4-1-2023

A Diffusion Network Event History Estimator

Jeffrey J. Harden
University of Notre Dame

Bruce A. Desmarais
Pennsylvania State University

Mark Brockway
Syracuse University

Frederick J. Boehmke
University of Iowa

Scott J. Lacombe
Smith College, slacombe@smith.edu

See next page for additional authors

Recommended Citation

Authors

Jeffrey J. Harden, Bruce A. Desmarais, Mark Brockway, Frederick J. Boehmke, Scott J. Lacombe, Fridolin Linder, and Hanna Wallach

A Diffusion Network Event History Estimator

Jeffrey J. Harden, University of Notre Dame
Bruce A. Desmarais, Pennsylvania State University
Mark Brockway, Syracuse University
Frederick J. Boehmke, University of Iowa
Scott J. LaCombe, Smith College
Fridolin Linder, Pennsylvania State University
Hanna Wallach, Microsoft Research

Research on the diffusion of political decisions across jurisdictions typically accounts for units' influence over each other with (1) observable measures or (2) by inferring latent network ties from past decisions. The former approach assumes that interdependence is static and perfectly captured by the data. The latter mitigates these issues but requires analytical tools that are separate from the main empirical methods for studying diffusion. As a solution, we introduce network event history analysis (NEHA), which incorporates latent network inference into conventional discrete-time event history models. We demonstrate NEHA's unique methodological and substantive benefits in applications to policy adoption in the American states. Researchers can analyze the ties and structure of inferred networks to refine model specifications, evaluate diffusion mechanisms, or test new or existing hypotheses. By capturing targeted relationships unexplained by standard covariates, NEHA can improve models, facilitate richer theoretical development, and permit novel analyses of the diffusion process.

Understanding patterns in the diffusion of decisions across governmental jurisdictions has been a topic of interest in the social sciences for over half a century. One critical component in this research agenda involves accounting for interdependence between units in the adoption of political choices. Scholars have devoted considerable attention to the shape and scope of these underlying connections. They have operationalized diffusion networks with observable characteristics, such as similar geopolitics, economies, cultures, or shared geographic borders (e.g., Brooks 2005; Kreitzer 2015).

While these options are intuitive, they are also restrictive; they typically assume that the interdependence is static or require ex ante researcher decisions. Recent work improves on measuring unit-to-unit influence solely with observed data by dynamically inferring latent networks with an arbitrary structure (Boehmke et al. 2020; Desmarais, Harden, and Boehmke 2015). But even this approach is problematic for the study of diffusion.

Specifically, to account for latent diffusion networks researchers must fit two separate models to the data—one to infer the networks and another to model the effects of

covariates on the choice to adopt. This requirement represents a key inconsistency. If covariates and latent networks both explain a diffusion process, then each of these two models is itself misspecified. Furthermore, to incorporate diffusion networks into the final model of adoption, researchers must include them as covariates in an ad hoc process that ignores the functional form used to infer the latent ties. Ultimately, such issues hinder researchers' ability to understand diffusion processes and may bias their substantive conclusions.

In this article we address these problems by developing a unified estimator that simultaneously infers latent diffusion ties and estimates the effects of covariates on adoption—all within the familiar functional form of a discrete-time event history analysis (EHA) model. We label this methodology network event history analysis (NEHA). Our NEHA estimator offers researchers an adoption modeling framework that draws on both the strength of conventional EHA (a simple and intuitive means of incorporating covariate effects derived from theory) and the strength of network inference (the ability to infer, describe, and adjust for unmodeled patterns of interdependence between units). NEHA can yield novel insight regarding questions such as the following: Are there diffusion patterns in the data beyond ideological similarity or shared borders? Does the structure of the diffusion patterns suggest competition between units, a learning process, or emulation of a small number of leaders? Are units with greater legislative capacity more likely to become policy leaders? In short, NEHA better accounts for interdependence and offers new substantive information for theory development and testing that conventional estimators lack.

We illustrate this new modeling framework with applications to public policy diffusion in the American states. The typical empirical approach in these studies is an EHA model that quantifies the role of several factors in states' decisions to adopt new policies. Scholars have produced a large volume of these studies in American politics, comparative politics, and international relations (see Graham, Shipan, and Volden 2013). Desmarais et al. (2015) introduced latent network inference (NetInf) to this literature to relax assumptions about the structure of interdependence. However, the standard output from NetInf is a network. Inferences on its parameters cannot be drawn directly within an EHA model as with the effect of a covariate. In contrast, NEHA is an EHA model that seamlessly allows for theory testing and network inference.

NEHA is also broadly applicable. It can be used in any context that includes multiple political decisions (e.g., policies, treaties, actions) spreading across a set of units (e.g., states, countries, organizations). We focus on policy diffusion in the American states because of the high prevalence of multipolicy studies, or pooled event history analysis (PEHA), in that domain.[1] But NEHA could be fruitfully used for other lines of inquiry in political science, such as the diffusion of human rights treaties (Wotipka and Tsutsui 2008) or liberal norms (Tallberg et al. 2020).

MODELING DIFFUSION

Researchers in a variety of fields have long been interested in the diffusion of ideas, norms, products, or policies across actors or jurisdictions. Rogers's (1962) early work on the spread of hybrid seed corn pioneered inquiry into the diffusion of innovations, and the concept has been widely applied in the time since, especially in the social sciences (see also Gray 1973). In political science, notable examples include Berry and Berry's (1990) examination of the diffusion of lotteries across the American states, Brooks's (2005) examination of the spread of pension privatization across countries, and Shipan and Volden's (2008) work on antismoking policy in cities.

Walker (1969) laid the foundation for studying the diffusion of policy innovations in the American states. He collected data on the timing of adoption for 88 policies to study patterns in states' innovativeness. He postulated the existence of "more or less stable patterns of diffusion of innovations among the American states" (888) that resembled a network with "pioneer" states at the top and other states "sorted out along branches of the tree according to the pioneer, or set of pioneers, from which they take their principal cues" (893). This theorizing emphasized the relational nature of diffusion and motivated subsequent work to better understand the interdependence of the units under study. The literature has evolved over time, but in general scholars have converged on a distinct set of theoretical forces underlying diffusion: learning, emulation, competition, and coercion (Gilardi 2016; Shipan and Volden 2008).

Despite the theoretical richness of the various processes that drive diffusion, empirical operationalizations of their roles are limited by data availability on relevant connections (Maggetti and Gilardi 2014). Many quantitative diffusion studies rely on measures of lagged adoptions by jurisdictions with similar geographic, social, or political characteristics (e.g., Grossback,

---

1. Researchers in American state politics have increasingly relied on pooled analyses of multiple policies to better understand the systematic factors that explain adoption across many decisions. A Google Scholar search for relevant terms reveals that this approach appears in one article before 2009, six from 2009 to 2015, and 11 articles after 2016. With the introduction of resources such as Boehmke et al.'s (2020) State Policy Innovation and Diffusion (SPID) database, analyzing multiple policies at once is becoming more empirically feasible and theoretically important.

Nicholson-Crotty, and Peterson 2004; Mallinson 2021a). Dyadic EHA is particularly well suited to this strategy because it can easily accommodate multiple measures of similarity among pairwise combinations of units (Gilardi and Füglister 2008; Hinkle 2015; Volden 2006). But in spite of its precision, this approach implicitly requires the assumptions that the interdependence between units is static and measured perfectly by the observed variables.

Latent network inference recently arose in this literature to improve scholars' tools for accounting for the role of other jurisdictions' behavior in the adoption process. The methodology does expand scholars' abilities to measure persistent diffusion patterns between units, but not without drawbacks of its own. Desmarais et al. (2015) present the NetInf algorithm as well as dynamic networks of the diffusion process inferred from a large number of states' past policy adoption decisions. These networks describe ties between states, indicating which states a particular state tends to follow in making policy adoption choices. They can be aggregated to produce an alternative measure of other states' decision-making that may be an improvement over observed covariates alone (Desmarais et al. 2015).

However, network inference in its current form is not a complete solution. Specifically, there are at least three major shortcomings of the NetInf methodology with respect to studying adoption with EHA. First, researchers must estimate two separate structural models on the same, or highly related, data sets—one to infer ties and another to estimate covariate effects, controlling for the ties estimated in the first model. Second, NetInf does not permit researchers to use the conventional functional form assumed for discrete-time adoption choices. Finally, even if researchers are willing to estimate two separate models, it is not immediately apparent how the latent diffusion ties from NetInf should be included in the EHA specification. Desmarais et al.'s (2015) proposed options simply construct another covariate to add to the model, which dilutes the improvement of NetInf over previous approaches.

Consequently, while researchers may wish to incorporate latent diffusion network ties into their models, they lack the empirical guidance to do so effectively. In what follows we motivate, define, and apply a modeling framework that overcomes all of these shortcomings by (1) directly extending conventional PEHA models to incorporate latent diffusion ties and (2) providing an algorithm according to which latent diffusion parameters and covariate effects are simultaneously estimated. This approach bridges the conceptual advancements offered by NetInf and the empirical realities of applied research. The result is a modeling framework that best reflects the data-generating process of diffusion and is comparatively easier to use than the existing alternatives.

EVENT HISTORY MODELING WITH LATENT DIFFUSION TIES

The statistical methods that have dominated this area of research generally fall into three classes: (1) unit-level (e.g., state or country) EHA (Berry and Berry 1990), (2) dyadic EHA (Gilardi and Füglister 2008; Hinkle 2015; Volden 2006), and (3) latent variable methods (Desmarais et al. 2015; Garrett and Jansa 2015; Linder et al. 2018). The first class is designed primarily to test hypotheses regarding the effects of unit-level variables on the likelihood of adoption (e.g., the role of wealth or political ideology). The dyadic approach models and tests hypotheses about the patterns associated with interaction between units. Both of these approaches are straightforward to incorporate into commonly used EHA models (Grossback et al. 2004). In contrast, the existing latent variable measurement methods are not; they must be applied separately from the models used to estimate the effects of covariates on adoption.

