Phytopathology® • 2024 • 114:2287-2299 • https://doi.org/10.1094/PHYTO-04-24-0127-R

Ecology and Epidemiology

What Explains Hop Growers' Fungicide Use Intensity and Management Costs in Response to Powdery Mildew?

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
[2] Department of Statistics, Oregon State University, Corvallis, OR 97331
[3] Department of Mathematics, City University of New York, New York City, NY 10031
[4] School of Economic Sciences, Washington State University, Pullman, WA 99163
Accepted for publication 4 August 2024.

Abstract

Methods for causal inference from observational data are common in human disease epidemiology and social sciences but are used relatively little in plant pathology. We draw upon an extensive data set of the incidence of hop plants with powdery mildew (caused by *Podosphaera macularis*) collected from yards in Oregon from 2014 to 2017 and associated metadata on grower cultural practices, cultivar susceptibility to powdery mildew, and pesticide application records to understand variation in and causes of growers' fungicide use and associated costs. An instrumental causal forest model identified growers' spring pruning thoroughness, cultivar susceptibility to two of the dominant pathogenic races of *P. macularis*, network centrality of yards during May–June and June–July time transitions, and the initial strain of the fungus detected as important variables determining the number of pesticide active constituents applied by growers and the associated costs they incurred in response to powdery mildew. Exposure-response function models fit after covariate weighting indicated that both the number of pesticide active constituents applied and their associated costs scaled linearly with the seasonal mean incidence of plants with powdery mildew. Although the causes of pesticide use intensity are multifaceted, biological and production factors collectively influence the incidence of powdery mildew, which has a direct exposure-response relationship with the number of pesticide active constituents that growers apply and their costs. Our analyses point to several potential strategies for reducing pesticide use and costs for management of powdery mildew on hop. We also highlight the utility of these methods for causal inference in observational studies.

Keywords: data science, disease control and pest management, epidemiology, fungal pathogens

The severity of plant diseases varies in response to numerous well-known factors related to host susceptibility, the virulence of the pathogen, and the favorability of the environment. In managed systems, disease severity also varies in response to a fourth vertex of the disease tetrahedron, control measures applied by humans (Agrios 2005). When, what, and how often growers apply pesticides can vary substantially from field to field and farm to farm (Andert et al. 2015; Jørgensen et al. 2017; Marsh et al. 2000; Nicholson and Williams 2021; Oakley et al. 2007). Identifying variates that explain and predict outcomes of production systems such as disease severity and pesticide use can provide insights into the individual risk factors or suites of factors underlying the efficiency of entire production systems. These factors could be related to variation in cultural practices, regional differences due to climate or crop diversity, crop value, or grower risk avoidance (Gent et al. 2012; Jørgensen et al. 2017; Lybbert et al. 2016; Mourtzinis et al. 2018, 2019; Nicholson and Williams 2021; Savary et al. 2000).

Our motivating pathosystem for this research is the disease powdery mildew of hop, caused by the fungus *Podosphaera macularis*. Hop (*Humulus lupulus*) is a long-living herbaceous perennial that is produced for its strobiles, colloquially referred to as cones or simply hops (Neve 1991). Powdery mildew is one of the most damaging and costly diseases for hop producers in the Western United States. This is in part because of the long period during the growing season when the disease must be managed, the rapid annual growth of the host, and market factors that prioritize brewing attributes over agronomic factors such as disease resistance (Gent et al. 2008; Mahaffee et al. 2009). Presently in the Western United States, *P. macularis* persists in hop yards in association with infected crown buds because one of the two mating types necessary for formation of ascocarps is absent from the region (Weldon et al. 2021a, b; Wolfenbarger et al. 2015). Bud infection may lead to shoots emerging from winter dormancy colonized by *P. macularis*, the so-called flag shoots (Gent et al. 2018). Bud perennation leading to flag shoot development is a rare event that occurs, on average, in approximately 6% of hop yards in Oregon (Laurie et al. 2023). From these primary infections, the pathogen is readily disseminated by wind within and between hop yards to infect leaves, stems, and other photosynthetic tissue. The spread of powdery mildew between hop yards can be conceptualized as a directed network at the landscape level, with each hop yard representing a node and edge weight representing the probability of disease transmission due to spread from other nodes (Gent et al. 2019a). Disease spread among yards is influenced by the source strength of the inoculum in affected yards, host susceptibility to the pathogen, distance from a source, and wind run (Gent et al. 2019a).

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

Funding: Support was provided by the U.S. Department of Agriculture-Agricultural Research Service (2072-21000-061-000-D), the U.S. Department of Agriculture-National Institute of Food and Agriculture (2021-51181-35901), and the National Science Foundation (DMS 2154564).

e-Xtra: Supplementary material is available online.

The author(s) declare no conflict of interest.

Conidia produced from foliar infections provide the inoculum for later infection of cones that begin to form just after the summer solstice (Royle 1978; Twomey et al. 2015). Cultural practices such as removal of the first flush of shoots (referred to as pruning), elimination of superfluous basal foliage, and moderation of nitrogen fertility delay disease onset or reduce epidemic velocity (Gent et al. 2012, 2016, 2019b, 2024; Probst et al. 2016; Royle 1978). Resistance to powdery mildew is available in certain cultivars (Mahaffee et al. 2009), but in many instances, host resistance has been short-lived when a given form of resistance is broadly deployed on the landscape (Block et al. 2021; Gent et al. 2017; Wolfenbarger et al. 2016). Resistance often is not available in cultivars demanded by brewers or markets. Consequently, regular applications of fungicides are required to suppress powdery mildew to maintain both hop yield and quality (Gent et al. 2014; Nelson et al. 2015; Royle 1978).

In 2010, hop producers in Oregon reported making on average 5.1 fungicide applications per year for powdery mildew in susceptible cultivars (Sherman and Gent 2014). There is not a single preferred fungicide program used by hop producers in the Western United States for powdery mildew (Sherman and Gent 2014). Rather, fungicide use intensity may vary among individual hop yards based on the susceptibility of given cultivar (Sherman and Gent 2014), cultural practices such as the thoroughness of spring pruning (Gent et al. 2012), and the presence of an overwintered inoculum (Gent et al. 2019b), among other factors. Representing only the population mean fungicide use intensity may obscure important treatment heterogeneity between yards or farms. Indeed, analysis of pesticide use intensity and patterns in other pathosystems typically reveals treatment heterogeneity that can illuminate practices or factors associated with more or less intensive pesticide use patterns (Andert et al. 2015; Lybbert et al. 2016; Nicholson and Williams 2021; Oakley et al. 2007).

There are several potential methods to understand how a complex suite of variates may influence disease development and, in turn, pesticide use by growers. For establishing causality, randomized control trials conducted over space and time could experimentally deduce the effect of common production practices on disease development and grower responses. Randomized control trials are commonly considered the gold standard for establishing cause-and-effect relationships (Hariton and Locascio 2018), but often, they are too costly to conduct because crop production systems are incredibly diverse and the number of potential factors of interest is impractical to assess in more than just a few combinations. Randomized control trials also suffer from external validity issues (Rothwell 2005). That is, conclusions from one experiment may not be generalizable to other situations because of differences in factors such as microclimate, management conditions, or other covariates. Some factors are also impossible to vary because of ethical considerations with human subjects or practical constraints. As an alternative to randomized control trials, process-based or agent-base models that simulate all aspects of disease development and grower responses to disease are useful tools for gaining insights into complex system behavior (Atallah et al. 2012; Babcock et al. 2022; Cunniffe et al. 2016; Murray-Watson et al. 2022). However, these models require many decisions and assumptions on what variables to include and the functional forms of relationships among variables. Data acquisition costs for their development and use may be cost prohibitive. Because of these limitations, analyses of large, observational data sets are common in ecology, the social sciences, and epidemiology. Numerous data-driven methods are available (Pichler and Hartig 2023).

There is rich and extensive literature in the social sciences on natural experiments or quasi-experiments that apply the Rubin casual model to observational data, which can be applied in an associational or causal inference framework. These methods are not without limitations when attempting to understand cause-and-effect relationships because of potential confounding from observed or unobserved variables (Athey and Wager 2019). In observational studies, in which the treatment assignment was out of the control of the investigator, much of the epidemiological literature assumes that the assignment of a treatment might be ignorable after controlling for potential confounders (Little and Rubin 2000). Observational causal inference methods usually start from a baseline observational model and then build on the treatment effect estimation framework under the auspices of the baseline model. In this paper, we start with a predictive model based on random forests (Breiman 2001).

Random forests are a widely used algorithm for supervised statistical learning that have found increasing application in plant disease contexts (Domingues et al. 2022; Shah et al. 2023). Random forest models are an ensemble method that generalize numerous individual regression trees for conditional mean estimation (Breiman 2001). Individual trees have low bias but high variance, which limits their out-of-sample prediction. The random forest method reduces variance through drawing bootstrap samples from a data set to fit an unpruned regression tree for each respective bootstrap sample. The variable selection for each split in a random forest is conducted from a small, random subset of independent variables to avoid both the path dependency problem and the "small n large p" or high-dimensionality problem (Strobl et al. 2007). From the final node, the response variable is predicted as an average or majority vote of the predictions of the ensemble of pruned trees. The random forest algorithm for regression trees uses mean squared error (MSE) to determine how the data branches from each node:

$$\mathrm{MSE} = \frac{1}{N}\sum_{i=1}^{N}(f_i - y_i)^2$$

where N is the number of observations, f_i is the value returned by the model, and y_i is the actual value for each observation. Random forest models are attractive for prediction problems with large, high-dimensional data sets because of their flexibility for handling multiple forms of data, robustness to outliers, and few hyperparameters. A random forest is a stepwise linear approximation that provides a global understanding of a treatment effect averaged over the entire sample population; the algorithm does not estimate heterogeneity in covariates (treatment heterogeneity), which may mask variability within subpopulations.

Generalized random forests are an extension of random forests built on local moment conditions that seeks to estimate heterogeneous treatment effects for causal inference (Athey et al. 2019). Instead of making a prediction of the outcome itself, generalized random forests enable prediction of a treatment effect of a specific covariate (subgroup) on a response variable. This is powerful for understanding how a predictor variable interacts with contextual covariates, for example, biophysical or management-related factors. Generalized random forests first calculate a causal regression tree from a low-dimensional representation of treatment effect heterogeneity (with respect to observable covariates) based on recursive partitioning and sampling without replacement to estimate the conditional average treatment effect (CATE). A subset of observations J are randomly drawn from N number of observations, and J is partitioned into two sets: J_1 to build trees that maximize the variance of CATE estimation in a leaf and J_2 to analyze the CATE. While building a tree, the m subset of covariates is utilized to split the observations to maximize heterogeneity in subgroups. For unit i let Y_i be the outcome and T_i be the treatment, then an estimate of the heterogeneous treatment effect $\hat{\tau}$ can be conducted in the presence of confounding variables as follows (Nie and Wager 2021):

$$\hat{\tau} = \frac{\sum_i a_i(x)\,(Y_i - \hat{m}^{(-i)}(X_i))(T_i - \hat{e}^{(-i)}(X_i))}{\sum_i a_i(x)\left(T_i - \hat{e}^{(-i)}(X_i)\right)^2}$$

where $a_i(x)$ is a data-adaptive kernel that weights how often a unit i in training has fallen into the same leaf conditional on covariates x, $\hat{m}^{(-i)}$ is the out-of-sample prediction of a conditional outcome $m(x) = \mathbb{E}[Y_i|X_i = x]$, and $\hat{e}^{(-i)}$ is the conditional probability of being treated $e(x) = \mathbb{E}[T_i|X_i = x]$ (Athey and Wager 2019). In brief, the generalized random forest algorithm first estimates $\hat{e}(X_i)$ and $\hat{m}(X_i)$ separately, along with a bootstrap-aggregated (out-of-bag) prediction. A residual treatment $T_i - \hat{e}^{(-i)}(X_i)$ and outcome $Y_i - \hat{m}^{(-i)}(X_i)$ are computed, and the generalized random forests are trained on these residuals (Athey and Imbens 2019).

