Multidimensional Responses to Disease Information: How Do Winegrape Growers React to Powdery Mildew Forecasts and To What Environmental Effect?

Travis J. Lybbert, Nicholas Magnan, and W. Douglas Gubler

Access to better information can enable adjustments that improve individual and social outcomes. In agriculture, accurate forecasts can be especially valuable to both producers and society at large. We analyze the impact of disease forecasts on the responses of growers along multiple margins of adjustment. Using high-resolution plot-level panel data, we analyze how California winegrape growers adjust powdery mildew control strategies in response to disease forecasts. While trials that allow only treatment timing adjustments suggest these forecasts bring large environmental benefits, in practice many growers simultaneously adjust what, when, and how much they spray. We find that the net environmental impact of this multidimensional response may actually be negative.

Key words: Forecast information; Pesticides; Disease; Agriculture; Environment; IPM; Grapes; California.

JEL codes: D8, Q1, Q5.

When stochastic shocks shape production outcomes, producers value reliable forecast information—and their response to new or better forecasts can have direct implications for production, market, and broader social or environmental outcomes. These links between uncertainty, information, producer responses, and outcomes are especially pronounced in agricultural contexts. How farmers manage production risk and adjust production strategies using weather or price forecasts can directly affect markets and therefore consumer welfare. Similarly, many integrated pest management (IPM) approaches aim to reduce pesticide usage and related environmental impacts by providing growers with better, more precise pest and disease information. Such spatially and temporally precise information may enable growers to refine their expectations and operations, but the fact that producers can adjust along several dimensions simultaneously makes characterizing their response to new information as challenging as it is important. In this article, we test how California winegrape growers alter their disease management strategies in response to powdery mildew forecasts and estimate the environmental effects of these responses.

Powdery mildew is one of the most important production risks that grape growers face. Their only real hope in the annual battle against powdery mildew is proper

Travis J. Lybbert (tlybbert@ucdavis.edu) is a professor, Agricultural and Resource Economics and Member, Giannini Foundation for Agricultural Economics, UC Davis, One Shields Ave, Davis, CA 95616; (530) 554-1393. Nicholas Magnan is an assistant professor, Agricultural and Applied Economics, University of Georgia. W. Douglas Gubler is a Professor, Plant Pathology, UC Davis.

This project was funded in part by the Giannini Foundation for Agricultural Economics. Several people have been instrumental in this research by generously sharing their time, expertise, and perspectives, namely, Jenny Broome, Joyce Strand, Christopher Janousek, Rick Melnicoe, Larry Wilhoit, Yu Zhan, Minghua Zhang, Sara Bonetti, Kimberly Steinmann, Olena Sambucci, and anonymous reviewers. We thank Terra Spase and AgUnlimited for providing PMI data for several regions in northern California. In central and southern California, we accessed the PMI or raw temperature data with the help of the University of California Statewide IPM Program and Jenny Broome. For further insights and suggestions, we thank seminar audiences at the California Department of Pesticide Regulation, UC Davis, UC Berkeley, the Western Weather Workgroup, the 2009 and 2010 meetings of SCC-76 (“Economics and Management of Risk in Agriculture and Natural Resources”), the 2010 Conference of the American Association of Wine Economists, and the 2010 meetings of the American Applied Economics Association.

Amer. J. Agr. Econ. 98(2): 383–405; doi: 10.1093/ajae/aav097

preventative management because powdery mildew can quickly outpace any eradication regimen when optimal temperature and humidity conditions prevail. Such rapid growth poses substantial production risks to growers: An entire season can be lost with a single poorly timed powdery mildew treatment. In response, growers apply heavy and frequent doses of sulfur products and relatively more toxic synthetic fungicides in their vineyards to avoid such outbreaks, essentially over-applying these chemicals as a form of insurance that is implicitly subsidized by the environment, which absorbs part of the costs (Mumford and Norton 1984). Better information about the risk of a powdery mildew outbreak is especially valuable in such a context.

Recent advances in remote sensing, telemetry, GPS, and computing power have improved the collection, processing, and rapid dissemination of high-resolution spatial data. In addition to enabling other precision agriculture technologies (Cassman 1999; Bongiovanni and Lowenberg-DeBoer 2004; McBratney et al. 2005), these advances set the stage for IPM approaches aimed at reducing pesticides, including decision rules based on economic thresholds (Fabre, Plantegenest, and Yuen 2007) and better forecast information. Several studies have documented the potential environmental effects of changes in pest management decisions due to disease and pest forecasts (Mumford and Norton 1984; Moffitt et al. 1986; Swinton and King 1994) and precision production technologies that enable more targeted practices (Khanna and Zilberman 1997; Bongiovanni and Lowenberg-DeBoer 2004). While better information about pests need not improve environmental outcomes in practice (Lichtenberg 2002), this is nonetheless a common theoretical result.

The powdery mildew forecast we study in this article emerged in the mid-1990s as one such information innovation in agriculture and as a high-profile IPM strategy promoted by University of California's Statewide IPM Program.[1] This Gubler-Thomas Powdery Mildew Index (PMI) (Thomas, Gubler, and Leavitt 1994; Weber, Gubler, and Derr 1996) is designed to help growers anticipate outbreaks so they can more precisely time preventative powdery mildew treatments and thereby reduce fungicide applications. In addition to private benefits internalized by producers, the social and environmental benefits of reduced fungicide use due to better treatment timing could be substantial. Since powdery mildew affects several crops beyond grapes, many producers in many places could reap similar benefits from these disease forecasts. Yet the purported value of the PMI to growers has been extrapolated from controlled field trials in which technicians mechanically adjust the timing of powdery mildew treatments according to the PMI. In reality, growers can adjust along multiple dimensions: when (how often) to spray, what to spray, and how much to spray. Little is known about how growers actually alter their pesticide use in response to information, or how these potentially multidimensional responses affect net pesticide usage.

To address these questions we estimate models of growers' disease management strategies using high-resolution temporal and spatial data collected at the grower- and plot-level. The temporal and spatial resolution of these data allows us to estimate the impacts of receiving disease forecasts on grower behavior at an unprecedented level of detail. Using a response function difference-in-difference approach, we exploit the temporal resolution of the data by comparing daily pesticide use between farmers using and not using the PMI, conditional on the information PMI recipients receive. Building on recent work to estimate the plot-specific environmental risks associated with a given pesticide regime (Zhan and Zhang 2012), we test the impact of receiving the PMI on environmental risks using an event study identification strategy.

Our results suggest that growers do adjust their disease management strategies based on the PMI, but that this response is multidimensional, nonlinear, and spatially heterogeneous. While some PMI recipients do tighten treatment intervals as the PMI increases as intended by its promoters, they tend to adjust other dimensions of powdery mildew treatment even more aggressively, including what product to use, whether to mix products in a single treatment, and dosage rates. These complex responses are also distinctly nonlinear, with high PMI values often inciting aggressive responses. They are also intimately interrelated such that primary responses seem to induce secondary

[1] See http://www.ipm.ucdavis.edu/WEATHER/index.html (accessed 9 December 2015).

adjustments as a by-product—adjustments that may appear counterintuitive when observed in isolation. For example, growers who shift to more potent synthetic products as the PMI increases stretch rather than tighten their treatment intervals because pesticide regulations require longer minimum intervals for synthetic products. We also find that growers' PMI responses vary dramatically across different grape growing regions. Since PMI recipients do not respond to the PMI by adjusting intervals alone, the net environmental effect of the PMI is ambiguous and open to empirical testing. Using plot-year level scores that assess the environmental risks associated with pesticide usage, we find evidence that receiving the PMI has a negative environmental impact: Receiving the PMI results in roughly a 40% increase in environmental risks as measured by these scores. In the concluding section, we discuss the implications of these results both for policy and the design and potential impact of IPM programs.

Background

Grapes contribute roughly 10% to California's annual $30 billion in farm sales and are the second most important crop in California.[2] Winegrapes constitute an important part of total grape production, and the California wine industry has become a major component of the state's dynamic agriculture sector (Heien and Martin 2003; Goodhue et al. 2008). For grape growers worldwide, powdery mildew (PM) control is arguably the most important single management practice. For high-value winegrapes, less than 5% disease coverage of berries at harvest can cause crucial off-flavors in wine (Stummer et al. 2005). California winegrape growers can allocate 20% of their total cultural costs to controlling PM and often suffer more losses to it than any other disease (Fuller, Alston, and Sambucci 2014). In their annual battle with PM, winegrape growers use a variety of preventative control options (Flaherty et al. 1992). PM is generally controlled using an integrated program with regular treatments occurring every 7–21 days. The default treatment is sulfur dust, which is relatively cheap and can be applied at faster speeds, or micronized dry flowable sulfur.[3] Sulfur is also acceptable for use on organically certified winegrapes. As conditions change throughout the season, growers often switch to more potent synthetic fungicides such as quinoxfen, demethylation inhibitors (DMI), or strobilurin fungicides. Sulfur is, however, commonly maintained in the program—either mixed in the tank or in rotation—to combat resistance or delay the onset of resistance to narrow spectrum synthetic fungicides.

Growth and development of PM is strongly and nonlinearly affected by climatic conditions.[4] When optimal temperatures prevail during critical windows, a mistimed treatment can have catastrophic effects on the value of production at harvest. In this context, the potential value of disease forecasts is substantial. The PMI, based on the Gubler-Thomas model, aims to provide such a forecast using the documented relationship between temperature, humidity, and ascospore release to predict initial disease onset and the subsequent reproductive rate of the pathogen (Thomas, Gubler, and Leavitt 1994; Weber, Gubler, and Derr 1996; Gubler et al. 1999).[5] The PMI ranges from 0 (no risk) to 100 (extreme risk): Low index values of 0–30 indicate the pathogen is not reproducing, moderate values of 40–50 imply a reproductive rate of approximately 15 days, and high index values of 60–100 indicate that the pathogen is reproducing rapidly (as fast as every 5 days) and that the risk for a disease epidemic to occur is extreme.[6] Since

[2] As Californian crops, grapes are second only to greenhouse and nursery products in terms of total sales.

[3] Micronized sulfur is processed and formulated to provide better coverage on vegetation and to oxidize more slowly.

[4] It thrives under dry conditions with moderate temperatures (21 to 30° C), but spores and mildew colonies can be killed by extended durations of temperatures above 32° C. The fungus can be destroyed completely when air temperatures rise above 32° C for 12 hours or more (Ypema and Gubler 1997). During continuous favorable temperature periods, the time between spore germination and production of spores by the new colony can be extremely rapid, occurring in as little as 5 days.

[5] The model and its forecast have been validated throughout the grape growing regions of California and other parts of the world. Similar disease forecasting models have been developed to predict the onset and severity of other plant diseases whose development is predictably influenced by climatic conditions, namely apple and pear scab, fireblight, botrytis bunch rot, wheat diseases, and tomato diseases.

