association between vulnerability factors and the probability of ignition using logistic regression, and the costs of an outage using linear regression. I report both separate results for each utility, and results from aggregating observations from each utility. The results from separate utilities are my preferred estimates, as this more closely models the setting where each firm relies on their own data and methods. All regressions include utility-by-year fixed effects, and all population or weather control variables are interacted with these fixed effects. I use four specifications: no controls (beyond fixed effects), only population, primary weather variables plus population, and all weather variables (including calculations of wildfire hazard) plus population.

The dependent variables of interest are the health risk index and socioeconomic factor index from CalEnviroScreen. In these indices, 0 is the least vulnerable and 100 is the most vulnerable. Increasing the socioeconomic or health risk index by 1 is equivalent to an average increase of 1 percentage point across the ranks of the sub-indices. A positive coefficient indicates that higher health risk (lower SES) circuits have a greater rate of PSPS shutoffs or ignition, or that the magnitude of costs from the outage is larger. For logistic regression, these coefficients represent the amount that log-odds change with an increase of 1 unit of the index. They can be approximately interpreted as the percentage change in likelihood given an increase in 1 unit of the index, as the coefficients are fairly close to 0. For example, an estimated coefficient of 0.01 indicates that PSPS events or ignitions are 1% more likely in circuits with 1 higher index. I refer to circuits where the population has a lower (higher) average health risk index as lower (higher) health risk circuits, and circuits where the population has a lower (higher) average SES index as lower (higher) SES circuits.

Figure 3 visualizes estimates of the association between vulnerability indices and log-odds of a PSPS shutoff. Figure 4 shows estimates for the association between vulnerability and three factors that contribute to the shutoff decision: ignition, CMI impacted, and number of customers impacted. Tables reporting the estimates are provided in Appendix C. Results from some alternate specifications are provided in Appendix D (for PSPS probability) and Appendix E (for ignition probability).

4.1 PSPS Probability

As described in Section 3.1, I use a subset of data during red flag warnings from October 2013 (the month of the first PSPS event) onward to estimate how vulnerability indices are associated with the probability of PSPS. Appendix D shows the results using the full sample; for firm-level estimates, the patterns are similar but the magnitude of effects is lower.

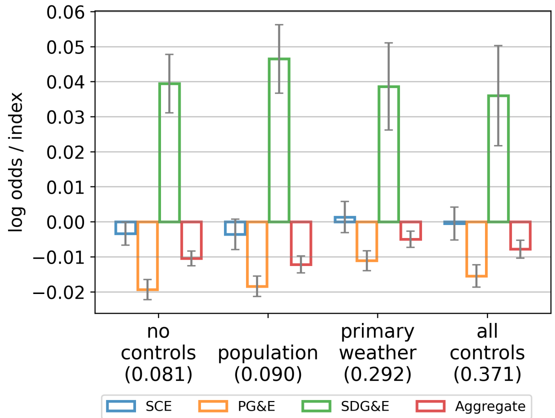

(a) PSPS, Health Risk Index

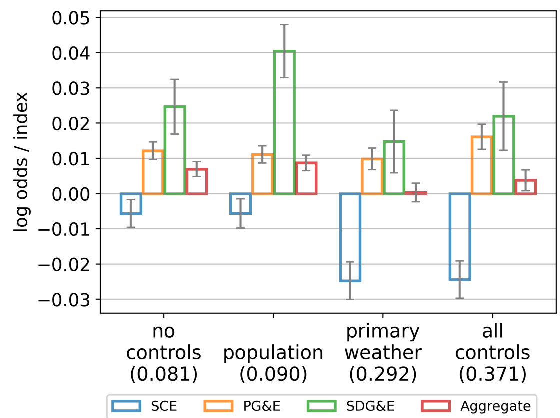

(b) PSPS, SES Index

Figure 3: Summarizing coefficients from logistic regression of PSPS decisions. Groups on the x axis collect results from regression with a given set of controls. Pseudo R squared values are shown in parentheses below each collection. Error bars show the 95% confidence interval. Each group of plots is ordered SCE, PG&E, SDG&E, Aggregate.

Without controlling for weather factors (column 1), higher health risk circuits are significantly less likely to have a PSPS in both SCE and PG&E, and more likely in SDG&E. This finding is significant at the $p = 0.001$ level for PG&E and SDG&E, and at the $p = 0.05$ level for SCE. Lower SES circuits are more likely to have a shutoff in PG&E and SDG&E, and less likely in SCE; this finding is significant at the $p = 0.001$ level for PG&E and SDG&E, and at the $p = 0.01$ level for SCE. These magnitudes are on the order of 0.01, so a 1 point increase in the index corresponds to roughly one percent difference in the likelihood of PSPS.

After controlling for population and weather variation (columns 2-4), model fit (shown by the pseudo-R squared in parentheses below each group of columns) improves but these patterns remain largely consistent. McFadden (1973) suggests that a pseudo-R squared of 0.2-0.4 suggests good model fit for logistic regression, indicating that this model acceptably fits the PSPS decisions after controlling for weather variation. The exception is that the coefficient on the health risk index for SCE is no longer significant, but the coefficient on socioeconomic factors for SCE is larger in magnitude and is statistically significant at the $p = 0.001$ level.

When using the aggregate sample, I find that there is a significant (at the $p = 0.001$ level) negative correlation between the health risk index and PSPS, but that the correlation between SES index and PSPS is only positively significant at the $p = 0.05$ level. As the utility-level estimates show, there is significant heterogeneity in these associations between firms. These aggregate values, which present an average experience of PSPS rates across all utilities, are therefore not representative of the experience of PSPS shutoff rates in any service area.

I do not observe the full set of relevant variation that firms have while making these decisions, and therefore these estimates may be susceptible to omitted variable bias. In linear models, Oster (2019) gives an approach to quantify the degree of omitted variable bias by comparing the stability of coefficients as the model fit improves. I am not aware of an analogous approach for logistic regression. Informally, the sign and magnitude of coefficients remain relatively stable as the model fit improves while including population and weather controls, indicating that these conclusions may be robust to incorporating additional variables.

4.2 Ignition Probability

As described in Section 3.2, I use data from years where utilities did not conduct PSPS to estimate the probability of ignition. This avoids the identification concern that when a firm conducts a PSPS, I do not observe whether an ignition would have occurred without that intervention. In Appendix E, I conduct a robustness exercise using the full set of data. I evaluate the coefficients assuming that each PSPS event would be an ignition, or that no PSPS event would be an ignition. The conclusions below still hold in both alternate specifications, suggesting that my findings hold regardless of any changes in the relationship between weather and ignition probability over time.

