Phytopathology® • 2026 • 116:239-251 • https://doi.org/10.1094/PHYTO-06-25-0199-R

Ecology and Epidemiology

How Do Growers Respond to Host Resistance? A Conditional Gaussian Bayesian Network for Causal Inference of Fungicide Cost Savings

Jae Young Hwang,[1] Sharmodeep Bhattacharyya,[2] Shirshendu Chatterjee,[3] Thomas L. Marsh,[4] Joshua F. Pedro,[3] and David H. Gent[1,†]

[1] U.S. Department of Agriculture-Agricultural Research Service, Forage Seed and Cereal Research Unit, Corvallis, OR 97331, U.S.A.
[2] Department of Statistics, Oregon State University, Corvallis, OR 97331, U.S.A.
[3] Department of Mathematics, City University of New York, New York City, NY 10031, U.S.A.
[4] School of Economic Sciences, Washington State University, Pullman, WA 99163, U.S.A.
Accepted for publication 24 September 2025.

Abstract

The economic value of cultivars resistant to disease is of great interest, but how growers change their fungicide use in response to host resistance may be nuanced. We draw upon a well-described data set of the incidence of hop plants with powdery mildew and associated production metadata and demonstrate the utility of Bayesian networks as a framework for quantifying causal relationships for fungicide use and cost in response to host resistance. Conditional Gaussian Bayesian network models applied to cultivars differing in race-specific resistance to powdery mildew revealed cultivar resistance to powdery mildew influenced disease levels in early spring, which had a causal effect on how often and what fungicides growers later applied. Annual costs depended on not only the number of applications made but also the specific types of fungicides growers selected. Fungicide costs were little changed on cultivars that possessed race-specific resistance to only one of two extant strains of the pathogen. For cultivars with resistance to both pathogen strains, annual costs of fungicides were reduced commensurate with the level of resistance. Predicted values from the Bayesian networks and simulation indicate that growers apply a baseline level of fungicide, independent of cultivar resistance. Fungicide cost savings result from how fungicide inputs differentially scale with the incidence of powdery mildew and the type of fungicides used. Our analyses indicate that for a high-value crop, deployment of disease resistance may cause complex and unexpected changes in growers' fungicide use patterns that may not be obvious in simplified randomized controlled trials.

Keywords: data science, disease resistance, economics, epidemiology

The goal of many studies is to understand the causal effect of an intervention on some outcome. In plant disease contexts, as examples, we might want to understand how the number of fungicide applications made by a grower change depending on their use of a particular cultural practice or when receiving new information from an extension educator. The intervention and outcome may be associated with other covariates in a causal network, and understanding these associations is key to enabling the prediction of the intervention subject to these covariates (Spirtes 2010). In an associational analysis, this is done through estimation of parameters of a distribution from samples that are extracted from that distribution and then estimating relationships defined in terms of a joint distribution of the observed variables. Examples of associational analyses are correlation, regression, dependence, conditional independence, and likelihoods. In contrast, a causal relationship is a concept that infers conditional probabilities under changing conditions that may be induced by a treatment or an external intervention (Pearl 2009, 2010). Causal relationships cannot be defined solely from the joint distribution of observed variables. Examples of causal inference methods are randomization, instrumental variables, and intervention. Distinguishing between association and causal concepts is important to understand the underlying assumptions of each framework to substantiate scientific claims (Pearl 2010). Pearl (2010) argues that causal analysis relies on assumptions that cannot be met, inferred from, or defined in terms of associations alone.

An outcome variable of interest typically is the result of numerous covariates that may interact in complex ways. For concreteness, the well-known interaction of host, pathogen, environment, and management intervention in the plant disease tetrahedron is a network of interactions that leads to a disease outcome. Effects on networks concern how the effect of an intervention may rely on the structure of the network of variables or how an intervention may affect the structure of the network itself (VanderWeele and An 2013). Controlled randomized experiments are designed to identify and estimate the causal effects of an intervention on an outcome variable, usually given a small set of covariates defined in a network. This concept can be extended to observational studies when important covariates are also measured.

There are many statistical methods for causal inference, yet only a handful have been applied in plant pathology contexts (Lankau et al. 2020). A Bayesian network is a straightforward machine learning model and well suited to understanding model uncertainty and causality. Bayesian networks are a type of graphical model that overcomes the limitations of structural equation modeling by representing a system as a network of interactions between predictors from a primary cause to an outcome variable with explicit

†Corresponding author: D. H. Gent; dave.gent@usda.gov

The use of trade, firm, or corporation names in this publication is for the information and convenience of the reader. Such use does not constitute an official endorsement or approval by the U.S. Department of Agriculture or the Agricultural Research Service of any product or service to the exclusion of others that may be suitable.

Funding: Funding was provided by the U.S. Department of Agriculture-Agricultural Research Service CRIS Project (2072-21000-061-000-D), the U.S. Department of Agriculture National Institute of Food and Agriculture Specialty Crop Research Initiative (award number 2021-51181-35901), and the National Science Foundation Division of Mathematical Sciences (2154564).

***e*-Xtra:** Supplementary material is available online.

The author(s) declare no conflict of interest.

cause-and-effect assumptions. The graphical structure $G = (V, A)$ of a Bayesian network is a directed acyclic graph (DAG), where V is the node set of random variables and A is the edge (or arc) set. A DAG is formed by nodes and edges connecting pairs of nodes through probabilistic relationships to build an overall structure between variables, where each edge is directed from one node to another without cyclic loops. An example of a DAG is presented in Figure 1 for our motivating pathosystem, hop powdery mildew caused by the fungus *Podosphaera macularis*.

The structural specification of a Bayesian network makes explicit our knowledge of how a system works through defined causal relationships between variables in the model. The way these variables connect is calculated from all possible configurations between nodes and edges, creating the network structure or the basic connections. The connections between nodes provide the basic building blocks of a Bayesian network graph and its probabilistic properties. These connections specify postulated cause-and-effect relationships, whether these relationships are direct or indirect.

Bayesian networks use Bayes' theorem and related Bayesian learning algorithms to calculate posterior probabilities of outcome variables (Marcot 2012). Bayesian networks are defined by a network structure, where each node $v_i \in V$ refers to a random variable X_i. The graphical structure of a Bayesian network is a DAG that defines a factorization of the joint probability distribution of $V = \{X_1, X_2, \ldots, X_v\}$. The form of the factorization is given by the Markov property of Bayesian networks (Scutari 2010), which indicates that all random variable X_i directly depend only on parent nodes Π_{X_i}:

$$P\{X_1, X_2, \ldots, X_v\} = \Pi_{i=1}^{v} P(X_i|\Pi_{X_i}) \text{ (for discrete variables)}$$

$$f\{X_1, X_2, \ldots, X_v\} = \Pi_{i=1}^{v} f(X_i|\Pi_{X_i}) \text{ (for continuous variables)}$$

where $\Pi_{X_i} = \{\text{parents of } X_i\}$. There are multiple, and possibly interactive, steps to learning the graphical structure of a Bayesian network. The first step is to construct a knowledge-based structure of the postulated causal relationships between variables present in available data using expert knowledge, theory, and existing literature. The postulated causal network determines which edges should be extant in the graph and relationships that should never exist in the real world. The edges presented in the whitelist (denoted as red lines in Figure 1) are included in every Bayesian network structure; edges presented in the blacklist are prohibited.

The second step is to learn parameters to deal with the estimation of the parameters of the global distribution. In Bayesian network parameter learning, a domain $\mathcal{D} = \{V, G, S\}$ comprises three components: (i) the random variables $V = \{X_1, X_2, \ldots, X_v\}$ corresponding to nodes of the Bayesian network; (ii) a DAG G encoding the conditional dependencies given the variables; and (iii) associated sample S (Zhou et al. 2016). Within the domain $\mathcal{D}$, the objective of parameter learning is to conduct parameter tuning for each conditional distribution $p(X_i|\Pi_{X_i})$. A conditional probability table related to all variables specifies the probability $p(X_i|\Pi_{X_i})$ of each value given the parent nodes defined by graph G. We provide further explanation of Bayesian network parameter learning in the Supplementary Material.

A disease intervention of singular importance across plant pathology and plant breeding is deployment of host resistance for disease control. Host resistance is often considered the intervention that is most efficient, is most economical, and has the least environmental impact (Hückelhoven and Pillen 2024). Although the benefits of host resistance seem clear, empirical evaluation of its benefits can be difficult in practice. The impact of host resistance could be based on various metrics, such as reductions in disease levels, improvements in product quality or safety, yield, profit, reduction of inputs, changes in poverty or welfare, or enhanced trade (Morris and Heisey 2003). Although each of these metrics is conceptually simple, their measurement is not trivial. Measurement may be biased or inaccurate because of issues related to the external validity of controlled replicated experiments, simultaneous changes in other management practices, or ignoring counterfactuals about outcomes in the absence of host resistance, among others (Mooney et al. 2022; Morris and Heisey 2003).

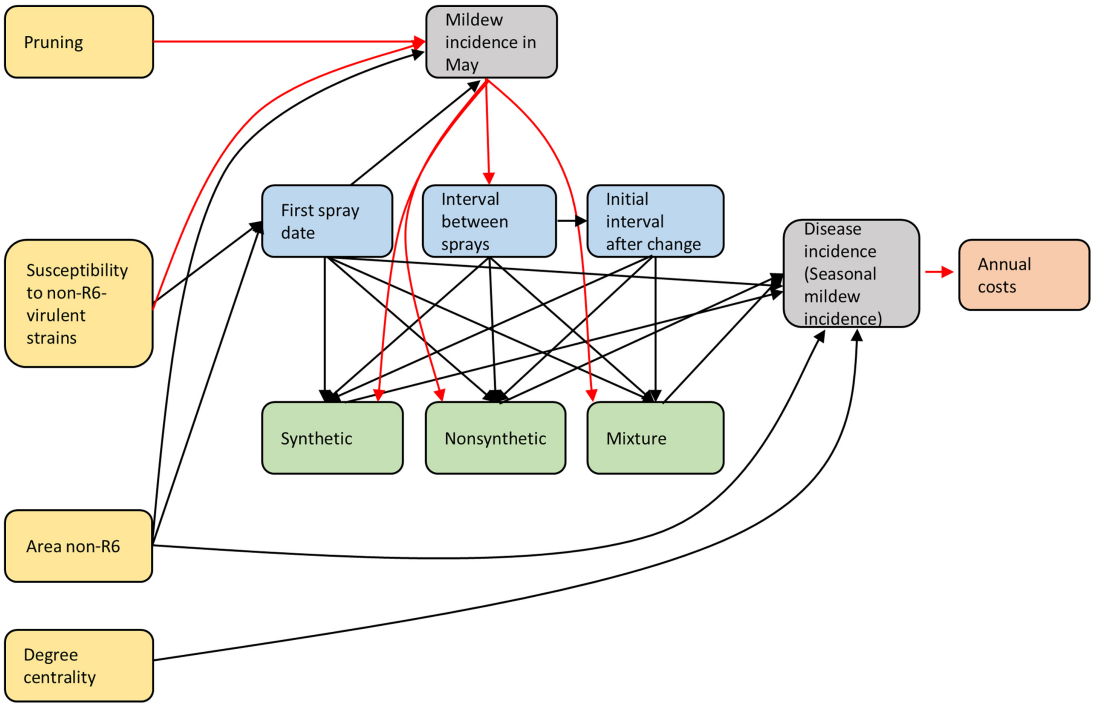

Fig. 1. Directed acyclic graph (DAG) representing causal paths for powdery mildew and management costs. Arrows are edges that indicate postulated causal relationships between variables. Red arrows indicate edges that are assumed always to be present based on prior knowledge. Variables are described in the text and Table 1. Colors indicate exogenous variables (gold), disease incidence (gray), spray date (blue), type of fungicide use (green), and annual costs (orange).

