Socioeconomic Vulnerability and Hurricane-Related Outages: Evidence from Hurricane María in Puerto Rico

John Mantus*

November 30, 2023

Abstract

Hurricane María made landfall in Puerto Rico on September 20, 2017 as a Category 4 storm and caused Puerto Rico's entire energy grid to fail, leading to the longest blackout in American history. The storm damaged critical infrastructure and reduced households' access to food, water, and medical care. Several studies of María find a significant, positive relationship between outage duration and several measures of social vulnerability. Formulating the proper policy response to address this inequality requires identifying the underlying causes, and recent studies of the Southeast United States suggest recovery procedures themselves contribute. In this paper, I apply spatial regression methods to data from power recovery crew deployments following Hurricane María and find, conditional on the impact of the storm itself, a significant, positive relationship between outage duration and socioeconomic vulnerability, but no statistically-significant relationship between two other forms of social vulnerability and outage duration, namely vulnerability defined by the quality of available housing and transportation and by household composition (e.g., 65+ population). This is consistent with recent studies of the American Southeast but contradicts another recent study of María which relies on the same data but uses methods ill-suited for the empirical setting. In addition to this primary analysis and unique to this paper, I obtain geospatial data on Puerto Rican infrastructure to explore alternative spatial weight matrices and test for potential biases caused by standard weighting practices in the literature. I find no evidence of such biases.

*I would like to thank Scott Ganz for his tremendous patience and assistance throughout this project.

1 Introduction

Hurricane María made landfall in Puerto Rico on September 20, 2017 as a Category 4 storm and caused Puerto Rico's entire energy grid to fail, leading to the longest blackout in American history (Kwasinski et al., 2019). The storm damaged critical infrastructure for the island's communications, transportation, water supply, and wastewater treatment, reducing households' access to food, water, and medical care (Fischbach et al., 2020). Estimates of the death toll reach as high as 4,645 (Kishore et al., 2018).

Several studies of María find these harms were disproportionately shouldered by socioeconomically vulnerable households. Roman et al. (2019) collect data from satellite imagery taken over a six-month period following María's landfall and find that poorer residents experienced longer outages. Tormos-Aponte et al. (2021) use data on the timings and locations of power restoration crew deployments by the Puerto Rico Electric Power Authority (PREPA)—the power company responsible for Puerto Rico's electric infrastructure—and similarly find a positive relationship between outage duration and socioeconomic vulnerability, as well as three other forms of social vulnerability. Studies of hurricanes elsewhere in the U.S. present robust statistical to support this claim, as well (Lievanos and Horne, 2017; Coleman et al., 2023). However, formulating the proper policy response to this finding requires identifying the underlying causes of these disparities. It is likely, for example, that socioeconomically vulnerable households experience longer outages because they are situated in more at-risk locations for hurricanes (Logan and Xu, 2015). Alternatively, recent studies of the Southeast United States suggest the recovery procedures used by electric utility organizations themselves contribute, as well (Mitsova et al., 2018; Best et al., 2023; Ganz et al., 2023). Typical residential recovery procedures used throughout the U.S (including Puerto Rico) prioritize restoring power to the most households as quickly as possible, regardless of any form of vulnerability (Edison Electric Institute, 2016; U.S. Department of Energy, 2018). This "colorblind" (Tormos-Aponte et al., 2021, p. 1) approach to household power restoration may inadvertently exacerbate existing inequalities across communities of varying

levels of socioeconomic vulnerability. These studies statistically identify the role of recovery procedures using spatial regression techniques which both control for the direct effect of the storm and account for various dependencies between neighboring areas.

This paper is the first to apply such methods in the context of Hurricane María. Using data from PREPA restoration crew deployments—the same data used in Tormos-Aponte et al. (2021)—I estimate a series of spatial lag models to test the relationship between socioeconomic vulnerability and power outage duration. I find a one-decile increase in socioeconomic vulnerability (as measured by the CDC/ATSDR's Social Vulnerability Index, described in detail below) is associated with a 3.8% increase in outage duration, conditional on the effect of the storm itself. For the median household in Puerto Rico, this translates to an additional 3 days without power. Unlike Tormos-Aponte et al. (2021), I find no statistically-significant relationship between two other forms of vulnerability ("Household Characteristics" and "Housing Type & Transportation") and outage duration. These results are robust to several alternative specifications. Further, using data on the physical locations of power lines and major roads, this paper explores alternative weighting schemes which may improve model fit and reveal potential biases caused by standard weighting schemes. When estimating the spatially lagged relationship between geographically adjacent counties' outages, current practices weight all neighbors equally regardless of whether or not a power line or roads runs between a focal tract and that geographic neighbor. This exploration of alternative weighting schemes which account for these dependencies suggests the standard weighting scheme actually improves the fit of the model compared to estimates produced by these alternatives.

This paper contributes to two literatures. First, it is the first paper to use spatial-lag models to study the socioeconomic effects of Hurricane María and the ensuing outage recovery efforts. In doing so, it improves our understanding of exactly how outage recovery procedures in Puerto Rico exacerbate outage inequalities caused by hurricanes. Second, this paper finds similar results to studies of hurricane recoveries elsewhere in the United States, contributing

to our understanding of American outage recovery procedures more broadly. Further, and unique to this paper, by testing a series of spatial weight matrices which incorporate inter-tract dependencies on power lines and major roads, I test and support the legitimacy of the current practice of using a purely geographic-adjacency weighting scheme in this setting. Future researchers can use these results to justify the use of this weighting scheme or may wish to repeat the exercises presented in this paper to conduct similar tests in the context of the storm they study.

