

DEUTSCHES ELEKTRONEN-SYNCHROTRON
Ein Forschungszentrum der Helmholtz-Gemeinschaft



DESY 21-150
HU-EP-21/33
arXiv:2201.07026
January 2022

Socioeconomic Disparities and COVID-19: The Causal Connections

T. Banerjee

Department of Economics, Auburn University, USA

A. Paul

Deutsches Elektronen-Synchrotron DESY, Hamburg

and

Institut für Physik, Humboldt-Universität, Berlin

V. Srikanth

BASIS Independent Silicon Valley, San Jose, USA

and

Stanford Online High School, Stanford, USA

I. Strümke

Department of Engineering Cybernetics, NTNU,

Trondheim, Norway

and

Department of Holistic Systems, SimulaMet, Oslo, Norway

ISSN 0418-9833

NOTKESTRASSE 85 - 22607 HAMBURG

DESY behält sich alle Rechte für den Fall der Schutzrechtserteilung und für die wirtschaftliche Verwertung der in diesem Bericht enthaltenen Informationen vor.

DESY reserves all rights for commercial use of information included in this report, especially in case of filing application for or grant of patents.

To be sure that your reports and preprints are promptly included in the
HEP literature database
send them to (if possible by air mail):

DESY Zentralbibliothek Notkestraße 85 22607 Hamburg Germany	DESY Bibliothek Platanenallee 6 15738 Zeuthen Germany
---	---

Socioeconomic disparities and COVID-19: the causal connections

Tannista Banerjee^{a,*}, Ayan Paul^{b,c,*}, Vishak Srikanth^{d,e,*}, and Inga Strümke^{f,g,*}

^aDepartment of Economics, Auburn University, 140 Miller Hall, Auburn, AL, 36849, USA

^bDESY, Notkestraße 85, D-22607 Hamburg, Germany

^cInstitut für Physik, Humboldt-Universität zu Berlin, D-12489 Berlin, Germany

^dBASIS Independent Silicon Valley, San Jose, CA

^eStanford Online High School, Stanford, CA

^fDepartment of Engineering Cybernetics, NTNU, 7034 Trondheim, Norway

^gDepartment of Holistic Systems, SimulaMet, 0167 Oslo, Norway

*e-mail: tz0018@auburn.edu ayan.paul@desy.de, ivishak@gmail.com, inga.strumke@ntnu.no,

ABSTRACT

The analysis of causation is a challenging task that can be approached in various ways. With the increasing use of machine learning based models in computational socioeconomics, explaining these models while taking causal connections into account is a necessity. In this work, we advocate the use of an explanatory framework from cooperative game theory augmented with *do* calculus, namely causal Shapley values. Using causal Shapley values, we analyze socioeconomic disparities that have a causal link to the spread of COVID-19 in the USA. We study several phases of the disease spread to show how the causal connections change over time. We perform a causal analysis using random effects models and discuss the correspondence between the two methods to verify our results. We show the distinct advantages a non-linear machine learning models have over linear models when performing a multivariate analysis, especially since the machine learning models can map out non-linear correlations in the data. In addition, the causal Shapley values allow for including the causal structure in the variable importance computed for the machine learning model.

Introduction

The early stages of the spread of COVID-19 in the USA laid bare how socioeconomic disparities bring about a disproportionate spread of the disease in certain parts of society [1–14]. This trend is a pattern that has been historically observed for several diseases caused by viruses like HIV [15], MERS-CoV, SARS-CoV, Ebola [16–18] etc. Even beyond variances in socioeconomic conditions, ethnicity is a contributing factor [19–22] for disparities in healthcare. A good understanding of the causal connection between socioeconomic disparities and the spread of a disease is necessary for implementing policies to mitigate disease spread amongst those who are the most vulnerable.

In a recent work, it was shown that the spread of COVID-19 is correlated to various socioeconomic metrics averaged at the county level [23] during the initial stages of the pandemic. The use of Shapley values was introduced with a machine learning framework to understand how socioeconomic disparities affect the spread of the disease. However, Shapley values decomposing a machine learning model probe the correlations between the endogenous and exogenous variables and do not take the causal connection in the data into account. In this work, we go a step further and introduce causal Shapley values to the analysis of disease spread.

We obtain a regression model by training a machine learning model in a supervised manner on COVID-19 prevalence data. As the resulting model is too complex to be considered interpretable, we use methods from the field of explainable AI to explain the model’s predictions. Specifically, we use a framework built on the game-theoretic solution concept of Shapley values [24], upon which the packages of Shapley additive explanations (SHAP) [25] and shapr [26] are based. The idea behind these was recently built upon by Ref. [27] to avoid independence assumptions in the data as well as to incorporate causal knowledge from causal chain graphs in the Shapley value calculation. Our objectives in this work can be delineated as:

- We study the spread of COVID-19 in three different regions of the USA broken down by the pattern of spread of the disease.
- We build socioeconomic metrics using data sources from the US Census Bureau at the county level to model the demographic distribution of the population.
- We use an ensemble of boosted decisions trees (BDTs) to build a regression model of the data.
- We calculate the causal Shapley values using these regression models to understand the causal connections between the socioeconomic metrics and the spread of COVID-19.
- We study two different times, the early stage the disease spread between February and July 2020 and the later stage of the disease spread between July 2020 and January 2021 to understand the variance in the causal connections.
- We compare the multivariate methods that we establish here with linear random effects models and discuss the relative advantage of using machine learning for highly correlated datasets.

To avoid the effects of causation being washed out due to large scale averaging [23], we divide our analysis into three regions. This is particularly important since the demographic nature of various parts of the USA are different. In addition, the pattern of disease spread were also observed to vary in different parts of the USA. The states with higher densities were significantly affected in the first wave of the pandemic while the southern states were significantly affected by the second wave. The states on the west coast were affected throughout the extent of the pandemic. Demographically, the northeastern states have several urban areas that are closely clustered together. However, in the southern states the urban areas are spread far apart and a large fraction of the population lives in rural areas where the socioeconomic conditions are different from the urban areas. The states on the west coast are a mix of urban and rural areas with a very different economic structure. Hence, we will focus on three different clusters of states.

- **High population density regions:** States with population density over 400 individuals per sq. km: District of Columbia, New Jersey, Rhode Island, Massachusetts, Connecticut, Maryland, Delaware and New York.
- **The southern states:** These include: Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia.
- **The west coast:** These include: California, Oregon and Washington.

We hope that by highlighting how the socioeconomic disparities that aggravated the spread of COVID-19 in of the community, and often amongst those who are at a social disadvantage, we can not only aid in future policy building but also establish a robust method for causal analysis that can be easily used for highly multivariate datasets.

Causal ordering

The framework advocated by Heskens et al. [27] for causal Shapley values accounts for indirect effects in order to take the causal structure of the data into account in when estimating the Shapley values. The approach recognizes that it is not practical to compute interventional probabilities to account for indirect effects and instead proposes a causal chain graph. The (partial) causal ordering in the data is represented as a causal chain graph by a directed acyclic graph (DAG). A causal ordering (CO) is then a nested list that maps to a sequence of the respective causal orderings of the socioeconomic metrics. Variables of equivalent causal importance are grouped into the same sub-list within the list. For example, [NW, [SC, Emp]] would represent non-white as the most significant cause which influences the next set of variables, here Senior Citizen and Employed, each with equal importance¹. More intuitively, the causal ordering depends on a hypothesis of how the socioeconomic metrics form subgroups that are causally connected. Within a subgroup, the metrics are assumed to have no causal connection but can be cyclically connected or share

¹The abbreviations and definitions of the metrics can be found in the Materials and Methods section

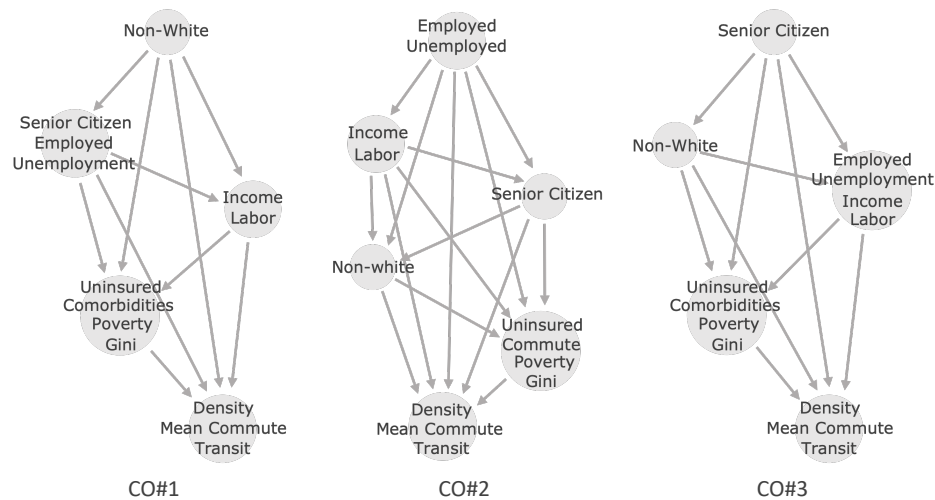


Figure 1. The three primary partial causal ordering that we consider in this work.

confounding variables. There is no fixed prescription for constructing these partial causal orderings. Rather, domain knowledge has to be used to assume how each metric can affect the others to build causal orderings. We investigate six different causal orderings, CO#1 - CO#3 described below and CO#4 - CO#6 included in the Supplementary Materials, and explain the rationale behind each choice.

CO#1: The first causal ordering we construct is [NW, [SC, Emp, Uemp], [Inc, Lab], [Uins, Com, Pov, GI],[Den, MC, Tran]]. The rationale for this baseline ordering is as follows.

- Areas with a higher proportion of non-white population tend to have higher unemployment rates and a lower proportion of employed workers [28].
- Areas with higher unemployment rates may have higher proportions of senior citizens if there is low economic mobility and younger generations move out of the area [29].
- Areas with higher unemployment and lower employment tend to have more people working in construction, service, delivery or production or are rust belt areas that have a declining manufacturing industry [30,31].
- Areas with higher unemployment tend to be poorer and have lower incomes. [32]
- Areas with more senior citizens who are on medicare or those with higher incomes individuals who receive insurance from their employer tend to have lesser uninsured fraction of the population [33,34].
- Counties with higher average incomes tend to have more income inequality and a higher Gini index [35,36].
- Areas with a higher percent of people working in construction, service, delivery, or production may have more poverty [37].
- Counties with more poverty are likely to have less single family homes and more areas zoned for multi-family homes or are likely to contain inner cities and thus have a greater density [38,39]. In such dense areas, there is likely to be shorter commutes and better transit due to these areas being more populous [40,41].

CO#2: The second causal ordering is [[Emp, Uemp], [Inc, Lab], SC, NW, [Uins, Com, Pov, GI], [Den, MC, Tran]]. This causal ordering has unemployment, employed, income per capita, and fraction working in manufacturing or manual labor as causing the proportion of senior citizens and non-white fraction of the population. This causal ordering takes into account that senior citizens are more likely to stay in areas with a declining manufacturing industry and little economic mobility [29]. Since senior citizens are less likely to be non-white, a higher fraction of senior citizens causes there to be a lower proportion of non-whites [42].

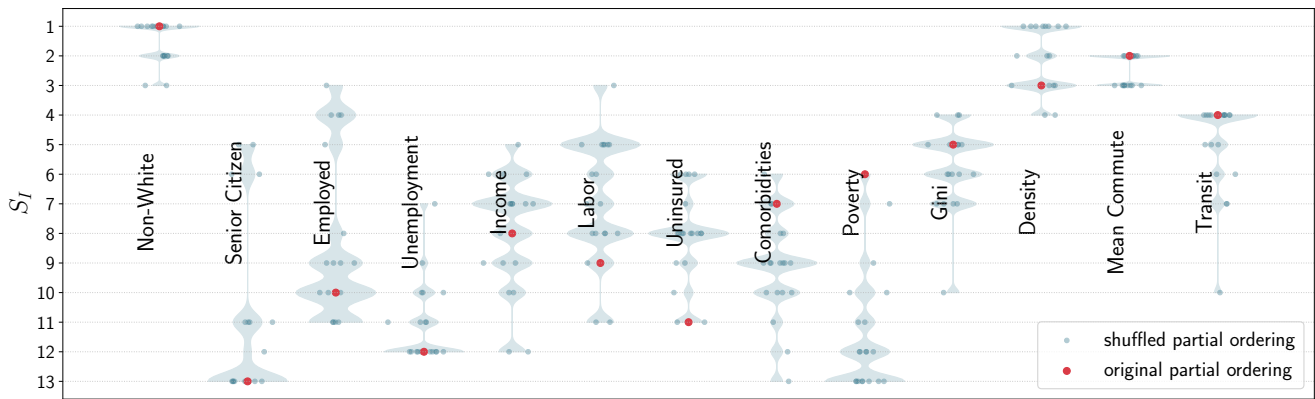


Figure 2. The Shapley Index, S_I , of all the metrics from 20 different causal ordering where the ordering of the subgroups were randomized starting from the causal ordering used in the left panel. The plot shows that the variation in S_I is small for the most important metrics with high $|\overline{S}_v|$ while larger variations in S_I can be seen in the less important metrics depending on the causal ordering. For more detail please see the text.