[6] After budbreak, there must be three consecutive days with a minimum of six consecutive hours of temperatures between 21 and 30° C for a powdery mildew epidemic to be initiated. The early season PMI therefore begins at 0 and increases by 20 points for each day with six or more consecutive hours in this optimal temperature range (e.g., after three consecutive days of six or more hours of optimal temperatures the PMI climbs to 60). If after one day of temperatures in this range optimal temperatures

the mid-1990s, the PMI has been available in many regions as either a specific value for a single location or as a contour map for a defined space—often via daily email messages. Increasingly, growers use their own on-site weather stations with integrated software to compute the PMI. In areas with substantial microvariability in climate—which characterizes high value winegrape areas such as Napa Valley—these spatially precise forecasts are particularly valuable.

Per its original motivation, the PMI can potentially enable growers to time their fungicide treatments more precisely with the actual disease risks that prevail in their vineyards. In particular, growers may postpone fungicide applications during extended periods with low PMI values. This potential value of the PMI has been demonstrated in field trials, which have shown that spraying according to the PMI can reduce fungicides "by 2–3 applications over the course of the growing season with equal or better disease control" (Gubler et al. 1999, 10). The social, economic, and environmental benefits of this reduction in fungicide use could be substantial. For example, it is estimated that the PMI could have reduced total sulfur applications by over one million pounds in 2003 (8%) if only a quarter of raisin growers followed the PMI (UC Agriculture and Natural Resources 2005).[7]

The magnitude of the realized environmental benefits associated with growers' use of the PMI depends on three important factors. First, the actual benefits from the PMI obviously hinge on PMI adoption among growers, especially those responsible for large shares of total fungicide applications. Second, how growers make PM treatment decisions in the absence of the PMI provides a baseline from which the PMI response must be assessed. Many in the industry assume that growers' baseline tendency is a calendar (or minimum interval) spray schedule. In this case, having more information on pest risk could only result in less frequent treatments. However, aggregate analysis of pesticide use reports in California suggest that these baseline schedules often deviate from a strict calendar spray regimen and may be partly conditioned on other factors (Epstein and Bassein 2003). For example, prior to the development and diffusion of the PMI, plant pathologists typically told growers, "If you like the weather outside (mild and dry), then so does powdery mildew." Finally, the field trials cited above used the PMI exclusively to adjust the timing of powdery mildew treatment. The realized benefits of growers' actual use of the PMI may depend on this unidimensional response. Simultaneously adjustment along multiple dimensions may complicate and reduce the net benefits of PMI usage.

Data

Our empirical analysis hinges on the high spatial and temporal resolution of the data we use to estimate growers' response to the PMI. By merging data from multiple sources, we construct a high-resolution panel data set that tracks daily fungicide use and yearly PMI use among wine-grape growers from 1996 to 2007 and includes spatially specific daily PMI forecasts for this period.

As the starting point for constructing this data set, we conducted an online survey of California winegrape growers in January and February 2008. The survey included questions on disease management generally and PM specifically—including whether the grower has received the PMI in each of the preceding 12 years[8]—on vineyard and vineyard manager characteristics, and on use of the PMI. Members of the California Association of Winegrape Growers and several other

do not persist for three consecutive days, the PMI reverts to zero. Once this early season requirement for three consecutive days of optimal temperatures is met, the index fluctuates between 0 and 100 based on daily temperatures for the remainder of the season: the PMI gains 20 points for each day of optimal temperatures and loses 10 points for each day that does not meet this six hour optimal temperature requirement. The PMI also loses 10 points if at any point during the day temperature rose to 35° C or higher for at least 15 min.

[7] Winegrape growers manage powdery mildew even more aggressively than raisin growers and, based on this simple extrapolation, PMI use might generate even greater reductions in sulfur usage.

[8] These questions included whether they had received the PMI in each of the preceding 12 years, their perceptions of the accuracy of the PMI, and how and why they use the information. Given the structure of our data, there is no way for us to know how many of a grower's neighbors also received the PMI or how widely this information may disseminate through social networks. Throughout the analysis, we implicitly assume that only growers who explicitly reported that they received the PMI daily in a given growing year had access to the PMI, which ignores any diffusion of this information through social networks. Given that the PMI is a daily forecast, casual conversation is unlikely to confer this information effectively, which makes "no social diffusion" a reasonable assumption.

state and local winegrape growers' associations were invited to participate, of which 110 growers completed the survey. The growers who responded to the survey have similar yields and receive similar prices per ton as other winegrape growers in their respective counties (table S1). Because we cannot test the representativeness of this sample across other dimensions—including those most relevant for our purposes (e.g., disease management tendencies)—extrapolating our results to growers who choose not to join a winegrape grower association, for example, may be problematic.[9] At the time of the survey, roughly half of the growers in our survey actively used the PMI to control PM. However, adoption rates steadily increased over time from 1996 to 2007. Adoption was highest (75%) among winegrape growers in Napa and Sonoma counties where production value is high, varieties are highly susceptible to mildew, and the importance of locational branding can amplify growers' sensitivities to the environmental impacts of their production practices (Friedland 2002; Warner 2007) as well as their interest in local partnerships for promoting sustainable viticulture practices (Broome and Warner 2008). At the time of the survey, nearly two-thirds of our surveyed growers had used or were currently using the PMI to some degree (Lybbert and Gubler 2008).

Next, we obtained daily PMI values for locations near our surveyed growers. In some cases, we reconstructed the PMI from raw hourly temperature data collected from weather stations near these growers. In either case, our PMI data are unlikely to be taken from exactly the same source as the PMI received by a particular grower but is likely to be highly correlated with the PMI he received. Most of these PMI data begin when the model was first used in 1996 and continue until 2007. We matched the plots of growers in our survey to the nearest station for which we have PMI data using linear distance measures. In the full sample, the mean distance between the plots and their matched weather stations is 22 km, but this varies widely across growing regions. Although the correlation in the PMI derived from geographically dispersed stations hovers around 50–60%, in all subsequent analyses that use these matched PMI data we exclude plots that are considered "too far" from their matched stations and conduct robustness tests that include only plots that are much closer.[10] Once we exclude plots that are "too far" from their closest weather station, the overall average distance falls to 13.2 km and to just over 3 km for the hilly, microclimate rich North Coast region.

Our final data layer introduces unprecedented daily resolution on pesticide management decisions. The data in this layer were collected as part of California's rigorous pesticide use reporting system administered by the Department of Pesticide Regulation (DPR). In this system, growers must obtain a pesticide use permit before applying any pesticides and then must file a pesticide use report (PUR) with their respective county agriculture commissioner each time they apply a pesticide. These PURs are collected across counties by the DPR and are publicly available via the DPR website at the plot level (see Epstein [2006] for more details about the PUR system). Each PUR contains the grower's grower ID, which allowed us to match our surveyed growers to their PURs, along with an impressive battery of other details, including the crop treated, the product used, its active ingredient, the application rate, the number of acres treated, and the total size and spatial location of each of their plots. In our analysis, we use plot-level PURs to understand growers' PM treatment decisions. One aggregate analysis of PUR data and the potential impact of IPM on pesticide usage precedes our grower-specific analysis (Epstein and Bassein 2003).[11] Our empirical approach affords much greater precision,

[9] To our knowledge, there are no publicly available survey data from representative samples of California winegrape growers to serve as the basis for representativeness tests across many dimensions. The USDA agricultural census for California provides some basic information for grape growers in general (including wine, raisin, and table grape growers) but does not include details that overlap with our survey instrument. For instance, the census does not report operation size in a way that is comparable to our survey data.

[10] Because winegrape growing areas that have highly variable microclimates require closer proximity to weather stations, these subjective determinations—which were made in consultation with plant pathologists—are different for each growing region. In Mendicino, "too far" is >15 km and "close" is <8 km. In Napa and Sonoma, "too far" is >8 km and "close" is <5 km. In San Luis Obispo, "too far" is >23 km and "close" is <13 km. In Fresno and Madera, "too far" is >60 km and "close" is <38 km. In San Joaquin, "too far" is >30 km and "close" is <12 km. Any noise that is due to the distance to a weather station will be absorbed in the error term and will consequently reduce the precision of our estimates.

[11] Several other inquiries—including both environmental and human health related—have used PUR data (e.g., Reynolds et al. 2002; Davidson 2004).

as we leverage the temporal resolution of the PUR data by integrating it with spatial daily PMI data and unique grower survey data that includes PMI usage by year. The resulting data set allows us to directly test the impact of the PMI on growers' PM strategies, pesticide usage, and environmental effect.

There are, however, some limitations with PUR data, two of which are worth mentioning here. First, PURs do not include *why* the grower applied the pesticide (e.g., the pest or disease targeted). Fortunately, PM treatment can be inferred quite accurately based on the product used since most of the fungicides used to control PM focus narrowly on this disease. We used data from the UC IPM Program on the efficacy of different fungicides for grapes[12] and consultations with plant pathologists to identify PM treatments. Second, growers choose their own plot labels and sometimes change labels from one year to the next, which can complicate tracking plot trends over years. Since there is complete consistency of plot labeling *within* a given growing season, we construct treatment intervals at the plot level[13] and use plots as our unit of analysis. To remedy inconsistencies in plot labeling *between* years, we use a combination of algorithmic and manual matching. This allows us to identify the vast majority of plots from year to year, which enables us to control for unobservable time-invariant plot characteristics.

Table 1 contains an overview of the relevant dimensions of our data. Our analysis of growers' response to PMI information includes growers from three major winegrape growing regions for which we had adequate PMI data: North Coast (Napa, Sonoma, and Mendocino counties), Central Coast (San Luis Obispo county) and Central Valley (Fresno, Madera, and San Joaquin counties). Since there are important within-region similarities and significant between region differences (e.g., in growing calendars, grape varieties, pest pressure, and production value), we disaggregate our analysis by these growing regions (Flaherty et al. 1992). While our survey sample includes 86 growers in these three regions, due to data constraints in the other layers of the database, the usable sample of growers varies across the analyses below and is typically less than this full sample.

Our estimation strategy leverages several key dimensions in our data, including growers, plots, years, and regions. While there are nearly a million grower-plot-days in our full sample (45% of which are in the high value North Coast region), we model product and dosage decisions conditional on having decided to spray for PM.[14] This approach uses each PM treatment as an observation. In the full sample, we have a total of 23,958 treatments by growers not using the PMI and 28,224 by those using the PMI. Since our data begin as early as 1996, the year the PMI was first released, many of our growers switch from not using to using the PMI during our data window, which enables us to test the robustness of our results by estimating models for these "switchers" only.

The empirical models that follow are based on treatment differences between those who receive the PMI in a given year with those who do not. The simple descriptive statistics in table 2 offer an unconditional comparison of the interval and fungicide use tendencies of the growers in our sample. There appear to be some differences between those receiving the PMI and those not receiving it. Those getting the PMI are slightly less likely to include synthetic products in their sprays for PM and seem to have longer intervals after spraying sulfur, but these intervals are statistically indistinguishable from those of their non-PMI receiving counterparts after spraying synthetic products. In addition to not controlling for any of the factors other than receiving the PMI that potentially affect PM treatments, these unconditional comparisons pool together growers from all three regions, which likely respond to the PMI in quite distinct ways.

