Quantifying secondary pest outbreaks in cotton and their monetary cost with causal inference statistics

Kevin Gross[1,3] and Jay A. Rosenheim[2]

[1]Biomathematics Program, North Carolina State University, Raleigh NC 27695

[2]Department of Entomology, University of California, Davis, CA 95616

[3]Corresponding author, email: kevin_gross@ncsu.edu

Abstract

Secondary pest outbreaks occur when the use of a pesticide to reduce densities of an unwanted target pest species triggers subsequent outbreaks of other pest species. Although secondary pest outbreaks are thought to be familiar in agriculture, their rigorous documentation is made difficult by the challenges of performing randomized experiments at suitable scales. Here, we quantify the frequency and monetary cost of secondary pest outbreaks elicited by early-season applications of broad-spectrum insecticides to control the plant bug *Lygus* spp. (primarily *L. hesperus*) in cotton grown in the San Joaquin Valley, California, USA. We do so by analyzing pest-control management practices for 969 cotton fields spanning 9 years and 11 private ranches. Our analysis uses statistical methods to draw formal causal inferences from non-experimental data that have become popular in public health and economics, but which are not yet widely known in ecology or agriculture. We find that, in fields that received an early-season broad-spectrum insecticide treatment for *Lygus*, 20.2% (s.e. 4.4%) of late-season pesticide costs were attributable to secondary pest outbreaks elicited by the early-season insecticide application for *Lygus*. In 2010 US dollars, this equates to an additional US$6.0 (s.e. US$1.3) per acre in management costs. To the extent that secondary pest outbreaks may be driven by eliminating pests' natural enemies, these figures place a lower bound on the monetary value of ecosystem services provided by native communities of arthropod predators and parasitoids in this agricultural system.

Keywords: Causal inference, cotton, ecosystem services, indirect effects, integrated pest management, potential outcomes, secondary pest outbreak.

Introduction

Secondary pest outbreaks, in which the use of a pesticide to reduce densities of an unwanted target pest species triggers subsequent outbreaks of other pest species, are a well-known phenomenon in agriculture (Ripper, 1956; Hardin et al., 1995; Dutcher, 2007). Several mechanisms can drive secondary pest outbreaks, including reduction of natural enemies that suppress densities of non-target pests, physiological changes in the plant or non-target species (hormoligosis), and reductions in competing arthropod species (Ripper, 1956; White, 1984; Hardin et al., 1995). Secondary pest outbreaks can be detrimental to the welfare of the farmer, as they may reduce profit by reducing yield and by necessitating costly additional pesticide applications (Horton et al., 2005; Dutcher, 2007). Secondary pest outbreaks are also of interest from the perspective of 'ecosystem services', because quantifying the loss in profit attributable to secondary pest outbreaks may arguably provide a lower bound on the monetary value of the regulation of economically injurious pest species provided by communities of natural enemies.

While the existence of secondary pest outbreaks is uncontroversial, rigorous documentation of secondary pest outbreaks is difficult (Hardin et al., 1995; Dutcher, 2007). Experimental demonstration of secondary pest outbreaks is often stymied by considerations of scale, because well-replicated, controlled experiments are often (but not always) infeasible at the spatial and temporal scales at which the ecological mechanisms driving or preventing secondary pest outbreaks operate. Consequently, most evidence for secondary pest outbreaks comes from so-called observational data collected outside an experimental framework. With traditional analyses, observational data do not provide the conclusive evidence for causation that experimental data allow.

This article investigates secondary pest outbreaks elicited by management for the plant bug *Lygus* spp. in cotton grown in the San Joaquin Valley of California, USA. *Lygus* is a key pest of cotton in California and throughout the southwestern US (Leigh et al, 1988;

Leigh & Goodell, 1996). Management of *Lygus* in cotton is thought to provide a prime candidate for secondary pest outbreaks, because cotton harbors a rich community of arthropod herbivores and natural enemies, and because, until very recently, only non-selective, broad-spectrum pesticides have been available for *Lygus* control (Rao et al., 2003; Dutcher, 2007). Indeed, some of the most convincing experimental demonstrations of secondary pest outbreaks that do exist come from early studies of *Lygus* control in California cotton, where repeated and heavy applications of broad-spectrum insecticides to control *Lygus* elicited outbreaks of armyworms and other lepidopteran larvae (Falcon et al., 1968, 1971; Eveleens et al., 1973). More recently, informal observations have suggested suggest that, under contemporary management practices, early-season insecticide applications to control *Lygus* can also trigger secondary pest outbreaks of other herbivorous arthropods such as spider mites (*Tetranychus* spp.) (University of California, 1996).

Here, we investigate secondary pest outbreaks in California cotton with an 'ecoinformatics' approach. With the generous cooperation of four professional pest-control advisors (PCAs), we have assembled data detailing management practices in cotton fields operated by 11 different ranches from 1997–2008. We have assembled these data in hopes that the breadth of management strategies that they span will allow us to measure agriculturally meaningful effects at scales that are pertinent to contemporary agriculture. We analyze these data for secondary pest outbreaks using statistical methods for causal inference from observational data that have been developed in the context of public health and economics. To our knowledge, these causal inference methods have not yet penetrated the ecological or agricultural literature (Plowright et al., 2008).

Thus, this paper has two primary goals. Our first goal is to determine if secondary pest outbreaks are caused by the application of contemporary broad-spectrum insecticides for *Lygus* pests in California cotton, and if so, to quantify the monetary cost of managing those outbreaks. Our second goal is to introduce statistical methods for causal inference

from observational data that are not yet widely known among ecologists. The remainder of this paper is structured as follows. We first introduce the *Lygus*–cotton system in more depth, and provide details about the data that we have assembled. We then provide a brief introduction to causal-inference statistics, and offer citations for further reading. We then use these causal-inference methods to estimate the effect of early-season insecticide treatment for *Lygus* on the number and cost of late-season insecticide applications for non-*Lygus* pests. Readers uninterested in the causal-inference framework may bypass the mathematical sections without loss.

Lygus in cotton

Introduction to the system

Cotton pest management in California's San Joaquin Valley is predicated upon the judicious and sparing use of pesticides so as to maximize the pest management services contributed by an abundant and diverse community of natural enemies (University of California, 1996). The primary threats to cotton production due to herbivorous arthropods change over the course of the growing season. *Lygus* spp., predominantly *L. hesperus* Knight (Hemiptera: Miridae) but also occasionally *L. elisus* Van Duzee (Hemiptera: Miridae), damages cotton by feeding on young flower buds, potentially eliciting their abscission. This damage is of most concern early during the reproductive phase of cotton's growth (late May through June), when cotton's ability to compensate for loss of flower buds appears to be particularly weak (unpublished data; see also Musser et al. (2009) for a parallel result). Because this window of crop sensitivity to *Lygus* is relatively brief, a single application of insecticides often suffices to suppress *Lygus* below damaging levels until the plant attains a developmental stage with enhanced capacity to compensate for loss of flower buds. *L. hesperus* and *L. elisus* are usually not distinguished in pest management.

Later during the growing season, other pests can become more significant. Spider mites (*Tetranychus* spp. Dufour [Acari: Tetranychidae]) are especially important during the hottest months (July, August), when their populations can grow explosively. Armyworms (mostly *Spodoptera exigua* [Hübner] [Lepidoptera: Noctuidae]) and other lepidopteran larvae are also more likely to emerge as pests late in the growing season (July, August). Aphids (*Aphis gossypii* Glover [Hemiptera: Aphididae]) are primarily a concern late during the growing season as well (September-October), because their populations grow most rapidly under cooler fall temperatures and because their excreta ('honeydew') can contaminate cotton lint, which is exposed once mature cotton fruits ('bolls') start to open as harvest approaches.