CO#3: The third causal ordering is [SC, NW, [Emp, Uemp, Inc, Lab], [Uins, Com, Pov, GI], [Den, MC, Tran]]. This causal ordering has the proportion of senior citizens causing the proportion of non-whites. Since senior citizens are less likely to be non-white, a higher fraction of senior citizens causes there to be a lower proportion of non-whites [42]. This causal ordering also has income per capita, the proportion of people in a county working in professions classified as Labor, unemployment, and employment on equal footing. This is the case since manual labor jobs are less temporary and pay lower wages [43].

Evidently, the decision on which causal ordering to choose is quite subjective and requires some prior knowledge of how the various metrics interact with each other. In the case of the problem we are addressing, the causal ordering can even change from region to region depending on the nature of the underlying socioeconomic fabric that governs the social structure. Thus it is important to check how robust the conclusion of our work is vis-à-vis a perturbation in the causal ordering. Ideally, the Shapley Index, S_I^2 , should depend not only on the data but also on the causal ordering. However, S_I should not be agnostic of either the data or the causal ordering since, after all, S_I should encapsulate the information from the causal ordering *and* the data.

In Figure 2, we show the result of shuffling the causal ordering CO#1. We randomly rearranging the subgroups into 20 different permutations. The violin plots show the variations in S_I for each metric. The red dot corresponds to the S_I for CO#1. We see that the three most important metric—non-white, density, and mean commute—do not show large variations in S_I . The metrics that are less causally connected to the confirmed case rate, however, show large variation. Nevertheless, there is a clear distinction between the three most important ones and the less important ones. This result shows that while S_I does depend on the causal ordering, the most important metrics are tied to the information that is gleaned from the data. If we restrict ourselves to picking the metric with the highest S_I , and hence, the largest causal connection to the outcome, small perturbations in the causal ordering will not alter our results. This observation displays the robustness of our approach making it not as subjective as it might appear to be on the surface.

Results

The question that we would like to pose is: are there any casual connection between the socioeconomic metrics and the spread of the disease and is this causal connection stable over time and geographies?

We will explore two slices of time in this section. The first one is the period when the disease had just started spreading through the population. In the beginning of 2020, COVID-19 affected the most densely populated parts of

²Shapley Index is defined in the Materials and Methods section.

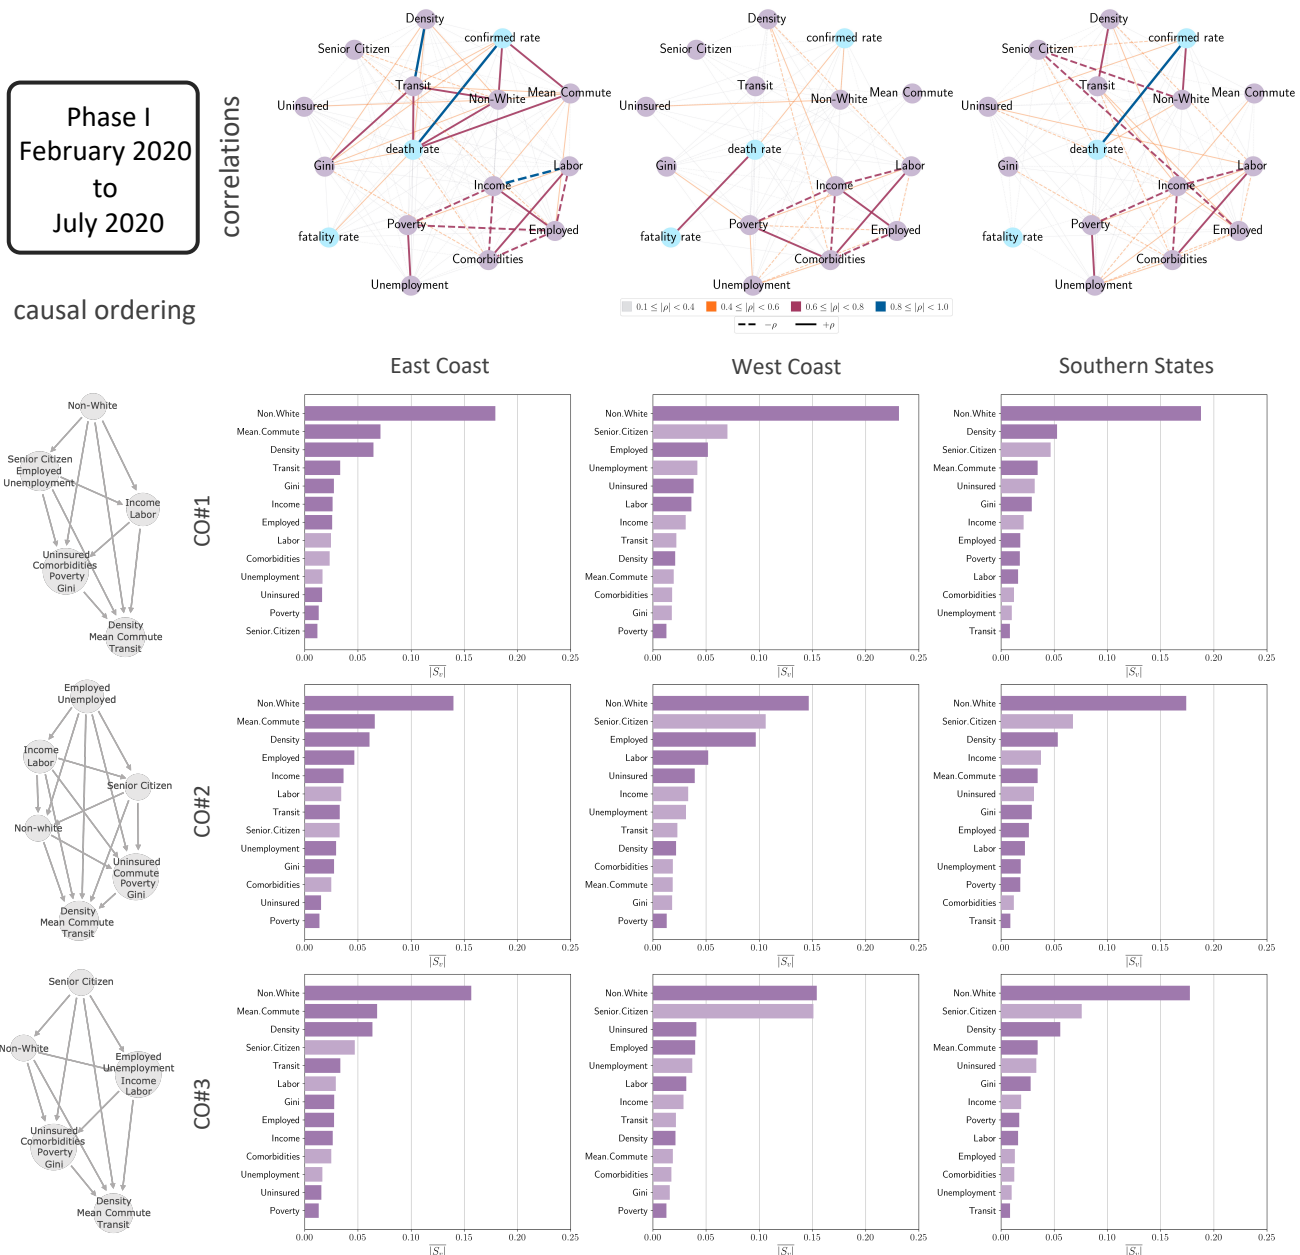


Figure 3. The causal connection between the socioeconomic metrics that we consider in this work and confirmed case rates in various regions of the USA from February 2020 to July 2020. The states included in the East Coast region are District of Columbia, New Jersey, Rhode Island, Massachusetts, Connecticut, Maryland, Delaware and New York. The ones in the West Coast region are California, Oregon and Washington. The states in the Southern States region are Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia. The networks in the top row show the correlations between the different metrics in the three regions. The graphs in the left column of the plot show the three partial causal orderings that we consider in this analysis. Each panel of the array of bar plots show the hierarchy of the metrics in terms of mean absolute causal SHAP values, $|\overline{S}_v|$, for a particular region and causal ordering. The lighter shaded bars indicate a negative effect of the metric on the confirmed case rates while the darker bars indicate a positive effect.

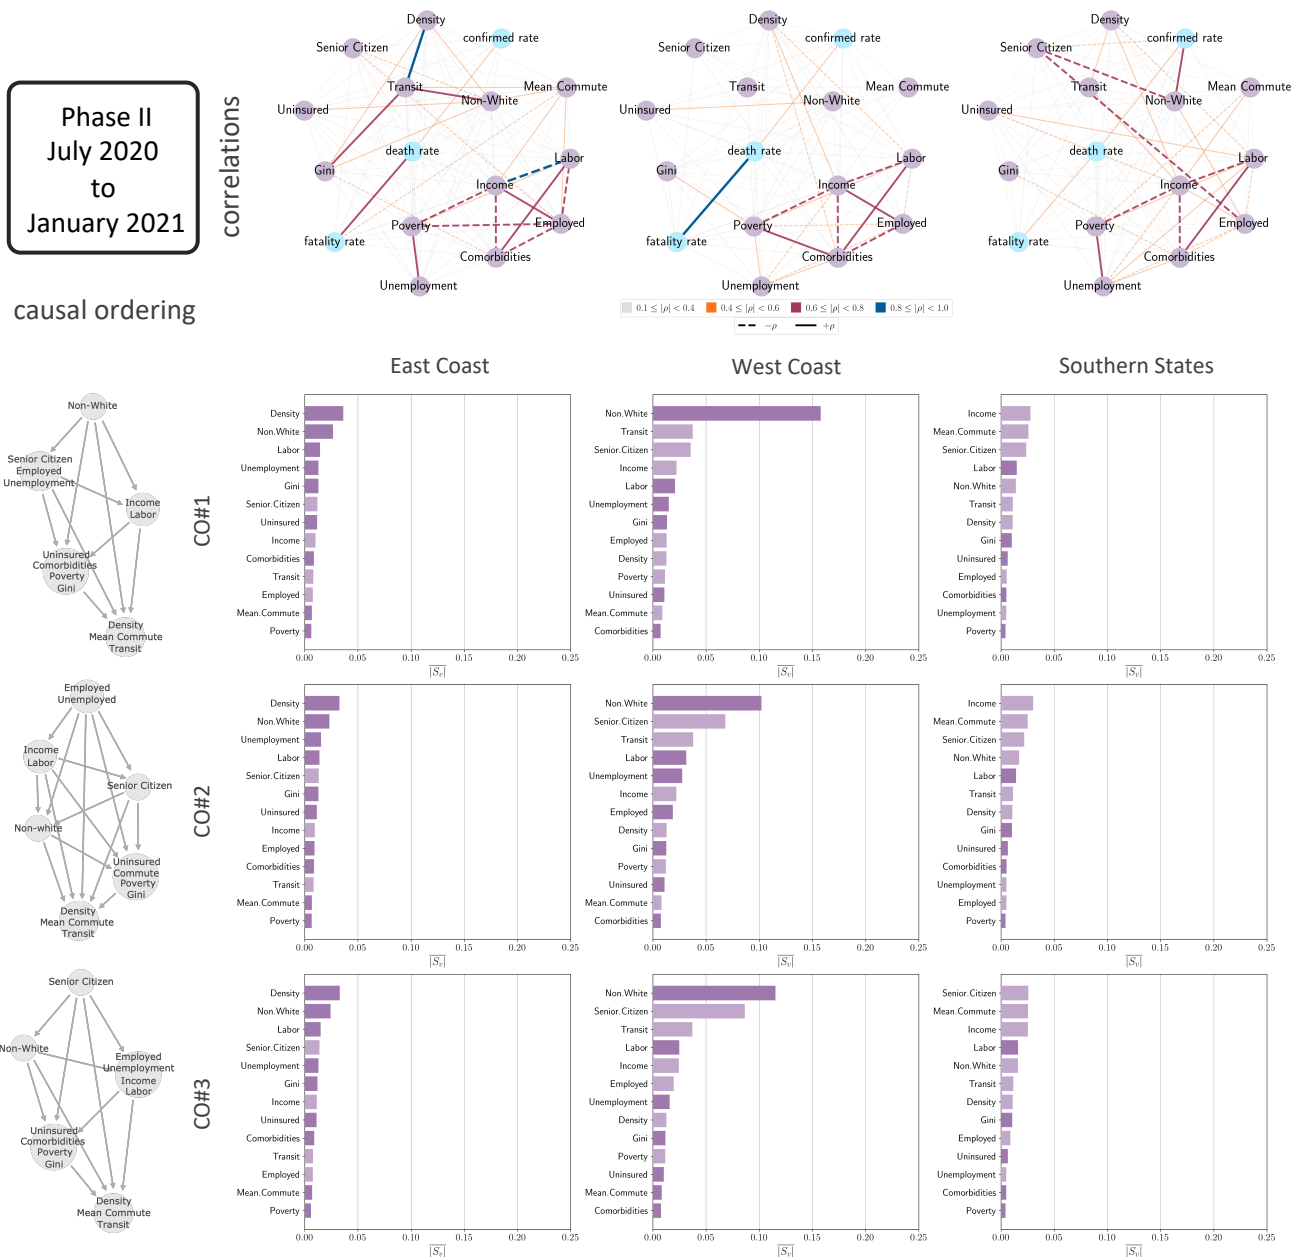


Figure 4. The causal connection between the socioeconomic metrics that we consider in this work and confirmed case rates in various regions of the USA for July 2020 to January 2021. The states included in the East Coast region are District of Columbia, New Jersey, Rhode Island, Massachusetts, Connecticut, Maryland, Delaware and New York. The ones in the West Coast region are California, Oregon and Washington. The states in the Southern States region are Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia. The networks in the top row show the correlations between the different metrics in the three regions. The graphs in the left column of the plot show the three partial causal orderings that we consider in this analysis. Each panel of the array of bar plots show the hierarchy of the metrics in terms of mean absolute causal SHAP values, $|S_v|$, for a particular region and causal ordering. The lighter shaded bars indicate a negative effect of the metric on the confirmed case rates while the darker bars indicate a positive effect.