[12] Available at http://www.ipm.ucdavis.edu/PMG/r302902111.html (accessed 10 March 2010).

[13] We reconstruct these treatment intervals as the number of days between consecutive powdery mildew treatments on a given plot. Often, growers treat over consecutive days (e.g., when plots are large and take more than one day to treat), in which case we used the day on which the highest proportion of the plot was treated as the center point of the treatment and compute the number of days between center points.

[14] In earlier versions of the analysis, we used the daily binary variable PM treatment = {yes, no}. We prefer instead to model the timing of PM treatment using the interval since the last PM treatment as it is a more accurate depiction of growers' decision making process and more consistent with the way the PMI is explained to growers.

Table 1. Temporal and Spatial Dimensions of Data

	Number of Surveyed Growers	Number of Plots	Grower-Plot-Years	Grower-Plot-Days[a]	Number of PM Treatments		Avg Distance to Weather Station (km)[b]
					No PMI	With PMI	
North Coast (subtotal)	**31**	**444**	**2,299**	**423,016**	**8,125**	**12,580**	
Napa	8	221	908	167,072	1,283	4,826	4.7
Sonoma	19	180	1,103	202,952	6,453	6,453	6.6
Mendocino	4	43	288	52,992	389	1,301	25.9
Central Coast							
San Luis Obispo	17	66	475	80,750	2,739	4,470	16.5
Valley (subtotal)	**19**	**229**	**1,534**	**244,680**	**6,547**	**5,587**	
San Joaquin	9	95	805	120,750	855	4,175	26.5
Fresno	5	55	406	69,020	4,680	33	68.9
Madera	5	79	323	54,910	1,012	1,379	69.4
TOTAL	**86**	**968**	**5,842**	**993,126**	**23,958**	**28,224**	**22.5**

[a]For each grower-plot this includes the number of days during the growing season that are at risk for PM. This PM "window" within the growing season is region-specific (see Flaherty et al. 1992).
[b]This average distance is based on the distance between the plots managed by our surveyed growers and the weather stations that collected the PMI data we use (to which these plots were matched).

Table 2. Descriptive Statistics for Fungicide Use and Intervals for the Growers in Our Survey

Variable	Combined	Get PMI	Do Not Get PMI
Percentage of PM sprays with sulfur	47.6%	47.1%	48.0%*
Percentage of PM sprays with synthetics (sterol inhibitors or strobulurins)	55.0%	53.3%	56.2%***
Interval after using sulfur	12.27	13.8	11.1***
Interval after using synthetic	15.0	14.6	15.2
Stretching past recommended interval (next treatment with any chemical) in days	−2.22	−2.06	−2.33

Note: Asterisks ***0.01, **0.05,*0.1 indicate statistical significant differences between those receiving the PMI and those not receiving it. Sum of percentages across fungicide types can sum to more than 100% due to combination sprays.

Empirical Model and Results

The objectives that guide our empirical analysis are twofold. First, we aim to test whether winegrape growers respond to the PMI by adjusting their PM treatment strategies along three potential margins of adjustment: treatment timing, product choice, and dosage rate. Second, we aim to assess the net environmental impact of any response to the PMI. Since reducing the environmental impact of pest management practices is fundamental to disease forecast models as a component of IPM, understanding this impact empirically is critical—but also complicated once we allow for simultaneous adjustment on more than one margin. PMI users may stretch intervals when the PMI is low, but they may also increase dosage rates when it is high. Furthermore, growers may respond asymmetrically to changes in the PMI. They may respond aggressively to increases in the index, while being relatively unresponsive to a low PMI.

In our pursuit of these empirical objectives, we leverage the panel structure of our data and control for unobservable determinants of PM management with plot and year fixed effects. In addition to comparing the treatment tendencies of growers using the PMI to those not using the PMI, we estimate response models for growers who initially did not use it but switched to using the PMI during the coverage of our data. With sufficient numbers of "switchers," growers can serve as their own counterfactual in a before-after identification approach with grower fixed effects to control for time invariant unobservables that may be correlated with PMI usage. We also estimate "placebo" models that complement these before-after results with "switchers": We estimate similar models for growers who never received the PMI and generate a random, false PMI receipt year for each.

How do growers respond to the PMI?

We characterize growers' response to the PMI along three margins of adjustment: (1) when to spray, (2) what to spray, and (3) how much to spray. PM treatment is conditioned on several factors in addition to expected disease pressure, including the growth phase of the vines. Furthermore, the evolution of PM pressure varies widely by year and across space, with PM outbreaks often flaring up at the plot level according to vagaries in the microclimate and to the susceptibility of different grape varieties. We use a variety of control variables and plot-year fixed effects to account for these factors. Conditional on these determinants of PM treatment, we aim to compare growers who receive the PMI with those who do not. While growers' PMI response may be nonlinear in important ways, we begin with simple approximations that assume that their response is piecewise linear. With these broad tendencies in mind, we then estimate more flexible specifications that accommodate nonlinearities in growers' PMI response. Since grape growing regions in California have distinct agroclimatic conditions, we estimate these specifications separately for each growing region.

We use variations of two similar specifications as the basis for modeling growers' responses. Both of these are essentially a difference-in-difference specification applied to growers' PMI response function (rather than to a scalar outcome variable). In the first, we pool PMI recipients and nonrecipients together and use interaction terms to test for PMI recipients' response to the PMI. In the second, we estimate separate response functions for recipients and nonrecipients. The first (pooled) specification imposes a linear response function—which we relax later[15]—to enable simple comparisons and is of following form:

$$(1) \quad R_{ijdt} = \alpha_0 + \beta_0 PMI_{ijdt} + \beta_1 \left(Z_{it} \times PMI_{ijdt}\right) + \beta_2 \left(Z_i^{Ever} \times PMI_{ijdt}\right) + \mathbf{x_{ijdt}}'\varphi + \boldsymbol{\lambda}_{\mathbf{ijd}} + \boldsymbol{\tau}_{\mathbf{t}} + \boldsymbol{\mu}_{\mathbf{j}} + \varepsilon_{ijdt}$$

where R_{ijdt} is a response variable that captures one of the three adjustment margins for grower i on plot j on day d in year t, $Z_{it} = \{0,1\}$ indicates whether or not grower i received the PMI in year t, and PMI_{ijdt} is the relevant PMI. We use the PMI for a specific day rather than a moving average because as a Markov process its value on a given day contains all relevant information about current PM risk.[16] The parameter β_1 is of primary interest and indicates how PMI recipients respond to the PMI when compared to nonrecipients, whose baseline responsiveness to disease pressure is captured by β_0. Because of the linear specification, β_1 indicates the overall (average) response function across the domain of PMI values. When R indicates growers' PM treatment interval, $\beta_1 < 0$ implies that PMI recipients tighten intervals as the PMI increases—or, alternatively, stretch intervals as the PMI decreases—and do so more aggressively as the PMI goes from low to moderate to high. Note that in this specification we use the PMI as a proxy for growers' baseline intuition about PM risks (e.g., based on the common pre-PMI advice, "if you like the weather, then so does PM"). While a simpler proxy for this intuition, such as average daily temperature, might be sufficient, using the PMI to proxy for baseline intuition makes this a strong test for recipients' response to the PMI.[17]

To better control for unobservable farmer characteristics that may influence both PMI

[15] We have estimated a similar linear spline specifications and other more flexible functional forms, but believe this simple linear response best characterizes growers' basic responses to PMI information.

[16] Since growers typically receive the PMI at the end of the day on which it is calculated or the next day, we assume that the relevant PMI value for our response models is lagged one day. In other words, we assume that growers adjust today's PM treatment based on yesterday's PMI. To simplify notation, we use to indicate the PMI value lagged one day rather than. Also note that because the timing of PM treatment varies widely by year and region according to weather and other factors, we define a plot-year specific PM window that begins with the first registered PM treatment and ends with the last.

[17] As a simpler proxy for baseline intuition, we have used the two-day change in PMI, which effectively strips the "memory" out of the PMI model. The results are qualitatively robust to this simpler proxy.

use and spraying outcomes in equation (1), we include several control variables, where $Z_i^{Ever} = \max_t(Z_{it}) = \{0,1\}$ indicates that grower i receives the PMI for at least one year. Thus, β_2 controls for any pre-receipt PMI response differences between recipients and nonrecipients and between the few growers who stop receiving the PMI and other growers. Specified this way, the coefficient β_1 is net of any pre- and post-receipt differences and indicates the additional PMI responsiveness that is attributable to receipt of the PMI. Lastly, $\mathbf{x_{ijdt}}$ are control variables related to characteristics of the preceding PM treatment and to annual average PM prevalence,[18] $\boldsymbol{\lambda_{ijd}}$ are growth phase fixed effects,[19] $\boldsymbol{\tau_t}$ are year fixed effects to control for common changes year-to-year in market conditions or PM treatment product offerings, and $\boldsymbol{\mu_j}$ are plot fixed effects to control for time-invariant unobservable features of a given plot. To account for correlations in PM management within grower-year combinations, we cluster the error term ε_{ijdt} at the grower-year level. Since we do not observe the PMI at the location of each plot and instead match plots to the closest weather station from which we can compute the PMI, we use proximity to weather station to weight the variance of each observation such that $\varepsilon_{ijdt} \sim N(0, \sigma_{jt}^2)$ where $\sigma_{jt}^2 = \sigma^2/(dist_{jt}^{-1})$.

While the linear response functions above offer easy parametric tests of growers' response to the PMI, a more flexible specification that allows for nonlinear differences between recipients and nonrecipients potentially adds richer insights—especially since the reported low, medium, and high ranges of the PMI with associated recommendations may introduce nonlinearities in PMI response functions. For this second specification, we focus on growers who switch from nonrecipient to recipient and model PMI responsiveness separately before and after receipt of the PMI using the cubic function:

$$(2) \qquad R_{ijdt} = \alpha_0^Z + \beta_{11}^Z PMI_{ijdt} + \beta_{12}^Z PMI_{ijdt}^2 + \beta_{13}^Z PMI_{ijdt}^3 + \mathbf{x_{ijdt}}'\boldsymbol{\varphi^Z} + \boldsymbol{\lambda_{ijd}^Z} + \boldsymbol{\tau_t^Z} + \boldsymbol{\mu_j^Z} + \upsilon_{ijdt}^Z$$

where Z again indicates receipt of the PMI and the error term is again clustered within grower-years with a variance weighted by proximity to weather station. In contrast to specification (1), this cubic specification uses a polynomial to fit nonlinearities in growers' response to the PMI. Since we estimate (2) separately for growers before and after they receive the PMI, the other coefficients are allowed to change with PMI receipt as well. Rather than displaying the estimated coefficients this specification, we depict the results graphically using marginal response functions across PMI values conditioning on the effect of the control variables and fixed effects on PM treatment response.