Spider mite, armyworm, and cotton aphid populations are potentially regulated by a diverse community of natural enemies in cotton fields. Generalist predators, including *Orius* spp. Wolff (Hemiptera: Anthocoridae), *Geocoris* spp. Fallen (Hemiptera: Lygaeidae), *Nabis* spp. Latreille (Hemiptera: Nabidae), *Zelus* spp. Fabr. (Hemiptera: Reduviidae), a complex of ladybeetles (family Coccinellidae), and a complex of common green lacewings (family Chrysopidae) are consumers of each of these herbivores. In addition, each herbivore has a complex of more specialized predatory and parasitic exploiters: spider mites are attacked by specialist predators, including *Frankliniella occidentalis* (Pergande) (Thysanoptera: Thripidae), *Scolothrips sexmaculatus* (Pergande) (Thysanoptera: Thripidae), and a complex of predatory mites (family Phytoseiidae); armyworms are attacked by a complex of hymenopteran parasitoids; and aphids are attacked by the parasitoid *Lysiphlebus testaceipes* (Cresson) (Hymenoptera: Braconidae) and a complex of predatory hover flies (family Syrphidae) and midges (family Cecidomyiidae) (van den Bosch & Hagen, 1966; University of California, 1996). The use of broad-spectrum insecticides to control early-season *Lygus* populations may impose mortality on any or all members of this diverse community of natural enemies.

Data base assembly

We compiled data from four PCAs who manage cotton for private ranchers in California's San Joaquin Valley. Each PCA provided data for 1–4 unique ranches and 5–10 years. Overall, our data span 11 ranches, and for each ranch we have data for some subset of the years from 1997–2008.

Our data consist of scouts' reports, pesticide applications, including all insecticides and acaricides, targets for each pesticide application, and yields for multiple fields at each ranch. (Here and throughout, we use the term 'field' to refer to a single year's planting on a physical parcel of land, not as the land itself.) Scouting data typically include weekly or biweekly counts of the average number of *Lygus* individuals captured in multiple standard sweep-net samples (50 sweeps). Some PCAs also recorded square shed. Pests other than *Lygus* were not routinely sampled. We calculated the monetary cost of each late-season pesticide application for secondary pests by adding the price of the pesticide and standard application costs, using cost data from early 2010. Detailed methods for our cost calculations appear in the appendix.

Because our data do not include secondary pest densities, we use the number of late-season pesticide applications for secondary pests as a proxy for secondary pest outbreaks. Instead of attempting to estimate the effects of broad-spectrum insecticide applications throughout the growing season, we simplified the analysis by partitioning the growing season into 'early' and 'late' phases, using July 1 as the first day of the late-growing season. We chose July 1 as our separation point because cotton is most vulnerable to yield loss from *Lygus* herbivory from planting through June, and thus it is during this period when farmers may need to suppress *Lygus* populations aggressively. Thus, we will specifically ask how pre-July 1 insecticide application for *Lygus* affects the number of post-July 1 pesticide applications for non-*Lygus* pests.

Late in our study period, some fields were treated with the *Lygus*-selective insecticide

flonicamid, which suppresses populations of *Lygus* and cotton aphids, but has few effects on beneficial insects. Because flonicamid is not expected to impact native arthropod communities as severely as broad-spectrum insecticides, fields treated with flonicamid were excluded from the analysis.

Causal inference for observational data

In this section, we introduce the statistical methods for drawing formal causal inferences from observational data. Causal inference methods have become popular in scientific disciplines that study human welfare, namely public health (Little & Rubin, 2000) and economics (Imbens & Wooldridge, 2009; Gangl, 2010), where it is unfeasible, unethical, or impractical to subject human subjects to randomized, controlled experiments. Because causal inference methods are relatively unknown in the natural sciences, we provide a basic introduction to the underlying logic here. Of course, causal inference methodology extends far beyond the material presented below. Readers interested in a deeper exposition of causal inference methods may consult Imbens & Wooldridge (2009) and Gangl (2010); we find the former to be particularly readable yet comprehensive.

Potential outcomes and treatment effects

We adopt the perspective of defining causal effects via potential outcomes (also referred to as 'counterfactuals'; Rubin (2005)), and consider only the simple case of estimating a causal effect with a binary treatments and a single outcome. This scenario is illustrated in table 1. Consider estimating the causal effect of applying an early-season broad-spectrum insecticide for *Lygus* (the 'treatment') on the number of late-season insecticide applications for secondary pests (the 'response'). In notation, let $A \in \{0, 1\}$ denotes the treatment, with $A = 1$ indicating early-season insecticide application for *Lygus* and $A = 0$ indicating no such insecticide application. Let Y denote the response. Now envision the value of Y that

would result if a field receives treatment $A = 0$, and the value that would result if the field receives treatment $A = 1$. Denote these so-called potential outcomes as $Y^\star(0)$ and $Y^\star(1)$, respectively. Because each field receives only one treatment, we are not able to observe both $Y^\star(0)$ and $Y^\star(1)$ for any given field; instead, we only observe one potential outcome for each field.

For a given field, define the 'unit-level treatment effect' as the simple difference $Y^\star(1) - Y^\star(0)$ (table 1). Of course, this treatment effect is never observable for any field. We define population-level treatment effects as averages, or expectations, of the unit-level treatment effects. We will examine two population-level treatment effects. The average treatment effect (ATE) is just the expectation of the unit-level treatment effects, that is

$$ATE = \mathrm{E}\left[Y^\star(1) - Y^\star(0)\right] = \mathrm{E}\left[Y^\star(1)\right] - \mathrm{E}\left[Y^\star(0)\right]. \tag{1}$$

In the context of cotton, we can think of the ATE as the expected difference in the average response if all fields were treated for early-season *Lygus*, vs. the average response if no fields were treated for early-season *Lygus*. In addition to the ATE, we can also define the population-level treatment effect for the subset of fields that actually were treated for early-season *Lygus*. This quantity is typically referred to as the average treatment effect on the treated (ATT), and is defined as

$$ATT = \mathrm{E}\left[Y^\star(1) - Y^\star(0)|A = 1\right] = \mathrm{E}\left[Y^\star(1)|A = 1\right] - \mathrm{E}\left[Y^\star(0)|A = 1\right]. \tag{2}$$

The ATT is a more appropriate measure of the causal effect than the ATE if it does not make sense to contemplate the potential outcomes under treatment $A = 1$ for untreated fields. Here, we argue that the ATT is the most appropriate estimator of the average causal effect of early-season *Lygus* treatment in cotton, because treating a cotton field with low *Lygus* densities bears questionable relevance towards estimating secondary pest outbreaks in fields with sufficient *Lygus* densities to merit insecticide treatment.

Before proceeding, we note that although $Y^\star(1) - Y^\star(0)$ is the most commonly considered treatment effect, other treatment effects can be defined. For example, we could consider the

treatment effect to be the proportional change $Y^\star(1)/Y^\star(0)$. One could also consider different population-level summaries of treatment effects, such as the median unit-level treatment effect, or the proportion of units for which $Y^\star(1) > Y^\star(0)$.

Estimating treatment effects from data

We now discuss estimating population-level treatment effects with data. Although we are ultimately interested in estimating treatment effects with observational data, it is helpful to first discuss estimating the ATE in the context of randomized experiments.

All of the arguments below require a technical assumption that the outcome observed equals the potential outcome for the treatment received. That is, for $a = 0, 1$, if the unit received treatment $A = a$, then the observed outcome $Y = Y^\star(a)$. Rubin and colleagues call this the stable-unit treatment value assumption (SUTVA; Rubin (1980)). The primary implication of SUTVA is that the outcome observed for any unit is not influenced by the treatment received by any other unit. In essence, SUTVA is an assumption of independence among the data.

