

How to Stop Contagion: Applying Network Science to Evaluate the Effectiveness of COVID-19 Vaccine Distribution Plans

Olga V. Chyzh, University of Toronto

President Trump's haphazard decision to delegate COVID-19 vaccine distribution to US states set up conditions for evaluating state-level vaccine prioritization policies using a quasi-experimental design. Despite agreement on the goal, state-formulated vaccine distribution plans diverged beyond initial priority groups: some prioritized using mortality risks only (i.e., age), while others also included several high-exposure risk groups. After establishing that this divergence was driven by stochastic rather than systematic factors, I leverage it as an identification strategy to test a key insight from network theory: reducing contagion requires disabling the transmission potential of the most connected actors. Given this, I argue that early prioritization of high-exposure risk groups, especially public-facing essential workers, led to a greater reduction in COVID-19 cases than prioritization based solely on mortality risks. Analysis of daily COVID-19 data in a matched sample of Oregon and California counties shows strong support for this hypothesis.

The 2020 COVID-19 pandemic put public health policy making under the microscope of public scrutiny. Fueled by the anguish of social isolation brought about by closures and travel restrictions, the public poured over every aspect of each new policy, from economic impact, to fairness, to public safety. The combination of urgency and public outcry resulted in wild policy oscillation, with mask mandates, curfews, and quarantines implemented and lifted, often seemingly at random. Once vaccines became available, health authorities were faced with a yet more pressing policy decision on how to distribute the initially scarce vaccine supplies. In this article, I show how the policy-making processes around early vaccine distribution in the United States led to a divergence in vaccine prioritization, resulting in substantial variation in public health outcomes.

Combined with President Trump's haphazard decision to delegate COVID-19 vaccine distribution to US states, these conditions set up a unique opportunity for evaluating al-

ternative vaccine prioritization policies, as well as the general mechanisms of preventing contagion in social and political networks. Lacking experience and expertise in drafting and implementing new vaccine distribution protocols, states were largely unprepared for the task. Despite a general agreement that vaccine prioritization plans should be designed so as to minimize deaths and hospitalizations, public health officials lacked a clear understanding or consensus on the best way to achieve these goals. Rather broad guidelines from the Centers for Disease Control (CDC) forced states to define vaccine priority groups using their own discretion. As a result, prioritization lists diverged as soon as vaccination campaigns moved beyond groups with the highest mortality risk.¹

The broad strategy was to prioritize vaccine access on the basis of vulnerability, proxied by age, from oldest to youngest.² Some US states, however, also prioritized one or more groups with the highest risk of exposure: K-12 educators and staff, incarcerated individuals, and agriculture and food

Olga V. Chyzh (olga.chyzh@utoronto.ca) is an associate professor in the Department of Political Science at the University of Toronto, Toronto, ON M5S 3G3, Canada.

Replication files are available in the *JOP* Dataverse (<https://dataverse.harvard.edu/dataverse/jop>). The empirical analysis has been successfully replicated by the *JOP* replication analyst. An online appendix with supplementary material is available at <https://doi.org/10.1086/726933>.

1. Priority was given to medical workers, the elderly, and residents of long-term care facilities.
2. Some states also gave priority to individuals with specific preexisting conditions.

Published online December 7, 2023.

The Journal of Politics, volume 86, number 1, January 2024. © 2023 Southern Political Science Association. All rights reserved. Published by The University of Chicago Press for the Southern Political Science Association. <https://doi.org/10.1086/726933>

ooo

employees. Thus, the primary difference among US states' vaccine priority approaches was in the relative place in the queue for these three groups.

While this divergence is apparent in hindsight, a combination of statistical and qualitative analyses reveals little evidence that it was either intentional or a result of systematic drivers, such as political interests. The results of these analyses suggest that, rather than customizing vaccine prioritization to their local political or demographic context, states simply did their best to adhere to the national guidelines, however vague. The analyses also suggest that variation in state-formulated priority lists was due to stochastic, rather than systematic, factors, such as variation in interpreting the CDC guidelines, a failure to anticipate the pace of increases in vaccine supply, and idiosyncratic delays in state-level decision-making.

Using these analyses, I argue that the divergence in state-formulated vaccine priority plans sets up conditions for testing a key network theory insight—that the bulk of transmission through a network is disproportionately channeled through only a handful of highly connected or central actors (Box-Steffensmeier et al. 2018; Granovetter 1973; Kirkland 2011; Padgett and Ansell 1993). In a service-oriented economy, such as the United States, grocery store employees are the largest occupational group engaged in the highest number of face-to-face interactions. In contrast to other service sectors, grocery stores had to remain open throughout the pandemic, when most other retail businesses and even schools reduced hours or switched to remote operation.³ As such, grocery employees function as central nodes within a transmission network of a virus that diffuses most effectively among individuals in close contact in indoor spaces.

The handful of US states that happened to assign grocery employees a higher priority in the vaccine eligibility queue essentially reduced or disabled the virus transmission potential of the most central actors in the network of human interactions.⁴ In contrast, states that gave lower priority to grocery employees, especially if they chose to prioritize on the basis of age only, effectively allocated their vaccines to the most isolated actors within the network.⁵ Individuals with

the fewest nonelective interactions could more effectively reduce the likelihood of both contracting and spreading the virus, compared to public-facing essential workers, such as cashiers at grocery or convenience stores (Milligan et al. 2021).⁶ Therefore, I expect to see lower rates of contagion in states that prioritized grocery employees earlier in the vaccine eligibility queue.

I test this prediction using a matched sample of counties from two contiguous, Democratic-governed states—Oregon and California. Both states implemented similar policy responses to the COVID-19 pandemic in terms of school closures, stay-at-home orders, and mandatory mask mandates (Adeel et al. 2020). Even more importantly for this study, both California and Oregon prioritized grocery workers in phase 1b, before the general public. The practical difference, however, is that, due to a confluence of idiosyncratic factors, Oregon lagged California in terms of the actual eligibility date for grocery employees by about a month. Consistent with the predictions of network theory, county-level statistical analysis shows that prioritizing grocery store employees led to a substantial reduction in new COVID-19 infections, and this effect grew stronger over time.

With the COVID-19 pandemic still underway, this research provides a theory- and data-informed cost-benefit analysis of giving higher priority to public-facing essential workers, such as grocery store employees. A key nuance is that, while carrying a disproportionately high potential for spreading the virus, grocery employees make up only a small fraction—less than 1%—of the population. To put this into perspective, even given the initial vaccine scarcity, vaccinating every single grocery employee would have delayed vaccine access for other groups by less than a week—a negligible delay given the substantial effect on reducing case numbers shown here.⁷

interactions exhibit a clique-ish structure: individuals within groups come in frequent contact with one another but have few contacts outside of the group (Milligan et al. 2021). Should a single member of a clique get infected, the entire clique comes under high risk. Because cliques have limited interactions with nonclique members, however, infecting a clique only marginally increases the risk of infection for the general public.

6. The pandemic, and the ensuing policy measures (e.g., stay-at-home orders), intensified two known patterns of social interaction: assortativity by age (most of daily interactions happen among individuals within the same age group) and a drastic decrease in the number of interactions, the higher one's age. Early in the pandemic, individuals of all ages reported few nonhousehold contacts; although younger adults reported an increase in work- and school-related interactions by September 2020, the elderly continued to report the highest level of social isolation (Brankston et al. 2021; Feehan and Mahmud 2021).

7. There are 313,045 grocery employees in California and 37,491 in Oregon (US Department of Labor 2020). In December 2020–January 2021, California and Oregon received about 500,000 and 50,000 vaccine doses per

3. Pickup and delivery services were unfeasible for many grocers because of added costs, a lack of experience with online platforms, and customer preference for in-store shopping.

4. Although not available at the start of vaccine prioritization, evidence that COVID-19 vaccines prevent or substantially reduce transmission of the virus has since become available (Eyre et al. 2022; Harris et al. 2021; Richerman, Meyerowitz, and Cevik 2022).

5. From the network theory perspective, prioritizing other high-exposure groups, such as K–12 students or the incarcerated, should have a more limited effect on contagion among the general public. These groups'

By fleshing out the trade-offs of different priority sequences, this article opens an informed conversation about the benefits and costs of public health policies, as they relate to political trust and participation (Mattila 2020), inequality (Lynch 2020), long-term institutional development (Gingrich and Vogler 2021), and international cooperation (Norrlof 2020). Beyond vaccine distribution, this article contributes to the study of the core political problem of the distribution of scarce resources more broadly. Remarkably, the analyses show, and explain, how the most politically charged and controversial decision-making of the pandemic—vaccine prioritization—produced a seemingly apolitical outcome. As I demonstrate, vaccine priority lists did not clearly align with the usual political cleavages. Despite the high levels of political polarization throughout Trump’s presidency and the pandemic, governors appeared to have prioritized the goal of effective policy response and demonstrating competency in a time of crisis, rather than distributing the scarce vaccine so as to buy off political supporters. With this example as a starting point, future research could explore the conditions that induce politicians to act on behalf of the entire electorate rather than cater to more narrow partisan interests.

This research also contributes to the general understanding of contagion and its pathways. Within studies of diffusion and network analysis, scholars have long applied models of disease contagion to study the spread of information (Lohmann 1994), censorship (King, Pan, and Roberts 2013), protests and repression (Siegel 2011; Van Belle 1996), social movements (Ayoub, Page, and Whitt 2021), and diffusion of policy innovation (Desmarais, Harden, and Boehmke 2015). While many studies emphasize the importance of central nodes for network diffusion processes, empirical tests of the posited causal effects are often impeded by the lack of data to approximate the counterfactual outcome (i.e., had the central nodes been removed from the network). For example, one cannot observe what a policy diffusion process across US states would look like in the absence of New York and California or what the post–World War II international alliance network would look like without the United States. The variation in the COVID-19 vaccine distribution among US states, however, allows for a unique opportunity to empirically evaluate and isolate the causal effect of a handful of highly connected actors on the transmission rate within a

week, respectively (Botkin 2021; CVAC 2021c, 6). In contrast, other essential employees prioritized by both states amounted to much larger numbers: medical employees (2.4 million in California and 0.4 million in Oregon), K–12 teachers and staff (1.3 million in California and 0.1 million in Oregon), and food and agriculture workers (3.4 million in California and 0.2 million in Oregon; Botkin 2021; CVAC 2020a, 17; CVAC 2020b, 18; CVAC 2020d, 15; OHA 2020, 2022).

network by varying the central actors’ transmission ability between the treatment group (networks with the vaccinated central actors) and the control group (networks with the unvaccinated central actors), before and after the start of the treatment (vaccine eligibility for grocery employees), while holding all else constant (via matching). This design permits a direct comparison between the observations in the treatment and control groups.