the USA, namely, the states in the east coast and some states in the west coast. However other parts of the USA, especially the rural parts of the USA, remained unaffected by large. It was the second wave in mid-2020 that affected the parts that were spared by the first wave with vast majority of the southern states seeing a surge in the number of infected individual. So a study of socioeconomic conditions that might have affected the differential advent of COVID-19 across the social strata should include the months from February 2020 to July 2020. We shall refer to this as Phase I of our study. Consequently, we shall deem Phase II as the next phase of the disease spread from July 2020 to January 2021 keeping it safely before when vaccinations could possibly have any effects. We will contrast Phase I against Phase II to understand the role socioeconomic disparities played in the spread of COVID-19.

In [Figure 3](#) we show the correlations and causation between the socioeconomic factors, population density and the confirmed case rate during Phase I of the disease spread. The network plots in the top panel display the linear correlations between the exogenous and endogenous variable(s). From the correlations in the plots being distinct for each region, one can expect the causal connections between the variables should be different too. This is indeed confirmed by the bar plots, the columns corresponding to the different regions and the rows corresponding to the different causal ordering. A very clear pattern emerges independent of the causal ordering. In the East Coast region during Phase I, the primary causation is driven by the fraction of non-white population, the mean commute and the population density. While the fraction of non-white population has by far the largest causal connection with the confirmed case rates, the mean commute and density variables interchange positions depending on the causal ordering used. However, it is clear that these three are the most important causation.

In the West Coast region, the fraction of non-white and the fraction of senior citizen in a county are the top causation. The fraction employed in a county complete the top three list in the CO#1 and CO#2 while loses importance in CO#3 leading us to believe that its importance is more sensitive to the causal ordering. In the Southern States the conclusion is similar with the fraction of non-white population, the fraction of senior citizen and the population density having the strongest causal connections with the confirmed case rate with non-white fraction being the primary causation.

Moving on to Phase II the pattern is quite different. In [Figure 4](#) we see that the value of $|\overline{S}_v|$ for all the variables are much lower than the corresponding ones in Phase I for both the East Coast States and the Southern States leading us to conclude that the variables have less effect in causing the spread of the disease which is more homogeneous geographically. The case for the West Coast region is somewhat different with the non-white fraction still holding some causal connection with the spread of the disease along with the fraction of senior citizen for CO#2 and CO#3. In addition, the correlation patterns shown in the network plots in [Figure 4](#) differ from those in [Figure 3](#) and seem to be more sparse for the East Coast region and the Southern States while being relatively unchanged for the West Coast region.

Econometric analysis

The machine learning framework along with the causal Shapley values are novel method for analyzing causalities in socioeconomic data. To validate its strengths and to highlight its differences with a more traditional linear causal analysis we use a simple linear model to analyze the effect of socioeconomic metrics on the spread of COVID-19 across the USA.

We use weekly balanced panel data, and we safely assume that standard errors in the balanced panel data are not independent. To begin with, we perform the White test and Breusch–Pagan test to detect the presence of heteroscedasticity. For each of the six datasets (three regions and two phases) we get a p -value < 0.05 from both the tests which leads us to reject the null hypothesis of homoscedasticity and we conclude that heteroscedasticity is present in the residual of this empirical model specification. Furthermore, we perform the Durbin-Watson test for auto-correlation and get test results of < 1 for all datasets leading us to the conclusion that there are significant positive auto-correlations in the datasets. Therefore, we need to use a random effect model with clustered standard error at the county level. These clustered standard errors allow for the presence of heteroscedasticity and correlation in the error term within a cluster. A fixed effect model cannot be used given the socioeconomic metrics do not change over the times during which the COVID-19 prevalence data is collected. The random effect model estimates

Phase I			
Variables	East Coast	West Coast	Southern States
Density	0.161 (0.037) ¹	0.114 (0.028) ¹	0.170 (0.009) ¹
Unemployment	0.016 (0.024)	-0.094 (0.024) ¹	-0.007 (0.009)
Income	0.036 (0.027)	-0.062 (0.027) ²	-0.043 (0.010) ¹
Non-White	0.096 (0.034) ¹	0.226 (0.023) ¹	0.160 (0.008) ¹
constant	1.094 (0.019) ¹	0.675 (0.019) ¹	0.889 (0.007) ¹
Observations	3,036	2,926	30,668
Phase II			
Variables	East Coast	West Coast	Southern States
Density	0.111 (0.046) ²	0.142 (0.031) ¹	0.1006 (0.006) ¹
Unemployment	0.039 (0.030)	-0.083 (0.027) ¹	-0.0188 (0.006) ¹
Income	-0.035 (0.034)	-0.177 (0.031) ¹	-0.1103 (0.007) ¹
Non-White	0.038 (0.042)	0.226 (0.027) ¹	0.0067 (0.006)
constant	1.814 (0.024) ¹	1.844 (0.021) ¹	2.1862 (0.006) ¹
Observations	3,588	3,458	36,224
no. of counties	138	133	1,394

Table 1. Results of the causal analyses performed with random effects models. The central value and the robust standard errors (in brackets) are given for the variables that were considered in the analyses. The p -values are marked as: ¹ $p < 0.01$, ² $p < 0.05$.

the effects of time-invariant socioeconomic metrics as presented in the data collected from the US Census Bureau. We build the model with only a subset of the socioeconomic metrics that we consider in the previous analysis since highly correlated metrics cannot be used to fit a linear random effect model. Hence we chose four of the metrics that are not highly correlated and perform the analysis so as to reasonably compare it with the results obtained from the machine learning framework.

The result of the empirical analysis for Phase I is presented in [Table 1](#). The dependent variable is the \log_{10} of weekly COVID-19 confirmed cases per 100,000 individuals within a county. The first row of [Table 1](#) shows that for all the regions, higher population density increases the probability of the spread of the number of COVID-19 cases in the region. The effect is positive and significant (p -value < 0.05). The second row shows that in regions with higher unemployment rate, the spread of COVID-19 is relatively less as seen in the numbers from the west coast (p -value < 0.05). Third and fourth row have the expected results that higher income region and regions with a smaller fraction of non-white population sees a decrease in the probability of the spread of COVID-19 (p -value < 0.05). The effect is significant for all three regions.

The Phase II results in the lower half of [Table 1](#) has similar results as Phase I and shows that socioeconomic metrics affect the spread of COVID-19 significantly, especially in the West Coast states. The primary difference between Phase I and Phase II is that some of the effects of the socioeconomic metrics on the spread of COVID-19 decreases in Phase II compared to Phase I as presented in our previous analysis, especially for the East coast and the Southern States regions. These results highlight two implications of the causal analyses. Firstly, the initial health care policy intervention of some of the US states and the federal government seem to have helped in reducing the differential in the spread of the pandemic from what was seen at the initial stage and was disproportionately affecting the more vulnerable. Secondly, socioeconomic metrics prove to be important in shaping the possible policy interventions to control the spread of the pandemic.

Comparison between methods

Having used two very different frameworks for this analysis which point towards a very similar interpretation of the data, we would like to highlight the major differences between the frameworks themselves in an attempt to allow the readers to judge the merits of the novel machine learning framework that we propose for causal analyses.

Non-linearities

Linear models like the random effect model fail to capture non-linearities in the data and their use is further complicated by correlations present amongst the exogenous variables. This leads to the necessity of considering only those variables that are not highly correlated which leaves open the possibility of ignoring a confounding variable. On the other hand, the framework based on ensembles of BDTs that we use naturally models non-linearities in the data taking into account all correlations and, hence, provide a better regression of the data. This allows for the consideration of a larger set of exogenous variables, which in turn, allows for a better modeling of the endogenous variables.

Interpretability

The drawback of using machine learning (beyond linear regression or other simple models) is that they are not easily interpretable and often end up as black-boxes. This is, however, not the case for linear models like the random effects model which allows for a very clear interpretation of the model parameters. The lack of interpretability is, in fact, a major hurdle in the use of machine learning in computational socioeconomics. We address this problem by using Shapley values to interpret the regression model we create with the ensemble of BDTs and do so in a manner in which the causality in the data is probed, in effect, adding interpretability to a black-box model.

Summary

The work that we present here has two primary components. Firstly, we establish the tenets of a multivariate causal analysis using interpretable machine learning which aims at consistently taking into account non-linearities and correlations and build upon Shapley values, taken from coalition game theory, to extract causal information from the data. Our use of an ensemble of BDTs, instead of other non-linear machine learning models, is motivated by the ease of computation of Shapley values from BDTs using available software packages and should not be taken as a limitation of the causal analysis we propose on a choice of the machine learning models that can possibly be used. The causal Shapley values we calculate are based on the hypotheses about the causal connections between the exogenous variables which are represented by causal chain graphs built on assumptions of the partial causal orderings between the variables. We have tested several hypotheses of these causal orderings to ascertain the robustness of the analysis framework and its dependence on such hypotheses.

Secondly, we use this analysis framework to study the effects of socioeconomic disparities on the spread of COVID-19 in the USA. We use population demographics data gathered from the US Census Bureau and COVID-19 prevalence data from the Johns Hopkins University to fit regression models using the interpretable machine learning framework. The extraction of causal Shapley values from these regression models allow us to infer on the causal connections between the socioeconomic metrics and the prevalence of COVID-19 at the county level. To compare the results of our analysis with a more traditional causal analysis we use a weekly balanced panel data to fit a random effects model with a reduced set of exogenous variables and find reasonable agreement with the results from the interpretable machine learning analysis.

From these analyses we conclude that:

- The effects of socioeconomic disparities on the spread of COVID-19 was more pronounced at the beginning of the pandemic than at the later stages.
- While in the parts of the USA with higher population density, the spread was driven partially by the population density, socioeconomic metrics like the fraction of non-white population in a county also show significant causal connection with the spread of COVID-19. In fact, population density was not causally connected to the spread of the disease in the West coast region.

- Of particular note is the causal connection between the fraction of senior citizen in a county and the spread of the disease in the West Coast region and the Southern States. The Shapley values being anti-correlated to the COVID-19 confirmed case rate implies that counties with a younger population saw a larger spread of the disease. As we know that the probability of COVID-19 infection increases drastically with age, our results imply that while the older fraction of the population were being differentially affected to a larger extent by the disease it was being spread more widely by the younger and more mobile fraction of the population.

While we have not addressed the question of confounding variables in any detail, it is possible to assume the presence of confounding amongst the variable clustered in a partial ordering. We present this analysis in the Supplementary Information. Assuming the presence of confounding variables does not change the results of our analysis.

With this work, we hope that our proposed methods for causal analysis finds some utility in computational socioeconomics much beyond the application that we have proposed. The conclusions that we draw about the effects of socioeconomic disparities on the spread of COVID-19 in the USA, especially during the onset of the pandemic reinforces what has been observed with clinical data and population level analyses and points to the necessity for restructuring of the crisis response system to nullify the causes of such disparities. We hope our work will lead to more detailed thoughts, insights and actions that will prove to be useful in the near future.

Methods

The data for this work comes from three primary sources:

- The 2019 American Community Survey (ACS) 5-years supplemental update to the 2011 Census found in the USA Census Bureau database for constructing the socioeconomic metrics.
- The population density data that reflects the 2019 estimates of the US Census Bureau.
- Data on COVID-19 prevalence and death rate is obtained from the Johns Hopkins University, Center for Systems Science and Engineering database [44] sourced from github.com/CSSEGISandData/COVID-19.