The coefficient estimates from this regression are shown in Figure 4a and Figure 4d. Many of the utility-level coefficient estimates are statistically indistinguishable from 0. In PG&E circuits, higher health risk circuits are significantly (at $p = 0.001$ level) less likely to have an ignition and lower SES circuits are more likely to have an ignition. This finding is robust to including population and weather variables. At the $p = 0.01$ level, ignitions in SDG&E lines are positively correlated with higher vulnerability indices, although these relationships are not significant after controlling for weather factors.

In the aggregate, these coefficient estimates are significant (at $p = 0.001$ level). I find that ignitions are positively correlated with the SES index, and negatively correlated with the health risk index. This shows that both lower SES circuits and higher health risk circuits have higher rates of ignition, after controlling for weather and population differences.

Some patterns from the coefficient estimates are similar to those of the PSPS decisions, although less precisely estimated. Without controlling for weather variation, I find that lower SES circuits have higher rates of ignition in PG&E and SDG&E, and lower rates of ignition in SCE. I find that higher health risk circuits have higher rates of ignition in SDG&E, and lower rates in PG&E. Controlling for population and weather variation, the only significant associations that remain are that lower SES circuits in PG&E have higher rates of ignition and that higher health risk circuits in PG&E have lower rates of ignition. This is similar to the findings from the regression on PSPS results, although there is greater uncertainty. This suggests that population and weather variation are able to explain much of the observed differences in ignitions between more and less vulnerable communities. These coefficient results generally show that the same patterns of unequal treatment occur in ignitions and in PSPS decisions. Both observations could be explained by unequal conditions along electric distribution lines, although I do not have data on conditions along electric lines that would be necessary to evaluate that hypothesis.

4.3 Costs of Outages

I use two proxies to find how vulnerability indices correlate with the cost of an outage, as calculated by the utilities. This cost is the value the utility uses when weighing the costs and benefits of a shutoff; it reflects the estimated size of the disruption from declaring a PSPS event. As discussed

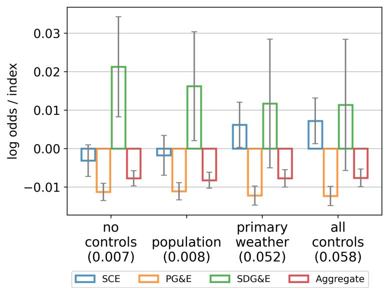

(a) Ignition, Health Risk Index

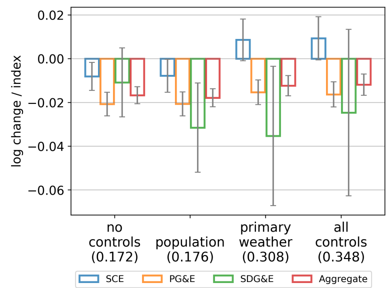

(b) CMI, Health Risk Index

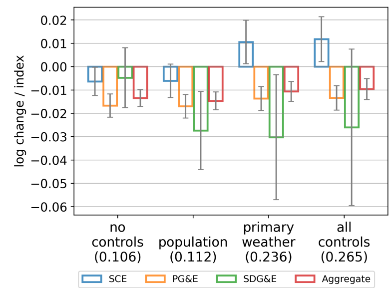

(c) Customers, Health Risk Index

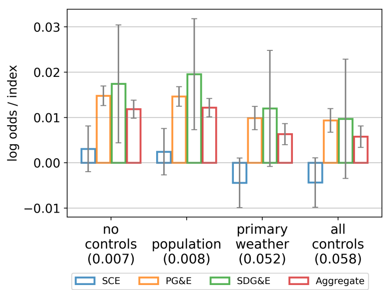

(d) Ignition, SES Index

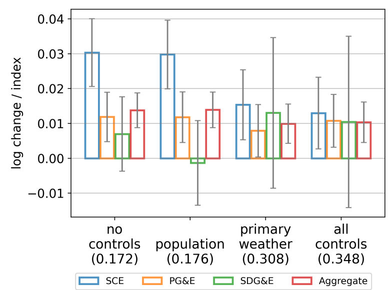

(e) CMI, SES Index

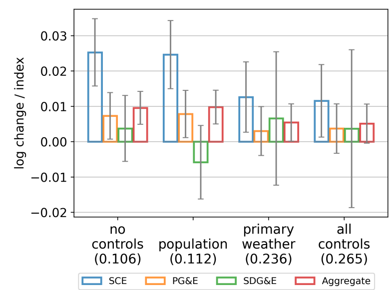

(f) Customers, SES Index

Figure 4: Summarizing coefficients from regressions, with various controls. (a) and (d) show logistic regression of ignitions, (b) and (e) show regression of log customer minutes interrupted, and (c) and (f) show log number of customers impacted. Groups on the x axis collect results from regression with a given set of controls. R squared values are shown in parentheses below each collection of regression results. For ignition, the average pseudo R squared from each first-stage regression is reported. Error bars show the 95% confidence interval. Each group of plots is ordered SCE, PG&E, SDG&E, Aggregate.

in Section 3.3, I use customer minutes interrupted (CMI) and number of customers impacted.

The results of these regressions are shown in Figure 4b and Figure 4e (for CMI) and Figure 4c and Figure 4f (for number of customers impacted). The patterns are generally similar between regression of CMI and number of customers impacted, although both are relatively noisy. Estimates with log CMI are generally larger in magnitude and more precisely estimated than those using log number of customers. Without controlling for weather variation, there is a significant positive correlation between low SES and the cost proxy for SCE (p value < 0.001) and PG&E (p value 0.001 for CMI, 0.029 for customers). There is a significant negative correlation for the health risk index for PG&E (p value < 0.001). After controlling for weather variation, these observations are largely similar.

After aggregating data from all utilities, I am able to get more precise estimates. I find that by both measures, the magnitude of disruption is significantly lower in high-health risk circuits (p value < 0.001). I find that CMI is positively associated with low-SES circuits (p value < 0.001), although I do not find a significant association between the number of customers impacted and SES.

Rules that determine the cost of PSPS may disadvantage low-SES or high-health risk populations if the number of customers is negatively correlated with these indices. This occurs regardless of whether the rules intend to discriminate based on these characteristics; that is, it is an example of statistical rather than taste-based discrimination (Guryan and Charles, 2013). The utility's decision rule places more weight on circuits with higher historical customer outages. If circuits with a higher share of vulnerable individuals are less impacted by shutoffs, the utility's rule calculates a lower cost from shutoffs in those circuits. These findings indicate that PG&E's decision rules may disadvantage high health risk circuits. Utilities can adjust their rules to avoid this potential discrimination, and some already are. SCE already scales part of their PSPS risk score by the size of populations with medical needs (Valdberg, Tozer, and Kilberg, 2021, p. 16), and SDG&E has plans to implement a similar program (SDG&E, 2021a, p. 30).