The article proceeds in the following way. After describing alternative vaccine prioritization plans, and the process that led to their formulation, I contextualize each within network theory. Next, I compare the effectiveness of vulnerability-based prioritization to an approach that targets public-facing essential workers, using a simulation experiment. I then introduce the data and research design, present and discuss the statistical analysis, and conclude.

THE POLITICS BEHIND VACCINE ALLOCATION

As the Food and Drug Administration gave emergency approval to two COVID-19 vaccines in late November 2020, numerous groups—from teachers to fast food employees, to morticians, to commercial pilots—made their bids for early vaccine priority. With the Trump administration delegating vaccine distribution to states, all eyes turned to the governors’ offices.⁸ The CDC (2021a) issued a set of broad recommendations: (1) “decrease death and serious disease as much as possible,” (2) “preserve functioning of society,” and (3) “reduce the extra burden COVID-19 had on people already facing disparities.” Given the high stakes and climbing death rates, these guidelines were interpreted as a justification for offering vaccine priority to the most at-risk individuals, such as medical workers and the elderly.

Once the vaccination campaign moved beyond these groups, however, the CDC guidelines as to further prioritization lists were less clear. Notably, in light of persistent vaccine shortages, the first two goals were somewhat contradictory: decreasing deaths required continuing to prioritize vaccine access for the most vulnerable, whereas maintaining a functioning society necessitated inoculating individuals on the basis of occupational risk of exposure.⁹ Without any additional clarifications from the CDC, states formulated vaccine priority lists using their best judgment.

One strategy was to continue prioritizing vaccine distribution on the basis of vulnerability, using age, from oldest to

8. In the United States, governors assumed the main decision-making power related to vaccine distribution (Adeel et al. 2020).

9. In the United States, vaccine supplies remained scarce until around May 1, 2020, when almost all states expanded eligibility to all adults (*American Journal of Managed Care* 2021).

youngest, as a proxy for mortality risk. Governor Holcomb of Indiana, for instance, justified his state’s strict adherence to vulnerability-based prioritization, pointing to a strong correlation between age and hospitalization and death rates in his state (Salameh 2021). At the time, individuals age 50 and older, while accounting for just over 35% of the Indiana’s population, made up 80% of the COVID-19 hospitalizations and 98% of all COVID-19 deaths (Darling 2021).¹⁰

An alternative strategy was to also prioritize individuals with the highest risk of occupational exposure. In particular, this strategy raised priority for essential frontline workers—individuals whose jobs were crucial to maintaining essential services and who, because of the nature of their work, could not maintain a safe distance from their coworkers or other individuals.¹¹ Vaccinating people within these sectors ensured a continuation of essential services while reducing the risk of infection for the individuals performing these services.¹² Some critiqued this prioritization plan on the grounds that individuals in these occupational groups tend to skew younger and face lower overall health risks from contracting the virus.

Devising an effective, inclusive, and equitable vaccine prioritization plan became a test of a governor’s competency, with the expectations rising as vaccine scarcity became more apparent. In these conditions, most governors created numerous vaccine task forces, including vaccine advisory committees that held regular public meetings, so as to enhance the transparency of decision-making. In contrast to earlier pandemic-related health policies, wrought with deep partisan divisions (Neelon et al. 2021; VanDusky-Allen and Shvetsova 2021), COVID-19 vaccine distribution plans exhibited a high degree of consensus. At least at the administrative level, both Democratic- and Republican-led states agreed that the immediate goal was to distribute vaccines in a way that would minimize hospitalizations and deaths, with the broader goal of controlling the spread of the virus by vaccinating the largest number of people within the shortest possible time frame.

Under these constraints, the outcome of the most scrutinized decision-making of the pandemic—which groups will receive priority in the vaccine line—turned out to look re-

10. These numbers were consistent with national statistics (CDC 2020).

11. The CDC defines frontline essential workers as distinct from essential health care workers. Specifically, frontline essential workers are “the subset of essential workers likely at highest risk for work-related exposure . . . because their work-related duties must be performed on-site and involve being in close proximity (<6 feet) to the public or to coworkers” (Dooling et al. 2021, 1657). Examples include grocery and manufacturing workers.

12. Among essential workers, public-facing employees, such as grocery workers, carry a uniquely high risk of spreading the virus to the general public (Milligan et al. 2021).

Table 1. Comparison of States with Early/Late Vaccine Prioritization for Grocery Workers

Variable	On/Before March 1	After March 1
Republican governor	.31	.61 ⁺
State GDP/capita, logged	4.13	3.97 [*]
Median age	37.72	38.62
State population, logged	15.32	15.11
Median income, thousands	65.60	62.24
Unemployment rate	.35	.36
Percentage BA degree	.20	.19
Urbanization	.64	.59
Percentage Black	.14	.10
Percentage Latin	.16	.11
Biden’s margin	.06	−.03
Cumulative COVID/1,000 (on December 16, 2020)	12.20	11.86
χ^2		15.8
df		12
<i>p</i>		.20

⁺ *p* < .1.

^{*} *p* < .05.

markably nonpartisan, that is, orthogonal to constituency preferences of decision makers. For instance, table 1 displays a mean comparison between the states that opened vaccine eligibility to grocery employees before March 1, 2021 (the date of eligibility in California), and those with later eligibility dates, on key demographic and other variables. The demographic variables here proxy possible political divides, such as urban/rural, education, ethnic/racial, economic, and age. Had vaccine priority lists been designed so as to deliver direct benefits (vaccines) to the most likely voters, one would be able to discern predictable systematic differences across these dimensions. The two groups, however, look very similar. An overall (χ^2) balance test fails to reach statistical significance (*p* = .2), which means that we cannot reject the null that there are no differences between the two groups.¹³

The two groups differ on only two variables, neither of which is demographic: GDP/capita, which is a proxy for state capacity, and the party of the governor. The difference in the GDP/capita is controlled for in the statistical analysis; in the matched sample, counties are also matched on this variable. Had the difference between Republican- and Democratic-led states stemmed from a strategy to reward, or electorally

13. In app. C (apps. A–F are available online), I replicate these results using survival analyses that treat time to eligibility for (a) grocery employees and (b) individuals age 65 and above as the dependent variables.

target, constituents, we would expect differences in terms of the demographic variables.¹⁴ Moreover, the timing of vaccine access for grocery workers does not appear to result in large partisan benefits or costs: grocery employees are not a sizable voting bloc, are characterized by low levels of political participation, and are not organized/unionized in most states (US Department of Labor 2022).¹⁵

Qualitative analysis of vaccine advisory committee materials in Oregon and California, presented in appendixes D and E, further reaffirms this somewhat surprising result, as well as offers more insight as to the process that led to it. In particular, the case studies reveal a high degree of similarity in the two states' decision-making, a high degree of reliance on health authorities, and a general reluctance to deviate from the national guidelines (CDGW 2021; CVAC 2020c, 2021a, 2021b, 2021d; OHA 2021a, 2021b).

Early in the process, for example, both states considered implementing a combination of an age-based and an exposure-based approach. In the end, the determining factor turned out to be the timing of a sudden change in the CDC guidelines with respect to the progress in each state's planning. With no advance notice to the states, on January 12, 2021, the CDC issued a revised guideline that the states prioritize the vaccine on the basis of age alone and expand eligibility to individuals age 65 and above (Brown 2021b; CVAC 2021c, 3). Both states complied with this revised recommendation. At the time of the change, however, California happened to be slightly ahead of Oregon in its planning and progression through the vaccine priority list: in fact, California's Governor Newsom had already announced eligibility dates for several groups, including grocery employees, and several small counties had already been contacting these groups for a few days.

14. This difference is not well explained by an electoral impetus to prioritize the elderly, as the two groups are not significantly different in terms of median age. Moreover, in the 2020 presidential election, Republican vote share advantage among individuals age 65 and over was 4 percentage points (between the two main parties)—the smallest of the four age groups as broken down by Igielnik, Keeter, and Hartig (2021). Alternatively, this difference may have stemmed from a belief among Republican governors that prioritizing on the basis of age minimized deaths. Since prioritizing individuals over 64 is uncorrelated with state demographics, in particular, median age, this type of potential selection into treatments is akin to an example, given by Imbens and Rubin (2015, 265), in which a doctor uses information on the amount of insurance coverage when deciding whether to prescribe a patient drug A or B. The treatment is exogenous, as long as insurance was purchased before the diagnosis. In other words, unconfoundedness still holds, as long as the unobserved differences that resulted in unit assignment to the treatment/control group "are independent of the potential outcomes, conditional on observed covariates" (265).

15. Since both California and Oregon are Democratic governed and have similar proportions of grocery workers, this difference is irrelevant for the main analysis.

As a result, while Oregon responded to the change in the CDC guidelines by simply moving to an age-based approach, California kept the already announced priority level for several exposure-based groups, although in practice, these groups were still pushed back in line behind the newly added age-based priority group. In other words, the choice of prioritization approach was not as set in stone as it may appear without context. Had the change in the CDC guidelines not taken place, we may well have seen more states, including Oregon, give earlier priority to grocery employees.

In summary, the evidence suggests that vaccine distribution was less about divvying up resources among political supporters and more about demonstrating competence, or minimizing political fallout, in a situation with high political risk and uncertainty (Ashworth, Bueno de Mesquita, and Friedenber 2018; Gasper and Reeves 2011; Quiroz Flores and Smith 2013).

NETWORK THEORY AND VACCINE PRIORITIZATION

From the perspective of network theory, the nature of contagion—its speed, reach, and main pathways of transmission—depends on the local structures within the network (Chyzh and Kaiser 2019). Rather than simply a function of network density, the rate of contagion depends on the presence of a few highly connected actors with crosscutting connections to otherwise disconnected parts of the network (Feld 1991; Granovetter 1973; Padgett and Ansell 1993). What ultimately determines an actor's transmission potential is its level of connectedness or centrality.¹⁶ It follows that, as long as there is some level of heterogeneity in actors' centrality (i.e., some actors have more direct or indirect connections than others), disabling the transmission potential of the network's most central actors is the fastest and most effective way to reduce or stop contagion.