To study the relevance of socioeconomic condition in the spread of COVID-19, we focus on various metrics that can characterize these at a county level including factors such as per capita income, poverty, the employed fraction of the population, and the unemployment rate. The latter two are not fully correlated since they add up to the fraction of the population that is employable which varies from county to county. We also include factors that relate to the mobility such as the mean commute time in any county along with the fraction of the population that uses a transit system³. We also include the fraction of senior citizen in a county which can also affect the mobility patterns in the county. For example, a population with a lower median age will, on an average, have higher mobility than a population with a much higher median age. Since studies have already shown that COVID-19 tends to affect the elderly preferentially, the fraction of a population that fall into the senior citizen category (over 65 years of age) is a good measure of quantifying both the mobility and the degree to which age plays a role in determining the spread of COVID-19. To account for professions that put individuals at higher risk of exposure to a larger number of people such as those in service industry, construction, delivery, labor etc., which increases their chance of being infected by COVID-19, we employ a metric that quantifies the fraction of the population that work in these industries per county. Lastly, we include the fraction of individuals in a county that do not have health insurance as a metric to assess if it affects the spread of the disease.

The socioeconomic metrics for each county collected from the US Census Bureau 2019 5-year ACS data used in this work are the following:

- **Population Density (Den):** The population density data taken from 2019.
- **Non-White (NW):** The fraction of non-white population in any county including Hispanics and Latinos.

³We do not consider the relative reduction in mobility due to partial lock-downs or closures of institutions.

- **Income (Inc)**: The income per capita as defined by the US Census Bureau.
- **Poverty (Pov)**: The fraction of the population deemed as being below the poverty line.
- **Unemployment (Uemp)**: The unemployment rate as defined by the US Census Bureau.
- **Uninsured (Uins)**: The fraction of the population that does not have health insurance.
- **Employed (Emp)**: The fraction of the population that is employed.
- **Labour (Lab)**: The fraction of the population working in construction, service, delivery or production.
- **Transit (Tran)**: The fraction of the population who take the public transportation system or carpool excluding those who drive or work from home.
- **Mean Commute (MC)**: The mean commute distance for a person living in a county in minutes.
- **Senior Citizen (SC)**: The fraction of the population that is above 65 years of age.
- **Gini index, or Gini coefficient (GI)**: A measure of the distribution of income across a population used as a gauge of economic inequality.

The abbreviations in parentheses are what we use to refer to these variables while discussing partial causal ordering and the results. For the data on COVID-19 prevalence we look at the total number of confirmed cases and deaths up until the 15th of January 2021 and define the following rates:

- **Confirmed Case Rate**: The total number of confirmed cases per 100,000 individuals in any county.
- **Death Rate**: The total number of deaths per 100,000 individuals in any county.

We do not use the fatality rate (number of deaths per confirmed case) as a measure of the disease spread since in [23] it was shown the fatality rates are mostly uncorrelated with the socioeconomic metrics pointing to the fact that fatality rates were not disproportionately high in certain socioeconomic strata of the society.

Comorbidities are known to aggravate COVID-19 infections leading to higher chances of a symptomatic infection and hospitalization. We include the data provided in a study conducted by the Centers for Disease Control and Prevention [45] which provides the distribution of comorbidities in all the counties in the USA to study the possible effects of comorbidities in determining the extent to which COVID-19 spreads in any region. We define the corresponding metric as:

- **Comorbidities (Com)**: The total fraction of population with one or more pre-existing chronic conditions.

Analysis Framework

In this work we aim at extracting the causal connections between the exogenous variables (socioeconomic metrics) and endogenous variables (COVID-19 prevalence and death rate) by using machine learning for modeling underlying multivariate distributions from the data, and then calculating causal Shapley values. The procedure we follow can be delineated as:

- Use an ensemble of boosted decisions trees (BDTs) which act as weak learners that perform reliably in a statistical ensemble. This allows us to build a non-linear model of COVID-19 case and death rates in terms of the socioeconomic metrics and population density at the county level.
- Define the causal flow of the variables by setting up partial causal ordering graphs.
- Calculate the causal Shapley values that allow us to infer upon the causal connections between the exogenous variables and the endogenous ones.

To clarify the analysis procedures we delve into some details of the machine learning framework that we use and the definition of causal Shapley values. Of particular importance is our discussion of how we set up the partial causal ordering graphs that determine the causal flow of the variables. This is the part where an understanding of the interactions between the variables is necessary and we use several hypothesis to make sure we properly address the subjective nature of designing the partial causal orderings.

Machine learning framework

To perform a regression of the data we use an ensemble of boosted decisions trees (BDTs) taking either disease prevalence or death rate as the endogenous variable while keeping the socioeconomic metrics and population density as exogenous variables. This allows us to not assume a functional form for the model but rely on the data alone. We emphasize here that we are not trying to build a predictive model using machine learning but rather to perform regression in a model-agnostic manner. We use XGBoost [46], a scalable end-to-end boosting system for decision trees that is particularly suitable for sparse data. For data augmentation we use an ensemble of BDTs as weak learners trained on random selections of the sample with replacement split 70/30 into training and testing sets. This also gives us a stable measure of the accuracy with which we can model the data. We use the coefficient of determination, R^2 , calculated from the test sample set aside for each BDT, as our regression accuracy measure and estimate the error of estimating R^2 from the ensemble. It is normal for weak learners to differ in their predictions (it is a consequence of bagging within the ensemble). The prediction of an ensemble of weak learners is taken as the average in a regression problem. The accuracy of the ensemble is represented by the mean R^2 .

Causal Shapley values

The Shapley value, first introduced in [24], is a solution concept from cooperative game theory for transferable utility games. The game is characterised by a pair (C, F) , where $F = \{1, \dots, d\}$ is a set representing players in the game, and the *characteristic function* $C : 2^F \rightarrow \mathbb{R}$ assigns a non-negative real value $C(S)$ to every coalition $S \subseteq F$, and zero to the empty coalition, i.e. $C(\emptyset) = 0$. The Shapley decomposition is provably (cf. [47, Thm. 2]) the only solution concept satisfying the four favorable axioms of *efficiency*, *additivity*, *symmetry* and *null player*, well exposed in [48]. This uniqueness has contributed to the popularity of the Shapley value in the literature of explainable artificial intelligence (XAI). However, different incarnations of Shapley value based explanation methods make different assumptions and use of different approximations, rendering Shapley value based explanations less “unique”.

Using Shapley values in a machine learning context amounts to interpreting the players F as features in a model, and the characteristic function C as either the model itself or an evaluation measure of the model’s performance. Lundberg et al. [25] identified that several existing explanation methods, including Shapley value based ones, belong to a class of “additive feature attribution methods”, and unified Shapley values with the solution concept of local interpretable model-agnostic explanations (LIME) [49], introducing the SHAP model with simplified inputs, i.e. inclusion/exclusion of features. As machine models assume an input shape defined during training, evaluating it in the absence of a feature is not possible in general. Hence, the SHAP characteristic function approximates the (counterfactual) model prediction under removal of features via the expected value of the prediction conditional upon the values of the included features.

In the [25] iteration of so-called Kernel SHAP, this conditional expectation is approximated via the marginal distribution, see eq. (11) in [25], which amounts to an assumption of feature independence. This is done for computational efficiency, and not inherent to the SHAP framework. The Python SHAP package [50], first released along with [25], is maintained. In [26], Aas et al. suggest improving Kernel SHAP by obtaining the values of the out of coalition features by conditioning upon the excluded features when calculating the expected values. As the conditional distribution is in general not known, Ref. [26] presents four ways of approximating it. One of these is via the empirical distribution of the data, weighing data instances based on a scaled Mahalanobis distance [51]. The authors of Ref. [26] refer to this as conditional Kernel SHAP, which they implement in the R-package *shapr* [52].

In Ref. [53], Janzing et al. argue that unconditional, rather than conditional, expectations provide the right notion of excluding features, which contradicts the theoretical justification of [25], used in the SHAP software packages. Drawing from Pearl’s seminal work on causality [54]⁴, Janzing et al. stress that the *interventional* distribution is represented by the marginal distribution. Leaning heavily on the work of Datta et al. [56], Janzing et al. thus conclude that the marginal distribution – not the conditional distribution – should be used for excluded features in the SHAP calculation.

Shortly after, Heskens et al. [27] show that this marginal, respectively interventional, SHAP calculation only represents direct effects, meaning that ‘root causes’ with strong indirect effects are ignored in the feature attribution.

⁴The interested reader is referred to [55] for an introduction to the *do*-calculus

They introduce *causal* Shapley values by explicitly including the causal relationships between the data, in the SHAP value calculation. They provide an extension to the shapr package, where the user specifies the causal ordering as well as possible confounding, available at [57]⁵.

In our work, we will use the form of causal Shapley values defined by Heskes et al. [27] through the definition of partial causal ordering graphs. On one hand, this makes the analysis of the causal connection somewhat subjective and based on our decision of the causal connections between the variables (and the presence of confounding variables). On the other hand, this enforcement of the causal ordering can be seen as a “prior” with which we inform the computation of the causal Shapley values the information that might not be directly gleaned from the data. We will, in particular, use two quantities to examine the data. The first one is the mean of the absolute Shapley values, $|\overline{S}_v|$, that serves as a measure of global importance of a variable [58]. Higher the value of $|\overline{S}_v|$ of a variable, the greater effect the variable has in determining the outcome. We shall also define S_I as the integer index corresponding to the position in a hierarchy that a certain variable holds when ordered by the corresponding $|\overline{S}_v|$ with $S_I \in [1, n]$ for a n -variable problem.

Data Availability

All the data used in this work is sources from public databases that have been cited in the main text. All codes that have been used to process and analyze the data can be found at <https://github.com/talismanbrandi/Causal-Inference-IML-C19>.

References

1. Millett, G. A. *et al.* Assessing differential impacts of covid-19 on black communities. *Annals Epidemiol.* **47**, 37–44, DOI: [10.1016/j.annepidem.2020.05.003](https://doi.org/10.1016/j.annepidem.2020.05.003) (2020).
2. Yancy, C. W. COVID-19 and African Americans. *JAMA* **323**, 1891–1892, DOI: [10.1001/jama.2020.6548](https://doi.org/10.1001/jama.2020.6548) (2020).
3. Chastain, D. B. *et al.* Racial disproportionality in covid clinical trials. *New Engl. J. Medicine* **383**, e59, DOI: [10.1056/NEJMp2021971](https://doi.org/10.1056/NEJMp2021971) (2020).
4. Moore J T, R. C. E., Ricaldi J N *et al.* Disparities in incidence of covid-19 among underrepresented racial/ethnic groups in counties identified as hotspots during june 5-18, 2020 - 22 states, february–june 2020. *MMWR Morb Mortal Wkly Rep.* DOI: [10.15585/mmwr.mm6933e1](https://doi.org/10.15585/mmwr.mm6933e1) (2020).
5. Martinez, D. A. *et al.* SARS-CoV-2 Positivity Rate for Latinos in the Baltimore-Washington, DC Region. *JAMA* **324**, 392–395, DOI: [10.1001/jama.2020.11374](https://doi.org/10.1001/jama.2020.11374) (2020).
6. DiMaggio, C., Klein, M., Berry, C. & Frangos, S. Black/african american communities are at highest risk of covid-19: spatial modeling of new york city zip code-level testing results. *Annals Epidemiol.* **51**, 7–13, DOI: [10.1016/j.annepidem.2020.08.012](https://doi.org/10.1016/j.annepidem.2020.08.012) (2020).
7. Khanijahani, A. Racial, ethnic, and socioeconomic disparities in confirmed covid-19 cases and deaths in the united states: a county-level analysis as of november 2020. *Ethn. & Heal.* **26**, 22–35, DOI: [10.1080/13557858.2020.1853067](https://doi.org/10.1080/13557858.2020.1853067) (2021). PMID: 33334160.
8. Pareek, M. *et al.* Ethnicity and covid-19: an urgent public health research priority. *The Lancet* **395**, 1421–1422, DOI: [10.1016/S0140-6736\(20\)30922-3](https://doi.org/10.1016/S0140-6736(20)30922-3) (2020).
9. Laurencin, C. T. & McClinton, A. The covid-19 pandemic: a call to action to identify and address racial and ethnic disparities. *J. Racial Ethn. Heal. Disparities* **7**, 398–402, DOI: [10.1007/s40615-020-00756-0](https://doi.org/10.1007/s40615-020-00756-0) (2020).
10. Goyal, M. K. *et al.* Racial and/or ethnic and socioeconomic disparities of sars-cov-2 infection among children. *Pediatrics* DOI: [10.1542/peds.2020-009951](https://doi.org/10.1542/peds.2020-009951) (2020).

⁵This patch has not yet been included in the official shapr package.