Generalized random forests can be implemented using an instrumental variable to resolve inherent endogeneity or induced endogeneity of measurement error. An instrumental variable is a proxy correlated with a predictor variable conditionally on other covariates but conditionally uncorrelated with a response variable. An instrumental variable produces more stable and unbiased results when an explanatory variable is correlated with the residuals due to confounding from unobserved covariates, non-random measurement error, or when a response variable may change the value of an explanatory variable (Greene 2018). In an instrumental forest, the conditional local average treatment effect, τ, is identified using instruments. Therefore, the instrumental forest is followed as

$$(\tau) = \frac{Cov[Y,\ Z|X = x]}{Cov[W,\ Z|X = x]}$$

where W is treatment and Z is the instrument (Tibshirani et al. 2023). We are unfamiliar with application of generalized random forests or instrumental forests in plant disease epidemiology.

In this research, we draw upon an extensive data set of the incidence of hop powdery mildew collected in a census sample of commercial hop yards in Oregon from 2014 to 2017 and associated metadata on certain grower cultural practices, cultivar susceptibility to powdery mildew, and pesticide application records (Gent et al. 2019a). Our overall objectives were to summarize variation in growers' fungicide use and the associated cost of their fungicide programs in response to powdery mildew. We do this through fitting generalized random forests to predict fungicide use intensity and the associated costs of the fungicide programs. We then estimate heterogeneous treatment effects to understand factors related to variation in fungicide use intensity and its costs. We also confirmed the findings by fitting an exposure-response model on a covariate-weighted data set to quantify how fungicide use and costs change with exposure to powdery mildew.

Materials and Methods

Description of disease assessment and data set

A full description of the disease assessment methods and data is presented in Gent et al. (2019a), so we provide only an abbreviated summary here. In each year from 2014 to 2017, a census sample of the incidence of hop plants with powdery mildew was conducted monthly from April to July by sampling every hop yard on every hop farm in the eastern extent of the hop production region in Oregon. There were 8 to 10 farms sampled per year, spanning from near the cities of Silverton to Hubbard (maximum distance between yards of 26 km). In total, data were available for 103 yards assessed in 2014, 118 in 2015, 112 in 2016, and 125 in 2017. As this was a census sample, all cultivars were evaluated independent of their susceptibility to powdery mildew. The cultivars evaluated and their relative susceptibility to powdery mildew are detailed in Laurie et al. (2023).

The incidence of plants with powdery mildew was assessed using a modification of cluster sampling methods described previously (Turechek and Mahaffee 2004; Turechek et al. 2001). Each yard was divided into strata of 20 rows, and at least two strata per yard were sampled by evaluating 50 to 200 hills (referred to hereafter as plants) in one transect (row) per strata. The number of plants harboring a flag shoot or colonies of *P. macularis* was recorded, the former nearly always occurring during sampling in April and May.

Potential covariates were recorded during the time of disease assessments, after consultation with the cooperating growers, derived from literature or other sources, or calculated, as we describe below. These were generally related to host susceptibility to powdery mildew, the race (strain) of the pathogen present in a yard, various network centrality measures, selected cultural practices, and physical location of yards.

We assigned each cultivar an ordinal score for its susceptibility to two pathogenic races of *P. macularis* as described in Laurie et al. (2023). For background, in the Pacific Northwestern United States, there are three dominant pathogenic races of *P. macularis* that display differential host genotype adaptation (Gent et al. 2017; Wolfenbarger et al. 2016). Two of the three races were relevant at the time of our study, namely strains virulent to hop cultivars possessing the R-genes Rb, R3, or R5 (race Vb,V3,V5) and strains virulent on cultivars possessing Rb, R3, R4, R5, or R6 (race Vb,V3,V4,V5,V6) (Wolfenbarger et al. 2016). For brevity, we refer to race Vb,V3,V5 as non-V6 virulent and race Vb,V3,V4,V5,V6 as V6-virulent because a defining difference between these races is their ability to cause disease on the widely deployed powdery mildew resistance dubbed R6 (Henning et al. 2011; Wolfenbarger et al. 2016).

The initial strain of *P. macularis* present in each hop yard was determined as being virulent or not on cultivars possessing R6, as described previously (Gent et al. 2019a). In 13 instances, we could not obtain isolates or virulence data. In these instances, or when powdery mildew did not occur at any level, we coded the initial strain as 0. Otherwise, we coded the initial strain as 1 if the pathogen was non-V6-virulent and 2 if the pathogen was V6-virulent.

The thoroughness of spring pruning was rated using a 1-to-5 ordinal scale, as described previously (Laurie et al. 2023). Spring pruning is potentially relevant for powdery mildew development and pesticide use intensity because this practice can remove the overwintered inoculum associated with flag shoots and delay development of the disease (Gent et al. 2012, 2019b; Laurie et al. 2023; Probst et al. 2016; Turechek et al. 2001). In this ordinal scale, 1 represents the most thorough pruning, which removed all green leaves and stems from every plant. Each subsequent point represents an approximation of the incidence of plants with green foliage remaining, such that a 5 indicates that more than 80% of plants had green leaves and shoots remaining after pruning.

The centrality of each hop yard in the inferred disease spreading network was expressed using the outward degree centrality statistics reported in Gent et al. (2019a). The outward degree centrality of the ith node is the number of outward-directed edges stemming from that node:

$$d_k = \sum_i 1\,\{(k, i) \in E\}$$

where outward degree centrality is summed on the edges from k to i on edge set E. Outward degree centrality was calculated for monthly time transitions from May to June and June to July for each network of yards affected by V6-virulent or non-V6-virulent strains planted to cultivars that possess R6 and those that do not. For each network and month, degree centrality was dichotomized depending on whether the statistic was non-zero. We also calculated the midpoint between the centroid of all yards and divided the landscape into four equidistant quadrants from the central reference point to investigate whether general position in the landscape influenced pesticide use and costs.

Lastly, we assigned a dummy variable to each grower and year. One grower had two hop yards that were produced in only 2017. We removed data for this grower given the small sample size. Summed over the 4 years represented in the data set, the number of yards for the other growers ranged from 21 to 103, with a mean of 51.

Descriptive statistics for pesticide use intensity and costs

We used the total number of fungicide active constituents applied by growers in a given year and their estimated costs as the response variables in our analyses. We obtained pesticide application records from each grower for all yards sampled, digitized their records, and then tallied the number of active constituents applied per hectare with activity against powdery mildew. The 17 active constituents in 22 products we identified in the pesticide application records are given in Supplementary Table S1. Some subjectivity was involved in determining whether to consider certain pesticides and adjuvants as powdery mildew fungicides. For instance, some adjuvants may have some activity against powdery mildew diseases (Jibrin et al. 2021) even though they are typically not used specifically for powdery mildew suppression. We did not consider adjuvants in our tally because it would have artificially inflated the tallies because adjuvants are used with most pesticides and because most adjuvants have not been evaluated for suppression of powdery mildew diseases. As an exception to this rule, we did consider various mineral oils as active constituents for powdery mildew even though they may be used as general adjuvants because mineral oils have well-documented activity against powdery mildew and, in our experience, are routinely used specifically for powdery mildew suppression (Claassen et al. 2022). Similarly, we included copper-based fungicides in our tallies because these are recommended for powdery mildew suppression (Royle 1978), even though they are often applied for suppression of other diseases as well. For all pesticides, we considered only the number of active constituents (Essling et al. 2021) and not the dose or dose equivalency applied.

To estimate the annual costs of the powdery mildew fungicides, we requested price quotes for the pesticides from each of three vendors in western Oregon that service hop producers. Quotes were obtained from each vendor during October and November 2021 that reflected nominal prices for 2014 to 2022, where available. We estimated real prices using January 2022 as the base by adjusting the nominal price by the U.S. Bureau of Labor Statistics producer price index for farm products for each year, and then averaged over all available years to derive a single real price per unit. In some instances, a given mineral oil or sulfur-based product was a proprietary formulation only available from one source, so averaging across the three vendors was not possible. In these situations, we used the price quote from the sole source that carried it or, if none of the three vendors carried a product found in the pesticide records, we substituted an equivalent product.

Random forest model

We used the Python 3.10 *Scikit-learn* tools (Pedregosa et al. 2011) to fit and optimize a random forest through hyperparameter tuning. We trained on 80% of the data set and tested 20%, testing several thousand combinations of settings while growing the trees. We evaluated the performance of a random forest using mean absolute error (MAE), MSE, root mean squared error, and the coefficient of determination (R^2) in the training data sets, test data set, and validation data sets based on leave-one-out cross-validation.

As a check on the prediction accuracy of the random forest, we applied three other machine learning methods to the data set for comparison: ridge regression, LASSO (least absolute shrinkage and selection operator) regression, and a decision tree (Pichler and Hartig 2023). Ridge regression is an extension of linear regression that is more stable in the presence of multicollinearity that uses a modified loss function. Ridge regression can control overfitting and underfitting by a parameter and reduce variance in parameter estimates. LASSO regression is a regularization technique that applies a modified loss function to obtain a more accurate prediction using shrinkage, which allows estimated coefficient parameter values to shrink toward zero. A decision tree is a single model algorithm, which continues to split the data until it reaches a point where it cannot improve predictions. We again calculated mean absolute error, MSE, root mean squared error, and the R^2 for each of the training data, test data, and in cross-validation to measure in- and out-of-sample model performance and prediction, similarly to the random forest.

Generalized random forest

One of our aims was to examine heterogeneous treatment effects associated with the presence of powdery mildew for the number of pesticide active constituents applied and the annual cost of pesticides. We constructed histograms of conditional average treatment effects for the two response variables as a diagnostic to detect potential heterogeneity. Based on the approximately normal distribution of pesticide use intensity and costs, we suspected heterogeneity existed among potential subpopulations in the data set (Fig. 1).

We implemented an instrumental forest model using the incidence of powdery mildew in May as an instrumental variable for the seasonal mean incidence of plants with powdery mildew. We used this instrumental variable because the seasonal mean incidence of plants with powdery mildew is potentially related to the residuals, as growers that apply more or less pesticides for powdery mildew may in turn modify the seasonal mean incidence of plants with powdery mildew. We reasoned that the incidence of plants with powdery mildew in May could be a suitable instrument for the seasonal mean incidence of plants with powdery mildew for two primary reasons. First, the seasonal mean incidence of powdery mildew is correlated with primary inoculum dose and disease development in the earliest stages of the epidemic (Gent et al. 2019b; Turechek et al. 2001). Second, the incidence of plants with powdery mildew in May is correlated with this mean incidence of plants with powdery mildew over the season but is less strongly correlated with the response variables (Supplementary Fig. S1). We also considered other potential instruments, such as the incidence of plants with flag shoots. We used the incidence of plants with powdery mildew in May instead of flag shoots because most observations for flag shoots were zero.