Hop powdery mildew is our motivating system for exploring the impact of host resistance on the use and cost of fungicide inputs. Hop (*Humulus lupulus*) is a long-lived perennial produced for its cones (strobiles). Powdery mildew is one of the costliest diseases of this crop because of the rapid reproductive rate of the causal fungus, the long period during which disease control is necessary, and the potential for damage to both yield and quality (Gent et al. 2014; Royle 1978). Host resistance is a preferred means of disease management yet can be difficult to adopt because market factors other than disease resistance dictate the cultivars that growers produce (Royle 1978). Variation in susceptibility to powdery mildew exists among hop cultivars, often varying in a race-specific manner (Gent et al. 2017; Wolfenbarger et al. 2016). At present, only a few cultivars produced in the United States are completely resistant to all pathogenic races of *P. macularis* (Gent et al. 2017; Laurie et al. 2022; Wolfenbarger et al. 2016). Because of this, disease management often relies on a combination of cultural practices augmented with regular application of fungicides (Gent et al. 2016, 2019b; Hwang et al. 2024; Mahaffee et al. 2003; Royle 1978; Turechek et al. 2001). The aforementioned strategies are applied at some level to all hop cultivars, although heterogeneously depending on disease pressure (Hwang et al. 2024). How growers use fungicides may vary based on susceptibility to powdery mildew, but fungicide use appears complex and may be associated with factors beyond the disease susceptibility of a given cultivar (Hwang et al. 2024, 2025b).

In this research, we developed conditional Gaussian Bayesian network models to understand how host resistance to hop powdery mildew influences the cost of fungicide inputs. We drew upon a large extant data set on the incidence of hop powdery mildew collected from commercial hop yards in Oregon, associated production metadata, and grower pesticide application records (Hwang et al. 2024). We illustrate the utility of Bayesian networks in this context to understand probabilistic relationships between determinants of disease that differentially influence the annual costs of pesticides. We demonstrate through simulation that the annual costs for pesticides scale with exposure to powdery mildew in proportion to cultivar susceptibility, but for only one of the two relevant pathogenic races of the pathogen and with causes that are multifaceted.

Materials and Methods

Data acquisition

A detailed description of methods used for disease assessment is available elsewhere (Gent et al. 2019a). In brief, disease data were collected from a census sample of every hop yard in the eastern extent of the hop production region in Oregon during each year from 2014 to 2017. For the present analyses, we considered data from a total of 101 yards in 2014, 109 in 2015, 73 in 2016, and 119 in 2017. There were 7 to 10 farms included each year. The occurrence of hop plants with powdery mildew was assessed monthly from April to July using a modification of cluster sampling methods presented in Turechek et al. (2001) and Turechek and Mahaffee (2004). We used the monthly data to calculate mean seasonal disease incidence over all months.

Each cultivar in the data set was designated a 0 to 5 ordinal score for its susceptibility to two pathogenic races of *P. macularis*, as described previously (Laurie et al. 2022). These races are defined as having virulence Vb,V3,V5 or a more complex virulence of Vb,V3,V4,V5,V6. These two races were the dominant pathogenic races of *P. macularis* relevant at the time of data collection (Gent et al. 2017; Wolfenbarger et al. 2016). For brevity, we refer to race Vb,V3,V5 as non-R6-virulent and race Vb,V3,V4,V5,V6 as R6-virulent because a key distinction between these races is their ability to cause disease on the widely deployed powdery mildew resistance termed R6 (Henning et al. 2011; Wolfenbarger et al. 2016). We dichotomized cultivars into two susceptibility groups based on the ordinal scores: cultivars scored as ≤ 1 were designated resistant, and those with higher scores were designated susceptible. We subsequently calculated the yearly proportion of the acreage planted to cultivars possessing susceptibility to R6-virulent isolates of *P. macularis* as a correlate of potential pathogen pressure at the landscape level.

The data set also contained information on selected cultural practices that were relevant for disease management. Spring pruning thoroughness was rated using a 1 to 5 ordinal scale (Laurie et al. 2022), where larger values indicate that more green tissue remained after pruning, and therefore, *P. macularis* was more likely to perennate in that yard (Gent et al. 2019b; Laurie et al. 2022). The centrality of each hop yard in an inferred disease-spreading network was expressed using the weighted outward degree centrality (Pedro et al. 2025). In short, weighted outward degree centrality is a measure of inferred disease transmission potential calculated as an edge weight for each source yard to all possible target yards as the product of source yard area, average wind run in the direction of the source yard to the target yard, and a decreasing function of inter-yard distance. We considered the weighted outward degree centrality during the June to July monthly time transition, as this is the period when long-distance dispersal between yards is most probable (Gent et al. 2019a).

We had access to pesticide application records for all yards sampled. We calculated the total number of fungicide active ingredients relevant for powdery mildew management applied by growers in a given year and their estimated costs per hectare as described previously (Hwang et al. 2024). This involved obtaining price quotes in 2021 for the pesticides from three vendors in western Oregon that service hop producers. These price quotes reflected nominal prices for 2014, 2015, 2016, and 2017, where available. We estimated real prices using 2022 as the base by adjusting the nominal prices by the producer price index from the United States Bureau of Labor Statistics and then averaged over all available years to obtain a single real price per unit. We also calculated the mean application interval and the number of synthetic fungicides, nonsynthetic fungicides, or mixtures of the two that were applied per yard, which was considered only through July in order to match the temporal scale of the available covariates. Growers may switch the type of fungicide they apply (Hwang et al. 2025a). When switching occurred, we calculated the time in days from the switch to the subsequent application. For succinctness, we refer to this value as initial interval after change. We calculated simple descriptive statistics to understand the mean, standard errors, and range for the covariates we considered.

Bayesian network learning

We constructed and analyzed separate Bayesian networks for hop cultivar susceptibility to each of R6-virulent and non-R6-virulent strains of *P. macularis.* Our prior belief of the causal network of variables that are most likely to affect the annual cost of fungicide expenditures for powdery mildew is presented (Fig. 1). We specified whitelists of edges that should be present in every DAG based on well-documented determinants of disease or pesticide use. Whitelisted edges are indicated by the red lines and later DAGs. We also specified numerous edges that should never occur in a blacklist (Supplementary Table S1) because of impossibilities based on time of occurrence. For example, we specified that thoroughness of pruning in early spring (Pruning) and cultivar susceptibility to a given strain of *P. macularis* are antecedents of the incidence of plants that develop powdery mildew in May (Mildew May) (Gent et al. 2012, 2019b; Hwang et al. 2024; Turechek et al. 2001). From previous studies and experience, we further specified that the incidence of plants with powdery mildew in May is antecedent to growers' fungicide use pattern. Concretely, disease levels in May have a causal effect on whether and how often growers select a nonsynthetic fungicide, a synthetic fungicide, or a mixture of the two, as well as their mean reapplication interval (Hwang et al. 2025a, b). We also specified that mean seasonal disease incidence (disease

incidence) has a direct causal relationship with annual costs (Hwang et al. 2024).

From the computational programming perspective, our approach to constructing, parameterizing, and evaluating the resulting Bayesian networks was as follows:

1. Draw a knowledge-based structure of the hypothesized causal relationships between variables in available data based on our expert knowledge and literature.
2. Apply a structure-learning algorithm to learn the structure of networks from the data and constraints imposed in step 1 to generate bootstrap samples.
3. Perform model averaging on the bootstrapped network structures to obtain a single network model.
4. Compare the resulting network structure with the knowledge-based structure in step 1.
5. Incorporate further knowledge-based constraints into the algorithm and repeat steps 2 to 4 to obtain a final network structure consistent with both expert knowledge and the data.
6. Select the final network model from step 5 and learn the parameters of the edges.
7. Use the resulting Bayesian network for prediction, causal inference, and simulation of interventions.

We used the R package 'bnlearn' (Liman Harou et al. 2021) to fit a conditional Gaussian Bayesian network model to learn the network structure and then learn the parameters. An important assumption of a conditional Gaussian Bayesian network model is the normality of continuous variables. We explored various transformations of the quantitative variables we considered (Table 1) before settling on a square root transformation because some variables had many zeros. However, even after transformation, some variables we considered were not normally distributed (Fig. 2). Therefore, we estimated the parameters for nonnormalized data using a zero-inflated or a generalized linear model approach appropriate for the form of the data.

We applied a score-based method to learn the structure of the networks. During each of 200 iterations, the algorithm added, deleted, or reversed every edge and scored the resulting network structure (Adhitama and Saputro 2022). The algorithm stopped when the chosen score could not be further improved. We specifically used the hill climbing algorithm with scoring based on the Bayesian information criterion. Constraint-based learning methods are also available that are more computationally efficient and extend the network size for structure learning (Natori et al. 2015). We also applied constraint-based learning as a robustness check on the learned structure of the networks. Model averaging by bootstrapping is often conducted to reduce model uncertainty caused by random noise in a data set. We utilized the boot.strength() function in bnlearn to find the threshold for whether an edge should exist in the final (averaged) network. The averaged networks contain the edges whose strength is greater than the threshold attribute of the boot.strength() function. In the present data set, these were 0.5 for susceptibility to non-R6-virulent strains and 0.515 for susceptibility to R6-virulent strains. From the 200 bootstrapped replications drawn with the boot.strength() function, the relative frequency of each edge was calculated, and only edges that exceeded the fixed threshold were retained. We checked the number of true positives (edges that appeared in two directions) and false positives/negatives (edges that appeared in only one of the two directions) against the blacklists we constructed based on our knowledge of the biological system. Network edges that were learned but prohibited by the blacklist (i.e., false negatives) were not included.