By implication, individuals who are the most likely to limit their social interactions so as to protect themselves against the virus (e.g., due to a preexisting condition), effectively self-select into less central network positions. Before widespread vaccine availability, most of the general population drastically reduced their social interactions, with the highest level of social isolation reported by individuals of age 65 and older (Brankston et al. 2021; Feehan and Mahmud 2021). Since essential-service providers are the toughest social links to eliminate, giving them vaccine priority is also an effective way to reduce the risk of infection for the most vulnerable

16. Measures of centrality include degree, eigenvector, closeness, and betweenness centrality (Bonacich 1972; Patty and Penn 2017). See app. A for formal definitions.




Figure 1. Centrality in the interaction network

individuals.¹⁷ Conversely, age-based vaccine prioritization is less likely to reduce or stop the spread of the virus, as the key transmitters, individuals in entry-level positions, such as front-facing grocery employees, tend to fall within a younger age range.

To explore how these well-known theoretical insights work within this specific application, I compare the effects of vaccine prioritization plans (centrality vs. vulnerability based) using a simulation experiment based on a real-world human interaction network. I start by defining the nodes in a network as a set of n actors, $i \in \{1, 2, \dots, n\}$, and one or more interactions between each pair of nodes i and j as a link, $l_{ij} = 1$. To assess a node's transmission potential within a network, I use degree centrality, which is defined as the total number of its direct links, or $\sum_{j:j \neq i} l_{ij}$ (Bonacich 1972; Patty and Penn 2017).¹⁸

17. This logic also applies to reducing contagion for individuals with the highest vulnerability due to living arrangements, such as nursing homes. From a network perspective, individuals who work or live in close proximity to others are a clique, as discussed in n. 5. Assuming that there is not enough vaccine to inoculate every member of the clique, the highest risk to the clique comes from its members' external contacts (see app. B). Thus, the best protection is to vaccinate individuals who are indispensable for the cliques' functioning—medical workers and public-facing essential workers. This implication is even stronger given that, as we now know, the COVID-19 vaccines' effectiveness declines with individuals' vulnerability (Andrews et al. 2021; Nunes et al. 2021; Salmeron Rios et al. 2022).

18. Degree centrality is the most common network metric for modeling nodes' transmission potential. In app. B, I replicate the experiment using other

I begin with Coleman's (1964) widely used data set on a high school friendship network.¹⁹ In this network, two individuals are connected by a link if at least one of them named the other as someone with whom they frequently interacted.²⁰ Figure 1 shows the distribution of individuals' degree centralities in these data.

Figure 2 shows the node with the highest value on degree centrality in dark red. Without loss of generality, assume that a virus spreads on contact with certainty. If the dark red node is the initial carrier (patient 0), then its direct neighbors (shown in red) are the next to contract the virus, after which the virus spreads to the nodes that are reachable through a shortest path of length 2 (shown in orange) as part of the second round.²¹ The figure shows that, for this network, choosing the node with the highest degree centrality as patient 0

common centrality metrics, such as eigenvector, closeness, and betweenness centrality. The choice of centrality measure does not substantively affect the results, as in social interaction networks actors who rank high on one centrality measure tend to also rank high on other centrality measures.

19. The data set consists of a friendship network among 73 boys in a small high school in Illinois in spring 1958 (Coleman 1964).

20. I transform the original directed network data into a nondirected symmetric network.

21. For each pair of nodes, the shortest path, d_{ij} between i and each other node j in the network is the total number of links on the path between them. For example, the shortest path between two nodes that are connected by a link is of length 1. If i and j do not share a direct link but i is connected to l and l is connected to j , then the shortest path between i and j is of length 2.




Figure 2. Contagion in the interaction network

results in 40 infected individuals at the end of two rounds of contagion.

Now suppose there are 10 available vaccines. Without loss of generality, assume getting a vaccine makes an individual both immune to the virus and unable to transmit it. Figure 3 shows the spread of the virus under both the age- and exposure-based prioritization scenarios. Groups receiving priority based on age are less likely to be a part of the workforce and, hence, have the greatest flexibility to limit their risk of exposure by reducing their number of social interactions (i.e., they are

willing and able to do so).²² Therefore, to reflect this prioritization strategy within the simulation experiment, I define *vulnerable individuals* as those with the lowest degree centrality. Age-based prioritization, by implication, consists of allocating the 10 available vaccines on the basis of degree centrality, from lowest to highest. Vaccinated individuals are shown in blue. I refer to this part of the simulation as the Low Degree scenario.

To reflect the exposure-based prioritization strategy, I define the individuals with the greatest risk of exposure as those with the highest degree centrality. Occupational risk vaccine prioritization, therefore, consists of allocating the 10 vaccines to the most central nodes in the network (other than patient 0). This is the High Degree scenario.

The demonstration shows that the second vaccine prioritization strategy leads to a drastic reduction of contagion in the network: it results in nine fewer cases of contracting the virus (a 30% reduction). For a network of 69 individuals (after removing isolates), this is a difference between infecting 57% versus 43% of the nodes in the network. If we think of the nodes with the lowest degree centrality (marked with a green border color in fig. 3) as vulnerable individuals, this demonstration also shows that the second vaccine prioritization scenario fares well at protecting these individuals: only one such node contracted the virus under the second scenario.

22. The network science focus limits our insights to groups of individuals who are identifiable as a function of their network connections. Making separate inferences for vulnerable individuals who are either unable or unwilling to self-isolate by reducing their number of direct contacts is beyond the scope of this article.

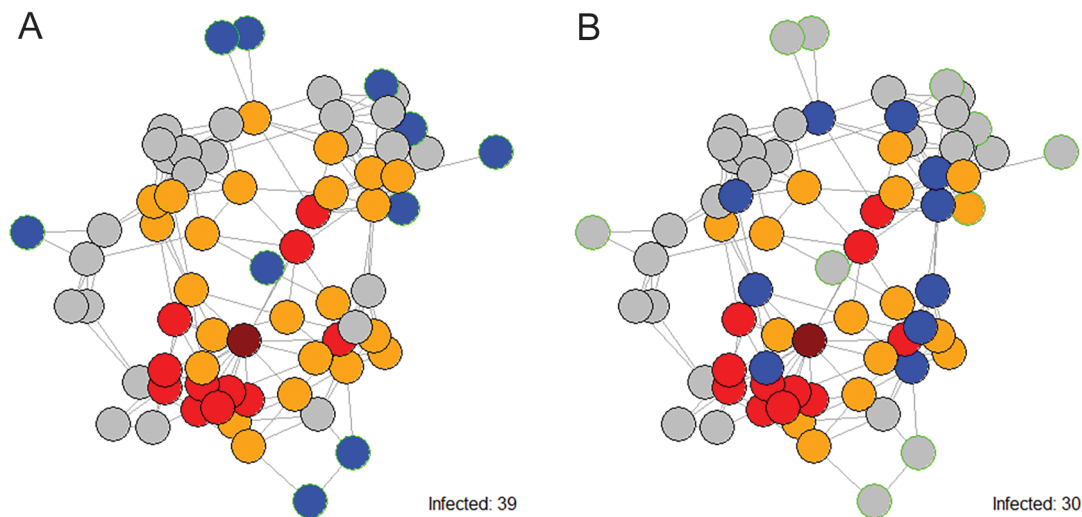


Figure 3. Vaccination scenarios: A, low degree; B, high degree

Table 2. Round 2 Summary for 10,000 Simulations of the Interaction Network among 73 Individuals

	Vaccinated	Infected	Not Infected	Vulnerable Infected
No vaccine	0	39.35 (5.40)	28.65 (5.40)	3.24 (1.51)
Low degree	10	36.04 (5.07)	21.96 (5.07)	0
High degree	10	24.13 (4.74)	33.87 (4.74)	2.38 (1.38)

Note. Means over 10,000 simulations. Numbers in parentheses are standard deviations. Not infected does not include vaccinated.

To generalize from this example, I perform the following Monte Carlo experiment. I start by estimating an exponential random graph model (ERGM; Wasserman and Faust 1994), in which the interaction network is the dependent variable, and the network parameters of interest are the baseline link probability (edges), the tendency toward open triangles (2 stars), and closed triangles (geometrically weighted edge-wise shared partnerships)—the most common ERGM specification that models the basic features of a human interaction network.²³ The estimates of this model are shown in appendix B. I then use the estimates from this model to simulate the interaction network 10,000 times and repeat the analysis done on the original interactions network on these simulated networks.

Table 2 shows the means (and standard deviations) of the number of individuals infected/not infected at the end of the second round of contagion (individuals connected to patient 0 directly or via one intermediary) in 10,000 simulations of the interaction network. Just as before, I denote the 10 individuals with the lowest number of direct (lowest degree centrality) connections as “vulnerable” individuals. These are the individuals vaccinated in the Low Degree scenario. The last column of table 2 shows the mean (and standard deviation) of the number of vulnerable individuals who are infected under each of the vaccination scenarios. I also perform a simulation for a No Vaccine scenario for comparison.

Under the Low Degree scenario of giving vaccine priority to the vulnerable individuals, the number of individuals infected at the end of the second round is only slightly lower than that under the No Vaccine scenario (roughly 36 vs. 40). Under the High Degree scenario, the number infected is substantially lower: 12 fewer than under the Low Degree scenario, or a 32% decrease.

These results, of course, are based on the simplifying assumption that the vaccine is 100% effective at preventing both virus contraction and transmission and that the effectiveness does not change with vulnerability. Given that, in actuality, vaccine effectiveness decreases with vulnerability (Andrews et al. 2021; Nunes et al. 2021; Salmeron Rios et al. 2022), the results error on the conservative side, downplaying the difference in effectiveness between the two vaccine priority plans. If we relaxed this assumption and allowed the probability of virus transmission to increase with vulnerability, even for the vaccinated, the difference between the numbers of individuals infected under each vaccine priority plan would be even larger (i.e., among the vaccinated, vulnerable individuals would contract the virus at a higher rate than the more central nodes).