We recognize that the response variables we include in these specifications may be jointly determined. For example, anecdotal evidence suggests that these joint responses may be recursive, with the grower first deciding when to treat, then choosing a treatment product and dosage rate. Although we have experimented with a system of equations estimation approach, we remain concerned about the strong assumptions required to identify the system. We therefore opt for an equation-by-equation estimation approach that should yield consistent estimates albeit with some potential loss of efficiency.

Identification strategy and statistical inference. With the main estimating equations for testing how growers respond to the PMI in mind, a few econometric issues of potential concern are worth discussing. The specifications above use comparisons between PMI recipients and nonrecipients to understand how growers adjust treatment strategies in response to disease forecasts. Standard concerns emerge if PMI receipt is endogenous. Although some growers actively sought out the PMI as it became available, many simply started receiving the forecast as part of a standard package of daily weather information—either from their own weather station or from a service provider. As mentioned previously, PMI receipt is therefore less problematic than PMI use, and we use

[18] These include the relative dosage rate of the preceding PM treatment, the product used in the treatment, whether multiple products were used, and—in the case of all non-timing response functions—the time since the last treatment. While these are clearly important control variables given the dynamics of disease management over the course of season, they are not our focus in this analysis, and their corresponding estimates are therefore not reported. We also include plot-year-specific average PMI as an indication of the prevailing PM conditions in a given year.

[19] These fixed effects control for the five growth phases—bud break, shoot growth, bloom, and veraison. The timing of these growth phases is assumed constant across years, but varies across growing region (see Flaherty et al. 1992).

PMI receipt throughout to distinguish PMI from non-PMI growers.

Still, we acknowledge that recipients may differ from their nonrecipient counterparts in systematic but unobservable ways. In our core analysis, our parameter of interest captures how PMI recipients respond to weather risk on a day-to-day basis compared to nonrecipients, and not average differences in disease treatment tendencies (e.g., average treatment intervals, average dosage rates, etc.), which we net out as fixed effects. However, we recognize the possibility that growers' daily reactions to weather risk may also be endogenous to PMI receipt. For instance, growers who are most concerned with powdery mildew may be both exceptionally responsive to their intuition about PM risk on a given day and also more likely to receive the PMI. Additionally, highly concerned growers or those particularly inclined to practice IPM may respond to weather risk with nonspraying activities (e.g., canopy thinning) that are correlated with spraying, and also be more likely to receive the PMI. We therefore formulate an identification strategy accordingly. In the absence of valid instruments, IV estimates may well be worse than any bias introduced by the quasi-endogeneity of PMI receipt, so we seek to improve our identification in several other ways.

First, the specifications above include plot fixed effects $\boldsymbol{\mu}_{\mathbf{j}}$. Thus, the PMI response parameters are estimated within plots, implying that systematic differences across plots that affect PM incidence or management and shape growers' decisions about when, what, and how much to spray (e.g., variety of grape, slope, soil type, etc.) do not confound our estimates of PMI responsiveness. In other words, when a grower begins to receive the PMI, his PMI responsiveness is assessed relative to his pre-receipt responsiveness on the same plot. Second, we include the term $\left(Z_i^{Ever} \times PMI_{ijdt}\right)$ in specifications that pool all growers—recipients and nonrecipients alike. This term enables us to test whether recipients' PMI responsiveness is systematically different than nonrecipients. In the presence of differences in pre-existing responsiveness, this term ensures that our estimate of primary interest, β_1, is net of any such differences and can be easily interpreted.

As a third identification tack, we estimate the above specifications including only growers who switch from nonrecipient to PMI recipient during our window of analysis. Results from these "switchers only" estimations continue to include plot fixed effects, implying that these within estimates compare recipients to their pre-receipt selves. Analogous to this "switchers only" approach is a simple placebo test for growers who never receive the PMI: we randomly assign each of these growers a PMI "false receipt" year and test whether false receipt has any effect on PMI responsiveness. This placebo test helps to ensure that our 'switcher only' results are not driven by the fact that most growers who begin receiving the PMI one year continue to receive it in the future, which could generate spurious correlations with trending changes in PM management. While these switcher only results provide a strong test of growers' response to the PMI, we consider this approach to be a complement to rather than a substitute for the pooled grower approach. We have many observations for switchers in the North Coast region but fewer in the Central Coast and Valley regions.

A final potential identification concern relates to grower decisions that pertain to PM management but are not captured in the PUR data. For example, canopy thinning can be an effective technique for lowering canopy humidity and thereby reducing the risk of a PM outbreak. We have no way of observing such IPM practices that fall outside the PUR data. If these nonspray practices are uncorrelated with PMI receipt, the fixed effects and controls in the specifications above will be sufficient to preserve the unbiasedness of our coefficient of interest. If, on the other hand, growers modify these practices upon receiving the PMI, our estimate of PMI responsiveness will reflect these correlated PM management changes. Since these nonspray practices are intended to reduce the frequency of growers' sprays, such practices are likely to reduce our estimated PMI responsiveness coefficient. In our mind, this does not represent a downward bias in the coefficient but rather underscores that the coefficient captures the total effect of PMI receipt on PMI responsiveness. If part of the mechanism behind this effect involves growers using the PMI to respond with nonspraying practices in addition to spraying, this is an important part of the effect we wish to capture. To keep these practices in perspective, it is important to note that while they are encouraged as part of broader IPM practices, they play a minor role in PM management in comparison to fungicide treatments.

Overall, we are confident that this identification strategy provides plausible causal estimates of the effect of PMI receipt on growers' treatment strategies. Standard complexities and caveats persist because we use observational—not experimental—data and therefore lack pure exogenous variation in access to the PMI. In the interpretation of the results that follows, we strive to keep these complexities in mind, while exploiting the richness of the PUR data.

A few econometric issues regarding statistical inference are also worth noting. As mentioned above, we cluster standard errors by grower-year, which cross-cuts our plot fixed effects. This allows errors to be correlated within grower-year to account for unobservables that shape a particular grower's PM treatment strategies from year to year (e.g., as he is exposed to new information, products, or training). We further adjust these errors for a specific form of heteroscedasticity that arises from our data structure. Because we reconstruct the PMI measures using available weather station data, we have range of distances between a given plot and the closest, matched PMI data series. For much of the analysis, we drop plots that are too far from their matched weather stations (see above and note 10). For the remaining plots, we weight the variance of the error term by the proximity of the weather station as described above. Introducing this structure accounts for the fact the PMI data are measured with error that is proportional to this distance.

A final inference issue relates to multiple inference. Because our aim is to test for multidimensional responses to disease forecast information, we necessarily estimate the specifications above for several different potential margins of adjustment, including when, what, and how much to spray for PM. Under some circumstances this would raise concerns about multiple hypothesis testing and justify corrections for multiple inference (e.g., Anderson 2008). Such corrections assume a degree of independence among outcomes, which makes them inappropriate in our case. The response dimensions we test as outcome variables are all interrelated: An adjustment along one dimension can directly induce an offsetting or complementary adjustment in another. We aim to identify how receiving the PMI impacts overall fungicide treatment strategies, which consist of several intertwined outcomes. Any correction for multiple inference in this case is complicated by (and the motivation for such a correction undermined by) the fact that growers simultaneous adjust along all the response dimensions we test.

Response Dimension 1: When to spray? To estimate these specifications for the first response dimension—when to spray—we use the interval since the last PM treatment within the disease window for each region (Flaherty et al. 1992).[20] By focusing on the treatment interval we restrict our focus to only PM treatment days (see number of PM treatments in table 1). As additional controls, $\mathbf{x_{ijdt}}$, we construct several variables related to the most recent prior PM treatment, including whether this last treatment included synthetic products or multiple products (i.e., "cocktail" treatment) and the dosage rate of this last treatment. To accommodate the wide range of dosage rates and units of measurement across products, we compute the percentile of the dosage rate for each product and each region separately. Where multiple products are used, we combine these dosage percentiles into a weighted average for the combined treatment with each product weighted by the share of the plot treated.

Table 3 contains the pooled linear response results estimated and shown separately for the three growing regions. Each column in this table represents a different response dimension. Throughout, we calibrate the range of the PMI from 0 to 1 so the point estimates reflect the response change induced by a change in the (normalized) PMI from 0 (no risk) to 1 (maximum risk). In the first column of table 3, the coefficient on PMI suggests that nonrecipients tend to respond to perceived higher PM risk by tightening intervals, but this is only weakly statistically significant in the North Coast region where an increase to the maximum PM risk causes growers to tighten intervals by 3 days. This weak baseline response is offset by the response of PMI recipients, who appear to stretch instead of tighten intervals as the PMI increases in both the North and Central Coast regions. While this runs counter to a uni-dimensional response model

[20] Adjusting this PM window at the grower-plot level is possible by excluding the first and last PM treatment of the season for each plot and only including the days between these first and last treatments, which we do.

Table 3. Pooled Regression Results by Growing Region with Plot Fixed Effects, Growth Phase Fixed Effects, and Other (Unreported) Controls (*p*-values Based on Grower-Year Clustering Shown in Parentheses; shading denotes statistical significance)

		Interval (days)	Sulfur {0,1}	Synthetic {0,1}			Dosage Rate (percentile)			Multi-Product {0,1}
				All	Strob.	Sterol.	All	Sul.	Syn.	
North Coast	PMI	−3.08	0.022	−0.077	−0.039	−0.038	−0.057	−0.074	0.036	−0.00086
		(0.140)	(0.670)	(0.170)	(0.180)	(0.470)	(0.140)	(0.160)	(0.650)	(0.980)
	PMI*Recipient	1.74	−0.035	0.015	0.095	−0.08	0.015	−0.026	0.042	0.11
		(0.018)	(0.410)	(0.730)	(0.001)	(0.071)	(0.590)	(0.440)	(0.330)	(0.033)
	PMI*Ever Receive	1.83	−0.0063	0.07	−0.013	0.083	0.065	0.091	−0.039	−0.11
		(0.390)	(0.910)	(0.260)	(0.650)	(0.180)	(0.120)	(0.099)	(0.640)	(0.076)
	N=	20,705	20,705	20,705	20,705	20,705	20,705	14,847	4,394	20,705
	J (# grower-years)=	264	264	264	264	264	264	258	198	264
	R Sq	0.028	0.270	0.230	0.057	0.190	0.095	0.130	0.120	0.077
Central Coast	**PMI**	−0.94	0.086	−0.054	0.033	−0.087	0.14	0.062	0.029	0.039
		(0.460)	(0.140)	(0.280)	(0.350)	(0.099)	(0.016)	(0.340)	(0.590)	(0.400)
	PMI*Recipient	2.84	0.0043	−0.00081	0.03	−0.029	−0.0036	0.11	−0.0031	−0.02
		(0.012)	(0.870)	(0.970)	(0.170)	(0.150)	(0.930)	(0.002)	(0.990)	(0.600)
	PMI*Ever Receive	−2.1	−0.065	0.031	−0.03	0.059	−0.11	−0.11	−0.05	0.04
		(0.220)	(0.280)	(0.540)	(0.410)	(0.290)	(0.110)	(0.110)	(0.720)	(0.450)
	N=	7,209	7,209	7,209	7,209	7,209	7,209	6,550	429	7,209
	J (# grower-years)=	106	106	106	106	106	106	97	51	106
	R Sq	0.046	0.45	0.39	0.17	0.2	0.18	0.25	0.42	0.24
Valley	**PMI**	−0.39	0.16	−0.16	−0.13	−0.025	−0.024	−0.043	−0.081	−0.11
		(0.720)	(0.042)	(0.054)	(0.048)	(0.620)	(0.490)	(0.079)	(0.110)	(0.031)
	PMI*Recipient	−0.46	−0.13	0.14	−0.0011	0.14	−0.078	0.0015	−0.1	−0.26
		(0.690)	(0.003)	(0.001)	(0.970)	(0.001)	(0.047)	(0.970)	(0.390)	(0.000)
	PMI*Ever Receive	1.85	−0.013	0.00091	0.13	−0.13	0.085	0.035	0.15	0.29
		(0.210)	(0.880)	(0.990)	(0.053)	(0.050)	(0.150)	(0.480)	(0.280)	(0.000)
	N=	12,128	12,128	12,128	12,128	12,128	12,128	10,447	1,549	12,128
	J (# grower-years)=	156	156	156	156	156	156	150	106	156
	R Sq	0.091	0.17	0.16	0.066	0.13	0.068	0.15	0.2	0.032