Having learned the network structure, we then learned the parameters on the network structure. Because we used both continuous variables and discrete variables, we estimated the Bayesian network as a conditional Gaussian Bayesian network. We estimated the parameters of the network using their maximum likelihood estimate, yielding a respective local distribution in a classic linear regression for each continuous daughter node against each of its parent nodes when the daughter node was normally distributed. When daughter nodes were not Gaussian, we fit a zero-inflated model or a generalized linear model appropriate for the form of the data. We used the Poisson function for count data if there was no evidence of overdispersion. For count variables with zero inflation, we fit a zero-inflated Poisson regression. Zero-inflated count models are mixture models with two components: (i) a point mass at zero and (ii) a Poisson distribution. Unlike the standard Poisson distribution, zero-inflated models have not one but two parameters. The first parameter is a probability π that the value was 0, which can be considered a Bernoulli distribution. The second parameter has a probability of $1 - \pi$ that the value was non-zero, in which

TABLE 1. Summary statistics and description of variables considered in conditional Gaussian Bayesian networks

Variable[a]	Mean	SD	Min.	Max.	Description
Pruning	2.9	1.45	1	5	Thoroughness of spring pruning; 1 to 5 ordinal scale
Susceptibility to R6-virulent strains	2.85	1.31	0	5	Susceptibility to pathogenic races of *Podosphaera macularis* with virulence Vb,V3,V4,V5,V6 (R6-virulent); 0 to 5 ordinal scale
Susceptibility to non-R6-virulent strains	1.77	1.53	0	4	Susceptibility to pathogenic races of *P. macularis* with virulence Vb,V3,V5 (non-R6-virulent); 0 to 5 ordinal scale
Area R6[a]	0.65	0.24	0.2	0.8	Yearly proportion of hop acreage planted to cultivars susceptible to R6-virulent strains
Area non-R6[a]	0.5	0.04	0.44	0.56	Yearly proportion of hop acreage planted to cultivars susceptible to non-R6-virulent strains
Mildew may	0.01	0.03	0	0.38	Incidence of plants with powdery mildew in May; continuous variable
First spray date	119.3	16.0	64	166	Day of year of the first fungicide application
Interval	16.4	10.4	1	61	Mean interval in days between fungicide applications
Initial interval after change	6.81	10.6	0	61	Interval in days of first spray after switching fungicide type
Disease incidence	0.10	0.22	0	1	Seasonal mean incidence of plants with powdery mildew during April to July; continuous variable
Degree centrality	0.5	0.34	0.03	2.08	Degree centrality of a yard, with edges weighted by wind run, area of the yard, and distance to other yards
Annual costs	538.63	308.66	6.36	1,733.36	Annual costs of pesticides for suppression of powdery mildew; expressed in real 2022 dollars per hectare
Synthetic	0.89	0.95	0	5	Number of nonsynthetic fungicides used per hop yard per year
Nonsynthetic	3.44	1.87	0	10	Number of synthetic fungicides used per hop yard per year
Mixture	0.91	1.13	0	4	Number of applications made using a mixture of at least one synthetic and nonsynthetic fungicide per hop yard per year
N		1,174			

[a] Ordinal values of disease susceptibility were dichotomized as 1 for cultivars with scores $\leq$1 and 0, otherwise to determine the proportion of hop yards that were planted to cultivars resistant or susceptible to a given strain of *P. macularis*.

case it follows a Poisson distribution with mean and variance λ. A generalized linear model was fit for proportion data or count data using the binomial or Poisson distribution through the glm function in R (Dunn and Smyth 2018). As a robustness check on the sign and strength of the learned parameters, we conducted conditional independence tests using constraint-based learning methods from the learned DAG. This test assumes that the main variables are independent, given covariates. We tested for independence between nodes with the ci.test() function in the bnlearn package.

Prediction and simulation

A Bayesian network model's output after parameter learning is easy to interpret, allowing the outcomes to be used for experimental analysis and for causal analysis. We extracted the predicted values from the fitted model of the conditional Gaussian distribution using the predict function in bnlearn. We performed queries of the network by Monte Carlo inference utilizing the cpdist function in bnlearn for simulation. A total of 10,000 random draws from the data set were made to compute the posterior probability distribution for each node in the network, which, for our interest, was ordinal scores of the susceptibility to non-R6-virulent or R6-virulent strains of *P. macularis*. We plotted the resulting estimates of annual costs of pesticides versus mean seasonal disease incidence and plotted fit linear regressions by cultivar susceptibility to summarize the data. In some simulations and predictions, estimated disease incidence was <0 or >1. We only report data bounded between 0 and 1, inclusive.

Results

Simple descriptive summaries for the variables we considered in the Bayesian networks are provided (Table 1). Of note to this analysis, powdery mildew was observed on a 0.10 proportion of the plants in the yards in the data set. Cultivar susceptibility to powdery mildew caused by R6-virulent strains ranged from 0 to 5 with a mean of 2.85, whereas susceptibility to non-R6-virulent strains was less, ranging from 0 to 4, with a mean of 1.77. The overall higher susceptibility to R6-virulent strains simply indicates that race-specific resistance is present.

The average annual cost for fungicides was $538.63/ha but had a wide range, from $6.36 to $1,733.36/ha, among individual hop yards. There was also large variation between most of the other variates in the data set that could influence the annual cost of fungicides, such as the thoroughness of spring pruning (range 1 to 5), the

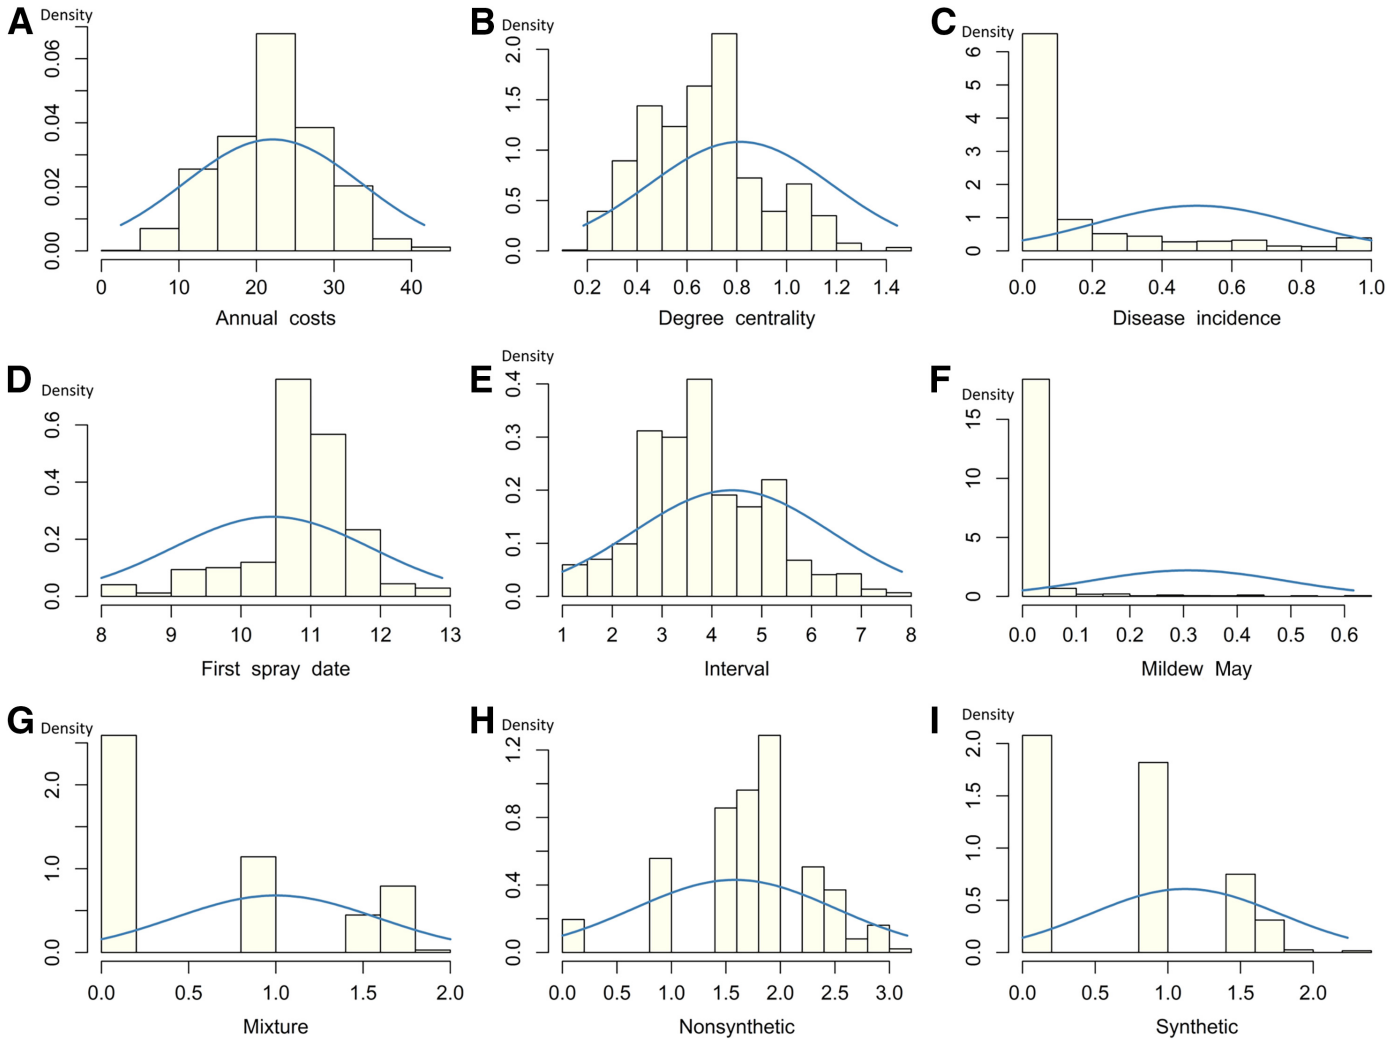

Fig. 2. Distribution of continuous variables considered in the Bayesian networks after square root transformation. Blue lines are the fit of a Gaussian distribution for reference in comparison to the histograms. Variables represented are the **A,** annual costs of fungicides; **B,** weighted outward degree centrality; **C,** mean seasonal incidence of plants with powdery mildew; **D,** day of year of the first fungicide application; **E,** mean fungicide application interval; **F,** incidence of plants with powdery mildew in May; **G,** number of fungicide applications made using a mixture of nonsynthetic and synthetic fungicides; **H,** number of nonsynthetic fungicide applications made; and **I,** number of synthetic fungicide applications made.

incidence of plants with powdery mildew in May (range 0 to 0.38), and the seasonal mean incidence of powdery mildew (range 0 to 1). A descriptive summary of potential covariates related to cultivar susceptibility to powdery mildew is provided (Supplementary Table S2).