The paper is organized as follows. Section 2 describes the data used in this study. Section 3 describes the model estimated in this paper and the method used to implement alternative spatial weight matrices. Section 4 presents the spatial model used, several baseline estimates using geographic-adjacency weighting, and the results from an exploration of alternative weighting schemes. Section 5 summarizes the findings of this paper and provides recommendations for possible reforms.

2 Data[1]

2.1 Primary Data

Power outage duration data is from the Puerto Rico Electric Power Authority (PREPA). PREPA kept detailed records of power restoration crew deployments in the aftermath of Hurricane María, including the geo-location and date of each deployment. These data cover 18,736 crew deployments across 777 of 945 Census tracts in Puerto Rico. Outage duration is computed as the median number of days between María's landfall (September 20, 2017) and the date of service crew deployments within each Census tract.

The Centers for Disease Control and Prevention and the Agency for Toxic Substances and Disease Registry (CDC/ATSDR) regularly compute a Social Vulnerability Index (SVI)

[1]All data used in this study are publicly-available; a complete replication package is available upon request.

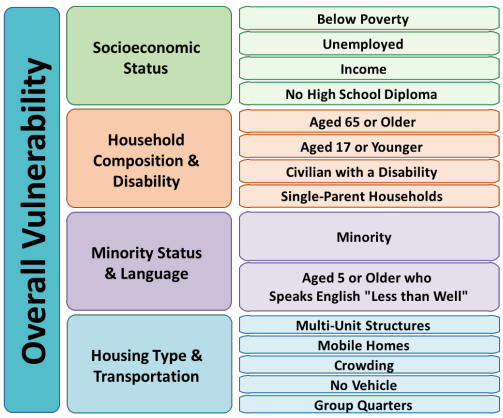

Figure 1: Variables from American Community Survey used to compute the Social Vulnerability Index and its four underlying themes.
Source: CDC SVI Documentation, 2016.

as a measure of social vulnerability at the tract level which is widely used to study the differential effects of natural disasters (see, e.g., Mitsova et al. (2018); Flores et al. (2023); Tormos-Aponte et al. (2021); Do et al. (2023); Ganz et al. (2023)). SVI is a percentile ranking of tracts according to four distinct themes: Socioeconomic Status, Household Characteristics, Racial and Ethnic Minority Status, and Housing Type & Transportation. The four themes are denoted henceforth as SES, HC, REM, and HTT vulnerabilities, respectively, and they are aggregated by the CDC/ATSDR to compute an overall SVI index (see Flanagan et al. (2018) for details). Each theme is computed using data from the American Community Survey (ACS) with the exact variables used to compute each theme found in Figure 1. The SES vulnerability theme is of primary interest in this paper, while HC and HTT are included as covariates. REM is omitted for reasons explained below. Greater SVI values indicate higher vulnerability.

As shown in Figure 1, the REM vulnerability theme is based on tract-level prevalence of English-speaking abilities and racial/ethnic minority populations, where minority populations are defined as those of the overall United States. This raises several flags in the context of Puerto Rico. First, 95% of Puerto Rican households speak a language other than English at home so the inability to speak English is unlikely to be a source of vulnerability. Second,

99% of Puerto Ricans self-identify as Hispanic, a minority population in the overall U.S. population. It is possible to recompute this index after omitting these variables, however, research from the fields of sociology and demography on the fluidity of racial identification indicates a severe, unresolved difficulty in accurately measuring race when the same survey is administered across racially and ethnically distinct regions, e.g. the mainland U.S. and Puerto Rico (Loveman and Muniz, 2007; Vargas-Ramos, 2015; Davenport, 2020). Notably, according to the 2010 Census, 76% of Puerto Ricans identified as 'White, alone' compared to 72% for the 50 states and the District of Columbia. In 2022, this figure was 51.1% in Puerto Rico but barely changed in the rest of the U.S.. Because of its unclear interpretation and volatility over time in Puerto Rico, the REM theme of vulnerability is omitted from the analyses in this paper.

Additional data is collected to serve as covariates in our model. Tract-level population data is from the 2017 ACS, obtained prior to Hurricane María. The coefficient on this variable will reflect the extent to which current restoration procedures restore power to the most individuals as quickly as possible, as stated in official written procedures. Local maximum wind speed comes from the National Oceanic and Atmospheric Administration's National Climatic Data Center and serves as another proxy for direct storm damage, alongside the peak number of outages experienced by a tract measured as the total number of crews deployed to that tract.

2.2 Geospatial Dependency Data

Geospatial Census tract geometries come directly from the U.S. Census Bureau. Data on major roads comes from the U.S. Geological Survey National Transportation Dataset for Puerto Rico. I define major roads as U.S. routes and Interstate highways, pictured in the top panel of Figure 2. Data on power lines comes from the Homeland Infrastructure Foundation-Level Data, a publicly-available dataset maintained by the Department of Homeland Security containing all power transmission lines in Puerto Rico with voltages ranging from 69 kV to

765 kV. These are pictured in the bottom panel of Figure 2. Each panel of the figure displays a map of all Census tracts in Puerto Rico with roads and power lines overlaid in red and blue, respectively. Tracts which intersect with either roads or power lines are shaded dark gray. These data are used to generate alternative sets of neighbors for a given census tract, as explained in Section 3.1.