With strong instruments, two-stage least squares and limited information maximum likelihood estimators are asymptotically unbiased. However, weak instruments can bias point estimates and distort test sizes (Nelson and Startz 1990). We tested whether the incidence of plants with powdery mildew in May was a weak instrument. Stock and Yogo (2005) suggest a robust test for weak instruments under the assumption of conditionally homoscedastic, serially uncorrelated model errors. The test rejects the null hypothesis of weak instruments when the critical value exceeds a given threshold. We used the STATA 'weakivtest' critical values for the null hypothesis that the two-stage least squares asymptotic bias exceeds 10% of a fraction τ of worst-case bias (Olea and Pflueger 2013). After establishing that the incidence of plants in May was not a weak instrument, we proceeded to fit the causal forest with this instrument.

Each hop yard was indexed by $i = 1, 2, \ldots, N$, and a vector of the covariates. The incidence of powdery mildew as a treatment is indicated as a continuous variable $T_i \in \{0,\ 0.6\}$, representing the observed range of disease incidence. Concretely in our application, treatment is the seasonal mean incidence of plants with powdery mildew in a given hop in a given year, which may have a causal effect on the number of active constituents applied and annual costs in the population of hop yards. The effect of powdery mildew occurrence on these response variables may be heterogeneous depending on various covariates, such as cultivar susceptibility to the disease or position in the landscape. The CATE was estimated as the average slope in the partial dependence for a continuous treatment variable (Chernozhukov et al. 2022; Tibshirani et al. 2023):

$$\tau(x) = E[(Cov[W, Y|X])/(Var[W|X])]$$

where W is the continuous treatment variable and X is covariates. The general tree-specific procedure followed these five steps:

i. Randomly draw a sample of hop yards and a subset of available covariates
ii. Randomly split the sample in half, creating a training set and an estimation set
iii. Use the training set to grow a tree

iv. Match hop yards in the estimation set to leaves of the tree according to the observed covariates
v. Estimate average treatment effects from observations determined by the final partition, and then estimate for each leaf or observation based on the associated observed covariates

For each of 2,000 trees, we predicted the average marginal treatment effect on pesticide use intensity or annual costs for all hop yards not used in growing the specific tree in step i above. Each of these predictions was aggregated into a single estimate of CATE for the number of active constituents or annual costs. We used an "honest" estimation approach based on recursive partitioning and subsampling to estimate CATE, as described in Athey and Imbens (2016). We excluded estimates of CATE for the cases when the initial strain was assigned a value of 0 or when the outward degree centrality variables were 0 because powdery mildew essentially did not occur in these categories. Thus, these variables might cause confounding bias with powdery mildew occurrence due to the nature of the covariates. Analyses were conducted using the default settings in the R package *grf* for an instrumental forest with a continuous treatment (Tibshirani et al. 2023). The number of variables during each split was $\sqrt{p}$ +20, where p is the number of variables considered during each split. The minimum number of observations in each tree leaf was 5. Each unit was given the same weight. Variable importance was calculated as the weighted sum of how many times a feature i was split on at each depth in the forest.

Covariate balancing and exposure-response function

To complement the analysis above, we further examined covariate balancing and estimated an exposure-response function. An important assumption for establishing a causal effect is that there must be no other confounders present that could have their own causal effect (Pearl 2010). One means to test this assumption is through calculation of a propensity score. A propensity score is the conditional probability of assignment to a treatment given a set of covariates, and this score is used to assure treatment and control groups have similar covariates when making causal inference (Rosenbaum and Rubin 1983). Given a propensity score, treatment effects can be estimated by appropriately matching covariates to reduce confounding through weighting, stratification, or adjustment of the propensity score. Propensity score approaches are well developed for binary treatments (Austin 2011), but statistical methodology is just now being developed to calculate propensity scores for treatments with continuous effect, such as exposure to disease in our data set (Khoshnevis et al. 2023).

Given the potential for confounding among the covariates in the instrumental forest we described above, we calculated a generalized propensity score and fit an exposure-response function as a further check on the instrumental forest analysis. The generalized propensity score is defined as the conditional density of the treatment level given confounders for a continuous treatment. The causal inference of a continuous exposure can be examined using generalized propensity scores for removing selection bias. The average causal exposure-response function is defined as the specific range of the exposure levels $w_i \in W$ and is followed as

$$\begin{aligned}\mu(w) &= E[Y_i(w)] \\ &= \lim_{\delta \to 0} E[E\{Y_i^{obs} | e(W_i, C_i), W_i \in [w-\delta, w+\delta]\}]\end{aligned}$$

where C_i represents the pre-exposure covariates for unit i, and δ is a constant for a given data set with sample size N. To estimate a generalized propensity score for this model, we first generated a pseudo-population by weighted observations to reduce pre-exposure covariate confounding across different levels of disease incidence, our exposure variable. We conducted a covariate balance test to assess how well the weighting balanced the covariates in the pseudo-population. The covariate balance test measures the absolute correlations between the exposure (mean seasonal disease incidence) and observed pre-exposure covariates. The correlation coefficients allow us to investigate how closely the distribution of observed pre-exposure covariates aligns across all levels of exposure (Khoshnevis et al. 2023). Second, we calculated generalized propensity scores using conditional density estimation to visualize

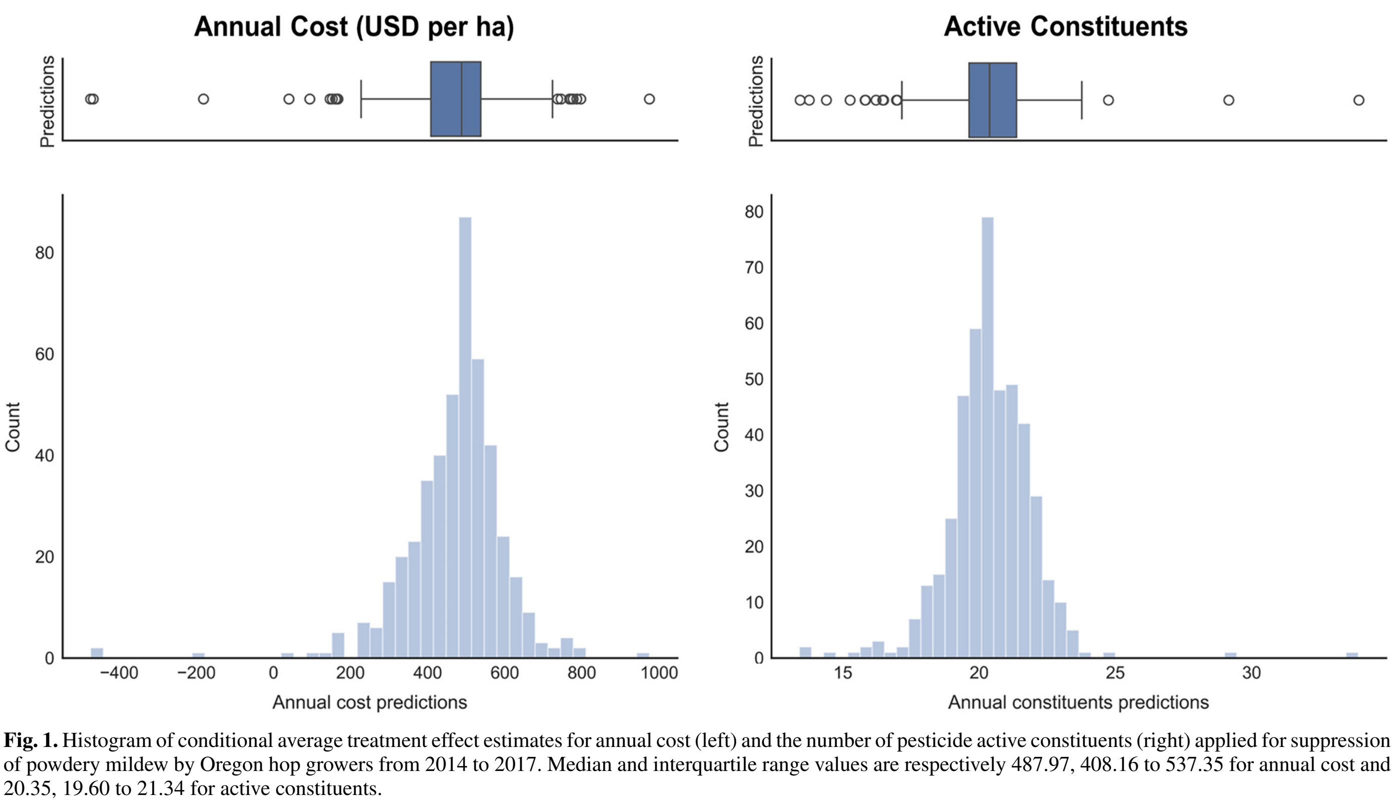

Fig. 1. Histogram of conditional average treatment effect estimates for annual cost (left) and the number of pesticide active constituents (right) applied for suppression of powdery mildew by Oregon hop growers from 2014 to 2017. Median and interquartile range values are respectively 487.97, 408.16 to 537.35 for annual cost and 20.35, 19.60 to 21.34 for active constituents.

the degree to which we could achieve covariate balance. We used a nonparametric kernel density based on a local polynomial to estimate the conditional density of the outcome variables given the covariates.

After confirming acceptable covariate balance, we then estimated the exposure response functions to quantify the effect of disease incidence on pesticide active constitutes applied and annual costs. We fit the exposure-response functions as a parametric local-linear regression model. We conducted these analyses using the R package *CausalGPS* (Khoshnevis et al. 2023).

Results

Descriptive statistics for pesticide use intensity and costs

Overall, the seasonal mean incidence of plants with powdery mildew was 0.03 (standard deviation 0.08) over all yards and years (Table 1). The mean number of fungicide active constituents applied for powdery mildew and the associated costs varied somewhat by year, but among years, there were similar monthly patterns when growers applied the most pesticides and when they incurred the most costs (Fig. 2). There were few to no pesticides applied in March, followed by a general pattern of increasing pesticide use up to June, and then a progressive decrease in July and August in most years (Fig. 2B). Whereas the most active constituents were applied in June, growers incurred the most costs in July or August, reflecting that the average cost per active constituent applied was greater in these months than in earlier months (Fig. 2A). There were also yearly trends in pesticide use and their costs, with the mean number of active constituents applied (independent of the covariates) varying from 7.86 in 2014 to 8.28 in 2017 (Fig. 2D). Correspondingly, mean annual costs per hectare varied from as low as \$314.58 in 2015 to \$587.53 in 2017 (Fig. 2C). Unsurprisingly, the number of active constituents applied and the associated annual costs were closely correlated (Spearman's rank correlation coefficient $\rho = 0.73$; Fig. 3).

