

Sustainable Pest Management, Beliefs, and Behavior: A Dynamic Bioeconomic Analysis¹

D. Adeline Yeh, Miguel I. Gómez, and C.-Y. Cynthia Lin Lawell

Abstract

Pest management is an important concern for agricultural producers. We develop a novel dynamic bioeconomic analysis framework that combines numerical dynamic optimization and dynamic structural econometric estimation, and apply it to analyze the sustainable management of Spotted Wing Drosophila (SWD) by lowbush blueberry producers in Maine. Our framework enables us not only to solve for the optimal management strategy, but also to understand the beliefs and perceptions of growers that underlie and rationalize their actual spraying and harvesting decisions, and therefore to propose possible improvements in management practices that align with their beliefs and perceptions. Results show that a sustainable pest control alternative -- early harvesting -- can be part of an optimal management strategy, and that spraying insecticide is not optimal in most cases when pest pressure is low. In contrast, data on the actual decisions of growers show that growers tend to spray earlier and more often than is optimal, and harvest later than is optimal. We find that the actual behavior of growers is rationalized by perceptions and beliefs about the spray cost and yield loss that differ from what economic data, expert opinion, and extension reports show to be the case. Furthermore, even if the spray cost and yield loss were what the growers believe and perceive them to be, the optimal strategy would still tend to include early harvesting and very little if any spraying. Our results suggest some possible ways to improve growers' actual pest management strategies and therefore grower welfare and sustainability.

Keywords: pest management, agriculture, sustainability, insecticide, pesticide, farmer behavior

JEL codes: Q12, Q57, Q16, Q18

This draft: December 2022

¹ Yeh: U.S. Department of Agriculture (USDA) Economic Research Service; adeline.yeh@usda.gov. Gómez: Cornell University; mig7@cornell.edu. Lin Lawell: Cornell University; clinlawell@cornell.edu. We thank Brian Adam, Arnab Basu, Lucas Davis, Limor Golan, Nathan Hendricks, Harry Kaiser, Francesco Lippi, Ariel Ortiz-Bobea, Jerrod Penn, Roderick Rejesus, Trisha Shrum, Vasiliki Skreta, and Chung Piau Teo for detailed and helpful comments and discussions. We also benefited from comments from seminar participants at Cornell University and at the U.S. Department of Agriculture (USDA) Economic Research Service (ERS); and conference participants at the Agricultural and Applied Economics Association (AAEA) Annual Meeting, at the Northeast Agricultural and Resource Economics Association (NAREA) Annual Meeting, and at the International Consortium on Applied Bioeconomy Research (ICABR) Conference. This research was made possible through financial support from the National Institute of Food and Agriculture, U.S. Department of Agriculture Specialty Crops, Research Initiative under Agreement No. 2015-51181-24252. Yeh received funding from the Taiwan Ministry of Education Government Scholarship for Study Abroad and the Taiwanese Overseas Pioneers Grant from the Taiwan Ministry of Science and Technology. Gómez and Lin Lawell are Faculty Fellows at the Cornell Atkinson Center for Sustainability. The views expressed are those of the authors and should not be attributed to the Economic Research Service or USDA. All errors are our own.

1. Introduction

Pest management is an important concern for agricultural producers: pests can cause costly and irreparable harm to crops, and damage from pests often results in vast economic consequences (U.S. Department of Agriculture 2022). Spotted Wing Drosophila (SWD), or *Drosophila suzukii*, is an invasive pest that infests soft-skinned fruit such as berries and has resulted in large pest management costs for berry growers in the U.S. since it was first detected in 2008 (Walsh et al. 2011; Cini, Ioriatti and Anfora 2012; Asplen et al. 2015). Unlike most of the *Drosophila* species, female SWD have a unique serrated ovipositor which allows them to lay eggs under the fruit skin, causing causes direct damage (Asplen et al. 2015). High-value fruit crops such as blueberries are the most affected by SWD. Wholesale fruit buyers generally have very low tolerance for infested fruit; in fresh or exporting markets, the whole shipment is rejected if any infestation is found. The economic loss from SWD has been severe -- SWD has been estimated to have caused over \$500 million in revenue losses for the West Coast berry industry alone (Bolda, Goodhue and Zalom 2010).

To protect their crops from SWD infestation, domestic growers have increased their use of insecticide applications (Bolda et al. 2010; Walsh et al. 2011; Cini et al. 2012). The management costs of intensive insecticide applications are high, however (Farnsworth et al. 2017; Drummond, Ballman and Collins 2019), and recent research suggests that growers may be over-spraying when the pest pressure is low (Yeh et al. 2020). Moreover, increased insecticide usage due to SWD may also lead to insecticide resistance (Farnsworth et al. 2017; Gress and Zalom 2019), as well as environmental and health concerns (Sexton 2007; Goeb et al. 2020; Chatzimichael, Genius and Tzouvelekas 2021).

The high production costs and environmental concerns from intensive insecticide applications have led the industry and researchers to evaluate alternative pest management strategies (Farnsworth et al. 2017). Possible sustainable SWD control alternatives include exclusion netting (Leach, Van Timmeren and Isaacs 2016; Del Fava, Ioriatti and Melegaro 2017; Stockton et al. 2020), frequent harvesting (Leach et al. 2018), mulching (Rendon et al. 2020), drip and overhead sprinkler irrigation (Rendon and Walton 2019), bagging infested waste berries (Leach et al. 2018), postharvest cold storage (Kraft et al. 2020), or early harvesting (Drummond et al. 2018; Drummond, Ballman and Collins 2019; Yeh et al. 2020), among other strategies. Nevertheless, given the high penalty associated with infestation, the majority of growers still tend

to cling to an intensive calendar-based spraying schedule to control SWD (Yeh et al. 2020), although it may not be the best management practice.

In this paper, we focus on lowbush blueberry production in Maine. Maine is the third largest state in domestic blueberry production volume (U.S. Department of Agriculture 2021), and produces over 90% of the lowbush blueberry produced in the U.S. Given the climate of the region, SWD is a mid-to-late-season pest. Unlike highbush blueberry production, which has multiple harvests per season, lowbush blueberry growers harvest only once per season. As a result, for lowbush blueberry growers in Maine, a unique sustainable pest control alternative to calendar-based insecticide applications that has been proposed by entomologists is to harvest earlier to avoid the higher pest pressure later in the season (Drummond et al. 2018; Drummond, Ballman and Collins 2019). The tradeoff for early harvesting, however, is that growers may incur revenue loss from any prematurely harvested unripe fruit that is sorted out in the processing line (Drummond et al. 2018). The optimal pest management strategy is further complicated by the uncertainty growers face regarding pest pressure and the corresponding infestation, which affects their profits.

In this paper we examine the following research questions: What is the optimal SWD management strategy? Are growers currently following the optimal SWD management strategy? Is sustainable pest control a part of the optimal strategy?

In order to answer these research questions, we develop a novel dynamic bioeconomic analysis framework to analyze sustainable SWD pest management under uncertainty. We apply our multi-stage dynamic bioeconomic analysis framework to a unique dataset of 92 lowbush blueberry farms in Maine to investigate the optimal SWD management strategy for lowbush blueberry production, the resulting welfare changes, and grower behavior.

Our dynamic bioeconomic analysis framework combines numerical dynamic optimization and dynamic structural econometric estimation, and consists of three stages. In the first stage, we construct a numerical bioeconomic model to solve for the dynamically optimal management strategy, and compare optimal decisions predicted by our model with actual decisions made by growers. In the second stage, we develop a dynamic structural econometric model that accounts for unobservable state variables that growers observe (but we do not observe) when they make their spraying and harvesting decisions, and estimate the structural parameters econometrically using the data. Our structural model enables us to ascertain what parameters, beliefs, and perceptions would rationalize the decisions growers have actually made. In the third stage, we

parameterize our numerical bioeconomic model from Stage 1 using our structural parameter estimates from Stage 2 in order to determine the optimal strategy conditional on growers' beliefs, and to assess whether the actual decisions made by growers are optimal given their beliefs as determined in Stage 2.

We model the grower's decision-making problem over a growing season as a finite-horizon stochastic dynamic optimization problem. Each week during the growing season of a fruit-bearing year, the grower makes decisions about whether to spray insecticide to control SWD and whether to harvest.² Spraying and harvesting are both decisions that are irreversible. Once the grower decides to harvest, the grower receives the revenues from harvest (which is the yield net of any loss from SWD, times price) at the time of harvest. The decision-making problem is dynamic because blueberry yields, SWD population, and infestation levels vary over time and are affected by previous spraying decisions made by the grower over the course of the season. The model is stochastic given the uncertainty faced by growers regarding SWD population, infestation, and weather. The grower's problem is therefore one of investment under uncertainty (Dixit and Pindyck 1994).

The numerical bioeconomic model we develop in the first stage is a finite-horizon stochastic dynamic programming problem that solves for the dynamically optimal management strategy. The state variables include the observed SWD larval and adult population, the interval of time since the last insecticide application, precipitation, and temperature. SWD larval population, SWD adult population, and weather all evolve stochastically. We non-parametrically estimate separate stochastic transition densities for the SWD population during the early-, mid-, and late-season using the data. The distributions for stochastic precipitation and temperature, which we allow to vary by time of season, are estimated using empirical averages in the actual data. The interval since the last insecticide application evolves deterministically. We solve the finite-horizon dynamic programming model via backwards iteration in order to determine the

² Lowbush blueberry production follows a two-year production cycle. In the first year (the vegetative year), growers mow or burn the harvested fields to allow new plants to emerge. In the second year (the fruit-bearing year), growers pollinate, irrigate, and harvest the blueberries. We focus on modeling the grower's decision-making problem over a growing season during a fruit-bearing year as a finite-horizon stochastic dynamic optimization problem. As the plants have already been planted by the beginning of the fruit-bearing year, and as the planting date in the previous vegetative year is not likely to influence a grower's decisions during the fruit-bearing year (University of Maine Cooperative Extension 2010), we do not model alternative planting dates for the grower, but instead focus on the spraying and harvesting decisions, as these are the relevant decisions for the fruit-bearing year.

dynamically optimal management strategy. We compare the optimal decisions predicted by our model with the actual decisions made by growers in the data, and assess whether and by how much growers can increase their welfare by employing the optimal strategy.

One key assumption we make when inferring optimality using the numerical bioeconomic model is that growers' perceptions and beliefs about spraying costs and infestation loss are the same as those we use in our model. Although we calibrate the spraying costs and infestation loss parameters in our numerical bioeconomic model based on actual data and SWD studies of infestation, it is possible that growers may have different perceptions and beliefs about spraying costs and infestation loss. If growers believe or perceive parameters to be different from what actual data or expert opinion shows them to be (and therefore different from what our model assumes), then the optimal strategy given their beliefs may differ from what is suggested by our numerical bioeconomic model.

To address this, in the second stage of our study we develop and apply a dynamic structural model to estimate the parameters that underlie the decision-making of Maine lowbush blueberry growers. Our dynamic structural econometric model builds upon our numerical bioeconomic model, and additionally accounts for unobservable state variables that growers observe (but we do not observe) when they make their spraying and harvesting decisions. We estimate several parameters econometrically using the data, including the yield losses due to SWD infestation and the perceived spray costs. Building on the nested fixed point maximum likelihood estimation technique developed by Rust (1987), we solve for the structural parameters using a maximum likelihood estimation that nests a backwards iteration to solve for the continuation values and conditional choice probabilities for each evaluation of the likelihood function. Our dynamic structural econometric model enables us to ascertain what parameters, beliefs, and perceptions would rationalize the decisions growers actually made.

In the third stage of our analysis, we apply the structural parameters estimated at the second stage to the numerical model built in the first stage in order to solve for the optimal strategies conditional on growers' beliefs. In other words, we solve for what the optimal strategy would be if the spray cost and yield loss were what the growers believe and perceive them to be. Then, taking growers' beliefs as given, we recalculate welfare using our third-stage composite model to assess whether and by how much growers can increase their welfare by employing the optimal strategy.

The contribution of this study is twofold. First, our dynamic bioeconomic analysis of the optimality of SWD management strategies as well as their welfare and sustainability have important implications for growers and policymakers. We not only assess the optimality of sustainable pest management alternatives for this severe pest issue, but also provide actionable results for growers to improve their welfare.

Our second contribution is methodological. Programming- or optimization-based bioeconomic models have the drawback of simplifying or assuming the behavior and beliefs of decision-makers (Janssen and van Ittersum 2007). To the best of our knowledge, our paper is one of the few studies in the field of bioeconomics to incorporate a dynamic structural econometric estimation of grower's behavior into numerical optimization. By including structural estimates of growers' beliefs and perceptions, our novel multi-stage dynamic bioeconomic analysis framework enables us to propose possible improvements in management practices that align with growers' beliefs and perceptions.

Our results show that early harvesting, a sustainable pest control option that has been proposed by entomologists (Drummond et al. 2018; Drummond, Ballman and Collins 2019), can be part of an optimal management strategy, and that spraying insecticide is not optimal in most cases when pest pressure is low. In contrast, data on the actual decisions of growers show that growers tend to spray earlier and more often than is optimal, and harvest later than is optimal. We find that the actual behavior of growers is rationalized by perceptions and beliefs about the spray cost and yield loss that differ from what economic data, expert opinion, and extension reports show to be the case. Furthermore, even if the spray cost and yield loss were what the growers believe and perceive them to be, the optimal strategy would still tend to include early harvesting and very little if any spraying. Our results suggest some possible ways to improve growers' actual pest management strategies and therefore grower welfare and sustainability. Our research has the potential to not only provide timely information to stakeholders regarding optimal management strategies, but also shed light on the literature on the bioeconomics of pest management.

The balance of our paper proceeds as follows. We review the previous literature in Section 2. We describe our data in Section 3. We present our numerical bioeconomic model in Section 4 and our dynamic structural econometric model in Section 5. We present and discuss our results in Section 6. Section 7 concludes.

2. Previous Literature

Our paper contributes to various strands of the literature. First, we build on the literature on the economics of SWD. Since SWD has now become an established pest in the U.S., there is a burgeoning literature examining the economic impacts SWD has brought to the industry (Bolda et al. 2010; Goodhue et al. 2011; Farnsworth et al. 2017; Yeh et al. 2020). Although pest management decisions are made at the farm level, however, there have heretofore been few economic studies of SWD at the farm level; to our knowledge the only two farm-level studies are Fan et al. (2020) and Yeh et al. (2020). Fan et al. (2020) develop a Bayesian bioeconomic model to examine whether monitoring-based integrated pest management strategies perform better in terms of minimizing costs. Yeh et al. (2020) use Monte Carlo simulations to compare the expected revenues under different management strategies for a typical wild blueberry farm in Maine. Both Fan et al. (2020) and Yeh et al. (2020) use simulation-based approaches to rank and compare how different pest control strategies perform. We build on the previous literature on the economics of SWD by developing and applying a novel dynamic bioeconomic analysis framework that combines numerical dynamic optimization and dynamic structural econometric estimation to analyze farm-level SWD control decisions.

Our paper also relates more generally to a broader literature on the economics of pest management. Production-oriented studies of pest management focus on estimating the productivity change or functional form from including pesticides as a damage-control input in the crop production function (Carrasco-Tauber and Moffitt 1992; Kuosmanen, Pemsil and Wessler 2006; Sexton 2007; Chambers, Karagiannis and Tzouvelekas 2010), while other studies focus on the welfare implications of pesticide usage, such as the environmental externalities and health impacts (Sunding and Zivin 2000; Sexton 2007; Grogan and Goodhue 2012; Waterfield and Zilberman 2012). Given the negative externalities associated with pesticide applications, researchers have also assessed how to incentivize pesticide reduction in various settings (Lohr, Park and Higley 1999; Falconer and Hodge 2000; Jacquet, Butault and Guichard 2011). For instance, Jacquet, Butault and Guichard (2011) simulate the effects of pesticide reduction and show that reducing pesticide use by 30% could be possible without reducing farmers' income. Bakker et al. (2021) analyze responses to an online survey to identify which social-psychological constructs determine farmers' intentions to decrease pesticide use, and find that farmers need

successful examples of how to decrease pesticide use, either via exchange with peer farmers or knowledge provisioning on alternative pest control methods.

As pest control decisions are highly related to ecological and environmental factors, bioeconomic models provide an integrated framework to evaluate optimal pest control decisions. Given the intertwined feedback links between the biological and economic systems, bioeconomic modeling is generally challenging and the modeling approach depends highly on the nature of the problem (Finnoff et al. 2005; Smith 2008; Kling et al. 2017). For the case of optimal farm-level pest management, the majority of bioeconomic studies are based on mathematical programming or optimization with a certain objective such as profit maximization (Falconer and Hodge 2000; Buysse, Van Huylenbroeck and Lauwers 2007; de Frahan et al. 2007; Mérel and Howitt 2014).

One caveat of this type of bioeconomic model is that it often neglects the behavioral factors of producers that may explain why producers may not desire to manage their production according to what the model deems to be optimal (Janssen and van Ittersum 2007). Previous research suggests the importance of acknowledging not only biological system parameters but also producer decision-making behavior in bioeconomic studies (Falconer and Hodge 2000; Smith 2008; Chen, Jayaprakash and Irwin 2012). Synthesis papers such as Kling et al. (2017) point out that there is a lack of economic models of decision-making coupled with the biophysical system to provide policy-relevant implications. Our integrated framework suggests a novel way of incorporating producers' perceptions into bioeconomic research, which sheds light on this strand of work.

Our paper also contributes to the literature on dynamic structural econometric models and its related applications. Dynamic structural econometric models provide great flexibility for researchers to estimate behavioral parameters under various circumstances, and have been applied to various economic research topics (Hotz and Miller 1993; Rust 1987; Adda and Cooper 2000; Timmins 2002; Iskhakov 2010; Arcidiacono and Miller 2011; Keane, Todd and Wolpin 2011; Duflo, Hanna and Ryan 2012; Gayle, Golan and Soytaş 2018; Blundell, Gowrisankaran and Langer 2020; Cook and Lin Lawell 2020; Reeling, Verdier and Lupi 2020; Agarwal et al. 2021; Li, Liu and Wei 2021; Anderson et al. 2021). Applications in agriculture include disease management (Carroll et al. 2019; Sambucci, Lin Lawell and Lybbert 2022; Carroll et al. 2022b; Carroll et al. 2022a), land use (Scott 2013), agroforestry (Oliva et al. 2020), and agricultural groundwater extraction (Sears, Lin Lawell and Walter 2022). For instance, Carroll et al. (2021b) estimate a dynamic structural model for lettuce crops disease control in California to compare

long-term and short-term decision-makers. We build on the previous literature using dynamic structural models by combining numerical dynamic optimization and dynamic structural econometric estimation to analyze farmers' within-season decisions and to ascertain what parameters, beliefs, and perceptions would rationalize the decisions growers actually made. Although farm managers generally make pest management decisions within the production season on a weekly basis, the use of dynamic structural econometric models to understand growers' within-season pest control decision-making of pest control is rare in the literature, which could be due to data limitations.

Previous research by Misra and Nair (2011) provides evidence that dynamic structural econometric models can help significantly improve decision-making and outcomes. In their study, Misra and Nair (2011) develop and apply a dynamic structural econometric model to data on the US sales force of a large contact lens manufacturer to design sales-force compensation schemes to increase the firm's profits. Their recommendations were then implemented at the firm, resulting in an increase in annual revenues of about \$12 million. Our research strives to similarly improve decision-making and outcomes by growers managing SWD.

3. Data

We obtain field data from 92 lowbush blueberry fields in Maine. Lowbush blueberry production follows a two-year production cycle. The first year is solely for field preparation (vegetative), while the second year is for harvesting (fruit-bearing). In our dataset, each farm reported data from one fruit-bearing year over the period 2012 to 2017. The adult SWD population was monitored with sticky traps in the field, while SWD larvae were observed using fruit sampling (see Drummond et al. (2019) for details on the data collection).

Given that each farm-year reported slightly different observation intervals and time lengths, we linearly interpolate the uneven observations into a total of 18 weeks (June 1 to September 30) per season, which is when the wild blueberries are generally susceptible to SWD infestation. We then subset the data into weekly discrete time steps for calibration and estimation, given that a grower usually makes weekly decisions in practice. We obtain weather data for each farm location from PRISM.³

³ <http://www.prism.oregonstate.edu/>

4. Numerical Bioeconomic Model

Our numerical bioeconomic model is a finite-horizon stochastic dynamic programming problem. Each week t , given the state of the system $state_t$ (which includes state variables measuring SWD population levels, weather, and the interval of time since the last insecticide application), the grower makes decisions on whether to spray insecticide to control SWD, and whether to harvest and sell the berries. If the grower decides to spray at time t , she pays the cost of spraying $spraycost$ at time t . If the grower decides to harvest at time t , the grower receives the crop revenue $\pi_t(\cdot)$ from selling the berries at time t . If the grower has not yet harvested before t , and neither sprays nor harvests at time t , then the grower's per-period payoff is zero at time t (i.e., the grower neither incurs any costs nor receives any revenue at time t). A grower can spray multiple times during the season, but can only harvest once during the season, after which time the grower no longer has any spraying or harvesting decisions to make. Both the spraying and harvesting decisions are irreversible. Thus, the choice variable a_t is a vector consisting of a dummy variable for spraying (s_t) and a dummy variable for harvesting (h_t), and takes one of three possible values: $a_t \equiv (s_t, h_t) \in A \equiv \{(1,0), (0,0), (0,1)\}$. Growers generally avoid applying insecticide at the same time of harvesting to comply with the regulated maximum pesticide residue levels; as a consequence, we do not observe growers harvesting and spraying during the same week in the data and therefore do not include harvesting while spraying as a possible action choice in our model.

The state tuple $state_t \equiv (y_{larva,t}, y_{adult,t}, int_t, precip_t, tmax_t)$ is a vector of state variables measuring SWD population levels (observed larva SWD $y_{larva,t}$ and observed adult SWD $y_{adult,t}$), weather variables (weekly accumulated precipitation $precip_t$ and weekly maximum temperature $tmax_t$), and the interval of time since the last insecticide application (int_t).

The SWD population levels in the state tuple $state_t$ are indicated by two state variables: observed larva SWD $y_{larva,t}$, which affects the expected yield loss $loss(\cdot)$ and is discretized to six levels; and observed adult SWD $y_{adult,t}$, which is a dummy variable indicating whether adult SWD was observed at time t .⁴ The expected yield loss due to larval infestation, $loss(\cdot)$, is a

⁴ Since observed larva SWD $y_{larva,t}$ affects the expected yield loss $loss(\cdot)$ and therefore the crop revenue from harvesting, and since our assumed values for the yield loss $loss(y_{larva,t})$ as a function of observed larval infestation $y_{larva,t}$, which is based on expert opinion and extension reports (Burrack 2014; De Ros et al. 2015; DiGiacomo et al.

function of observed larval infestation $y_{larva,t}$. The crop revenue $\pi_t(\cdot)$ from harvesting at week t is a function of the yield $yield_t$ at week t in the absence of infestation, the yield loss $loss(\cdot)$, the observed larval infestation $y_{larva,t}$, and the crop price $price$:

$$\pi_t(y_{larva,t}) = yield_t \cdot (1 - loss(y_{larva,t})) \cdot price, \quad (1)$$

where the yield $yield_t$ at week t in the absence of infestation is given by the yield at full maturity in the absence of infestation $yield_{mature}$ times the percentage of ripeness $percent_ripe_t$ at harvest time t :

$$yield_t = yield_{mature} * percent_ripe_t. \quad (2)$$

In our base case specification, we assume that the grower is risk neutral, and therefore that the grower's week- t utility $U(\cdot)$ from harvesting at time t is linear in crop revenue $\pi_t(\cdot)$:

$$U(\pi_t(y_{larva,t})) = \pi_t(y_{larva,t}). \quad (3)$$

Since crop revenue $\pi_t(\cdot)$ is uncertain and depends on an infestation loss probability which depends on stochastic SWD population levels, in an alternative specification we allow the grower to be risk averse, and use a constant relative risk aversion (CRRA) functional form for the grower's week- t utility $U(\cdot)$ from harvesting at time t as a function of crop revenue $\pi_t(\cdot)$:

$$U(\pi_t(y_{larva,t})) = \frac{\pi_t(y_{larva,t})^{1-\eta}}{1-\eta}, \quad (4)$$

where η is the coefficient of constant relative risk aversion. Since the cost of spraying each time t the grower sprays, $spraycost$, is deterministic, certain, and incurred in weeks prior to harvest, we assume that the grower's week- t utility from spraying at time t is linear in the certain spray cost $spraycost$ whether or not the grower is risk averse with respect to the risky and uncertain crop revenue.⁵

2019; Drummond et al. 2019; Yeh, Drummond and Gómez 2019; Yeh et al. 2020) and reported in Table 1, has five tiers, we discretize observed larva SWD $y_{larva,t}$ into six levels: one level for no observed larva SWD, plus five levels for each of the five tiers of the yield loss function $loss(y_{larva,t})$. In contrast to observed larva SWD $y_{larva,t}$, which affects crop revenue through its effect on expected yield loss, observed adult SWD $y_{adult,t}$ does not directly affect expected yield loss or crop revenue, but instead indirectly affects expected yield loss through its effect on the population dynamics of observed larva SWD $y_{larva,t}$ via the transition density. Thus, owing to state space constraints, we discretize observed adult SWD $y_{adult,t}$ into a binary dummy variable indicating whether adult SWD was observed at time t ; given that we already discretize observed larva SWD $y_{larva,t}$ into six levels, a finer discretization of observed adult SWD $y_{adult,t}$ is neither feasible or desirable and would lead to insufficient observations for many values of the state tuples.