5 Conclusions

I find that PSPS is used more frequently in low-SES circuits among two of California's major utilities, and among higher health risk circuits in one of the major utilities. After controlling for population and weather variation, model fit improves but these patterns remain largely consistent. This shows that the difference in rates of PSPS by vulnerability indices is largely unexplained by population or weather differences.

I find some evidence that factors in firms' PSPS decisions are also associated with vulnerability indices. I find that ignitions are more frequent in low-SES circuits and in lower health risk circuits in PG&E, but otherwise do not find significant associations within any individual utility's service area. After aggregating data from all utilities, I find that both the probability of ignition and the magnitude of PSPS disruption are higher in low-SES circuits, and in lower health risk circuits.

This work starts to explore a gap in the literature on empirically assessing the equity of adaptation mechanisms. More research is needed in this area more broadly, as well as to better understand the impacts of electric utilities' response to wildfire risk. This research agenda is challenging without better data about the firm's problem, particularly how the firm computes costs and benefits of PSPS. These data would allow researchers to explore a broader range of research questions, such as the explorations of systematic bias in Obermeyer et al. (2019) or Rambachan (2021). Future work should also explore how utilities invest to reduce future wildfire risk, and whether these investments are equitably distributed among communities with different vulnerability to wildfire hazards.

Data Availability

The data that support the findings of this study are openly available. They can be accessed at `https://doi.org/10.7910/DVN/AAFQRQ`. The code used to generate this file is also openly available, and can be accessed at `https://github.com/maxoboe/PSPS_equity`.

References

Abatzoglou, John (2013). "Development of gridded surface meteorological data for ecological applications and modelling". In: *International Journal of Climatology* 33.1, pp. 121–131.

Anderson, Sarah, Andrew Plantinga, Matthew Wibbenmeyer, et al. (2020). "Inequality in agency responsiveness: evidence from salient wildfire events". In: *Resources for the Future: Washington, DC, USA*, p. 36.

August, Laura, Komal Bangia, Laurel Plummer, Shankar Prasa, Kelsey Ranjbar, Andrew Slocombe, and Walker Wieland (2021). *CalEnviroScreen 4.0.* California Environmental Protection Agency, Office of Environmental Health Hazard Assessment.

Becker, Gary S (1957). *The economics of discrimination.* University of Chicago press.

Buolamwini, Joy and Timnit Gebru (2018). "Gender shades: Intersectional accuracy disparities in commercial gender classification". In: *Conference on fairness, accountability and transparency.* PMLR, pp. 77–91.

California Public Utilities Commission (2014). *Order instituting rulemaking to revise and clarify commission regulations relating to the safety of electric utility and communications infrastructure provider facilities. Decision 19-05-042.*

— (2019). *Decision adopting de-energization (Public Safety Power Shut-off) guidelines (Phase 1 Guidelines). Decision 19-05-042.*

Cinelli, Carlos, Andrew Forney, and Judea Pearl (2022). "A Crash Course in Good and Bad Controls". In: *Sociological Methods & Research.*

Coggins, Shaugn, Lea Berrang-Ford, Keith Hyams, Poshendra Satyal, James Ford, Jouni Paavola, Ingrid Arotoma-Rojas, and Sherilee Harper (2021). "Empirical assessment of equity and justice in climate adaptation literature: A systematic map". In: *Environmental Research Letters* 16.7, p. 073003.

D'Evelyn, Savannah M et al. (2022). "Wildfire, Smoke Exposure, Human Health, and Environmental Justice Need to be Integrated into Forest Restoration and Management". In: *Current environmental health reports*, pp. 1–20.

Gianfrancesco, Milena A, Suzanne Tamang, Jinoos Yazdany, and Gabriela Schmajuk (2018). "Potential biases in machine learning algorithms using electronic health record data". In: *JAMA internal medicine* 178.11, pp. 1544–1547.

Guliasi, Leslie (2021). "Toward a political economy of public safety power shutoff: Politics, ideology, and the limits of regulatory choice in California". In: *Energy Research & Social Science* 71, p. 101842.

Guryan, Jonathan and Kerwin Kofi Charles (2013). "Taste-based or statistical discrimination: The economics of discrimination returns to its roots". In: *The Economic Journal* 123.572, F417–F432.

Hill, L L, R Blythe, E M Krieger, A Smith, A McPhail, and S B C Shonkoff (2020). "The public health dimensions of California wildfire and Wildfire Prevention, Mitigation, and Suppression". In: *Technical Report; Physicians and Scientists for Healthy Energy.*

Holmes, Thomas P, Robert J Huggett, and Anthony L Westerling (2008). "Statistical analysis of large wildfires". In: *The economics of forest disturbances.* Springer, pp. 59–77.

IPCC (2022). *Climate Change 2022: Impacts, Adaptation, and Vulnerability. Working Group II Contribution to the IPCC Sixth Assessment Report.* Ed. by H.-O. Pörtner et al. Cambridge University Press.

Jain, Piyush, Sean CP Coogan, Sriram Ganapathi Subramanian, Mark Crowley, Steve Taylor, and Mike D Flannigan (2020). "A review of machine learning applications in wildfire science and management". In: *Environmental Reviews* 28.4, pp. 478–505.

Khan, Shakeeb and Elie Tamer (2009). "Inference on endogenously censored regression models using conditional moment inequalities". In: *Journal of Econometrics* 152.2, pp. 104–119.

Lang, Kevin and Ariella Kahn-Lang Spitzer (2020). "Race discrimination: An economic perspective". In: *Journal of Economic Perspectives* 34.2, pp. 68–89.

McFadden, D. (1973). "Conditional Logit Analysis of Qualitative Choice Behaviour". In: *Frontiers in Econometrics*. Ed. by P. Zarembka. New York, NY, USA: Academic Press New York, pp. 105–142.

Mehrabi, Ninareh, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan (2021). "A survey on bias and fairness in machine learning". In: *ACM Computing Surveys (CSUR)* 54.6, pp. 1–35.

Niemi, Ernest G. and Kristin S Lee (2001). "Wildfire and poverty: an overview of the interactions among wildfires, fire-related programs, and poverty in the Western States". In: The Center for Watershed and Community Health, Mark O. Hatfield School of Government, Portland State University.

Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan (2019). "Dissecting racial bias in an algorithm used to manage the health of populations". In: *Science* 366.6464, pp. 447–453.