The number of active constituents applied and annual costs varied depending on whether powdery mildew was present. In yards where powdery mildew was not detected, growers applied on average 5.75 fungicide active constituents, whereas growers applied 7.70 active constituents in yards where the disease was detected. In turn, associated costs for these pesticides were on average \$131.49 per hectare more expensive in yards with disease as compared with those without disease (\$516.44 versus \$384.95) (Table 1).

For the subset of covariates identified as important in the generalized random forest model described below, the means of the powdery mildew treatment subgroups (i.e., powdery mildew present or not) were significantly different based on *t* tests, with the exception of pruning thoroughness ($P = 0.963$). For reference, we provide summary statistics in Table 1 for most variables relevant later in the analysis but point out that treatment means for certain variables are trivial (e.g., outward degree centrality when powdery mildew was not present because, by definition, these yards would have outward degree centrality of 0).

Random forest model

Prediction of the number of active constituents applied, and annual costs were well captured, as measured by MSE and R^2, in out-of-sample test data and cross-validation by the random forest model as compared with other machine learning approaches (Table 2). Given the overall predictive accuracy of a random forest, our preferred model was a generalized random forest to understand sources of heterogeneity in outcomes associated with the powdery mildew status of yards.

Generalized random forest

The effective *F* statistics of the Montiel-Pflueger robust weak instrument test were 50.080 and 38.404 for annual cost and pesticide active constituents, respectively, at a 0.05 significance level (Table 3). For both response variables, the *F* statistic was greater than the worst-case bias of two-stage least square and limited information maximum likelihood. Thus, the instruments were not weak and were robust to heterogeneity in the data.

The distribution of quartiles of the CATE for active constituents and costs varied for all of the covariate groups and were particularly divergent in some case (Supplementary Table S2). For instance, for annual cost, only 15.11% of the CATE fell in the southeast quadrant in the lowest quantile, whereas 66.30% of the CATE fell in this quantile when the hop yards were located in the northeast quadrant. We also detected relatively large differences in the CATE distribution depending on cultivar susceptibility to V6 strains of *P. macularis* and pruning thoroughness. Qualitatively similar dif-

TABLE 1. Summary statistics and mean comparisons between hop yards where powdery mildew was detected or not for selected variables used in generalized random forest models

Variable[a]	Full sample		Powdery mildew not detected		Powdery mildew detected					Variable importance[b]	
	Mean	SD	Mean	SD	Mean	SD	Difference	*t* test	*P* value	Annual cost	Active constituents
Mean incidence of diseased plants (proportion)	0.03	0.08	0.00	0.00	0.08	0.12	0.08	10.29	<0.001	–	–
Mean incidence of diseased plants in May (proportion)	0.0032	0.0254	0.00	0.00	0.01	0.04	0.01	3.06	0.002	–	–
Pruning thoroughness (1 to 5)	2.98	1.53	2.99	1.65	2.98	1.38	–0.01	–0.04	0.963	0.186	0.180
Susceptibility to non-V6-virulent strains (0 to 5)	1.84	1.49	1.92	1.44	1.76	1.57	–0.16	–1.08	0.276	0.120	0.152
Susceptibility to V6-virulent strains (0 to 5)	2.73	1.31	2.35	1.36	3.21	1.10	0.86	7.35	<0.001	0.092	0.087
Degree centrality non-R6-cultivars May–June (count)	1.02	3.46	0.00	0.00	2.26	4.88	2.26	7.35	<0.001	0.124	0.063
Degree centrality non-R6-cultivars June–July (count)	0.49	1.88	0.00	0.00	1.10	2.69	1.10	6.51	<0.001	0.073	0.071
Initial strain V6-virulent (0/1/2)[c]	0.72	0.91	0.00	0.00	1.61	0.64	1.61	39.61	<0.001	0.029	0.032
Active constituents (count)	6.63	2.97	5.75	2.85	7.70	2.78	1.95	7.39	<0.001	–	–
Annual cost (USD/ha)	\$444.38	275.9	\$384.95	276.89	\$516.44	257.57	\$131.49	5.21	<0.001	–	–
Observations	458		251		207						

[a] Variable descriptions are as described in the Materials and Methods. The form of the data or units is noted parenthetically. Dummy variables associated with the quadrant of the landscape where yards were located are not presented because the summary statistics are not meaningful.

[b] Variable importance for selected variables was calculated as the weighted sum of how many times a feature *i* was split on at each depth in the forest. Dash marks indicate variable importance is not relevant.

[c] Means for this variable indicate the proportion of all yards where the initial strain of *Podosphaera macularis* was detected possessed V6-virulence. In the total of 207 yards where powdery mildew occurred, in 22% of yards, the initial strain was non-V6-virulent, and in 78%, the initial strain was V6-virulent.

ferences in the CATE distribution were observed for both pesticide active constituents and annual costs (Supplementary Table S2).

Variables identified as important for the annual costs of pesticides in the causal forest were, in descending order, spring pruning thoroughness (0.186), degree centrality of yards planted to non-R6-cultivars during the May–June transition (0.124), susceptibility to non-V6-virulent strains of the pathogen (0.120), susceptibility to V6-virulent strains of the pathogen (0.092), degree centrality of yards planted to non-R6-cultivars during the June–July transition (0.073), and the initial strain of the pathogen detected (0.029) (Table 1). Similarly, the most important variables associated with the number of active constituents applied were spring pruning thoroughness (0.180) and susceptibility to non-V6-virulent strains of the pathogen (0.152). Susceptibility to V6-virulent strains of the pathogen (0.087), network centrality of non-R6-cultivars during the June–July transition (0.071), network centrality of non-R6-cultivars during the May–June transition (0.063), and the initial strain of the pathogen detected (0.032) had varying importance (Table 1). The variables related to pesticide use intensity and annual costs in the generalized random forest were interrelated and correlated to varying degrees with the response variables (Fig. 3).

The estimated CATE calculated by the covariate subgroups confirmed the differences observed in the CATE distributions (Table 4). The number of active constituents applied and their associated costs depended on multiple factors. When the most thorough pruning was applied (ordinal value = 1), growers applied 6.79 fewer active constituents in response to powdery mildew that resulted in $81.73 lower costs than when more than 80% of plants had green leaves and shoots remaining (ordinal value = 5). The initial strain of *P. macularis* also impacted pesticide use intensity and cost differentials in yards with or without powdery mildew. Hop yards with a strain that was non-V6-virulent applied fewer additional active constituents (0.08; 95% CI [–57.05 to 57.2]) and incurred lesser additional costs ($175.87; 95% CI [$–1,673.29 to $2,025.03]) in response to powdery mildew occurrence compared with yards where the initial strain was V6-virulent (16.92; 95% CI [6.7 to 27.15] active constituents; $766.15 [$–691.53 to $2,223.85] in annual costs).

There was also a large degree of heterogeneity in the CATE for pesticide use intensity and annual costs that depended on the specific grower, indicating that individual growers varied their management to much different degrees in response to the occurrence of powdery mildew. The estimated CATE for annual costs ranged from $–1,093.72 (95% CI [$–3,126.28 to $938.85]) with Grower 3 to $2,407.88 (95% CI [$1,213.4 to $3,602.37]) with Grower 8, largely mirroring heterogeneity in their associated pesticide use (Table 4).

The seasonal differences in the CATE from 2014 to 2017 indicate that growers' management became more sensitive to whether powdery mildew was present or not (Table 4). In 2014, the differential in pesticide active constituents in yards with or without powdery mildew was 8.82 (95% CI [–16.28 to 33.93]), but this differential was 18.13 (95% CI [8.1 to 28.16]) in 2017 (Table 4). Correspondingly, the CATE for annual costs for pesticides increased from $112.80 (95% CI [$–1,134.62 to $1,360.24]) in 2014 to $882.46 (95% CI [$–595.94 to $2,360.88]) in 2017.

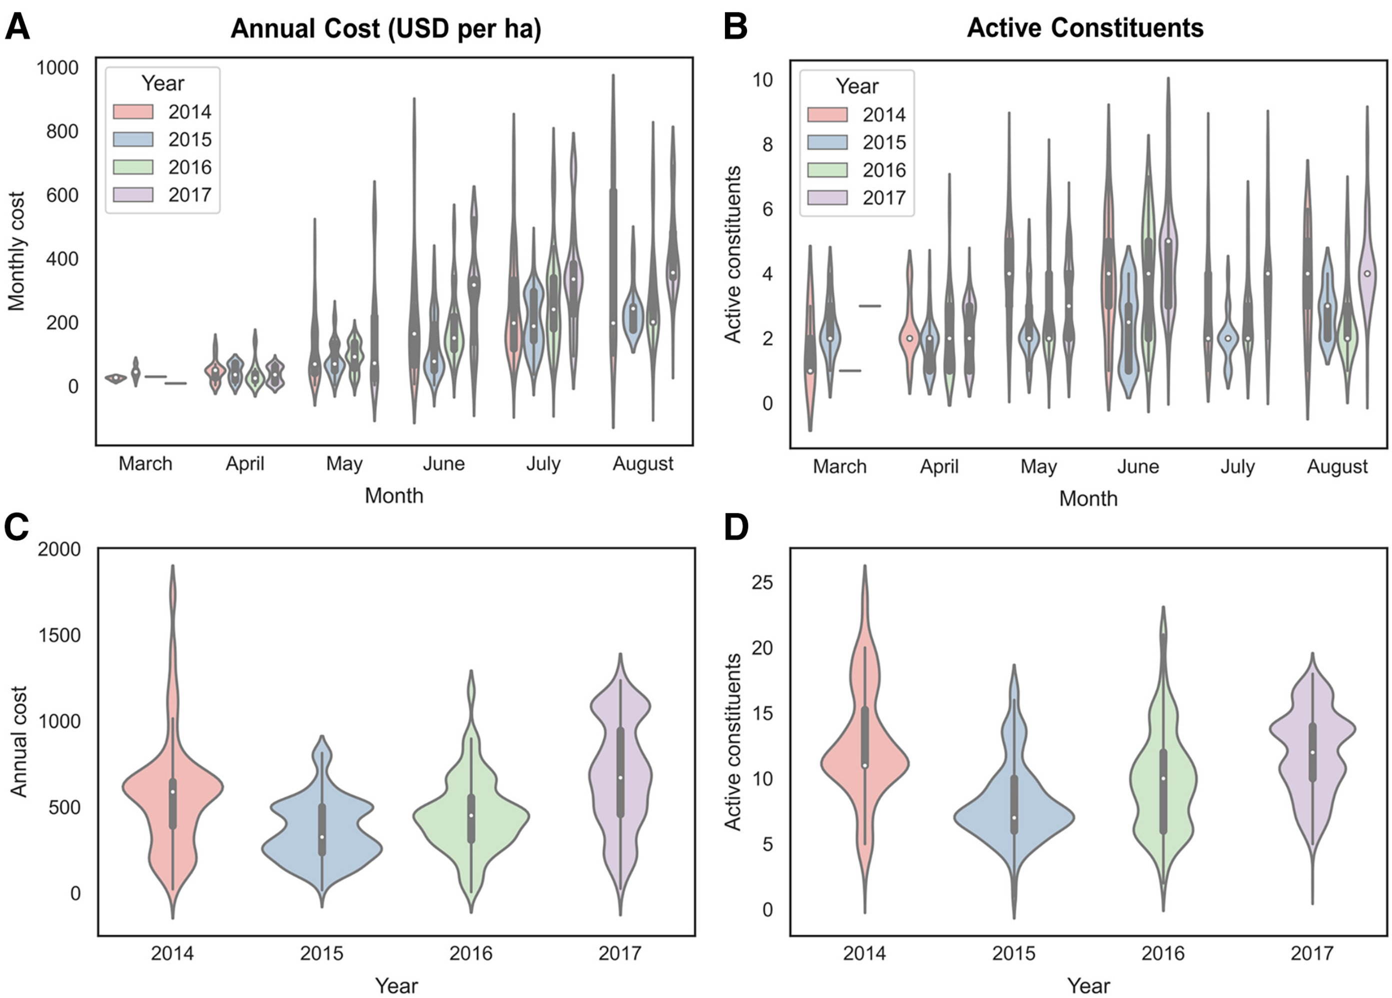

Fig. 2. Violin plots summarizing the distribution of **A and B,** monthly or **C and D,** annual costs of pesticides and the number of pesticide active constituents applied for suppression of powdery mildew by Oregon hop growers from 2014 to 2017.