⁵ Since the grower receives a negative payoff (equal to the negative of the spray cost) on the weeks when the grower sprays, it is not straightforward to make the grower risk averse with respect to the certain cost of spraying, since the

The state tuple $state_t$ also includes a discretized variable int_t for the interval of time since the last insecticide application, which can take the value of 1 if the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective; 2 if the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective; or 3 if the grower has not yet sprayed this season prior to week t . Since the insecticide is effective for a maximum of 7 days, the maximum number of weeks for which the insecticide is effective is 1.

The state tuple $state_t$ additionally includes two weather-related dummy variables, the weekly accumulated precipitation $precip_t$ and weekly maximum temperature $temp_t$, as studies have shown that precipitation may affect spraying efficacy while the temperature affects SWD population dynamics (Tochen et al. 2014; Wiman et al. 2014; Hamby et al. 2016; Gautam et al. 2016; Van Timmeren et al. 2017).⁶ We use 2 mm and 27 degrees Celsius as the thresholds for the dummy variables for accumulated precipitation $precip_t$ and the maximum temperature $temp_t$, respectively.

The grower’s finite-horizon dynamic optimization problem is to make spraying and harvesting decisions each week of the season in order to maximize the expected present discounted value of the entire stream of weekly payoffs. The value function $V_t(state_t)$ at time t , which is a function of the state variable tuple $state_t$ at time t , and which gives the optimized value of the expected present discounted value of the entire stream of per-period payoffs from time t until the remainder of the season when the spraying and harvesting decisions are chosen optimally, is given by:

CRRA utility function is defined only over non-negative numbers. Nevertheless, we tried an alternative specification of risk aversion in which we shifted all the per-period payoffs by some positive constant greater than the spray cost (with this positive constant representing weekly baseline income from sources other than the current blueberry growing season, such as a weekly withdrawal from their savings from the previous season, etc.), and then applied the CRRA utility function to the weekly payoffs for all weeks, including the certain cost of spraying for weeks when the farmer sprayed, and the uncertain crop revenue for the week when the farmer harvested. This alternative specification, in which the farmer was risk averse with respect to both the certain cost of spraying and the uncertain crop revenue, fit the data very poorly, however, likely because risk aversion arguably only matters if there is uncertainty (i.e., a risky payoff), and therefore does not apply to the certain cost of spraying. We therefore focus on the alternative specification for risk aversion in which the farmer is risk averse with respect to the uncertain crop revenue, but risk neutral with respect to the certain cost of spraying.

⁶ We focus on modeling the effects of weather on SWD population dynamics and spraying efficacy. While weather may also affect lowbush blueberry yields as well, the possible variation in yield due to variations in weather during the growing season is relatively small (Parent et al. 2020), especially compared to yield losses from SWD infestation. Thus, as SWD is a larger source of uncertainty and impact on yield than weather is for lowbush blueberry production in Maine, we focus on modeling the uncertainty that arises from SWD and the effects of weather on SWD population dynamics and spraying efficacy.

$$V_t(state_t) = \max_{(s_t, h_t) \in A} \left(\begin{array}{c} s_t \cdot (1-h_t) \cdot (-\text{spraycost} + \beta \cdot E_t[V_{t+1}(state_{t+1}) | state_t, s_t=1, h_t=0]), \\ (1-s_t) \cdot (1-h_t) \cdot \beta \cdot E_t[V_{t+1}(state_{t+1}) | state_t, s_t=0, h_t=0], \\ (1-s_t) \cdot h_t \cdot U(\pi_t(y_{larva,t})) \end{array} \right), \quad (5)$$

where β denotes the weekly discount factor; and where $E_t[V_{t+1}(state_{t+1}) | state_t, s_t, h_t = 0]$ is the continuation value from not harvesting, which is the expected value of the value function next period, conditional on not harvesting this period, and conditional on the state variables and spray decision this period, where the expectations are taken over next period's values of the stochastic SWD population levels and stochastic weather variables. We assume that the continuation value from waiting instead of harvesting in the final week T of the season is 0: $E_{T+1}[V_{T+1}(state_{T+1}) | state_T, s_T, h_T = 0] = 0$. We solve for the value function for each week t via backwards iteration.

We assume that the state variables follow a first-order Markov process. The transition densities $Pr_t(state_{t+1} | state_t, a_t)$ governing the evolution of state variables from one period to the next given the state variables and choice variable this period are specified as follows. For SWD population levels (i.e., $y_{larva,t}$ and $y_{adult,t}$), we non-parametrically estimate the transition densities for the stochastic SWD population levels conditional on the spraying decision, weather variables, the interval of time since last spray, and the time of season.⁷ The distributions for stochastic precipitation and temperature, which we allow to vary by time of season, are estimated using empirical averages in the actual data. The interval of time since the last spray evolves deterministically as a function of this period's action and the value of the previous interval of time since the last spray.

The values of key model parameters in our numerical dynamic bioeconomic model are values for a representative lowbush blueberry farm in Maine based on various sources. We assume that a single lowbush blueberry grower is a price taker in the blueberry market. As there are 485 commercial-scale wild blueberry growers in Maine (Calderwood, Yarborough and Tooley 2020), and as they compete in the industry with Canada, which also produces the fruit (Whittle 2021), it seems reasonable to assume that the lowbush blueberry industry in Maine is perfectly competitive and that a single lowbush blueberry grower in Maine is a price taker. We use the average price of \$0.26 per lb (U.S. Department of Agriculture 2019) for the crop price *price*. Lowbush blueberries are almost entirely sold to the processed market, wherein they can be preserved relatively long

⁷ We describe our methodology for estimating the transition densities for SWD dynamics in more detail in Appendix B.

compared to fresh blueberries, and thus the price does not fluctuate much within a growing season; as a consequence, the timing of harvest within the same season does not affect the price of lowbush blueberries (Yarborough 2012; Yeh et al. 2020). We assume an annual discount factor of 0.9 and calculate the corresponding weekly discount factor β accordingly.⁸ The yield at full maturity $yield_{mature}$ of a healthy berry field in the absence of infestation is 4,000 lbs per acre. The cost of spraying $spraycost$ is \$40 per acre (Esau 2019).

The percentage of ripeness $percent_ripe_t$ at harvest time t is based on a function estimated by Drummond et al. (2019) using 2012-2018 field data, and is given by:

$$percent_ripe_t = \begin{cases} 0 & \text{for } t < first_ripe_week \\ 100/(1 + e^{(30.903 - 0.159 \cdot JulianDate_t)}) & \text{for } t \geq first_ripe_week, \end{cases} \quad (6)$$

where $JulianDate_t$ is the Julian date of harvest time t ; and where the first week any of the berries begin to ripen, $first_ripe_week$, is week 9. According to this function for fruit ripeness, fruit reaches close to its maximum maturity around weeks 11-12 of the 18-week season.

Our assumed values for the yield loss $loss(y_{larva,t})$ as a function of observed larval infestation $y_{larva,t}$ are based on expert opinion and extension reports (Burrack 2014; De Ros et al. 2015; DiGiacomo et al. 2019; Drummond et al. 2019; Yeh, Drummond and Gómez 2019; Yeh et al. 2020), and are reported in Table 1.

As explained below, we conduct sensitivity analyses that vary the values of key model parameters, and find that the results of our numerical model are robust to a reasonable range of values for these parameters.

5. Dynamic Structural Econometric Model

To understand the beliefs and perceptions of growers that underlie and rationalize their spraying and harvesting decisions as revealed in the data, we nest our numerical bioeconomic model within a dynamic structural econometric model adapted from Rust (1987). Our dynamic structural econometric model builds upon our numerical bioeconomic model, and additionally accounts for unobservable state variables that growers observe (but we do not observe) when they make their spraying and harvesting decisions.

⁸ An annual discount factor of 0.9 is commonly assumed in the literature using dynamic models (see e.g., Ryan (2012); Lin (2013) ; Sears, Lim and Lin Lawell (2019); Cook and Lin Lawell (2020)).

The vector of structural parameters θ we estimate relates to costs induced by SWD, including the five tiers of infestation loss in the yield loss function $loss(y_{larva,t})$ and the spraying cost. In alternative specifications in which we allow growers to be risk averse, the parameters θ also include the coefficient of constant relative risk aversion η in the grower's CRRA utility from crop revenue.

To account for unobservable state variables that growers observe (but we do not observe) when they make their spraying and harvesting decisions, we next expand the per-period payoff to each choice a_{it} to include both a deterministic component $u_0(state_{it}, a_{it}, \theta)$ and a stochastic component $\epsilon_{a_{it}}$. In any given week t , the deterministic component $u_0(\cdot)$ of the per-period payoff is equal to the negative of the spray cost if the grower sprays that week; is equal to 0 if the grower neither sprays nor harvests that week; and is equal to the grower's utility $U(\pi_t(y_{larva,it}))$ from harvesting if the grower harvests that week. The stochastic component to the per-period payoff to each action is an unobserved shock $\epsilon_{a_{it}}$ associated with that action choice a_{it} that is assumed to be distributed i.i.d. extreme value across time t , farms i , and actions a_{it} . The value function incorporating these unobserved shocks $\epsilon_{a_{it}}$ is now given by:

$$V_t(state_{it}) = \max_{a_{it}} u_0(state_{it}, a_{it}, \theta) + \epsilon_{a_{it}} + \beta \cdot E_t[V_{t+1}(state_{i,t+1})|state_{it}, a_{it}], \quad (7)$$

which can be expanded out for each possible action choice as follows:

$$V_t(state_{it}) = \max_{(s_{it}, h_{it}) \in A} \begin{pmatrix} s_{it} \cdot (1-h_{it}) \cdot (-\text{spraycost} + \epsilon_{1it} + \beta \cdot E_t[V_{t+1}(state_{i,t+1})|state_{it}, s_{it}=1, h_{it}=0]), \\ (1-s_{it}) \cdot (1-h_{it}) \cdot (\epsilon_{2it} + \beta \cdot E_t[V_{t+1}(state_{i,t+1})|state_{it}, s_{it}=0, h_{it}=0]), \\ (1-s_{it}) \cdot h_{it} \cdot (U(\pi_t(y_{larva,it})) + \epsilon_{3it}) \end{pmatrix}. \quad (8)$$

The conditional choice probabilities $Pr_t(a_{it}|state_{it}, \theta)$ are given by:

$$Pr_t(a_{it}|state_{it}, \theta) = \frac{\exp(u_0(state_{it}, a_{it}, \theta) + \beta V_t^c(state_{it}, a_{it}, \theta))}{\sum_{\tilde{a}_{it}} \exp(u_0(state_{it}, \tilde{a}_{it}, \theta) + \beta V_t^c(state_{it}, \tilde{a}_{it}, \theta))}, \quad (9)$$

where $V_t^c(state_{it}, a_{it}, \theta)$ is the continuation value, which is the expected value of the value function next period given the states and actions this period:

$$V_t^c(state_{it}, a_{it}, \theta) = E_t[V_{t+1}(state_{i,t+1})|state_{it}, a_{it}]. \quad (10)$$

We use maximum likelihood estimation to find the parameters θ that maximize the log-likelihood function $L(\theta)$, which is the following function of the conditional choice probabilities $Pr_t(a_{it}|state_{it}, \theta)$:

$$L(\theta) = \ln \mathcal{L}(\theta) = \sum_i \sum_t \ln(Pr_t(a_{it} | state_{it}, \theta)). \quad (11)$$

Building on the nested fixed point maximum likelihood estimation technique developed by Rust (1987), our maximum likelihood estimation methodology nests a backwards iteration to solve for the continuation values and conditional choice probabilities for each time t at each evaluation of the likelihood function.

Identification of the parameters θ comes from the differences between per-period payoffs across different action choices, which in finite-horizon dynamic discrete choice models are identified when the discount factor β , the distribution of the choice-specific shocks ε_{it} , and the final period continuation value are fixed (Rust 1994; Magnac and Thesmar 2002; Abbring 2010). In particular, because the discount factor β and the distribution of the choice-specific shocks ε_{it} are fixed and the final period continuation value is zero, the parameters in our model are identified because each term in the deterministic component $u_0(state_{it}, a_{it}, \theta)$ of the per-period payoff depends on the action a_{it} being taken at time t , and therefore varies based on the action taken; as a consequence, the parameters do not cancel out in the differences between per-period payoffs across different action choices and are therefore identified. For example, the *spraycost* parameter is identified in the difference between the per-period payoff from choosing to spray and the per-period payoff from any action choice a_{it} that does not involve spraying.

Standard errors are formed by a non-parametric bootstrap. Farms are randomly drawn from the data set with replacement to generate 100 independent panels each with the same number of farms as in the original data set. The structural model is run on each of the new panels. The standard errors are then formed by taking the standard deviation of the parameter estimates from each of the panels.

6. Results and Discussion

6.1. SWD population dynamics

As described in detail in Appendix B, the results of the transition densities for the SWD population levels we estimate non-parametrically and use for our numerical model and structural econometric models are consistent with models of SWD dynamics. Our estimated transition densities show that, during the mid-season, if the farmer either has never sprayed yet or is past the effective periods of past insecticide application, there is a higher probability that the SWD larva

population will increase if the farmer does not spray. Also during the mid-season, spraying diminishes the probability of SWD getting worse, especially when precipitation is low, which is consistent with studies that show that higher rainfall reduces spraying efficacy because rain washes the insecticide away from the crop surface (Gautam et al. 2016; Van Timmeren et al. 2017). We also find that when the temperature is very high, this makes it more likely for SWD populations to decline, which is consistent with studies that show that the SWD mortality rate increases dramatically if the temperature exceeds 27 degrees Celsius (Tochen et al. 2014).

6.2. *Optimal management strategies and welfare*

Our numerical bioeconomic model solves for the policy function for each week, which specifies the optimal strategy for each week as a function of the state variables. As seen in Table B1 in Appendix B, under the base case specification, the policy function shows that for most weeks when adult SWD have not yet been observed, it is not optimal to spray insecticide. Spraying is only optimal in some cases when both adult SWD are observed and the discretized variable for the time interval since last spray is high, indicating that the grower either has not sprayed at all this season or that enough time has elapsed since the last spray that the last spray is no longer effective.

We also find that harvesting can be optimal as early as week 10 for a few states. These states tend to be states in which adult SWD have been observed and when the discretized variable for the time interval since last spray is high (indicating that the grower either has not sprayed at all this season or is past the effective periods of past insecticide applications). Thus, early harvesting can be part of an optimal SWD management strategy when SWD population levels are high. We also find that harvesting is optimal for almost all state-tuples starting week 12 onwards. In the base case specification, fruit reaches its maximum maturity around weeks 11-12.

Our results therefore suggest that a sustainable pest control alternative -- early harvesting -- can be part of an optimal management strategy, and that spraying insecticide is not optimal in most cases when pest pressure is low. We also find that the presence of adult SWD -- which do not themselves directly cause yield loss but may affect the population dynamics of larva SWD, which do -- may be an important and relatively intuitive and easy-to-implement indicator of whether the grower should engage in early harvesting.

We use our numerical model to compare the welfare from the optimal strategy predicted by our model with the welfare from the actual decisions made by growers in the data. We define

welfare as the present discounted value (PDV) of the entire stream of per-period payoffs over the entire season for the farm. Our numerical bioeconomic model solves for the value function at the beginning of the season, which specifies the maximum expected PDV of the entire stream of per-period payoffs the farm could have received that year if it followed the optimal strategy. Using our numerical bioeconomic model, we can thus infer actual welfare, optimal welfare, and deadweight loss. We calculate actual welfare using observed actions and states of each farm. The optimal welfare, as given by the value function at the beginning of the season evaluated at the initial observed states, is the maximum expected PDV of the entire stream of per-period payoffs the farm could have received that year if it followed the optimal strategy. We calculate the deadweight loss from the growers' actual actions as optimal welfare minus actual welfare. Under our base case specification, the average actual welfare is estimated at \$826.7 per acre (Table 2). The optimal welfare is on average \$963.6 per acre across all farms. The mean deadweight loss across all farms is \$136.9 per acre.⁹

We also use our numerical model to compare the optimal decisions predicted by our model with the actual decisions made by growers in the data. For each farm in our data set, we run 100 simulations of the optimal spray and harvest strategies in which we start with the actual observed states for that farm at the first week, and then forward simulate using our policy function and non-parametrically estimated transition densities. For each farm for each week of each simulation, we infer the optimal action using the policy function solved for from the numerical optimization, and we draw from the estimated transition densities to simulate the transition of the states from one week to the next. Thus, for each farm, the optimal decisions predicted by our model for the farm uses only the actual observed states for that farm in the first week of the season, our policy function, and our non-parametrically estimated transition densities. For each farm, the optimal decisions predicted by our model could therefore have been plausibly implemented by the farmer since these optimal strategies do not assume knowledge of future values of state variables, but instead account for the uncertainty that farmers face over the course of the season.

⁹ Since the optimal strategy maximizes the expected welfare, some farms have an actual welfare that exceeds the optimal (expected) welfare owing to actual realized draws of stochastic state variables. Nonetheless, the deadweight loss to actual decisions would be reduced and welfare increased on average and in expectation if the farmers pursued the optimal spray and harvest strategy instead.

As an example, Figure B1 in Appendix B compares the optimal probabilities of spraying and harvesting for one of the farms in our data set (Farm 90) with the actual spraying and harvesting decision of that farm. Farm 90 sprayed twice, and the actual SWD larva and SWD adult levels ended up fairly high. In contrast, the optimal strategy for Farm 90 (based on only the actual observed states for Farm 90 in the first week of the season, our policy function, and our non-parametrically estimated transition densities) would have had a lower probability of spray, and instead an early harvest, resulting in lower expected SWD larva levels, lower expected SWD adult levels, and higher expected farm welfare. The deadweight loss for Farm 90 from its actual spraying and harvesting decisions is \$352.65 per acre. In other words, Farm 90 could have increased its expected PDV by \$352.65 per acre by following the optimal strategy in lieu of the spraying and harvesting decisions it actually made.

Using the simulated trajectories from each farm, we plot the average optimal probabilities of spraying and harvesting at each time period, and compare with the actual spraying and harvesting decisions made in the data (Figure 1). The average probabilities of actual spray and actual harvest refer to the averages across the observed actions employed by the growers, while the average probabilities of optimal spray and optimal harvest are calculated by averaging over the 100 simulated trajectories for each farm. Both probabilities are averaged across the 92 farms. The results indicate that the probability of spraying in the observed data is much higher than is optimal. The actual (first) spray tends to be earlier than is optimal. The actual harvest, which take places around week 12 to week 15, is later than the optimal harvest timing, which is around week 10 to week 12.

We also run alternative specifications that vary the crop price, spray cost, the start dates used to define each week, or the start days of week used to define each week. As seen in Figure B2 in Appendix B, our results are robust across alternative specifications that vary crop price or spray cost; and our qualitative results are also fairly similar across alternative specifications that vary the start dates used to define each week or the start days of week used to define each week. Across these various alternative specifications, the optimal SWD strategy still tends to include early harvesting and very little if any spraying.

We also try an alternative specification of our numerical model in which we model organic farms only. Of the 92 farms in our data set, 19 are organic farms. For our organic-only numerical model, we non-parametrically estimate stochastic transition densities for the SWD population

using data from organic farms only, reflecting the SWD population dynamics on organic farms; and we use the organic crop price of \$6.00/lb. Since organic farmers generally do not spray, the transition densities for organic farms do not condition on whether a farmer sprays or on the interval of time since the last spray, since there is very little variation in these variables in the data for organic farms. Moreover, since organic farmers generally do not spray, precipitation, which decreases the effectiveness of spray, does not matter for the stochastic transition densities for the SWD population for organic farms so our transition densities for the SWD population on organic farms do not condition on precipitation. As seen in Figure B3 in Appendix B, the optimal strategy for organic farms places a higher probability on early harvesting than was pursued by the organic growers' in the data. The average actual welfare for organic farms is estimated at \$20,678.9 per acre (Table 2). The optimal welfare for organic farms is on average \$22,078.7 per acre across organic farms. The mean deadweight loss across organic farms is \$1,399.8 per acre.

Since crop revenue $\pi_t(\cdot)$ is uncertain and depends on an infestation loss probability which depends on stochastic SWD population values, in an alternative specification we allow the grower to be risk averse, and use a constant relative risk aversion (CRRA) functional form for the grower's week- t utility $U\left(\pi_t(y_{larva,t})\right)$ from the risky and uncertain crop revenue $\pi_t(\cdot)$. Figure 2 compares the optimal and actual average probabilities of spraying and harvesting for various values of the coefficient of constant relative risk aversion η . As η increases and the grower is more risk averse, the lower the optimal spraying probability and the earlier the optimal harvest timing. Thus, even if growers are risk averse, the optimal SWD strategy still involves early harvesting and little if any spraying.

6.3. *Growers' beliefs and perceptions*

Our structural model enables us to ascertain what parameters, beliefs, and perceptions would rationalize the decisions growers have actually made. The results for the structural parameters from the base-case specification of our structural econometric model are summarized in Table 3. Results show that growers perceive the cost associated with spraying to be very high, at \$2,966.17 per acre, which is significant at a 5% level. In contrast, the actual spray cost (i.e., that we use in our numerical model) is \$40 per acre (Esau 2019).

Figure 3 compares the actual yield loss $loss(y_{larva,t})$ as a function of observed larval infestation y_{larva} with growers' beliefs and perceptions about yield loss based on the base-case structural parameter estimates. The yield loss $loss(y_{larva,t})$ is the cumulative percentage of yield loss at each tier of observed larval infestation. The actual yield loss $loss(y_{larva,t})$ values are the assumed values based on data and expert opinion based on expert opinion and extension reports (Burrack 2014; De Ros et al. 2015; DiGiacomo et al. 2019; Drummond et al. 2019; Yeh et al. 2019; Yeh et al. 2020) reported in Table 1. The dashed blue lines indicate the 95% confidence interval for the growers' beliefs and perceptions, as calculated using the standard errors from our base-case specification of the structural model.

The results in Table 3 and Figure 3 show that growers perceive that there will be a 100% loss in yield if there is a medium-high level of SWD larva (i.e., between 5 and 10% observed larval infestation) or higher. In other words, growers do not perceive much yield loss when the larval infestation is under 5%, but a 100% loss when the infestation is over the 5% threshold. In contrast, as seen in Table 1 for the yield loss that we use in our numerical model, expert opinion and extension reports show that yield loss can still occur (with a cumulative yield loss of up to 30%) when the larval infestation is under 5%, and that the cumulative yield loss for a medium-high level of SWD larva (i.e., between 5 and 10% observed larval infestation) is only 50% (Burrack 2014; De Ros et al. 2015; DiGiacomo et al. 2019; Drummond et al. 2019; Yeh et al. 2019; Yeh et al. 2020). Thus, as seen in Figure 3, growers believe and perceive that yield loss is more concentrated at higher observed larval infestation levels than expert opinion and extension reports show to be the case.

Growers are therefore making decisions as if they face a spray cost (monetary or otherwise) of \$2,966.17 and a 100% loss in yield if the observed larval infestation is 5% or higher. In other words, growers' actual decisions are rationalized by a very high spray cost (monetary or otherwise) and high yield losses from medium-high levels of SWD larva.

We also try several alternative specifications for robustness. First, we try an alternative specification of our structural model in which we use coarser bins for larva and larval infestation. As seen in Table B2a in Appendix B, our results are robust. Once again, we find that growers perceive that there will be a 100% loss in yield if the observed larval infestation is 5% or higher. Our results also show once again that growers perceive the cost associated with spraying to be very high, this time at a statistically significant \$2,965.83 per acre.

Second, we try an alternative specification of our structural model in which we restrict our sample to non-organic (conventional) farms only. As seen in Table B2b in Appendix B, our results are robust: non-organic growers perceive that there will be a 99.9% loss in yield if the observed larval infestation is 5% or higher; and they perceive the cost associated with spraying to be very high, this time at a statistically significant \$2,775.31 per acre.

Third, we try alternative specifications of our structural model in which we restrict our sample to organic farms only, and in which we use for the crop price either the crop price of \$0.26/lb (U.S. Department of Agriculture 2019) as before, or the organic crop price of \$6.00/lb (Yarborough and D'Appollonio 2017).¹⁰ As seen in Table B2c in Appendix B, our results are robust: organic growers perceive that there will be a 100% loss in yield if the observed larval infestation is 5% or higher. Organic farmers also perceive the cost associated with spraying to be very high, at \$5,036.10 per acre, although the point estimate is not statistically significant at a 5% level, likely because of the lack of identifying variation in the spraying decision among organic farms (only one organic farm in our data set ever sprayed).