Here we introduce NEHA as a discrete-time event history estimator that models (1) unit-level effects on adoption decisions, (2) dyadic effects on the tendency of units to emulate each other, and (3) residual dyadic ties between units (i.e., latent diffusion networks).[2] In other words, it is a method that unifies the three classes of analysis noted above. It permits the analyst to identify which units are the innovative leaders, and which are followers, but also model the associations between covariates and adoption decisions as in conventional models.

Network inference via latent edge selection

Consider an asymmetric matrix of diffusion parameters, denoted Γ, in which the $\gamma_{k,i}$ cell gives the effect of a previous adoption by unit k on the conditional log odds that unit i adopts a policy at time t, given that i has not adopted before time t. The conventional approach of incorporating a diffusion covariate into EHA models can be understood as a special case of Γ—one in which element $\gamma_{k,i}$ is 0 if i and k do not share the trait and some constant value if they do share it.

The underlying diffusion process represented by NetInf, although not exactly a special case of Γ, reflects a related set of assumptions in which element (k, i) is 0 if there is not a diffusion tie to from k to i and some constant value if there is a tie from k to i. Suppose we (1) relax the assumption of constant effects (e.g., that all units sharing a trait have the same diffusion effects) and (2) relax the assumption that there exists a dichotomy such that some diffusion effects are 0. These choices would result in a matrix of $n \times (n - 1)$ diffusion parameters (2,450 in the case

---

2. NEHA is flexible in that it can accommodate multiple diffusion processes. As long as all relevant units appear in the data, the estimator is agnostic about the hierarchy of the ties. It can infer horizontal, top-down, or bottom-up paths of influence.

of the 50 American states)—far too many to estimate in the context of most data sets used in diffusion research. Moving forward, we retain the assumption that some cells in Γ are 0 but relax the assumption of constant effects. Our approach to extending conventional EHA models to estimate covariate effects or latent network effects is to include functions of previous adoptions by other units as covariates in the linear predictors of the models. These functions could include indicators of previous adoption or functions of the time since the previous adoption occurred. In conducting inference with NEHA, we use regularization to impose sparsity on the latent network effects, assuring that some are estimated to be exactly zero.

Additionally, in our standard implementation of NEHA we include the constraint that network effects must increase the likelihood of adoption. This assumption simplifies model estimation and matches scholars' traditional conceptualization of the diffusion process. However, it is straightforward to relax and estimate antidiffusive effects (i.e., a decrease in the likelihood of adoption). In the appendix we demonstrate this feature by allowing NEHA to infer negative ties between units. Doing so produces a negative diffusion network, or the set of units that a unit seeks to avoid when making decisions. Such a network is only possible with NEHA, and thus this use of the estimator could potentially break new theoretical and empirical ground in diffusion research. However, it also demands more from the data compared to networks that are constrained to positive ties.

Defining the model

We begin with the conventional discrete-time event history model, as implemented using logistic regression (see Box-Steffensmeier and Jones 2004). We work completely within the framework of the standard risk set, in which each unit can adopt a policy at most once. However, this approach could be extended to repeated event models or to binary time series cross-sectional modeling more generally. Let $y_{i,j,t} = 1$ if unit i has adopted decision j by time t and 0 otherwise. Let $\Pr(y_{i,j,t} = 1 | y_{i,j,T} = 0, \forall \ T < t) = \pi_{i,j,t}$. A discrete event history logit is implemented by including observations in the data set for time periods t that $y_{i,j,t} = 0$, and the minimum t for which $y_{i,j,t} = 1$, then removing them from the set in the second period for which $y_{i,j,t} = 1$ and beyond. The functional form of the discrete-time, unit-level, event history model is

$$\ln\left(\frac{\pi_{i,j,t}}{1 - \pi_{i,j,t}}\right) = \beta' x_{i,j,t},$$

where β and $x_{i,j,t}$ are p-element vectors of regression coefficients and covariates, respectively.

We further add directed dyadic covariates to this model that condition on the emulated unit in the dyad having adopted the choice under study (e.g., policy) at a previous time point (see, e.g., Grossback et al. 2004). We call this model variant a unit-level discrete event history model with dyadic effects because it departs from the functional form of the dyadic event history model that has been used in recent policy diffusion studies (e.g., Hinkle 2015). The dyadic effects are given by η in the model form

$$\ln\left(\frac{\pi_{i,j,t}}{1 - \pi_{i,j,t}}\right) = \beta' x_{i,j,t} + \eta' z_{\mathcal{K}_t,i,j,t},$$

where η is a q-element vector of model coefficients, and $z_{\mathcal{K}_t,i,j,t}$ is a q-element vector of covariates. Each covariate is defined as a function of $\mathcal{K}_t$, the set of units that have previously adopted the policy and unit i (e.g., the number of geographic neighbors of i that have adopted previously).

We add one more set of parameters to the model to incorporate latent diffusion ties that are not sufficiently modeled with the covariates,

$$\ln\left(\frac{\pi_{i,j,t}}{1 - \pi_{i,j,t}}\right) = \beta' x_{i,j,t} + \eta' z_{\mathcal{K}_t,i,j,t} + \sum_{k \in \mathcal{K}_t} \gamma_{k,i}.$$

The term $\gamma_{k,i}$ is a nonnegative real-valued parameter that gives the effect of unit k having previously adopted a policy on the log odds of unit i adopting the policy. The parameters β and η can be estimated through standard logistic regression. However, the γ parameters raise two estimation challenges. First, we do not expect that every unit affects every other unit's adoption probability. We refer to this as a sparsity assumption. Second, depending on the number of decisions in which k is an early adopter, there may be very little data from which to estimate the value of $\gamma_{k,i}$. We use a variable selection approach to implement sparsity for the γ parameters. We also consider two versions of NEHA concerning whether or not the values of γ vary in the network. A common assumption in the policy diffusion literature is that diffusion patterns (e.g., the effects of geographic neighbors) follow a single parameter value. NEHA can be parameterized in this way, or it can be parameterized such that γ varies across edges. Analysts looking for guidance on this choice can compare the fit of the two versions (see below).

As it is currently defined, NEHA assumes diffusion effects are constant over time. That is, based on $\gamma_{k,i}$, adoption by k has a constant effect on the odds that i adopts in the future, regardless of how long in the past k adopted. Such an assumption may not be warranted. Accordingly, we introduce structure to NEHA to relax this assumption and model exponential decay in the diffusion effect. Specifically, we extend the diffusion component of NEHA to be

$$\sum_{k \in \mathcal{K}_t} \gamma_{k,i} \exp(-\exp[\alpha] \times [t - t_k]),$$

where t_k is the time that unit k adopted (see the appendix for a visualization of this function). We implement a grid search

to tune α. Finally, it is important to note that the latent edge is unidentified if either $\gamma_{k,i} = 0$ or $\exp[\alpha] = 0$.

NEHA estimation

The structure of NEHA represents a straightforward extension of the regression models commonly used to study diffusion. However, estimation for NEHA presents a challenge in that there are far more potential directed ties between units than we would like to include in the latent diffusion network. Pragmatically speaking, there may be too many potential ties for all of their estimates to be simultaneously identified. Even if we do have enough data to estimate all of the diffusion tie effects, if we assume that the true diffusion network is relatively sparse, we have strong reason to suspect that estimating a parameter value for each potential tie could lead to a substantial loss in efficiency for estimating the true diffusion tie parameters as well as the covariate effects.

To estimate the parameters in NEHA while simultaneously pushing most of the diffusion tie parameters to exactly zero, we use a combination of data subsetting and variable selection. There are many potential approaches to this problem. In the appendix we discuss two prominent alternatives and explain why we do not use them. Instead, we construct an edge-selection algorithm that is custom tailored to the structure of diffusion data and centers around the consistency of the Bayesian information criterion (BIC) in variable selection.[3] Within the maximum likelihood framework, model selection using the BIC is consistent (Cook and Forzani 2009), meaning that selecting the model with the lowest BIC will, in large data sets, result in selecting the true model in expectation. In our approach to variable selection with NEHA, we fix the observed covariates and apply variable selection to the edge parameters. It is not computationally feasible to estimate and compare all models with 0 to $n \times (n - 1)$ edges, so we propose a doubly iterative technique to find the model with the optimal BIC while keeping the computing demands manageable.

We begin by estimating the model with no edges. We then subset the data set into n separate data sets—one for each node in the data, where each data set represents only those adoption decisions by a single node. For each subset, we then calculate the BIC for every one-edge model, conditional on the covariate effects estimated in the full data set. If the BIC of the best one-edge model is lower than that of the zero-edge model, we consider all two-edge models and continue to consider higher numbers of edges until the BIC on the subset of data fails to improve.[4]

After we identify the best subset of edges for each node, we reestimate the covariate effects on the full data set, including the edges that have been identified. At this point we use a grid search, tuned by tenfold cross-validation, to update α. Given updated covariate effect estimates, and an updated value of α, we repeat the BIC-based selection of edges for each node. We repeat the covariate effect estimation, α updates, and BIC-based edge selection until the set of edges estimated in two consecutive iterations does not change. All estimation is done using standard logistic regression.

Our approach to edge selection with NEHA exploits the fact that, when represented as a covariate, an edge variable for source and target i and k can only be nonzero for observations in which k is the potential adopting node. Speed gains can be realized by parallelizing the BIC-based edge selection across separate node-specific data sets. This method remains quite fast, as long as the number of edges sent to any one node remains relatively low (e.g., fewer than 10). If the BIC-based edge selection moves on to large subsets of edges, the process of estimating and comparing all models of larger subset sizes would make the computing demands prohibitive. Fortunately, the results presented in our applications below suggest that the number of edges sent to any one node is indeed quite low.[5]

Because the motivation for our approach rests on the consistency of subset selection using BIC and we do not have analytical understanding of its finite-sample properties, we recommend the use of parametric bootstrap methods (Lewis and Poole 2004) to evaluate the performance of NEHA on a given data set. That is, we recommend that researchers simulate synthetic data from the model reflected by NEHA and reestimate it to assess (1) its effectiveness in identifying edges in the given data set and model and (2) the likelihood that the edges identified would have been falsely discovered. We illustrate this use of the parametric bootstrap method below.[6]

NEHA works best as the sample of policies increases in size. This condition would have been problematic several decades

---

ommend a bootstrap-based hypothesis test to assess NEHA model fit (see below).