In an experiment, randomized treatment assignment implies that the potential outcomes $Y^\star(0)$ and $Y^\star(1)$ are independent of A for each unit. This independence plus SUTVA implies that the difference between the means of the treated and untreated groups is an unbiased estimator of the ATE. To see this, let n_a be the number of units that received treatment $A = a$, and write the expectation of the difference between the treatment-group means as

$$\mathrm{E}\left[\frac{1}{n_1}\sum_{i:A_i=1} Y_i - \frac{1}{n_0}\sum_{i:A_i=0} Y_i\right] = \mathrm{E}\left[Y|A=1\right] - \mathrm{E}\left[Y|A=0\right]. \tag{3}$$

Now, it suffices to show that $\mathrm{E}\left[Y|A=a\right]$, the expected outcome of a unit that received treatment $A = a$, is equal to $\mathrm{E}\left[Y^\star(a)\right]$, the expected potential outcome under $A = a$ for all

units. The proof proceeds as

$$\begin{aligned} \mathrm{E}\left[Y|A=a\right] &= \mathrm{E}\left[Y^{\star}(a)|A=a\right] \\ &= \mathrm{E}\left[Y^{\star}(a)\right] \end{aligned} \tag{4}$$

where the first equality follows by SUTVA, and the second by independence of the potential outcomes and the treatment under randomized treatment assignment. Plugging eq. 4 into eq. 3 yields

$$\begin{aligned} \mathrm{E}\left[Y|A=1\right] - \mathrm{E}\left[Y|A=0\right] &= \mathrm{E}\left[Y^{\star}(1)\right] - \mathrm{E}\left[Y^{\star}(0)\right] \\ &= ATE. \end{aligned}$$

In observational studies, treatment assignment is not random. Thus, the treatment assignment may not be independent of the potential outcomes, and thus the average response for fields that received treatment $A = a$ may not be an unbiased estimate of the expected potential outcome under $A = a$ across all fields. In particular, non-random treatment assignment introduces the possibility that confounding with one or more additional variables may produce spurious (i.e., non-causal) correlations between treatment and response. In cotton, such non-causal correlations may arise from (among other confounders) variation in PCAs' tendencies to recommend insecticide applications, and/or variation in the vigor of the cotton crop (more vigorous crops may attract arthropod herbivores of several species). This potential for spurious correlations between treatment A and response Y in observational data makes it impossible to assign a causal interpretation to the simple difference between treated and untreated fields.

How can we construct unbiased estimators of average treatment effects when the treatment assignment is not random? The key insight is this. Unbiased estimators of causal effects are obtainable if we measure a set of possible confounders, denoted $\mathbf{X}$, such that given knowledge of $\mathbf{X}$, the treatment assignment is independent of the potential outcomes. The assumption that A is conditionally independent of $Y^{\star}(0)$ and $Y^{\star}(1)$ given $\mathbf{X}$ is referred to

as the 'no unmeasured confounders' assumption, or (more awkwardly, but more accurately) the 'unconfoundedness' assumption.

We present one method for estimating causal effects under unconfoundedness below. First, however, we discuss the unconfoundedness assumption. Ultimately, unconfoundedness is an assumption on the science of the process being studied. Validation of the unconfoundedness assumption solely via data is (to our knowledge) impossible. As such, the unconfoundedness assumption requires careful scrutiny and thorough knowledge of the system. Moreover, the unconfoundedness assumption could easily be controversial, as two reasonably minded scientists could reach different conclusions regarding whether a set of covariates fully removes confounding between treatment and potential outcomes (Imbens & Wooldridge, 2009). Nonetheless, as Imbens & Wooldridge (2009) state, "there are many cases where there is no clearly superior alternative [to the unconfoundedness assumption], and the only alternative is to abandon the attempt to get precise inferences." Pearl (1995, 2000; summarized in Jewell (2004)) discusses graphical methods for identifying confounders, and we use these methods below.

A host of methods have been developed for estimating causal effects under the assumptions of SUTVA and unconfoundedness (Imbens & Wooldridge, 2009), and it is beyond the scope of this article to review them all here. Here, we estimate causal treatment effects using regression. Let $m_a(\mathbf{X}, \gamma_a)$ denote regression models for the observed outcome under $A = a$, that is, $m_a(\mathbf{X}, \gamma_a) = \mathrm{E}\left[Y|\mathbf{X}, A = a\right]$ where γ_a is a vector of parameters. The models m_a may be any type of regression, including a multiple regression, a generalized linear model, or a nonparametric regression.

To derive an unbiased estimator for ATE or ATT, it suffices to show that $m_a(\mathbf{X}, \gamma_a)$ equals to the expected potential outcome $Y^\star(a)$ given $\mathbf{X}$, that is,

$$m_a(\mathbf{X}, \gamma_a) = \mathrm{E}\left[Y^\star(a)|\mathbf{X}\right] \tag{5}$$

for $a = 0, 1$. The desired equality follows by

$$\begin{aligned} m_a(\mathbf{X}, \gamma_a) &= \mathrm{E}\left[Y|\mathbf{X}, A = a\right] \\ &= \mathrm{E}\left[Y^\star(a)|\mathbf{X}, A = a\right] \\ &= \mathrm{E}\left[Y^\star(a)|\mathbf{X}\right] \end{aligned}$$

where the first equality follows by definition, the second by SUTVA and the third by unconfoundedness. Taking expectations of eq. 5 with respect to $\mathbf{X}$ yields $\mathrm{E}_X\left[m_a(\mathbf{X}, \gamma_a)\right] = \mathrm{E}\left[Y^\star(a)\right]$.

Thus, to estimate causal treatment effects using regression, we build regression models $m_0(\mathbf{X}, \gamma_0)$ and $m_1(\mathbf{X}, \gamma_1)$ that regress the observed response Y on the confounders $\mathbf{X}$ using the data that received treatments $A = 0$ and $A = 1$, respectively. Then, an unbiased estimator for the ATE is

$$\widehat{ATE} = \frac{1}{n} \sum_{i=1}^{n} \left\{ m_1\left(\mathbf{X}_i, \hat{\gamma_1}\right) - m_0\left(\mathbf{X}_i, \hat{\gamma_0}\right) \right\}. \tag{6}$$

and an unbiased estimator for the ATT is

$$\widehat{ATT} = \frac{1}{n_1} \sum_{i:A_i=1} \left\{ m_1\left(\mathbf{X}_i, \hat{\gamma_1}\right) - m_0\left(\mathbf{X}_i, \hat{\gamma_0}\right) \right\}. \tag{7}$$

The estimators above are not identical to the naive estimate that one would obtain by the regressing Y on $\mathbf{X}$ and A, and then extracting the partial regression coefficient associated with A. In general, such a partial regression coefficient does not permit a causal interpretation. Standard errors for $\widehat{ATE}$ and $\widehat{ATT}$ can be approximated with a nonparametric bootstrap.