focused on interval adjustment, allowing for multiple margins of adjustment—which growers face in practice—reconciles this result. More aggressive responses along these other margins may necessitate this seemingly counterintuitive timing response. For example, more potent products typically come with longer minimum intervals. Finally, in the first column, note that the "ever receive" interaction term is insignificant, suggesting that there are no systematic responsiveness differences between recipients prereceipt and nonrecipients. Results for switchers only (table 4) show consistent interval responses—albeit with weak significance of the North Coast coefficient—suggesting that these patterns are not due to unobservable time-invariant differences between recipients and nonrecipients. This table also reports the placebo test results, which suggest that the interval response results along with all others are attributable to receipt of the PMI and not due to spurious correlation with temporal trends in disease management practices.

The marginal interval functions, depicted in figure 1, are based on the cubic response function in specification (2). These graphical results allow for direct comparison of the response function for PMI recipients and nonrecipients. These functions show the heterogeneity in growers' interval response across the three regions. The interval response of North Coast and Valley switchers is statistically comparable with and without the PMI, while Central Coast growers are more likely to stretch intervals when the PMI is either low or high after they begin receiving the PMI. At low PMI values, this may be precisely the intended response. At high PMI values, this is likely a by-product of more aggressive responses along another dimension.

Response Dimension 2: What to spray? We consider growers' product choice as the next response variable in the specifications above. Sulfur products are growers' default PM treatment and are relatively inexpensive; synthetic fungicides can offer greater protection but are more expensive. Although growers can also opt for biological and contact treatments, over 90% of PM treatments consist of sulfur or synthetic fungicides, which are made up of two classes of products: strobilurins and sterol inhibitors. In our analysis of this response dimension, we focus on the decision to spray sulfur and synthetic products as distinct from all other products. We estimate linear probability models of these product choices. The sulfur and synthetic columns of table 3 suggest that Valley recipients shift from sulfur to synthetics (especially sterol inhibitors) as the PMI increases. In contrast, North and Central Coast growers shift between synthetic products, specifically, from sterol inhibitors to strobilurins, which offer better protection against PM for many susceptible varieties. While the point estimates for switchers only (table 4) show largely consistent patterns, they are estimated less precisely—only the strobilurin response in the North Coast region continues to be statistically precise.

The marginal response functions in figure 2 show some similarities between Central Coast and Valley growers, who are more likely to rely on sulfur when the PMI is low or high. Again, this response at high PMI levels may be a by-product of other adjustments made when PM risk is high. In both regions, switchers' product choices are completely nonresponsive to the PMI before they begin receiving the forecast and become responsive afterwards. North Coast growers, on the other hand, strongly shift toward synthetic products when the PMI gets high, an effect that is masked by the linear response results in tables 3 and 4. Moreover, these growers' product choices are more responsive to PM risk even before they receive the PMI, which is likely a reflection of greater sophistication in baseline treatment strategies as justified by the significantly higher value-at-risk for these growers (see table 1). The marginal response functions for synthetic products show a strong shift towards strobilurins among North Coast switchers and continued nonresponsiveness of baseline synthetic product choices among Central Coast and Valley switchers (figure S1).

Response Dimension 3: How much to spray? As the third dimension of PMI response, we consider growers' decision of how much to spray. We capture this with two different variables: dosage percentiles and cocktails consisting of multiple products. Because product labels often provide a range of recommended dosage rates, growers have some flexibility to adjust their dosage rates within regulatory range. They are not, however, supposed to exceed the maximum dosage rate listed on the label. As explained above, we compute dosage percentiles across all the treatments in our data by product and by region. Where multiple products are used

Table 4. Pooled Regression Results by Growing Region with Plot Fixed Effects, Growth Phase Fixed Effects, and Other (Unreported) Controls Including Only Growers Who Switch from Not Receiving to Receiving the PMI (*p*-values Based on Grower-Year Clustering Shown in Parentheses; shading denotes statistical significance)

Switchers Only		Interval (days)	Sulfur {0,1}	Synthetic {0,1} All	Synthetic {0,1} Strob.	Synthetic {0,1} Sterol.	Dosage Rate (percentile) All	Dosage Rate (percentile) Sul.	Dosage Rate (percentile) Syn.	Multi-Product {0,1}
North Coast	**PMI**	−1.35	0.0029	−0.0017	−0.047	0.045	0.017	0.02	0.029	−0.088
		(0.024)	(0.930)	(0.960)	(0.004)	(0.280)	(0.370)	(0.290)	(0.340)	(0.065)
	PMI*Recipient	1.01	−0.031	0.015	0.075	−0.06	0.027	0.011	0.037	0.093
		(0.130)	(0.450)	(0.680)	(0.000)	(0.180)	(0.240)	(0.700)	(0.350)	(0.078)
	N=	15,105	15,105	15,105	15,105	15,105	15,105	10,598	3,376	13,501
	J (# grower-years)=	173	173	173	173	173	173	171	148	153
	R Sq	0.029	0.290	0.260	0.078	0.210	0.130	0.150	0.150	0.110
	Placebo	*1.68*	*−0.024*	*0.036*	*0.11*	*−0.073*	*−0.049*	*−0.11*	*0.023*	*0.04*
		(0.440)	*(0.770)*	*(0.700)*	*(0.140)*	*(0.120)*	*(0.320)*	*(0.032)*	*(0.810)*	*(0.620)*
Central Coast	**PMI**	−2.33	0.055	−0.046	−0.034	−0.015	0.03	0.028	0.078	0.15
		(0.140)	(0.160)	(0.098)	(0.160)	(0.470)	(0.340)	(0.400)	(0.540)	(0.001)
	PMI*Recipient	6.36	0.094	−0.1	−0.088	−0.0093	−0.069	−0.074	−0.76	0.095
		(0.000)	(0.360)	(0.270)	(0.250)	(0.760)	(0.160)	(0.190)	(0.023)	(0.390)
	N=	2,517	2,517	2,517	2,517	2,517	2,517	2,432	42	2,517
	J (# grower-years)=	29	29	29	29	29	29	27	11	29
	R Sq	0.061	0.24	0.3	0.15	0.19	0.17	0.16	0.63	0.29
	Placebo	*−0.5*	*0.03*	*−0.096*	*0.0086*	*0.018*	*0.072*	*0.025*	*0.077*	*−0.055*
		(0.680)	*(0.320)*	*(0.740)*	*(0.800)*	*(0.600)*	*(0.150)*	*(0.410)*	*(0.510)*	*(0.130)*
Valley	**PMI**	1.63	0.17	−0.16	−0.076	−0.088	0.037	0.041	−0.15	−0.032
		(0.320)	(0.004)	(0.004)	(0.018)	(0.034)	(0.400)	(0.440)	(0.370)	(0.790)
	PMI*Recipient	−1.04	−0.055	0.029	−0.0064	0.036	0.0047	0.0099	0.12	0.042
		(0.490)	(0.320)	(0.600)	(0.880)	(0.450)	(0.940)	(0.890)	(0.530)	(0.670)
	N=	1,379	1,379	1,379	1,379	1,379	1,379	1,182	176	1,159
	J (# grower-years)=	47	47	47	47	47	47	46	26	39
	R Sq	0.093	0.2	0.21	0.11	0.16	0.16	0.17	0.4	0.11
	Placebo	*−0.17*	*0.033*	*0.02*	*−0.0022*	*−0.018*	*0.058*	*0.028*	*0.13*	*−0.016*
		(0.900)	*(0.580)*	*(0.750)*	*(0.960)*	*(0.710)*	*(0.048)*	*(0.390)*	*(0.170)*	*(0.750)*

Placebo estimates are taken from a similar regression of nonrecipients in which they are each assigned a random, false PMI receipt year.

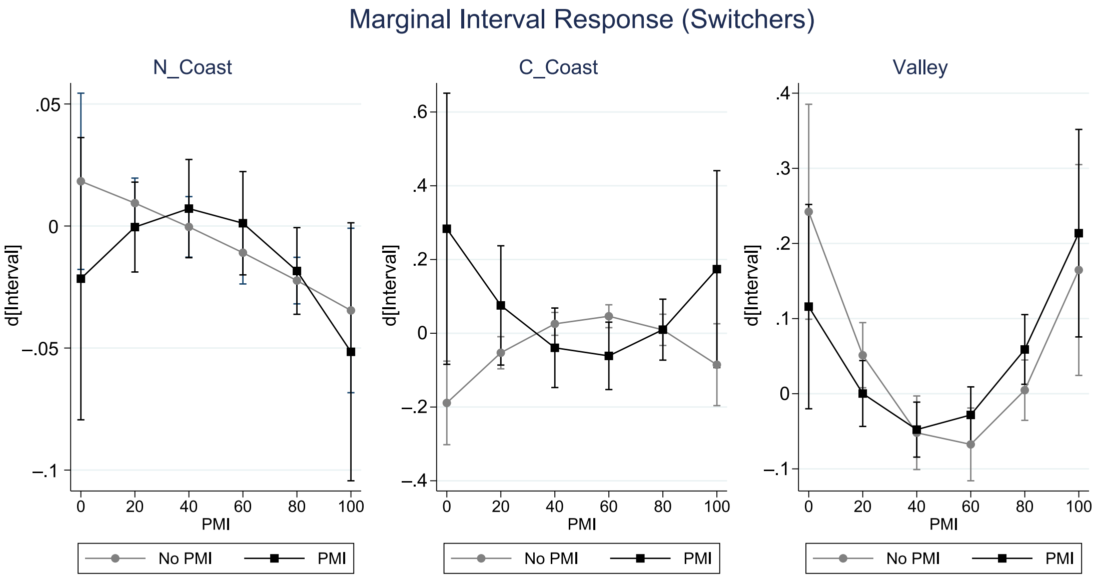

Figure 1. Marginal interval response functions for switchers before and after they received the PMI with 90% confidence intervals (based on cubic response specification)