3. We present this algorithm in detailed pseudocode in the appendix.

4. BIC must be decreased for edges to be added to the network. However, the magnitudes of improvements are not guaranteed. For instance, each edge could reduce BIC by 0.05, 0.5, or 5. Thus, we also rec-

5. Desmarais et al. (2015) use BIC to select the number of edges that optimizes an auxiliary event history model. However, they represent the latent edges as a single variable in the regression and do not penalize each edge separately. In our methodology, the BIC is penalized for each additional edge added. We see this choice—which results in many fewer edges being inferred—as a more accurate characterization of the degrees of freedom used in the network inference process.

6. We also recommend bootstrap methods to measure uncertainty of the inferred ties. The diffusion effects are subjected to both selection and a nonnegativity constraint, so we discourage researchers from interpreting the conventional standard errors for the diffusion effect parameters. The parametric bootstrap better accounts for both constraints. The covariate effects are never subject to regularization or variable selection. Thus, we expect their confidence intervals to exhibit proper coverage probabilities.

ago, but it is now quite common for researchers to study data sets of dozens or even hundreds of policies or more in a single model (e.g., Bricker and LaCombe 2021; Hinkle 2015; Mallinson 2021b; Parinandi 2020). That said, the method will not introduce problems if estimated on a smaller number of policies. Relevant edges could go unidentified in small data sets, but NEHA would not introduce new bias in such a case. Additionally, in single policy studies, in which it would not be feasible to infer diffusion ties, researchers could include measures of latent diffusion networks inferred via NEHA applied to larger data sets of policies, as can be done with NetInf. However, we limit our focus in the current study to simultaneous estimation of covariate and network effects using NEHA. Overall, we expect that NEHA can provide meaningful methodological and substantive improvements with as few as 10 policies in the data.[7]

As a method for network inference applied to diffusion data, NEHA offers several advantages in comparison with NetInf and other related tools.[8] First the variable selection step solves the problem of selecting the number of edges in the network, which must be done manually or through secondary methods with NetInf (Gomez-Rodriguez, Leskovec, and Krause 2010). Second, with NEHA, a coefficient can be estimated with each diffusion tie, whereas with NetInf it is assumed that each diffusion tie is equally strong. Third, NEHA is a discrete-time EHA model, which better matches the measurement precision for data on political choices such as policy adoption than does the continuous-time measurement assumption used in NetInf. Indeed, diffusion is often conceptualized substantively in discrete time because of the constraints of units' decision-making processes, such as legislative sessions.[9]

Substantive interpretation

Before proceeding to applications of NEHA, it is important to consider what the additional parameters that the method adds to an EHA specification mean in substantive terms. By jointly estimating the covariate effects and latent network parameters, NEHA decomposes the diffusion variance into the known covariate effects and the remaining systematic variance that cannot be explained by the covariates. The latent network parameters represent the persistent tendencies for one unit to influence another beyond any measured factors.[10] If the parameters are large and statistically significant, the analyst can conclude strong patterns of interdependence exist in the data, net of covariate effects.[11]

Indeed, researchers can look to the network parameters in NEHA models to evaluate the extent to which units systematically rely on their sources—other units that they tend to follow when making choices. Perhaps one unit follows another even though the two units' covariate profiles look quite different. The network parameters permit summary and visualization, providing a better understanding of how unit-level connections play a role in adoption decisions. Rather than attempting to capture all of the interdependence in observed covariates, NEHA network effects yield valuable insight into how a specific diffusion process unfolds.[12]

We recommend that researchers complete four tasks for a standard interpretation of NEHA results. First, the analyst should compare model fit between a conventional PEHA model and the two versions of NEHA distinguished by the γ parameter. Recall that standard NEHA makes the assumption of constant diffusion effects and estimates one γ. Alternatively, the method can estimate γ to vary across edges, a version we denote NEHA-S. The next step is to interpret the coefficients on the covariates from the chosen estimator (or all of them). This step is straightforward because NEHA is estimated with logistic regression, and thus the standard set of interpretive tools applies. Third, researchers should interpret results from the inferred network itself to maximize NEHA's substantive value to their work. The expansive toolkit of network analysis is available here and the available methods can be adapted to test new and existing hypotheses. Finally, we recommend a null simulation robustness check. This process involves generating fake data from NEHA without latent ties to see whether the number of ties inferred with the actual data is significantly more than would be expected if there were none in the data-generating process.

---

7. NEHA is also flexible with respect to the number of policy areas contained in the data. Researchers with specific expertise in a particular policy area may have better a priori expectations of what the inferred network will look like if they are only examining that area. But they can still use NEHA with a heterogeneous set of policies (see examples below).

8. If the covariates are omitted, NEHA simply represents an alternative to these algorithms.

9. There is one other important difference between NEHA and NetInf to note. Boehmke et al. (2020) and Desmarais et al. (2015) use data sets with hundreds of policies to apply NetInf to rolling temporal windows of adoptions, which infers dynamic networks. This procedure could also be done with NEHA, but we do not demonstrate that functionality here because of space constraints.

---

10. Importantly, the inferred network and its parameters are likely to vary across applications because of differences in model specification, sample selection, and other factors. Thus, the features of the network are not necessarily generalizable to other samples.

11. If there is correlation between the network effects and the covariates, differences will likely emerge between the PEHA and NEHA coefficient estimates. These differences will shrink as the correlation between the network and covariates decreases.

12. By systematically capturing diffusion effects through edge formation rather than covariate specification, NEHA may ultimately promote more parsimonious models of adoption processes.

Assessing diffusion mechanisms

One possible analysis worth additional discussion is the evaluation of diffusion mechanisms (Shipan and Volden 2008). These theoretical processes are conceptually straightforward but can be difficult to distinguish empirically (Maggetti and Gilardi 2014). Thus, researchers often triangulate evidence, using covariates, sample selection, or other research design choices to evaluate them (e.g., Boushey 2016; Karch et al. 2016). All of these options remain available with NEHA. However, the inferred network can provide important new information to assist in this objective as well. Specifically, the systematic variation encoded in the network can yield insight into mechanisms beyond any covariates the analyst included to test mechanisms.

Consider the three classes of diffusion mechanisms that Gilardi (2016, 9–11) discusses: learning, emulation, and competition. Learning involves adoptions that occur due to the observation of previous success, defined with respect to the policy itself, ease of implementation, or its political and electoral implications. Emulation is detached from policy success and involves units aspiring to look like a set of leaders who hold the status or capacity to set policy norms. Finally, competition is characterized by units reacting to one another to attract or retain resources. These unique processes suggest different patterns of interdependence. Competition, for instance, is typically localized to units sharing borders, while emulation implies a small number of leader units. Researchers interested in assessing mechanisms with NEHA should first develop expectations for the structure of the diffusion network based on these processes, then evaluate those expectations empirically with the inferred network. We demonstrate this approach below.

APPLYING NEHA TO DIFFUSION STUDIES

We next apply NEHA to data from four published studies on state policy adoptions: (1) Bricker and LaCombe (2021; 244 policies, various areas); (2) Boushey (2016; 44 policies, criminal justice); (3) Boehmke et al. (2017; 86 policies, various areas); and (4) Karch et al. (2016; 43 policies, interstate compacts). These studies allow us to compare results across different sample sizes, policy domains, and eras while illustrating the benefits of NEHA. We demonstrate heterogeneity in covariate effects, the structure of the inferred network, and even the method's performance relative to PEHA. This approach provides a complete picture of NEHA's role in applied research rather than only instances in which it is the preferred choice (see Harden, Sokhey, and Wilson 2019).

We consider both methodological and substantive aspects of the replication results. First, we consider the out-of-sample (i.e., predictive) performance of PEHA versus NEHA with either a single value of γ (NEHA) or varying γ values (NEHA-S). This quantity informs which version we report and work with throughout the rest of the replication analysis—whichever one performs best in predictive terms. Second, we consider whether adding the latent network terms coincides with any substantively notable changes in the estimates of covariate effects reported in the original results. Third, we present and discuss visualizations of the latent networks inferred and their effects on policy diffusion. Fourth, through a simulation study using the estimated NEHA models, we consider the performance of NEHA in recovering the correct edges and the estimates of the γ values. Fifth, through a similar simulation study, but with γ set to zero, we consider the likelihood that we would draw inferences with NEHA similar to what we estimated in the respective application if there were actually no latent edges in the true data-generating process. Finally, we demonstrate substantive analyses that can be conducted with NEHA but not PEHA.

Bricker and LaCombe (2021) replication

Bricker and LaCombe (2021) examine the adoption of 244 policies across a wide range of issue domains, to evaluate whether policy diffusion is associated with a measure of state similarity developed from citizens' perceptions. We replicate the model on which the authors focus in their substantive interpretation (table 3, col. 1, 384). The original analysis draws on data from policies in the SPID database (Boehmke et al. 2020) that diffused during 1990–2016. This replication has the most policies and observations among the four we consider. It also produces the most extensive latent network and the greatest differences between the NEHA and PEHA results.