Fourth, we try alternative specifications of our structural model in which we restrict our sample to organic farms only, in which we non-parametrically estimate stochastic transition densities for the SWD population using data from organic farms only, and in which we use for the crop price either the crop price of \$0.26/lb or the organic crop price of \$6.00/lb. Since organic farmers generally do not spray, the transition densities for organic farms do not condition on whether a farmer sprays or on the interval of time since the last spray, since there is very little variation in these variables in the data for organic farms. Moreover, since organic farmers generally do not spray, precipitation, which decreases the effectiveness of spray, does not matter for the stochastic transition densities for the SWD population for organic farms so our transition densities for the SWD population on organic farms do not condition on precipitation. Thus, in this alternative specification for organic farms, we remove spraying s_t from the choice set, and we remove the interval of time since the last spray and precipitation from the vector of state variables. After removing spraying s_t from the choice set, the spray cost is no longer a structural parameter

¹⁰ The organic crop price of \$6.00/lb is based on an internal survey to growers in Maine, on consultations with extension experts, and on the organic value-added Maine wild blueberry price in Yarborough and D'Appollonio (2017). The organic price is much higher than the conventional price because organic blueberries in Maine are mostly sold on the fresh market, whereas conventional blueberries are mostly sold on the processed market (Yeh et al. 2020).

to be estimated. As seen in Table B2d in Appendix B, our results are robust: organic growers perceive that there will be a 100% loss in yield if the observed larval infestation is 5% or higher.

Since crop revenue is uncertain and depends on an infestation loss probability which depends on stochastic SWD population values, in a fifth alternative specification we allow the grower to be risk averse, and include the coefficient of constant relative risk aversion η among the structural parameters to be estimated. As seen in Table 4, the coefficient of constant relative risk aversion η is estimated to be 0, the spray cost is estimated to be \$2,905, and growers perceive that there will be a 100% loss in yield if there is a medium-high level of SWD larva (i.e., between 5 and 10% observed larval infestation) or higher. Thus, growers are risk neutral and, as before, growers are therefore making decisions as if they face a very high spray cost (monetary or otherwise) and high yield losses from medium-high levels of SWD larva.¹¹

To help tease out the channels of beliefs and misperception, we estimate an alternative specification of our structural model in which we hold the infestation loss fixed at its assumed values based on expert opinion and extension reports in Table 1, and then estimate the spray cost. As seen in Table 5a, even if growers are assumed to have the correct beliefs about the infestation loss, they still make decisions as if they face a very high spray cost (monetary or otherwise) of nearly \$3,000 (in this case, \$2,917.77).

To better understand why such a high perceived spray cost better rationalizes the choices made by the growers, we plot the optimal probabilities of spraying and harvesting for large values of spray costs, holding the infestation loss fixed at its assumed values based on expert opinion and extension reports in Table 1. As seen in Figure B5 of Appendix B, as the spray cost increases, it is optimal to spray less and somewhat earlier, and to harvest somewhat later. This earlier spray and later harvest better fit the actual data, although the actual decisions of growers still differ from the optimal strategy even with very high spray costs. Thus, when growers spray earlier and harvest later than is optimal under the base case parameters, they are acting as if they perceive the spray costs to be much higher than our base case assumed value for spray costs.

¹¹ To allow for heterogeneity by year, we estimate the structural parameters by year in Table B2e in Appendix B. In this set of alternative specifications, we run the structural model for each year using data from that year only and using the average crop price for that particular year. The average annual crop price ranges from \$0.25-0.75/lb; earlier years have higher prices. As the sample size is much smaller when we estimate the structural model separately for each year, we put less weight on these results and our preferred specification is the base-case specification in Table 3 which instead pools the data across all years. Nevertheless, as seen in Figure B4 in Appendix B, even when using the structural parameter estimates by year, the optimal spraying strategy still tends to include early harvesting and less spraying earlier in the season.

To further help tease out the channels of beliefs and misperception, we also estimate an alternative specification of our structural model in which we hold the spray cost fixed at its assumed value of \$40 per acre (Esau 2019), and then estimate the infestation loss. As seen in Table 5b, according to a likelihood ratio test of this alternative specification versus the base-case structural model in Table 3, the data rejects this constrained alternative specification constraining growers to believe and perceive spray costs to be \$40 per acre, since the unconstrained base-case model which instead includes the spray cost as a parameter to be estimated produces a statistically significant improvement in the ability of the model to fit the data at a 0.1% level. Nevertheless, even when growers are assumed to have the correct beliefs about spraying costs, they still misperceive the infestation loss. Once again, as seen in Figure B6 of Appendix B, growers believe and perceive that yield loss is more concentrated at high observed larval infestation levels than expert opinion and extension reports show to be the case. In this case, they perceive that there will be no loss in yield if the observed larval infestation is under 5%, a low and statistically insignificant 12.1% in cumulative yield loss if SWD larva levels are not high (i.e., less than or equal to 10% observed larval infestation), and a 70.9% incremental loss in yield if the SWD larva is at its highest tier (i.e., greater than 10% observed larval infestation). In contrast, as seen in Table 1 for the yield loss that we use in our numerical model, expert opinion and extension reports show that the cumulative yield loss when the SWD larva level is not high (i.e., less than or equal to 10% observed larval infestation) can be up to 50%, and that the infestation loss (or incremental loss in yield) when the SWD larva level is at its highest tier (i.e., greater than 10% observed larval infestation) is only 30% (Burrack 2014; De Ros et al. 2015; DiGiacomo et al. 2019; Drummond et al. 2019; Yeh et al. 2019; Yeh et al. 2020).

To better understand the infestation loss beliefs and perceptions that rationalize the choices made by growers, we plot the optimal and actual average probabilities of spraying and harvesting for various specifications of infestation loss in which the yield loss is low at low observed larval infestation levels and high at high observed larval infestation levels, holding the spray cost fixed at its assumed value of \$40 per acre (Esau 2019). As seen in Figure B7 of Appendix B, as the yield loss becomes more concentrated at high observed larval infestation levels, it is optimal to spray less and somewhat earlier, and to harvest somewhat later. This earlier spray and later harvest better fit the actual data, although the actual decisions of growers still differs from the optimal strategy even when yield loss is concentrated at high observed larval infestation levels. Thus,

when growers spray earlier and harvest later than is optimal under the base case parameters, they are acting as if they perceive the yield loss to be more concentrated at high observed larval infestation levels than what expert opinion and extension reports show to be the case.

6.4. *Optimal strategies conditional on growers' beliefs and perceptions*

Using our structural estimates, we re-run our numerical model to solve for the optimal strategy for the growers given their beliefs and perceptions. Figure 4 plots the optimal strategies given the beliefs and perceptions of growers, and compares them with the actual spraying and harvesting decisions made by growers in the data. As the structural parameter estimates include high perceived spray costs, when we re-run our numerical model using the structural parameter estimates, the policy function indicates that spraying is never an optimal strategy under any state. We also find that harvesting is optimal for most state tuples starting from week 12 onwards, which tends to be earlier in the season than when growers actually harvest in the data. Thus, given growers' beliefs and perceptions, early harvesting can be part of the optimal SWD strategy while spraying is not.

Given that growers perceive spraying costs to be high, the average actual welfare is negative due to the very high perceived spraying cost, at -\$1,643.20. In contrast, the average optimal welfare given the growers' beliefs, which is the average welfare the growers could have received had they followed the optimal strategy conditional on their beliefs, is \$985.30. Thus, even conditional on growers' beliefs, growers experience an average deadweight loss of \$2,638.50 across farms. The deadweight loss indicates that, even under the grower's perceptions of SWD-related costs, their actual choices are still sub-optimal and could be improved (Table 6).

Figure B8 in Appendix B compares the actual decisions with the optimal strategy given the beliefs and perceptions of growers when also accounting for unobservable state variables that growers observe (but we do not observe) when they make their spraying and harvesting decisions, as given by the optimal choice probabilities. The qualitative results are similar: given growers' beliefs, the optimal mixed strategy places a very low and almost negligible probability on spraying, and a higher probability on early harvesting than was pursued by the growers' in the data.

As seen in Figure B9 of Appendix B, we obtain a similar result when using the growers' beliefs and perceptions about infestation loss from the alternative specification of our structural model in which growers are assumed to have the correct beliefs about spraying costs in Table 5b.

Given growers' beliefs about infestation loss, as estimated under the assumption that they have the correct beliefs about spraying costs, the optimal strategy still places a very low and almost negligible probability on spraying, and a higher probability on early harvesting than was pursued by the growers' in the data.

7. Conclusion

In this paper, we develop a novel dynamic bioeconomic analysis framework that combines numerical dynamic optimization and dynamic structural econometric estimation to analyze optimal SWD management for wild blueberry production in Maine.

Our numerical bioeconomic model provides valuable insights on the optimal pest management strategy. Results show that early harvesting, a sustainable pest control option that has been proposed by entomologists (Drummond et al. 2018; Drummond, Ballman and Collins 2019), can be part of an optimal management strategy. For weeks with low SWD population levels, insecticide application is not optimal in most cases. For weeks for high SWD population levels, for example if an SWD adult has been observed, it can be optimal for the grower to harvest early. The more risk averse the grower, the lower the optimal spraying probability and the earlier the optimal harvest timing.

We find that the actual spraying and harvesting decisions of growers are not optimal: growers tend to spray more often than is optimal, the actual (first) spray tends to be earlier than is optimal, and the actual harvest tends to be later than is optimal. Under the assumed values for the parameters -- which are based on economic data, expert opinion, and extension reports -- there is a deadweight loss to actual decisions that would be reduced on average and in expectation if the farmers pursued the optimal spray and harvest strategy instead.

We use a dynamic structural econometric model to understand growers' behavior, perceptions, and beliefs. We find that the actual behavior of growers is rationalized by perceptions and beliefs about the spray cost and yield loss that differ from what economic data, expert opinion, and extension reports show to be the case. In particular, the actual behavior of growers is rationalized by a very high perceived spray cost, much higher than the actual spray cost, possibly owing to other perceived costs of spraying in addition to the actual monetary cost of purchasing the insecticide; as well as a different (and possibly incorrect) perception that yield loss is more concentrated at higher observed larval infestation levels than what expert opinion and extension

reports to may be the case. Growers do not perceive much yield loss when the larval infestation is under 5%, but a 100% loss when the infestation is over the 5% threshold. According to expert opinion and extension reports (Burrack 2014; De Ros et al. 2015; DiGiacomo et al. 2019; Drummond et al. 2019; Yeh et al. 2019; Yeh et al. 2020), however, cumulative yield loss can be up to 30% when the larval infestation is under 5%, but is only 50% when the larval infestation is between 5% and 10%.

A possible reason why growers might not perceive much yield loss when the larval infestation is low is that they can still sell infested berries to the scrap market (for juices, etc.), and therefore might not worry so much when the larval infestation is low since they can still receive some revenue (albeit at a reduced price) for heavily infested berries. A possible reason why growers might perceive that yield loss is more concentrated at higher observed larval infestation levels is that they might be concerned about long-term reputational concerns and the loss of marketing channels if they have a heavily infested crop. Growers often work with the same buyer over time, so an infestation in one year could cause long-lasting damage to the grower's reputation in the following years, leading to the loss of the grower's marketing channels and access to buyers.

Nonetheless, even when we take the beliefs of the growers as given and solve for the optimal SWD strategy conditional on their beliefs, we find that there still exists a deadweight loss to the actual decisions being made by growers. In other words, there is still room for welfare improvement even if the spray cost and yield loss were what the growers believe and perceive them to be. In particular, given growers' beliefs and perceptions, and in contrast to their actual spraying and harvesting decisions, the optimal SWD strategy still tends to include early harvesting and very little if any spraying.

Our results suggest some possible ways to improve the optimality of growers' actual SWD management strategies and therefore grower welfare and sustainability. These include: providing growers with information about actual spray costs and percentage yield loss based on infestation; apprising growers of early harvest as an optimal management strategy; incentivizing growers to consider early harvest as a possible strategy;¹² and mitigating any possible barriers to the use of early harvest as a SWD management strategy, including labor shortages (if the farm needs to hire

¹² Potential actors, organizations, or entities who might be willing or interested in incentivizing growers to consider early harvest as a possible strategy, and paying for this incentive, include extension educators, integrated pest management (IPM) educators, and possible sustainable-driven buyers who are looking for more suppliers that do not spray extensively.

additional labor for harvest, they might need to schedule ahead of time), constraints to harvest machinery (if the farm is small scale and rents the machines), and not knowing that early harvesting is a viable strategy. For example, when growers spray earlier and harvest later than is optimal under the base case parameters, they are acting as if they perceive the yield loss to be more concentrated at higher observed larval infestation levels than what expert opinion and extension reports show to be the case. This suggests that giving growers better information about infestation loss may lead them to spray later and harvest earlier, closer to the optimal strategy.

In their study of which social-psychological constructs determine farmers' intentions to decrease pesticide use, Bakker et al. (2021) find that farmers need successful examples of how to decrease pesticide use, either via exchange with peer farmers or better information about alternative pest control methods. This suggests that in order to improve the optimality of actual SWD management strategies of farmers and therefore their welfare and sustainability, it is important to provide growers with information and knowledge about early harvest as an optimal management strategy, and to encourage farmers who successfully adopt the early harvest strategy to share their experiences with other neighboring and peer farmers.

Our study contributes to the literature on pest management and has valuable policy implications for sustainable pest control. Methodologically, our unique within-season farm-level bioeconomic study combines structural estimations of growers' behavior with numerical optimization. Our results provide actionable suggestions to growers on an optimal pest control strategy that can improve their welfare and sustainability. The results can also be used to provide information on incentivizing sustainable strategy and mitigating possible barriers of adoption.

References

- Abbring, J.H. 2010. "Identification of Dynamic Discrete Choice Models." *Annual Review of Economics* 2(1):367–394.
- Adda, J., and R. Cooper. 2000. "Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies." *Journal of Political Economy* 108(4):778–806.
- Agarwal, N., I. Ashlagi, M. Rees, P. Somaini, and D. Waldinger. 2021. "Equilibrium Allocations Under Alternative Waitlist Designs: Evidence from Deceased Donor Kidneys." *Econometrica* 89(1):37–76.
- Anderson, A., J. Rosen, J. Rust, and K.-P. Wong. 2021. "Disequilibrium Play in Tennis." Working paper, Georgetown University.
- Arcidiacono, P., and R.A. Miller. 2011. "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity." *Econometrica* 79(6):1823–1867.
- Asplen, M.K., G. Anfora, A. Biondi, D.-S. Choi, D. Chu, K.M. Daane, P. Gibert, A.P. Gutierrez, K.A. Hoelmer, W.D. Hutchison, R. Isaacs, Z.-L. Jiang, Z. Kárpáti, M.T. Kimura, M. Pascual, C.R. Philips, C. Plantamp, L. Ponti, G. Véték, H. Vogt, V.M. Walton, Y. Yu, L. Zappalà, and N. Desneux. 2015. "Invasion Biology of Spotted Wing *Drosophila* (*Drosophila suzukii*): A Global Perspective and Future Priorities." *Journal of Pest Science* 88(3):469–494.
- Bakker, L., J. Sok, W. van der Werf, and F.J.J.A. Bianchi. 2021. "Kicking the Habit: What Makes and Breaks Farmers' Intentions to Reduce Pesticide Use?" *Ecological Economics* 180:106868.
- Blundell, W., G. Gowrisankaran, and A. Langer. 2020. "Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations." *American Economic Review* 110(8):2558–2585.
- Bolda, M.P., R.E. Goodhue, and F.G. Zalom. 2010. "Spotted Wing *Drosophila*: Potential Economic Impact of a Newly Established Pest." *Agricultural and Resource Economics Update* 13:5–8.
- Burrack, H. 2014. "Spotted Wing *Drosophila* Impacts, 2014." NC State Extension, NC State University. Available at: <https://entomology.ces.ncsu.edu/spotted-wing-drosophila-impacts-2014/> [Accessed October 20, 2019].
- Buysse, J., G. Van Huylenbroeck, and L. Lauwers. 2007. "Normative, Positive and Econometric Mathematical Programming as Tools for Incorporation of Multifunctionality in Agricultural Policy Modelling." *Agriculture, Ecosystems & Environment* 120(1):70–81.
- Calderwood, L., D.E. Yarborough, and B. Tooley. 2020. "About Maine Wild Blueberry." Fact Sheet No. 220, University of Maine Cooperative Extension. Available at: <https://extension.umaine.edu/blueberries/about-2/> [Accessed May 23, 2022].

- Carrasco-Tauber, C., and L.J. Moffitt. 1992. "Damage Control Econometrics: Functional Specification and Pesticide Productivity." *American Journal of Agricultural Economics* 74(1):158.
- Carroll, C.L., C.A. Carter, R.E. Goodhue, and C.-Y.C. Lin Lawell. 2019. "Crop Disease and Agricultural Productivity: Evidence from a Dynamic Structural Model of Verticillium Wilt Management." In W. Schlenker, ed. *Agricultural Productivity and Producer Behavior*. Chicago: University of Chicago Press, pp. 217–249.
- Carroll, C.L., C.A. Carter, R.E. Goodhue, and C.-Y.C. Lin Lawell. 2022a. "Supply Chain Externalities and Agricultural Disease." Working paper, Cornell University.
- Carroll, C.L., C.A. Carter, R.E. Goodhue, and C.-Y.C. Lin Lawell. 2022b. "The Economics of Decision-Making for Crop Disease Control." Working paper, Cornell University.
- Chambers, R.G., G. Karagiannis, and V. Tzouvelekas. 2010. "Another Look at Pesticide Productivity and Pest Damage." *American Journal of Agricultural Economics* 92(5):1401–1419.
- Chatzimichael, K., M. Genius, and V. Tzouvelekas. 2021. "Pesticide use, health impairments and economic losses under rational farmers behavior." *American Journal of Agricultural Economics*, forthcoming.
- Chen, Y., C. Jayaprakash, and E. Irwin. 2012. "Threshold Management in a Coupled Economic–Ecological System." *Journal of Environmental Economics and Management* 64(3):442–455.
- Cini, A., C. Ioriatti, and G. Anfora. 2012. "A Review of the Invasion of *Drosophila suzukii* in Europe and a Draft Research Agenda for Integrated Pest Management." *Bulletin of Insectology* 65(1):149–160.
- Cook, J.A., and C.-Y.C. Lin Lawell. 2020. "Wind Turbine Shutdowns and Upgrades in Denmark: Timing Decisions and the Impact of Government Policy." *Energy Journal* 41(3):81–118.
- de Frahan, B.H., J. Buysse, P. Polomé, B. Fernagut, O. Harmignie, L. Lauwers, G. Van Huylenbroeck, and J. Van Meensel. 2007. "Positive Mathematical Programming for Agricultural and Environmental Policy Analysis: Review and Practice." In A. Weintraub, C. Romero, T. Bjørndal, R. Epstein, and J. Miranda, eds. *Handbook Of Operations Research In Natural Resources*. Boston, MA: Springer US, pp. 129–154.
- Del Fava, E., C. Ioriatti, and A. Melegaro. 2017. Cost-Benefit Analysis of Controlling the Spotted Wing *Drosophila* (*Drosophila suzukii* (Matsumura)) Spread and Infestation of Soft Fruits in Trentino, Northern Italy. *Pest Management Science* 73(11):2318–2327.
- De Ros, G., S. Conci, T. Pantezzi, and G. Savini. 2015. "The Economic Impact of Invasive Pest *Drosophila suzukii* on Berry Production in the Province of Trento, Italy." *Journal of Berry Research* 5(2):89–96.
- DiGiacomo, G., J. Hadrich, W.D. Hutchison, H. Peterson, and M. Rogers. 2019. "Economic Impact of Spotted Wing *Drosophila* (Diptera: Drosophilidae) Yield Loss on Minnesota Raspberry Farms: A Grower Survey." *Journal of Integrated Pest Management* 10(1):11.
- Dixit, R.K., and R.S. Pindyck. 1994. *Investment Under Uncertainty*. Princeton University Press.

- Drummond, F., E. Ballman, and J. Collins. 2019. "Population Dynamics of Spotted Wing *Drosophila* (*Drosophila suzukii* (Matsumura)) in Maine Wild Blueberry (*Vaccinium angustifolium* Aiton)." *Insects* 10(7):205.
- Drummond, F., E.S. Ballman, J.A. Collins, and D.E. Yarborough. 2018. "Spotted Wing *Drosophila*: Pest Biology and IPM Recommendations for Wild Blueberries." Fact Sheet No. 210, University of Maine Cooperative Extension. Available at: <https://extension.umaine.edu/blueberries/factsheets/insects/210-spotted-wing-drosophila/> [Accessed August 14, 2018].
- Duflo, E., R. Hanna, and S.P. Ryan. 2012. "Incentives Work: Getting Teachers to Come to School." *American Economic Review* 102(4):1241–1278.
- Esau, T. 2019. "Wild Blueberry Management Tool, V1.1 Maine (Excel)." Available at: <https://extension.umaine.edu/blueberries/factsheets/marketing-and-business-management/> [Accessed April 30, 2019].
- Falconer, K., and I. Hodge. 2000. "Using Economic Incentives for Pesticide Usage Reductions: Responsiveness to Input Taxation and Agricultural Systems." *Agricultural Systems* 63(3):175–194.
- Fan, X., M.I. Gómez, S.S. Atallah, and J.M. Conrad. 2020. "A Bayesian State-Space Approach for Invasive Species Management: The Case of Spotted Wing *Drosophila*." *American Journal of Agricultural Economics* 102(4):1227–1244.
- Farnsworth, D., K.A. Hamby, M. Bolda, R.E. Goodhue, J.C. Williams, and F.G. Zalom. 2017. "Economic Analysis of Revenue Losses and Control Costs Associated with the Spotted Wing *Drosophila*, *Drosophila suzukii* (Matsumura), in the California Raspberry Industry." *Pest Management Science* 73(6):1083–1090.
- Finnoff, D., J.F. Shogren, B. Leung, and D. Lodge. 2005. "The Importance of Bioeconomic Feedback in Invasive Species Management." *Ecological Economics* 52(3):367–381.
- Gautam, B.K., B.A. Little, M.D. Taylor, J.L. Jacobs, W.E. Lovett, R.M. Holland, and A.A. Sial. 2016. "Effect of Simulated Rainfall on the Effectiveness of Insecticides Against Spotted Wing *Drosophila* in Blueberries." *Crop Protection* 81:122–128.
- Gayle, G.-L., L. Golan, and M.A. Soytaş. 2018. "Estimation of Dynastic Life-Cycle Discrete Choice Models." *Quantitative Economics* 9(3):1195–1241.
- Goeb, J., A. Dillon, F. Lupi, and D. Tschirley. 2020. "Pesticides: What You Don't Know *Can* Hurt You." *Journal of the Association of Environmental and Resource Economists* 7(5):801–836.
- Goodhue, R.E., M. Bolda, D. Farnsworth, J.C. Williams, and F.G. Zalom. 2011. "Spotted Wing *Drosophila* Infestation of California Strawberries and Raspberries: Economic Analysis of Potential Revenue Losses and Control Costs." *Pest Management Science* 67(11):1396–1402.
- Gress, B.E., and F.G. Zalom. 2019. "Identification and Risk Assessment of Spinosad Resistance in a California Population of *Drosophila suzukii*." *Pest Management Science* 75(5):1270–1276.