Before moving on, we note that each term in the summations of eqq. 6-7 is a difference between a fitted value from one regression model and a prediction from a different regression. For example, if $A_i = 1$, then $m_1\left(\mathbf{X}_i, \hat{\gamma_1}\right)$ is a fitted value and $m_0\left(\mathbf{X}_i, \hat{\gamma_0}\right)$ is a prediction. Hence, standard cautions apply regarding predictions with regression models. In particular, regression predictions are only trustworthy for confounder values that lie within the support of the fitted regression model. Thus, the causal-effect estimators above are only reliable to

the extent that the $A = 0$ and $A = 1$ treatment groups have comparable distributions of confounders, or that we are willing to extrapolate the fitted regression models to values of the confounders beyond the support of the fitted models. This makes intuitive sense — if there is a confounder with values that do not overlap for the $A = 1$ and $A = 0$ groups, then separating the effect of the treatment from the confounder is impossible. Imbens & Wooldridge (2009) have also noted that, in treatment vs. control studies, it is common for the confounder values for the untreated or control ($A = 0$) group to span a broader range than the confounder values for the treated ($A = 1$) group. When this occurs (as it does to some extent for our data), then predicting $Y^\star(0)$ for treated units requires less extrapolation than predicting $Y^\star(1)$ for untreated units. Hence, because only $\widehat{ATE}$ requires predicting $Y^\star(1)$ for untreated units, the ATT can be estimated more robustly than the ATE.

Analysis of secondary pest outbreaks

Identification of confounders

We identified confounders using the causal graph methodology developed by Pearl (1995, 2000) and summarized in Jewell (2004) (see Plowright et al. (2008) for another ecological example). A causal graph consists of a network of putative cause-and-effect relationships between variables. Once a causal graph has been suggested, a set of covariates can be identified that, once controlled for, remove the confounding between the cause and effect variables of interest. (A full discussion of how causal graphs can be used to identify confounders is beyond the scope of this article. In brief, confounding arises via 'unblocked backdoor paths' that connect the putative cause and effect variables. A complete set of confounders is any set that eliminates all unblocked backdoor paths.) Our causal graph appears in figure 1.

The causal graph in fig. 1 embodies the assumptions on which our analysis rests, and so we justify those assumptions here. Both year and cotton type (Acala vs. Pima) are assumed

to affect densities of all pests throughout the growing season. (In fig. 1, year and cotton type are separate nodes, but we combine them because they are topologically equivalent. PCA and ranch are combined for the same reason.) Early-season *Lygus* densities affect early-season *Lygus*-insecticide applications, both of which in turn affect late-season *Lygus* densities. The same is true for non-*Lygus* pests. Neighboring source crops for *Lygus* such as safflower affect both early-season *Lygus* densities and preemptive *Lygus* insecticide treatment. Early-season *Lygus* feeding also triggers square shed, which is incorporated into PCA's treatment decisions and is exacerbated by whether the previous year's crop was fertilized with phosphorous.

For several reasons, both early- and late-season insecticide application for *Lygus* may be likely to increase the chance of early- and late-season pesticide application for non-*Lygus* pests, and vice versa. This may be because, first, the cost of applying several pesticides simultaneously as part of a 'tank mix' is less than the cost of applying the same pesticides separately. Second, early-season application of broad-spectrum pesticides for *Lygus* may decrease the abundance of natural enemies, or enhance the vigor of cotton plants (White, 1984), both of which may in turn impact the densities of non-*Lygus* pests. Finally, both PCAs and farmers may be more or less aggressive in their management styles.

We re-iterate that our causal graph in figure 1 is a hypothesis, and the causal relationships that it embodies are open to debate. It is our best working hypothesis, however, and it is the hypothesis on which our analysis below rests. Using the causal graph in fig. 1, a set of six covariates can be identified that confound the relationship between early-season insecticide application for *Lygus* and late-season pesticide application for secondary pests. These confounders are year, cotton type, early-season *Lygus* density, early-season pesticide treatments for non-*Lygus* pests, PCA, and ranch. In addition to these six covariates, we also include a PCA $\times$ year interaction, because there is evidence that the PCAs' management strategies changed across the years. Importantly, this set of confounders is identical regardless of the particular mechanism that drives secondary pest outbreaks (e.g., reduction

in natural enemies, changes in cotton physiology or reduced competition).

Data preparation

We calculated early-season *Lygus* density as the average number of *Lygus* individuals collected in a standard sweep net sample, averaged across all sampling occasions prior to July 1 or the first insecticide application for *Lygus*, whichever came first. We excluded data from ranches with 3 or fewer early-season *Lygus* applications, or from PCA $\times$ year combinations with 0 or 1 early-season *Lygus* applications, as in our judgment these ranches or PCA $\times$ year combinations did not have sufficient data to estimate regression parameters reliably. Cotton type was occasionally not recorded, and so we allowed three levels of the categorical variable for cotton type: Acala (56%), Pima (33%) and unknown (11%). We did not include fields planted in hybrid (Pima $\times$ Acala) cotton ($< 0.5\%$). All told, we used $n = 969$ unique fields for our regression modeling, spanning 11 ranches and 9 years. These fields ranged in size from 2.1 to 593.0 acres, with a median size of 76.0 acres (approx. s.e. 2.7 acres). Of these fields, $n_1 = 217$ received early-season insecticide application for *Lygus*.

Regression modeling

We used Poisson regression models because our response variable was a count. We did not add or remove confounders from $\mathbf{X}$ based on their statistical significance, because our choice of variables to include in $\mathbf{X}$ is based on our hypothesized causal graph. Indeed, variable selection in regression models for causal inference is an area of active research (Imbens & Wooldridge, 2009).

We estimated the statistical uncertainty in our estimated treatment effects by a non-parametric bootstrap with 500 bootstrap data sets. To avoid bootstrap data sets with ranches or PCA-year combinations with too few instances of treated or untreated fields, we used a conditional resampling scheme, in which records were resampled within each

ranch-year-treatment combination. Consequently, bootstrap inferences pertain only to these specific ranches and years. Moreover, the bootstrap relies on the assumption that late-season, non-*Lygus* pesticide applications are conditionally independent across fields within each ranch and year, given treatment and covariates. To evaluate this assumption quantitatively, we estimated the correlation among residuals for fields from the same ranch and year, using a Pearson correlation coefficient and deviance residuals from the Poisson regressions.

We also analyzed late-season pesticide applications for three of the most common non-*Lygus* pests: aphids, mites, and armyworms. In each case, we only used data from ranches and PCA × year combinations in which at least one late-season pesticide application for the particular secondary pest was recorded for both possible early-season *Lygus* treatments. We used a subset of 805, 666 and 453 records for our aphid, mite and armyworm analyses, respectively.

Results

Table 1 summarizes the estimated causal effects of early-season insecticide application for *Lygus*. Fields that were treated for early-season *Lygus* ($A = 1$) received an average of 2.25 (s.e. = 0.13) late-season pesticide treatments for non-*Lygus* pests, incurring an average cost of US$29.6 per acre (s.e. = US$1.9). We estimate that, had those same fields not been treated for early-season *Lygus*, they would have required $\widehat{ATT} = 0.45$ (= 20.2%) fewer late-season non-*Lygus* pesticide applications (bootstrap s.e. = 0.10; 95% bootstrap CI = (0.23, 0.64); one-tailed bootstrap $p < 0.002$). We estimate the cost of these late-season pesticide applications caused by early-season treatment for *Lygus* at US$6.0 per acre (s.e. = $1.3; 20.2% of the total cost).

Across all fields, the estimated average treatment effect (ATE) is negligible (= 0.00; b.s.e. = 0.11; 95% bootstrap CI = (-0.21,+0.22)). The ATE is less than the ATT because the estimated causal effect of early-season *Lygus* treatment for untreated fields is negative

(-0.14). That is, if untreated fields had been treated, we estimate that they would have required fewer late-season secondary pesticide applications. Although the estimated causal effects for treated and untreated fields are not equal in magnitude, the two effects cancel out in the population-level ATE because only one-quarter of the fields in our data set were treated for early-season *Lygus*.