11. Wright, A. L., Sonin, K., Driscoll, J. & Wilson, J. Poverty and Economic Dislocation Reduce Compliance with COVID-19 Shelter-in-Place Protocols. *Univ. Chicago, Becker Friedman Inst. for Econ. Work. Pap.* DOI: [10.2139/ssrn.3573637](https://doi.org/10.2139/ssrn.3573637) (2020).
12. Weill, J. A., Stigler, M., Deschenes, O. & Springborn, M. R. Social distancing responses to covid-19 emergency declarations strongly differentiated by income. *Proc. Natl. Acad. Sci.* **117**, 19658–19660, DOI: [10.1073/pnas.2009412117](https://doi.org/10.1073/pnas.2009412117) (2020).
13. Qiu, Y., Chen, X. & Shi, W. Impacts of social and economic factors on the transmission of coronavirus disease 2019 (COVID-19) in China. *J. Popul. Econ.* **33**, 1127–1172, DOI: [10.1007/s00148-020-00778-2](https://doi.org/10.1007/s00148-020-00778-2) (2020).
14. Stojkoski, V., Utkovski, Z., Jolakoski, P., Tevdovski, D. & Kocarev, L. The socio-economic determinants of the coronavirus disease (COVID-19) pandemic. *arXiv e-prints* arXiv:2004.07947 (2020). [arXiv:2004.07947](https://arxiv.org/abs/2004.07947).
15. Ransome, Y., Kawachi, I., Braunstein, S. & Nash, D. Structural inequalities drive late HIV diagnosis: The role of black racial concentration, income inequality, socioeconomic deprivation, and HIV testing. *Heal. & place* **42**, 148–158, DOI: [10.1016/j.healthplace.2016.09.004](https://doi.org/10.1016/j.healthplace.2016.09.004) (2016).
16. Farmer, P. Social inequalities and emerging infectious diseases. *Emerg. infectious diseases* **2**, 259–269, DOI: [10.3201/eid0204.960402](https://doi.org/10.3201/eid0204.960402) (1996).
17. Hosseini, P., Sokolow, S. H., Vandegrift, K. J., Kilpatrick, A. M. & Daszak, P. Predictive Power of Air Travel and Socio-Economic Data for Early Pandemic Spread. *PLOS ONE* **5**, 1–8, DOI: [10.1371/journal.pone.0012763](https://doi.org/10.1371/journal.pone.0012763) (2010).
18. Quinn, S. C. & Kumar, S. Health inequalities and infectious disease epidemics: a challenge for global health security. *Biosecurity bioterrorism : biodefense strategy, practice, science* **12**, 263–273, DOI: [10.1089/bsp.2014.0032](https://doi.org/10.1089/bsp.2014.0032) (2014).
19. Smedley, B. D., Stith, A. Y. & Nelson, A. R. (eds.) *Unequal Treatment: Confronting Racial and Ethnic Disparities in Health Care* (The National Academies Press, Washington, DC, 2003), DOI: [10.17226/12875](https://doi.org/10.17226/12875).
20. The Council on Ethical and Judicial Affairs, American Medical Association. Black-White Disparities in Health Care. *JAMA* **263**, 2344–2346, DOI: [10.1001/jama.1990.03440170066038](https://doi.org/10.1001/jama.1990.03440170066038) (1990).
21. Andrews, R. & Elixhauser, A. Use of major therapeutic procedures: are Hispanics treated differently than non-Hispanic Whites. *Ethn. & disease* **10**, 384–394 (2000). URL: <http://europepmc.org/abstract/MED/11110355>.
22. Harris, D., Andrews, R. & Elixhauser, A. Racial and gender differences in use of procedures for black and white hospitalized adults. *Ethn. & disease* **7**, 91–105 (1997). URL: <http://europepmc.org/abstract/MED/9386949>.
23. Paul, A., Englert, P. & Varga, M. Socio-economic disparities and covid-19 in the usa. *J. Physics: Complex.* DOI: [10.1088/2632-072X/ac0fc7](https://doi.org/10.1088/2632-072X/ac0fc7) (2021).
24. Shapley, L. S. A value for n -person games. In Kuhn, H. W. & Tucker, A. W. (eds.) *Contributions to the Theory of Games* (Princeton University Press, Princeton, 1953).
25. Lundberg, S. M. & Lee, S.-I. A unified approach to interpreting model predictions. In Guyon, I. *et al.* (eds.) *Advances in Neural Information Processing Systems 30*, 4765–4774 (Curran Associates, Inc., 2017). URL: <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>.
26. Aas, K., Jullum, M. & Løland, A. Explaining individual predictions when features are dependent: More accurate approximations to shapley values. *Artif. Intell.* **298**, 103502, DOI: [10.1016/j.artint.2021.103502](https://doi.org/10.1016/j.artint.2021.103502) (2021).
27. Heskes, T., Sijben, E., Bucur, I. G. & Claassen, T. Causal shapley values: Exploiting causal knowledge to explain individual predictions of complex models. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F. & Lin, H. (eds.) *Advances in Neural Information Processing Systems*, vol. 33, 4778–4789 (Curran Associates, Inc., 2020). URL: <https://proceedings.neurips.cc/paper/2020/file/32e54441e6382a7fbacbbaf3c450059-Paper.pdf>.
28. Immergluck, D. *Neighborhood Jobs, Race, and Skills* (Routledge, 2018), DOI: [10.4324/9781351045957](https://doi.org/10.4324/9781351045957).

29. Mills, B. & Hazarika, G. The migration of young adults from non-metropolitan counties. *Am. J. Agric. Econ.* **83**, 329–340, DOI: [10.1111/0002-9092.00159](https://doi.org/10.1111/0002-9092.00159) (2001).
30. Charles, K. K., Hurst, E. & Schwartz, M. The transformation of manufacturing and the decline in US employment. *NBER Macroecon. Annu.* **33**, 307–372, DOI: [10.1086/700896](https://doi.org/10.1086/700896) (2019).
31. Lambert, T. E., Mattson, G. A. & Dorriere, K. The impact of growth and innovation clusters on unemployment in US metro regions. *Reg. Sci. Policy & Pract.* **9**, 25–37, DOI: [10.1111/rsp3.12087](https://doi.org/10.1111/rsp3.12087) (2017).
32. DeFina, R. H. The impacts of unemployment on alternative poverty rates. *Rev. Income Wealth* **50**, 69–85, DOI: [10.1111/j.0034-6586.2004.00112.x](https://doi.org/10.1111/j.0034-6586.2004.00112.x) (2004).
33. Finkelstein, A. & McKnight, R. What did medicare do? the initial impact of medicare on mortality and out of pocket medical spending. *J. Public Econ.* **92**, 1644–1668, DOI: [10.1016/j.jpubeco.2007.10.005](https://doi.org/10.1016/j.jpubeco.2007.10.005) (2008).
34. Cunningham, P. J. & Ginsburg, P. B. What accounts for differences in uninsurance rates across communities? *INQUIRY: The J. Heal. Care Organ. Provision, Financing* **38**, 6–21, DOI: [10.5034/inquiryjrn1_38.1.6](https://doi.org/10.5034/inquiryjrn1_38.1.6) (2001).
35. Levernier, W., Partridge, M. D. & Rickman, D. S. Differences in metropolitan and nonmetropolitan u.s. family income inequality: A cross-county comparison. *J. Urban Econ.* **44**, 272–290, DOI: [10.1006/juec.1997.2070](https://doi.org/10.1006/juec.1997.2070) (1998).
36. Nielsen, F. & Alderson, A. S. The kuznets curve and the great u-turn: Income inequality in u.s. counties, 1970 to 1990. *Am. Sociol. Rev.* **62**, 12, DOI: [10.2307/2657450](https://doi.org/10.2307/2657450) (1997).
37. Adelman, R. M. & Jaret, C. Poverty, race, and us metropolitan social and economic structure. *J. Urban Aff.* **21**, 35–56, DOI: [10.1111/0735-2166.00002](https://doi.org/10.1111/0735-2166.00002) (1999).
38. Rothwell, J. T. & Massey, D. S. Density zoning and class segregation in u.s. metropolitan areas. *Soc. Sci. Q.* **91**, 1123–1143, DOI: [10.1111/j.1540-6237.2010.00724.x](https://doi.org/10.1111/j.1540-6237.2010.00724.x) (2010).
39. Kasarda, J. D. Inner-city concentrated poverty and neighborhood distress: 1970 to 1990. *Hous. Policy Debate* **4**, 253–302, DOI: [10.1080/10511482.1993.9521135](https://doi.org/10.1080/10511482.1993.9521135) (1993).
40. Bertaud, A. & Richardson, H. W. Transit and density: Atlanta, the united states and western europe. In *in Urban Sprawl in Western Europe and the United States, Urban Planning and Environment*, Ashgate (2004).
41. LEVINSON, D. M. & KUMAR, A. Density and the journey to work. *Growth Chang.* **28**, 147–172, DOI: [10.1111/j.1468-2257.1997.tb00768.x](https://doi.org/10.1111/j.1468-2257.1997.tb00768.x) (1997).
42. Bélanger, A. Diversity explosion: How new racial demographics are remaking america. *Can. Stud. Popul.* **43**, 166, DOI: [10.25336/p69s3v](https://doi.org/10.25336/p69s3v) (2016).
43. Ono, Y. & Sullivan, D. Manufacturing plants' use of temporary workers: An analysis using census microdata. *Ind. Relations: A J. Econ. Soc.* **52**, 419–443, DOI: [10.1111/irel.12018](https://doi.org/10.1111/irel.12018) (2013).
44. Dong, E., Du, H. & Gardner, L. An interactive web-based dashboard to track COVID-19 in real time. *The Lancet Infect. Dis.* **20**, 533–534, DOI: [10.1016/S1473-3099\(20\)30120-1](https://doi.org/10.1016/S1473-3099(20)30120-1) (2020).
45. Razzaghi H, L. H., Wang Y *et al.* Estimated county-level prevalence of selected underlying medical conditions associated with increased risk for severe covid-19 illness — united states, 2018. *MMWR Morb Mortal Wkly Rep.* **69**, 945–950, DOI: [10.15585/mmwr.mm6929a1](https://doi.org/10.15585/mmwr.mm6929a1) (2020).
46. Chen, T. & Guestrin, C. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, 785–794, DOI: [10.1145/2939672.2939785](https://doi.org/10.1145/2939672.2939785) (Association for Computing Machinery, New York, NY, USA, 2016).
47. Young, H. P. Monotonic solutions of cooperative games. *Int. J. Game Theory* **14**, 65–72, DOI: [10.1007/BF01769885](https://doi.org/10.1007/BF01769885) (1985).
48. Huettner, F. & Sunder, M. Axiomatic arguments for decomposing goodness of fit according to Shapley and Owen values. *Electron. J. Stat.* **6**, 1239 – 1250, DOI: [10.1214/12-EJS710](https://doi.org/10.1214/12-EJS710) (2012).

49. Ribeiro, M. T., Singh, S. & Guestrin, C. “why should i trust you?” explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144 (2016).
50. Lundberg, S. shap 0.39.0. <https://pypi.org/project/> (2021).
51. Mahalanobis, P. C. On the generalized distance in statistics. *Proc. Natl. Inst. Sci. (Calcutta)* **2**, 49–55 (1936).
52. Aas, K., Jullum, M. & Løland, A. shapr 0.2.0.9000. <https://rdr.io/github/NorskRegnesentral/shapr/> (2021).
53. Janzing, D., Minorics, L. & Bloebaum, P. Feature relevance quantification in explainable ai: A causal problem. In Chiappa, S. & Calandra, R. (eds.) *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, vol. 108 of *Proceedings of Machine Learning Research*, 2907–2916 (PMLR, 2020).
54. Pearl, J. Causal diagrams for empirical research. *Biometrika* **82**, 669–688 (1995).
55. Pearl, J. The do-calculus revisited. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, UAI’12, 3–11 (AUAI Press, Arlington, Virginia, USA, 2012).
56. Datta, A., Sen, S. & Zick, Y. Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *2016 IEEE Symposium on Security and Privacy (SP)*, 598–617, DOI: [10.1109/SP.2016.42](https://doi.org/10.1109/SP.2016.42) (2016).
57. Bucur, I. G. shapr. <https://gitlab.science.ru.nl/gbucur/shapr/> (2020).
58. Lundberg, S. M. *et al.* From local explanations to global understanding with explainable ai for trees. *Nat. Mach. Intell.* **2**, 56–67, DOI: [10.1038/s42256-019-0138-9](https://doi.org/10.1038/s42256-019-0138-9) (2020).

Acknowledgements

A.P. is funded in part by Volkswagen Foundation within the initiative “Corona Crisis and Beyond – Perspectives for Science, Scholarship and Society”, grant number 99091. I. S. is grateful for support received through the EXAIGON project from industrial partners and the Research Council of Norway (grant no. 304843). This research was supported in part through the Maxwell computational resources operated at DESY, Hamburg, Germany.