- Grogan, K.A., and R.E. Goodhue. 2012. "Spatial Externalities of Pest Control Decisions in the California Citrus Industry." *Journal of Agricultural and Resource Economics* 37(1):156–179.
- Hamby, K.A., D. E. Bellamy, J.C. Chiu, J.C. Lee, V.M. Walton, N.G. Wiman, R.M. York, and A. Biondi. 2016. "Biotic and Abiotic Factors Impacting Development, Behavior, Phenology, and Reproductive Biology of *Drosophila suzukii*." *Journal of Pest Science* 89(3):605–619.
- Hotz, V.J., and R.A. Miller. 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *Review of Economic Studies* 60(3):497–529.
- Iskhakov, F. 2010. "Structural Dynamic Model of Retirement With Latent Health Indicator." *Econometrics Journal* 13:S126–S161.
- Jacquet, F., J.-P. Butault, and L. Guichard. 2011. "An Economic Analysis of the Possibility of Reducing Pesticides in French Field Crops." *Ecological Economics* 70(9):1638–1648.
- Janssen, S., and M.K. van Ittersum. 2007. "Assessing Farm Innovations and Responses to Policies: A Review of Bio-Economic Farm Models." *Agricultural Systems* 94(3):622–636.
- Keane, M.P., P.E. Todd, and K.I. Wolpin. 2011. "Chapter 4 - The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications." In O. Ashenfelter and D. Card, eds. *Handbook of Labor Economics*. Elsevier, pp. 331–461.
- Kling, C.L., R.W. Arriitt, G. Calhoun, and D.A. Keiser. 2017. "Integrated Assessment Models of the Food, Energy, and Water Nexus: A Review and an Outline of Research Needs." *Annual Review of Resource Economics* 9(1):143–163.
- Kraft, L.J., D.A. Yeh, M.I. Gómez, and H.J. Burrack. 2020. "Determining the Effect of Postharvest Cold Storage Treatment on the Survival of Immature *Drosophila suzukii* (Diptera: Drosophilidae) in Small Fruits" D. H. Cha, ed. *Journal of Economic Entomology* 113(5):2427–2435.
- Kuosmanen, T., D. Pems, and J. Wesseler. 2006. "Specification and Estimation of Production Functions Involving Damage Control Inputs: A Two-Stage, Semiparametric Approach." *American Journal of Agricultural Economics* 88(2):499–511.
- Leach, H., J. Moses, E. Hanson, P. Fanning, and R. Isaacs. 2018. "Rapid Harvest Schedules and Fruit Removal as Non-Chemical Approaches for Managing Spotted Wing *Drosophila*." *Journal of Pest Science* 91(1):219–226.
- Leach, H., S. Van Timmeren, and R. Isaacs. 2016. "Exclusion Netting Delays and Reduces *Drosophila suzukii* (Diptera: Drosophilidae) Infestation in Raspberries." *Journal of Economic Entomology* 109(5):2151–2158.
- Li, S., Y. Liu, and C.D. Wei. 2021. "The Cost of Greening Stimulus: A Dynamic Analysis of Vehicle Scrappage Programs." Working paper, Cornell University.
- Lin, C.-Y.C. 2013. "Strategic Decision-Making with Information and Extraction Externalities: A Structural Model of the MultiStage Investment Timing Game in Offshore Petroleum Production." *The Review of Economics and Statistics* 95(5):1601–1621.

- Lohr, L., T. Park, and L. Higley. 1999. "Farmer Risk Assessment for Voluntary Insecticide Reduction." *Ecological Economics* 30(1):121–130.
- Magnac, T., and D. Thesmar. 2002. "Identifying Dynamic Discrete Decision Processes." *Econometrica* 70(2):801–816.
- Mérel, P., and R. Howitt. 2014. "Theory and Application of Positive Mathematical Programming in Agriculture and the Environment." *Annual Review of Resource Economics* 6(1):451–470.
- Misra, S., and H.S. Nair. 2011. "A Structural Model of Sales-Force Compensation Dynamics: Estimation and Field Implementation." *Quantitative Marketing and Economics* 9(3):211–257.
- Oliva, P., B.K. Jack, S. Bell, E. Mettetal, and C. Severen. 2020. "Technology Adoption Under Uncertainty: Take-up and Subsequent Investment in Zambia." *Review of Economics and Statistics* 102(3):617–632.
- Parent, S.-É., J. Lafond, M.C. Paré, L.E. Parent, and N. Ziadi. 2020. "Conditioning Machine Learning Models to Adjust Lowbush Blueberry Crop Management to the Local Agroecosystem." *Plants* 9(10):1401.
- Reeling, C., V. Verdier, and F. Lupi. 2020. "Valuing Goods Allocated via Dynamic Lottery." *Journal of the Association of Environmental and Resource Economists* 7(4):721–749.
- Rendon, D., K.A. Hamby, A.L. Arsenault-Benoit, C.M. Taylor, R.K. Evans, C.R. Roubos, A.A. Sial, M. Rogers, A. Petran, S. Van Timmeren, P. Fanning, R. Isaacs, and V. Walton. 2020. "Mulching as a Cultural Control Strategy for *Drosophila suzukii* in Blueberry." *Pest Management Science* 76(1):55–66.
- Rendon, D., and V.M. Walton. 2019. "Drip and Overhead Sprinkler Irrigation in Blueberry as Cultural Control for *Drosophila suzukii* (Diptera: Drosophilidae) in Northwestern United States." *Journal of Economic Entomology* 112(2):745–752.
- Rust, J. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica* 55(5):999.
- Rust, J. 1994. "Structural Estimation of Markov Decision Processes." In Robert F. Engle and Daniel L. McFadden, eds. *Handbook of Econometrics, Volume IV*. Elsevier, pp. 3081–3143.
- Ryan, S.P. 2012. "The Costs of Environmental Regulation in a Concentrated Industry." *Econometrica* 80(3):1019–1061.
- Sambucci, O., C.-Y.C. Lin Lawell, and T.J. Lybbert. 2022. "Pesticide Spraying and Disease Forecasts: A Dynamic Structural Econometric Model of Grape Growers in California." Working paper, Cornell University.
- Scott, P. 2013. "Dynamic Discrete Choice Estimation of Agricultural Land Use." Working paper, New York University.
- Sears, L., D. Lim, and C.-Y.C. Lin Lawell. 2019. "Spatial Groundwater Management: A Dynamic Game Framework and Application to California." *Water Economics and Policy* 5(1):1850019.

- Sears, L., C.-Y.C. Lin Lawell, and M.T. Walter. 2022. “Groundwater Under Open Access: A Structural Model of the Dynamic Common Pool Extraction Game.” Working paper, Cornell University.
- Sexton, S.E. 2007. “The Economics of Pesticides and Pest Control.” *International Review of Environmental and Resource Economics* 1(3):271–326.
- Smith, M.D. 2008. “Bioeconometrics: Empirical Modeling of Bioeconomic Systems.” *Marine Resource Economics* 23(1):1–23.
- Stockton, D.G., S.P. Hesler, A.K. Wallingford, T.C. Leskey, L. McDermott, J.E. Elsensohn, D.I.M. Riggs, M. Pritts, and G.M. Loeb. 2020. “Factors Affecting the Implementation of Exclusion Netting to Control *Drosophila suzukii* on Primocane Raspberry.” *Crop Protection* 135:105191.
- Sunding, D., and J. Zivin. 2000. “Insect Population Dynamics, Pesticide Use, and Farmworker Health.” *American Journal of Agricultural Economics* 82(3):527–540.
- Timmins, C. 2002. “Measuring the Dynamic Efficiency Costs of Regulators’ Preferences: Municipal Water Utilities in the Arid West.” *Econometrica* 70(2):603–629.
- Tochen, S., D.T. Dalton, N. Wiman, C. Hamm, P.W. Shearer, and V.M. Walton. 2014. “Temperature-Related Development and Population Parameters for *Drosophila suzukii* on Cherry and Blueberry.” *Environmental Entomology* 43(2):501–510.
- University of Maine Cooperative Extension. 2010. “Two-Year Wild Blueberry Production Cycle.” *Maine Wild Blueberries*. Available at: <https://extension.umaine.edu/blueberries/factsheets/production/wild-blueberry-2-year-production-cycle/> [Accessed March 22, 2021].
- U.S. Department of Agriculture. 2019. “2018 Maine Wild Blueberry Production.” Available at: [https://www.nass.usda.gov/Statistics by State/New England includes/Publications/Current News Release/2019/2018%20ME%20Wild%20Blueberry%20Release.pdf](https://www.nass.usda.gov/Statistics%20by%20State/New%20England/includes/Publications/Current%20News%20Release/2019/2018%20ME%20Wild%20Blueberry%20Release.pdf).
- U.S. Department of Agriculture. 2021. “The Quick Stats Database.” Available at: <https://data.nal.usda.gov/dataset/nass-quick-stats>.
- U.S. Department of Agriculture. 2022. “Pest Management.” Available at: <https://nifa.usda.gov/topic/pest-management>.
- Van Timmeren, S., L. Horejsi, S. Larson, K. Spink, P. Fanning, and R. Isaacs. 2017. “Diurnal Activity of *Drosophila suzukii* (Diptera: Drosophilidae) in Highbush Blueberry and Behavioral Response to Irrigation and Application of Insecticides.” *Environmental Entomology* 46(5):1106–1114.
- Walsh, D.B., M.P. Bolda, R.E. Goodhue, A.J. Dreves, J. Lee, D.J. Bruck, V.M. Walton, S.D. O’Neal, and F.G. Zalom. 2011. “*Drosophila suzukii* (Diptera: Drosophilidae): Invasive Pest of Ripening Soft Fruit Expanding its Geographic Range and Damage Potential.” *Journal of Integrated Pest Management* 2(1):G1–G7.
- Waterfield, G., and D. Zilberman. 2012. “Pest Management in Food Systems: An Economic Perspective.” *Annual Review of Environment and Resources* 37(1):223–245.

- Whittle, P. 2021. “Maine Blueberry Growers Had Bounce-Back Year, Aided By Rain.” *Associated Press News* 26 September 2021. Available at: <https://apnews.com/article/business-environment-and-nature-maine-92eb728017366db08cd4acc6f74490e4> [Accessed May 23, 2022].
- Wiman, N.G., V.M. Walton, D.T. Dalton, G. Anfora, H.J. Burrack, J.C. Chiu, K.M. Daane, A. Grassi, B. Miller, S. Tochen, X. Wang, and C. Ioriatti. 2014. “Integrating Temperature-Dependent Life Table Data into a Matrix Projection Model for *Drosophila suzukii* Population Estimation.” *PLOS ONE* 9(9):e106909.
- Yarborough, D.E. 2012. “Establishment and Management of the Cultivated Lowbush Blueberry.” *International Journal of Fruit Science* 12(1–3):14–22.
- Yarborough, D.E., and J. D’Appollonio. 2017. “Cooperative Extension: Maine Wild Blueberries, 260-Blueberry Enterprise Budget.” Fact Sheet No. 260, UMaine Extension No. 2016. University of Maine Cooperative Extension. Available at: <https://extension.umaine.edu/blueberries/factsheets/marketing-and-business-management/260-blueberry-enterprise-budget/> [Accessed April 30, 2019].
- Yeh, D.A., F.A. Drummond, and M.I. Gómez. 2019. “Economic Management Support Tool for Spotted Wing Drosophila (SWD) in Maine Wild Blueberry Production.” Available at: <https://extension.umaine.edu/blueberries/tools/> [Accessed June 30, 2020].
- Yeh, D.A., F.A. Drummond, M.I. Gómez, and X. Fan. 2020. “The Economic Impacts and Management of Spotted Wing Drosophila (*Drosophila Suzukii*): The Case of Wild Blueberries in Maine.” *Journal of Economic Entomology* 113(3):1262–1269.

Tables and Figures

Table 1. Yield loss as a function of observed larval infestation

Observed larval infestation	Index for y_{larva}	Yield loss $loss(y_{larva})$	Infestation loss
Less than 0.3%	0 and 1	0%	0%
Between 0.3% and 1%	2	10%	10%
Between 1% and 5%	3	30%	20%
Between 5% and 10%	4	50%	20%
More than 10%	5	80%	30%

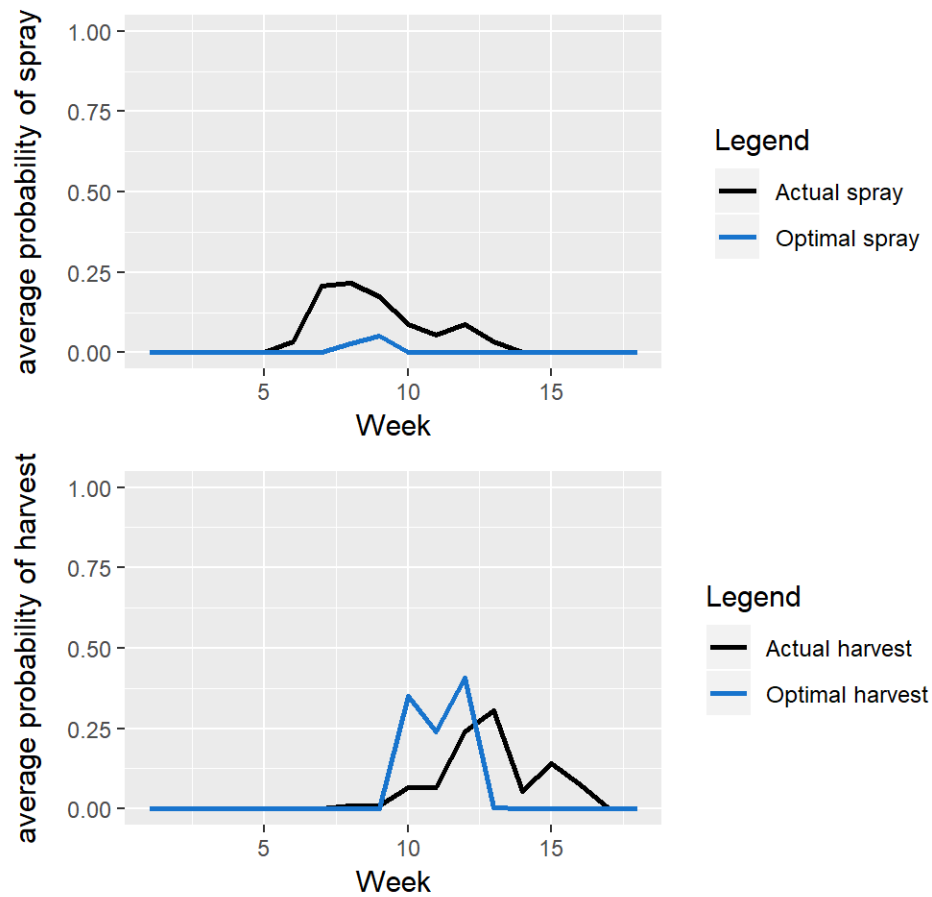
Notes: The yield loss $loss(y_{larva})$ as a function of observed larval infestation y_{larva} is based on expert opinion and extension reports (Burrack 2014; De Ros et al. 2015; DiGiacomo et al. 2019; Drummond et al. 2019; Yeh et al. 2019; Yeh et al. 2020). The infestation loss is the incremental yield loss at each tier of observed larval infestation, as calculated from the (cumulative) yield loss $loss(y_{larva})$.

Table 2. Average farm welfare

	All farms	Organic farms
Actual farm welfare (\$/acre)	826.7 (217.9)	20,678.9 (3,048.2)
Optimal farm welfare (\$/acre)	963.6 (0.2)	22,078.7 (0.0)
Deadweight loss (\$/acre)	136.9 (217.9)	1,399.8 (3,048.2)

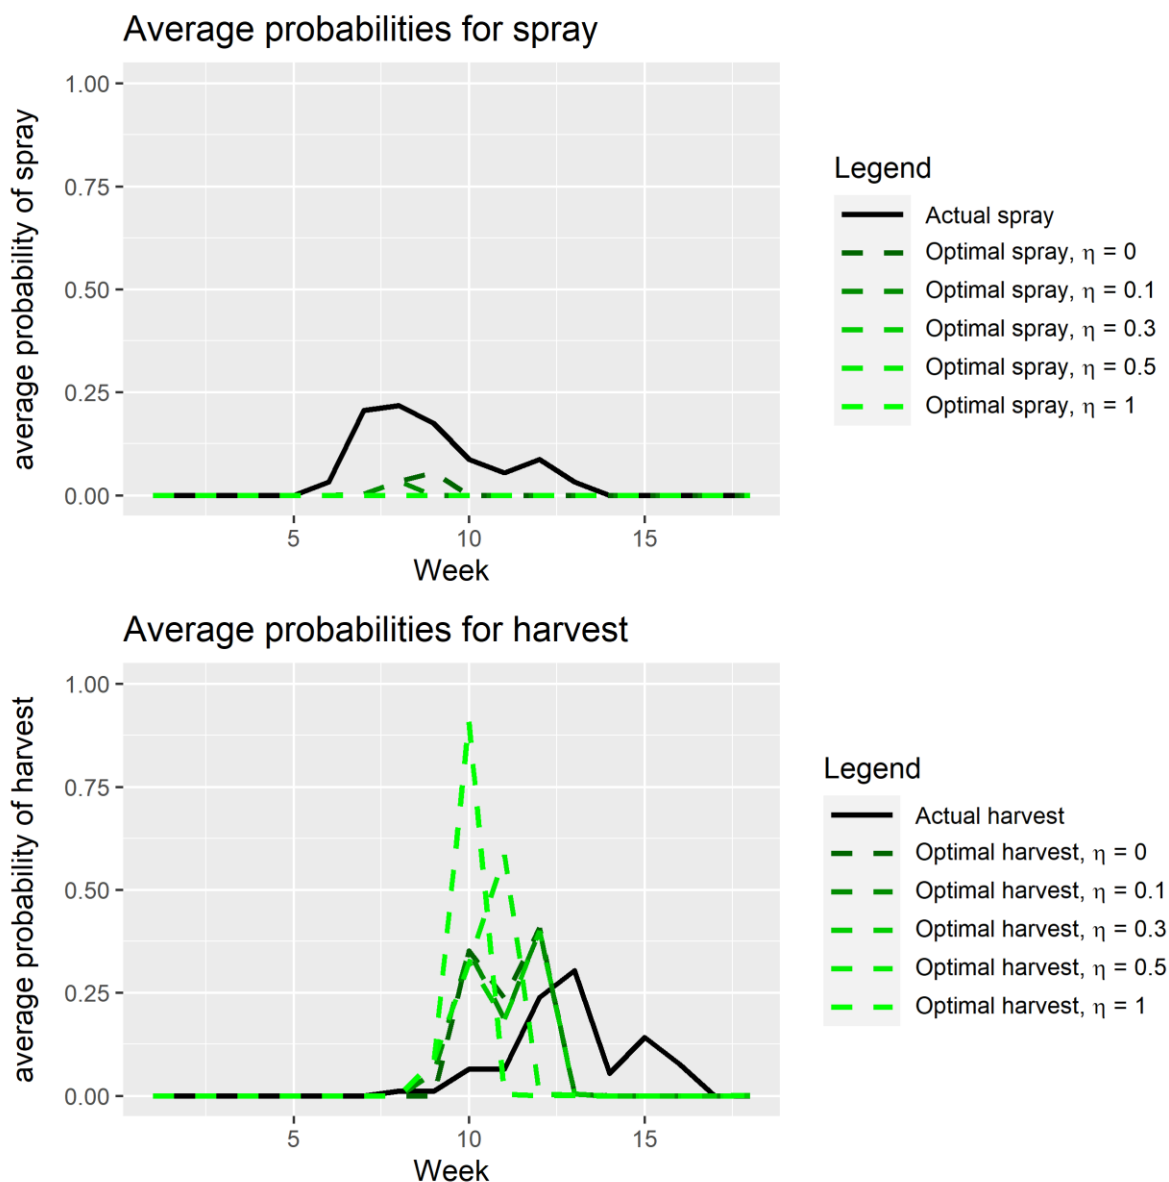
Notes: Table reports average farm welfare, as calculated using the assumed parameter values from our numerical model, under our base case specification. The assumed values of the parameters used in the numerical model are based on economic data, expert opinion, and extension reports. Standard deviations across farms are in parentheses. We define welfare as the present discounted value (PDV) of the entire stream of per-period payoffs over the entire season for the farm. We calculate actual welfare using observed actions and states of each farm. The optimal welfare, as given by the value function evaluated at on the initial observed states, is the maximum expected PDV of the entire stream of per-period payoffs the farm could have received that year if it followed the optimal strategy. We calculate the deadweight loss from the growers' actual actions as optimal welfare minus actual welfare. There are a total of 92 farms, of which 19 farms are organic.

Figure 1. Optimal vs. actual probabilities of spraying and harvesting



Notes: Figure compares the optimal and actual average probabilities of spraying and harvesting. The average probabilities of actual spray and actual harvest are calculated by averaging the observed actions over all 92 farms in the data set. The optimal average probabilities of spraying and harvesting are given by averaging over 100 simulations for each farm using the policy function from our numerical model and the non-parametrically estimated transition densities, and then averaging over all 92 farms.

Figure 2. Optimal vs. actual probabilities of spraying and harvesting when growers are risk averse



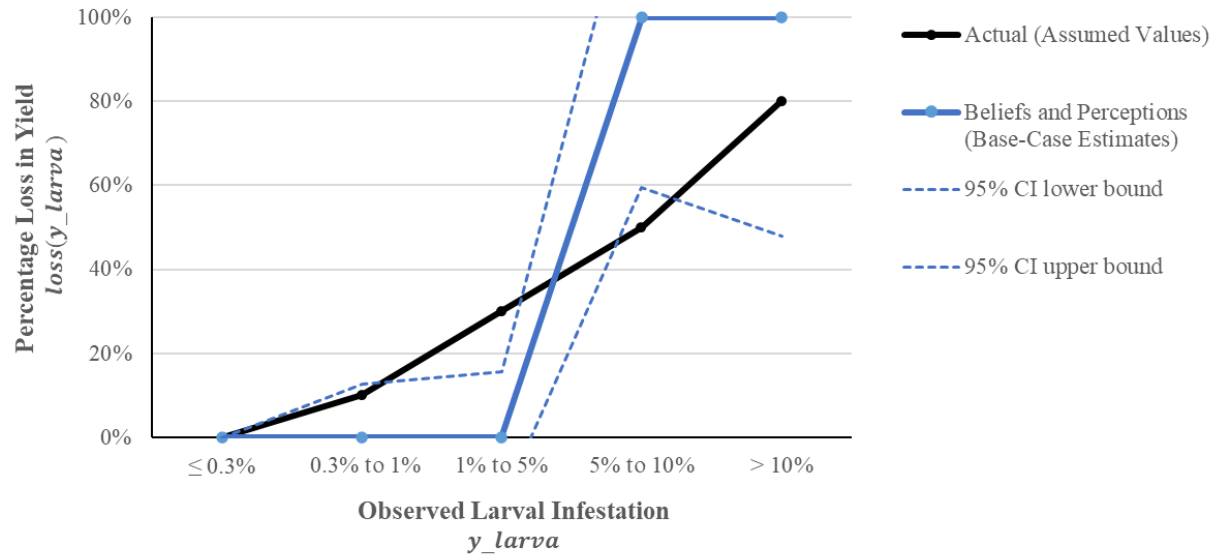
Notes: In this alternative specification, the grower utility from the uncertain crop revenue is a constant relative risk aversion (CRRA) utility function, where η is the coefficient of constant relative risk aversion. Figure compares the optimal and actual average probabilities of spraying and harvesting for various values of the coefficient of constant relative risk aversion η . The average probabilities of actual spray and actual harvest are calculated by averaging the observed actions over all 92 farms in the data set. For each value of η , the optimal average probabilities of spraying and harvesting are given by averaging over 100 simulations for each farm using the policy function from our numerical model and the non-parametrically estimated transition densities, and then averaging over all 92 farms.

Table 3. Structural parameter estimates

	Actual (Assumed Values)	Structural Parameter Estimates: Base-Case
Infestation loss at tier 1 (Infestation \leq 0.3%)	0	0.000 (0.000)
Infestation loss at tier 2 (0.3% < Infestation \leq 1%)	0.10	0.000 (0.065)
Infestation loss at tier 3 (1% < Infestation \leq 5%)	0.20	0.000 (0.046)
Infestation loss at tier 4 (5% < Infestation \leq 10%)	0.20	1.000*** (0.191)
Infestation loss at tier 5 (10% < Infestation)	0.30	0.000 (0.166)
Spray cost (\$ per acre)	40	2,966.17*** (96.123)
# Observations		1,656
# Farms		92

Notes: The actual values are the assumed values of the parameters used in the numerical model and are based on economic data, expert opinion, and extension reports. The structural parameter estimates are the parameter estimates from our base-case specification of the structural model. Infestation is the observed larval infestation. The infestation loss is the incremental yield loss at each tier of observed larval infestation; the cumulative value of infestation loss is the (cumulative) yield loss $loss(y_{larva})$. The assumed value of infestation loss is the value based on expert opinion and extension reports (Burrack 2014; De Ros et al. 2015; DiGiacomo et al. 2019; Drummond et al. 2019; Yeh et al. 2019; Yeh et al. 2020) reported in Table 1. Bootstrapped standard errors in parentheses. Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Figure 3. Growers' beliefs and perceptions about infestation loss



Notes: Figure compares the actual yield loss $loss(y_{larva})$ as a function of observed larval infestation y_{larva} with growers' beliefs and perceptions about yield loss. The yield loss $loss(y_{larva})$ is the cumulative percentage of yield loss at each tier of observed larval infestation. The actual yield loss $loss(y_{larva})$ values are the assumed values based on expert opinion and extension reports (Burrack 2014; De Ros et al. 2015; DiGiacomo et al. 2019; Drummond et al. 2019; Yeh et al. 2019; Yeh et al. 2020) reported in Table 1. The growers' beliefs and perceptions are based on the structural parameter estimates from our base-case specification of the structural model in Table 3. The dashed blue lines indicate the 95% confidence interval for the growers' beliefs and perceptions, as calculated using the standard errors from our base-case specification of the structural model in Table 3.