Next, we explored the distribution of predicted treatment effects for the out-of-bag samples generated from the causal forest for covariates with the greatest feature importance, spring pruning thoroughness, and cultivar susceptibility to non-V6-virulent strains of *P. macularis* (Fig. 4). The central tendency of the predicted treatment effects for spring pruning thoroughness were qualitatively similar for annual costs (Fig. 4A). However, the distribution of the number of active constituents applied was more variable, and likely to be larger, in response to the occurrence of powdery mildew as the thoroughness of pruning diminished from a rating of an ordinal value of 1 to 5 (Fig. 4B). The susceptibility of a cultivar to non-V6-virulent strains of *P. macularis* had more divergence on the predicted number of active constituents that were applied when powdery mildew was present (Fig. 4D). In cultivars resistant to these strains, growers incurred lesser costs and applied fewer fungicides in response to powdery mildew as compared with more susceptible cultivars. For cultivars rated as having powdery mildew susceptibility of 1 to 4, the annual costs they incurred in response to powdery mildew were generally similar (Fig. 4C), whereas the number of active constituents applied tended to decrease with increasing susceptibility (Fig. 4D). That is, growers applied more expensive fungicides, but not necessarily more total active constituents, in cultivars more susceptible to non-V6-virulent strains of the pathogen when powdery mildew was detected. Predicted treatment effects for other variables are given in Supplementary Figure S2.

Covariate balancing and exposure-response function

Figure 5 summarizes the covariate balance test, which shows the absolute correlations for each covariate in the original unadjusted data set (blue line) and the matched data set (orange line). The median absolute correlation was 0.153 after weighting by the generalized propensity score, only slightly above the nominal threshold

TABLE 3. Montiel-Pflueger robust weak instrument test for the annual cost model and the active constituents model[a]

Statistic	Annual cost		Active constituents	
Effective F statistic	50.080		38.404	
% of worst-case bias	TSLS	LIML	TSLS	LIML
$\tau = 5\%$	37.418	37.418	37.418	37.418
$\tau = 10\%$	23.109	23.109	23.109	23.109

[a] TSLS: two-stage least square. LIML: limited information maximum likelihood.

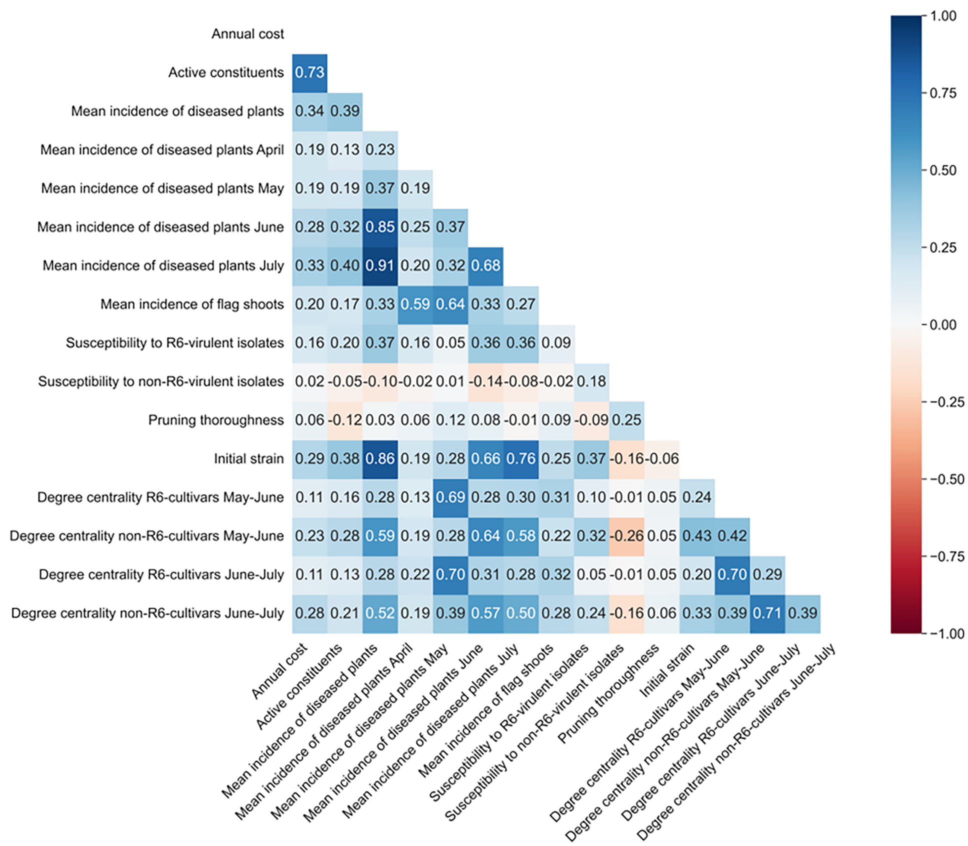

Fig. 3. Spearman's rank correlation coefficient matrix of selected important variables related to the number of active constituents applied for management of hop powdery mildew and the annual cost of pesticides.

TABLE 2. Predictive accuracy of selected machine learning regression models used to estimate the number of active constituents applied by hop growers in Oregon for powdery mildew (caused by *Podosphaera macularis*) and the annual cost of those active constituents[a]

	Annual cost						Active constituents					
Model	MAE	MSE	RMSE	R^2 training data	R^2 test data	Cross-validation	MAE	MSE	RMSE	R^2 training data	R^2 test data	Cross-validation
Ridge regression	199.47	263,037.2	512.87	0.42	–3.36	–1.47	1.54	4.99	2.23	0.89	0.36	0.53
LASSO regression	176.66	68,648.89	262.01	0.56	–0.14	–1.34	2.00	6.74	2.60	0.29	0.13	–0.05
Decision tree regression	204.49	78,739.20	280.61	0.78	–0.31	0.01	1.67	6.52	2.55	0.95	0.16	0.42
Random forest regression	166.28	52,766.63	229.71	0.72	0.13	0.24	1.41	3.79	1.95	0.91	0.51	0.64

[a] MAE is mean absolute error, MSE is mean squared error, RMSE is root mean squared error, and R^2 is the coefficient of determination.

value of 0.1 used to indicate that the confounding effect of the covariates is small (Zhu et al. 2015). We were unable to remove fully covariate imbalances for initial strain, susceptibility to V6 strains, and grower. Despite the imbalance in these specific covariates, the entire data set passed the covariate balance test (Fig. 5). The generalized propensity score densities for instances where powdery mildew was detected or not detected had their respective masses in overlapping regions (Fig. 6). Therefore, we verified the assumption of overlapping regions for the exposure-response function.

After assumptions of covariate balance and overlap were satisfied, we fit exposure-response function models for the annual costs of pesticides and the number of pesticide active constituents. A linear exposure-response relationship provided a good description of the average outcome value for the exposure level, with $y = 392 + 1{,}088.6x_i$ for annual cost and $y = 6.6 + 2.23x_i$ for the number of pesticide active constituents, where x_i is mean seasonal disease incidence. Therefore, growers' pesticide use and their associated costs scaled linearly with the incidence of powdery mildew.

Discussion

This research makes several contributions. We have identified covariates that predict and explain the number of active constituents applied and its economic costs for powdery mildew management among Oregon hop growers and causes of the observed heterogeneity. These factors include the cultural practice of spring pruning, cultivar susceptibility to two pathogenic races of *P. macularis*, position in the landscape, and centrality in the disease-spreading network. These factors collectively influence the incidence of powdery mildew, which has a direct exposure-response relationship to how growers respond in the overall number of pesticide applications they make and the annual costs of the pesticides they apply.

We were chiefly motivated to identify production practices and factors that explain why some growers may apply more or less pesticides than average and, in turn, incur varying costs. There is increasing scrutiny of the environmental and human health impacts of agriculture in general and pesticide use in particular (Verweij et al. 2009) and explicit policy objectives to reduce pesticide use (Donley 2019; Lamichhane et al. 2016; Skevas et al. 2013). It is of interest to note the timing of when pesticides were applied for powdery mildew and the seasonality of their costs. Growers applied the most pesticide active constituents during May and June, but months shouldering this period received notably fewer active constituents (Fig. 2). The seasonality of pesticide use largely reflects the development of the crop. Hop shoots begin to emerge following the boreal

TABLE 4. Conditional average treatment effect (CATE) estimated by subgroups for important variables used in generalized random forest models for predicting the number of pesticide active constituents applied for suppression of powdery mildew and the annual costs incurred by Oregon hop growers during 2014 to 2017[a]