Bayesian network learning

The average Markov blanket size and neighbor size were, respectively, 8.92 and 5.38 in the network for susceptibility to R6-virulent strains and 8.77 and 5.54, respectively, in the network for susceptibility to non-R6-virulent strains. The DAGs for both had an overall similar structure and strength of the edges, as there were 13 nodes in both networks and 35 or 36 directed edges (Table 2). A few differences were detected in the DAGs (Fig. 3). In the DAG for susceptibility to R6-virulent strains, cultivar susceptibility was not directly linked by an edge to the degree centrality of a yard, whereas cultivar susceptibility to non-R6-virulent strains was linked by a weak edge (edge strength 0.745; Table 2). In the DAG for susceptibility to non-R6-virulent strains, spring pruning was directly linked by a strong edge (edge strength 0.94) to mean seasonal disease incidence, whereas this edge was absent in the other network. The area of hop yards planted to cultivars susceptible to non-R6-virulent strains of *P. macularis* was not connected to any other nodes.

For both networks of cultivars, the annual cost of fungicide inputs had complex interrelationships with multiple variables related to disease levels, how often growers applied fungicides, the specific type(s) of fungicides they applied, and the degree centrality of yards. In general, the Bayesian networks depicted cultivar susceptibility as influencing disease levels in early spring and how often and what fungicides growers applied, which then determined seasonal disease levels and overall management strategies. Specifically, the annual cost of fungicides was directly related to the mean fungicide application interval, the degree centrality of a yard, the number of synthetic and nonsynthetic fungicide applications, and mean seasonal disease incidence (Fig. 3). Therefore, annual costs depended on not only the number of applications made but the specific types of fungicides used. We also found differences between the two networks in the determinants of variables with a direct relationship to annual costs of fungicides. As an example, in the network for R6-virulent strains, disease incidence was the product of cultivar susceptibility, the incidence of plants with disease in May, the initial interval after changing the fungicide type, the degree centrality of a yard, and the number of fungicide applications containing only a nonsynthetic fungicide or a mixture of fungicide types. In contrast, in the network based on susceptibility to non-R6-virulent strains, disease incidence also depended on spring pruning thoroughness but not the number of nonsynthetic fungicide applications made (Fig. 3).

We estimated the parameters of the learned DAG assuming a conditional Gaussian distribution, zero-inflated models, or generalized linear models depending on the nature of the data (Tables 3 and 4). The coefficients of the conditional Gaussian distribution were estimated from score-based learning. The predicted value of a given variable is interpreted as a regression with parameter estimates for the parent nodes in the path of a given daughter node. Concretely, when considering cultivar susceptibility to R6-virulent strains, the

TABLE 2. Edge strength of the conditional Gaussian Bayesian networks from structure learning

		Strength[a]	
Parent node	Daughter node	Susceptibility to R6-virulent strains	Susceptibility to non-R6-virulent strains
Pruning	Mildew May	1	1
	First spray date	1	1
	Disease incidence	–	0.94
	Synthetic	1	0.91
	Nonsynthetic	1	0.97
Susceptibility	Mildew May	1	1
	Disease incidence	1	1
	Degree centrality	–	0.745
	Synthetic	0.885	0.925
	Mixture	1	0.995
Area R6	Mixture	0.93	–
Mildew May	Interval	1	1
	Disease incidence	1	1
	Degree centrality	1	1
	Synthetic	1	1
	Nonsynthetic	1	1
	Mixture	1	1
First spray date	Interval	0.995	0.98
	Synthetic	0.88	0.71
Interval	Annual costs	0.975	0.955
Initial interval after change	Interval	1	1
	Disease incidence	0.51	0.555
	Synthetic	0.595	0.545
	Nonsynthetic	0.88	0.885
Disease incidence	Annual costs	1	1
	Disease incidence	1	1
	Annual costs	0.935	0.93
	Nonsynthetic	1	1
Synthetic	Interval	0.985	0.98
	Degree centrality	0.73	0.825
	Annual costs	1	1
Nonsynthetic	Interval	1	0.995
	Disease incidence	0.785	–
	Annual costs	0.995	0.985
Mixture	Interval	1	1
	Disease incidence	1	1
	Degree centrality	0.515	0.57

[a] Strength is the probability of including an edge (regardless of direction) between two nodes. Variable descriptions are given in Table 1.

annual cost of fungicides increased as the mean application interval decreased (coefficient −0.64, [95% confidence interval −12.38 to 11.10]), when the degree centrality of the yard increased (3.79 [−7.95 to 15.53]), when the number of applications of synthetic fungicides increased (1.88 [−9.86 to 13.62]), when the number of applications of nonsynthetic fungicides decreased (−0.71 [−12.45 to 11.03]), and when the mean seasonal disease incidence increased (7.0 [−4.74 to 18.74]). For count data or proportion data under this group of cultivars, coefficients of zero-inflated models and generalized linear models were estimated from the Poisson or binomial function. Growers were more likely to have had disease in May (log odds 5.20) and had a more than doubling of mean seasonal disease incidence (log odds 12.93) for each unit increase in cultivar susceptibility to powdery mildew. Growers had mixed results regarding fungicide use. Although they were less likely to use synthetic fungicides (log odds 0.77) as cultivar susceptibility increased, they had a 2.58-fold increase in the log odds of using mixture fungicides. The coefficient results of the conditional independence test

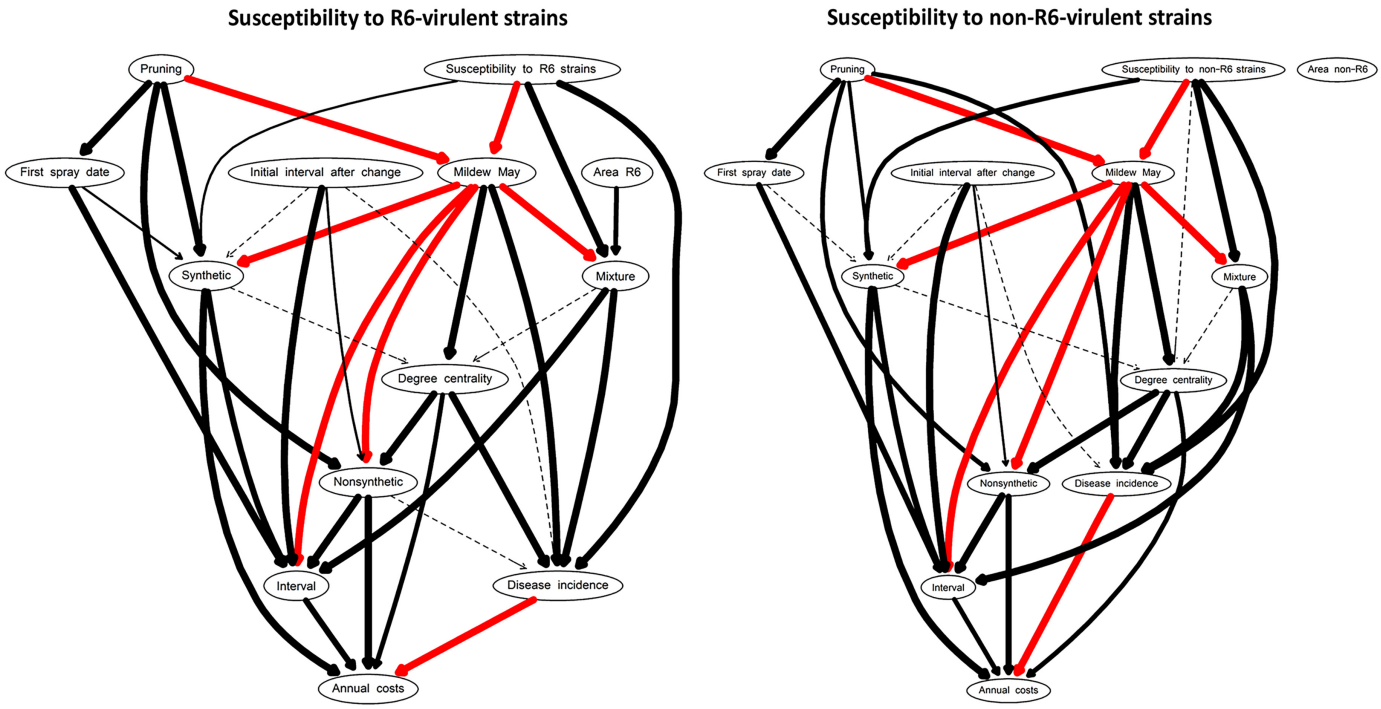

Fig. 3. Directed acyclic graph (DAG) from Bayesian network structure learning of the annual costs of fungicides applied to hop cultivars based on their susceptibility to R6-virulent strains (left) or non-R6-virulent strains of *Podosphaera macularis* (right). Red lines are edges in the whitelist that were always included in the network; black lines were learned during model fitting. Strength of the relationship is indicated by the width of the edge. Bold lines indicated a probability = 1, thin lines indicate probability between 0.85 and 1, and dashed lines indicate probability between the threshold and 0.85. Thresholds for an edge to appear are 0.51 for the network evaluating susceptibility to R6-virulent strains and 0.57 for the network evaluating susceptibility to non-R6-virulent strains.