Aphids, mites and armyworms comprised the preponderance (91%) of non-*Lygus* targets for late-season pesticides. Species-level analyses suggest that, in treated fields, early-season treatment for *Lygus* increased the number of late-season pesticide applications for aphids, mites and armyworms by 0.28 (b.s.e. = 0.06), 0.09 (b.s.e. = 0.11) and 0.09 (b.s.e. = 0.04), respectively (Table 3). (With Poisson regression, estimated treatment effects for individual species do not necessarily add together to equal the total treatment effect.) Considering the statistical precision of these estimates, the estimated treatment effects are roughly proportional to the relative frequencies with which each species occurred as a late-season target across all fields.

Estimated causal effects can be visualized by plotting predicted outcomes with and without early-season *Lygus* treatment (figure 2). This plot suggests that the effect of early-season *Lygus* insecticide is not uniform, but depends subtly on the expected number of late-season secondary-pest treatments. Early-season *Lygus* treatment appears to have the largest effect on secondary pest outbreaks when the expected number of late-season secondary pest treatments is small (1–4), but has a smaller (and possibly reversed) effect when the expected number of late-season secondary pest treatments is large (> 6). The banding in fig. 2 occurs because categorical predictors (ranch, PCA, year, and cotton type) were the dominant predictors in regression model m_0, while average early-season *Lygus* density explained more of the variation in the response in model m_1. Although the regression models m_0 and m_1 are not the focus of our analysis, we provide summaries of these models in the appendix. Analysis of deviance residuals from Poisson regressions suggested a mild but statistically

significant correlation among fields from the same ranch and year (common Pearson's correlation coefficient = 0.17, $p < .001$). Neither model m_0 nor m_1 suggested overdispersion relative to a Poisson distribution.

Discussion

Using ATT as the most relevant measure of causal effects, this analysis suggests that, for the cotton fields in this study, an early-season, broad-spectrum insecticide treatment for *Lygus* elicited secondary pest outbreaks that were responsible for 20% of late-season non-*Lygus* pesticide applications. Late-season pesticide applications to manage secondary pest outbreaks cost US$6.0 per acre, on average. To the extent that we can determine, secondary outbreaks of aphids, mites and armyworms occurred in roughly similar proportion to the overall frequencies with which each species appeared as a late-season pest.

We consider ATT to be a better measure of the causal effect of early-season insecticide treatment for *Lygus* than ATE. The ATT is an estimate of the causal effect of early-season *Lygus* treatment in fields that were actually treated. The ATE is an estimate of the overall causal effect if all fields had been treated for *Lygus* regardless of early-season conditions. Although the ATE is still informative, the ATT is a more relevant measure of secondary pest outbreaks elicited by early-season *Lygus* treatment under current management practices.

The difference between the estimated causal effects in treated and untreated fields might be explained by observing that the overall effect of an early-season broad-spectrum insecticide on secondary pests combines the direct suppressive effect of mortality from the pesticide with indirect, disruptive effects that promote secondary pest outbreaks (e.g., reduced abundances of natural enemies). We speculate that the indirect, disruptive effects exceeded the direct suppressive effects in fields that were treated for early-season *Lygus*, while the reverse would have been true (though to a lesser degree) in fields that were not treated for early-season *Lygus*. Although we don't know with certainty why this may be so, we observe that the

predicted number of late-season non-*Lygus* pesticide applications in treated fields, if those fields had not been treated (2-25-0.45=1.80), was less than the actual number of late-season non-*Lygus* pesticide applications in untreated fields (1.99). This difference may indicate that, prior to the application of an early-season insecticide, natural enemies were contributing more to pest suppression in fields that were ultimately treated.

Several possible (and non-exclusive) mechanisms may drive secondary pest outbreaks, and this analysis does not discriminate among them. However, to the extent that the secondary pest outbreaks observed are caused by the disruptive effects of killing arthropod predators and parasitoids, the dollar value of the cost of pesticide applications required to curb secondary pest outbreaks sets a lower bound on the ecosystem services provided by native communities of natural enemies in this system. A full accounting of the value of the ecosystem services provided by native natural enemies would require (at the least) measuring the cost of all pest outbreaks under the hypothetical scenario in which natural enemies were permanently absent. Such an accounting is beyond the scope of this analysis. Nevertheless, this lower bound may inform ongoing efforts to valuate ecosystem services in agriculture (Costanza et al., 1997; Zhang et al., 2007).

Without doubt, our analysis rests on a host of assumptions. The chief assumption is that the decision of whether or not to treat fields for early-season *Lygus* is conditionally independent of the potential outcomes (the number of late-season secondary pesticide applications), given the confounding variables in the regression models. We argue that this is a viable assumption, especially because our confounder data document the conditions (including *Lygus* densities) that were used to determine pesticide treatments. Nonetheless, if there are unmeasured confounders that are correlated with both the early and late-season pesticide recommendations, then those confounders would render this analysis suspect.

Careful consideration must also be given to the extent to which the hierarchical structure of these data compromise the assumption of conditional independence among fields.

Although the data consist of $n = 969$ fields (or, more accurately, field-years), these fields are nested within 11 ranches, and the ranches are in turn nested within 4 PCAs. Thus, it is reasonable to ask whether or not the statistical precision of the analysis is exaggerated by considering the fields as conditionally independent given treatment and covariates. An analysis of residuals from our regression models suggest that fields from the same ranch and year are indeed positively correlated, albeit mildly (+0.17). One likely explanation for this correlation is that population dynamics of arthropod communities could have a spatial aspect that exceeds the scale of a single field. For example, mowing an alfalfa field could trigger *Lygus* migrations into several nearby cotton fields. This may be particularly true for smaller cotton fields that have a greater edge-to-area ratio and are less buffered against arthropod dispersal. Thus, these data contain less information than $n = 969$ truly independent fields, and the statistical uncertainty in our estimates is slightly greater than the bootstrap calculations suggest. Conceivably, one could design a bootstrap procedure that accounts for this spatial correlation (e.g., Zhu & Morgan (2004)), although doing so in conjunction with the conditional resampling already required would be challenging.

This analysis is not intended as a management recommendation for cotton farmers. Integrated pest management in cotton must consider a host of additional factors, not the least of which is the yield of the cotton crop. Instead, the goal of this analysis is to quantify an ecological phenomenon that is difficult to document experimentally.

The statistical methods for causal inference used here are, to our knowledge, not yet broadly known or used in the natural sciences. On the one hand, the theory underlying causal inference methods has been rigorously developed, and the methods enjoy growing use in some realms of science (Rosenbaum, 2002; Rubin, 2005; Imbens & Wooldridge, 2009; Gangl, 2010). On the other hand, these methods have not withstood the test of time to the same extent as more conventional statistical approaches, and thus some healthy skepticism is warranted. However, causal inference methods may promise new analytical possibilities

for some types of ecological and/or agricultural studies, and their usefulness deserves to be investigated.

In our view, there are two primary challenges to using causal inference methods in ecology. First, these methods require enough data to support defensible statistical models for all of the possible treatments that one wishes to consider. As such, causal inference methods will be most useful in ecoinformatics settings, where considerable volumes of data can be gathered. Because these data should span a breadth of treatments or management strategies, the most promising settings will be ones in which decision-makers have attempted a diversity of approaches. Second, the 'unconfoundedness' assumption requires that data are available for covariates that confound treatment with response. While it is difficult to speculate broadly about the types of problems for which these data may exist, we suspect that appropriate data are more likely to be available in management settings such as IPM or natural resource management, where managers may document conditions that influenced management decisions.