	Annual cost (USD/ha)			Active constituents		
Variables	CATE (95% CI)	*F* statistic	*P* value	CATE (95% CI)	*F* statistic	*P* value
Pruning thoroughness						
1	471.22 (–184.39–1,126.83)	1.77	0.18	15.2 (8.49–21.91)	26.92	<0.001
2	–15.47 (–996.49–965.53)			15.07 (6.59–23.54)		
3	1,018.78 (–909.65–2,947.21)			17.59 (5.14–30.03)		
4	1,371.96 (–476.83–3,220.76)			2.86 (–10.64–16.37)		
5	552.95 (–133.02–1,238.93)			21.99 (1.12–42.86)		
Susceptibility to V6 strains						
0	2,183.03 (694.09–3,672.07)	29.27	<0.001	27.51 (15.55–39.48)	62.89	<0.001
1	317.49 (–286.71–921.69)			18.42 (11.25–25.58)		
2	836.27 (–932.54–2,605.1)			21.9 (4.68–39.12)		
3	577.23 (–41.26–1,195.74)			17.73 (10.07–25.39)		
4	659.97 (–613.56–1,933.51)			10.19 (–7.5–27.87)		
5	–1,356.82 (–9,491.92–6,778.28)			21.95 (–28.19–72.09)		
Susceptibility to non-V6 strains						
0	576.56 (–420.82–1,573.96)	5.84	0.01	13.08 (5.7–20.47)	75.15	<0.001
1	373.67 (–226.83–974.18)			18.46 (11.22–25.7)		
2	1,187.12 (–612.39–2,986.62)			20.21 (–3.52–43.95)		
3	577.23 (–41.26–1,195.74)			17.73 (10.07–25.39)		
4	914.76 (–1,627.76–3,457.28)			14.98 (–22.93–52.89)		
Initial strain						
Non-V6-virulent	175.87 (–1,673.29–2,025.03)			0.08 (–57.05–57.2)		
V6-virulent	766.15 (–691.53–2,223.85)			16.92 (6.7–27.15)		
Quadrant						
Northeast	916.02 (–1,232.34–3,064.4)	75.08	<0.001	4.22 (–19.44–27.89)	129.95	<0.001
Northwest	624.85 (49.54–1,200.16)			15.2 (7.35–23.05)		
Southeast	288.86 (–405.93–983.66)			24.4 (11.52–37.27)		
Grower						
1	456.83 (–406.23–1,319.92)	23.52	<0.001	26.59 (5.4–47.79)	44.27	<0.001
2	183.06 (–4,138.13–4,504.26)			1.96 (–45.01–48.94)		
3	–1,093.72 (–3,126.28–938.85)			0.42 (–16.65–17.48)		
4	829.69 (178.88–1,480.5)			19.13 (2.83–35.44)		
5	2,159.56 (1,065.97–3,253.15)			7.29 (–6.07–20.65)		
6	753.86 (–25.45–1,533.19)			25.9 (18.17–33.63)		
7	76.29 (-654.89 – 807.48)			16.16 (10.11–22.21)		
8	2,407.88 (1,213.4–3,602.37)			39.66 (24.93–54.39)		
9	504.92 (–781.29–1,791.13)			3.14 (–9.15–15.44)		
Year						
2014	112.80 (–1,134.62–1,360.24)	107.7	<0.001	8.82 (–16.28–33.93)	32.60	<0.001
2015	1,016.6 (582.90–1,450.3)			21.11 (15.67–26.55)		
2016	353.26 (–367.53–1,074.07)			15.9 (7.12–24.68)		
2017	882.46 (–595.94–2,360.88)			18.13 (8.1–28.16)		

[a] Conditional average treatment effect is the difference in response variables, annual cost, or the number of active constituents, conditional on whether a hop yard had powdery mildew or not (the treatment variable) given a set of covariates. CI, confidence interval.

winter dormancy in March to April and may grow rapidly (up to 15 cm per day) from emergence to bloom just after the summer solstice (Neve 1991). Hop growth habit is determinate, and vertical development of shoots and expansion of lateral branches largely ceases at bloom. Bloom and the juvenile stages of cone development are recognized as key periods for powdery mildew management (Gent et al. 2017; Royle 1978), and it was somewhat surprising that pesticide use tended to decline in July. This might reflect in part the limited efficacy of fungicide applications made relatively late in the season for reducing crop damage from powdery mildew (Gent et al. 2014, 2016). Whereas the number of active constituents applied in July and August decrease, the costs growers incur did not, pointing to fewer but more expensive fungicides being used at this time. This likely reflects the shift from sulfur-based fungicides to more potent and expensive synthetic fungicides, as is common in hop production because of considerations of organoleptic properties of the harvested hops, efficacy, and side effects of sulfur on arthropod pests (Nelson et al. 2015; Woods et al. 2012). Although pesticide use and annual costs are closely correlated, different portions of the cropping cycle could be considered when targeting strategic reductions in pesticide use versus costs.

The causes of pesticide use intensity are multifaceted and related to production practices that growers may have some control over but also other factors that are impossible to alter, such as market demand for certain cultivars or simply where yards are located. Our analyses point to several potential strategies for reducing pesticide use for powdery mildew on hop. The most important variable used for splitting leaves in the causal forest was spring pruning thoroughness (Table 1). Thorough spring pruning has long been advocated as a sanitation measure for reducing primary inoculum of *P. macularis* and delaying epidemic onset (Gent et al. 2008, 2019b; Probst et al. 2016; Royle 1978; Turechek et al. 2001). The novelty of the current study was not in suggesting that thorough spring pruning is important. Rather, we have explicitly linked pruning thoroughness to pesticide use intensity and costs and estimated these effects. In response to the presence of powdery mildew, growers generally increased pesticide use intensity the most as the thoroughness of their pruning diminished from the most to the least thorough categories. These differences were reflected in the conditional average treatment effects (Table 4) and demonstrate that growers who practice the most thorough pruning apply on average the least number of pesticide active constituents. The added costs of conducting more

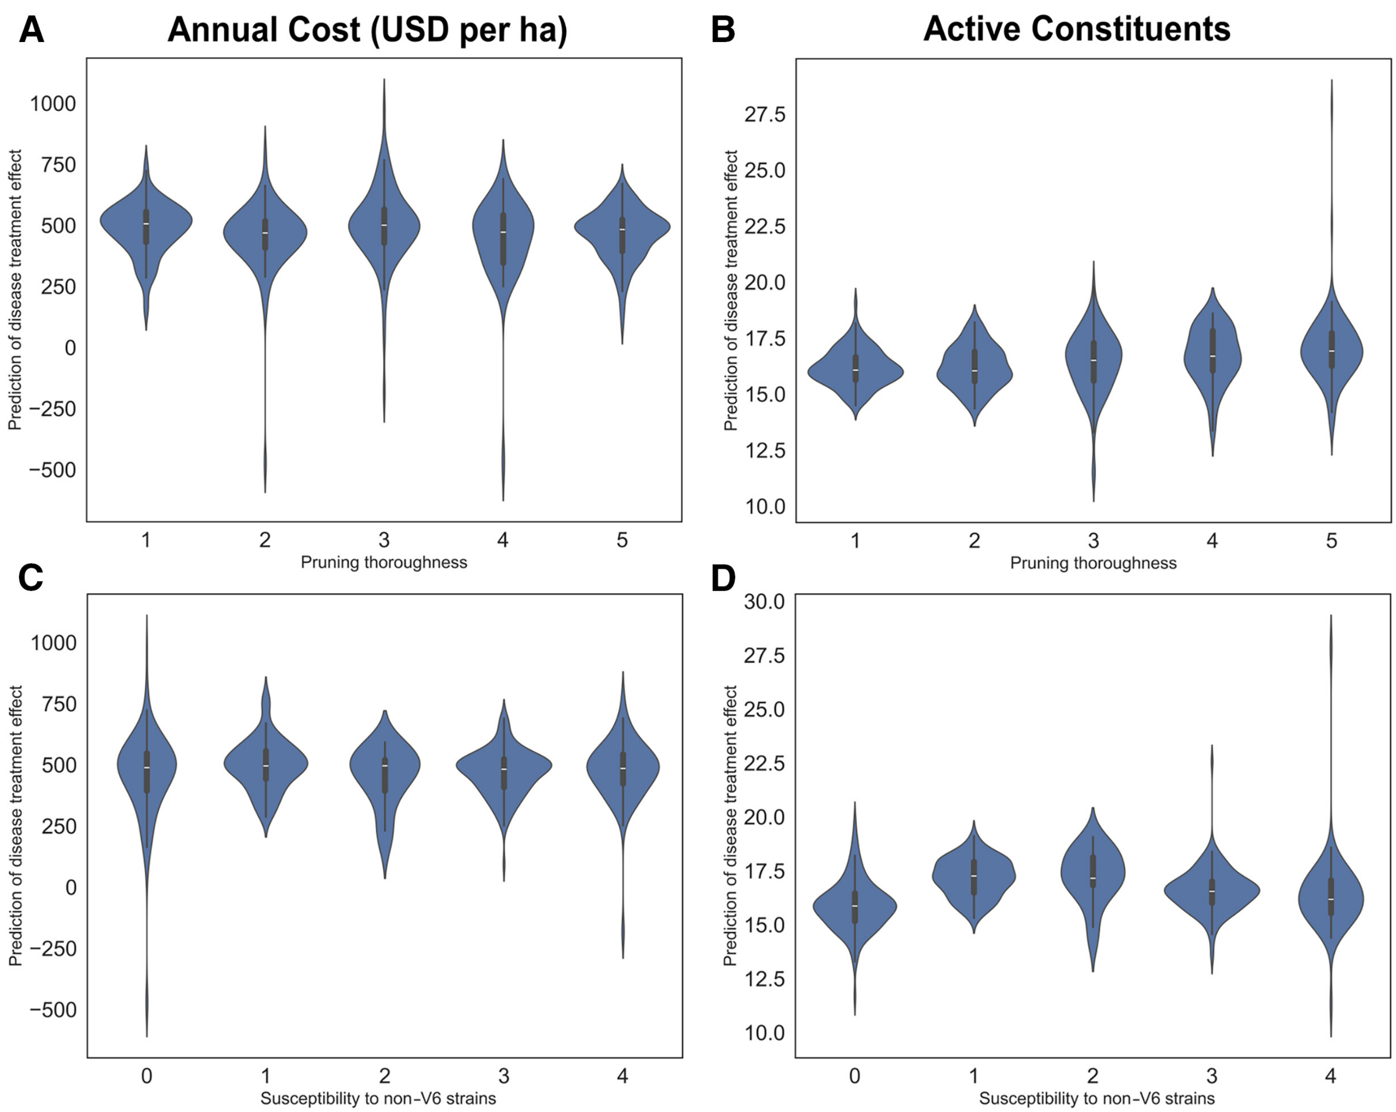

Fig. 4. Predicted treatment effects estimated from a causal forest plotted as a violin plot for two variables identified as most important in splitting: **A and B,** spring pruning thoroughness and **C and D,** susceptibility to non-V6-virulent strains of *Podosphaera macularis*. Annual costs per ha are presented in A and C, and pesticide active constituents are presented in B and D.

thorough spring pruning (Probst et al. 2016) are at least partially offset by savings due to use of fewer active constituent and less expensive ones.

Cultivar susceptibility to one of the dominant strains of *P. macularis* in the region was the next most important variable identified in the causal forest for pesticide use (Table 1). It is of interest to note that cultivar susceptibility to non-V6-virulent strains of the fungus was a more important variable than susceptibility to V6-virulent strains in the causal forest. These variables are closely related, as susceptibility to each strain is similar or identical for most cultivars with the exception of those that possess R6 (Wolfenbarger et al. 2016), which were expressed as 0 on the ordinal scale for disease susceptibility. Excluding this class of cultivars, the conditional average treatment effects are qualitatively similar for both pesticide active constituents and annual costs, indicating increasingly extreme management responses by growers on all cultivars having a non-zero level of susceptibility to either strain (Table 4). Our data also indicate that growers apply the fewest active constituents and incur the lowest costs in response to powdery mildew on cultivars that are more resistant (ordinal values 0 or 1). However, growers appear to respond to the disease by applying different mixtures or combinations of pesticide active constituents on cultivars with varying levels of susceptibility. This is evidenced by the approximately similar annual costs of pesticides applied, even though the number of active constituents varied (Fig. 4C and D). On moderately susceptible cultivars, growers tend to apply more active constituents that cost less per unit but with the net effect that this does not necessarily reduce their annual costs. A critical examination of pesticide programs on cultivars with intermediate susceptibility is warranted because the predicted costs of disease management on these cultivars might be unnecessarily high relative to more susceptible cultivars. We expected that cultivars with the highest levels of resistance to powdery mildew would have substantially smaller average treatment responses when powdery mildew is present for both pesticide use and cost. However, our data suggest that growers only marginally reduced spraying on such cultivars when the disease was present. This again suggests potential for more cultivar-specific management of powdery mildew that could reduce pesticide use and costs. Lybbert et al. (2016) noted that winegrape growers in California acted to minimize the risk of crop damage from powdery mildew in unexpected ways in response to disease forecasts. Certain winegrape growers applied a baseline level of fungicides independent of forecasted disease risk but increasingly applied mixtures of fungicides and more potent fungicides as forecasted disease risk increased. The present analysis suggests that hop growers could have an analogous management goal of mitigating the risk of powdery mildew on cultivars with partial or intermediate levels of resistance. That is, on cultivars with intermediate levels, growers may apply pesticides disproportionately relative to the risk of crop damage.