Table 4. Robustness: Structural parameter estimates with CRRA utility from crop revenue

	CRRA Utility From Crop Revenue
Infestation loss tier 1 (Infestation $\leq 0.3\%$)	0.000 (0.000)
Infestation loss tier 2 ($0.3\% < \text{Infestation} \leq 1\%$)	0.000 (0.000)
Infestation loss tier 3 ($1\% < \text{Infestation} \leq 5\%$)	0.000 (0.020)
Infestation loss tier 4 ($5\% < \text{Infestation} \leq 10\%$)	1.000*** (0.240)
Infestation loss tier 5 ($10\% < \text{Infestation}$)	0.000 (0.240)
Spray cost (\$ per acre)	2,904.999*** (0.007)
Coefficient of constant relative risk aversion η	0.000*** (0.000)
# Observations	1,656
# Farms	92

Notes: In this alternative specification, the grower utility from the uncertain crop revenue is a constant relative risk aversion (CRRA) utility function, where η is the coefficient of constant relative risk aversion. Infestation is the observed larval infestation. The infestation loss is the incremental yield loss at each tier of observed larval infestation; the cumulative value of infestation loss is the (cumulative) percentage of yield loss $loss(y_{larva})$. Bootstrapped standard errors in parentheses. Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 5a: Structural parameter estimates holding infestation loss fixed

	Parameter Estimate
Spray cost (\$ per acre)	2,917.77*** (96.743)
<i>Likelihood ratio test to compare this model with base-case structural model</i>	
<i>H0: This alternative model fits the data better than the base-case model in Table 3 does</i>	
LR Test statistic D	-0.868

Notes: In this alternative specification, infestation loss is held fixed at its assumed values based on expert opinion and extension reports in Table 1. Infestation is the observed larval infestation. The infestation loss is the incremental yield loss at each tier of observed larval infestation; the cumulative value of infestation loss is the (cumulative) yield loss $loss(y_{larva})$. There are 1,656 observations spanning 92 farms. Bootstrapped standard errors in parentheses. Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

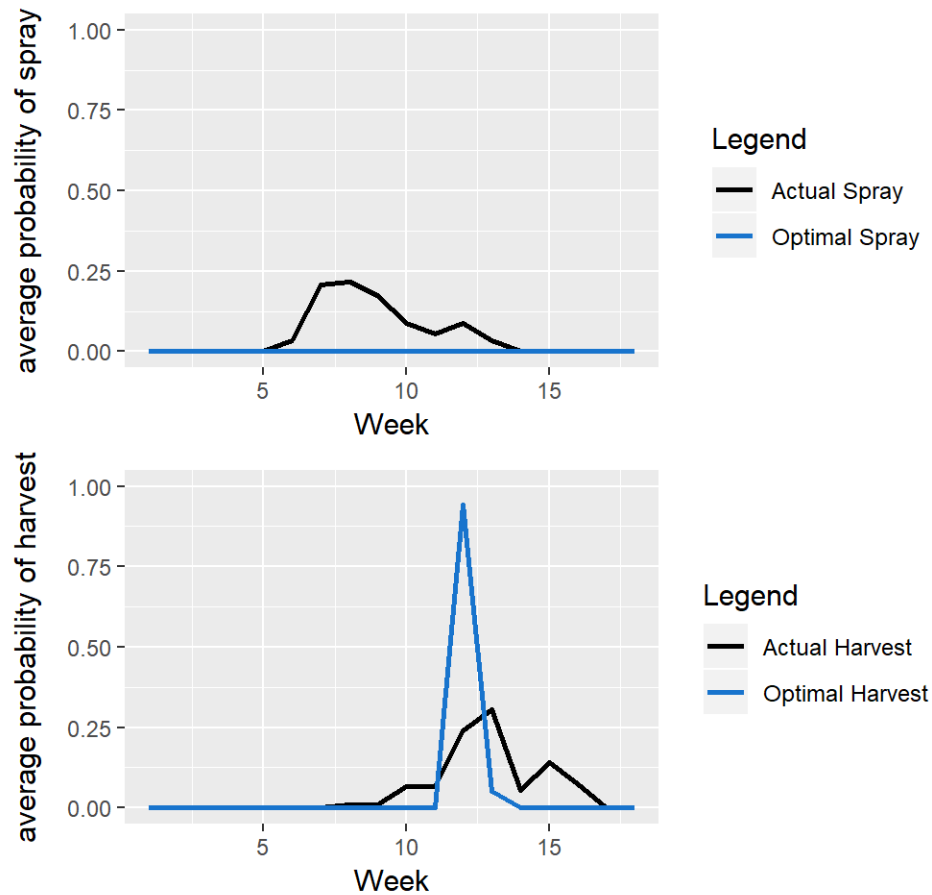
Table 5b: Structural parameter estimates holding spray costs fixed

	Parameter Estimate
Infestation loss at tier 1 (Infestation $\leq 0.3\%$)	0.000 (0.000)
Infestation loss at tier 2 ($0.3\% < \text{Infestation} \leq 1\%$)	0.000 (0.000)
Infestation loss at tier 3 ($1\% < \text{Infestation} \leq 5\%$)	0.000 (0.00)
Infestation loss at tier 4 ($5\% < \text{Infestation} \leq 10\%$)	0.121 (0.207)
Infestation loss at tier 5 ($10\% < \text{Infestation}$)	0.709*** (0.253)

Likelihood ratio test to compare this model with base-case structural model
H0: This alternative model fits the data better than the base-case model in Table 3 does
 LR Test statistic D 1,986.4***

Notes: In this alternative specification, spray cost is held fixed at its assumed value of \$40 per acre (Esau 2019). Infestation is the observed larval infestation. The infestation loss is the incremental yield loss at each tier of observed larval infestation; the cumulative value of infestation loss is the (cumulative) yield loss $loss(y_{larva})$. There are 1,656 observations spanning 92 farms. Bootstrapped standard errors in parentheses. Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Figure 4. Optimal probabilities of spraying and harvesting conditional on growers' beliefs and perceptions



Notes: Figure compares the optimal and actual average probabilities of spraying and harvesting, conditional on growers' beliefs from the base-case structural parameter estimates in Table 3. The average probabilities of actual spray and actual harvest are calculated by averaging the observed actions over all 92 farms in the data set. The optimal average probabilities of spraying and harvesting conditional on growers' beliefs use the parameters estimated from the structural model, and are given by averaging over 100 simulations for each farm using the policy function from solving our numerical model using the parameters estimated from the structural model and the non-parametrically estimated transition densities, and then averaging over all 92 farms.

Table 6. Average farm welfare conditional on growers' beliefs and perceptions

	All farms
Actual farm welfare conditional on growers' beliefs and perceptions (\$/acre)	-1,653.2 (2520.4)
Optimal farm welfare conditional on growers' beliefs and perceptions (\$/acre)	985.3 (0.0)
Deadweight loss (\$/acre)	2,638.5 (2520.4)

Notes: Table reports average farm welfare conditional on growers' beliefs from the base-case structural parameter estimates in Table 3. Standard deviations across farms are in parentheses. We define welfare as the present discounted value (PDV) of the entire stream of per-period payoffs over the entire season for the farm. We calculate actual welfare using observed actions and states of each farm. The optimal welfare, as given by the value function evaluated at on the initial observed states, is the maximum expected PDV of the entire stream of per-period payoffs the farm could have received that year if it followed the optimal strategy. We calculate the deadweight loss from the growers' actual actions as optimal welfare minus actual welfare. There are a total of 92 farms.

Appendix

Appendix A.

State Transition Densities

A.1. State Variables

The state tuple $state_t \equiv (y_{larva,t}, y_{adult,t}, int_t, precip_t, tmax_t)$ is a vector of state variables measuring SWD population levels (observed larva SWD $y_{larva,t}$ and observed adult SWD $y_{adult,t}$), weather variables (weekly accumulated precipitation $precip_t$ and weekly maximum temperature $temp_t$), and the interval of time since the last insecticide application (int_t).

The SWD population levels in the state tuple $state_t$ are indicated by two state variables: observed larva SWD $y_{larva,t}$, which affects the expected yield loss $loss(\cdot)$ and is discretized to six levels; and observed adult SWD $y_{adult,t}$, which is a dummy variable indicating whether adult SWD was observed at time t . Observed larva SWD $y_{larva,t}$ affects the expected yield loss $loss(\cdot)$ and therefore the crop revenue from harvesting. Based on expert opinion and extension reports (Burrack 2014; De Ros et al. 2015; DiGiacomo et al. 2019; Drummond et al. 2019; Yeh, Drummond and Gómez 2019; Yeh et al. 2020), our assumed values for the expected yield loss $loss(y_{larva,t})$ as a function of observed larval infestation $y_{larva,t}$ has five tiers, as reported in Table 1. Since the expected yield loss $loss(y_{larva,t})$ as a function of observed larval infestation $y_{larva,t}$ has five tiers, we discretize observed larva SWD $y_{larva,t}$ into six levels: one level for no observed larva SWD, plus five levels for each of the five tiers of the yield loss function $loss(y_{larva,t})$.

In contrast to observed larva SWD $y_{larva,t}$, which affects crop revenue through its effect on expected yield loss, observed adult SWD $y_{adult,t}$ does not directly affect expected yield loss or crop revenue, but instead indirectly affects expected yield loss through its effect on the population dynamics of observed larva SWD $y_{larva,t}$ via the transition density. Thus, owing to state space constraints, we discretize observed adult SWD $y_{adult,t}$ into a binary dummy variable indicating whether adult SWD was observed at time t ; given that we already discretize observed larva SWD $y_{larva,t}$ into six levels, a finer discretization of observed adult SWD $y_{adult,t}$ is neither feasible or desirable and would lead to insufficient observations for many values of the state tuples.

The state tuple $state_t$ also includes a discretized variable int_t for the interval of time since the last insecticide application, which can take the value of 1 if the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective; 2 if the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective; or 3 if the grower has not yet sprayed this season prior to week t . Since

the insecticide is effective for a maximum of 7 days, the maximum number of weeks for which the insecticide is effective is 1.

As studies have shown that precipitation may affect spraying efficacy while the temperature affects SWD population dynamics (Tochen et al. 2014; Wiman et al. 2014; Hamby et al. 2016; Gautam et al. 2016; Van Timmeren et al. 2017),¹³ the state tuple $state_t$ additionally includes two weather-related dummy variables: a dummy variable for weekly accumulated precipitation exceeding 2 mm ($precip_t$) and a dummy variable for weekly maximum temperature exceeding 27 degrees Celsius ($tmax_t$).

A.2. Methodology for Estimating Transition Densities

We assume that the state variables follow a first-order Markov process. The transition densities $Pr_t(state_{t+1}|state_t, a_t)$ governing the evolution of state variables from one period to the next given the state variables and choice variable this period are specified as follows.

For SWD population levels (i.e., $y_{larva,t}$ and $y_{adult,t}$), we non-parametrically estimate the transition densities for the stochastic SWD population levels conditional on the spraying decision, weather variables, the interval of time since last spray, and the time of season.

The distributions for stochastic precipitation $precip_t$ and temperature $tmax_t$, which we allow to vary by time of season, are estimated using empirical averages in the actual data.

The interval of time since the last spray (int_t) evolves deterministically as a function of this period's action and the value of the previous interval of time since the last spray.

We non-parametrically estimate the transition densities for the stochastic SWD population levels conditional on the spraying decision, weather variables, the interval of time since last spray, and the time of season. In particular, we estimate a total of 10 transition densities for the SWD dynamics. The detailed breakdown of the 10 transition densities is as follows. We divide the 18-week season into three equal periods, representing the early, mid- and late season, respectively. In our observed data, spraying only occurs mid-season. If the discretized variable for the time interval

¹³ We focus on modeling the effects of weather on SWD population dynamics and spraying efficacy. While weather may also affect lowbush blueberry yields as well, the possible variation in yield due to variations in weather during the growing season is relatively small (Parent et al. 2020), especially compared to yield losses from SWD infestation. Thus, as SWD is a larger source of uncertainty and impact on yield than weather is for lowbush blueberry production in Maine, we focus on modeling the uncertainty that arises from SWD and the effects of weather on SWD population dynamics and spraying efficacy.

since last spray (int_t) is equal to 1 (i.e., if the grower sprays this week), then the transition densities are further conditioned on precipitation $precip_t$, which affects the spraying efficacy. If the discretized variable for the time interval since last spray (int_t) is equal to 3 (i.e., the grower does not spray this period and is not within effective periods of past insecticide applications), then we further condition the transition densities on temperature $tmax_t$, and separately estimate the transition densities for early, mid-, and late season. If the discretized variable for the time interval since last spray (int_t) is equal to 2 (i.e., if the grower does not spray this period but is within effective periods of past insecticide applications), we estimate the transition densities separately for mid- and late season, but we do not further condition on other state variables given the limited data variation.

A.3. Non-Parametric Estimates of SWD Transition Densities

We non-parametrically estimate the transition densities for the stochastic SWD population levels conditional on the spraying decision, weather variables, the interval of time since last spray, and the time of season. In particular, we estimate a total of 10 transition densities for the SWD dynamics. Each of the 10 SWD transition densities therefore gives the distribution of SWD population levels next week ($y_{larva,t+1}, y_{adult,t+1}$) conditional on the SWD population levels this week ($y_{larva,t}, y_{adult,t}$), and conditional on the spraying decision, weather variables, the interval of time since last spray, and the time of season. In particular, each cell in each SWD transition matrix gives the probability that SWD population levels next week ($y_{larva,t+1}, y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}, y_{adult,t}$) being the values in the row, and conditional on the spraying decision, weather variables, the interval of time since last spray, and the time of season.

A.3.1. Early season

Table A1 reports the non-parametrically estimated transition density for the SWD population levels for the early part of the season (weeks $t \leq 5$), conditional on temperature $tmax_t$. Table A1a reports the estimated transition density for the early part of the season (weeks $t \leq 5$), conditional on temperature $tmax_t$ in week t being low; while Table A1b reports the estimated transition density for the early part of the season (weeks $t \leq 5$), conditional on temperature $tmax_t$

in week t being high. The transition density gives the distribution of SWD population levels next week $(y_{larva,t+1}, y_{adult,t+1})$ conditional on the SWD population levels this week $(y_{larva,t}, y_{adult,t})$, conditional on temperature $tmax_t$ in week t , and conditional on week $t \leq 5$. In particular, each cell in Table A1a gives the probability that SWD population levels next week $(y_{larva,t+1}, y_{adult,t+1})$ are the values given in the column, conditional on the SWD population levels this week $(y_{larva,t}, y_{adult,t})$ being the values in the row, conditional on temperature $tmax_t$ in week t being low (i.e., maximum temperature that week is less than or equal to 27 degrees Celsius), and conditional on week $t \leq 5$. Similarly, each cell in Table A1b gives the probability that SWD population levels next week $(y_{larva,t+1}, y_{adult,t+1})$ are the values given in the column, conditional on the SWD population levels this week $(y_{larva,t}, y_{adult,t})$ being the values in the row, conditional on temperature $tmax_t$ in week t being high (i.e., maximum temperature that week exceeds 27 degrees Celsius), and conditional on week $t \leq 5$.

For the rows in Table A1 that correspond to $(y_{larva,t}, y_{adult,t})$ tuples that were observed in the data during weeks $t \leq 5$ of the season (i.e., the early season), we report the non-zero values of the transition density; the remaining cells in these rows are all 0's. For $(y_{larva,t}, y_{adult,t})$ tuples that were never observed in the data during weeks $t \leq 5$ of the season (i.e., the early season), we highlight the row of that transition density for that $(y_{larva,t}, y_{adult,t})$ tuple in light grey.¹⁴

During the early season (weeks $t \leq 5$), we do not observe any spraying in the data, so we cannot additionally condition on int_t or on the grower's spraying decision s_t in week t during the first 5 weeks of the season, since there is no variation in int_t or s_t during the first 5 weeks of the season. In particular, $int_t = 3$ (grower has not yet sprayed this season prior to that week) and $s_t = 0$ (does not spray) for all observations in the early season (weeks $t \leq 5$). Thus, the transition densities in Tables A1a and A1b are estimated from pooling over all observations in the first 5 weeks of the season for which the temperature in week t is low or high, respectively, regardless of the values of int_t or s_t , since $int_t = 3$ (grower has not yet sprayed this season prior to that week) and $s_t = 0$ (does not spray) for all observations in the first 5 weeks of the season. Moreover,

¹⁴ For $(y_{larva,t}, y_{adult,t})$ tuples that were never observed in the data during weeks $t \leq 5$ of the season (i.e., the early part of the season), it does not matter what we put in their respective rows in the transition matrix since we would never end up at these $(y_{larva,t}, y_{adult,t})$ tuples. The value functions and policy functions for the early part of the season for $(y_{larva,t}, y_{adult,t})$ tuples we never see during the early part of the season therefore do not have any meaning. We also do not simulate forward starting in the first week from any state tuples that we never see in the data.

since we do not observe any spraying in the first 5 weeks of the season, and since the reason precipitation matters to the SWD transition densities is that it decreases the effectiveness of spray, precipitation does not matter in the first 5 weeks of the season so we do not need to condition on the precipitation $precip_t$ during the first 5 weeks of the season.

As seen in Table A1, SWD population levels are relatively low during the early part of the season (weeks $t \leq 5$). During the early season (weeks $t \leq 5$), no adult SWD was observed in the data for any observations (i.e., $y_{adult,t} = 1$ for all weeks $t \leq 5$), and the observed larva SWD levels in the data were only either none ($y_{larva,t} = 1$) or in the lowest tier ($y_{larva,t} = 2$). Moreover, when temperature $tmax_t$ in week $t \leq 5$ is low (Table A1a), the SWD levels remain the same the following week; but when temperature $tmax_t$ in week $t \leq 5$ is high (Table A1b) and the observed larva SWD level is in the lowest tier ($y_{larva,t} = 2$), there is a roughly 50 percent probability that the observed larva SWD level will decline to none ($y_{larva,t} = 1$) the following week. Thus, when the temperature $tmax_t$ in week $t \leq 5$ is high (i.e., maximum temperature that week exceeds 27 degrees Celsius), this makes it more likely for SWD populations to decline, which is consistent with studies that show that the SWD mortality rate increases dramatically if the temperature exceeds 27 degrees Celsius (Tochen et al. 2014).

A.3.2. Mid-season

Table A2 reports the non-parametrically estimated transition density for the SWD population levels for the mid-season (i.e., weeks $t \in (5, \frac{2}{3}T]$), conditional on being within the effective period of the last insecticide application (i.e., $int_t = 1$) and conditional on not spraying (i.e., $s_t = 0$). In particular, each cell in the table gives the probability that SWD population levels next week ($y_{larva,t+1}, y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}, y_{adult,t}$) being the values in the row, conditional on $int_t = 1$ (i.e., the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective), conditional on $s_t = 0$ (i.e., not spraying), and conditional on week $t \in (5, \frac{2}{3}T]$ (i.e., the mid-season).¹⁵ We do not distinguish between high and

¹⁵ For the rows in Table A2 that correspond to $(y_{larva,t}, y_{adult,t})$ tuples that were observed in the data during weeks $t \in (5, \frac{2}{3}T]$ of the season (i.e., the mid-season) when $int_t = 1$ (i.e., the number of weeks since the last spray is less

low temperature for the mid-season since, conditional on $int_t = 1$ and $s_t = 0$ (i.e., not spraying) there is very little difference in the transition densities for the mid-season (weeks $t \in (5, \frac{2}{3}T]$) when we additionally condition on $tmax_t$. Thus, we assume when $int_t = 1$ (i.e., the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective), $s_t = 0$ (i.e., not spraying), and it is the mid-season (i.e., weeks $t \in (5, \frac{2}{3}T]$), the SWD mortality rate does not depend on temperature.¹⁶

Our estimated transition density in Table A2 shows that, during the mid-season, if the farmer is within the effective period of the last insecticide application and does not spray this week, then the SWD larva population will remain at the same levels following next week.

Table A3 reports the non-parametrically estimated transition density for the SWD population levels for the mid-season (i.e., weeks $t \in (5, \frac{2}{3}T]$), conditional on $int_t \geq 2$ (i.e., either the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective or the grower has not yet sprayed this season prior to week t), conditional on $s_t = 0$ (i.e., not spraying), and conditional on temperature $tmax_t$. Table A3a conditions on temperature $tmax_t$ in week t being low; while Table A3b conditions on temperature $tmax_t$ in week t being high.¹⁷

than or equal to the maximum number of weeks for which the insecticide is effective) and $s_t = 0$ (i.e., not spraying), we report the non-zero values of the transition density; the remaining cells in these rows are all 0's.

¹⁶ For $(y_{larva,t}, y_{adult,t})$ tuples that were never observed in the data during weeks $t \in (5, \frac{2}{3}T]$ of the season (i.e., the mid-season) when $int_t = 1$ (i.e., the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective) and $s_t = 0$ (i.e., not spraying), we highlight the row of that transition density for that $(y_{larva,t}, y_{adult,t})$ tuple in light grey. For these rows, we assume that the SWD population levels next week $(y_{larva,t+1}, y_{adult,t+1})$ are the same as the SWD population levels this week $(y_{larva,t}, y_{adult,t})$. Thus, when there is missing data, we assume that even when the grower does not spray, the SWD population levels will stay the same for those SWD population levels that are missing when the grower does not spray.

¹⁷ For the rows in Table A3 that correspond to $(y_{larva,t}, y_{adult,t})$ tuples that were observed in the data during weeks $t \in (5, \frac{2}{3}T]$ of the season (i.e., the mid-season) when $int_t \geq 2$ (i.e., either the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective or the grower has not yet sprayed this season prior to week t) and $s_t = 0$ (i.e., not spraying), we report the non-zero values of the transition density; the remaining cells in these rows are all 0's. For rows highlighted in light grey corresponding to $(y_{larva,t}, y_{adult,t})$ tuples that were never observed in the data during weeks $t \in (5, \frac{2}{3}T]$ of the season (i.e., the mid-season) when $int_t \geq 2$ (i.e., either the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective or the grower has not yet sprayed this season prior to week t) and $s_t = 0$ (i.e., not spraying), we assume that the SWD population levels next week $(y_{larva,t+1}, y_{adult,t+1})$ are the same as the SWD population levels this week $(y_{larva,t}, y_{adult,t})$. Thus, when there is missing data, we assume that even when the grower does not spray, the SWD population levels will stay the same for those SWD population levels that are missing when the grower does not spray.

For tuples of SWD population levels this week ($y_{larva,t}, y_{adult,t}$) that are observed in the data during the mid-season (weeks $t \in (5, \frac{2}{3}T]$) during weeks when $int_t \geq 2$ (i.e., either the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective or the grower has not yet sprayed this season prior to week t) and $s_t = 0$ (i.e., not spraying), but not during weeks t when the temperature is high (i.e., maximum temperature that week exceeds 27 degrees Celsius), we fill in the row of that transition density for that ($y_{larva,t}, y_{adult,t}$) tuple (highlighted in light blue in Table A3b) with the respective row from the transition density for the mid-season (weeks $t \in (5, \frac{2}{3}T]$) during weeks when $int_t \geq 2$ and $s_t = 0$, conditional on temperature in week t being low (i.e., maximum temperature that week is less than or equal to 27 degrees Celsius) in Table A3a. Thus, for these missing observations, we make a conservative assumption that high temperature does not have any effect on SWD mortality.

Comparing the results of Table A3a and Table A3b. we find that, during the mid-season, when the temperature is very high (Table A3b), this makes it less likely for SWD populations to increase, which is consistent with studies that show that the SWD mortality rate increases dramatically if the temperature exceeds 27 degrees Celsius (Tochen et al. 2014).

Comparing the results of Table A3 with those of Table A2, we find that, during the mid-season, if the farmer does not spray, there is a higher probability that the SWD larva population will increase if the farmer either has never sprayed yet or is past the effective periods of past insecticide application (Table A3) than if the farmer is within the effective period of the last insecticide application (Table A2).

Table A4 reports the non-parametrically estimated transition density for the SWD population levels for the mid-season (i.e., weeks $t \in (5, \frac{2}{3}T]$), conditional on spraying (i.e., $s_t = 1$), and conditional on precipitation $precip_t$. Table A4a conditions on precipitation $precip_t$ in week t being low; while Table A4b conditions on precipitation $precip_t$ in week t being high. The transition density gives the distribution of SWD population levels next week ($y_{larva,t+1}, y_{adult,t+1}$) conditional on the SWD population levels this week ($y_{larva,t}, y_{adult,t}$), conditional on spraying in week t (i.e., $s_t = 1$), conditional on precipitation $precip_t$ in week t , and conditional on week $t \in (5, \frac{2}{3}T]$ (i.e., the mid-season). By not also conditioning on int_t when $s_t = 1$ (spray), we are assuming that if we spray (i.e., $s_t = 1$), it does not make any difference to the transition

density what the interval since the last spray int_t is; this is a reasonable assumption since once we spray it no longer matters whether the grower is within the effective period of any previous spray. By not also conditioning on temperature $tmax_t$ when $s_t = 1$ (spray), we are assuming that if we spray (i.e., $s_t = 1$), it does not make any difference to the transition density what the temperature $tmax_t$ is; this is a reasonable assumption since the reason temperature $tmax_t$ matters to SWD transition densities is that high temperatures may increase SWD mortality, and since spraying itself may increase SWD mortality.