As a final, technical note, we observe that for this analysis, the implementation of causal inference methods was complicated by the prevalence of categorical variables (e.g., ranch, PCA, year, cotton type) in our set of confounders. As we mention above, when using regression models to quantify causal effects, the distributions of the confounders need to be sufficiently comparable among treatment groups. This comparability is more challenging with categorical confounders because categorical confounders increase the dimensionality of the confounder space. To the extent that categorical confounders may be more common in the natural sciences than in econometrics or public health, high-dimensional confounder spaces may open a technical challenge in causal inference theory that is ripe for new progress.

Acknowledgments

We are especially grateful to the pest control advisors and farmers who generously shared their data for this project. We are also indebted to Lawrence Wilhoit for help compiling pesticide use records, to Jim Jones and Jason Kuechler for data base construction and management, to Anastasios Tsiatis for guidance with causal inference methods, and to reviewers for helpful suggestions. Funding for this work was provided by USDA-NRICGP grant 2006-01761, by the University of California Statewide IPM Program, and by Cotton Incorporated.

References

Costanza, R., d'Arge, R., et al. 1997. The value of the world's ecosystem services and natural capital. Nature 387: 253-260.

Dutcher, J.D. 2007. A review of resurgence and replacement causing pest outbreaks in IPM. in Ciancio, A., and Mukerji, K.G., eds. General Concepts in Integrated Pest and Disease Management. Springer, Dordrecht.

Eveleens, K. G., van den Bosch, R., and Ehler, L.E. 1973. Secondary outbreak induction of beet armyworm by experimental insecticide applications in cotton in California. Environmental Entomology 2:497-503.

Falcon, L.A., van den Bosch, R., Ferris, C.A., Stromberg, L.K., Etzel, L.K., Stinner, R.E., and Leigh, T.F. 1968. A comparison of season-long cotton-pest-control programs in California during 1966. Journal of Economic Entomology 61:633-642.

Falcon, L.A., van den Bosch, R., Gallagher, J., and A. Davidson. 1971. Investigation of the pest status of *Lygus hesperus* in cotton in Central California. Journal of Economic Entomology 64:56-61.

Gangl, M. 2010. Causal inference in sociological research. Annual Reviews in Sociology 36:21-47.

Hardin, M.R., Benrey, B., Coll, M., Lamp, W.O., Roderick, G.K., and Barbosa, P. 1995. Arthropod pest resurgence: an overview of potential mechanisms. Crop Protection 14:3-18.

Horton, D., Bellinger, B., Pettis, G.V., Brannen, P.M., and Mitchum, W.E. 2005. Pest management strategic plan for eastern peaches. USDA-ARS/CSREES.

Imbens, G.W. and Wooldridge, J.M. 2010. Recent developments in the econometrics of program evaluation. Journal of Economic Literature 47:5-86.

Jewell, N.P. 2004. Statistics for Epidemiology. Chapman & Hall, Boca Raton, FL, USA.

Leigh, T.F., Kerby, T..A., and Wynholds, P.F. 1988. Cotton square damage by the plant bug, *Lygus hesperus* (Hemiptera: Heteroptera: Miridae), and abscission rates. Journal of Economic Entomology 81:1328-1337.

Leigh, T.F., and Goodell, P.B. 1996. Insect management, pp. 260-293 in Hake, S.J., Kerby, T.A., and Hake, K.D., eds. Cotton Production Manual. University of California DANR Publ. 3352, Oakland.

Little, R.J., and Rubin, D.B. 2000. Causal effects in clinical and epidemiological studies via potential outcomes: concepts and statistical approaches. Annual Reviews in Public Health 21:121-145.

Lunceford, J.K. and Davidian, M. 2004. Stratification and weighting via the propensity score in estimation of causal treatment effects: a comparative study. Statistics in Medicine 23:2937-2960.

Musser, F. R., Lorenz, G.M., Stewart, S.D., Bagwell, R.D., Leonard, B.R., Catchot, A.L., Tindall, K.V., Studebaker, G.E., Akin, D.S., Cook, D.R., and Daves, C.A. 2009. Tarnished

plant bug (Hemiptera: Miridae) thresholds for cotton before bloom in the midsouth of the
554 United States. Journal of Economic Entomology 102:2109-2115.

Pearl, J. 1995. Causal diagrams for empirical research (with discussion). Biometrika 82,
556 669-710.

Pearl, J. 2000. Causality. Cambridge University Press, New York.

558 Plowright, R.K., Sokolow, S.H., Gorman, M.E., Daszak, P., and Foley, J.E. 2008. Causal inference in disease ecology: investigating ecological drivers of disease emergence. Frontiers
560 in Ecology and the Environment 6:420-429.

Rao, C.N., Shivankar, V.J., and Singh, S. 2003. Secondary pest resurgence. Encyclopedia of
562 Pest Management 1:1-5.

Ripper, W.E. 1956. Effect of pesticides on balance of arthropod populations. Annual Review
564 of Entomology 1:403-408.

Rosenbaum, P.R. 2002. Observational studies, 2e. Springer, New York.

566 Rubin, D.B. 1980. Discussion of "Randomization analysis of experimental data in the Fisher randomization test," by D. Basu, Journal of the American Statistical Association 75:591-
568 593.

Rubin, D.B. 2005. Causal inference using potential outcomes: design, modeling, decisions.
570 Journal of the American Statistical Association 100:322-331.

University of California. 1996. Integrated pest management for cotton in the western region
572 of the United States. Publication 3305, University of California, Division of Agriculture and Natural Resources, Oakland, CA.

574 van den Bosch, R. and Hagen, K.S. 1966. Predaceous and parasitic arthropods in California cotton fields. California Agricultural Experiment Station, Bulletin 820.

576 White, T.C.R. 1984. The abundance of invertebrate herbivores in relation to the availability of nitrogen in stressed food plants. Oecologia 63:90-105.

578 Zhang, W., Ricketts, T.H., Kremen, C., Carney, K., and Swinton, S.M. 2007. Ecosystem services and dis-services to agriculture. Ecological Economics 64: 253-260.

580 Zhu, J. and Morgan, G.D. 2004. Comparison of spatial variables over subregions using a block bootstrap. Journal of Agricultural, Biological, and Environmental Statistics 9:91-104.

Table 1: The potential outcomes framework for causal inference. Table adapted from Rubin (2005).

Experimental unit (field)	Treatment received	Covariates	Potential outcomes $A=1$	Potential outcomes $A=0$	Unit-level causal effect	Population-level causal effects
1	A_1	$\mathbf{X}_1$	$Y_1^\star(1)$	$Y_1^\star(0)$	$Y_1^\star(1) - Y_1^\star(0)$	
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$ATE = \mathrm{E}\left[Y^\star(1) - Y^\star(0)\right]$
i	A_i	$\mathbf{X}_i$	$Y_i^\star(1)$	$Y_i^\star(0)$	$Y_i^\star(1) - Y_i^\star(0)$	
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$ATT = \mathrm{E}\left[Y^\star(1) - Y^\star(0) \vert A=1\right]$
n	A_n	$\mathbf{X}_n$	$Y_n^\star(1)$	$Y_n^\star(0)$	$Y_n^\star(1) - Y_n^\star(0)$	

Table 2: Estimated effects of early-season, broad-spectrum insecticide application for *Lygus* on late-season pesticide applications for non-*Lygus* pests

	Treated fields ($A = 1$)	All fields
Sample size	$n_1 = 217$	$n = 969$
Average late-season applications (s.e.)[†]	2.25 (0.13)	2.05 (0.07)
Estimated causal effect	ATT = +0.45 (0.10)	ATE = -0.00 (0.11)
Estimated causal effect, percentage basis	20.2% (4.5%)	-0.2% (5.4%)
Average total cost in US$ per acre[‡]	$29.6 ($1.9)	$27.2 ($0.9)
Estimated causal effect, cost basis	$6.0 ($1.3)	-$0.1 ($1.5)

[†] Average total number of pesticide applications for arthropod pests other than *Lygus* on or after July 1.