We begin by assessing the performance of PEHA and NEHA in predicting held-out data. Doing so helps us to evaluate (1) whether the added complexity in NEHA is justified by improved model fit and (2) whether we should report the version of NEHA with a single constant value of γ or one in which γ varies with each edge. Specifically, we iteratively hold out and predict the observations representing each of the 244 policies. Using the observed values of the dependent variable and the predicted probabilities of adoption, we calculate the area under the precision-recall curve (AUC-PR). This metric represents a smooth average of precision (the proportion of predicted adoptions that are adoptions) and recall (the proportion of adoptions predicted to be adoptions), across all of the possible predicted probability values that could be used as thresholds for classifying a predicted adoption. It falls between 0 and 1, with higher values indicating better performance on both metrics across the spectrum of predicted probabilities.[13] Substantively, it

---

13. The AUC-PR is a particularly appropriate metric here because the data are highly imbalanced—the rate of positive cases is much different from .50 (Cranmer and Desmarais 2017). See the appendix for details on its use in improving NEHA (or PEHA) model specification.

can be compared to an intercept-only model, which in this case would predict .056, or just under 6%, of the data to be adoptions.

The AUC-PR values among the PEHA, NEHA with a constant γ value, and NEHA-S with varying γ values are .1422, .1493, .1484, respectively, which are 2.53–2.67 times better than the intercept-only model. To assess the statistical significance of the differences in AUC-PR values, we take 500 nonparametric bootstrap samples of policies and recalculate differences in predictive performance on each sample. The two-tailed p-values for the differences between NEHA and PEHA, NEHA-S and PEHA, and NEHA and NEHA-S are .016, .040, and .088, respectively. Between the two versions of NEHA, the model with a single value of γ outperforms that with varying γ, so we report on the former going forward.

Table 1 reports the coefficient estimates and in-sample fit for the Bricker and LaCombe (2021) replication. The BIC selects NEHA over PEHA, which is not surprising because the edges are selected to optimize BIC. Several differences in co-efficient estimates between NEHA and PEHA are noteworthy. Among the estimates that are statistically significant at the .05 level, many of the coefficient magnitudes are different by

Table 1. PEHA and NEHA Estimates for Bricker and LaCombe (2021)

	PEHA		NEHA	
	Est.	SE	Est.	SE
Similarity	.191*	.011	.164*	.012
Initiative process	.216*	.075	.427*	.077
Signatures	−.010	.008	−.037*	.009
Population	.076*	.020	.083*	.020
Citizen ideology	.019	.022	−.042	.023
Unified control	−.033	.033	.021	.033
Standard income	.066*	.027	.070*	.027
Legislative professionalism	−.066*	.027	−.058*	.027
Duration	.023	.026	−.063*	.026
Duration^2	−.000	.004	.006	.004
Duration^3	.000	.000	−.000	.000
Intercept	−4.103*	.201	−3.941*	.193
Policy variance	1.179*	.121	.973*	.102
γ	. . .	. . .	.872*	.037
α	. . .	. . .	−4.908	. . .
BIC	32,770		32,249	

Note. Logistic regression coefficients (est.) and standard errors (SE). The dependent variable is an indicator of policy adoption, with 4,838 adop-tions, 85,878 total cases, and 244 policies (intercept-only model: .056). The model includes year fixed effects (not reported) and policy random effects.

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

10% or more (a threshold for substantive significance sug-gested by Harden et al. [2019]). Additionally, the coefficient on citizen ideology is positive and nonsignificant with PEHA but negative and statistically significant at the .10 level with NEHA.

Most notably, the initiative process effect nearly doubles with NEHA, while the associated effect of the number of sig-natures required to qualify a measure more than triples and becomes statistically significant, bringing the results more in line with recent PEHA work on the effect of the initiative on policy adoption (LaCombe and Boehmke 2021). Figure 1 graphs the first difference in the probability of adoption (Y-axis) comparing a hypothetical state without the initiative process to an initiative state with the corresponding signature level (as a percentage of the voting population) on the X-axis. Both estimators indicate a positive effect for states with a low to average signature requirement, and both effects decrease so that the 95% confidence intervals overlap zero once the threshold reaches about 10%. But the first difference for the NEHA model has a steeper slope, which produces a larger change in the effect of the initiative as the associated signature requirement increases. The effect starts out over twice as large for NEHA as for PEHA and then drops more quickly, reach-ing an estimated effect of zero with a signature requirement around 11%. In contrast, the point estimate for PEHA never reaches zero over the observed range of the signature variable.

Moving to the network, NEHA identifies a total of 45 la-tent edges as depicted in figure 2A. We discuss the structure of this network below. Figure 2B combines the NEHA pa-rameters α and γ to compare the probability of adoption over time after a source's adoption to the average predicted probability of adoption when no source has previously adopted.[14] The full effect of γ comes in the first year, with a 7–8 percentage point increase in adoption probability that is statistically distinguishable from the baseline rate. The decay parameter (α) determines how quickly the immediate effect of the source's adoption fades over time. The graph shows fairly slow decay—a less than 3 percentage point decrease after 50 years. This pattern matches the common (although implicit) assumption in the literature that diffusion effects from neighbors' adoptions do not weaken over time.

We next employ simulations to assess NEHA's capacity to correctly identify true edges in the network.[15] We simulated

---

14. The shading indicates 95% confidence intervals, which are com-puted by bootstrapping at the policy level and reestimating the logit model with α fixed at its estimated value. This parameter could be bootstrapped as well, but doing so would create a major computational burden.

15. These simulations also demonstrate that NEHA returns more accurate coefficient estimates when network effects exist in the data-generating process (see the appendix).
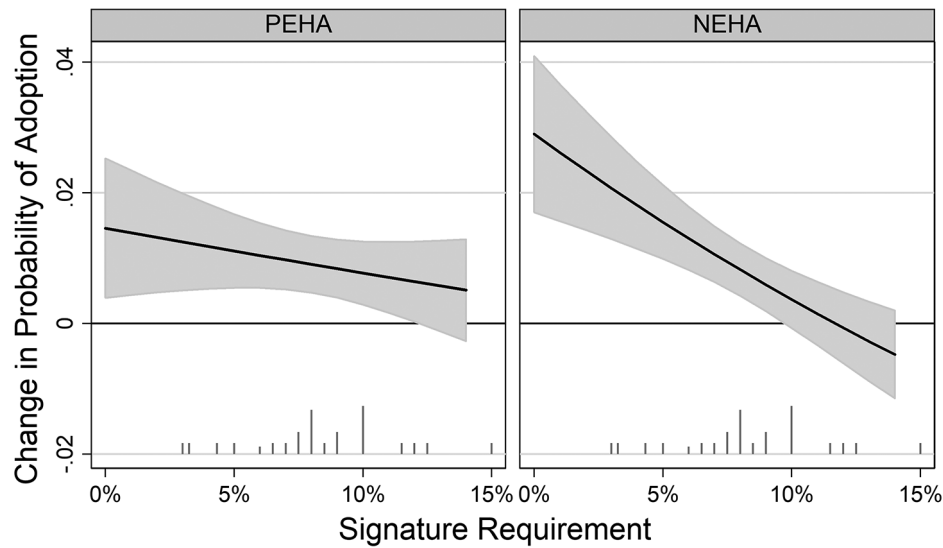
Figure 1. Change in the probability of adoption resulting from adding the initiative process with the given signature threshold (Bricker and LaCombe 2021).

50 data sets using the NEHA estimates and replication data and another 50 data sets using the NEHA estimates but with γ set to zero (a "null simulation"). We then estimated NEHA on all 100 simulated data sets. The overall precision and recall in recovering edges were .9879 and .9716, respectively. These results indicate that, under the data and modeling conditions represented in the replication, edges identified by NEHA have a high likelihood of being true edges. Moreover, NEHA identifies latent edges with high probability. Figure 3 displays the distribution of estimated γ values and indicates a high degree of accuracy in the estimation of γ, although we do see a tendency to slightly underestimate the magnitude of γ in this particular simulation. The maximum and median number of edges

identified in the null simulations are three and one, respectively, indicating that there would be an extremely low likelihood of identifying 45 latent edges if there were no diffusion ties in the true data-generating process. Overall, then, NEHA appears to perform exceptionally well in this application.

Substantive analyses with the inferred network
NEHA offers applied researchers the opportunity to test a richer set of substantive expectations derived from theory based on the inferred edges. We demonstrate three such examples with Bricker and LaCombe's (2021) data. First, we examine the mechanisms of diffusion (net of covariates) via the inferred network's structure, focusing specifically on the
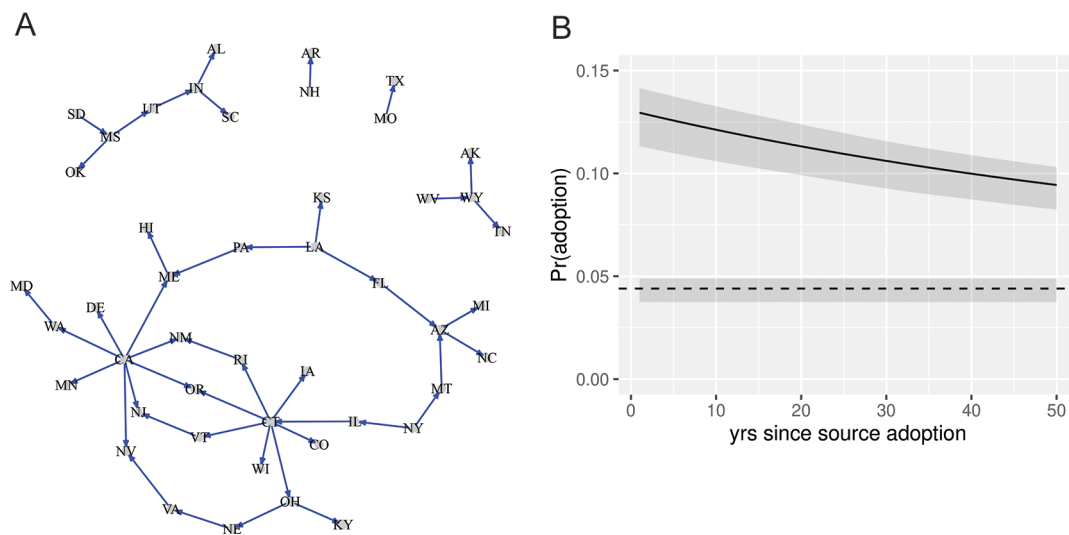