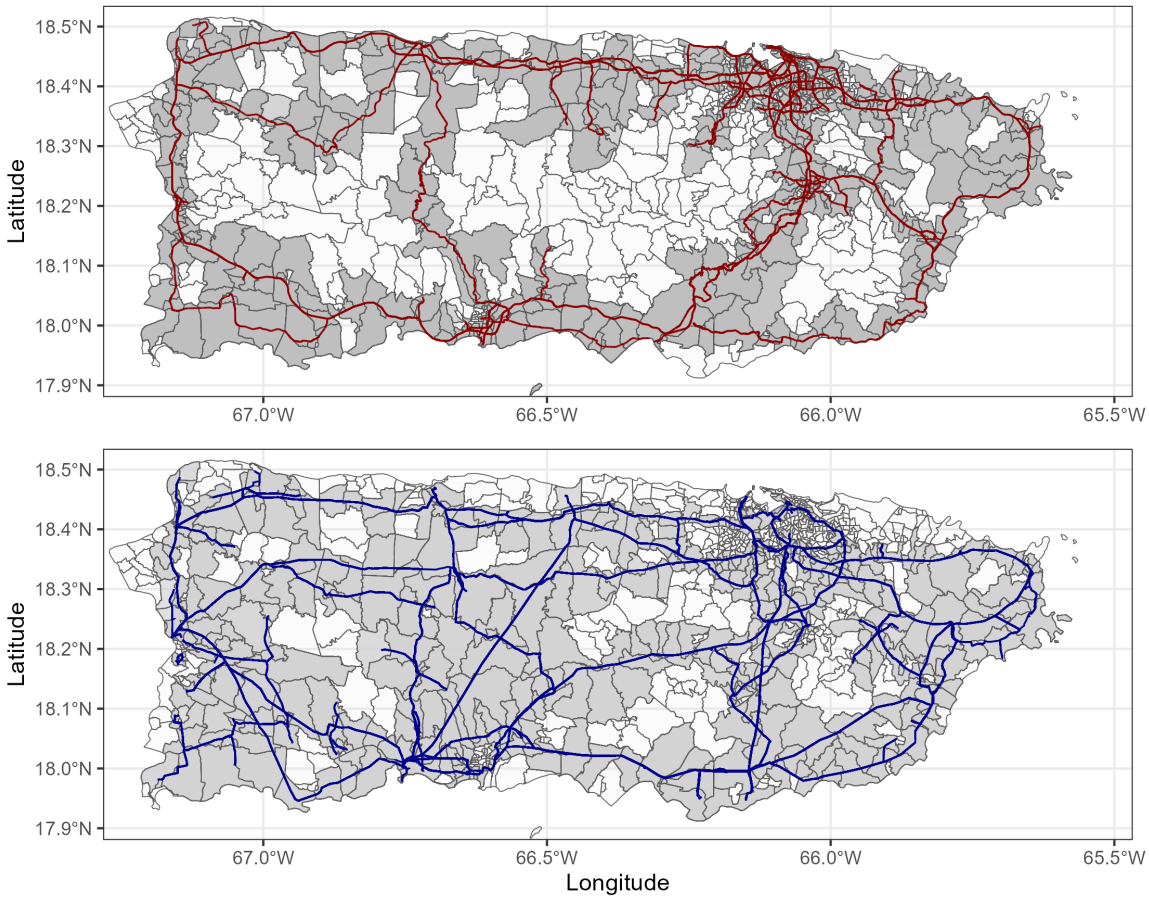

Figure 2: Puerto Rico Census tracts. Red (top) and blue (bottom) lines are roads and power lines, respectively. Tracts are shaded if they intersect with either roads or power lines.

3 Model

3.1 Spatial Weight Matrices

I now explain the role of spatial weight matrices in estimating a spatial-lag regression, with a particular focus on the different inter-tract dependencies explored in this paper. Baseline regression estimates which use the simple geographic-adjacency weighting scheme can be found in Section 4.

Estimating a spatial regression model requires the specification of a spatial weight matrix. Each row i of a $N \times N$ spatial weight matrix W contains information on tract i's neighbors, where, typically, $W_{ij} = 1$ if i and j neighbor one another, and 0 otherwise. To be considered a neighbor to i, tract j must share a common geographic border with i, or, equivalently, they must be geographically adjacent. All tracts satisfying this condition form a set of neighbors for i which is used to estimate the spatially lagged relationship between i's outcome and its neighbors' outcomes with each of these neighbors receiving equal weighting in this estimation.

Consider the set of tracts presented in Figure 3. Our focal tract, labeled A, is shaded gray while its neighbors are shaded either red, blue, or purple depending on whether or not they are connected to A via a road, power line, or both, respectively. Tract F has no shading because, while it is geographically adjacent to A, it is connected neither by a road nor a power line. Traditionally, each of these neighbors would receive equal weights of 1/5 when estimating a regression. This geographic-adjacency weighting scheme forms the weight matrix W^{adj}. However, we can define alternative schemes which reflect inter-tract road and power line dependencies. I define two such alternatives: W^{roads} and W^{lines}, which adds road and power line dependency criteria when defining a tract's neighbors. For example, only tracts C, D, and E would be considered A's neighbors in W^{roads} while only tracts B, C, and D would be considered A's neighbors in W^{lines}. In Section 4.2, I consider the space of possible weighted sums of these three weighting schemes to explore possible biases caused by using W^{adj} to estimate our model, as done in Section 4.1 and common in the literature.

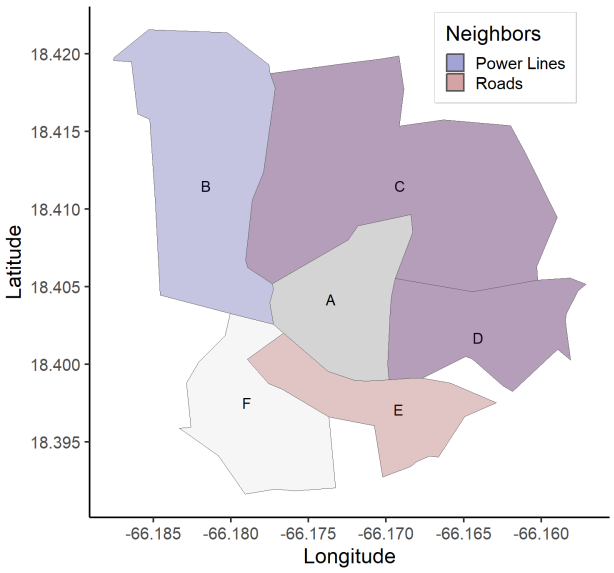

Figure 3: Focal Census tract and its neighbors as defined by geographic borders (light gray), power lines (blue), major roads (red), or all three (purple).