TABLE 3. Parameters for annual cost of fungicides applied to hop cultivars with varying susceptibility to R6-virulent strains of *Podosphaera macularis*[a]

Variable	Mildew May (GLM)	First spray date (GD)	Interval (GLM)	Degree centrality (GLM)	Synthetic (GLM)	Nonsynthetic (Zero)		Mixture (Zero)		Disease incidence (GLM)	Annual costs (GD)
						Count	Zero-inflated	Count	Zero-inflated		
Intercept	−9.04 (−13.21, −5.79)	11.51 (10.06, 12.96)	2.75 (2.29, 3.21)	0.64 (0.61, 0.66)	−2.13 (−3.27, −1.02)	0.26 (0.03, 0.50)	−63.80 (−374.11, 246.50)	−1.14 (−1.6, −0.69)	−10.75 (–)	−8.99 (−10.75, −7.41)	21.79 (10.01, 33.53)
Pruning	1.19 (0.12, 2.40)	−0.37 (−1.82, 1.08)	–	–	0.38 (0.22, 0.54)	−0.13 (−0.25, −0.03)	30.22 (−108.55, 169.01)	–	–	–	–
Susceptibility to R6-virulent strains	1.65 (0.32, 3.46)	–	–	–	−0.25 (−0.37, −0.13)	–	–	0.95 (0.57, 1.33)	−7.68 (−7,439.43, 7,424.05)	2.56 (1.83, 3.40)	–
Area R6	–	–	–	–	–	–	–	−1.48 (−2.55, −0.42)	−2.21 (–)	–	–
Mildew May	–	–	0.11 (−0.31, 0.52)	0.70 (0.51, 0.88)	0.52 (−0.44, 1.39)	−0.36 (−1.07, 0.34)	−20.88 (−297.40, 255.63)	1.64 (0.84, 2.44)	−8.44 (−546,749.7, 546,732.8)	7.50 (5.03, 10.19)	–
First spray date	–	–	−0.09 (−0.12, −0.05)	–	0.12 (0.03, 0.22)	–	–	–	–	–	–
Interval	–	–	–	–	–	–	–	–	–	–	−0.64 (−12.38, 11.10)
Initial interval after change	–	–	0.04 (0.03, 0.06)	–	0.07 (0.04, 0.10)	−0.01 (−0.04, 0.01)	−15.38 (−10,497.42, 10,466.65)	–	–	0.07 (0, 0.16)	–
Degree centrality	–	–	–	–	–	0.55 (0.35, 0.77)	−12.28 (−20.33, −4.24)	–	–	2.24 (1.36, 3.14)	3.79 (−7.95, 15.53)
Synthetic	–	–	−0.11 (−0.16, −0.05)	0.01 (0, 0.03)	–	–	–	–	–	–	1.88 (−9.86, 13.62)
Nonsynthetic	–	–	−0.18 (−0.24, −0.12)	–	–	–	–	–	–	0.19 (−0.22, 0.63)	−0.71 (−12.45, 11.03)
Mixture	–	–	−0.15 (−0.20, −0.10)	0.00 (−0.01, 0.02)	–	–	–	–	–	0.61 (0.33, 0.88)	–
Disease incidence	–	–	–	–	–	–	–	–	–	–	7.00 (−4.74, 18.74)

[a] Parameter estimates are based on square root-transformed variables, as noted in the text. Gaussian distribution (GD) for normal data, zero-inflated model (Zero), or generalized linear model (GLM) for count data or proportion data. Zero-inflated models have estimates for lambda (count models) and phi (zero-inflated models). Variable descriptions are given in Table 1.

from constraint-based learning had similar directions to the parameters from the score-based learning (Supplementary Tables S3 and S4).

Prediction and simulation

We estimated predicted values based on the fitted model from the conditional Gaussian distribution and made out-of-sample predictions for selected outcome variables that were approximately Gaussian distributed. We found an overlap in the distribution of annual costs of fungicides based on cultivar susceptibility to each of the two strains of *P. macularis* (Fig. 4A and B). Annual costs of fungicides were modestly reduced with increasing cultivar resistance to R6-virulent strains of the fungus, decreasing from a median of $566.05/ha (interquartile range [IQR] $514.33 to $904.60) to $449.65/ha (IQR $357.86 to $530.26) for cultivars scored as highly susceptible (5) to completely resistant (0). Maximum savings in fungicide costs required cultivars to have the highest levels of resistance. We did not find evidence that cultivar susceptibility to non-R6-virulent strains predicted the annual cost of fungicides (4B). The median for annual costs only ranged from $510.38 (IQR $395.34 to $634.58) to $520.61 (IQR $413.54 to $620.97) across the categories of cultivar susceptibility to these strains.

We next explored aspects of growers' fungicide use practices to understand how these were predicted to change in response to cultivar susceptibility. The predicted date of the first fungicide application (Fig. 4C to D) and mean application interval were largely unchanged with increasing host resistance to either strain of *P. macularis* (Fig. 4E to F). Rather, the primary change influencing the annual cost of fungicides was the type of fungicide growers applied, which was related to seasonal disease incidence (Fig. 5). Annual costs for fungicides increased with increasing mean seasonal disease incidence, which in turn was related to whether growers selected a nonsynthetic or synthetic fungicide, or applied a mixture of both types.

We then simulated annual costs for disease management, assuming different levels of cultivar susceptibility to powdery mildew, estimating the causal relationship between child nodes given the specific situation (Fig. 6). In the simulations considering cultivars affected by R6-virulent strains of *P. macularis*, the relationship between seasonal mean disease incidence and annual cost varied depending on cultivar susceptibility to powdery mildew. This effect was due to relatively small changes in fungicide costs in the absence of disease and, more importantly, how fungicide costs increased with mean seasonal disease incidence. The first effect is demonstrated by the generally increasing value of the y-intercept in the regressions at each level of cultivar susceptibility, ranging between $470 and $500 per hectare when susceptibility was 0 versus 5. The second effect is demonstrated by the generally increasing value of the slope of the regressions, with the slope increasing from $130 to $510/ha with each unit increase (i.e., from 0 to 1) in mean seasonal disease incidence as cultivar susceptibility increased.

Perhaps counterintuitively, when we considered cultivar susceptibility to non-R6-virulent strains of *P. macularis*, the intercept and the slope between annual costs and mean seasonal disease incidence were not significantly different with cultivar susceptibility (Fig. 6). Therefore, resistance to only one of the two cognate strains did not have a measurable effect on fungicide expenditures.

Discussion

The sum of this research demonstrates that the relationship between host resistance and the cost of growers' fungicide inputs is complex and multifaceted in the motivating pathosystem. The median predicted value of changing cultivar resistance to R6-virulent strains of the fungus from highly susceptible (5 on the ordinal scale in the present data set) to completely resistant (0 on the ordinal scale) reduces the annual costs of fungicides only modestly from $566.05 to $449.65/ha. These savings are not simply due to fewer applications. Rather, the most important changes are in the types of fungicides applied, shifting from relatively expensive synthetic fungicides or mixtures thereof to lesser expensive nonsynthetic fungicides. This discovery has important implications for assessing the economic value of host resistance and potential strategies for reducing fungicide usage, as we explain below.

The damage potential of crop diseases is clear and, even with active management, may be substantial (Oerke 2006). Considering this, the economic value of disease resistance is sometimes taken as a given, even though a precise estimate of the economic value is not simple to derive (Hogenboom 1993). Prospective analyses and randomized controlled trials often demonstrate large reductions in fungicide use in cultivars with improved disease resistance (Fuller et al. 2014). Eisenmann et al. (2023), for example, demonstrated that pesticide use could be reduced 60 to 90% in grapevine with cultivars with improved resistance to both grape powdery mildew

TABLE 4. Parameters learned for annual cost of fungicides applied to hop cultivars with varying susceptibility to non-R6-virulent strains of *Podosphaera macularis*[a]

Variable	Mildew May (GLM)	First spray date (GD)	Interval (GLM)	Degree centrality (GLM)	Synthetic (Zero)		Nonsynthetic (Zero)		Mixture (Zero)		Disease incidence (GLM)	Annual costs (GD)
					Count	Zero-inflated	Count	Zero-inflated	Count	Zero-inflated		
Intercept	−6.04 (−8.36, −4.13)	11.51 (10.06, 12.96)	2.75 (2.29, 3.21)	0.68 (0.65, 0.71)	−2.76 (−3.97, −1.56)	−156.98 (−4,447, 4,133.03)	0.26 (0.03, 0.5)	−63.80 (−374.11, 246.5)	−0.66 (−0.80, −0.54)	−12.13 (−112.56, 88.29)	−4.53 (−5.56, −3.55)	21.79 (10.05, 33.53)
Pruning	1.20 (0.10, 2.43)	−0.37 (−1.82, 1.08)	–	–	0.24 (0.06, 0.43)	−127.72 (−4,408.55, 4,153.09)	−0.13 (−0.25, −0.03)	30.22 (−108.55, 169.01)	–	–	0.60 (0.16, 1.05)	–
Susceptibility to non-R6-virulent strains	−0.10 (−0.67, 0.49)	–	–	−0.04 (−0.06, −0.02)	0.17 (0.08, 0.28)	−0.54 (−2.84, 1.75)	–	–	−0.09 (−0.19, 0.01)	−3.11 (−774.96, 768.73)	−0.77 (−1.01, −0.53)	–
Area non-R6	–	–	–	–	–	–	–	–	–	–	–	–
Mildew May	–	–	0.11 (−0.31, 0.52)	0.71 (0.52, 0.89)	0.67 (−0.31, 1.67)	179.36 (−4,538.27, 4,897)	−0.36 (−1.07, 0.34)	−20.88 (−297.4, 255.63)	1.97 (1.19, 2.76)	−8.36 (−2,620.82, 2,604.08)	8.09 (5.61, 10.85)	–
First spray date	–	–	−0.09 (−0.12, −0.05)	–	0.13 (0.04, 0.24)	24.45 (5.14, 43.77)	–	–	–	–	–	–
Interval	–	–	–	–	–	–	–	–	–	–	–	−0.64 (−12.38, 11.1)
Initial interval after change	–	–	0.04 (0.03, 0.06)	–	0.08 (0.05, 0.12)	−12.27 (−33,835.89, 33,811.34)	−0.01 (−0.04, 0.01)	−15.38 (−1,497.42, 10,466.65)	–	–	0.07 (−0.01, 0.15)	–
Degree centrality	–	–	–	–	–	–	0.55 (0.35, 0.77)	−12.28 (−20.33, −4.24)	–	–	2.28 (1.51, 3.06)	3.79 (−7.95, 15.53)
Synthetic	–	–	−0.11 (−0.16, −0.05)	0.02 (0, 0.04)	–	–	–	–	–	–	–	1.88 (−9.86, 13.62)
Nonsynthetic	–	–	−0.18 (−0.24, −0.12)	–	–	–	–	–	–	–	–	−0.71 (−12.45, 11.03)
Mixture	–	–	−0.15 (−0.20, −0.10)	0.001 (−0.01, 0.02)	–	–	–	–	–	–	0.78 (0.52, 1.04)	–
Disease incidence	–	–	–	–	–	–	–	–	–	–	–	7 (−4.74, 18.74)

[a] Parameter estimates are based on the square root-transformed variables, as noted in the text. Gaussian distribution (GD) for normal data, zero-inflated model (Zero), or generalized linear model (GLM) for count data or proportion data. Zero-inflated models have estimates for lambda (count models) and phi (zero-inflated models). Variable descriptions are given in Table 1.