Comparing our estimated transition densities in Table A4 (when the grower sprays) with the results from Table A3 (when the grower does not spray), we find that, during the mid-season, if the grower either has never sprayed yet or is past the effective periods of past insecticide application, spraying diminishes the probability of SWD getting worse, especially when precipitation is low (Table A4a). Our finding (from comparing Table A4a and Table A4b) that spraying is more likely to diminish the probability of SWD getting worse when precipitation is low is consistent with studies that show that higher rainfall reduces spraying efficacy because rain washes the insecticide away from the crop surface (Gautam et al. 2016; Van Timmeren et al. 2017).

A.3.3. Late season

Table A5 reports the estimated transition density for the late season (i.e., weeks $t > \frac{2}{3}T$), and conditional on either being within the effective period of the last insecticide application (i.e., $int_t = 1$) or spraying this week t (i.e., $s_t = 1$). By conditioning on either $int_t = 1$ (i.e., the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective) or $s_t = 1$ (spray), we are assuming that, in the late season, spraying this week has the same effect as having sprayed last week. By not also conditioning on either int_t when $s_t = 1$ (spray), we are assuming that if we spray, it does not make any difference to the transition density what the interval since the last spray was; this is a reasonable assumption since once we spray it no longer matters whether the grower is within the effective period of any previous spray. We do not distinguish between high and low temperature for the late season since we do not see any observations of temperature in week t being high (i.e., maximum temperature that week exceeds 27 degrees Celsius) when it is the late season, $int_t = 1$, and $s_t = 0$ (not spray). Thus, we assume when $int_t = 1$ (i.e., the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective) or $s_t = 1$ (spray), and it

is the late season, the SWD mortality rate does not depend on temperature this is a reasonable assumption since the reason temperature $tmax_t$ matters to SWD transition densities is that high temperatures may increase SWD mortality, and since spraying itself may increase SWD mortality. We do not distinguish between high and low precipitation for the late season since we do not see any observations of precipitation in week t being high (i.e., weekly accumulated precipitation exceeding 2 mm) when it is the late season (i.e., weeks $t > \frac{2}{3}T$), and $s_t = 1$ (spray). Thus, we assume when $int_t = 1$ (i.e., the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective) or $s_t = 1$ (spray), and it is the late season, higher rainfall after spraying does not wash the insecticide away and make spraying less effective.

Our estimated transition density in Table A5 shows that, during the late season, if the grower either is within the effective period of the last insecticide application (i.e., $int_t = 1$) or sprays this week (i.e., $s_t = 1$), then the SWD population levels will not increase the following week.

Table A6 reports the non-parametrically estimated transition density for the SWD population levels for the late season (i.e., weeks $t > \frac{2}{3}T$), conditional on $int_t \geq 2$ (i.e., either the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective or the grower has not yet sprayed this season prior to week t), conditional on $s_t = 0$ (i.e., not spraying), and conditional on temperature $tmax_t$. Table A6a conditions on temperature $tmax_t$ in week t being low; while Table A6b conditions on temperature $tmax_t$ in week t being high.^{18,19}

¹⁸ For rows highlighted in light grey corresponding to $(y_{larva,t}, y_{adult,t})$ tuples that were never observed in the data during weeks $t > \frac{2}{3}T$ of the season (i.e., the late season) when $int_t \geq 2$ (i.e., either the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective or the grower has not yet sprayed this season prior to week t) and $s_t = 0$ (i.e., not spraying), we assume that the SWD population levels next week $(y_{larva,t+1}, y_{adult,t+1})$ are the same as the SWD population levels this week $(y_{larva,t}, y_{adult,t})$. Thus, when there is missing data, we assume that even when the grower does not spray, the SWD population levels will stay the same for those SWD population levels that are missing when the grower does not spray.

¹⁹ For tuples of SWD population levels this week $(y_{larva,t}, y_{adult,t})$ that are observed in the data during the late season (weeks $t > \frac{2}{3}T$) during weeks when $int_t \geq 2$ (i.e., either the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective or the grower has not yet sprayed this season prior to week t) and $s_t = 0$ (i.e., not spraying), but not during weeks t when the temperature is high (i.e., maximum temperature that week exceeds 27 degrees Celsius), we fill in the row of that transition density for that $(y_{larva,t}, y_{adult,t})$ tuple (highlighted in light blue in Table A6b) with the respective row from the transition density for the late season (weeks $t > \frac{2}{3}T$) during weeks when $int_t \geq 2$ and $s_t = 0$, conditional on temperature in week t being low (i.e.,

Comparing the results from Table A6a and Table A6b, we find that when the temperature is very high, this makes it more likely for SWD populations to decline, which is consistent with studies that show that the SWD mortality rate increases dramatically if the temperature exceeds 27 degrees Celsius (Tochen et al. 2014).

A.3. Results of SWD Transition Densities

The results of the transition densities for the SWD population levels that we estimate non-parametrically and use for our numerical model and structural econometric models are consistent with models of SWD dynamics. Our estimated transition densities show that, during the mid-season, if the farmer either has never sprayed yet or is past the effective periods of past insecticide application, there is a higher probability that the SWD larva population will increase if the farmer does not spray. Also during the mid-season, spraying diminishes the probability of SWD getting worse, especially when precipitation is low, which is consistent with studies that show that higher rainfall reduces spraying efficacy because rain washes the insecticide away from the crop surface (Gautam et al. 2016; Van Timmeren et al. 2017). We also find that when the temperature is very high, this makes it more likely for SWD populations to decline, which is consistent with studies that show that the SWD mortality rate increases dramatically if the temperature exceeds 27 degrees Celsius (Tochen et al. 2014).

As seen in Table A1, SWD population levels are relatively low during the early part of the season (weeks $t \leq 5$). During the early season (weeks $t \leq 5$), no adult SWD was observed in the data for any observations (i.e., $y_{adult,t} = 1$ for all weeks $t \leq 5$), and the observed larva SWD levels in the data were only either none ($y_{larva,t} = 1$) or in the lowest tier ($y_{larva,t} = 2$). Moreover, when temperature $tmax_t$ in week $t \leq 5$ is low (Table A1a), the SWD levels remain the same the following week; but when temperature $tmax_t$ in week $t \leq 5$ is high (Table A1b) and the observed larva SWD level is in the lowest tier ($y_{larva,t} = 2$), there is a roughly 50 percent probability that the observed larva SWD level will decline to none ($y_{larva,t} = 1$) the following week. Thus, when the temperature $tmax_t$ in week $t \leq 5$ is high (i.e., maximum temperature that week exceeds 27 degrees Celsius), this makes it more likely for SWD populations to decline, which

maximum temperature that week is less than or equal to 27 degrees Celsius) in Table A6a. Thus, for these missing observations, we make a conservative assumption that high temperature does not have any effect on SWD mortality.

is consistent with studies that show that the SWD mortality rate increases dramatically if the temperature exceeds 27 degrees Celsius (Tochen et al. 2014).

Our estimated transition density in Table A2 shows that, during the mid-season, if the farmer is within the effective period of the last insecticide application and does not spray this week, then the SWD larva population will remain at the same levels following next week.

Comparing the results of Table A3a and Table A3b. we find that, during the mid-season, when the temperature is very high (Table A3b), this makes it less likely for SWD populations to increase, which is consistent with studies that show that the SWD mortality rate increases dramatically if the temperature exceeds 27 degrees Celsius (Tochen et al. 2014).

Comparing the results of Table A3 with those of Table A2, we find that, during the mid-season, if the farmer does not spray, there is a higher probability that the SWD larva population will increase if the farmer either has never sprayed yet or is past the effective periods of past insecticide application (Table A3) than if the farmer is within the effective period of the last insecticide application (Table A2). Comparing our estimated transition densities in Table A4 (when the grower sprays) with the results from Table A3 (when the grower does not spray), we find that, during the mid-season, if the grower either has never sprayed yet or is past the effective periods of past insecticide application, spraying diminishes the probability of SWD getting worse, especially when precipitation is low (Table A4a). Our finding (from comparing Table A4.a and Table A4b) that spraying is more likely to diminish the probability of SWD getting worse when precipitation is low is consistent with studies that show that higher rainfall reduces spraying efficacy because rain washes the insecticide away from the crop surface (Gautam et al. 2016; Van Timmeren et al. 2017).

Our estimated transition density in Table A5 shows that, during the late season, if the grower either is within the effective period of the last insecticide application (i.e., $int_t = 1$) or sprays this week (i.e., $s_t = 1$), then the SWD population levels will not increase the following week.

Comparing the results from Table A6a and Table A6b, we find that when the temperature is very high, this makes it more likely for SWD populations to decline, which is consistent with studies that show that the SWD mortality rate increases dramatically if the temperature exceeds 27 degrees Celsius (Tochen et al. 2014).

Table A1. Transition Densities for SWD Population Levels for the Early Season

(a) Transition density for SWD population levels for early season conditional on temperature in week t being low

SWD in week $t + 1$		$y_{larva,t+1}$	1	2	3	4	5	6	1	2	3	4	5	6
$t + 1$		$y_{adult,t+1}$	1	1	1	1	1	1	2	2	2	2	2	2
SWD in week t														
$y_{larva,t}$	$y_{adult,t}$													
1	1		1											
2	1			1										
3	1													
4	1													
5	1													
6	1													
1	2													
2	2													
3	2													
4	2													
5	2													
6	2													

Notes: Table reports the non-parametrically estimated transition density for the early season (weeks $t \leq 5$), conditional on temperature in week t being low. Each cell in the table gives the probability that SWD population levels next week ($y_{larva,t+1}$, $y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}$, $y_{adult,t}$) being the values in the row, conditional on temperature $tmax_t$ in week t being low (i.e., maximum temperature that week is less than or equal to 27 degrees Celsius), and conditional on week $t \leq 5$. For ($y_{larva,t}$, $y_{adult,t}$) tuples that were never observed in the data during weeks $t \leq 5$ of the season when temperature is low, we highlight the row of that transition density for that ($y_{larva,t}$, $y_{adult,t}$) tuple in light grey.

(b) Transition density for SWD population levels for early part of season conditional on temperature in week t being high

SWD in week $t + 1$ $y_{larva,t+1}$ $y_{adult,t+1}$		1	2	3	4	5	6	1	2	3	4	5	6
		1	1	1	1	1	1	2	2	2	2	2	2
SWD in week t $y_{larva,t}$ $y_{adult,t}$													
1	1	0.99	0.01										
2	1	0.50	0.50										
3	1												
4	1												
5	1												
6	1												
1	2												
2	2												
3	2												
4	2												
5	2												
6	2												

Notes: Table reports the non-parametrically estimated transition density for the early part of the season (weeks $t \leq 5$), conditional on temperature in week t being high. Each cell in the table gives the probability that SWD population levels next week ($y_{larva,t+1}$, $y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}$, $y_{adult,t}$) being the values in the row, conditional on temperature $tmax_t$ in week t being high (i.e., maximum temperature that week exceeds 27 degrees Celsius), and conditional on week $t \leq 5$. For ($y_{larva,t}$, $y_{adult,t}$) tuples that were never observed in the data during weeks $t \leq 5$ of the season when temperature is high, we highlight the row of that transition density for that ($y_{larva,t}$, $y_{adult,t}$) tuple in light grey.

Table A2. Transition Densities for SWD Population Levels for the Mid-Season Conditional on $int_t = 1$ and $s_t = 0$

(a) Transition density for mid-season conditional on being within the effective period of the last insecticide application ($int_t = 1$) and conditional on not spraying this week t ($s_t = 0$)

SWD in week $t + 1$		$y_{larva,t+1}$	1	2	3	4	5	6	1	2	3	4	5	6
		$y_{adult,t+1}$	1	1	1	1	1	1	2	2	2	2	2	2
SWD in week t														
$y_{larva,t}$	$y_{adult,t}$													
1	1	1												
2	1													
3	1													
4	1													
5	1													
6	1													
1	2	1												
2	2													
3	2													
4	2													
5	2													
6	2													

Notes: Table reports the estimated transition density for the mid-season (i.e., weeks $t \in (5, \frac{2}{3}T]$), conditional on being within the effective period of the last insecticide application (i.e., $int_t = 1$) and conditional on not spraying (i.e., $s_t = 0$). Each cell in the table gives the probability that SWD population levels next week ($y_{larva,t+1}, y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}, y_{adult,t}$) being the values in the row, conditional on $int_t = 1$ (i.e., the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective), conditional on $s_t = 0$ (i.e., not spraying), and conditional on week $t \in (5, \frac{2}{3}T]$ (i.e., the mid-season). For $(y_{larva,t}, y_{adult,t})$ tuples that were never observed in the data during weeks $t \in (5, \frac{2}{3}T]$ of the season when $int_t = 1$ and $s_t = 0$, we highlight the row of that transition density for that $(y_{larva,t}, y_{adult,t})$ tuple in light grey.

Table A3. Transition Densities for SWD Population Levels for the Mid-Season Conditional on $int_t \geq 2$ and $s_t = 0$

(a) Transition density for mid-season conditional on the grower either has never sprayed yet this season or is past the effective period of the last insecticide (i.e., $int_t \geq 2$), conditional on not spraying this week t (i.e., $s_t = 0$), and conditional on temperature in week t being low

SWD in week $t+1$	$y_{larva,t+1}$	1	2	3	4	5	6	1	2	3	4	5	6
$y_{adult,t+1}$		1	1	1	1	1	1	2	2	2	2	2	2
SWD in week t	$y_{larva,t}$	$y_{adult,t}$											
1	1		0.785	0.004	0.007			0.20	0.004				
2	1		0.50	0.25				0.25					
3	1		0.67		0.33								
4	1												
5	1												
6	1												
1	2		0.05					0.74	0.08	0.11	0.02		
2	2			0.11					0.11	0.44	0.33		
3	2									0.12	0.88		
4	2				0.2						0.6	0.2	
5	2												
6	2												

Notes: Table reports the estimated transition density for the mid-season (i.e., weeks $t \in (5, \frac{2}{3}T]$), conditional on the grower either has never sprayed yet this season or is past the effective period of the last insecticide (i.e., $int_t \geq 2$), conditional on not spraying (i.e., $s_t = 0$), and conditional on temperature in week t being low (i.e., $tmax_t = 0$). Each cell in the table gives the probability that SWD population levels next week ($y_{larva,t+1}, y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}, y_{adult,t}$) being the values in the row, conditional on $int_t \geq 2$ (i.e., either the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective or the grower has not yet sprayed this season prior to week t), conditional on $s_t = 0$ (i.e., not spraying), conditional on temperature in week t being low (i.e., maximum temperature that week is less than or equal to 27 degrees Celsius), and conditional on week $t \in (5, \frac{2}{3}T]$ (i.e., the mid-season). For $(y_{larva,t}, y_{adult,t})$ tuples that were never observed in the data during weeks $t \in (5, \frac{2}{3}T]$ of the season when $int_t \geq 2$ and $s_t = 0$, we highlight the row of that transition density for that $(y_{larva,t}, y_{adult,t})$ tuple in light grey.

(b) Transition density for mid-season conditional on the grower either has never sprayed yet this season or is past the effective period of the last insecticide (i.e., $int_t \geq 2$), conditional on not spraying this week t (i.e., $s_t = 0$), and conditional on temperature in week t being high

SWD in week $t + 1$	$y_{larva,t+1}$	1	2	3	4	5	6	1	2	3	4	5	6
$y_{adult,t+1}$		1	1	1	1	1	1	2	2	2	2	2	2
SWD in week t	$y_{larva,t}$	$y_{adult,t}$											
1	1		0.87					0.13					
2	1												
3	1												
4	1												
5	1												
6	1												
1	2		0.045					0.864	0.045	0.045			
2	2								1				
3	2										1		
4	2												
5	2												
6	2												2

Notes: Table reports the estimated transition density for the mid-season (i.e., weeks $t \in (5, \frac{2}{3}T]$), conditional on the grower either has never sprayed yet this season or is past the effective period of the last insecticide (i.e., $int_t \geq 2$), conditional on not spraying (i.e., $s_t = 0$), and conditional on temperature in week t being high (i.e., $tmax_t = 1$). Each cell in the table gives the probability that SWD population levels next week ($y_{larva,t+1}, y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}, y_{adult,t}$) being the values in the row, conditional on $int_t \geq 2$ (i.e., either the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective or the grower has not yet sprayed this season prior to week t), conditional on $s_t = 0$ (i.e., not spraying), conditional on temperature in week t being high (i.e., maximum temperature that week exceeds 27 degrees Celsius), and conditional on week $t \in (5, \frac{2}{3}T]$ (i.e., the mid-season). For $(y_{larva,t}, y_{adult,t})$ tuples that were never observed in the data during weeks $t \in (5, \frac{2}{3}T]$ of the season when $int_t \geq 2$ and $s_t = 0$, we highlight the row of that transition density for that $(y_{larva,t}, y_{adult,t})$ tuple in light grey. For tuples of SWD population levels this week ($y_{larva,t}, y_{adult,t}$) that are observed in the data during the mid-season (weeks $t \in (5, \frac{2}{3}T]$) during weeks when $int_t \geq 2$ and $s_t = 0$, but not during weeks t when the temperature is high, we fill in the row of that transition density for that $(y_{larva,t}, y_{adult,t})$ tuple (highlighted in light blue) with the respective row from the transition density for the mid-season (weeks $t \in (5, \frac{2}{3}T]$) during weeks when $int_t \geq 2$ and $s_t = 0$, conditional on temperature in week t being low in Table A3a.

Table A4. Transition Densities for SWD Population Levels for the Mid-Season Conditional on Spraying ($s_t = 1$)

(a) Transition density for mid-season conditional on spraying this week t , and conditional on precipitation in week t being low

SWD in week $t + 1$	$y_{larva,t+1}$ $y_{adult,t+1}$	1	2	3	4	5	6	1	2	3	4	5	6
SWD in week t $y_{larva,t}$ $y_{adult,t}$		1	1	1	1	1	1	2	2	2	2	2	2
1	1	0.67						0.33					
2	1												
3	1												
4	1												
5	1												
6	1												
1	2							1					
2	2												
3	2												
4	2												
5	2												
6	2												

Notes: Table reports the estimated transition density for the mid-season (i.e., weeks $t \in (5, \frac{2}{3}T]$), conditional on spraying ($s_t = 1$) and conditional on precipitation in week t being low ($precip_t = 0$). Each cell in the table gives the probability that SWD population levels next week ($y_{larva,t+1}$, $y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}$, $y_{adult,t}$) being the values in the row, conditional on $s_t = 1$ (spray), conditional on precipitation in week t being low (i.e., weekly accumulated precipitation that week is less than or equal to 2 mm), and conditional on week $t \in (5, \frac{2}{3}T]$ (i.e., the mid-season). For $(y_{larva,t}, y_{adult,t})$ tuples that were never observed in the data during weeks $t \in (5, \frac{2}{3}T]$ of the season when $s_t = 1$, we highlight the row of that transition density for that $(y_{larva,t}, y_{adult,t})$ tuple in light grey.

(b) Transition density for mid-season conditional on spraying this week t , and conditional on precipitation in week t being high

SWD in week $t + 1$ $y_{larva,t+1}$ $y_{adult,t+1}$		1	2	3	4	5	6	1	2	3	4	5	6
		1	1	1	1	1	1	2	2	2	2	2	2
SWD in week t $y_{larva,t}$ $y_{adult,t}$													
1	1	0.88						0.12					
2	1												
3	1												
4	1												
5	1												
6	1												
1	2							0.80	0.13		0.07		
2	2												
3	2												
4	2												
5	2												
6	2												

Notes: Table reports the estimated transition density for the mid-season (i.e., weeks $t \in (5, \frac{2}{3}T]$), conditional on spraying ($s_t = 1$) and conditional on precipitation in week t being high ($precip_t = 1$). Each cell in the table gives the probability that SWD population levels next week ($y_{larva,t+1}, y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}, y_{adult,t}$) being the values in the row, conditional on $s_t = 1$ (spray), conditional on precipitation in week t being high (i.e., weekly accumulated precipitation that week exceeds 2 mm), and conditional on week $t \in (5, \frac{2}{3}T]$ (i.e., the mid-season). For $(y_{larva,t}, y_{adult,t})$ tuples that were never observed in the data during weeks $t \in (5, \frac{2}{3}T]$ of the season when $s_t = 1$, we highlight the row of that transition density for that $(y_{larva,t}, y_{adult,t})$ tuple in light grey.

Table A5. Transition Densities for SWD Population Levels for the Late Season Conditional on either $int_t = 1$ or $s_t = 1$

(a) Transition density for late season conditional either on being within the effective period of the last insecticide application ($int_t = 1$), or on spraying this week t ($s_t = 1$)

SWD in week $t + 1$	$y_{larva,t+1}$	1	2	3	4	5	6	1	2	3	4	5	6
$y_{adult,t+1}$		1	1	1	1	1	1	2	2	2	2	2	2
SWD in week t	$y_{larva,t}$	$y_{adult,t}$											
1	1	1											
2	1	1											
3	1	1											
4	1	1											
5	1	1											
6	1	1											
1	2	2											
2	2	2											
3	2	2											
4	2	2											
5	2	2											
6	2	2											

Notes: Table reports the estimated transition density for the late season (i.e., weeks $t > \frac{2}{3}T$), and conditional on either being within the effective period of the last insecticide application (i.e., $int_t = 1$) or spraying this week t (i.e., $s_t = 1$). Each cell in the table gives the probability that SWD population levels next week ($y_{larva,t+1}$, $y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}$, $y_{adult,t}$) being the values in the row, conditional on either $int_t = 1$ (i.e., the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective) or $s_t = 1$ (spray), and conditional on week $t > \frac{2}{3}T$ (i.e., the late season). For ($y_{larva,t}$, $y_{adult,t}$) tuples that were never observed in the data during weeks $t > \frac{2}{3}T$ of the season when either $int_t = 1$ or $s_t = 1$, and conditional on week $t > \frac{2}{3}T$, we highlight the row of that transition density for that ($y_{larva,t}$, $y_{adult,t}$) tuple in light grey.

Table A6. Transition Densities for SWD Population Levels for the Late Season Conditional on $int_t \geq 2$ and $s_t = 0$

(a) Transition density for late season conditional on the grower either has never sprayed yet this season or is past the effective period of the last insecticide (i.e., $int_t \geq 2$), conditional on not spraying this week t (i.e., $s_t = 0$), and conditional on temperature in week t being low

SWD in week $t + 1$	$y_{larva,t+1}$	1	2	3	4	5	6	1	2	3	4	5	6
$y_{adult,t+1}$		1	1	1	1	1	1	2	2	2	2	2	2
SWD in week t	$y_{larva,t}$	$y_{adult,t}$											
1	1		0.98					0.02					
2	1												
3	1			1									
4	1				0.60	0.40							
5	1					0.75	0.25						
6	1						1						
1	2		0.04					0.86	0.07	0.02	0.01		
2	2								0.47	0.37	0.16		
3	2				0.03				0.03	0.42	0.52		
4	2				0.03				0.01	0.04	0.70	0.22	
5	2						0.04				0.04	0.58	0.33
6	2											0.11	0.89

Notes: Table reports the estimated transition density for the late season (i.e., weeks $t > \frac{2}{3}T$), conditional on the grower either has never sprayed yet this season or is past the effective period of the last insecticide (i.e., $int_t \geq 2$), conditional on not spraying (i.e., $s_t = 0$), and conditional on temperature in week t being low (i.e., $tmax_t = 0$). Each cell in the table gives the probability that SWD population levels next week ($y_{larva,t+1}, y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}, y_{adult,t}$) being the values in the row, conditional on $int_t \geq 2$ (i.e., either the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective or the grower has not yet sprayed this season prior to week t), conditional on $s_t = 0$ (i.e., not spraying), conditional on temperature in week t being low (i.e., maximum temperature that week is less than or equal to 27 degrees Celsius), and conditional on week $t > \frac{2}{3}T$ (i.e., the late season). For $(y_{larva,t}, y_{adult,t})$ tuples that were never observed in the data during weeks $t > \frac{2}{3}T$ of the season when $int_t \geq 2$ and $s_t = 0$, we highlight the row of that transition density for that $(y_{larva,t}, y_{adult,t})$ tuple in light grey.