3.2 Spatial-Lag Model

The analyses in this paper use regression models from the field of spatial econometrics. As explained below, the primary model used is a spatial autoregressive (SAR) model, specified as follows:

$$y_i = \rho W_c \boldsymbol{y} + X_i \beta + \epsilon, \tag{1}$$

where, for a tract i, y_i is the log median outage duration and X_i is a matrix of independent variables, including the three SVI themes, the logarithm of peak outages, the logarithm of tract population, the logarithm of tract area in square-miles, and maximum wind speed. $\boldsymbol{y}$ is a vector of log median outage duration for all Census tracts. The log-transformations address the left-censored quality of these variables and is consistent with past work (e.g., Ganz et al. (2023)). The parameter ρ captures the extent of the spatial dependence with respect to the outcome variable. ϵ is the random error term.

W_c is our spatial weight matrix. For the baseline results presented below, $W_c = W_{\text{adj}}$. In later specifications, W_c is the weighted sum of the three spatial weight matrices described

above: $\boldsymbol{W}$, a vector containing matrices W_{adj}, W_{roads}, and W_{lines}, is multiplied with a corresponding vector of scalar weights $\boldsymbol{\pi}$, containing π_{adj}, π_{roads}, and π_{lines}, where each weight lies strictly between 0 and 1 and sum to unity, to obtain $W_c = \boldsymbol{\pi}\boldsymbol{W}$.

I estimate other variations of this model, as well, to show my findings are robust to model selection (see LeSage and Pace (2009) for details on these models). A spatial Durbin model (SDM) includes an additional vector of spatially lagged explanatory variables on the right hand side of the equation while a spatial error model (SEM) does not include the ρ term and instead just incorporates spatial dependence in errors, estimated as a parameter λ.

4 Results

4.1 Baseline Estimates

Table 1 contains regression estimates of several plausible models for identifying our relationship of interest. In the far right column, we see estimates of a multiple linear regression (MLR) model which does not account for possible spatial dependencies. Lagrange multiplier tests indicate that the MLR is misspecified, exhibiting both spatially autocorrelated errors and a dependence on spatial lags (see Anselin (1988)), thus I estimate three spatial econometric models which account for these sources of spatial dependence. All three spatial models—SAR, SDM, and SEM—fit the data similarly well. Though the SDM produces the largest log-likelihood statistic, pairwise likelihood ratio tests reveal this difference is not significant. I thus settle on the SAR model which includes a spatial lag and accounts for spatially correlated errors. Note that any given coefficient estimate is not sensitive to model selection.

Because precise interpretation of the coefficient estimates is complicated by the feedback loops generated by spatial spillovers, I will begin with an overview of our SAR estimates. The spatial lag coefficient ρ is positive and significant; a focal tract's outage duration increases with the outage duration of its neighbors. Consistent with previous research, I also find a

significant, positive relationship between SES vulnerability and outage duration. However, contrary to the findings presented in Tormos-Aponte et al. (2021), I find no significant relationship between the other vulnerability themes and outage duration. This is consistent, however, with Ganz et al. (2023) who apply similar methods to data from the Southwest United States. Finally, contrary to the intentions of current power restoration protocol, there is a significant, positive relationship between the population size and outage duration. As we will see below, however, this finding is not robust to alternative specifications.

Table 1: Regression model estimates using complete model specification

	Dependent variable:			
	log(outage duration)			
	SDM	SEM	SAR	MLR
SES vuln.	0.334**	0.340**	0.323**	0.351**
	(0.130)	(0.111)	(0.106)	(0.107)
HC vuln.	0.038	−0.014	−0.023	−0.032
	(0.108)	(0.098)	(0.094)	(0.095)
HTT vuln.	−0.079	−0.061	−0.059	−0.071
	(0.097)	(0.095)	(0.094)	(0.095)
log(peak outage)	0.059**	0.063**	0.063**	0.068**
	(0.027)	(0.027)	(0.027)	(0.027)
log(area)	0.035	0.017	0.009	0.013
	(0.035)	(0.025)	(0.023)	(0.023)
log(population)	0.157**	0.164**	0.164**	0.165**
	(0.075)	(0.072)	(0.071)	(0.072)
Max. wind speed	0.010	−0.001	−0.0003	−0.001
	(0.032)	(0.005)	(0.004)	(0.004)
ρ	0.147**		0.141**	
	(0.053)		(0.052)	
λ		0.148**		
		(0.053)		
Observations	777	777	777	777
Log Likelihood	-818.797	-820.064	-820.256	
LM lag test statistics				7.498**
LM error test statistics				7.845**
LR test statistic		2.534	2.919	

Note: *p $<$ 0.1 **p $<$ 0.05. Estimates of lagged explanatory variables omitted for SDM.

The common partial derivative interpretation of MLR coefficient estimates does not apply to spatial model coefficients because each focal tract is a second-order neighbor of itself. LeSage and Pace (2009) provide a method for computing average direct and indirect effects across all Census tracts which account for these feedback loops. According to the authors, the interpretation of the average direct effect is "similar in spirit" to the partial derivative interpretation of MLR coefficients. Conditional on the impact of the storm itself, a one decile increase in SES vulnerability is directly associated with a 3.3% increase in outage duration and a 3.8% increase after accounting for indirect impacts. For the median household, which experienced a 78 day outage, these equate to roughly an additional 3 days without access to power.