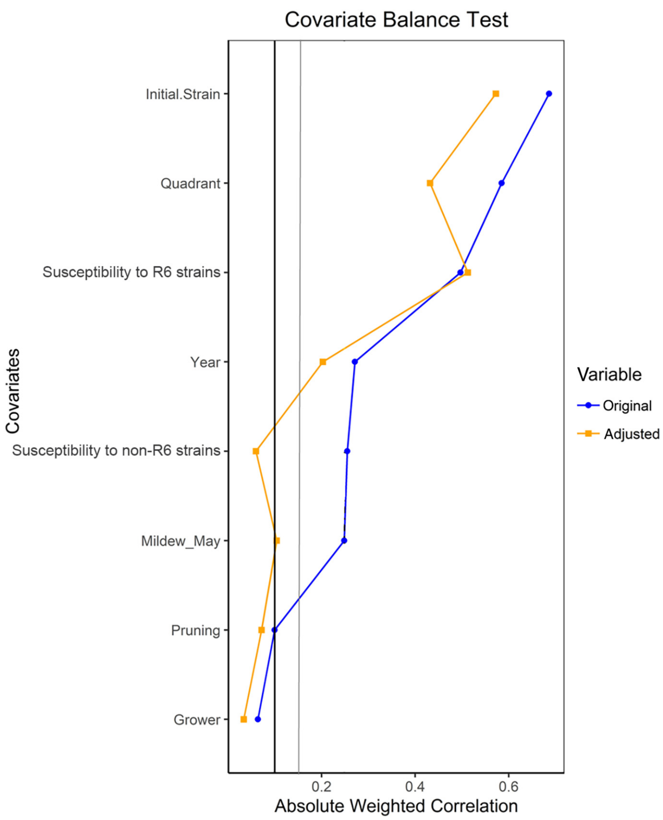

Fig. 5. Covariate balance plot with covariates sorted in descending order according to absolute correlation (imbalance) in the original data set. The gray vertical line represents the median absolute correlation value for the matched data, and the black vertical line is a nominal covariate balance threshold of 0.1. Absolute correlations are between the exposure (mean seasonal disease incidence) and observed pre-exposure covariates.

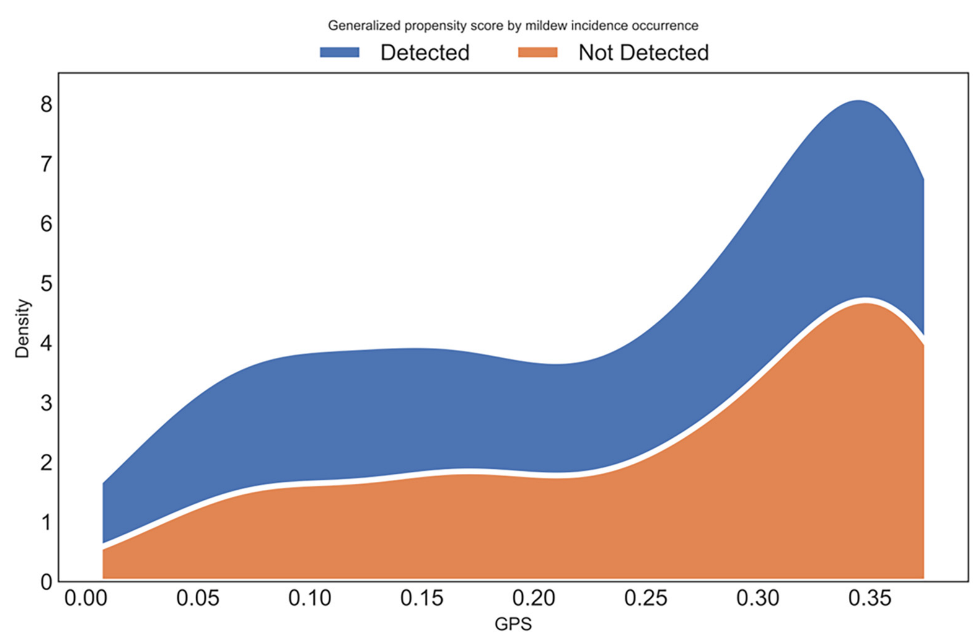

Fig. 6. Generalized propensity scores (GPS) based on the weighted covariates for hop yards where powdery mildew was not detected ($W = 0$) or was detected ($W = 1$). The generalized propensity score densities have the respective masses in regions of overlap, indicating covariates were successfully balanced.

General position in the landscape and centrality in the disease transmission network of hop yards also influenced pesticide use in response to powdery mildew. These factors are correlates of overall disease pressure, which we demonstrate has a linear relationship with pesticide use through the exposure-response modeling. The probability of disease transmission from a source yard, expressed as edge weights in the network, is a product of pathogen source strength, wind run, and distance to another yard (Gent et al. 2019a). Of these factors, management intervention can reduce only source strength. Source strength could be weakened by reducing the initial inoculum or increasing pesticide use. Another possible means of reducing source strength is strategic deployment of host resistance in a subset of yards or farms whose characteristics and position in the landscape would lead to high network centrality should disease occur. Host diversification at multiple spatial scales has long been recognized as a means of disease suppression for human, animal, and plant diseases (Boudreau 2013; Burdon et al. 2014; Keeling 1999; Plantegenest et al. 2007; Rimbaud et al. 2021). Where host resistance is most needed depends on contact structure in the early stages of an epidemic (Jeger et al. 2007), which is determined by dispersal characteristics of the pathogen, spatial arrangement of the host and pathogen, and physical factors such as wind that drive dispersal (Mikaberidze et al. 2016; Mundt and Brophy 1988). When highly connected individuals are known, immunization or deployment of host resistance is optimally targeted to these nodes in scale-free networks (Jeger et al. 2007; Pastor-Satorras and Vespignani 2002; Shaw and Pautasso 2014). Host resistance is available in hop cultivars that is effective against the dominant pathogenic races of *P. macularis* in the Western United States (Gent et al. 2017; Wolfenbarger et al. 2014, 2016). Future studies are warranted to understand where such resistance would be best placed to have maximal impact on disease suppression and pesticide use reduction.

Beyond the immediate applications to the motivating pathosystem, our present study also serves as a case example of how statistical methods for causal inference can be better utilized in plant disease epidemiology contexts. As experiments scale from local scales, such as plots, to landscapes (Meentemeyer et al. 2012), many of the methods we access herein are well developed and could serve to move research from simple description and prediction to causal inference.

Acknowledgments

We thank the many individuals who provided technical support for this project and to participating grower cooperators.

Literature Cited

Agrios, G. N. 2005. Plant Pathology. Elsevier, New York.

Andert, S., Bürger, J., and Gerowitt, B. 2015. On-farm pesticide use in four Northern German regions as influenced by farm and production conditions. Crop Prot. 75:1-10.

Atallah, S. S., Gomez, M. I., Conrad, J. M., and Nyrop, J. P. 2012. An agent-based model of plant disease diffusion and control: Grapevine leafroll disease. In: Proceedings of Agricultural & Applied Economics Association's 2012 AAEA Annual Meeting, Seattle, Washington.

Athey, S., and Imbens, G. 2016. Recursive partitioning for heterogeneous causal effects. Proc. Natl. Acad. Sci. U.S.A. 113:7353-7360.

Athey, S., and Imbens, G. W. 2019. Machine learning methods that economists should know about. Annu. Rev. Econ. 11:685-725.

Athey, S., Tibshirani, J., and Wager, S. 2019. Generalized random forests. Ann. Stat. 47:1148-1178.

Athey, S., and Wager, S. 2019. Estimating treatment effects with causal forests: An application. Obs. Stud. 5:37-51.

Austin, P. C. 2011. An introduction to propensity score methods for reducing the effects of confounding in observational studies. Multivar. Behav. Res. 46:399-424.

Babcock, B., McRoberts, N., and Figuera, G. S. 2022. Efficacy of coordinated area-wide treatments to control HLB. Citrograph 13:38-43.

Block, M., Wiseman, M. S., and Gent, D. H. 2021. Characterization of *Podosphaera macularis* derived from the hop cultivar 'Strata' and Strata's resistance to powdery mildew in Oregon. Plant Health Prog. 22: 154-156.

Boudreau, M. A. 2013. Diseases in intercropping systems. Annu. Rev. Phytopathol. 51:499-519.

Breiman, L. 2001. Random Forests. Mach. Learn. 45:5-32.

Burdon, J. J., Barrett, L. G., Rebetzke, G., and Thrall, P. H. 2014. Guiding deployment of resistance in cereals using evolutionary principles. Evol. Appl. 7:609-624.

Chernozhukov, V., Newey, W. K., Quintas-Martinez, V., and Syrgkanis, V. 2022. RieszNet and ForestRiesz: Automatic debiased machine learning with neural nets and random forests. arXiv 2110.03031. http://arxiv.org/abs/2110.03031 (accessed December 19, 2023).

Claassen, B. J., Wolfenbarger, S. N., and Gent, D. H. 2022. Fungicide physical mode of action: Impacts on suppression of hop powdery mildew. Plant Dis. 106:1244-1252.

Cunniffe, N. J., Cobb, R. C., Meentemeyer, R. K., Rizzo, D. M., and Gilligan, C. A. 2016. Modeling when, where, and how to manage a forest epidemic, motivated by sudden oak death in California. Proc. Natl. Acad. Sci. U.S.A. 113:5640-5645.

Domingues, T., Brandão, T., and Ferreira, J. C. 2022. Machine learning for detection and prediction of crop diseases and pests: A comprehensive survey. Agriculture 12:1350.

Donley, N. 2019. The USA lags behind other agricultural nations in banning harmful pesticides. Environ. Health 18:44.

Essling, M., McKay, S., and Petrie, P. R. 2021. Fungicide programs used to manage powdery mildew (*Erysiphe necator*) in Australian vineyards. Crop Prot. 139:105369.

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., Block, M., Massie, S. T., Phillips, C. L., Richardson, B. J., Shellhammer, T. H., Trippe, K. M., and Wiseman, M. S. 2024. Nitrogen and sulfur fertility practices: Influences on hop chemistry, aroma, and nitrate accumulation. J. Am. Soc. Brew. Chem. 82:50-60.

Gent, D. H., Claassen, B. J., Twomey, M. C., Wolfenbarger, S. N., and Woods, J. L. 2018. Susceptibility of hop crown buds to powdery mildew and its relation to perennation of *Podosphaera macularis*. Plant Dis. 102: 1316-1325.