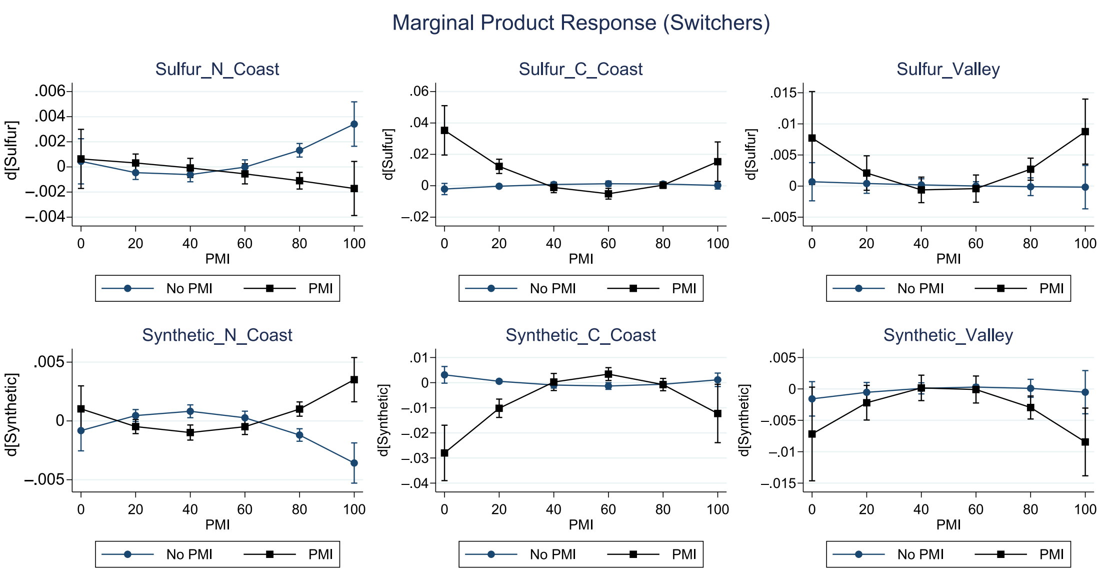

Figure 2. Marginal response functions for spraying sulfur fungicides for switchers with 90% confidence intervals (based on cubic response specification)

in a single treatment, we compute a weighted average percentile using percent of the plot treated for each product as the weight.

Table 3 provides mixed evidence of response along this dimension across the three regions. North Coast growers are more likely to mix products for better protection as the PMI increases. Central Coast growers seem to increase sulfur dosage rates with the PMI. In contrast, Valley growers appear to reduce dosage rates and product mixing as the PMI increases, a response that again may be a by-product of aggressive responses along other dimensions (e.g., shifting to synthetic products). Results for switchers (table 4) are robust for North Coast growers, who show a consistent multiproduct response before and after PMI receipt as they do when

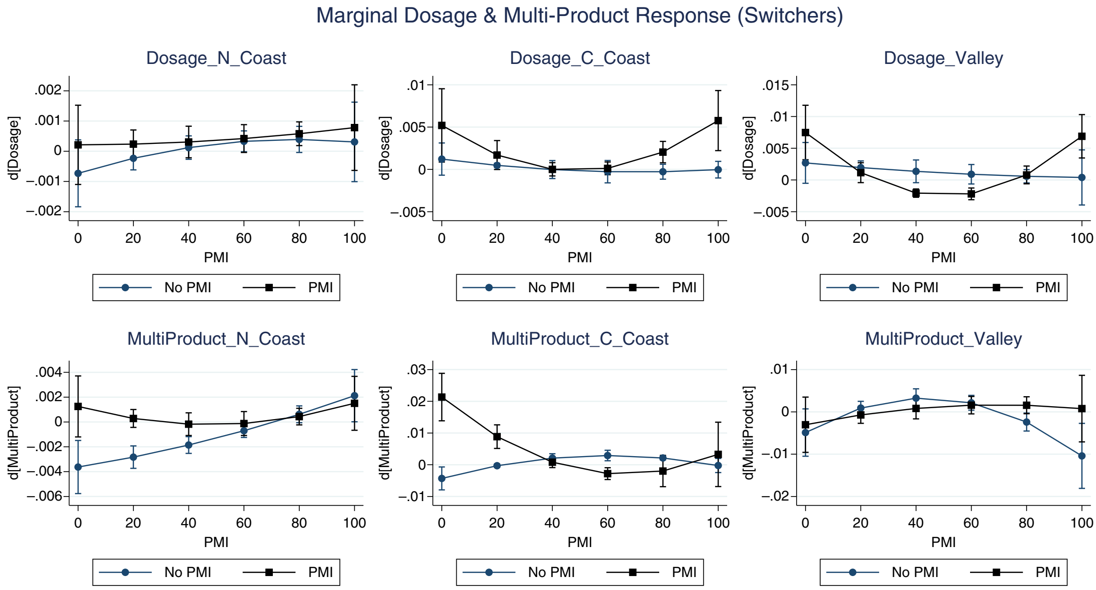

Figure 3. Marginal response functions for dosage rate and multiproduct responses for switchers with 90% confidence intervals (based on cubic response specification)

pooled with nonrecipients. Central Coast and Valley switchers, in contrast, show less responsiveness than the pooled results. While this may imply that some of the parallel results in table 3 reflect pre-existing differences between recipients and nonrecipients, the "ever receive" interaction terms in that table suggest that this may only be the case for the multiproduct response in the Valley. The cubic response functions in figure 3 show that while switchers are generally nonresponsive or linearly responsive to the PMI before receipt, they are more likely to adjust responses to low and high PMI values after receipt. Higher dosage rate responses to high PMI values are also evident in Central Coast and Valley regions, responses that are masked by the linear specification in table 4. In the North Coast region, where we have the most switchers, we can breakdown the dosage rate response by product category and find that dosage rates increase with the PMI for sulfur and decrease for synthetic products (figure S2).

Taken as a whole, these results convey four broad messages. First, growers seem to take the PMI seriously and adjust PM treatment strategies in response to receiving the forecast information. Second, treatment timing is but one response dimension for PMI recipients and a seemingly minor one at that. This contradicts the conventional wisdom among those who promote the PMI as an IPM tool that growers respond by adjusting treatment intervals and undermines the seemingly straightforward logic that the PMI reduces overall fungicide usage to the benefit of the environment. Third, because the response dimensions are interrelated components of growers' overall disease management strategies, PMI responses in one dimension appear to induce other responses in other dimensions: what appears to be a counterintuitive PMI response when considering one dimension in isolation of the others may simply be a by-product of sensible responses in these other dimensions. Fourth, the responses we detect show substantial heterogeneity across winegrape regions of California. This reflects not only differences in PM risks and susceptibility of different varieties to PM but also pronounced differences in the value of the winegrapes produced in these regions.

How heterogeneous and interrelated are response dimensions? The preceding analyses pool growers within growing regions based on the observation that disease management and other viticulture practices vary more between than within these regions. This captures a major source of heterogeneity in PMI responses, but sources of grower-level heterogeneity are likely to be pronounced as well. This is particularly true when it comes to *interactions* between different margins of

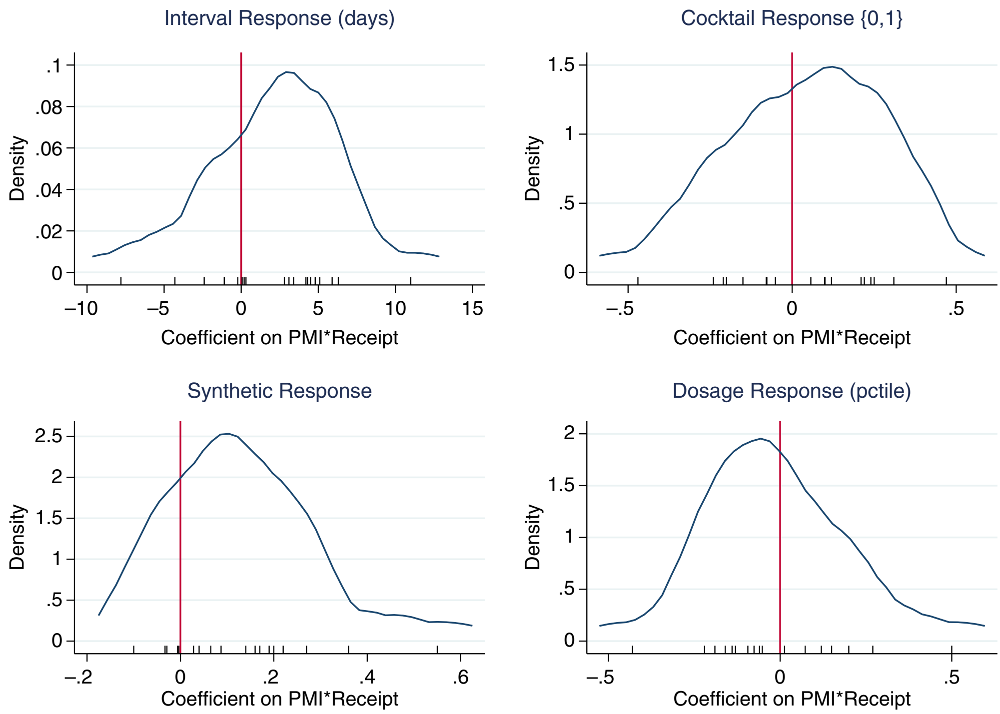

Figure 4. Kernel densities of estimated coefficients on linear *PMI*Recipient* interaction from grower-specific regressions for 19 switchers with more than 50 observations. Rugplot along bottom shows distribution of statistically significant coefficients (*p*-values<.10)

adjustment: growers formulate PM management strategies that consist of a bundle of adjustments—and the PMI likely affects the composition of these strategy bundles.

Since the temporal resolution of our data is high and many growers manage multiple plots, we can estimate separate grower-specific regressions for nearly half of our growers. Although such an empirical approach has limitations, it also complements the pooled-grower results presented thus far. Figure 4 shows distributions of estimated coefficients on term PMI*Receipt (i.e., akin to β_1 in equation) from these grower-specific regressions.[21] While some of the trends that emerge in the pooled regressions above also come through in these grower-specific results, these distributions also suggest pronounced heterogeneity in these responses.

Grower-specific results enable us to characterize the interrelationships between response dimensions. To jointly evaluate how these response dimensions function as either complements or substitutes for each other, we use factor analysis to identify dominant combinations of responses based on these grower-specific coefficients. Two distinct response regimes emerge from this analysis. First, some growers tend to respond to an increasing PMI by mixing multiple products in order to broaden their protection against a PMI outbreak and by increasing dosage rates of these products without adjusting spray intervals.[22] In the second response regime, growers tend to shift to synthetic products, increase dosage rates, and lengthen intervals to reflect the increased toxicity of these products.[23] These predominant response combinations are not mutually exclusive, but growers' factor scores for these two response combinations suggest that they tend to adopt one approach or the other rather than a mix of the two. Although these responses are likely to be correlated within growing regions, residual heterogeneity in response

[21] We only estimate grower-specific regressions for growers with more than fifty observations, but many of the forty-three growers with at least this many observations have well over 100 observations. Of these growers, nineteen have received the PMI during some year covered by our data.