In summary, the results of the simulation experiment confirm that, in terms of reducing contagion, a centrality-based (i.e., exposure-based) vaccine prioritization plan substantially outperforms that of allocating vaccines based solely on vulnerability. The simulation highlights the effectiveness of vaccinating essential frontline workers for reducing the spread of COVID-19, compared to other vaccine prioritization strategies, under the conditions of vaccine scarcity or distributional constraints. States that prioritized vaccine access for essential frontline workers were able to reduce transmission by eliminating a key source of contagion—inoculating the individuals who came into contact with both the greatest number of people and societal groups. This leads to the following hypothesis:

Research Hypothesis. Administrative units that prioritize individuals with higher centrality will experience fewer infection cases.

RESEARCH DESIGN

To test the research hypothesis, I leverage variation in vaccine priority lists between California and Oregon, two neighboring Democratic-governed states that were similar in their prevaccine COVID-19 policies (Adeel et al. 2020). While early

23. ERGMs are an estimation approach for modeling the probability of observing a network with a given set of endogenous statistics, such as the total number of edges, open or closed triangles, or other network features (Hunter et al. 2008; Robins et al. 2007).

on the two states followed a similar vaccine distribution strategy as the rest of the United States (i.e., prioritizing the elderly and medical workers), California was among the first states to extend vaccine eligibility to grocery store employees (on March 1), whereas in Oregon, grocery store employees did not become eligible for vaccination until almost a month later (on March 29). The unit of analysis is the county-day. The dependent variable is a logged seven-day rolling average of the number of new COVID-19 cases, obtained from Johns Hopkins University (Gassen 2021).²⁴

The independent variable is an interaction between California and Day of Treatment, in which California is a binary variable that equals 1 for California and 0 for Oregon. Day of Treatment is a count variable that starts at 1 on March 14, two weeks since grocery employees became eligible for vaccination in California—the date when those who had received the first dose would have achieved partial immunity (between 50% and 80%; Bernal et al. 2021; Polack et al. 2020). The estimation equation is

$$\begin{aligned} \log(\text{New Cases}) = & \beta_0 + \delta_0 \text{Day of Treatment} \\ & + \beta_1 \text{California} + \delta_1 \text{California} \\ & \times \text{Day of Treatment} + \text{other factors.} \end{aligned}$$

This, of course, is a textbook example of the difference-in-difference design (Wooldridge 2015, 407–12). The treatment here is measured as the day since the first grocery workers had developed partial immunity in California (Day of Treatment). The estimation parameter δ_0 is the average difference between the periods before and after the start of the treatment for the control group, β_1 is the average difference between the two groups before the treatment, $\beta_1 + \delta_1$ is the average difference between the two groups after the start of the treatment, and δ_1 is the difference-in-difference coefficient that gives the average difference attributable to the treatment (i.e., the average effect of expanding vaccine eligibility to grocery employees in California).

I control for a number of pretreatment county-level demographic variables that may influence COVID-19 contagion, including the logged cumulative number of reported COVID-19 cases per 1,000 residents (as recorded on December 16, 2020—the day before the first vaccines were administered in the two states), logged GDP (2019 USD), logged population, unemployment rate, percentage of the population that holds at least a bachelor's degree, urbanization, percentage of black and other racial minorities, percentage of Hispanic/

Latino population, percentage of foreign population, Biden's percentage margin in the two-party vote in the 2020 election, county-level proportion of residents of age 65 and above, and an indicator variable of whether a county was under a ban on indoor dining, bars, gyms, hair salons, and related services before the start of the vaccine roll out.²⁵ Data on county-level economic outcomes were obtained from the Bureau of Economic Analysis, and the demographic variables are from the most recent US Census American Community Survey (2015–19 averages). Unemployment data are 2019 numbers obtained from the most recently available decennial census. To address temporal autocorrelation in the data, I include a lagged value of new COVID-19 cases per 1,000 residents, logged.²⁶

To tighten the causal claims, I also implement a matched-sample design.²⁷ I matched California counties (treatment group) with those in Oregon (control group) on all pretreatment variables (other than California and its interaction with Day of Treatment) using coarsened exact matching (CEM; Iacus, King, and Porro 2012).²⁸ CEM consists of identifying exact matches (observations with the same values on all covariates) after coarsening (dichotomizing or multichotomizing variables into discrete categories) any continuous and ordinal variables. Per Iacus et al. (2012), I selected the number of cutpoints to multichotomize the variables using the empirical knowledge of the data, whenever possible. To maximize sample size, I allowed for multiple matches for each observation.

Table 3 displays the list of matched counties: 12 counties from California and 11 from Oregon. The matched sample

25. Before the vaccination campaign, the two states did not differ on other major COVID-19 prevention policy measures in a meaningful way (Adeel et al. 2020). Both states were under mask mandates (Hubbard 2022). Both states' primary, secondary, and tertiary educational establishments shifted to virtual instruction at the start of the pandemic (Brown 2021a; Mays 2021) and did not resume hybrid or in-person instruction until March 2021 (Brown 2021a; Mays 2021).

26. Diagnostics support using a three-period lag of the logged seven-day average in cases.

27. A key benefit of matching is that, when applied to pretreatment outcome levels, it ensures both balance between the treatment and control groups as well as parallel trends in the pretreatment period (Lindner and McConnell 2019, 129–30). When used in combination with a difference-in-difference design, however, a poorly specified matching model may introduce bias (Lindner and McConnell 2019). If the observable variables used in the matching model do not account for a significant portion of the variation in the outcome, then what one may interpret as the treatment effect in fact may be due to the unobserved and unmodeled confounders. Model fit statistics, such as the large adjusted coefficient of determination, indicate that this is not a concern in this application. Despite this, I follow best practices by reporting both adjusted (matched sample) and unadjusted (full sample) results (Lindner and McConnell 2019).

28. The matching strategy accounts for vaccine availability by matching on county population—a key criterion used for vaccine allocation.

24. The Johns Hopkins data contain cumulative cases by day and county. I calculated daily cases by first-differencing cumulative cases. Because of data corrections, for a small number of observations, this resulted in negative values, which were recoded to 0.

Table 3. Matched Counties

Oregon	California
Wheeler	Sierra
Wallowa	Mariposa
Union	Modoc
Linn, Klamath	Calaveras
Deschutes, Jackson, Polk, Yamhill	El Dorado, Lake, Nevada, Siskiyou
Marion	Placer, San Luis Obispo
Multnomah	Santa Barbara, Sonoma

includes a balance of coastal and inland counties. While it excludes several outlier counties on the basis of population size (e.g., Los Angeles County in California), it includes a mix of counties of various levels of population size and urbanization.

To assess the balance in the matched sample, table 4 displays the standardized mean differences for each covariate, along with the corresponding *p*-values. California and Oregon counties in the matched sample are not statistically significantly different on any of the matching variables; the overall χ^2 is also not statistically significant.

I estimate the model, on both full and matched samples, using ordinary least squares regression. The full sample consists of daily observations for all counties (58 in California, 36 in Oregon) between December 17, 2020 (the first day of vaccine administration in both states), and April 11, 2021 (two weeks after Oregon authorized giving vaccines to food

Table 4. Balance between Oregon and California in the Matched Sample

Variable	Standard Difference	<i>p</i>
County GDP, logged	-.13	.74
County population, logged	-.03	.94
Unemployment rate	-.48	.24
Percentage BA degree	-.02	.97
Urbanization	.15	.71
Percentage Black	-.24	.55
Percentage Latino	-.67	.10
Percentage other race	-.51	.21
Percentage foreign	-.53	.20
Biden’s margin	-.34	.41
Proportion age 65+	-.49	.23
Indoor dining ban	.04	.93
Cumulative COVID, logged	-.11	.78
χ^2		15.7
df		13
<i>p</i>		.26

processing employees, including grocery employees), for a total of 10,904 nonmissing observations, while the matched sample consists of 2,668 observations.

RESULTS

Table 5 shows the results of the statistical analysis for the full sample (model 1) and the matched sample (model 2). Since the dependent variable, New COVID-19 Cases, is measured on a logged scale, the model coefficients are interpretable as percentage changes (Wooldridge 2015). Thus, the coefficient of 0.23 on California is statistically significant at $\alpha = .05$ (two-tailed test) and indicates that, before opening up COVID-19 vaccination to grocery store employees, the state of California has had, on average, about 26% ($(\exp(0.23) - 1) \times 100\%$) more new daily COVID-19 cases than the state of Oregon in the full sample. In the matched sample, however, this effect is no longer statistically significant at $\alpha = .05$ (two-tailed test).

The coefficient of -0.009 (-0.011 in the matched sample) on the Day of Treatment gives the average difference in new COVID-19 cases in Oregon counties before and after California opened up vaccination to grocery employees. This coefficient is statistically significant, indicating that Oregon experienced about a 1% decrease per day, on average, in new daily COVID-19 cases in the period since March 14, 2021, compared to the period between December 17, 2020, and March 14, 2021.

The coefficient of -0.04 (-0.03 in the matched sample) on the interaction term is the difference-in-difference coefficient. It indicates a 4% decrease (3% in the matched sample) in new COVID-19 cases in California (compared to Oregon), since California opened vaccine eligibility to grocery employees. This coefficient, of course, gives us only the average daily effect, as Day of Treatment is measured on an integer scale.²⁹ The marginal effect of California by Day of Treatment is shown in figure 4. As one can see, the effect of vaccinating grocery employees results in about a 43% decrease in cases ($(\exp(-.56) - 1) \times 100\%$; about 46% in the matched sample) 20 days after the start of the treatment period. These results provide strong support for the research hypothesis. Similar results also hold for the outcomes of COVID-19-related hospitalizations and deaths.³⁰

VALIDATING MODEL ASSUMPTIONS

Interpreting the model’s estimates in causal terms rests on several assumptions, such as the stable unit treatment value

29. Results are robust to coding Day of Treatment on a nominal scale or including polynomials. Diagnostics favor the model presented in table 5.