Figure 2. Diffusion network inferred for Bricker and LaCombe (2021) (A, edge structure), and the probability of adoption with one source adopting (B, edge effects). The solid line in B is the effect of one source adopting. The dashed line is the average baseline predicted probability of adoption without a source adopting. Shading indicates 95% confidence intervals.
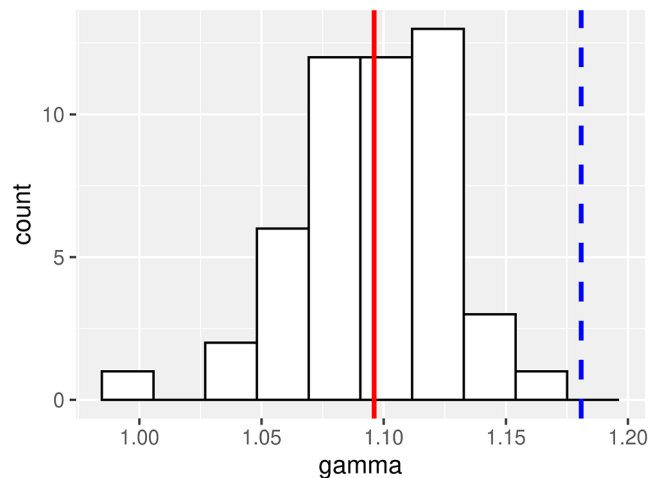
Figure 3. Histogram of estimates of γ in 50 simulated data sets from Bricker and LaCombe (2021). The true value and average estimate are depicted by the dashed and solid lines, respectively.

hierarchy of its ties. Second, in the appendix we report a covariate-based test in which we model tie formation over time to test a hypothesis from the literature about an independent variable of interest (legislative professionalism). Finally, also in the appendix, we use this replication to demonstrate NEHA's ability to infer antidiffusive ties between units in a negative diffusion network.

The inferred network can provide useful insight into diffusion mechanisms. The first step is to establish expectations for the network's structure based on each process. Learning involves adoption under particular success conditions, which suggests that it is a measured response, reflecting precise definitions of success. Accordingly, under this mechanism we would expect a relatively large number of "seeding" states, each sending outgoing ties to a small set of followers seeking a specific type of success. In contrast, emulation refers to aspirational imitation of a few perceived major players; this mechanism would thus be characterized by a hierarchical network with a small number of leaders and a large group of followers. Finally, support for the competition mechanism would appear if many of the targeted relationships in the network involve neighboring states.

Informally, the network appears to be quite hierarchical (fig. 2A). There are no reciprocal ties—either state A influences state B or B influences A, but none of the ties flow in both directions. Furthermore, two states that are known historically as exerting a great deal of influence—New York and California—both sit at the origins of influence pathways that reach many other states. New York reaches many other states through its pathway of influence on Illinois and then Connecticut, while California has direct influence paths sent

to many states. Connecticut itself also emerges as a highly influential state.[16]

We formally evaluate our learning versus emulation expectations with a Gini coefficient calculated on the number of ties sent by each state—a common method for measuring inequality in the distribution of ties across nodes in a network (Hu and Wang 2005). We simulate 50,000 random networks holding the distribution of dyads from the inferred network constant. For each simulated network we then compute the Gini coefficient of the out-degree distribution, which we plot in figure 4. We compare the random networks' Gini coefficients to that of the inferred network (*dashed line*). This procedure tests for concentration of the number of influence ties sent, controlling for the number of edges sent and the degree of reciprocity in the network.

We informally observed that the network appears hierarchical because it includes a small number of historically influential states at the origin of several ties. Figure 4 provides formal evidence on this point. In only one of the 50,000 randomly drawn networks ($p = .00002$) do we observe a Gini coefficient that exceeds our observed Gini coefficient. Thus, we conclude that sent ties in the inferred network are highly unequal in their distribution, reflecting a system in which only a small number of states are influential. This pattern in the data is consistent with the emulation mechanism but not learning. It aligns with the positive coefficient on citizens' perceptions of state similarity (table 1), which could be interpreted as a covariate-based test of the emulation mechanism. Moreover, there is little support for competition. Bricker and LaCombe (2021) do not include a geographic neighbors covariate. Yet even with this potential systematic variation left for the network to absorb, only three of the 45 ties (7%) involve states that share a border. In sum, the structure of the inferred network provides meaningful information for or against the diffusion mechanisms. In this particular case the weight of the evidence suggests support for emulation.

Boushey (2016) replication

Boushey (2016) studies the diffusion of criminal justice policies, finding that laws benefiting large, powerful groups or restricting marginalized groups diffuse more quickly. His PEHA analysis includes data on 44 policies that diffused from 1960 to 2008. We replicate the PEHA model from his table 2, column 1 (206), then estimate the same specification with NEHA. This replication falls on the lower end of our four cases with respect to the number of policies and illustrates a case in which NEHA and PEHA estimates are more similar. Nonetheless, as we show it

---

16. California and Connecticut exhibit the highest out-degree centrality in this network.
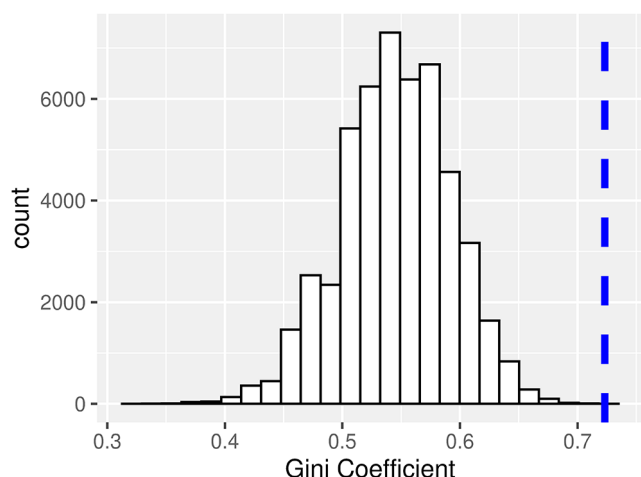
Figure 4. Conditional uniform graph test of out-degree distribution in Bricker and LaCombe (2021). Distribution of Gini coefficients on the sent ties in 50,000 randomly drawn networks with dyad distribution equal to the inferred network. The dashed line is the Gini coefficient from the inferred network.

still offers substantive value with the addition of the inferred network.

Before presenting results, we first evaluate the performance of PEHA and NEHA using the AUC-PR values. Here we find that PEHA, NEHA with a constant γ, and NEHA-S (varying γ) are similar—.2038, .1993, and .1987, respectively—and all improve over the intercept-only model (.058). The two-tailed p-values for the differences between NEHA and PEHA, NEHA-S and PEHA, and NEHA and NEHA-S are .312, .284, and .616, respectively. In this case, then, NEHA yields about the same fit as PEHA. The version of NEHA with a constant value of γ performs the best in terms of AUC-PR, so we report on that one here.

Table 2 reports the estimates and BIC, which favors NEHA over PEHA. There are moderate differences in covariate effects but no major alterations to inferences. Notably, the estimate for Boushey's (2016) main variable of interest, policy congruence, does not change meaningfully in terms of either the effect magnitude or statistical significance. In fact, none of the hypothesis test results changes in terms of statistical significance at the .05 level. There are, however, a few coefficients that exhibit noticeable differences in magnitude or direction. For instance, the PEHA and NEHA estimates for legislative session differ by about 14%, and the coefficient on legislative professionalism changes sign and doubles in magnitude.

Figure 5A shows that NEHA identifies 10 latent network edges, with half being sent by Massachusetts. Boushey (2016) does not discuss Massachusetts specifically, but previous research would have generated the expectation that the state is a leader in this domain. Scholarship from various disciplines

describes it as innovative in domestic violence (Putnam 2003), juvenile justice (Heilbrun, Goldstein, and Redding 2005), auto theft prevention (Cook and MacDonald 2010), and expungement policies (Silva 2010). But none of these studies (or Boushey 2016) uses data or methods to provide supporting evidence of its status as a significant criminal justice policy source. The NEHA structure test described above can do so by evaluating whether Massachusetts sends a disproportionate amount of ties to other states. This expectation is supported: the p-value for the inferred network's out-degree Gini coefficient is .00008.

These latent network ties, while few in number, still contribute to the probability of adoption. Figure 5B presents the

Table 2. PEHA and NEHA Estimates for Boushey (2016)

	PEHA		NEHA	
	Est.	SE	Est.	SE
Policy congruence	.308*	.056	.308*	.056
Gubernatorial election year	−.082	.086	−.032	.086
Off election year	−.004	.072	.026	.072
National crime salience	1.097*	.383	1.118*	.385
Democratic Party strength	.005*	.002	.005*	.002
Democratic governor	.040	.066	.012	.066
Legislative session	1.883*	.217	2.145*	.219
Neighbors	2.579*	.099	2.574*	.100
Ideological distance	−.041*	.003	−.043*	.003
Legislative professionalism	−.154	.323	.312	.318
Political ideology	−.006*	.003	−.007*	.003
Crime control spending	−.001	.002	−.001	.002
Crime control spending^2	.000	.000	.000	.000
Violent crime rate	.020	.018	.006	.018
% white	.006	.006	−.001	.006
Per capita income	.031*	.009	.030*	.009
Log population	.043	.054	−.023	.051
Time	−.018	.022	−.028	.022
Time^2	−.001	.001	−.001	.001
Time^3	.000	.000	.000	.000
Intercept	−6.455*	1.094	−5.144*	1.051
γ	. . .	. . .	1.506*	.110
α	. . .	. . .	−7.302	. . .
BIC	10,207		10,059	