To demonstrate the robustness of these results, I now present a second set of SAR estimates which vary the selection criteria of Census tracts and the included covariates. These results are presented in Table 2 and can be compared directly to the SAR column in Table 1. Column 1 omits tracts found within San Juan, reducing our sample size to 665 tracts. San Juan, the capital of Puerto Rico, is the most densely populated area of the island, measured both by population and number of Census tracts. The high tract density of San Juan may mask spatial dependencies found elsewhere on the island. However, the relationship between each vulnerability theme and outage duration is not sensitive to the exclusion of San Juan: SES is still similarly positively related to outage duration and the other themes remain insignificantly different from zero.

Columns 2, 3, and 4 remove two of the vulnerability themes as covariates, leaving just SES, HC, and HTT vulnerability in the model, respectively. The relationship between SES vulnerability and outage duration weakens when other vulnerability themes are excluded from the model but remains significant and positive. Here, the average direct effect of a one decile increase in SES vulnerability is a 3.0% increase in outage duration and including the indirect effects raises this to a 3.5% increase. Columns 3 and 4 indicate that the statistically insignificant relationships between HC and HTT vulnerabilities and outage duration found

Table 2: Alternative SAR estimates of Equation (1)

	Dependent variable:			
	log(outage duration)			
	(1)	(2)	(3)	(4)
SES vulnerability	0.296**	0.295**		
	(0.117)	(0.094)		
HC vulnerability	0.096		0.080	
	(0.106)		(0.088)	
HTT vulnerability	−0.083			0.033
	(0.100)			(0.090)
log(peak outage)	0.063**	0.062**	0.059**	0.057**
	(0.028)	(0.026)	(0.027)	(0.027)
log(area)	0.023	0.014	0.038*	0.037*
	(0.026)	(0.022)	(0.021)	(0.021)
log(population)	0.204**	0.160**	0.097	0.099
	(0.077)	(0.070)	(0.068)	(0.068)
Max. wind speed	−0.001	−0.0004	−0.001	−0.001
	(0.004)	(0.004)	(0.004)	(0.004)
ρ	0.149**	0.143**	0.156**	0.156**
	(0.055)	(0.052)	(0.052)	(0.052)
Observations	665	777	777	777
Log Likelihood	-691.534	-820.487	-824.930	-825.272

Note: *p $<$ 0.1 **p $<$ 0.05.

previously remain so if they are the only vulnerability covariates included in the model. The coefficient on the log of tract population remains significant and positive in columns 1 and 2 but become insignificant in columns 3 and 4.

4.2 Alternative Spatial Weight Matrices

We can now see if my baseline impact estimates are biased by the selection of our spatial weight matrix. As described in Section 3.1, we can construct alternative weighting schemes which incorporate inter-tract dependencies on power lines and major roads. If we find that an alternative weighting scheme provides a better fit and a different coefficient estimate (either higher or lower than our baseline of 0.323), then there is likely a bias introduced by ignoring

these alternative spatial dependencies. I address this question by estimating 171 regressions, each using a unique weighted combination of our three spatial weight matrices. Weights π_{adj}, π_{roads}, and π_{lines} are selected from $[0.05, 0.90]$ with intervals of 0.05 such that they sum to unity. These weights are used to compute unique weight matrices W_c, as described in Section 3.1. I judge the fit of the model estimates obtained by these alternative weighting schemes by their log-likelihoods and use the described grid search to find the log-likelihood-maximizing weighting scheme. I repeat this exercise a second time while excluding tracts in San Juan where the high tract-density may render road and power line dependencies inconsequential. In Appendix A, I repeat this exercise yet again on a series of simulated datasets to demonstrate its usefulness in detecting optimal weighting schemes when the underlying data-generating dependencies change.

The two sets of results are presented in Figure 4. Each ternary plot displays the complete range of log-likelihoods from each weighting scheme. π_{lines} is measured on the bottom edge, increasing from left to right. π_{roads} is measured on the right edge, increasing from bottom to top. π_{adj} is measured on the left edge, increasing from top to bottom. A black dot is placed where the set of weights is log-likelihood-maximizing. Darker (lighter) shading indicates better (worse) fits. The left plot uses the full tract-level dataset of Puerto Rico and the right plot excludes San Juan. These two plots indicate the best-fitting estimates are produced by schemes which weight each geographically adjacent tract similarly, regardless of road and power line dependencies. This best-fitting model estimate uses $\pi_{\text{adj}} = 0.9$ and $\pi_{\text{roads}} = \pi_{\text{lines}} = 0.05$ and, when using all tracts, has a log-likelihood of -820.307, indicating a worse fit than when using $W_c = W_{\text{adj}}$ as shown in Table 1. In both datasets, all model estimates obtained from weighting schemes which emphasize either road or power line dependencies increase the coefficient on SES vulnerability but worsen the fit of the model. This suggests that these alternative dependencies do not contribute meaningfully to the recovery process and that there are likely yet more alternative, unexplored spatial dependencies which can improve model fit.