Ojerio, Ryan S (2008). "Equity in Wildfire Risk Management: Does Socioeconomic Status Predict Involvement in Federal Programs to Mitigate Wildfire Risk?" MA thesis. Department of Planning, Public Policy & Management, University of Oregon.

Oster, Emily (2019). "Unobservable selection and coefficient stability: Theory and evidence". In: *Journal of Business & Economic Statistics* 37.2, pp. 187–204.

Pacific Gas and Electric Company (2020). *2020 risk assessment and mitigation phase report*. Tech. rep.

— (June 2021). *2021 Wildfire mitigation plan – revised*. Tech. rep.

Parikh, Ravi B, Stephanie Teeple, and Amol S Navathe (2019). "Addressing bias in artificial intelligence in health care". In: *Jama* 322.24, pp. 2377–2378.

Plantinga, Andrew J, Randall Walsh, and Matthew Wibbenmeyer (2022). "Priorities and effectiveness in wildfire management: evidence from fire spread in the western United States". In: *Journal of the Association of Environmental and Resource Economists* 9.4, pp. 603–639.

Rambachan, Ashesh (2021). "Identifying prediction mistakes in observational data". In: *Working Paper, Harvard University Department of Economics.*

Rhodes, Noah, Lewis Ntaimo, and Line Roald (2020). "Balancing wildfire risk and power outages through optimized power shut-offs". In: *IEEE Transactions on Power Systems* 36.4, pp. 3118–3128.

San Diego Gas and Electric Company (Feb. 2021a). *2020-2022 Wildfire mitigation plan update*. Tech. rep.

— (2021b). *San Diego Gas and Electric Company: risk assessment mitigation phase*. Tech. rep.

Satyal, Poshendra, Morten Fibieger Byskov, and Keith Hyams (2021). "Addressing multi-dimensional injustice in indigenous adaptation: the case of Uganda's Batwa community". In: *Climate and Development* 13.6, pp. 529–542.

Sheller, Mimi and Yolanda M Leon (2016). "Uneven socio-ecologies of Hispaniola: Asymmetric capabilities for climate adaptation in Haiti and the Dominican Republic". In: *Geoforum* 73, pp. 32–46.

Short, K C (2014). *A spatial database of wildfires in the United States, 1992-2011.*

Short, Karen C (2021). *Spatial wildfire occurrence data for the United States, 1992-2018.*

Smit, Barry and Olga Pilifosova (2003). "Adaptation to climate change in the context of sustainable development and equity". In: *Sustainable Development* 8.9, p. 9.

Southern California Edison Company (June 2021). *2021 Wildfire mitigation plan update (revision)*. Tech. rep.

— (2022). *Southern California Edison Company: risk assessment mitigation phase*. Tech. rep.

Taylor, Steve W, Douglas G Woolford, CB Dean, and David L Martell (2013). "Wildfire prediction to inform fire management: statistical science challenges". In: *Statistical Science* 28.4, pp. 586–615.

Thomas, David SG and Chasca Twyman (2005). "Equity and justice in climate change adaptation amongst natural-resource-dependent societies". In: *Global environmental change* 15.2, pp. 115–124.

Valdberg, Anna, Andrea Tozer, and Elena Kilberg (2021). *Southern California Edison Company's (U 338-E) public safety power shutoff post-event report for November 21, 2021 de-energization event*. Tech. rep.

Westerling, Anthony L and BP Bryant (2008). "Climate change and wildfire in California". In: *Climatic Change* 87.1, pp. 231–249.

Wibbenmeyer, Matthew and Molly Robertson (2022). "The distributional incidence of wildfire hazard in the western United States". In: *Environmental Research Letters* 17.6, p. 064031.

Wigtil, Gabriel, Roger B Hammer, Jeffrey D Kline, Miranda H Mockrin, Susan I Stewart, Daniel Roper, and Volker C Radeloff (2016). "Places where wildfire potential and social vulnerability coincide in the coterminous United States". In: *International Journal of Wildland Fire* 25.8, pp. 896–908.

Wong-Parodi, Gabrielle (2020). "When climate change adaptation becomes a "looming threat" to society: Exploring views and responses to California wildfires and public safety power shutoffs". In: *Energy Research and Social Science* 70.

Xi, DD, Stephen W Taylor, Douglas G Woolford, and CB Dean (2019). "Statistical models of key components of wildfire risk". In: *Annual Review of Statistics and Its Application* 6.1, pp. 197–222.

A Weather Summary Statistics

In this section, I include additional summary statistics about relevant variables from the GridMET dataset. Table 3 gives the average of each primary and derived weather variable, for all observations in my dataset, observations during a red flag warning, and circuits during a PSPS event. Red Flag Warnings are drier, hotter, and more elevated than the full sample, and PSPS events occur in windier and drier conditions and in higher locations. The derived variables Burning Index and Energy Release Component, two measures of the potential for a large fire, are higher during Red Flag Warnings and during PSPS events.

Figure 5 shows how some of these key weather variables have changed over time. Relative humidity has declined in all service areas over this sample, and maximum temperature has increased overall in California and in SCE and SDG&E's service areas. The Burning Index and Energy Release Component are increasing in all service areas.

	(1)	(2)	(3)
Max Air Temperature (C)	23.85	28.26	23.45
	(8.039)	(6.359)	(6.404)
Min Air Temperature (C)	10.01	11.85	9.892
	(5.573)	(5.104)	(4.879)
Precipitation Amount (daily mm)	1.240	0.0587	0.00353
	(5.506)	(0.772)	(0.0723)
Specific Humidity (kg/kg)	0.00641	0.00523	0.00396
	(0.00217)	(0.00252)	(0.00168)
Wind Velocity at 10 m (m/s)	3.320	3.674	5.332
	(1.590)	(1.669)	(2.123)
Wind From Direction (Degrees past North)	233.8	228.4	224.6
	(83.69)	(92.18)	(112.8)
Mean Vapor Pressure Deficit (kPa)	1.273	1.946	1.572
	(0.948)	(0.840)	(0.699)
Max Relatively Humidity (%)	78.15	57.56	51.27
	(19.43)	(21.52)	(19.49)
Min Relatively Humidity (%)	33.84	16.63	15.35
	(18.50)	(11.88)	(9.919)
Surface Downwelling Shortwave Flux (W/m^2)	223.8	232.1	190.1
	(96.92)	(74.97)	(47.95)
Burning Index (Derived)	36.06	54.76	68.36
	(20.94)	(16.75)	(19.70)
Energy Release Component (Derived)	45.87	66.71	69.75
	(23.76)	(15.08)	(14.25)
Potential Evapotranspiration (Derived, mm)	4.204	5.450	5.114
	(2.414)	(2.149)	(1.803)
Reference Evapotranspiration (Derived, mm)	5.770	8.009	8.015
	(3.353)	(3.100)	(2.824)
Dead Fuel Moisture 100 hr (Derived, %)	12.96	8.608	7.879
	(5.108)	(3.107)	(2.656)
Dead Fuel Moisture 1000 hr (Derived, %)	14.08	9.844	9.371
	(5.491)	(2.672)	(2.728)
Elevation	223.4	293.8	468.8
	(319.2)	(377.6)	(401.3)
Observations	16303343	652310	3333

Table 3: Summary statistics of GridMET data from October 2013 through 2021. Columns separate the full sample, the sample during a red flag warning, and the sample during a PSPS event. Observations are weighted by the length of each circuit segment. Standard errors of the mean for each column are in parentheses.