Note. Logistic regression coefficients (est.) and standard errors (SE) clustered by state year. The dependent variable is an indicator of policy adoption, with 1,543 adoptions, 26,479 total cases, and 44 policies (intercept-only model: .058).
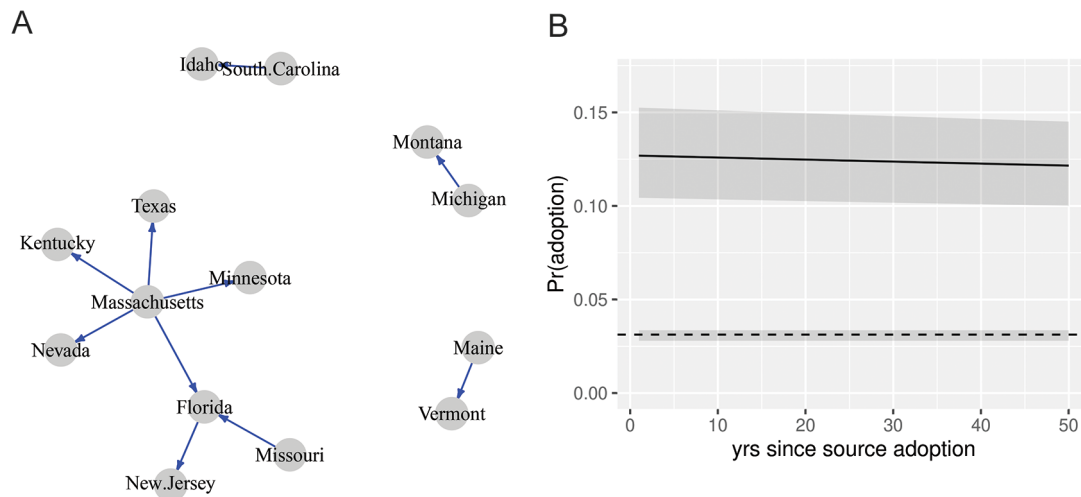* $p < .05$ (two tailed).

A

B

Figure 5. Diffusion network inferred for Boushey (2016) (A, edge structure), and the probability of adoption with one source adopting (B, edge effects). The solid line in B is the effect of one source adopting. The dashed line is the average baseline predicted probability of adoption without a source adopting. Shading indicates 95% confidence intervals.

effect of adoption by a latent source over time. Compared to the previous example, a larger value of γ produces a larger initial increase from approximately .03 for PEHA to .13 for NEHA, while the value of α indicates negligible decay over time.

Considering the performance of NEHA in the simulations, we see precision and recall values of .8355 and .7560, respectively. The distribution of estimated γ values given in figure 6 again indicates a high degree of accuracy in the estimation of γ. The maximum and median number of edges identified is three and one, respectively, implying a very low likelihood of identifying 10 latent edges if there were no true edges. As in the prior case, edges identified by NEHA have a high likelihood of being true edges, and NEHA identifies latent edges with high probability.

Boehmke et al. (2017) replication

Boehmke et al. (2017) estimate a PEHA model to obtain parameters from which to simulate the spread of policies seeded in different states. Their analysis uses data from Boehmke and Skinner (2012) and includes 86 policies covering 1960–2009. The contribution to predictive performance from NEHA is particularly notable in this replication, as the original specification includes a measure of latent networks inferred from NetInf (lagged source adoptions) and a geographic contiguity variable (lagged neighbor adoptions). As such, the improvement in fit from NEHA represents a contribution that goes beyond that of two means of explicitly accounting for interdependence via covariates. The AUC-PR values for the PEHA, NEHA, and NEHA-S with varying γ are .1430, .14716, and .1455, respectively (intercept-only model: .053). The two-tailed p-values for the differences between NEHA and PEHA, NEHA-S and

PEHA, and NEHA and NEHA-S are .092, .304, and .220, respectively. The NEHA with constant γ again fits best, so we proceed with analyzing and interpreting that version.

Table 3 reports the coefficient estimates and in-sample fit (BIC), which favors NEHA. The estimates show some magnitude changes between PEHA and NEHA but not statistical significance. For instance, the coefficient on legislative professionalism is nearly 25% smaller, while the coefficient on unified Republican control more than doubles in magnitude with NEHA. The estimates on state-to-state interdependence (sources and neighbors) are similar across the models.
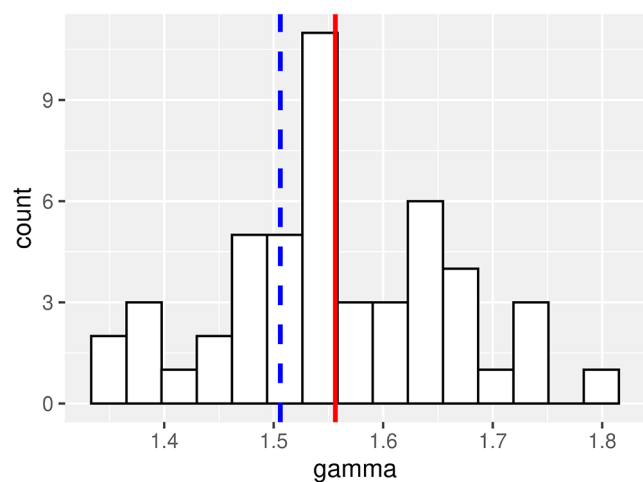

Figure 6. Histogram of estimates of γ in 50 simulated data sets from Boushey (2016). The true value and average estimate are depicted by the dashed and solid lines, respectively.

Table 3. PEHA and NEHA Estimates for Boehmke et al. (2017)

	PEHA		NEHA	
	Est.	SE	Est.	SE
Lagged source adoptions	8.527*	.438	8.174*	.437
Lagged neighbor adoptions	.393*	.022	.380*	.022
Personal income	.574*	.075	.563*	.075
Population	.091*	.028	.093*	.028
Legislative professionalism	−1.089	.687	−.827	.684
State citizen ideology	.010*	.004	.010*	.004
Unified Republican control	−.020	.076	−.056	.077
Unified Democratic control	.063	.066	.061	.067
Time	−.135*	.018	−.149*	.018
Time²	.007*	.001	.008*	.001
Time³	−.000*	.000	−.000*	.000
Intercept	−5.411*	.278	−5.364*	.276
γ	. . .	. . .	1.604*	.113
α	. . .	. . .	−7.124	. . .
BIC	16,757		16,573	

Note. Logistic regression coefficients (est.) and standard errors (SE) clustered by state policy. The dependent variable is an indicator of policy adoption, with 2,222 adoptions, 44,457 total cases, and 86 policies (intercept-only model: .053). The model includes state fixed effects (not reported).
* $p < .05$ (two tailed).

Nonetheless, the NEHA model still infers 11 latent ties, which are displayed in figure 7A. Figure 7B shows that adoption by a diffusion tie source increases the probability of adoption by a target state, with a slow decay over time. The number of latent edges and the magnitude of the effect of latent ties are very similar to our findings from the Boushey (2016) replication (see fig. 5). However, none of the individual ties is the same, which emphasizes the fact that the unexplained patterns of emulation are likely to depend on the specific policy domains and model specifications considered in a given study.

Substantively, no single state emerges as a clear leader in this application; the state with the most outgoing ties is Washington, which is a source state for just two other states. Instead, we mostly find several pairs of states, one as a source for the other. Recall that the model controls for interdependence via policy source states and contiguity. Thus, these ties likely represent targeted relationships between specific dyads that researchers would have difficulty capturing with another covariate. This structure underscores an advantage of NEHA; the method can account for interdependence while remaining agnostic on its source.

NEHA performs very well in the simulations with precision and recall over the 50 simulation iterations of .9812 and .9236, respectively. The distribution of the estimated values of γ, reported in figure 8, is very close to the true value. In the null simulations, where the value of γ is set to zero, the median number of edges inferred is zero, and the maximum over the 50 simulation iterations is one. These results suggest that, especially in data sets with a larger number of policies, there is little potential cost with respect to incorrect inferences with NEHA. Aside from somewhat longer computing times, the method is highly accurate in identifying latent diffusion ties, and in the event that there are no latent diffusion ties, NEHA is likely to reduce to PEHA.
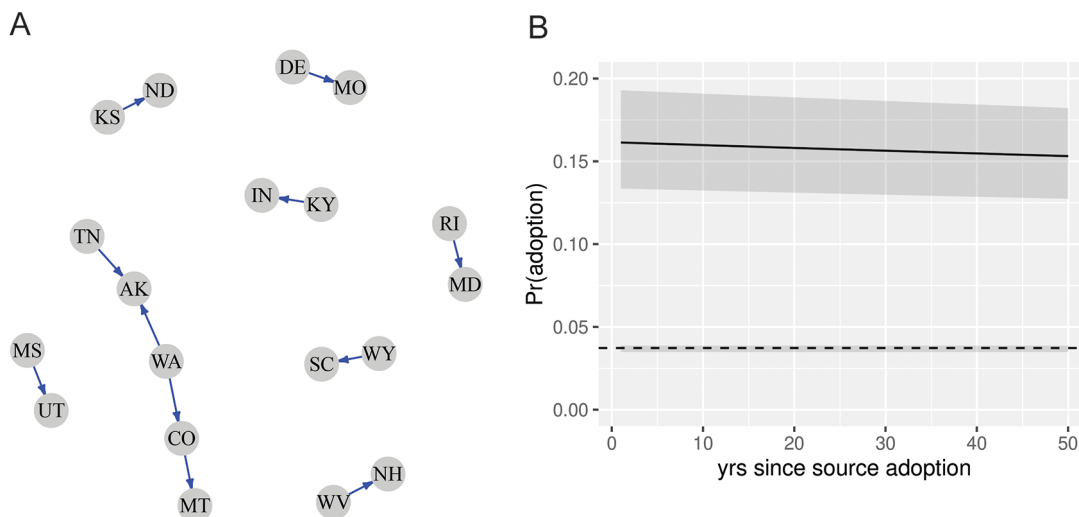

Figure 7. Diffusion network inferred for Boehmke et al. (2017) (A, edge structure), and the probability of adoption with one source adopting (B, edge effects). The solid line in B is the effect of one source adopting. The dashed line is the average baseline predicted probability of adoption without a source adopting. Shading indicates 95% confidence intervals.
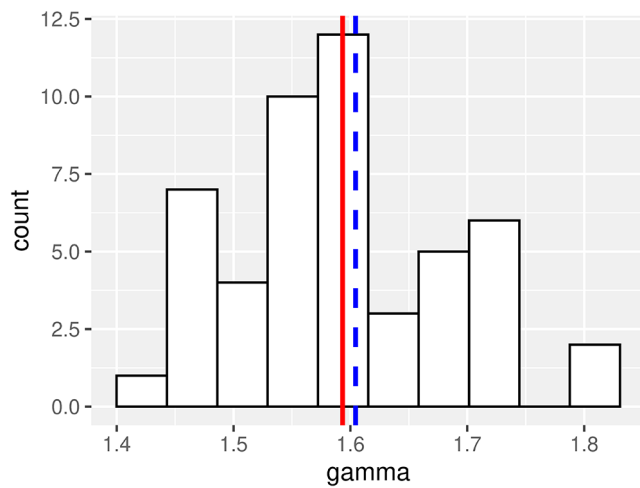
Figure 8. Histogram of estimates of γ in 50 simulated data sets from Boehmke et al. (2017). The true value and average estimate are depicted by the dashed and solid lines, respectively.