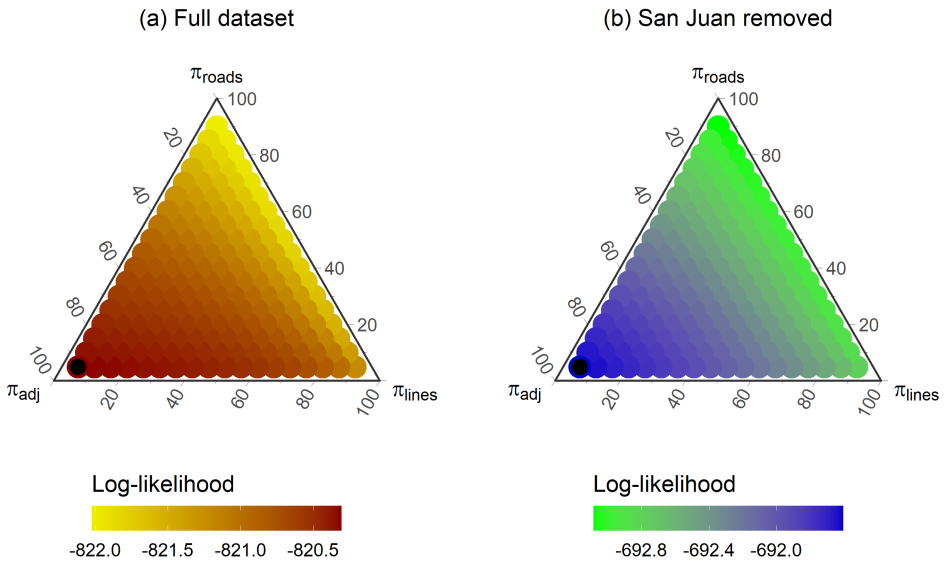

Figure 4: Log-likelihoods of regression estimates using various weighting schemes. The coordinates reflect different coefficients on weight matrices W_{adj}, W_{roads}, and W_{lines}. Plots (a) and (b) use outage data from all of Puerto Rico and just tracts which lie outside of San Juan, respectively.

To confirm the result implied by the grid search, I use a genetic optimization algorithm as part of the R package *rgenoud* to find the log-maximizing weight vector (Mebane, Jr. and Sekhon, 2011). Let $\text{LogLik}(M(\boldsymbol{\pi}))$ be the log-likelihood function applied to a model estimate M of Equation (1) which uses the weighting scheme $\boldsymbol{\pi}$ to compute the spatial weight matrix $W_c = \boldsymbol{\pi}\boldsymbol{W}$. I then solve the following maximization problem: $\max(\text{LogLik}(M(\boldsymbol{\pi})))$ subject to the restriction that each weight in $\boldsymbol{\pi}$ lies strictly between 0 and 1. I do this for both the full set of counties and again after excluding San Juan. The solution obtained in both cases indicates the best fitting model uses $\pi_{\text{adj}} \approx 1$, consistent with the grid-search results presented in Figure 4. This, in combination with the grid search, shows the unique $\boldsymbol{\pi}$ which maximizes the log-likelihood of the model is that which puts full weight on π_{adj}.

5 Discussion

This paper contributes to our understanding of the disparate outage durations experienced by communities of different SES vulnerability after Hurricane María hit Puerto Rico in 2017. Unlike previous empirical efforts in this setting, I use spatial regression methods which account for both the direct impact of the storm and spatial dependencies between neighboring areas to estimate the relationship between SES vulnerability and outage duration conditional on the direct impact of the storm itself. Consistent with evidence from elsewhere in the United States, I find a one decile increase in SES vulnerability produces a 3.8% increase in outage duration. I also find this estimate is robust to several alternative specifications. These results suggest current recovery procedures exacerbate outage inequalities caused by hurricanes across communities of different SES vulnerability levels.

Outage recovery procedures typically prioritize high-density areas in order to restore power to the most households as quickly as possible. In doing so, however, these restoration efforts appear to extend the outage durations experienced by communities of higher levels of SES vulnerability. This growing literature suggests policymakers must reassess current procedures. Shifting recovery priorities to communities which are more vulnerable is a possible solution but has the likely tradeoff of increasing overall outage durations (Ganz et al., 2023). Given that poorer households are typically less prepared for hurricanes and less able to respond to hurricanes after they have hit (Afsharinejad et al., 2021; Hong et al., 2021; Flanagan et al., 2018), this may be preferable from a total cost perspective but remains an open question for future research. Alternative solutions include improving the existing infrastructure in vulnerable communities to mitigate the direct effects of the hurricanes themselves (Brockway et al., 2021).

This paper also provides something of a robustness check for past and future research in this area which utilizes spatial regression methods. Past research in this area uses a spatial weighting scheme based purely on the geographic adjacency between a focal unit of observation and surrounding observations. The results in this paper suggest this weighting

scheme does not produce a biased estimate of the relationship between vulnerability and outage duration when compared to alternative weighting schemes which incorporate power line and road dependencies. This does not, however, mean that such a weighting scheme is optimal for such estimation efforts. Future research should explore yet more alternative spatial weighting schemes to improve model fit. In addition, the same weighting schemes tested in this paper can and should be tested in other geographic settings.

Hurricanes and other major weather events will likely increase in both probability and severity as a result of continuing changes to Earth's climate (Seneviratne et al., 2021; Emanuel, 2021). In turn, the number of proposals to reform power restoration procedures and empirical studies of existing procedures will increase, as well. This study serves to improve the quality of both.

References

Afsharinejad, A. H., Ji, C., and Wilcox, R. (2021). Large-scale data analytics for resilient recovery services from power failures. *Joule*, 5(9):2504–2520. Publisher: Elsevier.

Anselin, L. (1988). Lagrange Multiplier Test Diagnostics for Spatial Dependence and Spatial Heterogeneity. *Geographical Analysis*, 20(1):1–17. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4632.1988.tb00159.x.

Best, K., Kerr, S., Reilly, A., Patwardhan, A., Niemeier, D., and Guikema, S. (2023). Spatial regression identifies socioeconomic inequality in multi-stage power outage recovery after Hurricane Isaac. *Natural Hazards*, 117(1):851–873.