[‡] Average total cost of pesticide applications for arthropod pests other than *Lygus* on or after July 1.

Table 3: Pesticide applications for non-*Lygus* pests on or after July 1, by species

	Treated fields ($A = 1$)		All fields	
	ATT (s.e.)	Total (s.e.)	ATE (s.e.)	Total (s.e.)
Aphids[†]	0.28 (0.06)	1.14 (0.07)	0.00 (0.06)	1.00 (0.03)
Mites[‡]	0.09 (0.11)	0.45 (0.05)	0.04 (0.06)	0.40 (0.02)
Armyworms[§]	0.09 (0.04)	0.45 (0.05)	-0.10 (0.07)	0.40 (0.03)

[†] ATT and ATE calculated for a subset of $n = 805$ data records.

[‡] ATT and ATE calculated for a subset of $n = 666$ data records.

[§] ATT and ATE calculated for a subset of $n = 453$ data records.

Figure Legends

Figure 1. A causal graph for the *Lygus*-cotton system. Each directed edge represents a putative cause-and-effect relationship between variables. Nodes in boxes are the putative treatment (early-season insecticide applications for *Lygus*) and response (late-season insecticide applications for secondary pests) variables. Nodes in ovals form a set of confounders for the causal relationship between treatment and response.

Figure 2. Fitted or predicted number of late-season pesticide applications for secondary pests with early-season insecticide treatment for *Lygus* ($A = 1$) vs. without early-season insecticide treatment for *Lygus* ($A = 0$). Left panel: fields that were treated for early-season *Lygus* ($A = 1$). Right panel: all fields. Diagonals are lines of equality.

Figure 1

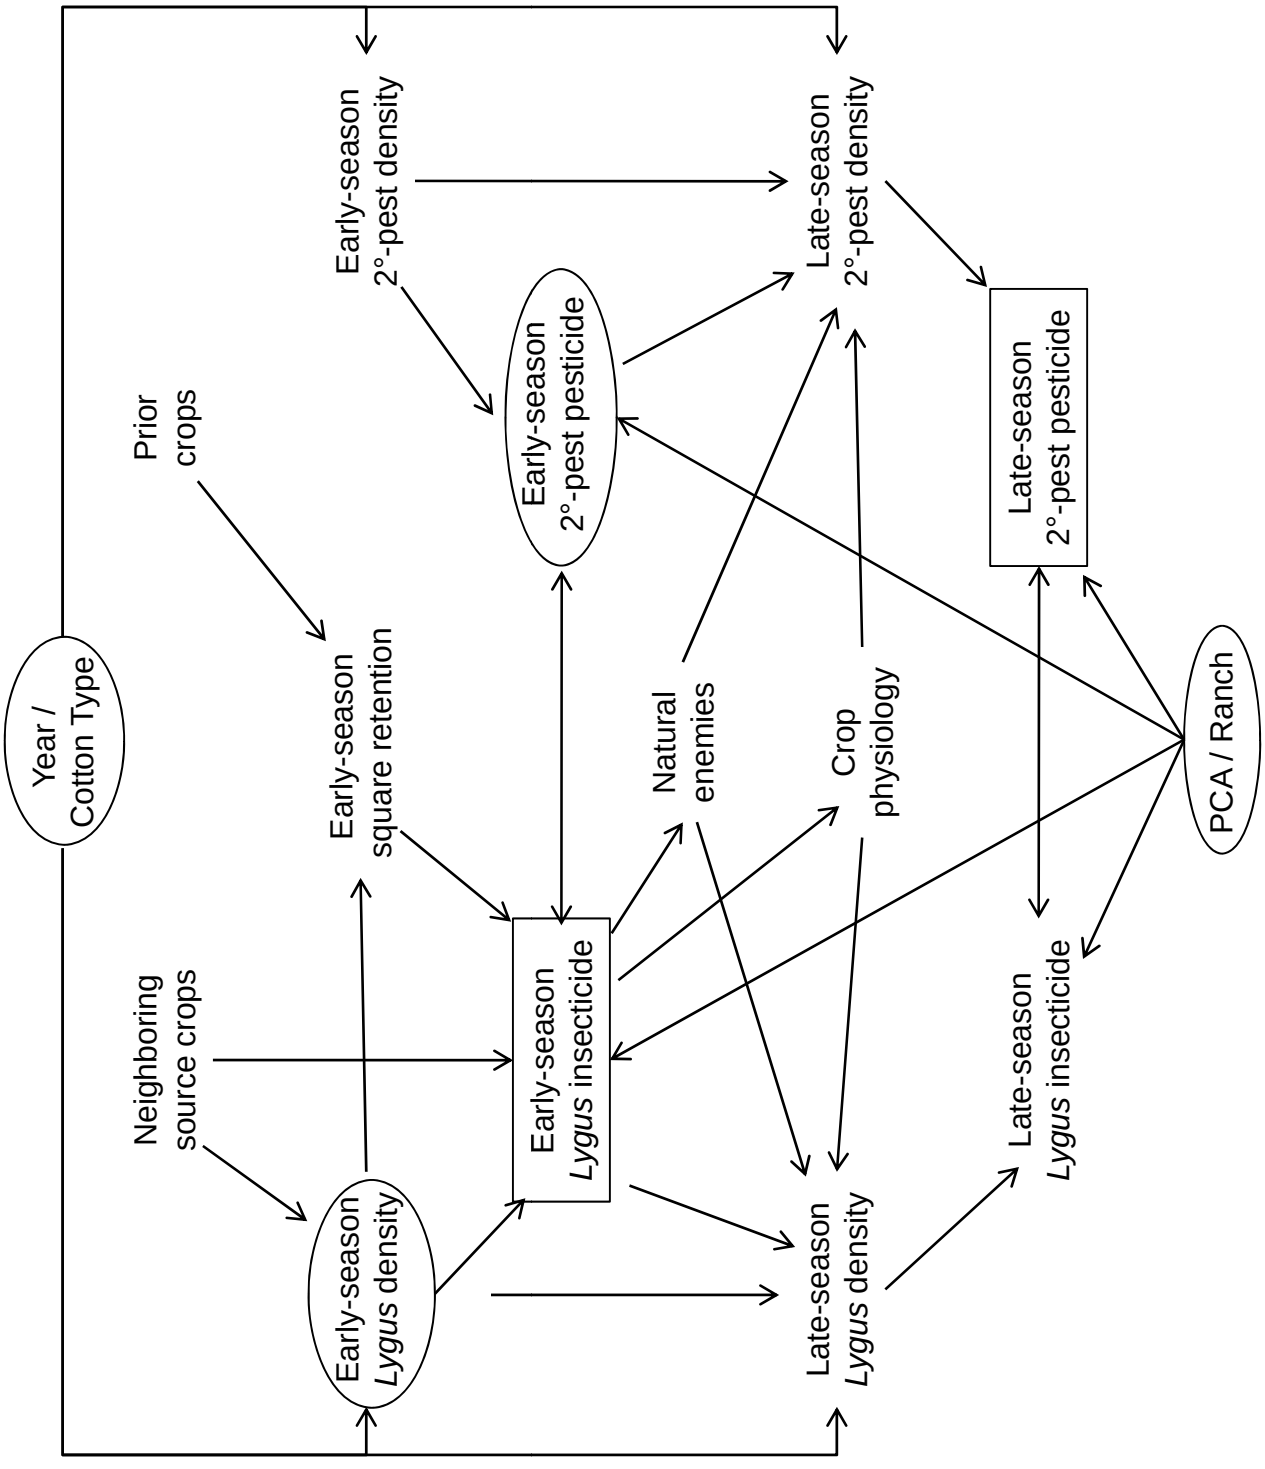
Year / Cotton Type
Prior crops
Neighboring source crops
Early-season square retention
Early-season Lygus density
Early-season Lygus insecticide
Early-season 2°-pest density
Early-season 2°-pest pesticide
Natural enemies
Crop physiology
Late-season density
Late-season Lygus density
Late-season Lygus insecticide
Late-season 2°-pest density
Late-season 2°-pest pesticide
PCA / Ranch