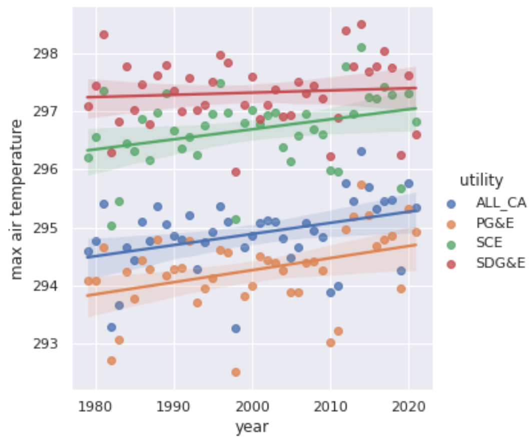

(a) Maximum temperature

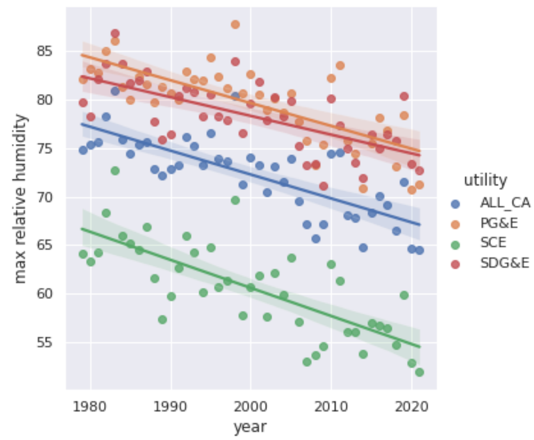

(b) Maximum humidity

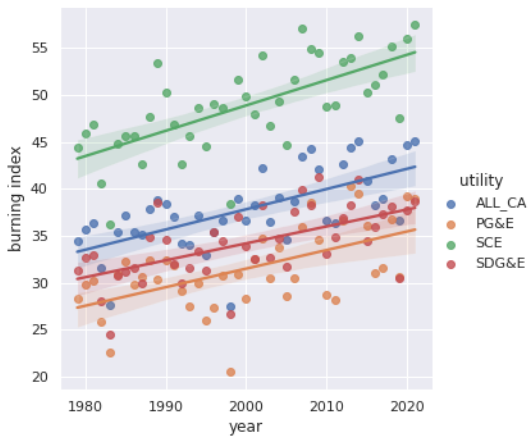

(c) Burning index

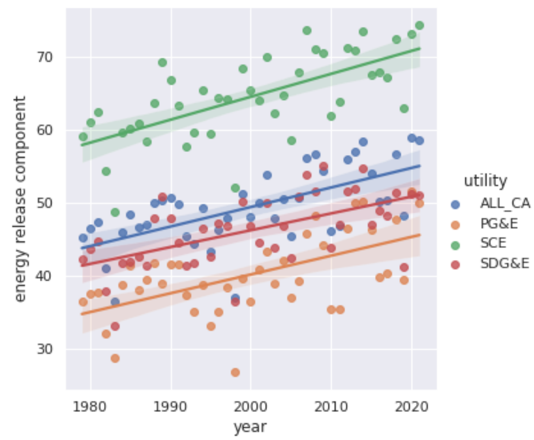

(d) Energy release component

Figure 5: Selected weather variables over time. Average values in all of California, and in the service area of the three major utilities.

B Predicting Fire Size

I use public data to attempt to predict fire size given weather covariates. From 1992-2018, comprehensive records of fire size are available from the US Forest Service (K C Short, 2014; Karen C Short, 2021). From 2019-2021, I include records from the National Interagency Fire Service.[3] Records include the date, fire size, and latitude and longitude of ignition. The final database includes 240,239 records within California. I then merge these data with my weather observations from GridMET.

I model the problem of predicting catastrophic fires both as a regression and classification

[3] From `https://data-nifc.opendata.arcgis.com/datasets/nifc::wfigs-current-wildland-fire-perimeters/about`, accessed 11 January 2022.

	Regression R squared	Fire size ¿= 300 acres Specificity	Sensitivity	Fire size ¿= 500 acres Specificity	Sensitivity	Fire size ¿= top 2% Specificity	Sensitivity
Linear	0.06311	0.6907	0.6455	0.7017	0.6624	0.6691	0.6115
Linear Interacted	0.08927	0.723	0.6612	0.7302	0.6893	0.6888	0.6271
Random Forest	0.08446	0.02007	0.9994	0.01766	0.9994	0.02622	0.9991

Table 4: Results from random forest and linear regression at predicting large fires.

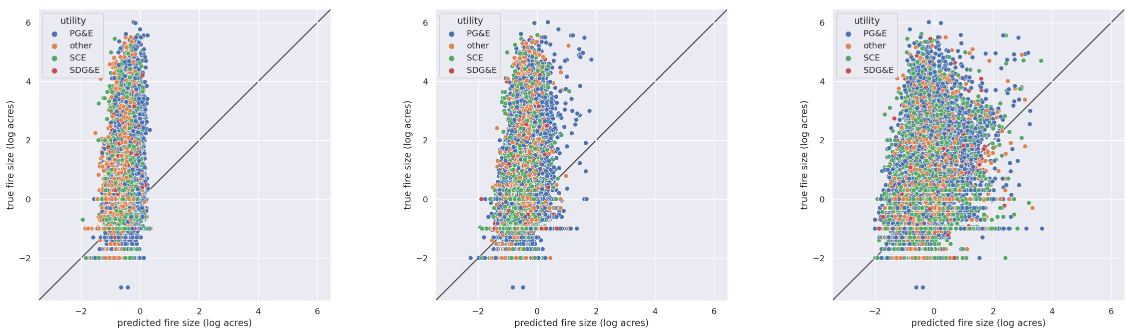

(a) Linear

(b) Linear interacted

(c) Random forest

Figure 6: Predicted vs actual fire size, using various regression methods.

problem. To predict fire size, I regress the log of fire size against the full set of weather variables from GridMET, as well as yearly fixed effects and fixed effect terms per utility's service area. For classification trials, I use three definitions of "large fire": top 0.02 quantile (my definition), larger than 300 acres[4], and larger than 500 acres.[5]. In each classification trial, I weight each observation by the inverse frequency of its class to predict the relatively rare event of a large fire. I consider a linear set of weather variables, linear regression with interactions between weather variables, and random forests with 5-folds cross validation.