Karch et al. (2016) replication

Karch et al. (2016) study the adoption of 43 interstate compacts during 1951–2012. They focus on the possibility of pro-innovation bias in extant studies, leveraging the fact that their database of compacts includes several that were adopted by small numbers of states. This replication includes the fewest policies among the four. We replicate their primary model in table 2 (90) and find that NEHA outperforms PEHA and PEHA outperforms NEHA-S in the hold-one-policy-out prediction. The AUC-PR values among the PEHA, NEHA, and NEHA-S models are .0455, .0459, and .0422, respectively (intercept-only model: .018). The two-tailed p-value for the difference between NEHA and PEHA, NEHA-S and PEHA, and NEHA and NEHA-S are .848, .068, and .000, respectively. Between the two versions of NEHA, the model with a single value of γ outperforms that with varying γ. Neither estimator clearly emerges as the better out-of-sample fit. However, a researcher would still be justified in using NEHA in this case if he or she were interested in the substantive implications of the inferred network.

Table 4 reports the coefficient estimates and BIC, which indicates that NEHA outperforms PEHA. The coefficient estimates show moderate differences between the two models. Several of the coefficient magnitudes differ by 10% or more (e.g., complexity, Republican governor, and legislative professionalism). Additionally, for one of the interaction covariates—the effect of government expenditures per capita on the adoption of national-scale compacts (traditional × expenditures per capita), the PEHA coefficient is not statistically significant at the .05 level but is significant with NEHA.

In addition to these differences, NEHA identifies 14 latent edges (fig. 9A). This network provides the opportunity to ex-

amine Karch et al.'s (2016) theoretical framework in a manner that is not readily available with PEHA. The authors posit that pro-innovation bias leads researchers to inflate the role of

Table 4. PEHA and NEHA Estimates for Karch et al. (2016)

	PEHA		NEHA	
	Est.	SE	Est.	SE
Traditional	5.689*	.800	5.770*	.799
Neighbors	1.434*	.207	1.400*	.207
Previous adopters	3.999*	.7689	4.213*	.775
Complexity	−.402	.272	−.285	.280
Interest group role	1.690*	.281	1.624*	.287
Republican governor	−.127	.208	−.152	.207
Unified	.068	.205	.080	.205
Democratic legislature	−.299*	.115	−.337*	.115
Income per capita	−.387	.302	−.275	.282
Expenditures per capita	−.566*	.271	−.661*	.276
Population	−.385	.141*	−.388*	.143
% college educated	.345	.207	.299	.207
% urban	.135	.115	.142	.117
Legislative professionalism	−.074	.114	−.058	.110
Traditional × neighbors	−.992*	.215	−.910*	.215
Traditional × previous adopters	−3.808*	.776	−3.910*	.782
Traditional × complexity	−.485	.370	−.666	.380
Traditional × interest group role	−1.690*	.299	−1.635*	.305
Traditional × Republican governor	.257	.228	.270	.227
Traditional × unified	.032	.224	.074	.224
Traditional × Democratic legislature	.257*	.122	.249*	.121
Traditional × income per capita	.396	.327	.380	.311
Traditional × expenditures per capita	.550	.288	.655*	.291
Traditional × population	.287	.151	.262	.154
Traditional × % college educated	−.590*	.240	−.606*	.241
Traditional × % urban	.006	.127	.045	.129
Traditional × legislative professionalism	−.047	.125	−.146	.122
Intercept	−9.211	.788*	−9.502*	.788
γ	. . .	. . .	2.073*	.146
α	. . .	. . .	−6.734	. . .
BIC	7,135		6,925	

Note. Logistic regression coefficients (est.) and standard errors (SE) clustered by state year. The dependent variable is an indicator of policy adoption, with 783 adoptions, 42,469 total cases, and 43 policies (intercept-only model: .018).
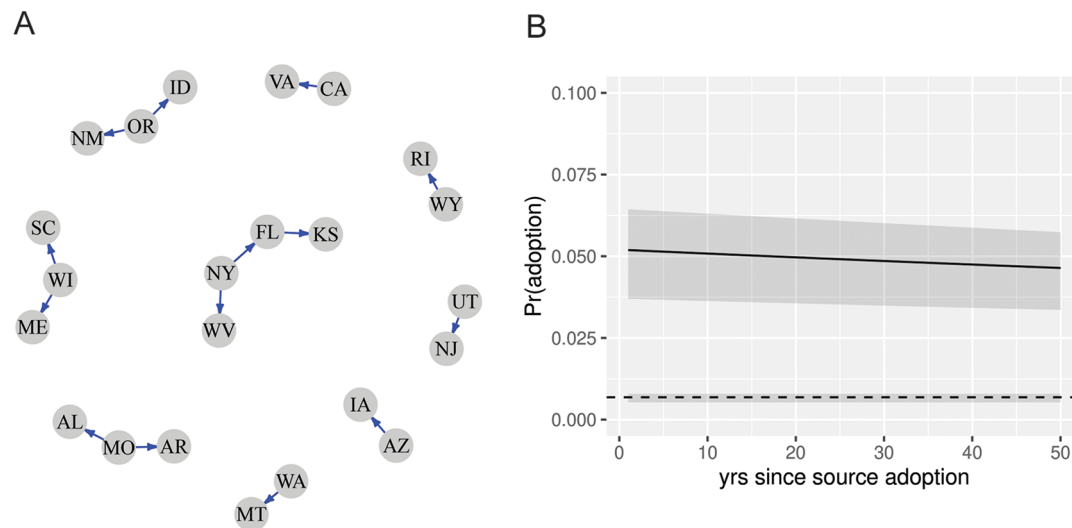
* $p < .05$ (two tailed).

Figure 9. Diffusion network inferred for Karch et al. (2016) (*A*, edge structure), and the probability of adoption with one source adopting (*B*, edge effects). The solid line in *B* is the effect of one source adopting. The dashed line is the average baseline predicted probability of adoption without a source adopting. Shading indicates 95% confidence intervals.

geography in policy diffusion and overlook the extent to which states learn from previous adopters, regardless of proximity (88). Thus, we hypothesize that an effect of noncontiguous source states on adoption probability remains after controlling for neighbors.

Figure 9*A* suggests support for this expectation. The network is composed of several noncontiguous dyads as well as four states (Oregon, New York, Missouri, and Wisconsin) that send ties to multiple mostly noncontiguous states. In all, only two of the 14 edges connect geographic neighbors. Moving to figure 9*B*, note that the base adoption rate is quite low in this application, but the relative impact of a source adopting on adoption probability is four to five times higher and statistically distinguishable from the base rate. This finding is important in light of Karch et al.'s (2016) theoretical point. Although the model shows that contiguity matters in the diffusion of interstate compacts (see table 4), the strong effect (and slow decay) of latent sources—most of which are not contiguous—supports the authors' claim that learning from source states, regardless of their geography, is important when considering policies that may not diffuse to many states.

Considering the performance of NEHA in the simulations, we see precision and recall values of .9028 and .8243 over the 50 iterations, respectively. The distribution of estimated γ values is given in figure 10, and it again indicates a high degree of accuracy in the estimation of γ. The maximum and median number of edges identified in the null simulations are three and one, respectively, indicating that there would be a very low likelihood of identifying 14 latent edges if there was no diffusion tie in the true data-generating process. In general, NEHA

appears to perform well even in a data set with comparatively fewer policies.

CONCLUSION

Interdependence between states or countries is a hallmark of political processes, and the diffusion of ideas, innovations, or policies is no exception. Research on decisions by political actors across the social sciences has long paid close attention to features of the actors or the innovations themselves. However, empirically addressing interdependence in this work is often a
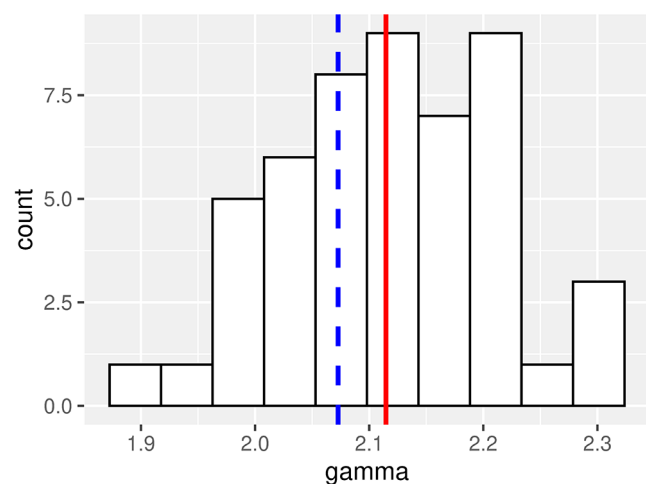

Figure 10. Histogram of estimates of γ in 50 simulated data sets from Karch et al. (2016). The true value and average estimate are depicted by the dashed and solid lines, respectively.

lower priority, typically resulting in the use of measures of similarity on various dimensions that are readily observable. But these operationalizations of diffusion impose a static, pre-existing structure on the ties that connect units. As recent advances in network inference show, such interdependence often takes on a more complex, dynamic structure.