30. These additional results are presented in app. F.

Table 5. Effect of Vaccine Eligibility to Grocery Employees on New Daily COVID-19 Cases (Logged)

	Full Sample	Matched Sample
Day of Treatment	-.009*** (.001)	-.011*** (.002)
California	.229*** (.021)	-.036 (.043)
California × day of treatment	-.039*** (.002)	-.030*** (.003)
Cumulative cases, logged	-.117*** (.022)	-.295** (.098)
County GDP, logged	.350*** (.023)	-.571*** (.095)
County population, logged	.376*** (.022)	1.150*** (.120)
Unemployment rate	-.018*** (.002)	-.045*** (.013)
Percentage BA Degree	.021*** (.003)	.052*** (.006)
Urbanization	.008*** (.001)	.011*** (.001)
Percentage Black	.044*** (.004)	.213*** (.043)
Percentage Latino	.030*** (.001)	.066*** (.007)
Percentage other race	.026*** (.001)	.024*** (.006)
Percentage foreign	-.045*** (.002)	-.059*** (.014)
Biden's margin	-.010*** (.001)	-.010*** (.002)
Proportion age 65+	3.716*** (.266)	12.078*** (1.874)
Indoor dining ban	.084** (.026)	-.814*** (.142)
New cases/1,000 residents, three-day lag	.213*** (.004)	.132*** (.007)
Constant	-7.185*** (.177)	-3.403*** (.585)
<i>N</i>	10,904	2,668
Adjusted <i>R</i> ²	.89	.87

Note. Standard errors in parentheses.

** $p < .05$ (two-tailed).

*** $p < .01$ (two-tailed).




Figure 4. Marginal effect of vaccine eligibility to grocery workers: A, full sample (model 1); B, matched sample (model 2). Error bars represent 95% confidence intervals.

assumption (SUTVA), cross-sectional and temporal treatment exogeneity, and parallel trends. In this section, I discuss the validity of each assumption in the context of the current application.

SUTVA

The first assumption, SUTVA, requires that treatment assignments for other units do not affect the outcome for unit i and that each treatment defines a unique outcome for each county (Imbens and Rubin 2015, 33). The first part of SUTVA—that participants cannot interfere with the assigned treatments—is met as long as Oregon grocery workers cannot travel to California to receive the vaccine ahead of their turn. Although technically possible (California did not require proof of residency for vaccine appointments), any such instances would be costly: an employee would have to learn how to register for an appointment in a different state and take time off work and travel, in many cases, a significant distance. Hence, any such violations would be rare. Were this type of interference pervasive, this would induce downward bias in the estimate of the average treatment effect, that is, make it more difficult to detect a difference between the treatment and the control group.

A secondary part of SUTVA is that each unit in the treatment group receives the same “dose” of the treatment. This assumption would be violated, for example, if counties varied in terms of vaccine take-up among grocery employees. I account for this and other possible unobserved differences by including control variables and, in the matched sample, by matching on the observable sources of variation. A weaker variant of the “equal dose” assumption is that the unobservable differences that may induce the variation in the treatment dose are correlated with the observable control variables included in the model (Imbens and Rubin 2015, 9–13; Stuart 2010, 3).

Treatment exogeneity

The second assumption is unconfoundedness or treatment exogeneity. Exogeneity implies that the treatment and control groups are equal, on average, on all observed and unobserved variables that may affect the outcome variable, with the exception of the treatment and confounders for which the researcher controls (Imbens and Rubin 2015; Sekhon and Titiunik 2012, 36). While in randomized controlled experiments this assumption is a function of random assignment of research participants to the treatment and control group by the researcher, natural experiments lack a comparable iron-clad validity guarantee. In quasi-experimental design, the validity of the exogeneity assumption hinges on whether the researcher can provide a compelling justification.

I justify this assumption using three alternative and self-reinforcing strategies: with statistical evidence; via the research design, which includes matching on observables (Imbens and Rubin 2015); and analytically, by using evidence from case studies of California and Oregon’s decision-making (Dunning 2008). The statistical evidence consists of the analysis of the national determinants of vaccine prioritization, discussed earlier, and placebo tests, presented below.

In addition to implementing a matched design, the research design helps alleviate unit self-selection concerns by decoupling the treatment assignment (at the state level) from the unit of observation (counties). No general one-size-fits-all prioritization plan, even the one tailored to the demographics of the state, would apply equally well at the county level, especially in states as diverse as Oregon and California. Oregon’s Multnomah County, which includes Portland, has demographics very different from any of the counties east of the Cascades. Self-selection at the state level, in other words, would not translate into self-selection at the county level, as long as the demographics of individual counties do not precisely match those of the entire state.³¹

For space consideration, the full analytical justification is presented in appendix D. In brief, the novelty of the virus, and the lack of information regarding its spread, effectively set up a natural experiment of policy making under conditions of incomplete information. Earlier in the pandemic, these conditions resulted, for example, in many states implementing controversial and, in retrospect, unnecessary measures of closing down beaches and national parks.³² Likewise, most state-level decisions related to vaccine prioritization were also made under conditions of very limited information, as the first vaccines were given emergency approval before various aspects of their effects were fully evaluated. In particular, early distribution plans were made in the absence of reliable information on whether vaccines reduced transmission of the virus, prevented disease, or merely ameliorated the symptoms; a sufficient body of evidence showing that vaccines were effective at reducing transmission did not become available until mid-March 2021, long after the prioritization plans were finalized (CDC 2021b; Christie, Mbaeyi, and Walensky 2021).³³

31. Individual counties had very limited discretion over vaccine prioritization, mainly with the aim of avoiding wasting doses (CVAC 2021c, 12).

32. These policies were shown to be unnecessary, once the evidence emerged that the virus transmission is significantly reduced in unconfined spaces.

33. Although there was a working hypothesis among medical researchers that vaccinations may reduce transmission—owing to fewer of the more contagious symptomatic cases and reduced viral loads among vaccinated individuals (Levine-Tiefenbrun et al. 2021; Mallapaty 2021)—there was little direct evidence in late 2020, when prioritization distribution plans were being

A systematic correlation between the treatment assignment and the outcome requires expertise—a reasonable expectation of which treatment would work best for each unit. In the absence of such expertise, however, the assignment mechanism is not that different from tossing a coin or rolling a die. In other words, an uninformed decision maker approximates a randomization mechanism.

Further evidence that the governor offices did not act on a set agenda comes from both states’ strict adherence to the national guidelines and willingness to quickly modify previously formulated plans in response to sudden changes in national guidelines, such as the unexpected January 12 revision to the previous recommendation of prioritizing on the basis of both age and exposure. In response to that change, both states shifted to the now-recommended age-based prioritization plan, with the exception of groups whose eligibility had already been announced, such as Oregon’s K–12 employees and California’s grocery employees.

Temporal treatment exogeneity

Related to unconfoundedness is the assumption that the treatment does not influence the pretreatment population, that is, that opening vaccine eligibility to grocery workers does not affect COVID-19 cases in the preceding time period (Lechner 2011). An example of this would be California grocery workers taking extra precautions to avoid contracting COVID-19, for example, by wearing masks, in anticipation of getting the vaccine. They might have reasoned, for example, that it is worthwhile to incur some additional inconveniences for several weeks, so as to avoid contracting the virus right before getting the vaccine. Had this been systematically the case, we would observe evidence of the treatment effect before the start of the treatment. The placebo tests, discussed below, show no evidence of such an effect.

Parallel trends assumption

Finally, the model relies on the parallel trends assumption. This assumption requires that, after controlling for observable confounders, changes in expected nontreatment outcomes between the time periods before and after the treatment are mean independent of the treatment assignment (Lechner 2011). If this assumption holds, then difference-in-difference estimation recovers unbiased estimates of the average treatment effect for the treated.

drafted. Moreover, there was no consensus opinion within the medical community: some researchers offered a countering view, arguing that vaccinated individuals could still carry and spread the virus, even if they themselves were largely protected (e.g., Bleier, Ramanathan, and Lane 2021). Even as evidence changed, most states stuck by their predetermined prioritization plans.

Figure 5 shows the temporal trends in the raw data (fig. 5A) and in the matched sample (fig. 5B). Both subfigures show that the daily number of new COVID-19 cases followed a similar trend in the two states between the start of vaccinations in December 2020 and the middle of March 2021: both experience a brief period of increasing cases between early to mid-January 2021 followed by a consistent decline up until early March. In the matched sample, the aggregate difference in new COVID-19 cases between the two states was about 460 daily new cases on December 17, 2020. This difference narrows to about 222 new cases by February 15, 2021, and further to about 86 by March 1 and to 12 cases by the start of the treatment on March 15.

In the treatment period, as California expanded vaccine eligibility to grocery store employees on March 1, 2021, whereas Oregon did not, the two trends diverge. The point of divergence falls somewhere between March 1 and March 15—the latter date is when the grocery employees who had been vaccinated on March 1 would have achieved between 50% and 80% immunity. After this time, the trend of new COVID-19 cases in California continues to decline, whereas Oregon starts observing an increase. At this point Oregon surpasses California in the aggregate cases in the matched sample, its lead reaching about 69 daily new cases by April 11, 2021 (two weeks after Oregon also opened vaccine eligibility to grocery workers).

While the pretreatment trends look similar, especially between mid-January and mid-March, whether they are indeed parallel is ultimately a subjective judgment. Although one cannot directly test the null hypothesis that the slopes of the two trends are the same, it is possible to test whether the difference between the slopes is statistically different from zero. To do this, I estimated the following model:

$$\begin{aligned} \log(\text{New Cases}) = & \beta_0 + \sum_{t=1}^{T-1} \delta_{0t} \text{Day}_t + \beta_1 \text{California} \\ & + \sum_{t=1}^{T-1} \delta_{1t} \text{Day}_t \times \text{California} \\ & + \text{other factors,} \end{aligned}$$

where Day_t is a set of $T - 1$ indicator variables for each day $t \in \{1, 2, \dots, T - 1\}$ included in the sample, δ_{0t} is a set of daily intercepts, and the coefficients δ_{1t} on the interactions between California and each of the daily dummies are estimates of the daily changes in the difference between California and Oregon in the sample.