[22] Scoring coefficients for this first factor are interval = 0.09, cocktail = 0.36, synthetic = −0.06, and dosage = 0.38.

[23] Scoring coefficients for this second factor are interval = 0.31, cocktail = −0.15, synthetic = 0.31, and dosage = 0.20.

functions complicate the detection of overall response tendencies per the above analysis.

What is the net environmental impact of multidimensional PMI responses?

Our analysis allows for growers to respond to the PMI by adjusting when, what, and how much to spray for PM. We find evidence that growers' response to the PMI involves adjustments along all three dimensions and some shifting between dimensions as the PMI passes from low to moderate to high values. This multidimensional and nonlinear response makes the environmental impact of PMI receipt difficult to predict. Initial projections of the potential environmental impact of the PMI assumed that treatment intervals were growers' only margin of adjustment. As shown in the empirical analysis above, once we allow for additional margins of adjustment, these unambiguous predictions of environmental benefits may be misleading. In this section, we test the net environmental impact of PMI usage knowing that growers' response is both multidimensional and nonlinear.

To do this, we use the Pesticide Use Risk Evaluation (PURE) model, which merges plot level PUR data with physical, soil, topographical, and meteorological characteristics to compute per acre pesticide use risk scores for five dimensions: surface water, groundwater, soil, air, and pollinators (bees) (Zhan and Zhang 2012). Since our primary focus is on PM treatment strategies, we construct risk scores separately for PM treatments (i.e., applications of products used for PM control) and all other treatments. Following Zhan and Zhang (2012), we use a log transformation of these scores to construct PURE indexes that range from 0 to 100 for each dimension and growing region and construct an integrated PURE risk index based on the most at-risk environmental dimension for each plot-year combination.[24] We have annual PURE risk indexes for 1,344 plots managed by growers who participated in our survey.[25] In the final data we use, we observe these indexes for an average of six years for each plot for a total of 7,718 observations. Before proceeding with the analysis, it is important to note that several of the environmental factors included in these PURE risk indexes lie outside growers' control, which potentially dilutes the relationship between receipt of the PMI and the indexes and thereby strengthens the test for a difference.

Table 5 shows mean (plot-year) PURE risk indexes for each environmental dimension for PMI recipients and nonrecipients. These indexes are further broken down by PM treatments and non-PM (i.e., all other) treatments. PMI recipients have significantly higher integrated PURE indexes than nonrecipients based on PM treatments, but their integrated PURE scores are statistically indistinguishable for non-PM treatments. The comparison by environmental dimension is broadly consistent with this pattern, albeit with some variation. Three of four significant differences for PM treatments indicate increasing risk. Interestingly, we see evidence of spillover effects on non-PM treatments when using the PURE scores by dimension. While these unconditional mean comparisons may indicate that the net environmental impact of growers' multidimensional PMI responses is negative, they are obviously only suggestive and do not control for annual and regional variation in PM treatment strategies.

To build on the unconditional t-tests in Table 5 and assess the impact of receiving and using the PMI on these risk scores more carefully, we restrict our focus to switchers who begin receiving the PMI during the years covered by our dataset and use an event study approach based on the following specification

$$(3) \qquad Risk^x_{ijct} = \beta_0 + \beta_1 Z^{years}_{it} + \beta_2 \left(Z^{years}_{it}\right)^2 + \boldsymbol{\gamma}_t + \boldsymbol{\mu}_c + (v_j + \varepsilon_{ijt})$$

where $x = \{$integrated, surface water, groundwater, soil, air, pollinator$\}$, Z^{years}_{it} is the number of years before/after grower i first received the PMI, $\boldsymbol{\gamma}_t$ is a year fixed effect, $\boldsymbol{\mu}_c$ is a county fixed effect, and v_j

[24] We follow Zhan and Zhang (2012) in the construction of this integrated index, which is taken as the highest PURE index of the five dimensions for each plot-year observation. In their words, this is meant to showcase "the most vulnerable environmental compartment" (105).

[25] Note that we are able to include several surveyed growers who are not included in the response analysis above due to PMI data limitations. In the case of this PURE analysis, we do not use PMI data and are therefore not constrained by PMI limitations. Note further that errors in reported plot area treated with a given application in the PUR data create a handful of extreme values in raw PURE scores. We trim these scores at the 99^{th} percentile to avoid extreme value problems in the subsequent analysis.

Table 5. Mean PURE Risk Scores for PMI Nonrecipients and Recipients with p-values Based on t-test of Different Means Assuming Equal Variances in Parentheses

		Treatments Included			
		Powdery Mildew		Non-Powdery Mildew	
Risk Score	Received PMI (plots)	Means	Difference (*p*-value)	Means	Difference (*p*-value)
Integrated	**No (N=4,300)**	41.9	3.1	54.3	0.6
	Yes (N=3,149)	45	(0.000)	54.9	(0.24)
Surface Water	**No**	7.3	−0.2	15.5	−4.2
	Yes	7.1	(0.58)	11.3	(0.000)
Groundwater	**No**	14.5	6.5	30.6	−2.4
	Yes	21	(0.000)	28.2	(0.001)
Soil	**No**	31.9	5.4	41.1	5.1
	Yes	37.3	(0.000)	46.2	(0.000)
Air	**No**	28.5	−3.3	41.2	1.4
	Yes	25.2	(0.000)	42.6	(0.014)
Pollinator	**No**	13	5.7	14.3	−0.8
	Yes	18.7	(0.000)	13.5	(0.10)

is a plot random effect. We estimate this model separately for before and after PMI receipt and—to provide greater flexibility for regional differences in this net environmental effect—also separately estimate the model by growing region (and drop μ_c).

It is important to note that our identification strategy here is different than the strategy we use above to understand the impact of the PMI on PM management. For this event study estimation, we cannot estimate how receipt of the PMI affects responses to daily changes in PM risk because we cannot use daily variation in the PMI as a separate source of information to distinguish responses of recipients from nonrecipients. Identification therefore relies on the assumption that growers' receipt of the PMI is uncorrelated with the adoption of other disease management practices that could affect treatment for PM. Based on anecdotal reports from growers and extension agents, we are confident that the majority of our estimated effects are attributable to direct responses to PMI receipt and not due to other correlated management changes, but we interpret the results with some caution because we are unable to completely rule out this possibility.

The results of this event study estimation are depicted graphically as conditional plots in figures 5 and 6. In the first figure, which pools together growers from all growing regions and includes county fixed effects, it is clear that the dominant environmental effect associated with PMI receipt is *negative* (i.e., higher risk indexes). The integrated PURE index trends slightly downward pre-receipt, but increases sharply in the years after first receipt. A similar pattern is evident for other risk dimensions with the possible exception of the air index (consistent with table 5). When conducting this event study analysis separately by growing region, the integrated PURE risk index shows a similar pattern (figure 6), albeit with heterogeneity that is consistent with the regional PMI response heterogeneity shown above.[26]

Since an interpretation of absolute changes of these risk scores is complicated by the construction of the PURE model, a relative interpretation must suffice. Measured by the integrated index, environmental risks increase 40% in the seven years after initial PMI receipt. This effect may not be due entirely to growers' responses to PMI information for the reasons discussed above, but the earlier evidence that PMI recipients

[26] The plot of these conditional functions disaggregated by both region and PURE environmental dimension show consistent patterns. These are available upon request.

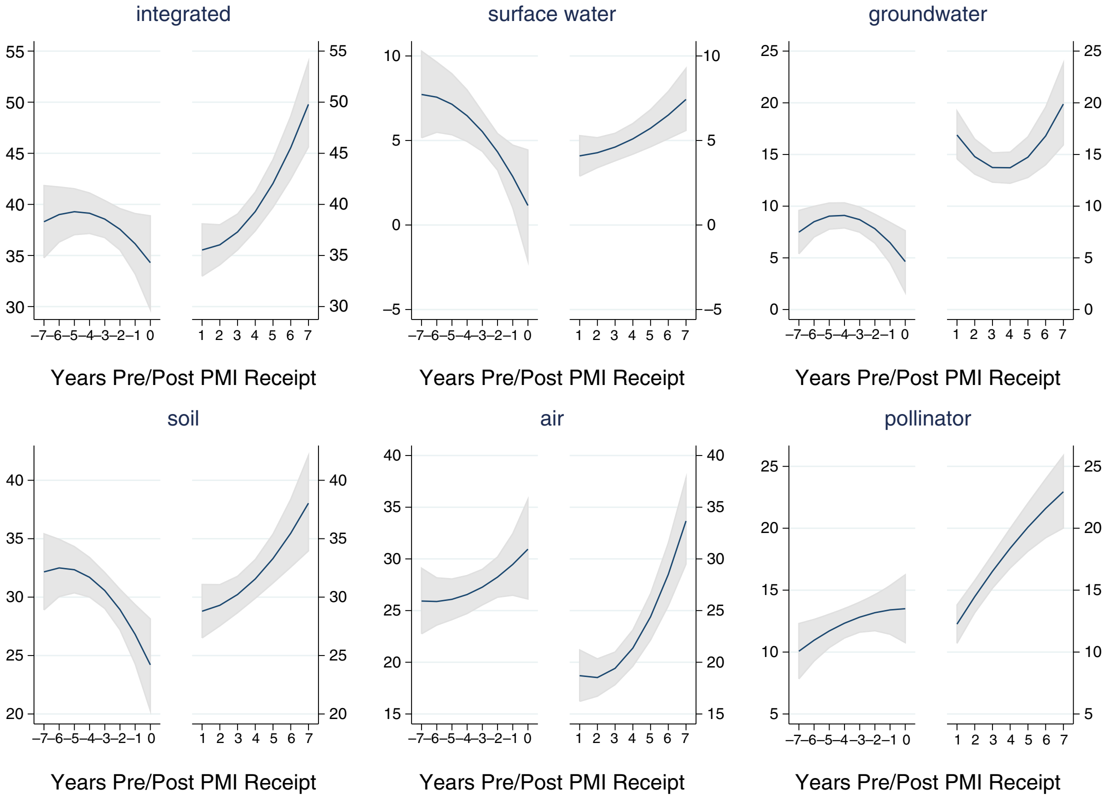

Figure 5. Event study of impact PMI receipt on PURE risk indexes for different environmental dimensions controlling for year and county fixed effects and plot random effects including only switchers. Gray bands indicate 90% confidence intervals

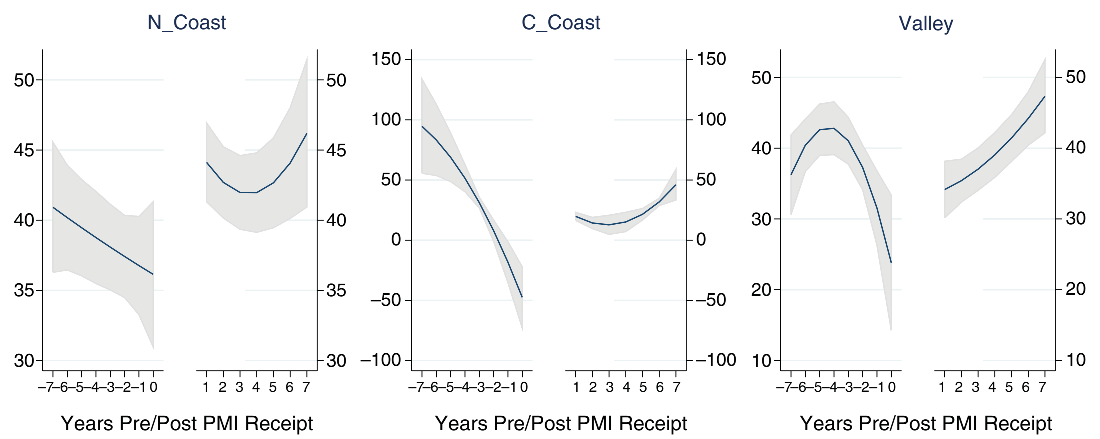

Figure 6. Event study of impact PMI receipt on the integrated PURE risk index estimated separately by growing region and controlling for year fixed effects and plot random effects including only switchers. Gray bands indicate 90% confidence intervals

shift to more aggressive treatment strategies along several dimensions in response to this information is consistent with the association between PMI receipt and environmental risk being driven by growers' direct responses to PMI information.