While the incorporation of latent diffusion networks into statistical models of political adoption decisions is a promising addition to these studies, the existing methodology is inadequate for doing so. Most notably, methods that infer latent diffusion ties connecting units require the analyst to estimate two separate statistical models, both of which are misspecified if the analyst believes diffusion and observed covariates influence adoption. These methods do not facilitate the use of common functional forms in conventional EHA models, nor do they easily permit the inclusion of network-based covariates. In contrast, NEHA offers researchers a methodological framework within which to simultaneously estimate the effects of observed covariates and infer latent diffusion networks. We show that, when latent ties are present, NEHA tends to improve model fit and can alter coefficient estimates and their standard errors. Furthermore, in simulations calibrated to our replication studies, we show that there is little risk in using NEHA rather than PEHA, as it is likely to reduce to PEHA in the event that there are no latent ties.

Importantly, beyond NEHA's ability to appropriately account for network structure in adoption models, the networks that it infers can provide novel and useful substantive information by revealing which units tend to be leaders or followers, conditional on the covariates. The method can assist applied researchers in refining model specifications, adjudicating between diffusion mechanisms, or testing new or existing hypotheses by analyzing network ties and structure. In short, beyond its methodological contributions, NEHA holds the potential to enrich the substantive insights scholars can draw from data on the spread of political decisions.

Our central conclusion from this research is that NEHA is a useful addition to the set of tools employed by scholars of diffusion. The estimator advances researchers' ability to understand diffusion processes on methodological and substantive dimensions without burdening them with additional complexity. Indeed, it is straightforward to apply in our accompanying R package (see the appendix) and is no more difficult to interpret than a standard EHA model. NEHA carries essentially no disadvantages that would render it suboptimal compared to the PEHA estimator. Instead, it provides a more unified, comprehensive approach to testing hypotheses and addressing latent diffusion processes in applied research on decisions made by governmental jurisdictions.

REFERENCES

Berry, Frances Stokes, and William D. Berry. 1990. "State Lottery Adoptions as Policy Innovations: An Event History Analysis." *American Political Science Review* 84 (2): 395–415.

Boehmke, Frederick J., Mark Brockway, Bruce A. Desmarais, Jeffrey J. Harden, Scott LaCombe, Fridolin Linder, and Hanna Wallach. 2020. "SPID: A New Database for Inferring Public Policy Innovativeness and Diffusion Networks." *Policy Studies Journal* 48 (2): 517–45.

Boehmke, Frederick J., Abigail Matthews Rury, Bruce A. Desmarais, and Jeffrey J. Harden. 2017. "The Seeds of Policy Change: Leveraging Diffusion to Disseminate Policy Innovations." *Journal of Health Politics, Policy and Law* 42 (2): 285–307.

Boehmke, Frederick J., and Paul Skinner. 2012. "State Policy Innovativeness Revisited." *State Politics and Policy Quarterly* 12 (3): 303–29.

Boushey, Graeme. 2016. "Targeted for Diffusion? How the Use and Acceptance of Stereotypes Shape the Diffusion of Criminal Justice Policy Innovations in the American States." *American Political Science Review* 110 (1): 198–214.

Box-Steffensmeier, Janet M., and Bradford S. Jones. 2004. *Event History Modeling: A Guide for Social Scientists*. Cambridge: Cambridge University Press.

Bricker, Christine, and Scott LaCombe. 2021. "The Ties That Bind Us: The Influence of Perceived State Similarity on Policy Diffusion." *Political Research Quarterly* 74 (2): 377–87.

Brooks, Sarah M. 2005. "Interdependent and Domestic Foundations of Policy Change: The Diffusion of Pension Privatization around the World." *International Studies Quarterly* 49 (2): 273–94.

Cook, Philip J., and John MacDonald. 2010. "Public Safety through Private Action: An Economic Assessment of Bids, Locks, and Citizen Cooperation." NBER Working paper 15877, National Bureau of Economic Research, Cambridge, MA.

Cook, R. Dennis, and Liliana Forzani. 2009. "Likelihood-Based Sufficient Dimension Reduction." *Journal of the American Statistical Association* 104 (485): 197–208.

Cranmer, Skyler J., and Bruce A. Desmarais. 2017. "What Can We Learn from Predictive Modeling?" *Political Analysis* 25 (2): 145–66.

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.

Garrett, Kristin N., and Joshua M. Jansa. 2015. "Interest Group Influence in Policy Diffusion Networks." *State Politics and Policy Quarterly* 15 (3): 387–417.

Gilardi, Fabrizio. 2016. "Four Ways We Can Improve Policy Diffusion Research." *State Politics and Policy Quarterly* 16 (1): 8–21.

Gilardi, Fabrizio, and Katharina Füglister. 2008. "Empirical Modeling of Policy Diffusion in Federal States: The Dyadic Approach." *Swiss Political Science Review* 14 (3): 413–50.

Gomez-Rodriguez, Manuel, Jure Leskovec, and Andreas Krause. 2010. "Inferring Networks of Diffusion and Influence." Presented at the 16th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, July 24–28.

Graham, Erin R., Charles R. Shipan, and Craig Volden. 2013. "The Diffusion of Policy Diffusion Research in Political Science." *British Journal of Political Science* 43 (3): 673–701.

Gray, Virginia. 1973. "Innovation in the States: A Diffusion Study." *American Political Science Review* 67 (4): 1174–85.

Grossback, Lawrence J., Sean Nicholson-Crotty, and David A. M. Peterson. 2004. "Ideology and Learning in Policy Diffusion." *American Politics Research* 32 (5): 521–45.

Harden, Jeffrey J., Anand E. Sokhey, and Hannah Wilson. 2019. "Replications in Context: A Framework for Evaluating New Methods in Quantitative Political Science." *Political Analysis* 27 (1): 119–25.

Heilbrun, Kirk, Naomi E. Sevin Goldstein, and Richard E. Redding. 2005. *Juvenile Delinquency: Prevention, Assessment, and Intervention.* New York: Oxford University Press.

Hinkle, Rachael K. 2015. "Into the Words: Using Statutory Text to Explore the Impact of Federal Courts on State Policy Diffusion." *American Journal of Political Science* 59 (4): 1002–21.

Hu, Hai-Bo, and Lin Wang. 2005. "The Gini Coefficient's Application to General Complex Networks." *Advances in Complex Systems* 8 (1): 159–67.

Karch, Andrew, Sean C. Nicholson-Crotty, Neal D. Woods, and Ann O'M. Bowman. 2016. "Policy Diffusion and the Pro-innovation Bias." *Political Research Quarterly* 69 (1): 83–95.

Kreitzer, Rebecca J. 2015. "Politics and Morality in State Abortion Policy." *State Politics and Policy Quarterly* 15 (1): 41–66.

LaCombe, Scott J., and Frederick J. Boehmke. 2021. "The Initiative Process and Policy Innovation in the American States." *State Politics and Policy Quarterly* 21 (3): 286–305.

Lewis, Jeffrey B., and Keith T. Poole. 2004. "Measuring Bias and Uncertainty in Ideal Point Estimates via the Parametric Bootstrap." *Political Analysis* 12 (2): 105–27.

Linder, Fridolin, Bruce A. Desmarais, Matthew Burgess, and Eugenia Giraudy. 2018. "Text as Policy: Measuring Policy Similarity through Bill Text Reuse." *Policy Studies Journal* 48 (2): 546–74.

Maggetti, Martino, and Fabrizio Gilardi. 2014. "Problems (and Solutions) in the Measurement of Policy Diffusion Mechanisms." *Journal of Public Policy* 36 (1): 1–21.

Mallinson, Daniel J. 2021a. "Policy Innovation Adoption across the Diffusion Life Course." *Policy Studies Journal* 49 (2): 335–58.

Mallinson, Daniel J. 2021b. "Who Are Your Neighbors? The Role of Ideology and Decline of Geographic Proximity in the Diffusion of Policy Innovations." *Policy Studies Journal* 49 (1): 67–88.

Parinandi, Srinivas C. 2020. "Policy Inventing and Borrowing among State Legislatures." *American Journal of Political Science* 64 (4): 852–68.

Putnam, Jan. 2003. "Battered Mothers Fight to Survive the Family Court System." *Research and Action Report* 25 (1): 7–9.

Rogers, Everett. 1962. *Diffusion of Innovation.* New York: Free Press.

Shipan, Charles R., and Craig Volden. 2008. "The Mechanisms of Policy Diffusion." *American Journal of Political Science* 52 (4): 840–57.

Silva, Lahny R. 2010. "Clean Slate: Expanding Expungements and Pardons for Non-violent Federal Offenders." *University of Cincinnati Law Review* 79:155–205.

Tallberg, Jonas, Magnus Lundgren, Thomas Sommerer, and Theresa Squatrito. 2020. "Why International Organizations Commit to Liberal Norms." *International Studies Quarterly* 64 (3): 626–40.

Volden, Craig. 2006. "States as Policy Laboratories: Emulating Success in the Children's Health Insurance Program." *American Journal of Political Science* 50 (2): 294–312.

Walker, Jack L. 1969. "The Diffusion of Innovations among the American States." *American Political Science Review* 63 (3): 880–99.

Wotipka, Christine Min, and Kiyoteru Tsutsui. 2008. "Global Human Rights and State Sovereignty: State Ratification of International Human Rights Treaties, 1965–2001." *Sociological Forum* 23 (4): 724–54.