Figure 6 displays the coefficients, δ_{1t} , on the interaction terms for each day. The daily changes are not statistically different from zero for most of the pretreatment period. The pretreatment mean is 0.09; it drops to -0.13 shortly after California

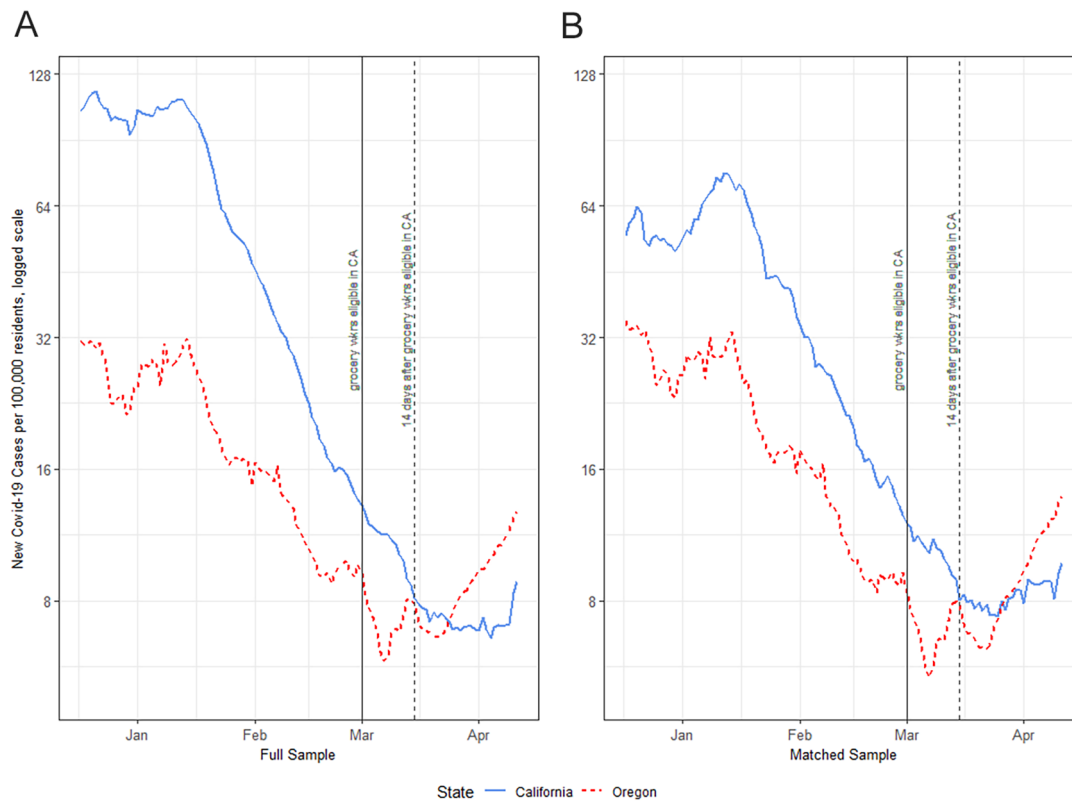


Figure 5. Temporal trends (state averages) in COVID-19 Cases, December 2020–April 2021

expands vaccine eligibility to grocery workers and further to -0.45 in the treatment period. I interpret this as (indirect) support for the parallel trends assumption. As additional support for the research hypothesis, the coefficients become statistically significant for the majority of the posttreatment period, and the trend line is increasing in magnitude. In summary, the above placebo tests show no discernible pretreatment trends in the matched sample.

I also perform placebo analysis by treating the dates that correspond to other differences between the two states' vaccine priority lists as pseudotreatments. Other than food and agriculture workers (which included grocery employees), the two states differed on the sequence and timing of several other groups, most prominently teachers and the incarcerated (see app. E).³⁴ California also had a less expansive definition for phase 1a, which put it 12 days ahead of Oregon in opening phase 1b.³⁵ Taking into the account the earlier prioritization of

K–12 employees, this delayed Oregon's opening date for vaccinating the elderly until February 7, 2021 (whereas California, which prioritized the elderly ahead of K–12 employees, was able to start vaccinating the elderly on January 13, 2021).

None of these differences matter from the theoretical point of view—since, out of all prioritized groups, only grocery workers act as central nodes that connect large segments of society.³⁶ None of the pseudotreatment test dates are associated with prominent shifts in the slope of the trend.³⁷

CONCLUSION

Unlike epidemiological studies that rely on aggregated population-level measures (e.g., Bubar et al. 2021), network

care facilities (CVAC 2020e, 6). Oregon included several additional groups, such as individuals with developmental disabilities, employees of early learning centers, and individuals working in death care services (OVAC 2021).

36. In network analysis terms, K–12 employees and the incarcerated are cliques—network clusters with large numbers of repeat interactions but few new ones.

37. One caveat is that statistical analysis does not allow for isolating the effect of grocery employees from that of other food and agriculture workers, as all three of these groups were included in the same “agriculture and food” category in both states. In technical terms, this means that the estimated effect is a “bundled” effect. Theoretically, agricultural workers or restaurant employees (e.g., short-order cooks) have network positions

34. Oregon opened vaccine eligibility for teachers on January 25, 2021, and to the incarcerated on February 3, 2021. In California, in contrast, teachers were not eligible until March 1 and the incarcerated until March 15.

35. While California's definition of health workers was limited to any personnel with direct interaction with patients, Oregon had no such qualifications and also included any nonmedical personnel of health care facilities, such as cleaners and food servers (OVAC 2021). In addition to health care workers, both states' phase 1a included residents of long-term

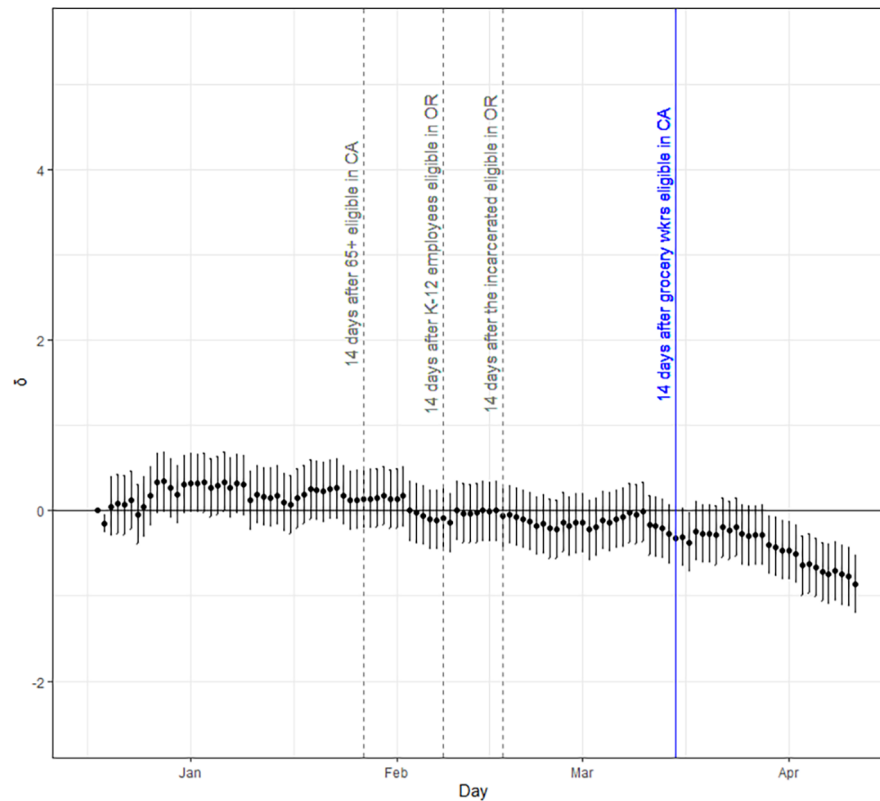


Figure 6. Daily changes in the difference between California and Oregon, δ_{t_i} , matched sample. Error bars represent 90% confidence intervals

science provides a more nuanced localized understanding of contagion in various types of network structures. Network tools allow for developing a more targeted and efficient approach to reducing contagion by focusing on the nodes with various transmission properties. This study demonstrates that using insights from network analysis to inform vaccine prioritization plans may have substantial effects on reducing contagion and, by implication, hospitalizations and mortality rates.

While highly effective as a means to reduce transmission, onset, symptoms, and mortality, COVID-19 vaccines have proven to be heterogeneous in their effectiveness. In particular, they are the least effective for individuals who face the highest risk of adverse effects from the virus, such as the immunocompromised and the elderly. Somewhat counterintuitively, this study demonstrates that the most effective strategy to protect these groups may be indirect. Rather than directly allocating vaccines to the individuals for whom these vaccines are the least effective, it may be more efficient, under conditions of vaccine scar-

similar to those of teachers and the incarcerated—they interact primarily with other members of their own professional setting but are relatively isolated from the general public, especially as a result of pandemic-related indoor dining bans. The placebo tests show no discernible effects for the incarcerated or K-12 educators, which implies that the treatment effect is driven by the grocery employees rather than other groups that happened to become eligible at the same time.

city, to minimize the probability of these individuals' exposure to the virus by inoculating their most frequent contacts. Allocating scarce vaccine resources to the individuals who act as central nodes in the network of human interactions—such as grocery workers—may be a more effective means to reducing community spread and a more effective means at protecting the most vulnerable.

Outside of the specific application to COVID-19, this article has implications for studies using insights from network science. Now that network analysis has become commonplace for studying political processes, the next step is to move from correlational to causal network analysis, and this article takes a first step. Theoretically, it offers a number of testable implications for the study of diffusion more broadly, such as information cascades, diffusion of policy innovations, political mobilization, and contagion of political violence. For instance, targeting the central nodes in a network of information transmission is key to controlling the spread of information, such as misinformation, radical speech, or antiregime rhetoric. The reverse strategy—structuring the network so as to ensure multiple transmission pathways, rather than relying on a small number of central nodes—is key to maintaining communication despite attempts at disruption, such as governments' attempts at censorship. Likewise, if the goal is to encourage the spread of policy innovation, from treaty ratification to regime change,

then a focus on a small number of central nodes—key innovators—may be more cost efficient than individual targeting of the entire pool of potential adopters. Similarly, advocacy groups may get a better return on investment by focusing on a handful of influential actors, rather than allocating resources across a wide set of government actors. This may hold for specific agencies within governments, as well.

Beyond the focus on network centrality, political science research can explore other types of network dependencies, such as indirect ties or the average length of the shortest paths (e.g., Chyzh 2016). For example, and related to the previous point, advocacy groups may not want to focus solely on actors who are central in the entire network but also on actors who connect otherwise disconnected network clusters. Explicitly modeling known dependencies among political actors and testing their implications within a causal network framework adds nuance to established political science findings.