Brockway, A. M., Conde, J., and Callaway, D. (2021). Inequitable access to distributed energy resources due to grid infrastructure limits in California. *Nature Energy*, 6(9):892–903. Publisher: Nature Publishing Group UK London.

Coleman, N., Esmalian, A., Lee, C.-C., Gonzales, E., Koirala, P., and Mostafavi, A. (2023). Energy inequality in climate hazards: empirical evidence of social and spatial disparities in managed and hazard-induced power outages. *Sustainable Cities and Society*, 92:104491. Publisher: Elsevier.

Davenport, L. (2020). The Fluidity of Racial Classifications. *Annual Review of Political Science*, 23(1):221–240. _eprint: https://doi.org/10.1146/annurev-polisci-060418-042801.

Do, V., McBrien, H., Flores, N. M., Northrop, A. J., Schlegelmilch, J., Kiang, M. V., and Casey, J. A. (2023). Spatiotemporal distribution of power outages with climate events and social vulnerability in the USA. *Nature Communications*, 14(1):2470.

Edison Electric Institute (2016). Understanding the electric power industry's response and restoration process. Technical report, Edison Electric Institute.

Emanuel, K. (2021). Atlantic tropical cyclones downscaled from climate reanalyses show

increasing activity over past 150 years. *Nature communications*, 12(1):1–8. Publisher: Nature Publishing Group.

Fischbach, J., Warren May, L., Whipkey, K., Shelton, S., Vaughan, C., Tierney, D., Leuschner, K., Meredith, L., and Peterson, H. (2020). *After Hurricane Maria: Predisaster Conditions, Hurricane Damage, and Recovery Needs in Puerto Rico*. RAND Corporation.

Flanagan, B., Hallisey, E. J., Adams, E., and Lavery, A. (2018). Measuring community vulnerability to natural and anthropogenic hazards: the Centers for Disease Control and Prevention's social vulnerability index. *Journal of Environmental Health*, 80(10):34–36.

Flores, N. M., McBrien, H., Do, V., Kiang, M. V., Schlegelmilch, J., and Casey, J. A. (2023). The 2021 Texas power crisis: distribution, duration, and disparities. *Journal of Exposure Science & Environmental Epidemiology*, 33(1):21–31.

Ganz, S. C., Duan, C., and Ji, C. (2023). Socioeconomic vulnerability and differential impact of severe weather-induced power outages. *PNAS Nexus*, 2(10):pgad295.

Hong, B., Bonczak, B. J., Gupta, A., and Kontokosta, C. E. (2021). Measuring inequality in community resilience to natural disasters using large-scale mobility data. *Nature Communications*, 12(1):1870.

Kishore, N., Marqués, D., Mahmud, A., Kiang, M. V., Rodriguez, I., Fuller, A., Ebner, P., Sorensen, C., Racy, F., Lemery, J., Maas, L., Leaning, J., Irizarry, R. A., Balsari, S., and Buckee, C. O. (2018). Mortality in Puerto Rico after Hurricane Maria. *New England Journal of Medicine*, 379(2):162–170. Publisher: Massachusetts Medical Society _eprint: https://doi.org/10.1056/NEJMsa1803972.

Kwasinski, A., Andrade, F., Castro-Sitiriche, M. J., and O'Neill-Carrillo, E. (2019). Hurricane Maria Effects on Puerto Rico Electric Power Infrastructure. *IEEE Power and Energy Technology Systems Journal*, 6(1):85–94. Conference Name: IEEE Power and Energy Technology Systems Journal.

LeSage, J. P. and Pace, R. K. (2009). *Introduction to spatial econometrics*. Statistics, textbooks and monographs. CRC Press, Boca Raton.

Lievanos, R. S. and Horne, C. (2017). Unequal resilience: the duration of electricity outages. *Energy Policy*, 108:201–211.

Logan, J. R. and Xu, Z. (2015). VULNERABILITY TO HURRICANE DAMAGE ON THE U.S. GULF COAST SINCE 1950. *Geographical Review*, 105(2):133–155.

Loveman, M. and Muniz, J. O. (2007). How Puerto Rico Became White: Boundary Dynamics and Intercensus Racial Reclassification. *American Sociological Review*, 72(6):915–939. Publisher: SAGE Publications Inc.

Mebane, Jr., W. R. and Sekhon, J. S. (2011). Genetic optimization using derivatives: The rgenoud package for R. *Journal of Statistical Software*, 42(11):1–26.

Mitsova, D., Esnard, A.-M., Sapat, A., and Lai, B. S. (2018). Socioeconomic vulnerability and electric power restoration timelines in Florida: the case of Hurricane Irma. *Natural Hazards*, 94(2):689–709.

Roman, M. O., Stokes, E. C., Shrestha, R., Wang, Z., Schultz, L., Carlo, E. A. S., Sun, Q., Bell, J., Molthan, A., Kalb, V., Ji, C., Seto, K. C., McClain, S. N., and Enenkel, M. (2019). Satellite-based assessment of electricity restoration efforts in Puerto Rico after Hurricane Maria. *PLOS ONE*, 14(6):e0218883.

Seneviratne, S., Zhang, X., Adnan, M., Badi, W., Dereczynski, C., Di Luca, A., Ghosh, S., Iskandar, I., Kossin, J., Lewis, S., Otto, F., Pinto, I., Satoh, M., Vicente-Serrano, S.,

Wehner, M., and Zhou, B. (2021). Weather and climate extreme events in a changing climate. In Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S., Péan, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M., Huang, M., Leitzell, K., Lonnoy, E., Matthews, J., Maycock, T., Waterfield, T., Yelekçi, O., Yu, R., and Zhou, B., editors, *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, pages 1513–1766. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA. Type: Book Section.