Author contributions statement

All authors contributed in their own capacities in performing the analysis and writing the draft. The contribution of V.S. deserves a special mention. He contributed significantly to the conceptual design and implementation of the causal analysis, the building of the causal orderings and data analysis.

Supplementary Information: Socioeconomic disparities and COVID-19: the causal connections

Tannista Banerjee, Ayan Paul, Vishak Srikanth, and Inga Strümke

Confounding variables

Given the complexity of the problems that we address the the relative difficulty in engineering all the variables necessary to address it, there is always a possibility of the existence of confounding variables in the analysis. The causal Shapley value framework, as described in the Materials and Methods section, allow for the assumption of the existence of confounding variables within a group of variables that are not causally connected or are at the same level in the directed acyclic graph. To test whether assuming the presence of confounding variables change the results of our analysis we performed all the analysis for the confirmed case rate (two phases, three regions and six causal ordering) assuming the presence of confounding variables in all the groups. The results of our analysis is shown in [Figure S1](#) and [Figure S2](#). It can be seen that assuming the presence of confounding variables do not change the final result of the analysis at all. Hence, we can say with confidence that we safely ignore the presence of confounding variables. That being said, one can argue that variables like the average level of education in a county are highly correlated with several of the socioeconomic metrics that we consider and maybe it can replace one of the metrics that we consider here. However, since we test the causal connections between the variables under a certain hypothesis represented by the directed acyclic graphs, studying the details of all possible hypothesis is beyond the scope of our work and we would like to motivate future works to consider a different set of variable to understand better, what the true causations in the real-world are.

Additional causal orderings

We have considered three causal orderings in the main text. To give some more possible scenarios for the causal dependence within the socioeconomic metrics we propose three more causal ordering. We will try to motivate them here and provide the results in the figures that follow

CO#4: The fourth causal ordering is [NW, SC, [Emp, Uemp], [Inc, Lab], [Uins, Com, Pov, GI], [Den, MC, Tran]]. This causal ordering builds upon the first causal ordering CO#1, except that the proportion of senior citizens is no longer on equal footing as unemployment and employed, but moved up in the causal chain. This causal ordering takes into account the fact that senior citizens who are looking for employment may not be able to find employment due to a skills mismatch with the required skills today and senior citizens may be structurally unemployed [1, 2].

CO#5: The fifth causal ordering is [NW, [SC, Emp, Uemp, Inc, Lab], [Uins, Com, Pov, GI], [Den, MC, Tran]]. This causal ordering has income per capita, the proportion of a county working in construction, service, delivery or production, unemployment, employed, and the proportion of senior citizens on equal footing. The rationale here is that income per capita and the proportion of a county working in construction, service, delivery or production are on equal footing with unemployment and employment since construction, service, delivery, or production jobs are (a) more likely to be temporary (b) pay low wages and counties with people working in those job profiles tend to have fewer permanent jobs and a higher unemployment rate due to the decline of manufacturing in the last half century [3].

CO#6: The sixth causal ordering is [SC, NW, [Emp, Uemp], [Inc, Lab], [Uins, Com, Pov, GI], [Den, MC, Tran]]. This causal ordering has the proportion of senior citizens causing the proportion of non-whites, and is otherwise the same as the first causal ordering CO#1. Since senior citizens are less likely to be non-white, a higher fraction of senior citizens causes there to be a lower proportion of non-whites [4].

The results obtained by assuming these causal ordering between the socioeconomic metrics can be found in [Figure S3](#) and [Figure S4](#). The results from these causal orderings are quite similar to the ones obtained from the CO#1, CO#2 and CO#3, pointing to the same sets of variables as the most causally connected ones in Phase I. In Phase II we see the diminished impact of the socioeconomic metrics in all regions except for the west coast states.

Causal analyses for death rates

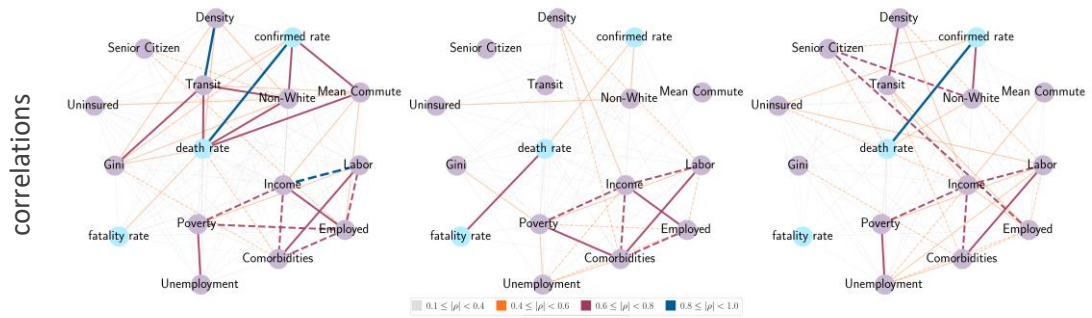
The analysis in the main paper focuses only on understanding the causal connections between the socioeconomic metrics and the confirmed case rate. A similar analysis can be performed using the death rate as the endogenous variable. We show the

	Den	NW	Pov	Inc	Uemp	Uins	Emp	Lab	Tran	MC	SC	GI	Com
mean	273.14	23.75	15.11	28071.83	5.29	9.64	55.11	47.53	15.20	23.85	18.79	0.45	46.79
std	1797.52	20.22	6.33	6782.95	2.65	5.11	8.27	7.32	6.55	5.86	4.66	0.04	6.38
min	0.04	0.00	2.40	10388.00	0.00	0.00	14.70	14.20	0.00	5.10	3.20	0.30	22.00
25%	16.52	7.52	10.60	23613.75	3.60	5.80	49.80	43.40	11.91	19.80	15.80	0.42	42.90
50%	44.84	16.26	14.20	27301.50	4.90	8.70	55.70	48.30	14.03	23.70	18.45	0.44	47.20
75%	118.55	35.56	18.48	31245.25	6.50	12.10	61.20	52.40	16.76	27.60	21.20	0.47	51.00
max	71874.13	99.31	55.50	76592.00	27.20	46.30	80.00	78.70	92.43	52.60	56.70	0.71	66.20

Table S1. Various statistics of the data collected from the US Census Bureau and Centers for Disease Control and Prevention. *Den: Population Density, NW: Non-White, Pov: Poverty, Inc: Income, Uemp: Unemployment, Uins: Uninsured, Emp: Employed, Lab: Labor, Tran: Transit, MC: Mean Commute, SC: Senior citizen, GI: Gini Index, Com.: Comorbidities. Density is measured per sq. km, Income in US dollars, Transit in minutes and the rest are percentages.*

results of this analysis in [Figure S5](#) to [Figure S8](#). We see that the landscape of the causal dependence is a bit different for the death rates as compared with the confirmed case rates. There seems to be a higher causal connected between a larger set of socioeconomic metrics and death rates in Phase I. The top three causations are Density, Non-White and Mean Commute for the East Coast region irrespective of the causal ordering assumed. For the West Coast region the Senior Citizen metric seems to almost always appear on the top. For the Southern States region, Non-White and Density are always the top two causations. In Phase II the causal connection are significantly diluted between the exogenous and endogenous variables in all regions except that we still see some causal connection between Senior Citizen and death rates in the West Coast region and as always it is anti-correlated pointing to the fact that the death rates are higher in the counties in which the fraction of senior citizen is lower and the overall population is younger.

Phase I
February 2020
to
July 2020



causal ordering

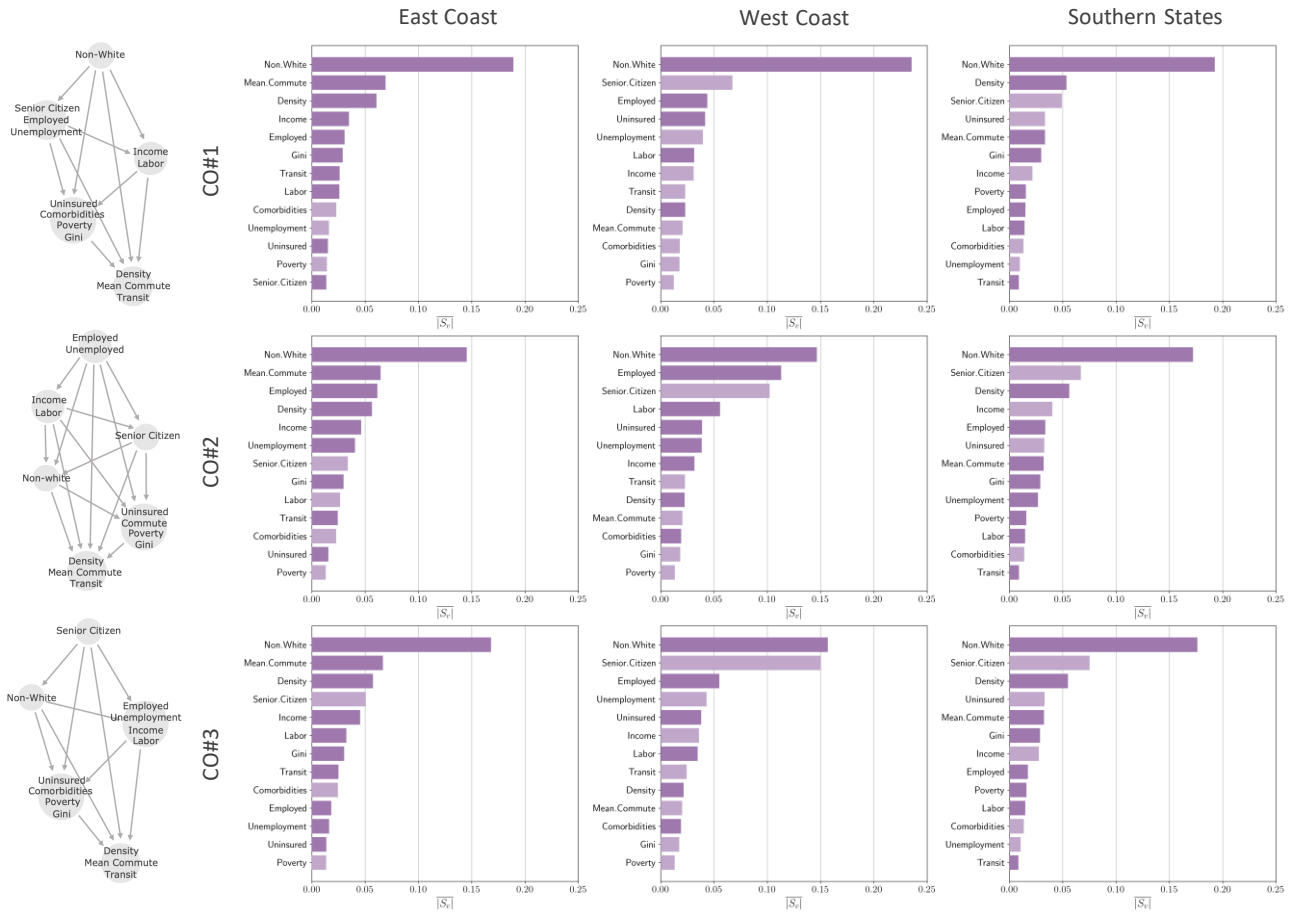


Figure S1. This analysis is performed assuming the presence of confounding variables in each group of variables in the causal ordering. For details please see text. The causal connection between the socioeconomic metrics that we consider in this work and confirmed case rates in various regions of the USA from February 2020 to July 2020. The states included in the East Coast region are District of Columbia, New Jersey, Rhode Island, Massachusetts, Connecticut, Maryland, Delaware and New York. The ones in the West Coast region are California, Oregon and Washington. The states in the Southern States region are Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia. The networks in the top row show the correlations between the different metrics in the three regions. The graphs in the left column of the plot show the three partial causal orderings that we consider in this analysis. Each panel of the array of bar plots show the hierarchy of the metrics in terms of mean absolute SHAP values, $|S_v|$, for a particular region and causal ordering. The lighter shaded bars indicate a negative effect of the metric on the confirmed case rates while the darker bars indicate a positive effect.