REFERENCES

- Adeel, Abdul Basit, Michael Catalano, Olivia Catalano, Grant Gibson, Ezgi Muftuoglu, Tara Riggs, Mehmet Halit Sezgin, Olga Shvetsova, Naveed Tahir, Julie VanDusky-Allen, Tianyi Zhao, and Andrei Zhirnov. 2020. “COVID-19 Policy Response and the Rise of the Sub-national Governments.” *Canadian Public Policy* 46 (4): 565–84.
- American Journal of Managed Care*. 2021. “A Timeline of COVID-19 Vaccine Developments in 2021.” <https://www.ajmc.com/view/a-timeline-of-covid-19-vaccine-developments-in-2021>.
- Andrews, Nick, Elise Tessier, Julia Stowe, Charlotte Gower, Freja Kirsebom, Ruth Simmons, Eileen Gallagher, Meera Chand, Kevin Brown, Shamez N. Ladhani, Mary Ramsay, and Jamie Lopez Bernal. 2021. “Vaccine Effectiveness and Duration of Protection of Comirnaty, Vaxzevria and Spikevax Against Mild and Severe COVID-19 in the UK.” medRxiv. <https://doi.org/10.1101/2021.09.15.21263583>.
- Ashworth, Scott, Ethan Bueno de Mesquita, and Amanda Friedenberg. 2018. “Learning about Voter Rationality.” *American Journal of Political Science* 62 (1): 37–54.
- Ayoub, Phillip M., Douglas Page, and Sam Whitt. 2021. “Pride Amid Prejudice: The Influence of LGBT Rights Activism in a Socially Conservative Society.” *American Political Science Review* 115 (2): 467–85.
- Bernal, Jamie Lopez, Nick Andrews, Charlotte Gower, Julia Stowe, Chris Robertson, Elise Tessier, Ruth Simmons, Simon Cottrell, Richard Roberts, Mark O’Doherty, Kevin Brown, Claire Cameron, Diane Stockton, Jim McMenamin, and Mary Ramsay. 2021. “Early Effectiveness of COVID-19 Vaccination with BNT162b2 mRNA Vaccine and ChAdOx1 Adenovirus Vector Vaccine on Symptomatic Disease, Hospitalisations and Mortality in Older Adults in England.” medRxiv. <https://www.medrxiv.org/content/10.1101/2021.03.01.21252652v1.full.pdf>.
- Bleier, Benjamin S., Murugappan Ramanathan, and Andrew P. Lane. 2021. “COVID-19 Vaccines May Not Prevent Nasal SARS-CoV-2 Infection and Asymptomatic Transmission.” *Otolaryngology—Head and Neck Surgery* 164 (2): 305–7.
- Bonacich, Phillip. 1972. “Factoring and Weighting Approaches to Status Scores and Clique Identification.” *Journal of Mathematical Sociology* 2 (1): 113–20.
- Botkin, Ben. 2021. “Governor Kate Brown Mounts Defense over Decision to Vaccinate Teachers before Seniors.” *Lund Report*, January 22. www.thelundreport.org/content/gov-kate-brown-mounts-defense-over-decision-vaccinate-teachers-seniors%C2%A0.
- Box-Steffensmeier, Janet M., Benjamin W. Campbell, Dino P. Christenson, and Zachary Navabi. 2018. “Role Analysis Using the Ego-ERGM: A Look at Environmental Interest Group Coalitions.” *Social Networks* 52:213–27.
- Brankston, Gabrielle, Eric Merkley, David N. Fisman, Ashleigh R. Tuite, Zvonimir Poljak, Peter J. Loewen, and Amy L. Greer. 2021. “Quantifying Contact Patterns in Response to COVID-19 Public Health Measures in Canada.” *BMC Public Health* 21 (1): 1–10.
- Brown, Kate. 2021a. “Governor Kate Brown Announces Statewide School Closure for Students in Oregon from Monday, March 16 through Tuesday, March 31.” Press release, March 12. <https://www.myoregon.gov/2020/03/12/governor-kate-brown-announces-statewide-school-closure-for-students-in-oregon-from-monday-march-16-through-tuesday-march-31/>.
- Brown, Kate. 2021b. “Governor Kate Brown Expands COVID-19 Vaccination to All Oregonians 65 and Older.” Press release, January 12. <https://www.oregon.gov/newsroom/pages/NewsDetail.aspx?newsid=63423>.
- Bubar, Kate M., Kyle Reinholt, Stephen M. Kissler, Marc Lipsitch, Sarah Cobey, Yonatan H. Grad, and Daniel B. Larremore. 2021. “Model-Informed COVID-19 Vaccine Prioritization Strategies by Age and Serostatus.” *Science* 371 (6532): 916–21.
- CDC (Centers for Disease Control). 2020. “Covid-19 Data Tracker.” <https://covid.cdc.gov/covid-data-tracker/#datatracker-home>.
- CDC (Centers for Disease Control). 2021a. “How CDC Is Making COVID-19 Vaccine Recommendations.” <https://www.cdc.gov/coronavirus/2019-ncov/vaccines/recommendations-process.html>.
- CDC (Centers for Disease Control). 2021b. “Science Brief: Background Rationale and Evidence for Public Health Recommendations for Fully Vaccinated People.” https://www.cdc.gov/coronavirus/2019-ncov/science/science-briefs/fully-vaccinated-people.html?CDC_AA_refVal=https%3A%2F%2Fwww.cdc.gov%2Fcoronavirus%2F2019-ncov%2Fmore%2Ffully-vaccinated-people.html.
- CDGW (California Drafting Guidelines Workgroup). 2021. “Guidelines to California’s Health Departments: Allocation of COVID-19 Vaccine during Phase 1A.” January 12. <https://www.cdph.ca.gov/programs/cid/dcdc/cdph%20document%20library/immunization/draftingguidelines1asummary.pdf>.
- Christie, Athalia, Sarah A. Mbaeyi, and Rochelle P. Walensky. 2021. “CDC Interim Recommendations for Fully Vaccinated People.” *Journal of the American Medical Association* 325 (15): 1501–2.
- Chyzh, Olga V. 2016. “Dangerous Liaisons: An Endogenous Model of International Trade and Human Rights.” *Journal of Peace Research* 53 (3): 409–23.
- Chyzh, Olga V., and Mark S. Kaiser. 2019. “Network Analysis Using a Local Structure Graph Model.” *Political Analysis* 27 (4): 397–414.
- Coleman, James Samuel. 1964. *Introduction to Mathematical Sociology*. New York: Free Press.
- CVAC (Community Vaccine Advisory Committee). 2020a. Meeting summary, December 9. <https://www.cdph.ca.gov%2Fprograms%2Fcid%2Fdc%2Fcdph%2520document%2520library%2F%2Fmeeting%2Fsummary120920.pdf>.
- CVAC (Community Vaccine Advisory Committee). 2020b. Meeting summary, December 16. <https://www.cdph.ca.gov%2Fprograms%2Fcid%2Fdc%2Fcdph%2520document%2520library%2F%2Fmeeting%2Fsummary121620.pdf>.
- CVAC (Community Vaccine Advisory Committee). 2020c. Meeting summary, December 23. <https://www.cdph.ca.gov/programs/cid/dcdc/cdph%20document%20library/covid-19/meetingsummarycvac122320.pdf>.
- CVAC (Community Vaccine Advisory Committee). 2020d. Meeting summary, November 25. www.cdph.ca.gov%2Fprograms%2Fcid%2Fdc%2Fcdph%2520document%2520library%2F%2Fmeeting%2Fsummary112520.pdf.