(*Erysiphe necator*) and downy mildew (*Plasmopara viticola*) depending on the level of resistance and seasonal effects. An analysis of adopters and nonadopters of dry bean cultivars with disease resistance to multiple diseases found an 11 to 29% reduction in unit variable costs, mostly associated with fungicide inputs, with the variation depending on disease pressure (Mooney et al. 2022). It is therefore striking in the present study that annual costs of fungicides with powdery mildew activity were still relatively large even with cultivars with high levels of resistance and where powdery mildew was never detected.

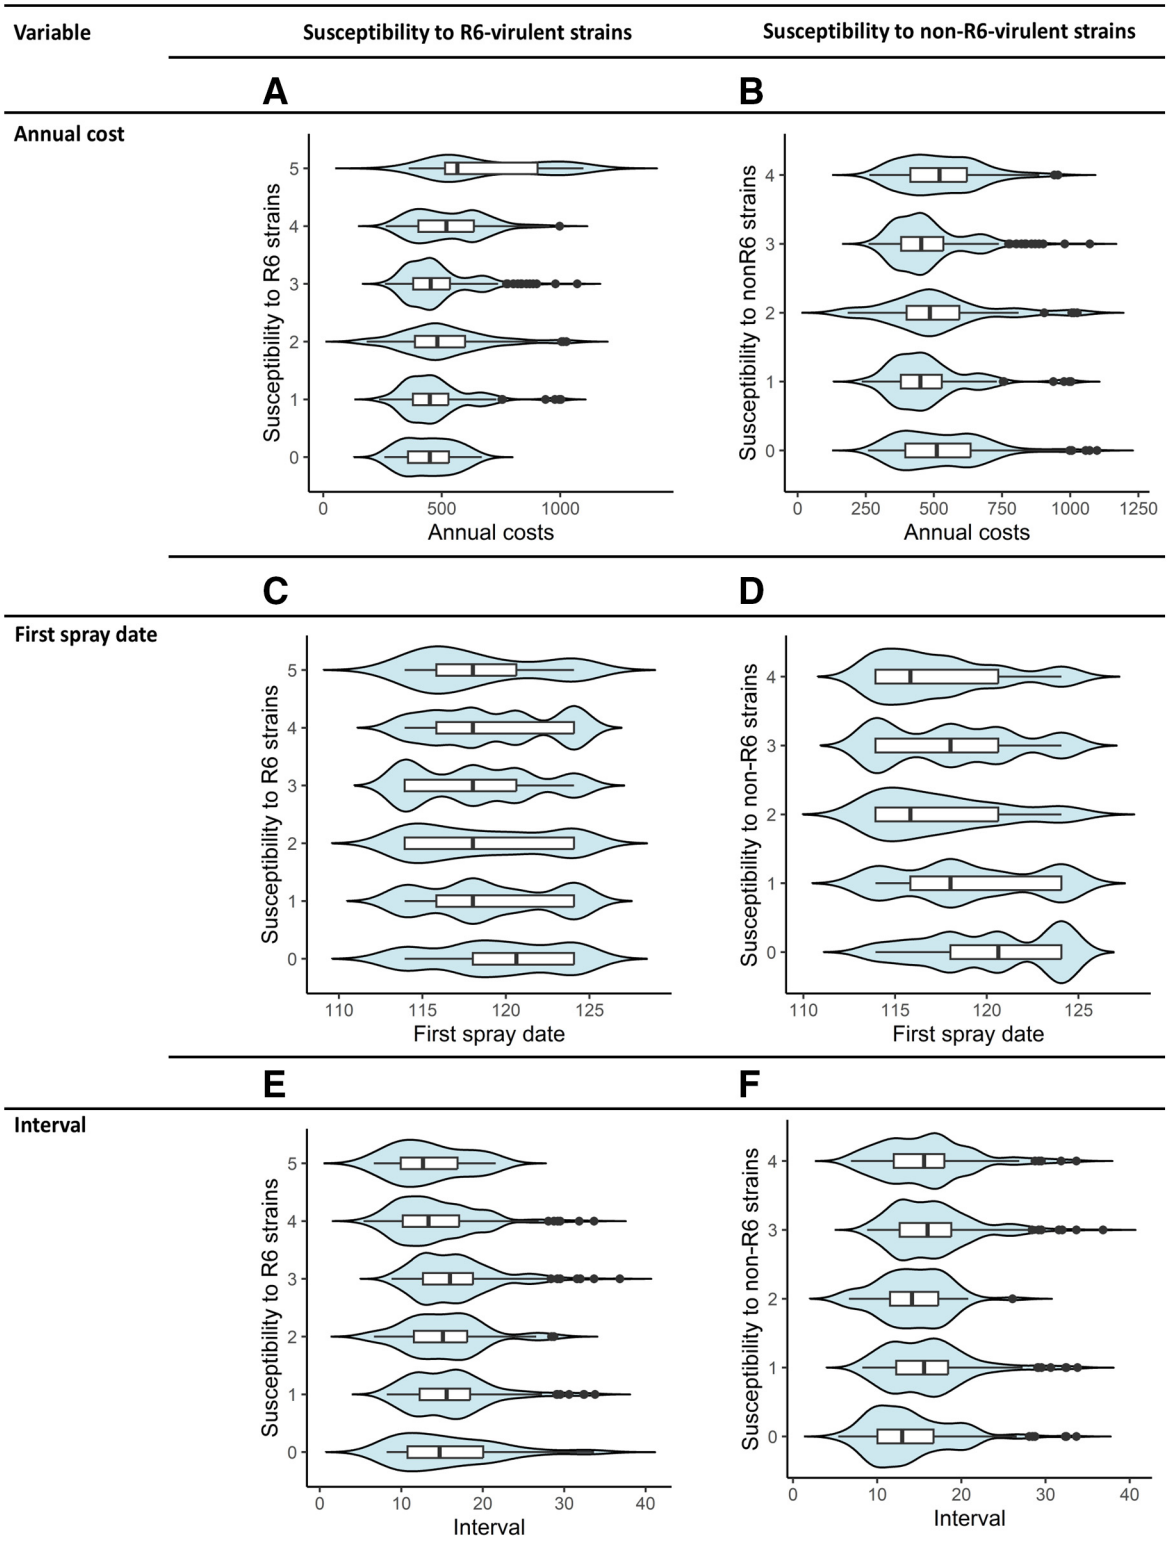

Fig. 4. Violin plots showing the distribution of predicted values for **A and B,** annual costs of fungicide expenditures; **C and D,** the first spray date; and **E and F,** mean fungicide application interval by hop cultivar susceptibility to (A, C, and, E) R6-virulent or (B, D, and F) non-R6-virulent strains of *Podosphaera macularis* at each ordinal value of cultivar susceptibility.

Our previous analysis of the same data set found that hop growers' response to powdery mildew could be modeled as an exposure-response function, with annual pesticide costs scaling linearly with mean seasonal disease incidence. However, these changes are nuanced and heterogeneous depending on cultivar susceptibility to the disease (Hwang et al. 2024). In that study, for cultivars rated as having a powdery mildew susceptibility of 1 to 4, the annual costs growers incurred in response to powdery mildew were generally similar, but the number of active constituents applied tended to decrease with increasing susceptibility. In other words, growers responded to powdery mildew by applying more expensive fungicides, but not necessarily more total active ingredients, in susceptible cultivars.

The present study confirms and extends these findings with several key insights. First, through the simulation studies, we show that how the annual cost of fungicides scales with seasonal mean disease incidence depends on the specific strain of *P. macularis* and cultivar susceptibility. When only considering disease phenotype to non-R6-virulent strains of the fungus, there is not a meaningful difference in fungicide costs in the absence or presence of powdery mildew. This is likely because cultivars that possess the R6-form of powdery mildew resistance are completely resistant (ordinal value 0) to non-R6-virulent strains but potentially quite susceptible to R6-virulent strains. A specific example is the cultivar 'Nugget' (Wolfenbarger et al. 2016). Resistance to a single strain with simple virulence when multiple strains exist in the pathogen population therefore does not meaningfully reduce fungicide expenditures. When considering R6-virulent strains with more complex virulence, we found that growers had a relatively similar baseline level of fungicide costs independent of cultivar susceptibility to powdery mildew. Their total annual cost of fungicides differed between cultivar susceptibility classes because of the steeper rate of change in costs with increasing mean seasonal disease incidence in more susceptible cultivars (Fig. 6).

Second, the conditional Gaussian Bayesian networks we fit indicate that the reasons annual costs vary are multifaceted and are attributable to the different types of fungicides growers use, not simply growers' total use of fungicides (as measured by mean application interval). This is consistent with a recent analysis of switching behavior by hop producers that found that producers switch from nonsynthetic fungicides to more potent and more expensive synthetic fungicides or mixtures of fungicide types in response to observed powdery mildew and crop developmental stage (Hwang et al. 2025a). The present study indicates that switching behavior, as inferred from the cost of fungicide inputs, likely occurs across cultivars, even those with relatively high levels of disease resistance.

An obvious question is, why do growers continue to apply fungicides of any type to cultivars with high levels of resistance to powdery mildew?

One plausible answer is aversion to economic risk. Pesticides often are applied because of their reliability and therefore reduction of economic risk (Bakker et al. 2021; Wearing 1988), particularly in high-value crops with highly asymmetrical cost ratios for false-negative and false-positive treatment decisions (Gent et al. 2013; McRoberts et al. 2011). If driven by risk aversion, there is a role for extension education, as better information and advice from trusted experts are factors that may reduce unnecessary pesticide use (Bakker et al. 2021; Sherman and Gent 2014; Wearing 1988).