(b) Transition density for late season conditional on the grower either has never sprayed yet this season or is past the effective period of the last insecticide (i.e., $int_t \geq 2$), conditional on not spraying this week t (i.e., $s_t = 0$), and conditional on temperature in week t being high

SWD in week $t+1$		1	2	3	4	5	6	1	2	3	4	5	6
$y_{larva,t+1}$ $y_{adult,t+1}$		1	1	1	1	1	1	2	2	2	2	2	2
SWD in week t													
$y_{larva,t}$	$y_{adult,t}$												
1	1	1											
2	1												
3	1												
4	1												
5	1												
6	1												
1	2							0.94	0.06				
2	2	0.33									0.67		
3	2										1		
4	2										0.86	0.14	
5	2												
6	2	1											

Notes: Table reports the estimated transition density for the late season (i.e., weeks $t > \frac{2}{3}T$), conditional on the grower either has never sprayed yet this season or is past the effective period of the last insecticide (i.e., $int_t \geq 2$), conditional on not spraying (i.e., $s_t = 0$), and conditional on temperature in week t being high (i.e., $tmax_t = 1$). Each cell in the table gives the probability that SWD population levels next week ($y_{larva,t+1}, y_{adult,t+1}$) are the values given in the column, conditional on the SWD population levels this week ($y_{larva,t}, y_{adult,t}$) being the values in the row, conditional on $int_t \geq 2$ (i.e., either the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective or the grower has not yet sprayed this season prior to week t), conditional on $s_t = 0$ (i.e., not spraying), conditional on temperature in week t being high (i.e., maximum temperature that week exceeds 27 degrees Celsius), and conditional on week $t > \frac{2}{3}T$ (i.e., the late season). For ($y_{larva,t}, y_{adult,t}$) tuples that were never observed in the data during weeks $t > \frac{2}{3}T$ of the season when $int_t \geq 2$ and $s_t = 0$, we highlight the row of that transition density for that ($y_{larva,t}, y_{adult,t}$) tuple in light grey. For tuples of SWD population levels this week ($y_{larva,t}, y_{adult,t}$) that are observed in the data during the late season (weeks $t > \frac{2}{3}T$) during weeks when $int_t \geq 2$ and $s_t = 0$, but not during weeks t when the temperature is high, we fill in the row of that transition density for that ($y_{larva,t}, y_{adult,t}$) tuple (highlighted in light blue) with the respective row from the transition density for the late season (weeks $t > \frac{2}{3}T$) during weeks when $int_t \geq 2$ and $s_t = 0$, conditional on temperature in week t being low in Table A6a.

Appendix B.

Supplemental Tables and Figures

Table B1. Policy Function

(a) Policy function for subset of states in which precipitation and temperature are both low

State variables			Policy function for week:																	
y_{larva}	y_{adult}	int	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
1	1	1																		
2	1	1																		
3	1	1																		
4	1	1																		
5	1	1																		
6	1	1																		
1	2	1																		
2	2	1																		
3	2	1																		
4	2	1																		
5	2	1																		
6	2	1																		
1	1	2																		
2	1	2																		
3	1	2																		
4	1	2																		
5	1	2																		
6	1	2																		
1	2	2																		
2	2	2																		
3	2	2																		
4	2	2																		
5	2	2																		
6	2	2																		
1	1	3																		
2	1	3																		
3	1	3																		
4	1	3																		
5	1	3																		
6	1	3																		
1	2	3																		
2	2	3																		
3	2	3																		
4	2	3																		
5	2	3																		
6	2	3																		

Notes: Table reports the policy function for each week of the season for the subset of state tuples in which precipitation and temperature are both low. For each week, the policy function for that week gives the optimal choice of the action for that week as a function of the state variables; states for which the optimal choice is to spray are indicated in light red, and states for which the optimal choice is to harvest are indicated in light green. The state variables measure SWD population levels, weather variables, and the interval of time since the last insecticide application. The SWD population levels are indicated by two state variables: observed larva SWD y_{larva} , which is discretized to six levels; and observed adult SWD y_{adult} , which is a binary variable indicating whether adult SWD was observed that week (1=no, 2=yes). The state variables also include a discretized variable int for the interval of time since the last insecticide application, which can take the value of 1 if the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective; 2 if the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective; or 3 if the grower has not yet sprayed this season prior to that week. The state variables additionally include a dummy variable for weekly accumulated precipitation exceeding 2 mm and a dummy variable for weekly maximum temperature exceeding 27 degrees Celsius; this table reports the policy function for the subset of states for which precipitation is low (i.e., accumulated precipitation that week is less than or equal to 2 mm) and temperature is low (i.e., maximum temperature that week is less than or equal to 27 degrees Celsius).

(b) Policy function for subset of states in which precipitation is low and temperature is high

State variables			Policy function for week:																	
y_{larva}	y_{adult}	int	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
1	1	1																		
2	1	1																		
3	1	1																		
4	1	1																		
5	1	1																		
6	1	1																		
1	2	1																		
2	2	1																		
3	2	1																		
4	2	1																		
5	2	1																		
6	2	1																		
1	1	2																		
2	1	2																		
3	1	2																		
4	1	2																		
5	1	2																		
6	1	2																		
1	2	2																		
2	2	2																		
3	2	2																		
4	2	2																		
5	2	2																		
6	2	2																		
1	1	3																		
2	1	3																		
3	1	3																		
4	1	3																		
5	1	3																		
6	1	3																		
1	2	3																		
2	2	3																		
3	2	3																		
4	2	3																		
5	2	3																		
6	2	3																		

Notes: Table reports the policy function for each week of the season for the subset of state tuples in which precipitation is low and temperature is high. For each week, the policy function for that week gives the optimal choice of the action for that week as a function of the state variables; states for which the optimal choice is to spray are indicated in light red, and states for which the optimal choice is to harvest are indicated in light green. The state variables measure SWD population levels, weather variables, and the interval of time since the last insecticide application. The SWD population levels are indicated by two state variables: observed larva SWD y_{larva} , which is discretized to six levels; and observed adult SWD y_{adult} , which is a binary variable indicating whether adult SWD was observed that week (1=no, 2=yes). The state variables also include a discretized variable int for the interval of time since the last insecticide application, which can take the value of 1 if the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective; 2 if the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective; or 3 if the grower has not yet sprayed this season prior to that week. The state variables additionally include a dummy variable for weekly accumulated precipitation exceeding 2 mm and a dummy variable for weekly maximum temperature exceeding 27 degrees Celsius; this table reports the policy function for the subset of states for which precipitation is low (i.e., accumulated precipitation that week is less than or equal to 2 mm) and temperature is high (i.e., maximum temperature that week exceeds 27 degrees Celsius).

(c) Policy function for subset of states in which precipitation is high and temperature is low

State variables			Policy function for week:																	
y_{larva}	y_{adult}	int	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
1	1	1																		
2	1	1																		
3	1	1																		
4	1	1																		
5	1	1																		
6	1	1																		
1	2	1																		
2	2	1																		
3	2	1																		
4	2	1																		
5	2	1																		
6	2	1																		
1	1	2																		
2	1	2																		
3	1	2																		
4	1	2																		
5	1	2																		
6	1	2																		
1	2	2																		
2	2	2																		
3	2	2																		
4	2	2																		
5	2	2																		
6	2	2																		
1	1	3																		
2	1	3																		
3	1	3																		
4	1	3																		
5	1	3																		
6	1	3																		
1	2	3																		
2	2	3																		
3	2	3																		
4	2	3																		
5	2	3																		
6	2	3																		

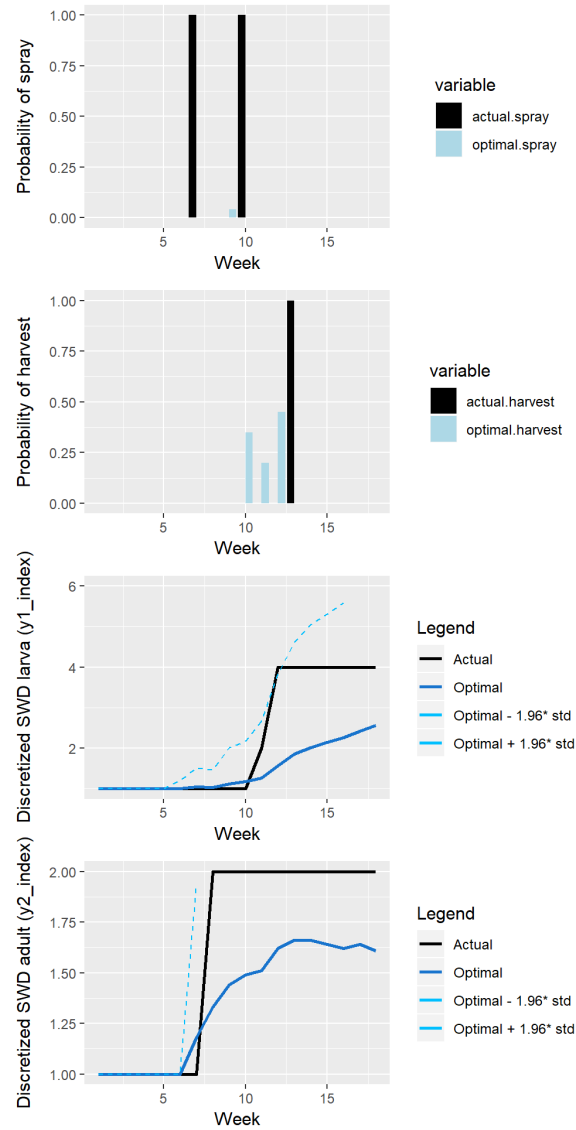
Notes: Table reports the policy function for each week of the season for the subset of state tuples in which precipitation is high and temperature is low. For each week, the policy function for that week gives the optimal choice of the action for that week as a function of the state variables; states for which the optimal choice is to spray are indicated in light red, and states for which the optimal choice is to harvest are indicated in light green. The state variables measure SWD population levels, weather variables, and the interval of time since the last insecticide application. The SWD population levels are indicated by two state variables: observed larva SWD y_{larva} , which is discretized to six levels; and observed adult SWD y_{adult} , which is a binary variable indicating whether adult SWD was observed that week (1=no, 2=yes). The state variables also include a discretized variable int for the interval of time since the last insecticide application, which can take the value of 1 if the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective; 2 if the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective; or 3 if the grower has not yet sprayed this season prior to that week. The state variables additionally include a dummy variable for weekly accumulated precipitation exceeding 2 mm and a dummy variable for weekly maximum temperature exceeding 27 degrees Celsius; this table reports the policy function for the subset of states for which precipitation is high (i.e., accumulated precipitation that week exceeds 2 mm) and temperature is low (i.e., maximum temperature that week is less than or equal to 27 degrees Celsius).

(d) Policy function for subset of states in which precipitation and temperature are both high

State variables			Policy function for week:																	
y_{larva}	y_{adult}	int	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
1	1	1																		
2	1	1																		
3	1	1																		
4	1	1																		
5	1	1																		
6	1	1																		
1	2	1																		
2	2	1																		
3	2	1																		
4	2	1																		
5	2	1																		
6	2	1																		
1	1	2																		
2	1	2																		
3	1	2																		
4	1	2																		
5	1	2																		
6	1	2																		
1	2	2																		
2	2	2																		
3	2	2																		
4	2	2																		
5	2	2																		
6	2	2																		
1	1	3																		
2	1	3																		
3	1	3																		
4	1	3																		
5	1	3																		
6	1	3																		
1	2	3																		
2	2	3																		
3	2	3																		
4	2	3																		
5	2	3																		
6	2	3																		

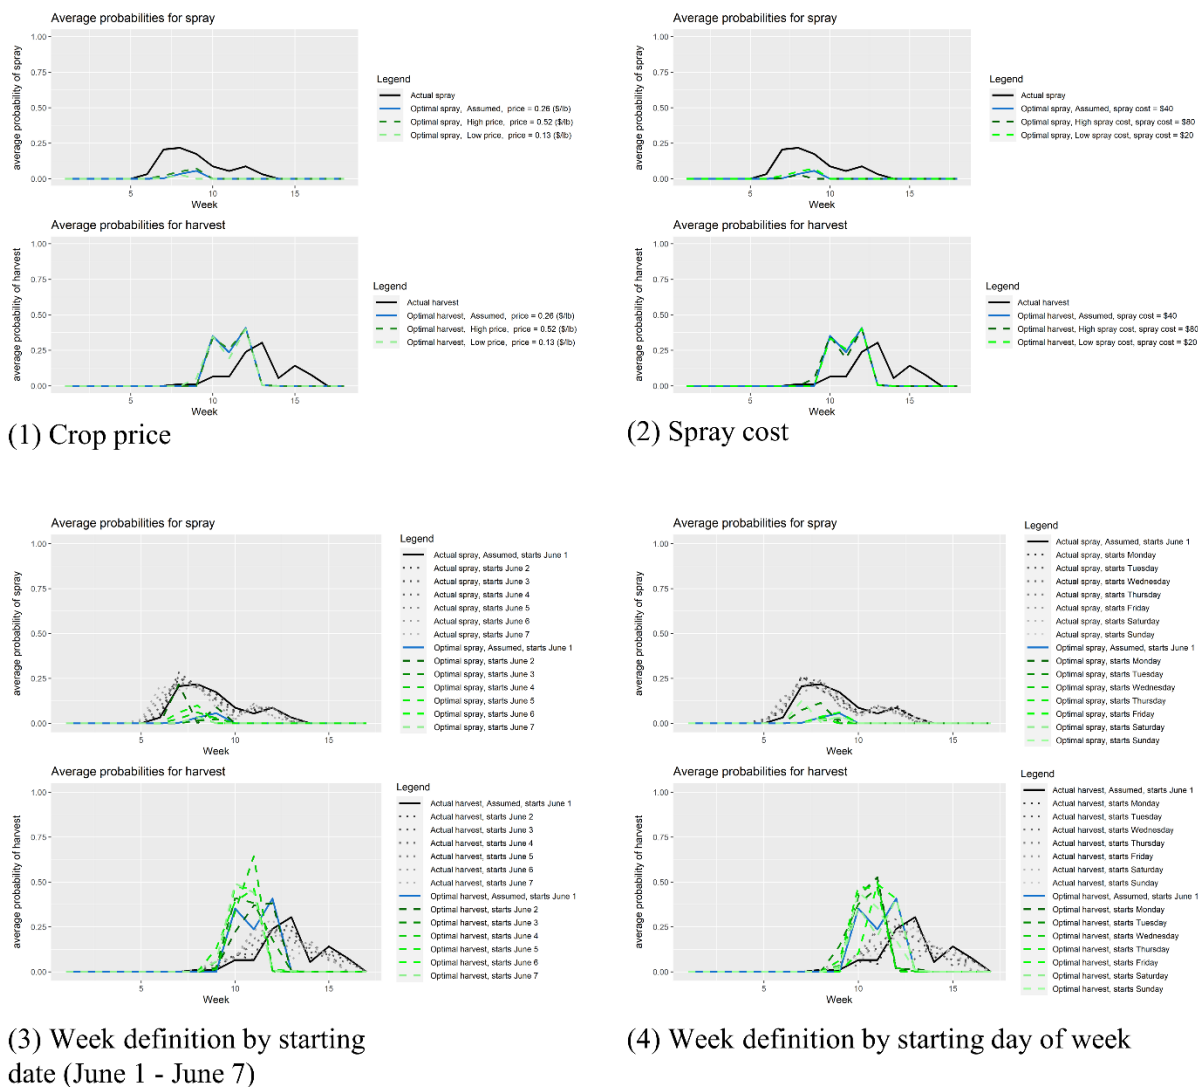
Notes: Table reports the policy function for each week of the season for the subset of state tuples in which precipitation and temperature are both high. For each week, the policy function for that week gives the optimal choice of the action for that week as a function of the state variables; states for which the optimal choice is to spray are indicated in light red, and states for which the optimal choice is to harvest are indicated in light green. The state variables measure SWD population levels, weather variables, and the interval of time since the last insecticide application. The SWD population levels are indicated by two state variables: observed larva SWD y_{larva} , which is discretized to six levels; and observed adult SWD y_{adult} , which is a binary variable indicating whether adult SWD was observed that week (1=no, 2=yes). The state variables also include a discretized variable int for the interval of time since the last insecticide application, which can take the value of 1 if the number of weeks since the last spray is less than or equal to the maximum number of weeks for which the insecticide is effective; 2 if the number of weeks since the last spray exceeds the maximum number of weeks for which the insecticide is effective; or 3 if the grower has not yet sprayed this season prior to that week. The state variables additionally include a dummy variable for weekly accumulated precipitation exceeding 2 mm and a dummy variable for weekly maximum temperature exceeding 27 degrees Celsius; this table reports the policy function for the subset of states for which precipitation is high (i.e., accumulated precipitation that week exceeds 2 mm) and temperature is high (i.e., maximum temperature that week exceeds 27 degrees Celsius).

Figure B1. Actual vs. Optimal Strategy for Farm 90



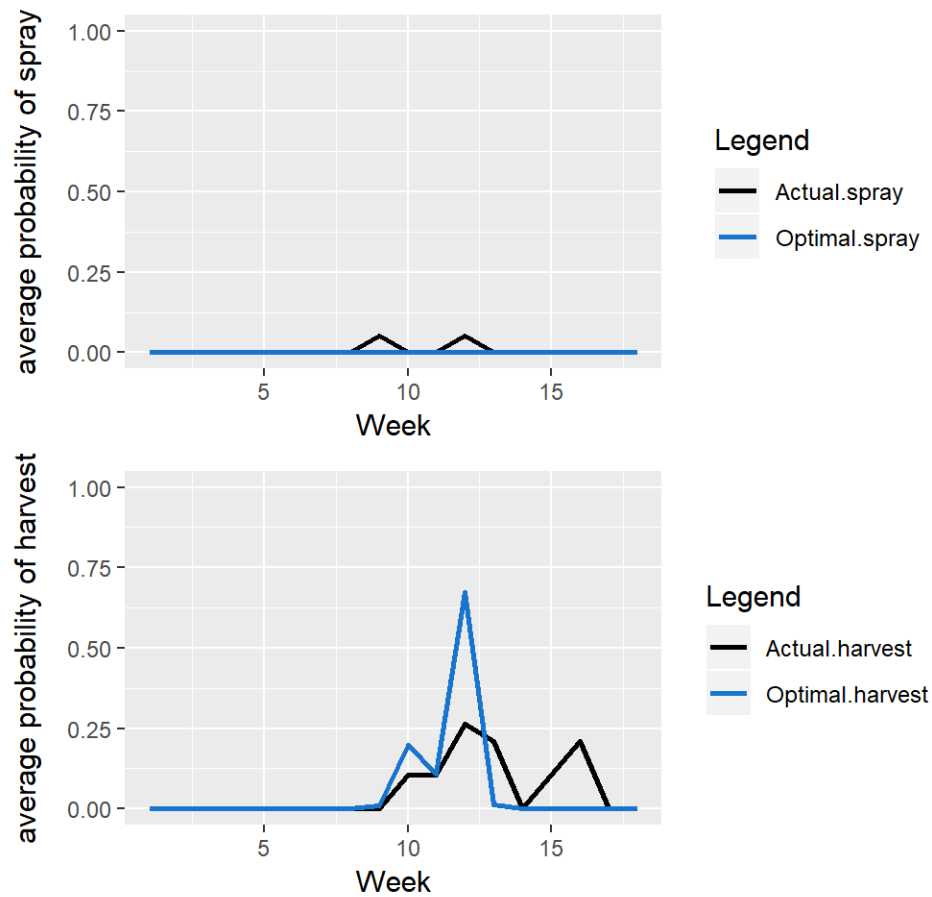
Notes: Figure compares the optimal probabilities of spraying and harvesting for Farm 90 with the actual spraying and harvesting decision of Farm 90. The optimal average probabilities of spraying and harvesting are given by averaging over 100 simulations for farm 90 using the policy function from our numerical model and the non-parametrically estimated transition densities.

Figure B2. Numerical Model Robustness: Optimal vs. actual probabilities of spraying and harvesting



Notes: We run alternative specifications of our numerical model that vary the crop price, spray cost, the start dates used to define each week, or the start days of week used to define each week. Figure compares the optimal and actual average probabilities of spraying and harvesting across these various alternative specifications. The assumed values are the assumed values of the parameters used in the numerical model and are based on economic data, expert opinion, and extension reports. The average probabilities of actual spray and actual harvest are calculated by averaging the observed actions over all 92 farms in the data set. The optimal average probabilities of spraying and harvesting are given by averaging over 100 simulations for each farm using the policy function from our numerical model and the non-parametrically estimated transition densities, and then averaging over all 92 farms.

Figure B3. Optimal vs. actual probabilities of spraying and harvesting for organic farms



Notes: In this alternative specification, we non-parametrically estimate stochastic transition densities for the SWD population using data from organic farms only, reflecting the SWD population dynamics on organic farms; and we use for the crop price the organic crop price of \$6.00/lb. Figure compares the optimal and actual average probabilities of spraying and harvesting for organic farms. The average probabilities of actual spray and actual harvest are calculated by averaging the observed actions over the 19 organic farms in the data set. The optimal average probabilities of spraying and harvesting are given by averaging over 100 simulations for each organic farm using the policy function from our numerical model and the non-parametrically estimated transition densities, and then averaging over all 19 organic farms.

Table B2a. Robustness: Structural parameter estimates with coarser bins for larva and larval infestation

	Coarser Bins
Infestation loss tier 1 (Infestation \leq 0.3%)	0.000 (0.000)
Infestation loss tier 2 (0.3% < Infestation \leq 1%)	0.000 (0.171)
Infestation loss tier 3 (1% < Infestation \leq 5%)	0.000 (0.045)
Infestation loss tier 4 (5% < Infestation)	1.000*** (0.177)
Spray cost (\$ per acre)	2,965.83*** (96.121)
# Observations	1,656
# Farms	92

Notes: In this alternative specification, coarser bins for larva and larval infestation are used. Infestation is the observed larval infestation. The infestation loss is the incremental yield loss at each tier of observed larval infestation; the cumulative value of infestation loss is the (cumulative) yield loss $loss(y_{larva})$. Bootstrapped standard errors in parentheses. Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table B2b. Robustness: Structural parameter estimates with non-organic (conventional) farms only

	Non-Organic Only
Infestation loss tier 1 (Infestation \leq 0.3%)	0.000 (0.000)
Infestation loss tier 2 (0.3% < Infestation \leq 1%)	0.000 (0.028)
Infestation loss tier 3 (1% < Infestation \leq 5%)	0.000 (0.115)
Infestation loss tier 4 (5% < Infestation \leq 10%)	0.999*** (0.240)
Infestation loss tier 5 (10% < Infestation)	0.001 (0.217)
Spray cost (\$ per acre)	2,775.31*** (92.887)
# Observations	1,314
# Farms	73

Notes: In this alternative specification, the sample is restricted to non-organic (conventional) farms only. Infestation is the observed larval infestation. The infestation loss is the incremental yield loss at each tier of observed larval infestation; the cumulative value of infestation loss is the (cumulative) yield loss $loss(y_{larva})$. Bootstrapped standard errors in parentheses. Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table B2c. Robustness: Structural parameter estimates with organic farms only

	Crop Price (\$0.26/lb)	Organic Crop Price (\$6.00/lb)
Infestation loss tier 1 (Infestation \leq 0.3%)	0.000 (0.000)	0.000 (0.000)
Infestation loss tier 2 (0.3% < Infestation \leq 1%)	0.000 (0.334)	0.000 (0.334)
Infestation loss tier 3 (1% < Infestation \leq 5%)	0.000 (0.000)	0.000 (0.000)
Infestation loss tier 4 (5% < Infestation \leq 10%)	1.000*** (0.353)	1.000*** (0.353)
Infestation loss tier 5 (10% < Infestation)	0.000 (0.117)	0.000 (0.117)
Spray cost (\$ per acre)	5,036.10 (6,141.70)	5,036.10 (6,141.70)
# Observations	342	342
# Farms	19	19

Notes: In both these alternative specifications, the sample is restricted to organic farms only. For the crop price, we use either the crop price of \$0.26/lb as before, or the organic crop price of \$6.00/lb. Infestation is the observed larval infestation. The infestation loss is the incremental yield loss at each tier of observed larval infestation; the cumulative value of infestation loss is the (cumulative) yield loss $loss(y_{larva})$. Bootstrapped standard errors in parentheses. Significance codes: *** p<0.001, ** p<0.01, * p<0.05

Table B2d. Robustness: Structural parameter estimates with organic farms only and organic transition density

	Crop Price (\$0.26/lb)	Organic Crop Price (\$6.00/lb)
Infestation loss tier 1 (Infestation \leq 0.3%)	0.000 (0.000)	0.000 (0.000)
Infestation loss tier 2 (0.3% < Infestation \leq 1%)	0.000 (0.376)	0.000 (0.376)
Infestation loss tier 3 (1% < Infestation \leq 5%)	0.000 (0.000)	0.000 (0.000)
Infestation loss tier 4 (5% < Infestation \leq 10%)	1.000*** (0.385)	1.000*** (0.385)
Infestation loss tier 5 (10% < Infestation)	0.000 (0.138)	0.000 (0.138)
# Observations	342	342
# Farms	19	19