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., 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., George, A. E., Grove, G. G., Mahaffee, W. F., Ocamb, C. M., Barbour, J. D., Peetz, A., and Turechek, W. W. 2008. A decade of hop powdery mildew in the Pacific Northwest. Plant Health Prog. 9. https://doi.org/10.1094/PHP-2008-0314-01-RV

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.

Greene, W. H. 2018. Econometric Analysis. 8th ed. Pearson, New York.

Hariton, E., and Locascio, J. J. 2018. Randomised controlled trials – the gold standard for effectiveness research. BJOG Int. J. Obstet. Gynaecol. 125: 1716.

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.

Jeger, M. J., Pautasso, M., Holdenrieder, O., and Shaw, M. W. 2007. Modelling disease spread and control in networks: Implications for plant sciences. New Phytol. 174:279-297.

Jibrin, M. O., Liu, Q., Jones, J. B., and Zhang, S. 2021. Surfactants in plant disease management: A brief review and case studies. Plant Pathol. 70:495-510.

Jørgensen, L. N., van den Bosch, F., Oliver, R. P., Heick, T. M., and Paveley, N. D. 2017. Targeting fungicide inputs according to need. Annu. Rev. Phytopathol. 55:181-203.

Keeling, M. J. 1999. The effects of local spatial structure on epidemiological invasions. Proc. R. Soc. Lond. B Biol. Sci. 266:859-867.

Khoshnevis, N., Wu, X., and Braun, D. 2023. CausalGPS: An R package for causal inference with continuous exposures. Papers 2310.00561. https://ideas.repec.org//p/arx/papers/2310.00561.html (accessed November 6, 2023).

Lamichhane, J. R., Dachbrodt-Saaydeh, S., Kudsk, P., and Messéan, A. 2016. Toward a reduced reliance on conventional pesticides in European agriculture. Plant Dis. 100:10-24.

Laurie, R. W., Richardson, B. J., Ross, C. J., and Gent, D. H. 2023. Yard age, cultivar susceptibility, and spring pruning practices as risk factors for overwintering of *Podosphaera macularis* on hop. PhytoFrontiers 3:390-398.

Little, R. J., and Rubin, D. B. 2000. Causal effects in clinical and epidemiological studies via potential outcomes: Concepts and analytical appraoches. Annu. Rev. Public Health 21:121-145.

Lybbert, T. J., Magnan, N., and Gubler, W. D. 2016. Multidimensional responses to disease information: How do winegrape growers react to powdery mildew forecasts and to what environmental effect? Am. J. Agric. Econ. 98: 383-405.

Mahaffee, W. F., Engelhard, B., Gent, D. H., and Grove, G. G. 2009. Powdery mildew. Pages 25-31 in: Compendium of Hop Diseases and Pests. W. F. Mahaffee, S. J. Pethybridge, and D. H. Gent, eds. American Phytopathological Society, St. Paul, MN.

Marsh, T. L., Huffaker, R. G., and Long, G. E. 2000. Optimal control of vector-virus-plant interactions: The case of potato leafroll virus net necrosis. Am. J. Agric. Econ. 82:556-569.

Meentemeyer, R. K., Haas, S. E., and Václavík, T. 2012. Landscape epidemiology of emerging infectious diseases in natural and human-altered ecosystems. Annu. Rev. Phytopathol. 50:379-402.

Mikaberidze, A., Mundt, C. C., and Bonhoeffer, S. 2016. Invasiveness of plant pathogens depends on the spatial scale of host distribution. Ecol. Appl. 26:1238-1248.

Mourtzinis, S., Krupke, C. H., Esker, P. D., Varenhorst, A., Arneson, N. J., Bradley, C. A., Byrne, A. M., Chilvers, M. I., Giesler, L. J., Herbert, A., Kandel, Y. R., Kazula, M. J., Hunt, C., Lindsey, L. E., Malone, S., Mueller, D. S., Naeve, S., Nafziger, E., Reisig, D. D., Ross, W. J., Rossman, D. R., Taylor, S., and Conley, S. P. 2019. Neonicotinoid seed treatments of soybean provide negligible benefits to US farmers. Sci. Rep. 9:11207.

Mourtzinis, S., Rattalino Edreira, J. I., Grassini, P., Roth, A. C., Casteel, S. N., Ciampitti, I. A., Kandel, H. J., Kyveryga, P. M., Licht, M. A., Lindsey, L. E., Mueller, D. S., Nafziger, E. D., Naeve, S. L., Stanley, J., Staton, M. J., and Conley, S. P. 2018. Sifting and winnowing: Analysis of farmer field data for soybean in the US North-Central region. Field Crops Res. 221:130-141.

Mundt, C. C., and Brophy, L. S. 1988. Influence of number of host genotype units on the effectiveness of host mixtures for disease control: A modeling approach. Phytopathology 78:1087-1094.

Murray-Watson, R. E., Hamelin, F. M., and Cunniffe, N. J. 2022. How growers make decisions impacts plant disease control. PLoS Comput. Biol. 18:e1010309.

Nelson, C. R., and Startz, R. 1990. The distribution of the instrumental variables estimator and its *t*-ratio when the instrument is a poor one. J. Bus. 63:S125-S140.

Nelson, M. E., Gent, D. H., and Grove, G. G. 2015. Meta-analysis reveals a critical period for management of powdery mildew on hop cones. Plant Dis. 99:632-640.

Neve, R. A. 1991. Hops. Chapman and Hall, London, U.K.

Nicholson, C. C., and Williams, N. M. 2021. Cropland heterogeneity drives frequency and intensity of pesticide use. Environ. Res. Lett. 16:074008.

Nie, X., and Wager, S. 2021. Quasi-oracle estimation of heterogeneous treatment effects. Biometrika 108:299-319.

Oakley, E., Zhang, M., and Miller, P. R. 2007. Mining pesticide use data to identify best management practices. Renew. Agric. Food Syst. 22:260-270.

Olea, J. L. M., and Pflueger, C. 2013. A robust test for weak instruments. J. Bus. Econ. Stat. 31:358-369.

Pastor-Satorras, R., and Vespignani, A. 2002. Immunization of complex networks. Phys. Rev. E 65:036104.

Pearl, J. 2010. An introduction to causal inference. Int. J. Biostat. 6:7.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. 2011. Scikit-learn: Machine learning in Python. J. Mach. Learn. Res. 12: 2825-2830.

Pichler, M., and Hartig, F. 2023. Machine learning and deep learning—A review for ecologists. Methods Ecol. Evol. 14:994-1016.

Plantegenest, M., Le May, C., and Fabre, F. 2007. Landscape epidemiology of plant diseases. J. R. Soc. Interface 4:963-972.

Probst, C., Nelson, M. E., Grove, G. G., Twomey, M. C., and Gent, D. H. 2016. Hop powdery mildew control through alteration of spring pruning practices. Plant Dis. 100:1599-1605.

Rimbaud, L., Fabre, F., Papaïx, J., Moury, B., Lannou, C., Barrett, L. G., and Thrall, P. H. 2021. Models of plant resistance deployment. Annu. Rev. Phytopathol. 59:125-152.

Rosenbaum, P. R., and Rubin, D. B. 1983. The central role of the propensity score in observational studies for causal effects. Biometrika 70:41-55.

Rothwell, P. M. 2005. External validity of randomised controlled trials: "To whom do the results of this trial apply?" Lancet 365:82-93.

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

Savary, S., Willocquet, L., Elazegui, F. A., Teng, P. S., Van Du, P., Zhu, D., Tang, Q., Huang, S., Lin, X., Singh, H. M., and Srivastava, R. K. 2000. Rice pest constraints in tropical Asia: Characterization of injury profiles in relation to production situations. Plant Dis. 84:341-356.

Shah, D. A., De Wolf, E. D., Paul, P. A., and Madden, L. V. 2023. Into the trees: Random forests for predicting Fusarium head blight epidemics of wheat in the United States. Phytopathology 113:1483-1493.

Shaw, M. W., and Pautasso, M. 2014. Networks and plant disease management: Concepts and applications. Annu. Rev. Phytopathol. 52:477-493.

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.

Skevas, T., Lansink, A. G. J. M. O., and Stefanou, S. E. 2013. Designing the emerging EU pesticide policy: A literature review. NJAS Wagening. J. Life Sci. 64-65:95-103.

Stock, J. H., and Yogo, M. 2005. Asymptotic distributions of instrumental variables statistics with many instruments. Pages 109-120 in: Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg. D. W. K. Andrews and J. H. Stock, eds. Cambridge University Press, New York.

Strobl, C., Boulesteix, A.-L., Zeileis, A., and Hothorn, T. 2007. Bias in random forest variable importance measures: Illustrations, sources and a solution. BMC Bioinform. 8:25.

Tibshirani, J., Athey, S., Friedberg, R., Hadad, V., Hirshberg, D., Miner, L., Sverdrup, E., Wager, S., and Wright, M., 2023. grf: Generalized Random Forests. https://cran.r-project.org/web/packages/grf/index.html (accessed September 26, 2023).

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. https://doi.org/10.1094/PHP-2001-0313-01-RS

Twomey, M. C., Wolfenbarger, S. N., Woods, J. L., and Gent, D. H. 2015. Development of partial ontogenic resistance to powdery mildew in hop cones and its management implications. PLoS One 10:e0120987.

Verweij, P. E., Snelders, E., Kema, G. H. J., Mellado, E., and Melchers, W. J. G. 2009. Azole resistance in *Aspergillus fumigatus*: A side-effect of environmental fungicide use? Lancet Infect. Dis. 9:789-795.

Weldon, W. A., Gent, D. H., and Gadoury, D. M. 2021a. Management of hop powdery mildew in the context of recent advances in pathogen ecology and population genetics. Plant Health Prog. 22:450-458.

Weldon, W. A., Knaus, B. J., Grünwald, N. J., Havill, J. S., Block, M. H., Gent, D. H., Cadle-Davidson, L. E., and Gadoury, D. M. 2021b. Transcriptome-derived amplicon sequencing markers elucidate the U.S. *Podosphaera macularis* population structure across feral and commercial plantings of *Humulus lupulus*. Phytopathology 111:194-203.

Wolfenbarger, S. N., Eck, E. B., and Gent, D. H. 2014. Characterization of resistance to powdery mildew in the hop cultivars Newport and Comet. Plant Health Prog. 15:55-56.

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.

Wolfenbarger, S. N., Twomey, M. C., Gadoury, D. M., Knaus, B. J., Grünwald, N. J., and Gent, D. H. 2015. Identification and distribution of mating-type idiomorphs in populations of *Podosphaera macularis* and development of chasmothecia of the fungus. Plant Pathol. 64:1094-1102.

Woods, J. L., Dreves, A. J., Fisher, G. C., James, D. G., Wright, L. C., and Gent, D. H. 2012. Population density and phenology of *Tetranychus urticae* (Acari: Tetranychidae) in hop is linked to the timing of sulfur applications. Environ. Entomol. 41:621-635.

Zhu, Y., Coffman, D. L., and Ghosh, D. 2015. A boosting algorithm for estimating generalized propensity scores with continuous treatments. J. Causal Inference 3:25-40.