Another plausible explanation is that growers apply fungicides targeting multiple diseases. Our motivating pathosystem was hop powdery mildew, but this is only one of two common foliar diseases of hop. Downy mildew (caused by *Pseudoperonospora humuli*) is ubiquitous in hop production regions in Oregon and also routinely managed with fungicides (Purayannur et al. 2021; Sherman and Gent 2014), some that may be specific to oomycete pathogens but others that have overlapping activity on powdery mildew diseases (Hwang et al. 2024). Our specific interest and available data were for resistance to powdery mildew, which is distinct from the complex, multigenic resistance to downy mildew (Henning et al. 2016; Olatoye et al. 2023; Purayannur et al. 2021). Evaluation of cultivars with multiple-disease resistance is warranted in future studies to understand how growers' fungicide use practices changes with single versus multiple disease resistance.

Beyond the idiosyncratic elements of the motivating pathosystem, we highlight the attractiveness of a Bayesian network framework for causal inference. Bayesian networks can simultaneously allow for exploring multiple characteristics, as drawing on direct dependencies means that each characteristic is explained by a local distribution depending only on the parent nodes directly linked to a child node (Scutari et al. 2017). The present analysis also points out a potential limitation that we expect will be common to other

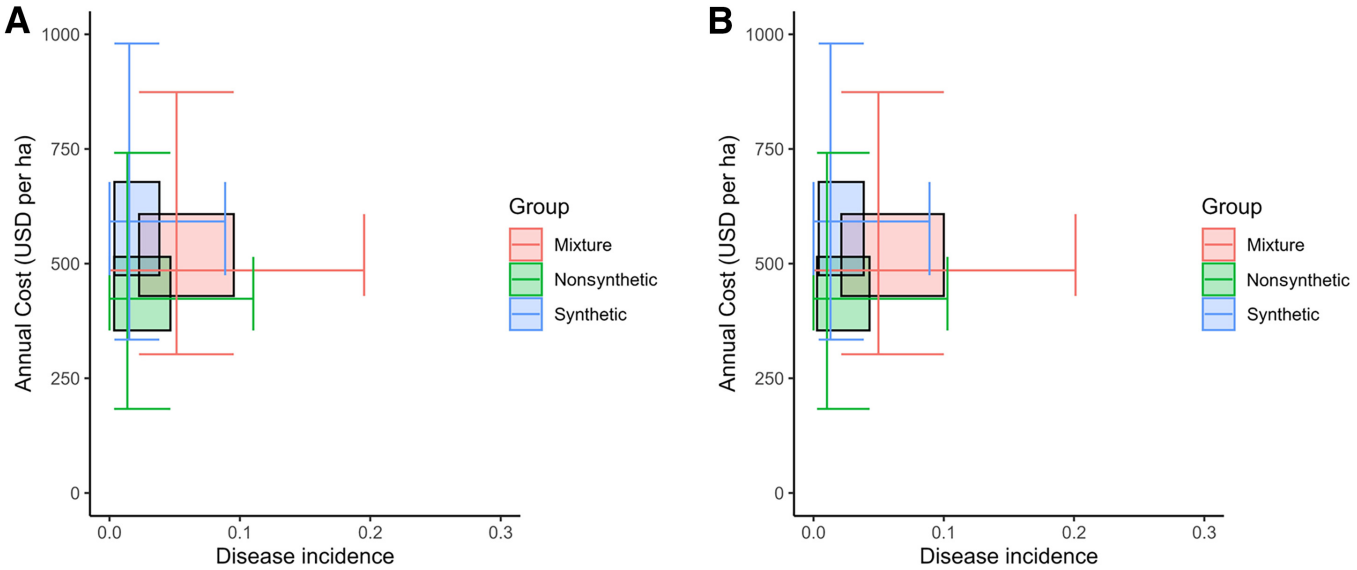

Fig. 5. Double boxplot expressing the distribution from predicted values of the seasonal mean incidence of hop powdery mildew (as a proportion) and the associated annual costs by fungicide type associated with cultivars susceptible to **A,** R6-virulent or **B,** non-R6-virulent strains of *Podosphaera macularis*. Note that for both groups of cultivars, growers are predicted to preferentially apply mixtures of nonsynthetic and synthetic fungicides as disease incidence increases.

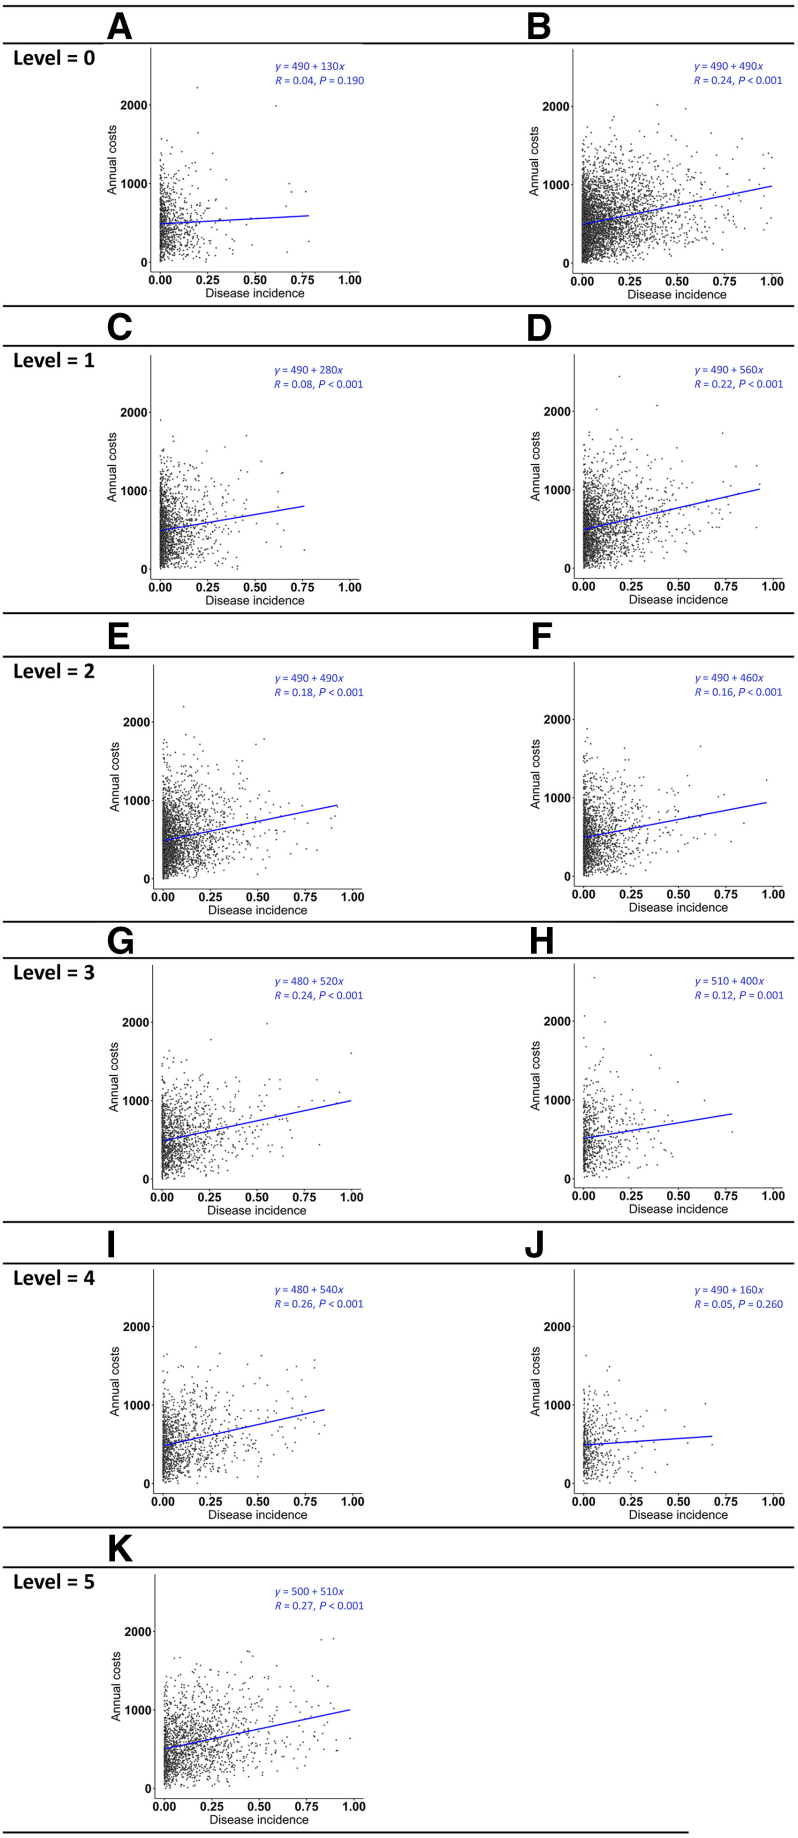

Fig. 6. Simulations for conditional distribution of mean seasonal disease incidence and annual costs of fungicides for hop powdery mildew control at each level of cultivar susceptibility to **A, C, E, G, I, and K,** R6-virulent or **B, D, F, H, and J,** non-R6-virulent strains of *Podosphaera macularis*. Blue lines indicate a linear regression fit to the simulated values. The coefficient of determination and *P* value for the slope are given below the equation for the linear regression. Note the increasing value of the slope in the regressions for host susceptibility to R6-virulent strains in panels A, C, E, G, I, and K.

complex plant disease data sets, the assumption of conditional Gaussian multivariate normality. Plant disease incidence ranges from 0 to 1 and often is positively skewed due to many zero observations. Transformation of such data may yield a Gaussian distribution that is only marginally normally distributed. A zero-inflated model or a generalized linear model might be more appropriate for prediction and simulation of such data after the structure of the network is learned. With a large data set, data that are only marginally normally distributed may not be a serious concern because sample means drawn from a large population should be normally distributed according to the central limit theorem, even though some variables are not strictly Gaussian. We also note that other transformations may be appropriate depending on the observed distribution of the data. Continuous data are often dichotomized when fitting a Bayesian network, although dichotomization causes loss of information and may influence accuracy and bias (Flores et al. 2011).

Calls for greater accountability and justification of public expenditures in science show no sign of ebbing (McRoberts et al. 2024), and it is important to document the impact of plant breeding investments. Our application of Bayesian networks provides a coherent framework for understanding the multifaceted interaction of these factors that influence the eventual outcome of interest and their causal effects. Our analyses indicate that for a high-value crop, resistance to one disease results in complex and unexpected changes in growers' fungicide use patterns that may not be obvious in simplified randomized controlled trials. Grower aversion to risk and the need to control other diseases may lead to savings in fungicide inputs that may be far less than expected. This points to the critical need for extension education efforts to communicate and demonstrate the efficacy of disease resistance and appropriately match fungicide inputs to the actual disease threat. Furthermore, our results suggest that, in practice, breeding for multiple disease resistance may be needed to reduce fungicide inputs, a hypothesis that should be explored in this and other pathosystems.