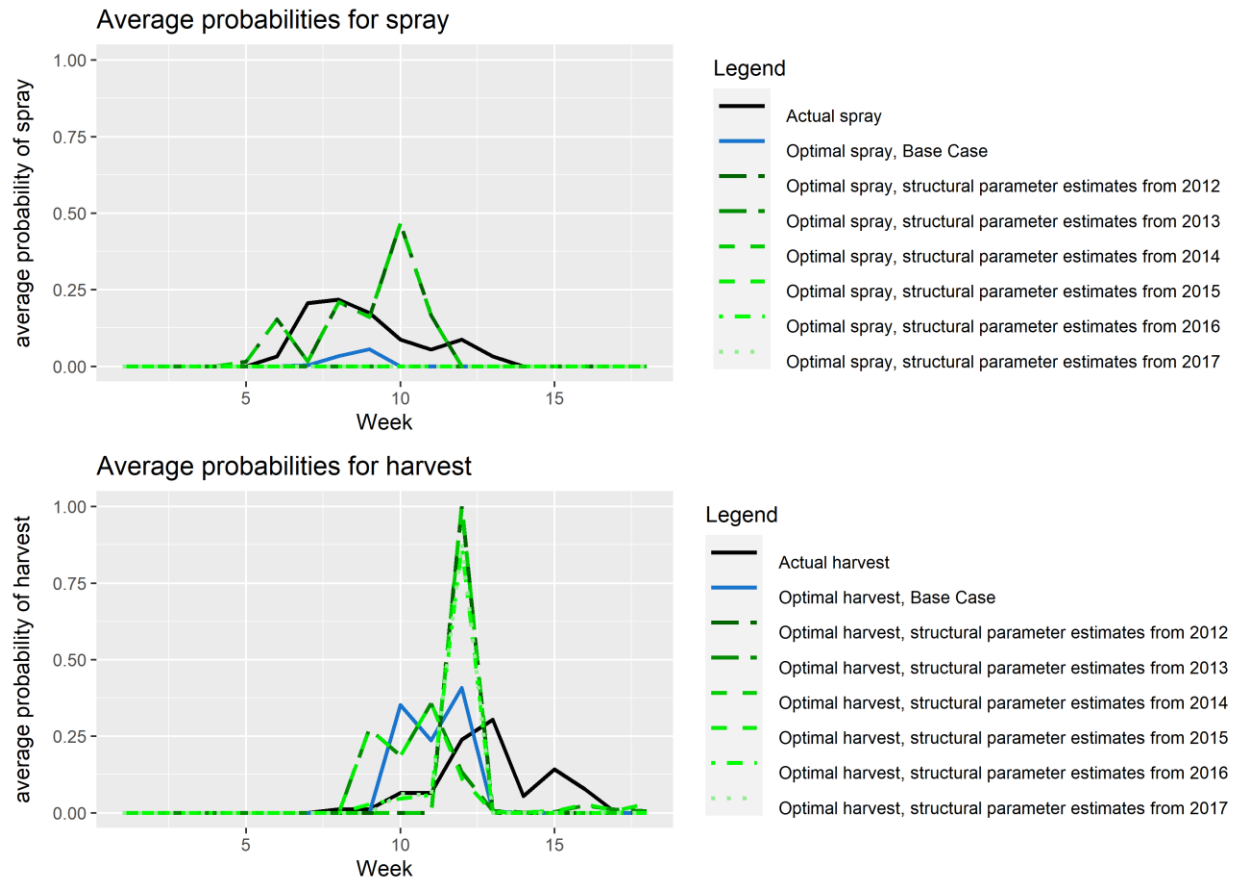
Notes: In both of these alternative specifications, the sample is restricted to organic farms only. For the crop price, we use either the crop price of \$0.26/lb as before, or the organic crop price of \$6.00/lb. Infestation is the observed larval infestation. The infestation loss is the incremental yield loss at each tier of observed larval infestation; the cumulative value of infestation loss is the (cumulative) yield loss $loss(y_{larva})$. Bootstrapped standard errors in parentheses. Significance codes: *** p<0.001, ** p<0.01, * p<0.05

Table B2e. Robustness: Structural parameter estimates by year

	2012	2013	2014	2015	2016	2017
Infestation loss at tier 1 (Infestation \leq 0.3%)	0.000 (0.000)	0.000** (0.000)	0.000 (0.000)	0.000*** (0.000)	0.000 (0.000)	0.000 (0.000)
Infestation loss at tier 2 (0.3% < Infestation \leq 1%)	0.000 (0.000)	1.000** (0.435)	0.000 (0.431)	1.000*** (0.000)	0.000 (0.000)	0.000 (0.435)
Infestation loss at tier 3 (1% < Infestation \leq 5%)	0.000 (0.000)	0.000 (0.000)	0.000 (0.038)	0.000 (0.000)	0.000 (0.184)	0.000 (0.000)
Infestation loss at tier 4 (5% < Infestation \leq 10%)	0.690*** (0.168)	0.000 (0.435)	0.064 (0.405)	0.000 (0.000)	1.000*** (0.184)	0.209*** (0.075)
Infestation loss at tier 5 (10% < Infestation)	0.310* (0.168)	0.000 (0.000)	0.936*** (0.057)	0.000 (0.000)	0.000 (0.000)	0.791* (0.429)
Spray cost (\$ per acre)	3,862.62 (4,045.171)	2,616.37*** (263.525)	3,578.02 (3,029.28)	3,042.76*** (169.67)	3,081.74*** (298.26)	3,106.77*** (339.46)
# Observations	360	306	252	306	252	180
# Farms	20	17	14	17	14	10

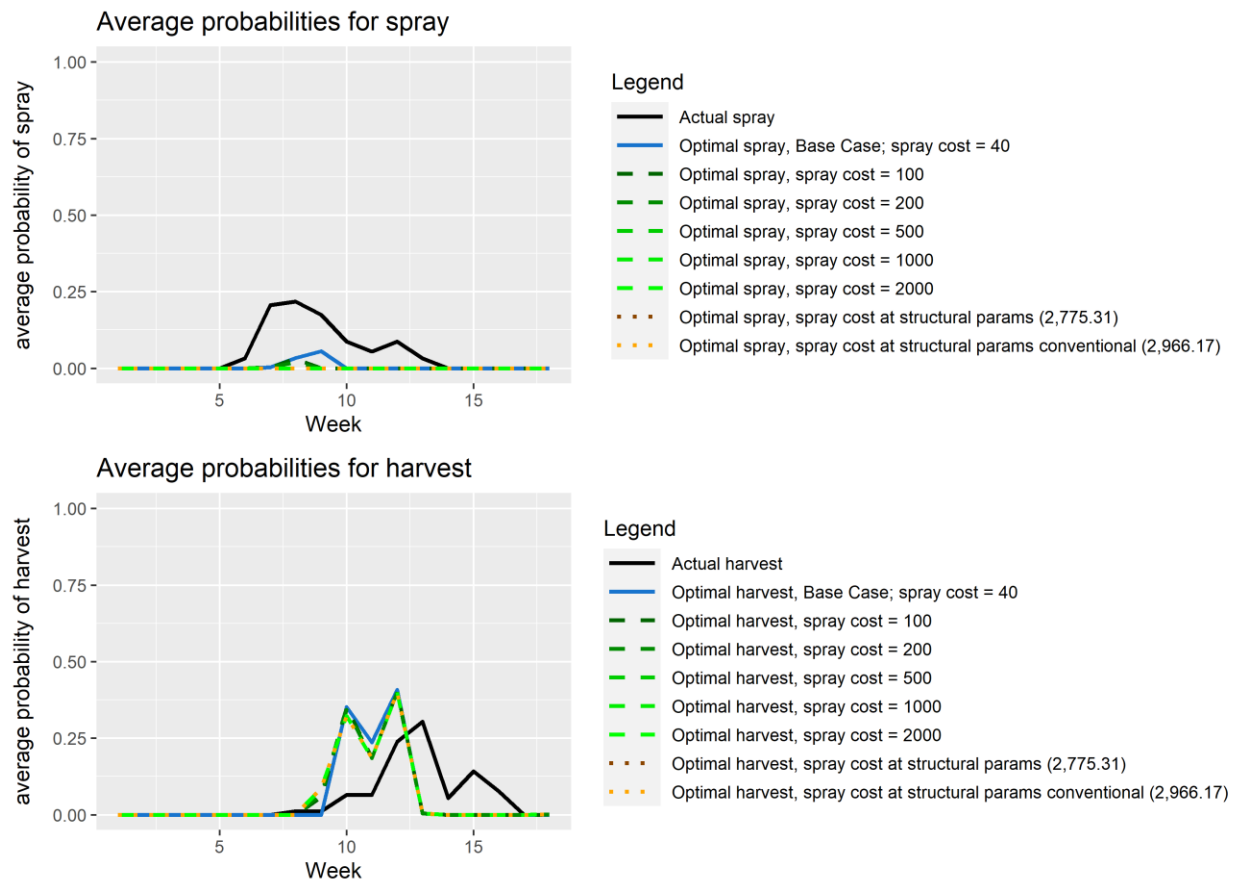
Notes: In this set of alternative specifications, we run the structural model for each year using data from that year only and using the average crop price for that particular year for the crop price. Infestation is the observed larval infestation. The infestation loss is the incremental yield loss at each tier of observed larval infestation; the cumulative value of infestation loss is the (cumulative) yield loss $loss(y_{larva})$. Bootstrapped standard errors in parentheses. Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Figure B4. Optimal vs. actual probabilities of spraying and harvesting conditional on growers' beliefs and perceptions based on structural parameter estimates by year




Notes: Figure compares the optimal and actual average probabilities of spraying and harvesting, conditional on growers' beliefs from the structural parameter estimates by year in Table B2e in Appendix B. The average probabilities of actual spray and actual harvest are calculated by averaging the observed actions over all 92 farms in the data set. For each set of structural parameters estimates by year, the optimal average probabilities of spraying and harvesting conditional on growers' beliefs use the parameters estimated from the structural model for that year, and are given by averaging over 100 simulations for each farm using the policy function from solving our numerical model using the parameters estimated from the structural model for that year and the non-parametrically estimated transition densities, and then averaging over all 92 farms.

Figure B5. Spray cost beliefs and perceptions: Optimal probabilities of spraying and harvesting for large values of spray costs



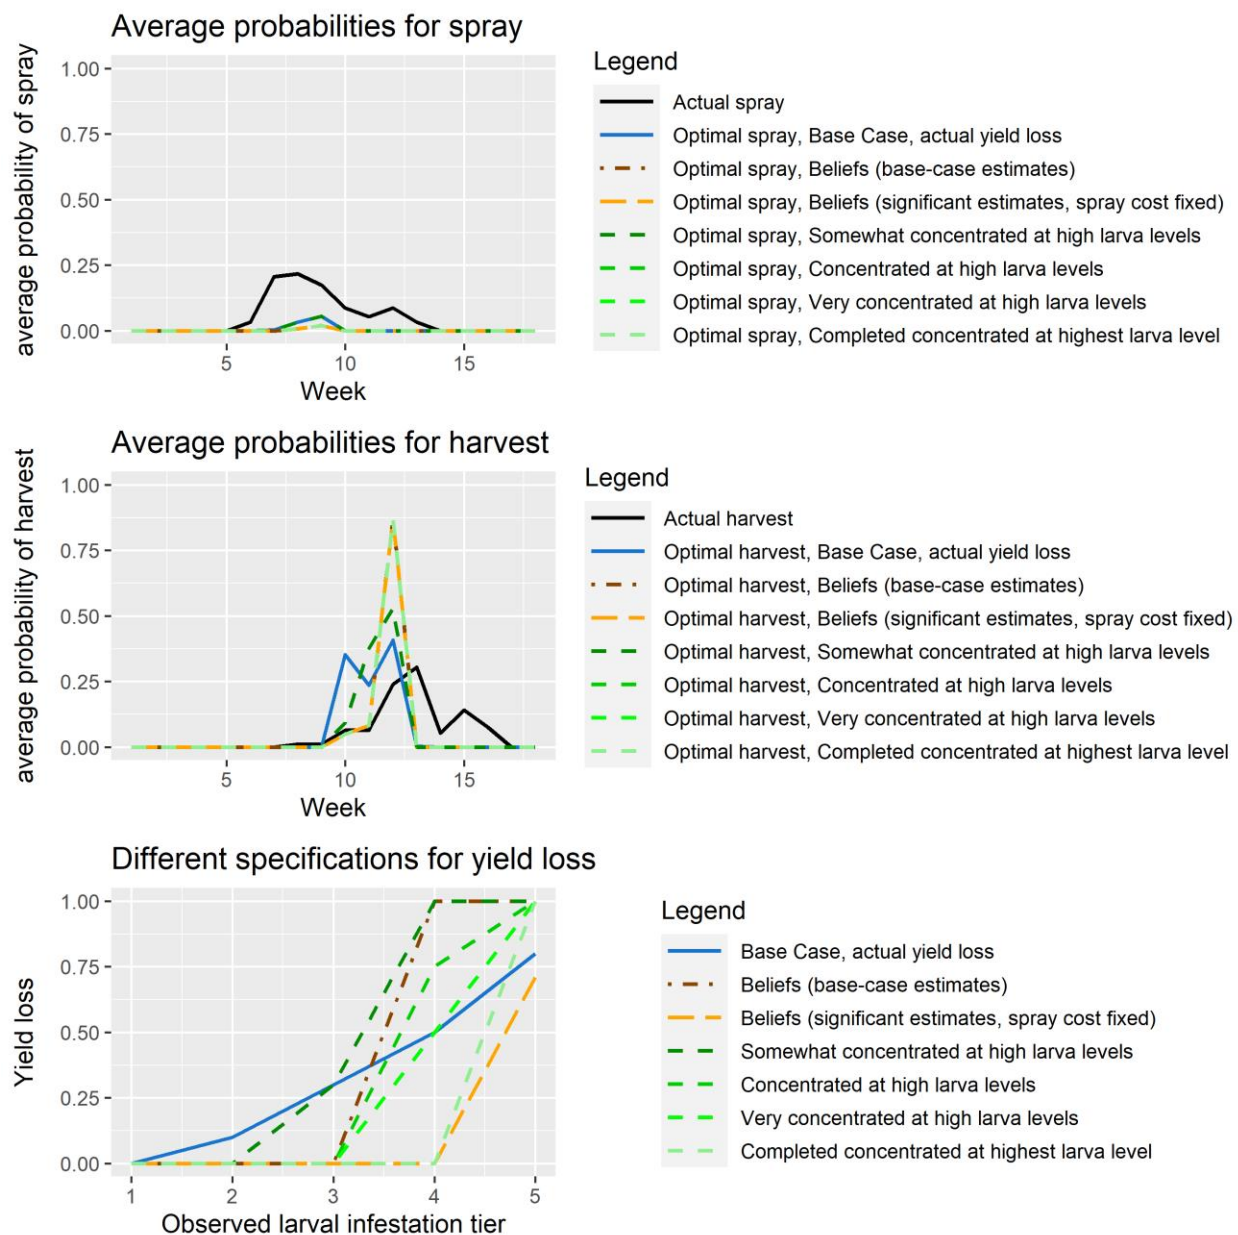
Notes: To better understand why a high perceived spray cost better rationalizes the choices made by the growers, this Figure compares the optimal and actual average probabilities of spraying and harvesting for large values of spray costs, holding the infestation loss fixed at its assumed values based on expert opinion and extension reports in Table 1. The average probabilities of actual spray and actual harvest are calculated by averaging the observed actions over all 92 farms in the data set. The optimal average probabilities of spraying and harvesting are given by averaging over 100 simulations for each farm using the policy function from our numerical model and the non-parametrically estimated transition densities, and then averaging over all 92 farms.

Figure B6. Growers' beliefs and perceptions about infestation loss when holding spray costs fixed at base case value



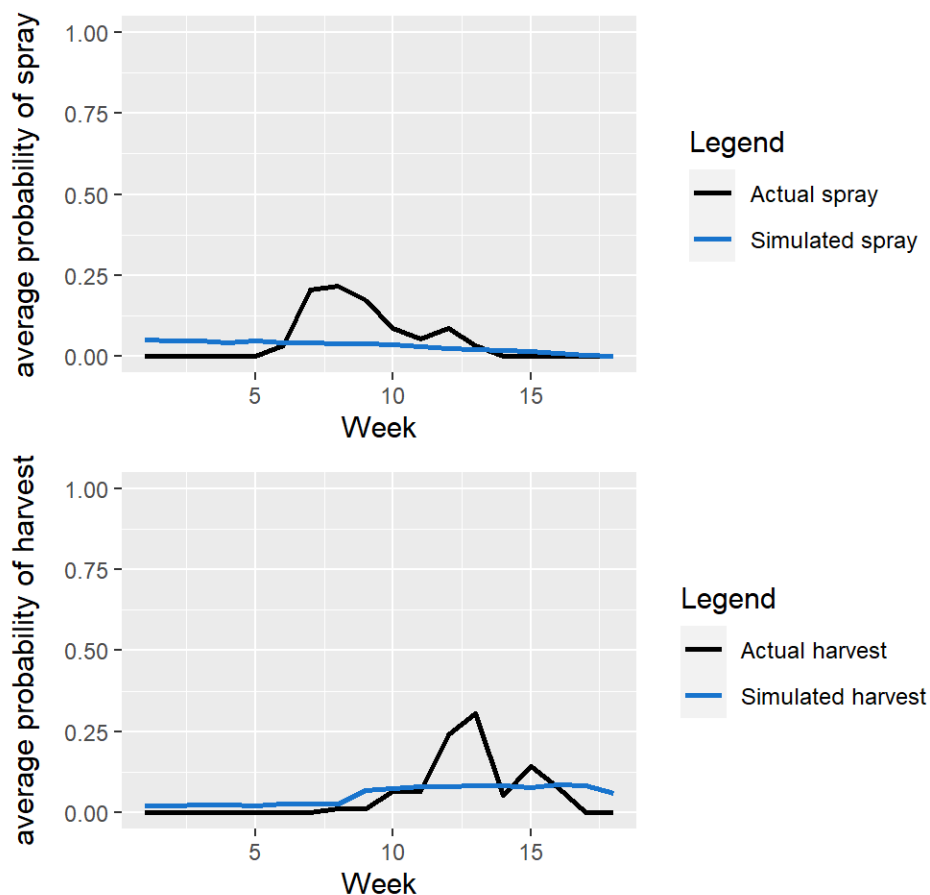
Notes: Figure compares the actual yield loss $loss(y_{larva})$ as a function of observed larval infestation y_{larva} with growers' beliefs and perceptions about yield loss from the alternative structural parameter estimates in Table 5b in which spray cost is held fixed at \$40 per acre. The yield loss $loss(y_{larva})$ is the cumulative percentage of yield loss at each tier of observed larval infestation. The actual yield loss $loss(y_{larva})$ values are the assumed values based on data and expert opinion based on expert opinion and extension reports (Burrack 2014; De Ros et al. 2015; DiGiacomo et al. 2019; Drummond et al. 2019; Yeh et al. 2019; Yeh et al. 2020) reported in Table 1. The growers' beliefs and perceptions are based on the alternative structural parameter estimates in Table 5b in which spray cost is held fixed at \$40 per acre. The dashed blue lines indicate the 95% confidence interval for the growers' beliefs and perceptions, as calculated using the standard errors from the alternative structural parameter estimates in Table 5b in which spray cost is held fixed at \$40 per acre.

Figure B7. Infestation loss beliefs and perceptions: Optimal probabilities of spraying and harvesting when yield loss is concentrated at high observed larval infestation levels



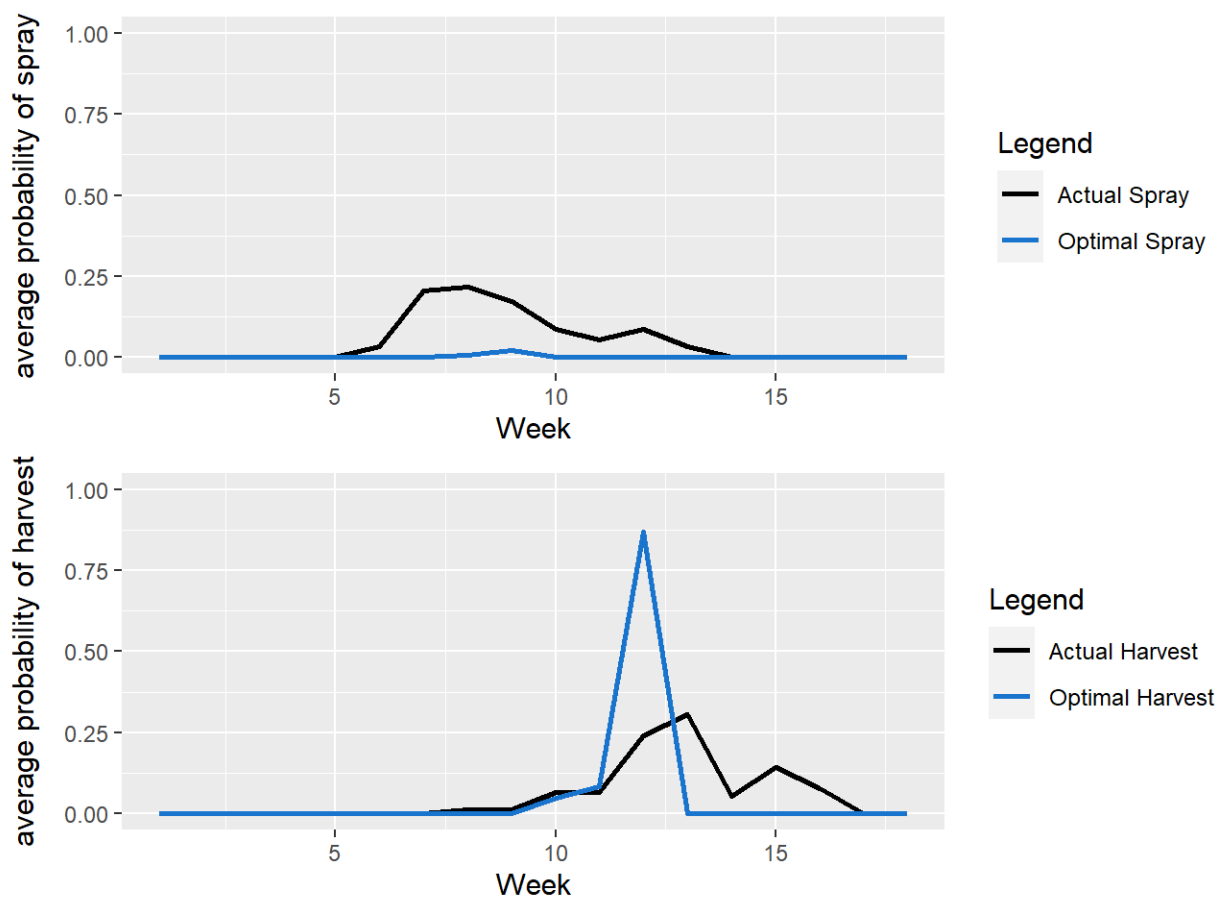
Notes: Figure compares the optimal and actual average probabilities of spraying and harvesting for various specifications of infestation loss in which the yield loss is low at low observed larval infestation levels and high at high observed larval infestation levels, holding the spray cost fixed at its assumed value of \$40 per acre (Esau 2019). For the growers' beliefs about yield loss in brown and orange, we use the parameters that are statistically significant at a 5% level from, respectively, the base-case specification of the structural model in Table 3 (brown) and the alternative specification of the structural model Table 5b in which spray cost is held fixed at \$40 per acre (orange). The optimal average probabilities of spraying and harvesting are given by averaging over 100 simulations for each farm using the policy function from our numerical model and the non-parametrically estimated transition densities, and then averaging over all 92 farms.

Figure B8. Optimal probabilities of spraying and harvesting conditional on growers' beliefs and perceptions when also accounting for unobservable state variables



Notes: Figure compares the optimal and actual average probabilities of spraying and harvesting, conditional on growers' beliefs from the base-case structural parameter estimates in Table 3 and when also accounting for unobservable state variables. The average probabilities of actual spray and actual harvest are calculated by averaging the observed actions over all 92 farms in the data set. The optimal average probabilities of spraying and harvesting conditional on growers' beliefs use the parameters estimated from the structural model, and are given by averaging over 100 simulations when using the base-case structural parameter estimates in Table 3 and when also accounting for unobservable state variables, and then averaging over all 92 farms.

Figure B9. Optimal probabilities of spraying and harvesting conditional on growers' beliefs and perceptions when holding spray costs fixed at base case value



Notes: Figure compares the optimal and actual average probabilities of spraying and harvesting, conditional on growers' beliefs from the alternative structural parameter estimates in Table 5b in which spray cost is held fixed at its assumed value of \$40 per acre (Esau 2019). The average probabilities of actual spray and actual harvest are calculated by averaging the observed actions over all 92 farms in the data set. The optimal average probabilities of spraying and harvesting conditional on growers' beliefs use the parameters estimated from the structural model, and are given by averaging over 100 simulations for each farm using the policy function from solving our numerical model using the parameters estimated from the structural model and the non-parametrically estimated transition densities, and then averaging over all 92 farms.