Tormos-Aponte, F., García-López, G., and Painter, M. A. (2021). Energy inequality and clientelism in the wake of disasters: from colorblind to affirmative power restoration. *Energy Policy*, 158:112550.

U.S. Department of Energy (2018). Puerto Rico Grid Restoration.

Vargas-Ramos, C. (2015). Migrating race: migration and racial identification among Puerto Ricans. In *Race, Migration and Identity*. Routledge. Num Pages: 21.

A Appendix

In this Appendix, I repeat the exercise found in Section 4.2 using several simulated datasets constructed using various weights on geographic adjacency, power line, and road dependencies. This serves to demonstrate that alternative spatial weighting schemes used to estimate our model can plausibly improve the fit of our estimate if they reflect data-generating spatial dependencies. I will first explain how I constructed this simulated data then present the results.

Let I be the subset of Puerto Rican Census tracts found in the power outage data ($|I| = 777$) and let J_i be the set of geographically adjacent neighbor tracts for a focal tract i. The spatial weight matrices W^{adj}, W^{roads}, and W^{lines} are those described and used elsewhere in the paper. I first assign a shock $\theta_i \sim N(0, 1)$ for each $i \in I$ and set our simulated outcome $y_i = \theta_i$ in time period 0. For each tract in each time period $t > 0$, I add a weighted sum of neighboring outcomes in the following way to form our simulated outcome variable y_i:

$$y_{i,t} = y_{i,t-1} + \sum_{j \in J_i} y_{j,t} w_{ij},$$
$$\text{where } w_{ij} = \gamma_{\text{adj}} W_{ij}^{\text{adj}} + \gamma_{\text{roads}} W_{ij}^{\text{roads}} + \gamma_{\text{lines}} W_{ij}^{\text{lines}}$$

I repeat this process $T = 10$ times to allow several iterations of spatial spillovers to occur between neighbors. The selection of γ parameters determines the relative dependence of a focal tract's outcome on those of its neighbors depending on the extent of their spatial dependencies. The γ values are strictly between 0 and 1 and sum to unity. If, e.g., γ_{roads} is

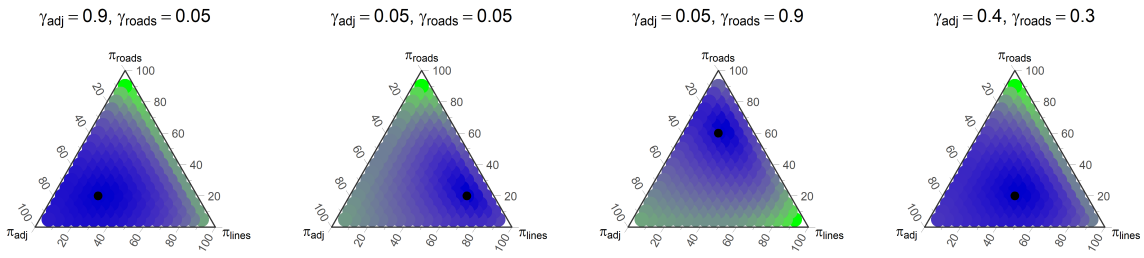

Figure A1: Ternary plots of model log-likelihoods using simulated spatially-lagged data. Each model is labeled by the weights used to compute the outcome variable. Dark blue indicates strong fits while green indicates weaker fits. The black dots indicate the best fitting weighting scheme.

very high, we would expect the log-likelihood-maximizing spatial weight matrix to have a high π_{roads} value. Below I use the following weighting schemes to generate four distinct simulated datasets: $\gamma_{\text{adj}} = 0.9, \gamma_{\text{roads}} = 0.05$; $\gamma_{\text{adj}} = 0.05, \gamma_{\text{roads}} = 0.05$; $\gamma_{\text{adj}} = 0.05, \gamma_{\text{roads}} = 0.9$; and, $\gamma_{\text{adj}} = 0.4, \gamma_{\text{roads}} = 0.3$. The explanatory variables used to predict y_i are randomly selected from a standard normal distribution. Note that this exercise is insufficient to estimate meaningfully different SDMs since there is no spatial lag dependence in explanatory variables. Such data could be constructed using similar methods used here but is unnecessary for this paper.

As before, I grid search the space of possible spatial weighting schemes. Figure A1 displays a panel of ternary plots displaying log-likelihoods under different spatial weighting schemes for each set of γ values, where darker shades indicate better fits of the data. The black dots indicate the best fitting model. The left panel treats all geographically adjacent neighbor similarly. As expected, the best fitting weighting schemes put the greatest weight on W^{adj}. The second and third plots weight power-line-adjacent neighbors and road-adjacent neighbors most heavily, respectively. Again, the best-fitting models use schemes which put greatest weight on W^{lines} and W^{roads}, respectively. The final plot uses data defined by $\gamma_{\text{adj}} = 0.4$, $\gamma_{\text{roads}} = 0.3$, and $\gamma_{\text{lines}} = 0.3$. The best-fitting model here uses $\pi_{\text{adj}} = 0.4$, $\pi_{\text{roads}} = 0.2$, and $\pi_{\text{lines}} = 0.4$, i.e. somewhere in between the previous three.

The results here suggest that if there were indeed underlying dependencies on power lines and/or roads in the power restoration process, the log-likelihood-reliant methods used in this paper would pick up on them. While the data-generating and log-likelihood-maximizing weighting schemes are not equivalent, they appear to be correlated. This supplementary exercise adds validity to the claim that estimates obtained using $W_c = W_{\text{adj}}$ are not misspecified in this way. A more extensive exploration of simulated data could reveal the required extent of dependencies on roads and power lines in the data generating process for *ex post* estimates to detect meaningful differences between alternative weighting schemes.