Acknowledgments

We thank the many individuals that provided technical support for this project and the participating grower cooperators that allowed access to their farms and production records.

Literature Cited

Adhitama, R. P., and Saputro, D. R. S. 2022. Hill climbing algorithm for Bayesian network structure. In: AIP Conference Proceedings. AIP Publishing, Surakarta, Indonesia.

Bakker, L., Sok, J., van der Werf, W., and Bianchi, F. J. J. A. 2021. Kicking the habit: What makes and breaks farmers' intentions to reduce pesticide use? Ecol. Econ. 180:106868.

Dunn, P. K., and Smyth, G. K. 2018. Generalized Linear Models with Examples in R. Springer, New York.

Eisenmann, B., Wingerter, C., Dressler, M., Freund, C., Kortekamp, A., and Bogs, J. 2023. Fungicide-saving potential and economic advantages of fungus-resistant grapevine cultivars. Plants 12:3120.

Flores, M. J., Gámez, J. A., Martínez, A. M., and Puerta, J. M. 2011. Handling numeric attributes when comparing Bayesian network classifiers: Does the discretization method matter? Appl. Intell. 34:372-385.

Fuller, K. B., Alston, J. M., and Sambucci, O. S. 2014. The value of powdery mildew resistance in grapes: Evidence from California. Wine Econ. Policy 3:90-107.

Gent, D. H., Bhattacharyya, S., and Ruiz, T. 2019a. Prediction of spread and regional development of hop powdery mildew: A network analysis. Phytopathology 109:1392-1403.

Gent, D. H., Grove, G. G., Nelson, M. E., Wolfenbarger, S. N., and Woods, J. L. 2014. Crop damage caused by powdery mildew on hop and its relationship to late season management. Plant Pathol. 63:625-639.

Gent, D. H., Mahaffee, W. F., McRoberts, N., and Pfender, W. F. 2013. The use and role of predictive systems in disease management. Annu. Rev. Phytopathol. 51:267-289.

Gent, D. H., Mahaffee, W. F., Turechek, W. W., Ocamb, C. M., Twomey, M. C., Woods, J. L., and Probst, C. 2019b. Risk factors for bud perennation of *Podosphaera macularis* on hop. Phytopathology 109:74-83.

Gent, D. H., Massie, S. T., Twomey, M. C., and Wolfenbarger, S. N. 2017. Adaptation to partial resistance to powdery mildew in the hop cultivar Cascade by *Podosphaera macularis*. Plant Dis. 101:874-881.

Gent, D. H., Nelson, M. E., Grove, G. G., Mahaffee, W. F., Turechek, W. W., and Woods, J. L. 2012. Association of spring pruning practices with severity of powdery mildew and downy mildew on hop. Plant Dis. 96: 1343-1351.

Gent, D. H., Probst, C., Nelson, M. E., Grove, G. G., Massie, S. T., and Twomey, M. C. 2016. Interaction of basal foliage removal and late-season fungicide applications in management of hop powdery mildew. Plant Dis. 100: 1153-1160.

Henning, J. A., Gent, D. H., Twomey, M. C., Townsend, M. S., Pitra, N. J., and Matthews, P. D. 2016. Genotyping-by-sequencing of a bi-parental mapping population segregating for downy mildew resistance in hop (*Humulus lupulus* L.). Euphytica 208:545-559.

Henning, J. A., Townsend, M. S., Gent, D. H., Bassil, N., Matthews, P., Buck, E., and Beatson, R. 2011. QTL mapping of powdery mildew susceptibility in hop (*Humulus lupulus* L.). Euphytica 180:411-420.

Hogenboom, N. G. 1993. Economic importance of breeding for disease resistance. Pages 5-9 in: Durability of Disease Resistance. T. H. Jacobs and J. E. Parlevliet, eds. Springer Netherlands, Dordrecht, the Netherlands.

Hückelhoven, R., and Pillen, K. 2024. Genetics of plant disease and resistance. Pages 133-159 in: Agrios' Plant Pathology, 6th ed. R. P. Oliver, R. Hückelhoven, E. M. Del Ponte, and A. Di Pietro, eds. Academic Press, London, U.K.

Hwang, J. Y., Bhattacharyya, S., Chatterjee, S., Marsh, T., Pedro, J., and Gent, D. H. 2025a. Fungicide selection, disease risk, and grower switching behavior. Plant Dis. 109:2525-2534.

Hwang, J. Y., Bhattacharyya, S., Chatterjee, S., Marsh, T., Pedro, J., and Gent, D. H. 2025b. The diversity and determinants of fungicide programs for hop powdery mildew. Plant Dis. 109:2514-2524.

Hwang, J. Y., Bhattacharyya, S., Chatterjee, S., Marsh, T. L., Pedro, J. F., and Gent, D. H. 2024. What explains hop growers' fungicide use intensity and management costs in response to powdery mildew? Phytopathology 114:2287-2299.

Lankau, E. W., Xue, D., Christensen, R., Gevens, A. J., and Lankau, R. A. 2020. Management and soil conditions influence common scab severity on potato tubers via indirect effects on soil microbial communities. Phytopathology 110:1049-1055.

Laurie, R. B., Richardson, B. J., Ross, C. J., and Gent, D. H. 2022. Yard age, cultivar susceptibility, and spring pruning practices as risk factors for overwintering of *Podosphaera macularis* on hop. PhytoFrontiers 3:390-398.

Liman Harou, I. L., Whitney, C., Kung'u, J., and Luedeling, E. 2021. Crop modelling in data-poor environments – A knowledge-informed probabilistic approach to appreciate risks and uncertainties in flood-based farming systems. Agric. Syst. 187:103014.

Mahaffee, W. F., Thomas, C. S., Turechek, W. W., Ocamb, C. M., Nelson, M. E., Fox, A., and Gubler, W. D. 2003. Responding to an introduced pathogen: *Podosphaera macularis* (hop powdery mildew) in the Pacific Northwest. Plant Health Prog. 4:21.

Marcot, B. G. 2012. Metrics for evaluating performance and uncertainty of Bayesian network models. Ecol. Model. 230:50-62.

McRoberts, N., Brinker, S., and Coleman, K. 2024. Theories for understanding the effect of impact assessment and project evaluation on the practice of science. Annu. Rev. Phytopathol. 62:385-400.

McRoberts, N., Hall, C., Madden, L. V., and Hughes, G. 2011. Perceptions of disease risk: From social construction of subjective judgments to rational decision making. Phytopathology 101:654-665.

Mooney, D. F., Swinton, S. M., Subía, C., and Peralta, E. 2022. Returns to disease resistance research when pest management is an option. Sustainability 14:2859.

Morris, M. L., and Heisey, P. W. 2003. Estimating the benefits of plant breeding research: Methodological issues and practical challenges. Agric. Econ. 29:241-252.

Natori, K., Uto, M., Nishiyama, Y., Kawano, S., and Ueno, M. 2015. Constraint-based learning Bayesian networks using Bayes factor. Pages 15-31 in: Advanced Methodologies for Bayesian Networks. J. Suzuki and M. Ueno, eds. AMBN 2015 Lecture Notes in Computer Science. Springer, Cham, Switzerland.

Oerke, E.-C. 2006. Crop losses to pests. J. Agric. Sci. 144:31-43.

Olatoye, M. O., Wiseman, M., Gent, D. H., Henning, J. A., and Altendorf, K. R. 2023. Genetic characterization of downy mildew resistance from the hop (*Humulus lupulus* L.) line USDA 64035M. Crop Sci. 63:1082-1091.

Pearl, J. 2009. Causal inference in statistics: An overview. Stat. Surv. 3:96-146.

Pearl, J. 2010. The foundations of causal inference. Sociol. Methodol. 40:75-149.

Pedro, J., Bhattacharyya, S., Chatterjee, S., Marsh, T. L., Hwang, J. Y., and Gent, D. 2025. Inoculum dose, diversity, dispersal, and damage: Simulating optimal economic control of an aerially-dispersed plant pathogen. SSRN Preprint. http://dx.doi.org/10.2139/ssrn.5845162

Purayannur, S., Gent, D. H., Miles, T. D., Radišek, S., and Quesada-Ocampo, L. M. 2021. The hop downy mildew pathogen *Pseudoperonospora humuli*. Mol. Plant Pathol. 22:755-768.

Royle, D. J. 1978. Powdery mildew of the hop. Pages 381-409 in: The Powdery Mildews. D. Spencer, ed. Academic Press, New York.

Scutari, M. 2010. Learning Bayesian networks with the bnlearn R package. J. Stat. Sw. 35:1-22.

Scutari, M., Auconi, P., Caldarelli, G., and Franchi, L. 2017. Bayesian networks analysis of malocclusion data. Sci. Rep. 7:15236.

Sherman, J., and Gent, D. H. 2014. Concepts of sustainability, motivations for pest management approaches, and implications for communicating change. Plant Dis. 98:1024-1035.

Spirtes, P. 2010. Introduction to causal inference. J. Mach. Learn. Res. 11:1643-1662.

Turechek, W. W., and Mahaffee, W. F. 2004. Spatial pattern analysis of hop powdery mildew in the Pacific Northwest: Implications for sampling. Phytopathology 94:1116-1128.

Turechek, W. W., Mahaffee, W. F., and Ocamb, C. M. 2001. Development of management strategies for hop powdery mildew in the Pacific Northwest. Plant Health Prog. 2:8.

VanderWeele, T. J., and An, W. 2013. Social networks and causal inference. Pages 353-374 in: Handbook of Causal Analysis for Social Research, Handbooks of Sociology and Social Research. S. L. Morgan, ed. Springer Netherlands, Dordrecht, the Netherlands.

Wearing, C. H. 1988. Evaluating the IPM implementation process. Annu. Rev. Entomol. 33:17-38.

Wolfenbarger, S. N., Massie, S. T., Ocamb, C., Eck, E. B., Grove, G. G., Nelson, M. E., Probst, C., Twomey, M. C., and Gent, D. H. 2016. Distribution and characterization of *Podosphaera macularis* virulent on hop cultivars possessing *R6*-based resistance to powdery mildew. Plant Dis. 100: 1212-1221.

Zhou, Y., Hospedales, T. M., and Fenton, N. 2016. When and where to transfer for Bayesian network parameter learning. Expert Syst. Appl. 55:361-373.