- CVAC (Community Vaccine Advisory Committee). 2020e. Meeting summary, November 30. <https://www.cdph.ca.gov/programs/cid/dcdc/cdph%20document%20library/covid-19/meetingsummary113020.pdf>.
- CVAC (Community Vaccine Advisory Committee). 2021a. Meeting summary, February 3. www.cdph.ca.gov%2fprograms%2fcid%2fdcdc%2fcdph%2520document%2520library%2fcovid-19%2fmeetingsummary112520.pdf.
- CVAC (Community Vaccine Advisory Committee). 2021b. Meeting summary, January 6. <https://www.cdph.ca.gov/Programs/cid/dcdc/cdph%20Document%20Library/COVID-19/MeetingSummaryCVAC010621.pdf>.
- CVAC (Community Vaccine Advisory Committee). 2021c. Meeting summary, January 12. <https://www.cdph.ca.gov/Programs/cid/dcdc/cdph%20Document%20Library/COVID-19/MeetingSummaryCVAC01122.pdf>.
- CVAC (Community Vaccine Advisory Committee). 2021d. Meeting summary, March 17. https://www.cdph.ca.gov/programs/cid/dcdc/cdph%20document%20library/covid-19/meetingsummary_cvac_031721.pdf.
- Darling, Anna. 2021. "Essential Workers Still Not Included in Indiana's Vaccine Distribution." WFLI. <https://www.wfli.com/content/news/Essential-workers-still-not-included-in-Indianas-vaccine-distribution-573914071.html>.
- Desmarais, Bruce A., Jeffrey J. Harden, and Frederick J. Boehmke. 2015. "Persistent Policy Pathways: Inferring Diffusion Networks in the American States." *American Political Science Review* 109 (2): 392–406.
- Dooling, Kathleen, Mona Marin, Megan Wallace, Nancy McClung, Mary Chamberland, Grace M. Lee, H. Keipp Talbot, José R. Romero, Beth P. Bell, and Sara E. Oliver. 2021. "The Advisory Committee on Immunization Practices' Updated Interim Recommendation for Allocation of COVID-19 Vaccine—United States, December 2020." *MMWR: Morbidity and Mortality Weekly Report* 69 (5152): 1657–60.
- Dunning, Thad. 2008. "Improving Causal Inference: Strengths and Limitations of Natural Experiments." *Political Research Quarterly* 61 (2): 282–93.
- Eyre, David W., Donald Taylor, Mark Purver, David Chapman, Tom Fowler, Koen B. Pouwels, A. Sarah Walker, and Tim E. A. Peto. 2022. "Effect of Covid-19 Vaccination on Transmission of Alpha and Delta Variants." *New England Journal of Medicine* 386:744–56.
- Feehan, Dennis M., and Ayesha S. Mahmud. 2021. "Quantifying Population Contact Patterns in the United States during the COVID-19 Pandemic." *Nature Communications* 12 (1): 1–9.
- Feld, Scott L. 1991. "Why Your Friends Have More Friends than You Do." *American Journal of Sociology* 96 (6): 1464–77.
- Gasper, John T., and Andrew Reeves. 2011. "Make It Rain? Retrospection and the Attentive Electorate in the Context of Natural Disasters." *American Journal of Political Science* 55 (2): 340–55.
- Gassen, Joachim. 2021. "Tidycovid19: Download, Tidy and Visualize Covid-19 Related Data." <https://joachim-gassen.github.io/tidycovid19/>.
- Gingerich, Daniel W., and Jan P. Vogler. 2021. "Pandemics and Political Development: The Electoral Legacy of the Black Death in Germany." *World Politics* 73 (3): 393–440.
- Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78 (6): 1360–80.
- Harris, Ross J., Jennifer A. Hall, Asad Zaidi, Nick J. Andrews, J. Kevin Dunbar, and Gavin Dabrera. 2021. "Effect of Vaccination on Household Transmission of SARS-CoV-2 in England." *New England Journal of Medicine* 385 (8): 759–60.
- Hubbard, Kaia. 2022. "These States Have COVID-19 Mask Mandates." *US News and World Report*, March 28. <https://www.usnews.com/news/best-states/articles/these-are-the-states-with-mask-mandates#cali>.
- Hunter, David R., Mark S. Handcock, Carter T. Butts, Steven M. Goodreau, and Martina Morris. 2008. "ERGM: A Package to Fit, Simulate and Diagnose Exponential-Family Models for Networks." *Journal of Statistical Software* 24 (3): 1–29.
- Iacus, Stefano M., Gary King, and Giuseppe Porro. 2012. "Causal Inference without Balance Checking: Coarsened Exact Matching." *Political Analysis* 20 (1): 1–24.
- Igelnik, Ruth, Scott Keeter, and Hannah Hartig. 2021. "Behind Biden's 2020 Victory." Pew Research Center, June 30.
- Imbens, Guido W., and Donald B. Rubin. 2015. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge: Cambridge University Press.
- King, Gary, Jennifer Pan, and Margaret E. Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107 (2): 326–43.
- Kirkland, Justin H. 2011. "The Relational Determinants of Legislative Outcomes: Strong and Weak Ties between Legislators." *Journal of Politics* 73 (3): 887–98.
- Lechner, Michael. 2011. "The Estimation of Causal Effects by Difference-in-Difference Methods." *Foundations and Trends in Econometrics* 4 (3): 165–224.
- Levine-Tiefenbrun, Matan, Idan Yelin, Rachel Katz, Esmá Herzel, Ziv Golan, Licitá Schreiber, Tamar Wolf, Varda Nadler, Amir Ben-Tov, Jacob Kuint, Sivan Gazit, Tal Patalon, Gabriel Chodick, and Roy Kishony. 2021. "Decreased SARS-CoV-2 Viral Load following Vaccination." *Nature Medicine* 27:790–92.
- Lindner, Stephan, and K. John McConnell. 2019. "Difference-in-Differences and Matching on Outcomes: A Tale of Two Unobservables." *Health Services and Outcomes Research Methodology* 19:127–44.
- Lohmann, Susanne. 1994. "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989–91." *World Politics* 47:42–101.
- Lynch, Julia. 2020. *Regimes of Inequality*. Cambridge: Cambridge University Press.
- Mallapaty, Smriti. 2021. "Can COVID Vaccines Stop Transmission? Scientists Race to Find Answers." *Nature* 10:1–6.
- Mattila, Mikko. 2020. "Does Poor Health Mobilize People into Action? Health, Political Trust, and Participation." *European Political Science Review* 12 (1): 49–65.
- Mays, Mackenzie. 2021. "Newsom Strikes School Reopening Deal with California Lawmakers." Politico, March 1. <https://www.politico.com/states/california/story/2021/03/01/newsom-strikes-school-reopening-deal-with-california-lawmakers-1366270>.
- Milligan, William R., Zachary L. Fuller, Ipsita Agarwal, Michael B. Eisen, Molly Przeworski, and Guy Sella. 2021. "Impact of Essential Workers in the Context of Social Distancing for Epidemic Control." *PLoS ONE* 16 (8): e0255680.
- Neelon, Brian, Fedelis Mutiso, Noel T. Mueller, John L. Pearce, and Sara E. Benjamin-Neelon. 2021. "Associations between Governor Political Affiliation and COVID-19 Cases, Deaths, and Testing in the US." *American Journal of Preventive Medicine* 61 (1): 115–19.
- Norrlof, Carla. 2020. "Is COVID-19 the End of US Hegemony? Public Bads, Leadership Failures and Monetary Hegemony." *International Affairs* 96 (5): 1281–303.
- Nunes, Baltazar, Ana Paula Rodrigues, Irina Kislaya, Camila Cruz, André Peralta-Santos, João Lima, Pedro Pinto Leite, Duarte Sequeira, Carlos Matias Dias, and Ausenda Machado. 2021. "mRNA Vaccine Effectiveness against COVID-19-Related Hospitalisations and Deaths in Older Adults: A Cohort Study Based on Data Linkage of National Health Registries in Portugal, February to August 2021." *Eurosurveillance* 26 (38).
- OHA (Oregon Health Authority). 2020. "Phase 1a of Vaccine Plan Targets Wide Range of Health Settings." December 18. <https://www.oregon.gov/oha/ERD/Pages/Phase1aOfVaccinePlanTargetsWideRangeOfHealthSettings.aspx>.
- OHA (Oregon Health Authority). 2021a. "COVID-19 Vaccine Advisory Committee Meeting." January 7. <https://www.oregon.gov/oha/ph/preventionwellness>

- /vaccinesimmunization/immunizationpartnerships/covid19vac/packet-covid-19-vac-2021-01-07.pdf.
- OHA (Oregon Health Authority). 2021b. "Vaccine Advisory Committee Minutes." January 21. <https://www.oregon.gov/oha/ph/preventionwellness/vaccinesimmunization/immunizationpartnerships/covid19vac/minutes-covid-19-vac-2021-01-21.pdf>.
- OHA (Oregon Health Authority). 2022. "Vital Role of Migrants in the State Economy." <https://www.oregon.gov/oha/HPA/HP-PCO/Pages/Migrant-Health.aspx#:~:text=The%20State%20of%20Oregon%20recognizes,multi%2Dbillion%20dollar%20agricultural%20industry>.
- OVAC (Oregon Vaccine Advisory Committee). 2021. "COVID-19 Vaccine Advisory Committee Meeting." Video, January 14. <https://www.youtube.com/watch?v=JFPsGuRgkK4&t=3612s>.
- Padgett, John F., and Christopher K. Ansell. 1993. "Robust Action and the Rise of the Medici, 1400–1434." *American Journal of Sociology* 98 (6): 1259–319.
- Patty, John W., and Elizabeth Maggie Penn. 2017. "Network Theory and Political Science." In *The Oxford Handbook of Political Networks*. Oxford: Oxford University Press.
- Polack, Fernando P., et al. 2020. "Safety and Efficacy of the BNT162b2 mRNA Covid-19 Vaccine." *New England Journal of Medicine* 383 (27): 2603–15.
- Quiroz Flores, Alejandro, and Alastair Smith. 2013. "Leader Survival and Natural Disasters." *British Journal of Political Science* 43 (4): 821–43.
- Richterman, Aaron, Eric A. Meyerowitz, and Muge Cevik. 2022. "Indirect Protection by Reducing Transmission: Ending the Pandemic with SARS-CoV-2 Vaccination." In *Open Forum Infectious Diseases*, vol. 9. Oxford: Oxford University Press.
- Robins, Garry, Philippa Pattison, Yuval Kalish, and Dean Lusher. 2007. "An Introduction to Exponential Random Graph (p) Models for Social Networks." *Social Networks* 29:173–91.
- Salameh, Tiffany. 2021. "Exclusive: Governor Holcomb Talks about the Expanded Vaccine Rollout and His Plans for 2021." ABC57, January 8.
- Salmeron Rios, Sergio, Elisa Belen Cortes Zamora, Almudena Avendano Cespedes, Luis Romero Rizos, Pedro Manuel Sánchez-Jurado, Ginés Sánchez-Nievas, Marta Mas Romero, María Teresa Tabernero Sahuquillo, José Joaquín Blas Señalada, Antonio Murillo Romero, Inmaculada García Nogueras, Juan de Dios Estrella Cazalla, Fernando Andrés-Pretel, Volker Martin Lauschke, Justin Stebbing, and Pedro Abizanda. 2022. "Immunogenicity After 6 Months of BNT162b2 Vaccination in Frail or Disabled Nursing Home Residents: The COVID-a Study." *Journal of the American Geriatrics Society* 70 (3): 650–58.
- Sekhon, Jasjeet S., and Rocio Tituniuk. 2012. "When Natural Experiments Are Neither Natural nor Experiments." *American Political Science Review* 106 (1): 35–57.
- Siegel, David A. 2011. "When Does Repression Work? Collective Action in Social Networks." *Journal of Politics* 73 (4): 993–1010.
- Stuart, Elizabeth A. 2010. "Matching Methods for Causal Inference: A Review and a Look Forward." *Statistical Science* 25 (1): 1–21.
- US Department of Labor. 2020. "The Economics Daily." Bureau of Labor Statistics, April 20.
- US Department of Labor. 2022. "Union Membership (Annual) 2021." Press release. Bureau of Labor Statistics.
- Van Belle, Douglas A. 1996. "Leadership and Collective Action: The Case of Revolution." *International Studies Quarterly* 40 (1): 107–32.
- VanDusky-Allen, Julie, and Olga Shvetsova. 2021. "How America's Partisan Divide over Pandemic Responses Played Out in the States." Conversation, May 12.
- Wasserman, Stanley, and Katherine Faust. 1994. *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.
- Wooldridge, Jeffrey M. 2015. *Introductory Econometrics: A Modern Approach*. 6th ed. Boston: Cengage Learning.