Table 4 summarizes the results of these trials. For classification trials, I report the specificity (share of negative outcomes that are correctly predicted) and sensitivity (share of positive outcomes that are correctly predicted) of each prediction method, for each "large fire" definition. For regression, I report the R^2 value. Figure 6 shows the scatter plots of predicted fire size vs. actual fire size.

Overall, these results indicate poor performance at predicting fire size. I have limited ability to extrapolate fire size, meaning I cannot construct informative bounds on the missing data as

[4]Definition from `https://www.nps.gov/olym/learn/management/upload/fire-wildfire-definitions-2.pdf`, accessed 1 April 2022.

[5]Definition from Holmes, Huggett, and Westerling (2008).

required to identify counterfactuals in Rambachan (2021).

C Regression Results in Tabular Form

Here, I include tabular forms of the visualizations showing the results of the main regression. The regression methods are described in Section 3, and these results are discussed in more detail in Section 4. Separate regression tables are provided for aggregate and utility-level analyses.

Table 5 and Table 6 show the results from logistic regression of PSPS shutoffs on the vulnerability indices and additional controls. Before aggregating data across utilities, higher health risk circuits are significantly less likely to have a PSPS in both SCE and PG&E, and more likely in SDG&E. This finding is significant at the $p = 0.001$ level for PG&E and SDG&E, and at the $p = 0.05$ level for SCE. Lower SES circuits are more likely to have a shutoff in PG&E and SDG&E, and less likely in SCE; this finding is significant at the $p = 0.001$ level for PG&E and SDG&E, and at the $p = 0.01$ level for SCE. After aggregating data, and controlling for weather variation, there is not a very significant relationship between SES and PSPS shutoffs. This is due to the heterogeneity in effects at the utility level. However, PSPS is significantly less likely in high-health risk circuits.

Table 7 and Table 8 show the results from logistic regression of ignitions on the vulnerability indices and additional controls. The pseudo-R squared value is relatively low, even for the model with both primary and derived weather covariates. Before aggregating data across utilities, results are generally quite noisy, and many coefficients are not statistically distinguishable from 0. After aggregating data, ignition probability is significantly lower in low-health risk circuits and significantly higher in low-SES circuits.

Table 9 and Table 10 show the coefficient estimates from linear regression of the log of CMI (columns 1-4) and the log number of customers impacted (columns 5-8). Before aggregating data across utilities, results are generally quite noisy, and many coefficients are not statistically distinguishable from 0. After aggregating data, these outcomes are significantly higher in low-health risk circuits, and the CMI is significantly higher in low-SES circuits.

	(1)	(2)	(3)	(4)
	psps	psps	psps	psps
SCE x Health	-0.00335*	-0.00355	0.00136	-0.000467
	(0.00167)	(0.00222)	(0.00226)	(0.00239)
SCE x SES	-0.00565**	-0.00562**	-0.0247***	-0.0244***
	(0.00202)	(0.00211)	(0.00271)	(0.00270)
PG&E x Health	-0.0193***	-0.0184***	-0.0111***	-0.0154***
	(0.00145)	(0.00147)	(0.00144)	(0.00163)
PG&E x SES	0.0122***	0.0111***	0.00987***	0.0161***
	(0.00129)	(0.00124)	(0.00157)	(0.00182)
SDG&E x Health	0.0394***	0.0465***	0.0386***	0.0360***
	(0.00425)	(0.00498)	(0.00635)	(0.00728)
SDG&E x SES	0.0247***	0.0404***	0.0148**	0.0220***
	(0.00397)	(0.00385)	(0.00452)	(0.00495)
Observations	375064	375064	370411	370411
Pseudo R^2	0.081	0.090	0.292	0.371
Population		X	X	X
Primary			X	X
Derived				X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5: Results from logistic regression of PSPS events. Perfectly predicted failures are omitted. A positive coefficient indicates that higher health risk (lower SES) circuits have a greater rate of PSPS shutoffs.

	(1)	(2)	(3)	(4)
	psps	psps	psps	psps
Health Index	-0.0104***	-0.0122***	-0.00495***	-0.00780***
	(0.00107)	(0.00124)	(0.00118)	(0.00130)
SES Index	0.00695***	0.00872***	0.000340	0.00377*
	(0.00109)	(0.00112)	(0.00135)	(0.00150)
Observations	375064	375064	370411	370411
Pseudo R^2	0.075	0.081	0.287	0.366
Population		X	X	X
Primary			X	X
Derived				X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 6: Regression of PSPS shutoffs, after aggregating data from all utilities. Perfectly predicted failures are omitted. A positive coefficient indicates that higher health risk (lower SES) circuits have a greater rate of PSPS shutoffs.

	(1) ignition	(2) ignition	(3) ignition	(4) ignition	(5) ignition	(6) ignition	(7) ignition	(8) ignition	(9) ignition
Health Index	-0.00314	-0.0113***	0.0213**	-0.00178	-0.0111***	0.0162*	0.00720*	-0.0123***	0.0114
	(0.00209)	(0.00114)	(0.00665)	(0.00264)	(0.00115)	(0.00723)	(0.00303)	(0.00127)	(0.00871)
SES Index	0.00308	0.0148***	0.0174**	0.00244	0.0146***	0.0196**	-0.00436	0.00936***	0.00969
	(0.00258)	(0.00110)	(0.00664)	(0.00262)	(0.00110)	(0.00625)	(0.00278)	(0.00133)	(0.00672)
Observations	1950168	4939641	700344	1950168	4939641	700344	1950168	4939641	700344
Pseudo R^2	0.006	0.007	0.013	0.007	0.007	0.017	0.047	0.059	0.080
Utility	SCE	PG&E	SDG&E	SCE	PG&E	SDG&E	SCE	PG&E	SDG&E
Population				X	X	X	X	X	X
Primary							X	X	X
Derived							X	X	X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 7: Results from logistic regression of PSPS events. A positive coefficient indicates that higher health risk (lower SES) circuits have a greater rate of ignitions. Robust standard errors are reported.