Conclusions

How agents respond to new or better forecast information often has direct implications for market and broader social outcomes. In agriculture, improved forecasts can have direct welfare implications for producers and consumers at both local and societal levels. Many integrated pest management (IPM) approaches, for example, aim to reduce pesticide usage and related environmental impacts by providing growers with better, more precise pest and disease information. Better forecast information can certainly benefit growers and improve their capacity to manage disease and pests effectively, but we must consider the net effect of growers' adjustments along multiple dimensions in order to properly understand broader environmental benefits associated with this information.

Using the case of California winegrape growers and high resolution panel data that includes plot-level powdery mildew treatments, we characterize growers' response to a popular PM risk model that generates forecasts in the form of a daily risk index (PMI). This index was designed and continues to be promoted as a way for growers to optimize the timing of their powdery mildew treatment. We find that growers indeed modify their treatment strategies in response to the PMI, but adjustments to the timing of treatment are but one of several response dimensions, including what and how much to spray conditional on spraying. We also find substantial spatial heterogeneity, which is sensible given the dramatic spatial differences that exist among winegrape regions (e.g., differential susceptibilities of grape varieties to powdery mildew, differential harvest values, etc.). This heterogeneity and the strong interrelationships between response dimensions complicate the estimation of growers' adjustments of disease management strategies in response to the PMI. Estimating the net environmental impact of these complex, multidimensional responses is, however, more straightforward. These results suggest that the net environmental impact of the PMI—as mediated by growers' response functions—may actually be negative. Specifically, for growers who are receiving the PMI, plot-level environmental risks increase markedly as compared to these same growers before they receive the forecast.

Two caveats to this analysis are worth noting. First, although we have been able to characterize growers' response to the PMI, we are unable to document the impact of their response on disease control. Based on widespread evidence from field trials, the PMI seems to enable growers to more effectively control powdery mildew in their vineyards. Indeed, our survey evidence strongly suggests that growers using the PMI value it as an important tool in their decision making. With plot-level measures of the efficacy of powdery mildew management, much more could be done to assess whether the risk responses we document in this analysis are rational given their effect on disease control. We suspect they are. Ultimately, such an assessment could also evaluate the private versus public costs and benefits associated with producers' responses to forecast information – and generate important insights into the value of risk information. Second, all of the growers included in our sample grow winegrapes. While the sample spans high and low value winegrape regions in California—and the analysis suggests this heterogeneity shapes growers' response to disease forecasts—there is yet greater production heterogeneity between winegrape, table grape, and raisin growers. A broader sample of growers drawn from across these distinct grower types might support more explicit modeling of the factors that influence how growers respond to disease forecast information.

A final conjecture suggested by this analysis seems to parallel the "alarmist" response to risk information documented by Viscusi (1997). The negative net environmental impact of improved disease forecasts we find may be due in part to a double "risk response asymmetry": first, in the risk responses of model builders and, second, in the risk responses of growers. In the calibration of forecasting models, model builders rationally fear false negative predictions more than false positive predictions. This rational response emerges because false negatives can bring dire consequences when the underlying stochastic process is characterized by nonlinearities that can lead to explosive risk episodes, while false positives have more mundane effects and modest costs. Moreover, a basic observability problem pushes model builders to choose conservative calibration of their forecast models: false negatives are readily observable by PMI users but false positives are not. The second asymmetric risk response occurs with growers, who appear to respond aggressively when the PMI is high, but may continue with "business as usual" when it is low. As long as growers and modelers do not take into account each other's asymmetric risk response, these have the effect of magnifying the overall risk response relative to what it would be under risk neutrality. At a time when increasingly sophisticated forecasts are flourishing in many fields—often with the intent of improving market and societal outcomes—this double asymmetry merits attention from economists and psychologists alike.

Supplementary Material

Supplementary material is available at http://oxfordjournals.our_journals/ajae/online.

References

Anderson, M.L. 2008. Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American Statistical Association* 103 (484): 1481–95.

Bongiovanni, R., and J. Lowenberg-DeBoer. 2004. Precision Agriculture and Sustainability. *Precision Agriculture* 5 (4): 359–87.

Broome, J., and K. Warner. 2008. Agro-environmental Partnerships Facilitate Sustainable Wine-Grape Production and Assessment. *California Agriculture* 62 (4): 133–41.

Cassman, K.G. 1999. Ecological Intensification of Cereal Production Systems: Yield Potential, Soil Quality, and Precision Agriculture. *Proceedings of the National Academy of Sciences* 96 (11): 5952–59.

Davidson, C. 2004. Declining Downwind: Amphibian Population Declines in California and Historical Pesticide Use. *Ecological Applications* 14 (6): 1892–902.

Epstein, L. 2006. California's Pesticide Use Reports and Trends in Pesticide Use. *Outlooks on Pest Management* 17 (4): 148–54.

Epstein, L., and S. Bassein. 2003. Patterns of Pesticide Use in California and the Implications for Strategies for Reduction of Pesticides. *Annual Review of Phytopathology* 41: 351–75.

Fabre, F., M. Plantegenest, and J. Yuen. 2007. Financial Benefit of Using Crop Protection Decision Rules Over Systematic Spraying Strategies. *Phytopathology* 97 (11): 1484–90.

Flaherty, D.L., L.P. Christiansen, W.T. Lanini, J.J. Marois, P.A. Phillips, and L.T. Wilson. 1992. *Grape Pest Management*. 2nd ed. Oakland: University of California, Division of Agriculture and Natural Resources.

Friedland, W. 2002. Agriculture and Rurality: Beginning the "Final Separation"? *Rural Sociology* 67 (3): 350–71.

Fuller, K.B., J.M. Alston, and O. Sambucci. 2014. The Value of Powdery Mildew Resistance in Grapes: Evidence from California. *Wine Economics and Policy* 3 (2): 90–107.

Goodhue, R., R. Green, D. Heien, and P. Martin. 2008. Current Economic Trends in the California Wine Industry. *ARE Update* 3 (3): 7–9.

Gubler, W.D., M.R. Rademacher, S.J. Vasquez, and C.S. Thomas. 1999. Control of Powdery Mildew Using the UC Davis Powdery Mildew Risk Index. *APS Net: Plant Pathology On-Line*.

Heien, D., and P. Martin. 2003. California's Wine Industry Enters New Era. *California Agriculture* 57 (3): 71–75.

Khanna, M., and D. Zilberman. 1997. Incentives, Precision Technology and Environmental Protection. *Ecological Economics* 23 (1): 25–43.

Lichtenberg, E. 2002. Agriculture and the Environment. *Handbook of Agricultural Economics* 2: 1249–313.

Lybbert, T.J., and W.D. Gubler. 2008. California Wine Grape Growers' Use of Powdery Mildew Forecasts. *ARE Update* 11 (3): 11–14.

McBratney, A., B. Whelan, T. Ancev, and J. Bouma. 2005. Future Directions of Precision Agriculture. *Precision Agriculture* 6 (1): 7–23.

Moffitt, L.J., R.L. Farnsworth, L.R. Zavaleta, and M. Kogan. 1986. Economic Impact of Public Pest Information: Soybean

Insect Forecasts in Illinois. *American Journal of Agricultural Economics* 68 (2): 274–79.

Mumford, J.D., and G.A. Norton. 1984. Economics of Decision Making in Pest Management. *Annual Reviews in Entomology* 29 (1): 157–74.

Reynolds, P., J. Von Behren, R. Gunier, D. Goldberg, A. Hertz, and M. Harnly. 2002. Childhood Cancer and Agricultural Pesticide Use: An Ecologic Study in California. *Environmental Health Perspectives* 110 (3): 319.

Stummer, B., I. Francis, T. Zanker, K. Lattey, and E. Scott. 2005. Effects of Powdery Mildew on the Sensory Properties and Composition of Chardonnay Juice and Wine when Grape Sugar Ripeness Is Standardised. *Aust. J. Grape Wine Res* 11 (1): 66–76.

Swinton, S.M., and R.P. King. 1994. The Value of Pest Information in a Dynamic Setting: The Case of Weed Control. *American Journal of Agricultural Economics* 76 (1): 36–46.

Thomas, C.S., W.D. Gubler, and G. Leavitt. 1994. Field Testing of a Powdery Mildew Disease Forecast Model on Grapes in California. *Phytopathology* 84: 1070.

UC Agriculture and Natural Resources. 2005. *UC Cooperative Extension Helps Farmers Reduce Fungicide Use in the San Joaquin Valley*.

Viscusi, W.K. 1997. Alarmist Decisions with Divergent Risk Information. *The Economic Journal* 107 (445): 1657–70.

Warner, K. 2007. The Quality of Sustainability: Agroecological Partnerships and the Geographic Branding of California Winegrapes. *Journal of Rural Studies* 23 (2): 142–55.

Weber, E., W.D. Gubler, and A. Derr. 1996. Powdery Mildew Controlled with Fewer Fungicide Applications. *Practical Winery & Vineyard* Jan/Feb.

Ypema, H., and W. Gubler. 1997. Long-Term Effect of Temperature and Triadimefon on Proliferation of *Uncinula Necator*: Implications for Fungicide Resistance and Disease Risk Assessment. *Plant Disease* 81 (10): 1187–92.

Zhan, Y., and M. Zhang. 2012. PURE: A Web-Based Decision Support System to Evaluate Pesticide Environmental Risk for Sustainable Pest Management Practices in California. *Ecotoxicology & Environmental Safety* 82: 104–13.