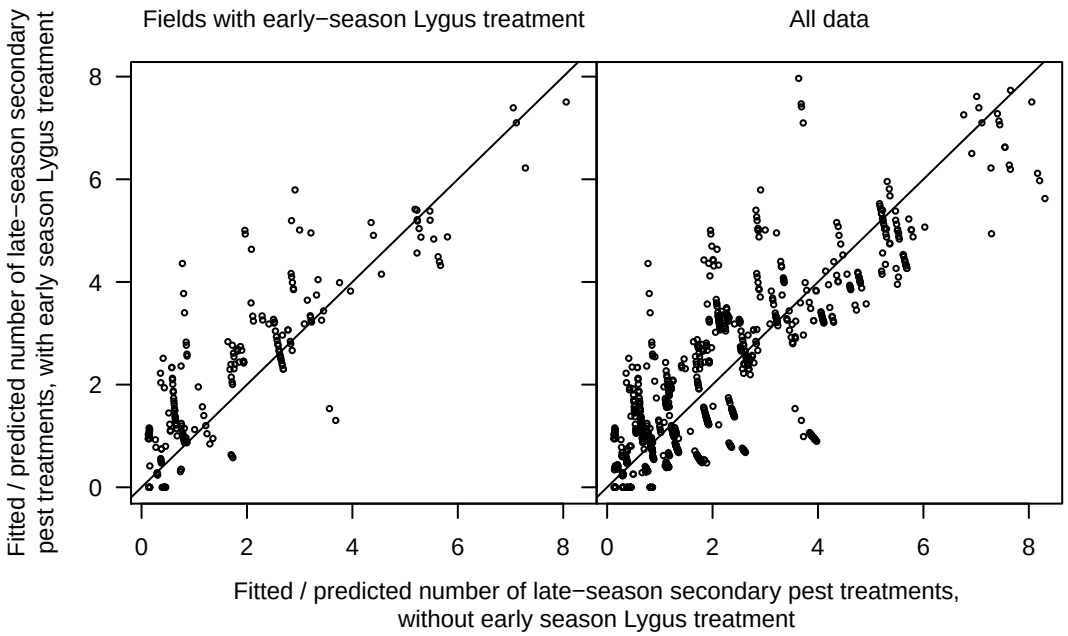
Fields with early−season Lygus treatment
All data
Fitted / predicted number of late−season secondary pest treatments, with early season Lygus treatment
Fitted / predicted number of late−season secondary pest treatments, without early season Lygus treatment
0
2
4
6
8

On-line appendix: Estimating costs of pesticide applications

We estimated the cost of pesticide applications made after July 1 as follows:

1. *Estimating the cost of the pesticides, as purchased from agricultural chemical suppliers.* Our database documented the applications of 44 different insecticides or acaricides (i.e., distinct active ingredients) to cotton. In many cases, different agricultural chemical companies market the same active ingredient under different trade names, which are sold at very similar prices; we therefore combined all brands of a given active ingredient when estimating cost. We used two sources of information to estimate the cost of the insecticides. First, we obtained list prices for 40 of the most common pesticides from a leading agricultural chemical company that sells to California farmers. The remaining four chemicals (chlorfenpyr, tebufenozide, carbofuran, and amitraz), which were applied to cotton only very rarely, were assigned a cost equal to the average for the 40 pesticides for which data were available ($10.56 per application; see below). Second, because California cotton is typically grown on a large scale, and because customers who purchase large quantities of agricultural chemicals can negotiate substantial price discounts, we solicited from our cooperating growers the actual prices they paid in recent (early 2010) chemical purchases. We obtained contract prices paid for 16 of the 40 pesticides, including most of the commonly used materials; the mean discount realized was 57%. For the remaining chemicals, we applied this mean discount to the list prices to estimate the price likely to be paid by commercial cotton farmers.

2. *Estimating the quantities typically used in a single application to cotton.*

We reviewed records in our database to obtain recent data (2007-2008) on the amounts of pesticides typically applied in a single application to a cotton crop. Standard rates can vary somewhat, within the legal bounds established by pesticide label requirements, across years, farmers, and pest control advisors. In all cases, we attempted to identify the most common or mid-range application rate. In a few cases, application rates varied as a function of (i) the pest target, (ii) the time of year, or (iii) whether or not the pesticide was being applied in combination with another pesticide. In each case, we adjusted our cost estimates accordingly.

Bringing together our data on pesticide costs and application rates, we estimated the mean chemical cost for a single application per acre of cotton of $10.56 (range: $1.73 - $37.80 across the 44 different pesticides)

3. *Estimating the costs of applying the pesticide to the cotton crop.*

Nearly all pesticide applications are made by custom applicators, rather than by the farmers themselves, facilitating estimation of the cost of applying the materials. Most pesticides are applied to cotton by air, with fixed-wing aircraft, at a rate of 10 gallons water (carrier) per acre. The cost of such an application was estimated as $9.00 per acre, as confirmed by cooperating pest control advisors and local custom applicators. Some pesticides were applied by air, but using a higher gallonage (15 gallons/acre; cost = $10.90/acre) or using ground spray rigs ($9.00/acre). Finally, one of the commonest pesticides, aldicarb, is frequently applied during the planting operation; in this case we assigned no additional application cost.

4. *Rules for assigning application costs when multiple agricultural chemicals were applied together ('tank mixes').*

It is common for farmers to apply several chemicals simultaneously as a ‘tank mix.’ Such mixes may include multiple pesticides, or may include combinations of pesticides and other agricultural chemicals, including plant growth regulators (e.g., mepiquat chloride, which restrains vegetative growth) or defoliants (which must be applied at least once prior to harvest). We used the following rules in assigning application costs to pesticides used for secondary pests when applied as tank mixes: (i) we assigned zero application cost in cases where a plant growth regulator or a defoliant was present in the tank mix; we reasoned that applications of plant growth regulators or defoliants are often obligatory for the farmer, and thus including pesticides in tank mixes with such chemicals entails no additional application cost; (ii) in the remaining cases, if N pesticides were applied together, we assigned $1/N$ of the full application cost to each of the N pesticides.

On-line appendix table: Residual deviance for model terms for Poisson regression models[†]

Predictor	df	Residual deviance reduction: Model m_0[‡] (untreated fields)	Residual deviance reduction: Model m_1[§] (treated fields)
Ranch	7	64.31	31.17
Cotton type	2	21.28	4.97
Average early-season *Lygus* density	1	0.24	6.74
Number of early-season non-*Lygus* treatments	1	1.37	0.01
Year[††]	8	188.40	10.98
PCA[‡‡]	3	169.37	101.35
Year × PCA	8	76.88	39.36

[†]Difference in residual deviance between the model with all terms included and the model without the given term; analogous to a Type III SS in linear models.

[‡]Total residual deviance 620.02 on 721 df

[§]Total residual deviance 148.58 on 186 df

[††]Residual deviance reduction relative to a model without year or year × PCA interaction

[‡‡]Residual deviance reduction relative to a model without ranch [nested in year] or year × PCA interaction