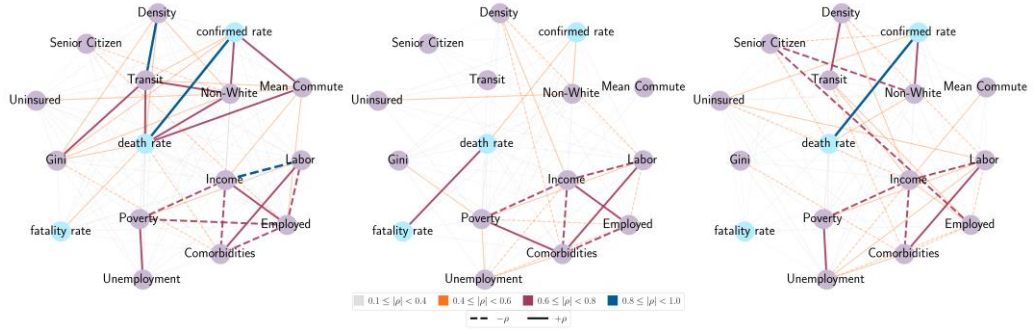


Figure S2. The analysis is performed assuming the presence of confounding variables in each group of variables in the causal ordering. For details please see text. The causal connection between the socioeconomic metrics that we consider in this work and confirmed case rates in various regions of the USA from July 2020 to January 2021. The states included in the East Coast region are District of Columbia, New Jersey, Rhode Island, Massachusetts, Connecticut, Maryland, Delaware and New York. The ones in the West Coast region are California, Oregon and Washington. The states in the Southern States region are Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia. The networks in the top row show the correlations between the different metrics in the three regions. The graphs in the left column of the plot show the three partial causal orderings that we consider in this analysis. Each panel of the array of bar plots show the hierarchy of the metrics in terms of mean absolute causal SHAP values, $|S_v|$, for a particular region and causal ordering. The lighter shaded bars indicate a negative effect of the metric on the confirmed case rates while the darker bars indicate a positive effect.

Phase I
February 2020
to
July 2020

correlations



causal ordering

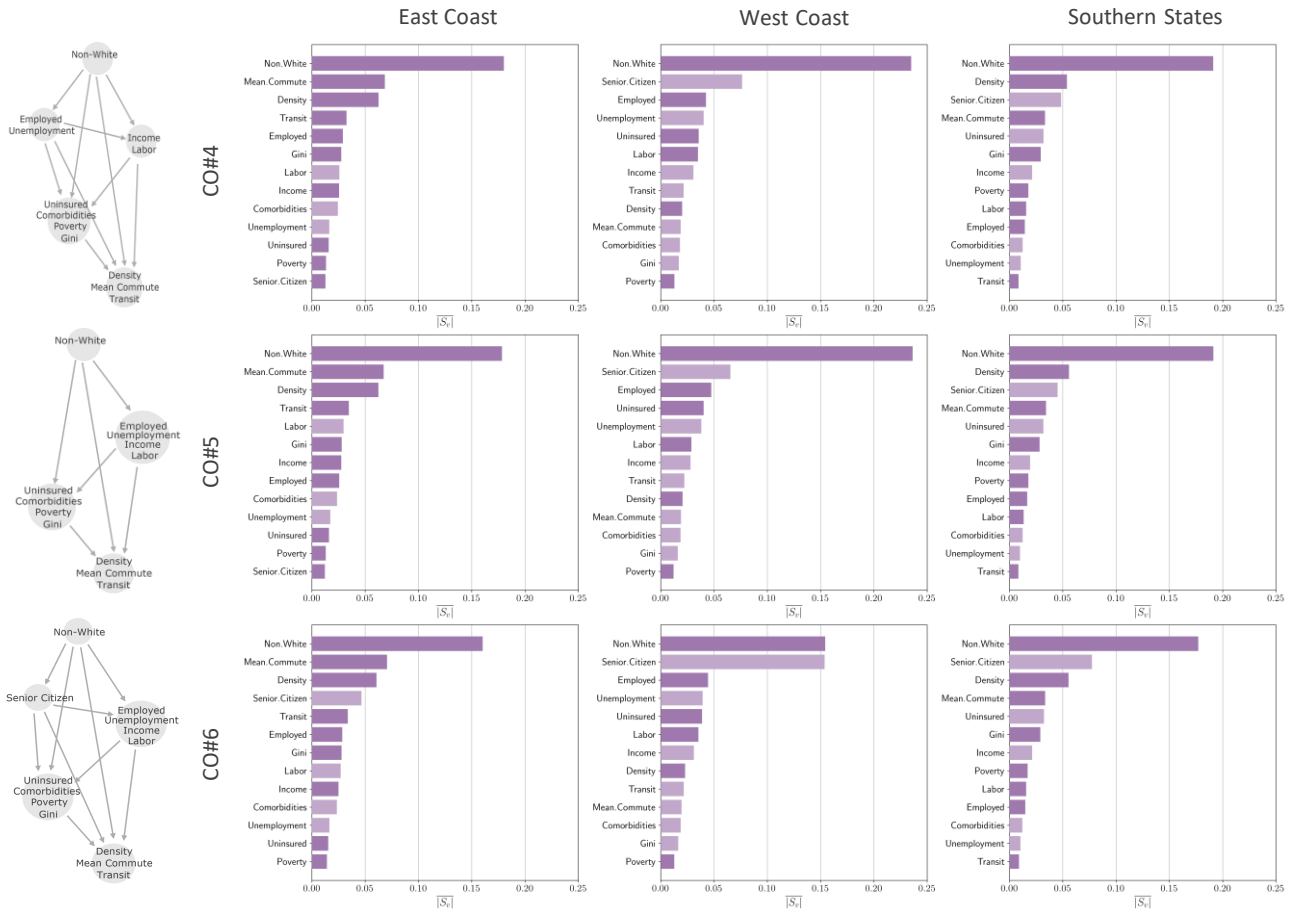
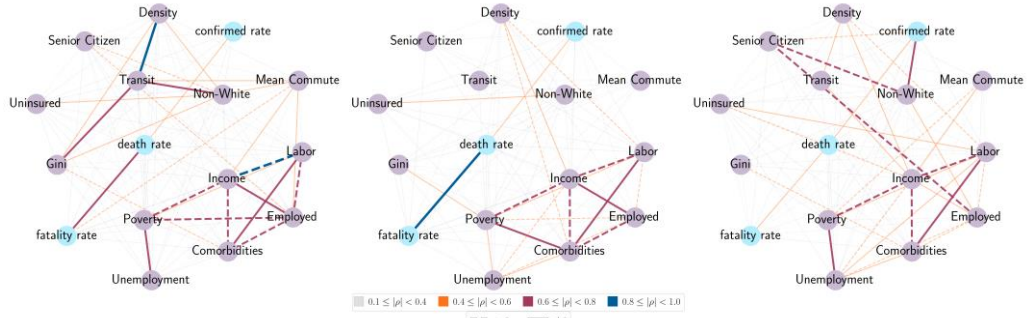


Figure S3. The causal connection between the socioeconomic metrics that we consider in this work and confirmed case rates in various regions of the USA from February 2020 to July 2020. The states included in the East Coast region are District of Columbia, New Jersey, Rhode Island, Massachusetts, Connecticut, Maryland, Delaware and New York. The ones in the West Coast region are California, Oregon and Washington. The states in the Southern States region are Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia. The networks in the top row show the correlations between the different metrics in the three regions. The graphs in the left column of the plot show the three partial causal orderings that we consider in this analysis. Each panel of the array of bar plots show the hierarchy of the metrics in terms of mean absolute causal SHAP values, $|S_v|$, for a particular region and causal ordering. The lighter shaded bars indicate a negative effect of the metric on the confirmed case rates while the darker bars indicate a positive effect.

Phase II
July 2020
to
January 2021

correlations



causal ordering


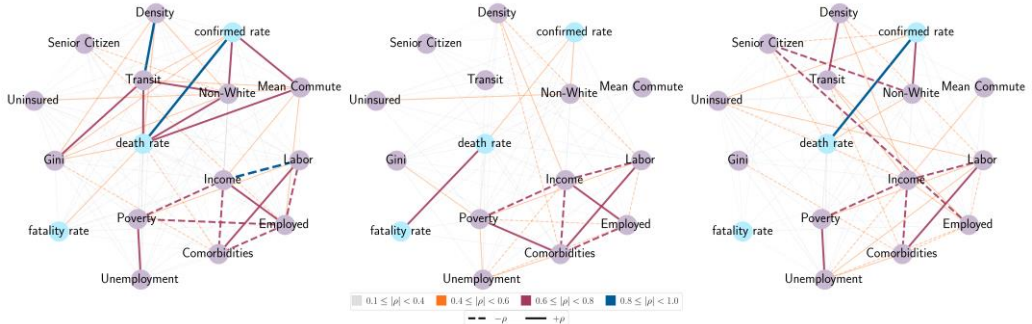


Figure S4. The causal connection between the socioeconomic metrics that we consider in this work and confirmed case rates in various regions of the USA from July 2020 to January 2021. The states included in the East Coast region are District of Columbia, New Jersey, Rhode Island, Massachusetts, Connecticut, Maryland, Delaware and New York. The ones in the West Coast region are California, Oregon and Washington. The states in the Southern States region are Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia. The networks in the top row show the correlations between the different metrics in the three regions. The graphs in the left column of the plot show the three partial causal orderings that we consider in this analysis. Each panel of the array of bar plots show the hierarchy of the metrics in terms of mean absolute causal SHAP values, $|S_v|$, for a particular region and causal ordering. The lighter shaded bars indicate a negative effect of the metric on the confirmed case rates while the darker bars indicate a positive effect.

Phase I
February 2020
to
July 2020

correlations



causal ordering

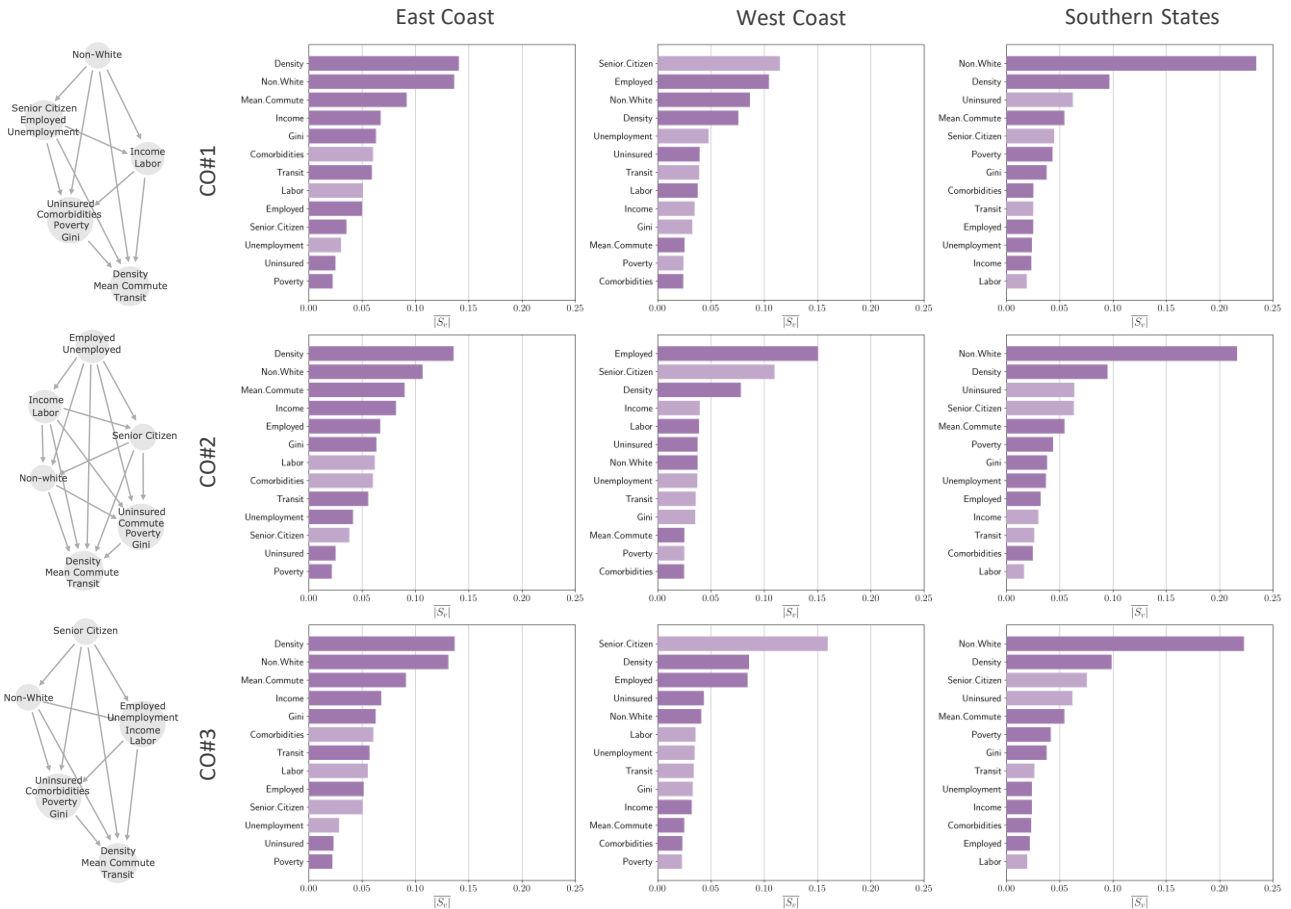
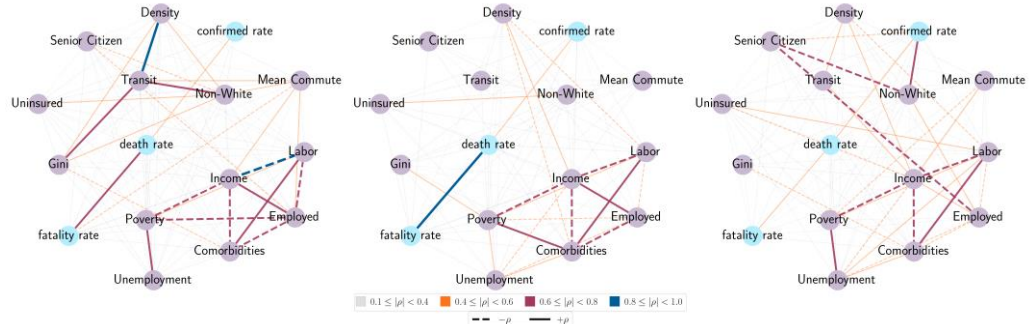


Figure S5. The causal connection between the socioeconomic metrics that we consider in this work and death rates in various regions of the USA from February 2020 to July 2020. The states included in the East Coast region are District of Columbia, New Jersey, Rhode Island, Massachusetts, Connecticut, Maryland, Delaware and New York. The ones in the West Coast region are California, Oregon and Washington. The states in the Southern States region are Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia. The networks in the top row show the correlations between the different metrics in the three regions. The graphs in the left column of the plot show the three partial causal orderings that we consider in this analysis. Each panel of the array of bar plots show the hierarchy of the metrics in terms of mean absolute causal SHAP values, $|S_v|$, for a particular region and causal ordering. The lighter shaded bars indicate a negative effect of the metric on the confirmed case rates while the darker bars indicate a positive effect.