	(1) ignition	(2) ignition	(3) ignition	(4) ignition
Health Index	-0.00773***	-0.00824***	-0.00775***	-0.00763***
	(0.000998)	(0.00106)	(0.00116)	(0.00117)
SES Index	0.0118***	0.0122***	0.00632***	0.00577***
	(0.00102)	(0.00103)	(0.00119)	(0.00120)
Observations	7590153	7590153	7590153	7590153
Pseudo R^2	0.009	0.010	0.054	0.059
Population		X	X	X
Primary			X	X
Derived				X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 8: Results from logistic regression of ignition, after aggregating data from all utilities. A positive coefficient indicates that higher health risk (lower SES) circuits have a greater rate of ignitions. Robust standard errors are reported.

	(1) Log CMI	(2) Log CMI	(3) Log CMI	(4) Log CMI	(5) Log Cust	(6) Log Cust	(7) Log Cust	(8) Log Cust
SCE x Health	-0.00808*	-0.00780*	0.00861	0.00932	-0.00633*	-0.00606	0.0106*	0.0118*
	(0.00327)	(0.00384)	(0.00487)	(0.00505)	(0.00309)	(0.00364)	(0.00475)	(0.00491)
SCE x SES	0.0303***	0.0298***	0.0154**	0.0130*	0.0253***	0.0246***	0.0126*	0.0115*
	(0.00497)	(0.00502)	(0.00512)	(0.00525)	(0.00486)	(0.00493)	(0.00507)	(0.00523)
PG&E x Health	-0.0208***	-0.0207***	-0.0153***	-0.0163***	-0.0167***	-0.0170***	-0.0137***	-0.0134***
	(0.00275)	(0.00280)	(0.00286)	(0.00294)	(0.00255)	(0.00258)	(0.00261)	(0.00268)
PG&E x SES	0.0119**	0.0118**	0.00790*	0.0107**	0.00731*	0.00783*	0.00302	0.00372
	(0.00363)	(0.00369)	(0.00384)	(0.00385)	(0.00335)	(0.00340)	(0.00354)	(0.00356)
SDG&E x Health	-0.0108	-0.0315**	-0.0353*	-0.0246	-0.00472	-0.0274**	-0.0302*	-0.0260
	(0.00803)	(0.0104)	(0.0162)	(0.0194)	(0.00654)	(0.00856)	(0.0137)	(0.0171)
SDG&E x SES	0.00697	-0.00131	0.0130	0.0104	0.00375	-0.00582	0.00659	0.00370
	(0.00545)	(0.00620)	(0.0110)	(0.0125)	(0.00476)	(0.00531)	(0.00963)	(0.0114)
Observations	3270	3270	3270	3270	3270	3270	3270	3270
R^2	0.172	0.176	0.308	0.348	0.106	0.112	0.236	0.265
Population		X	X	X		X	X	X
Primary			X	X			X	X
Derived				X				X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 9: Results from linear regression of proxies to PSPS cost, CMI (columns 1-4) and number of customers (columns 5-8). A positive coefficient indicates that higher health risk (lower SES) circuits have higher average CMI or number of customers impacted. Robust standard errors are reported. Outcome variable values of 0 are omitted.

	(1) Log CMI	(2) Log CMI	(3) Log CMI	(4) Log CMI	(5) Log Cust	(6) Log Cust	(7) Log Cust	(8) Log Cust
Health Index	-0.0167***	-0.0179***	-0.0123***	-0.0119***	-0.0134***	-0.0146***	-0.0106***	-0.00958***
	(0.00198)	(0.00210)	(0.00235)	(0.00248)	(0.00184)	(0.00195)	(0.00216)	(0.00228)
SES Index	0.0138***	0.0139***	0.00992***	0.0103***	0.00957***	0.00979***	0.00544*	0.00513
	(0.00254)	(0.00260)	(0.00287)	(0.00297)	(0.00237)	(0.00244)	(0.00270)	(0.00282)
Observations	3270	3270	3270	3270	3270	3270	3270	3270
R^2	0.164	0.169	0.301	0.343	0.098	0.104	0.227	0.257
Population		X	X	X		X	X	X
Primary			X	X			X	X
Derived				X				X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10: Results from linear regression of proxies to PSPS cost, CMI (columns 1-4) and number of customers (columns 5-8). A positive coefficient indicates that higher health risk (lower SES) circuits have higher average CMI or number of customers impacted. Robust standard errors are reported. Outcome variable values of 0 are omitted.

D PSPS regression with full sample

In the main text, I estimate PSPS probability using data during red flag warnings. As a robustness exercise, I estimate the same relationships using the full sample. I only include observations from years where a utility is using PSPS; this is without loss of generality as the model includes year-by-utility fixed effects.

Table 11 and Table 12 show the results from logistic regression of PSPS shutoffs on the vulnerability indices and additional controls, with all data. Before aggregating data across utilities, the results are generally similar to those from the main estimation, although magnitudes of effect sizes are smaller. Higher health risk circuits are significantly less likely to have a PSPS in both SCE and PG&E, and more likely in SDG&E. This finding is significant at the $p = 0.001$ level for PG&E and SDG&E, and at the $p = 0.05$ level for SCE. Lower SES circuits are more likely to have a shutoff in PG&E and SDG&E, and less likely in SCE; this finding is significant at the $p = 0.001$ level for SDG&E, and at the $p = 0.01$ level for SCE and PG&E. After aggregating data, and controlling for weather variation, there is a significant negative correlation between both indices and likelihood of PSPS shutoff.

	(1)	(2)	(3)	(4)
	psps	psps	psps	psps
SCE x Health	-0.00226	-0.00275	-0.000788	-0.00441*
	(0.00165)	(0.00213)	(0.00224)	(0.00223)
SCE x SES	-0.0100***	-0.00988***	-0.0298***	-0.0301***
	(0.00196)	(0.00208)	(0.00270)	(0.00265)
PG&E x Health	-0.0128***	-0.0118***	-0.00465***	-0.00683***
	(0.00101)	(0.00101)	(0.00115)	(0.00125)
PG&E x SES	0.00848***	0.00757***	-0.00118	0.00365**
	(0.000902)	(0.000858)	(0.00119)	(0.00142)
SDG&E x Health	0.0429***	0.0349***	0.0359***	0.0312***
	(0.00352)	(0.00385)	(0.00523)	(0.00557)
SDG&E x SES	0.0189***	0.0314***	0.0147***	0.0195***
	(0.00421)	(0.00366)	(0.00429)	(0.00430)
Observations	7335132	7335132	6658739	6658739
Pseudo R^2	0.051	0.059	0.357	0.445
Population		X	X	X
Primary			X	X
Derived				X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 11: Results from logistic regression of PSPS events, using the full sample. Perfectly predicted failures are omitted. A positive coefficient indicates that higher health risk (lower SES) circuits have a greater rate of PSPS shutoffs.