Phase II
July 2020
to
January 2021

causal ordering

correlations



East Coast

West Coast


Southern States



Figure S6. The causal connection between the socioeconomic metrics that we consider in this work and death rates in various regions of the USA from July 2020 to January 2021. The states included in the East Coast region are District of Columbia, New Jersey, Rhode Island, Massachusetts, Connecticut, Maryland, Delaware and New York. The ones in the West Coast region are California, Oregon and Washington. The states in the Southern States region are Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia. The networks in the top row show the correlations between the different metrics in the three regions. The graphs in the left column of the plot show the three partial causal orderings that we consider in this analysis. Each panel of the array of bar plots show the hierarchy of the metrics in terms of mean absolute causal SHAP values, $|S_v|$, for a particular region and causal ordering. The lighter shaded bars indicate a negative effect of the metric on the confirmed case rates while the darker bars indicate a positive effect.

Phase I
February 2020
to
July 2020

correlations



causal ordering

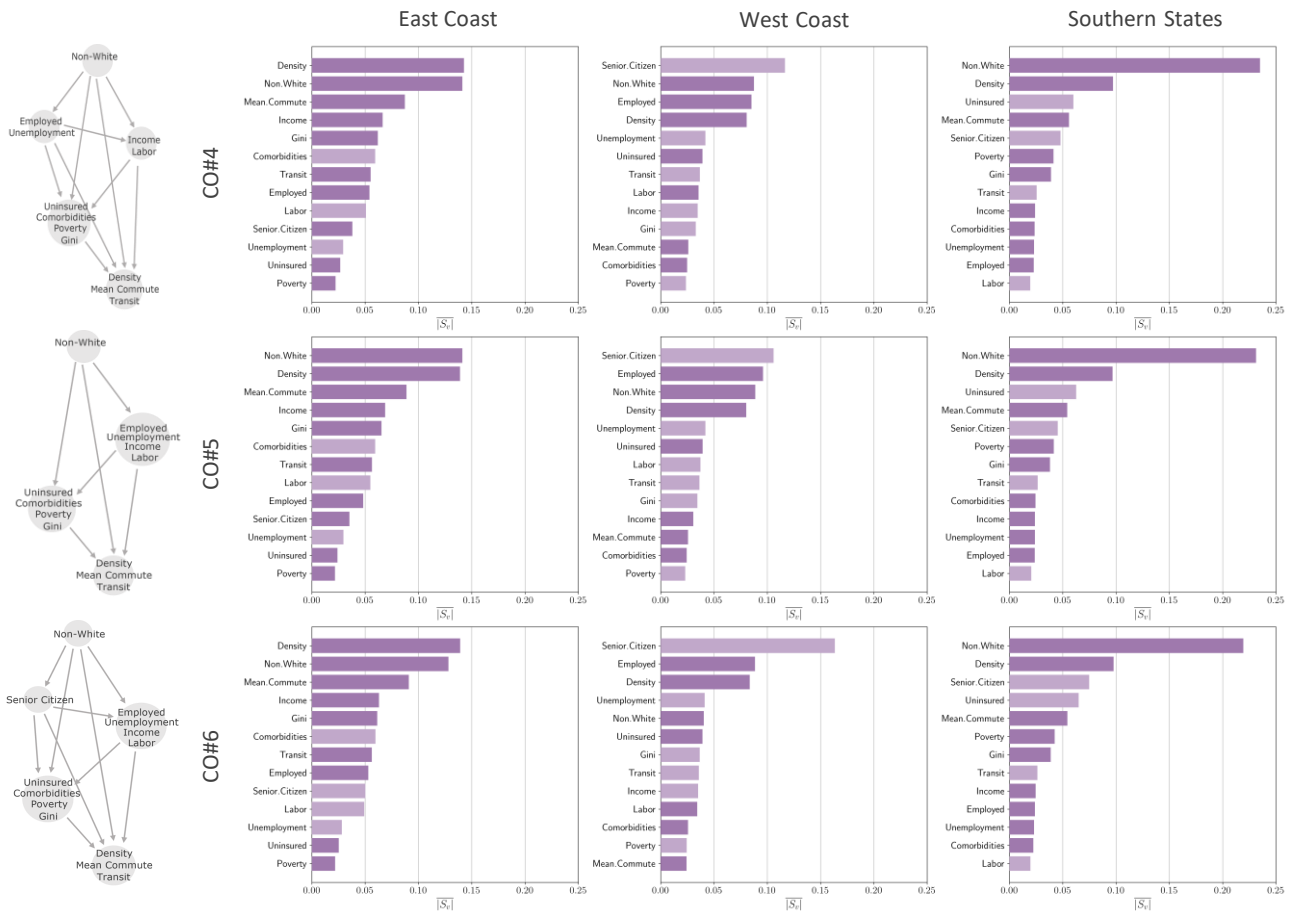



Figure S7. The causal connection between the socioeconomic metrics that we consider in this work and death rates in various regions of the USA from February 2020 to July 2020. The states included in the East Coast region are District of Columbia, New Jersey, Rhode Island, Massachusetts, Connecticut, Maryland, Delaware and New York. The ones in the West Coast region are California, Oregon and Washington. The states in the Southern States region are Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia. The networks in the top row show the correlations between the different metrics in the three regions. The graphs in the left column of the plot show the three partial causal orderings that we consider in this analysis. Each panel of the array of bar plots show the hierarchy of the metrics in terms of mean absolute causal SHAP values, $|S_v|$, for a particular region and causal ordering. The lighter shaded bars indicate a negative effect of the metric on the confirmed case rates while the darker bars indicate a positive effect.

Phase II
July 2020
to
January 2021

correlations



causal ordering

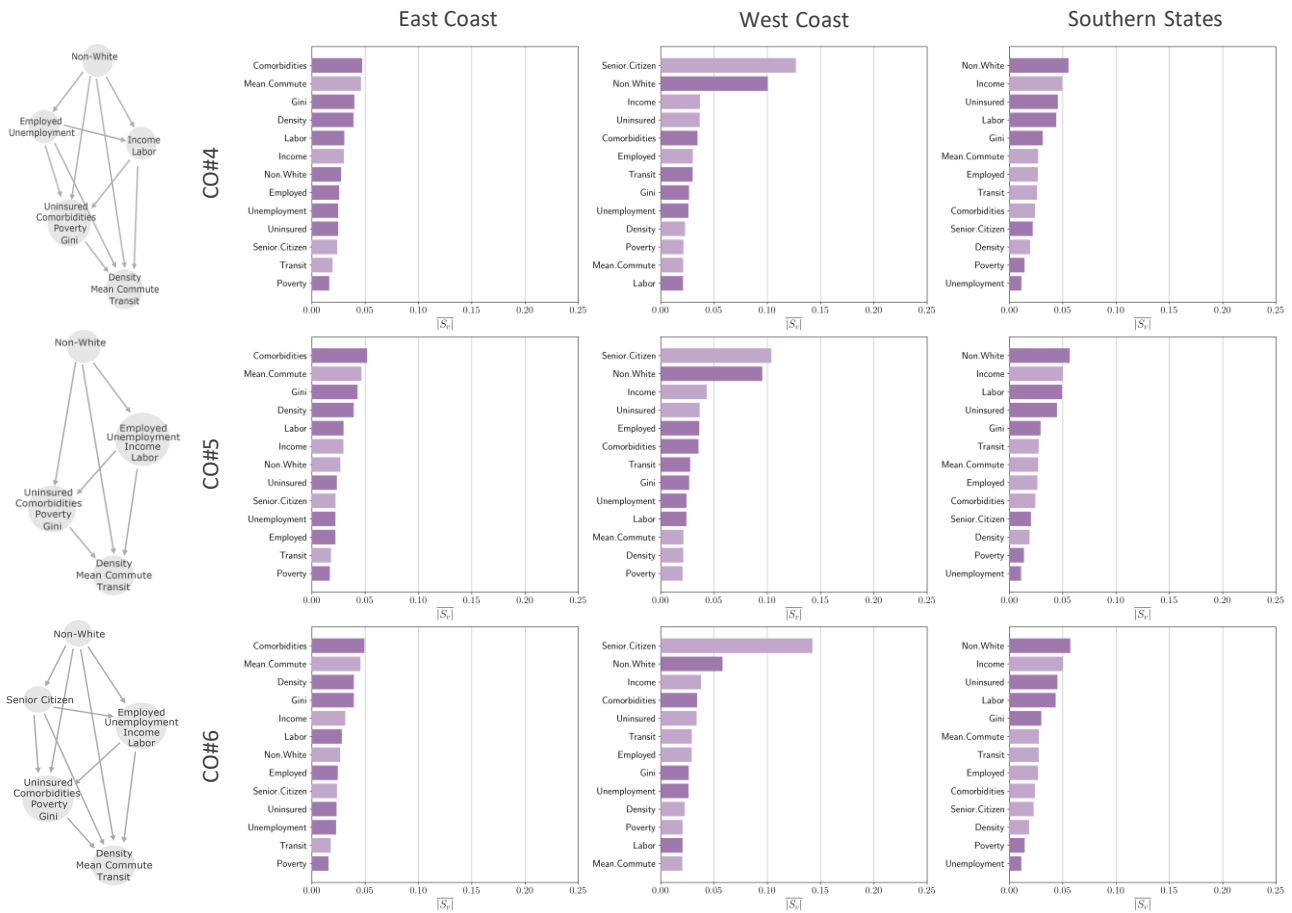


Figure S8. The causal connection between the socioeconomic metrics that we consider in this work and death rates in various regions of the USA from July 2020 to January 2021. The states included in the East Coast region are District of Columbia, New Jersey, Rhode Island, Massachusetts, Connecticut, Maryland, Delaware and New York. The ones in the West Coast region are California, Oregon and Washington. The states in the Southern States region are Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia. The networks in the top row show the correlations between the different metrics in the three regions. The graphs in the left column of the plot show the three partial causal orderings that we consider in this analysis. Each panel of the array of bar plots show the hierarchy of the metrics in terms of mean absolute causal SHAP values, $|S_v|$, for a particular region and causal ordering. The lighter shaded bars indicate a negative effect of the metric on the confirmed case rates while the darker bars indicate a positive effect.

References

1. Rust, J. & Phelan, C. How social security and medicare affect retirement behavior in a world of incomplete markets. *Econometrica* **65**, 781, DOI: [10.2307/2171940](https://doi.org/10.2307/2171940) (1997).
2. Peterson, C. L. & Murphy, G. Transition from the labor market: Older workers and retirement. *Int. J. Heal. Serv.* **40**, 609–627, DOI: [10.2190/hs.40.4.c](https://doi.org/10.2190/hs.40.4.c) (2010).
3. Ono, Y. & Sullivan, D. Manufacturing plants' use of temporary workers: An analysis using census microdata. *Ind. Relations: A J. Econ. Soc.* **52**, 419–443, DOI: [10.1111/irel.12018](https://doi.org/10.1111/irel.12018) (2013).
4. Bélanger, A. Diversity explosion: How new racial demographics are remaking america. *Can. Stud. Popul.* **43**, 166, DOI: [10.25336/p69s3v](https://doi.org/10.25336/p69s3v) (2016).