	(1)	(2)	(3)	(4)
	psps	psps	psps	psps
Health Index	-0.00684***	-0.00820***	-0.00165	-0.00380***
	(0.000840)	(0.000931)	(0.00101)	(0.00106)
SES Index	0.00407***	0.00559***	-0.00631***	-0.00328**
	(0.000839)	(0.000849)	(0.00112)	(0.00126)
Observations	7335132	7335132	6658739	6658739
Pseudo R^2	0.046	0.054	0.354	0.442
Population		X	X	X
Primary			X	X
Derived				X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 12: Regression of PSPS shutoffs, after aggregating data from all utilities, using the full sample. Perfectly predicted failures are omitted. A positive coefficient indicates that higher health risk (lower SES) circuits have a greater rate of PSPS shutoffs.

E Ignitions regression with alternate sample

In the main text, I estimate ignition probability using data from years where firms do not declare PSPS, to avoid a potential data censoring problem. As discussed in Section 3.2, this choice if the relationship between weather variables and ignition probability is not consistent between years when utilities do and do not use PSPS. As a robustness exercise, I include results using the full time period, with two different assumptions on the observed ignitions and shutoffs: that, absent a shutoff, each circuit with a PSPS event would either have an ignition (Table 13) or no ignition (Table 15). Overall, these results suggest that my conclusions in the main analysis are robust to including ignition data from the full sample.

This provides suggestive evidence about the partially identified set that contains the true parameter. To fully characterize that set, I could enumerate all possible potential realizations of the missing data and repeat the estimation procedure for each potential outcome. Due to the immense computational cost of such a procedure, I only repeat the two-stage estimation procedure for these two scenarios.

Table 13 and Table 14 show the results from the using the full sample and treating each missing value as a true positive. The findings that are statistically significant (at the $p = 0.001$ level) from Table 7 in the main analysis are also significant in these regressions: low-SES circuits in PG&E (and in the aggregated regression) have higher rates of ignition, as do lower health risk circuits. This regression finds additional significant associations, although many of these are not robust to an alternate assumption on the missing data. Note that the pseudo R squared is much higher in this sample - this is because the sample now includes PSPS shutoffs, where logistic regression has a higher accuracy than with ignitions.

Table 15 shows the results from using the full sample and treating each missing value as a true negative. Again, the statistically significant findings from the main analysis are confirmed in these regressions. Many of the additional significant associations from treating each missing value as a true positive are not significant in this exercise, although both find that lines in low-SES circuits in SDG&E are significantly more likely to have an ignition.

	(1) ignition	(2) ignition	(3) ignition	(4) ignition	(5) ignition	(6) ignition	(7) ignition	(8) ignition	(9) ignition
Health Index	-0.000153	-0.0123***	0.0322***	0.00128	-0.0117***	0.0243***	0.00287	-0.0118***	0.0187***
	(0.00113)	(0.000615)	(0.00272)	(0.00141)	(0.000617)	(0.00303)	(0.00152)	(0.000693)	(0.00372)
SES Index	-0.00299*	0.0134***	0.0197***	-0.00359**	0.0129***	0.0253***	-0.0150***	0.00963***	0.0106***
	(0.00134)	(0.000583)	(0.00292)	(0.00138)	(0.000575)	(0.00255)	(0.00164)	(0.000744)	(0.00290)
Observations	3120696	9879282	1925307	3120696	9879282	1925307	3120696	9879282	1896380
Pseudo R^2	0.034	0.029	0.045	0.035	0.030	0.069	0.198	0.165	0.329
Utility	SCE	PG&E	SDG&E	SCE	PG&E	SDG&E	SCE	PG&E	SDG&E
Population				X	X	X	X	X	X
Primary							X	X	X
Derived							X	X	X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 13: Results from logistic regression of ignitions, using the full sample, with PSPS results counted as ignitions. This is the assumption that all censored results would have been true positives.

	(1) psps	(2) psps	(3) psps	(4) psps
Health Index	-0.00688***	-0.0077***	-0.00616***	-0.00712***
	(0.000534)	(0.000573)	(0.000611)	(0.000619)
SES Index	0.00934***	0.0100***	0.00318***	0.00419***
	(0.000541)	(0.000546)	(0.000663)	(0.000686)
Observations	14925285	14925285	14896358	14896358
Pseudo R^2	0.031	0.034	0.159	0.183
Population		X	X	X
Primary			X	X
Derived				X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 14: Results from logistic regression of ignitions, using the full sample, with PSPS results counted as ignitions. Results are aggregated across all utilities. This is the assumption that all censored results would have been true positives.

	(1) ignition	(2) ignition	(3) ignition	(4) ignition	(5) ignition	(6) ignition	(7) ignition	(8) ignition	(9) ignition
Health Index	0.00144	-0.0119***	0.0179***	0.00460*	-0.0117***	0.0133**	0.00889***	-0.0137***	0.00860
	(0.00155)	(0.000769)	(0.00430)	(0.00187)	(0.000776)	(0.00461)	(0.00207)	(0.000871)	(0.00553)
SES Index	0.00298	0.0167***	0.0205***	0.00178	0.0166***	0.0224***	-0.00362	0.0125***	0.0137***
	(0.00180)	(0.000759)	(0.00404)	(0.00182)	(0.000761)	(0.00376)	(0.00200)	(0.000924)	(0.00409)
Observations	3120696	9879282	1867158	3120696	9879282	1867158	3120696	9879282	1846371
Pseudo R^2	0.007	0.007	0.013	0.008	0.007	0.019	0.052	0.060	0.086
Utility	SCE	PG&E	SDG&E	SCE	PG&E	SDG&E	SCE	PG&E	SDG&E
Population				X	X	X	X	X	X
Primary							X	X	X
Derived							X	X	X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 15: Results from logistic regression of ignitions, using the full sample, with no PSPS results counted as ignitions. This is the assumption that all censored results would have been true negatives.

	(1) ignition	(2) ignition	(3) ignition	(4) ignition
Health Index	-0.00694***	-0.00757***	-0.00862***	-0.00854***
	(0.000691)	(0.000726)	(0.000796)	(0.000801)
SES Index	0.0131***	0.0135***	0.00875***	0.00854***
	(0.000704)	(0.000710)	(0.000827)	(0.000839)
Observations	14867136	14867136	14846349	14846349
Pseudo R^2	0.010	0.011	0.057	0.062
Population		X	X	X
Primary			X	X
Derived				X

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 16: Results from logistic regression of ignitions, using the full sample, with no PSPS results counted as ignitions. Results are aggregated across all utilities. This is the assumption that all censored results would have been true negatives.