The Economic Costs of Climate Change from Ecosystem Disruptions: Evidence from Forest Defoliation Outbreaks in the United States*

Sushant Banjara[†1], Greg Dwyer[‡2], Eyal G. Frank[§1,3,4], and Andrew M. Liebhold[¶3]

[1] *Harris School of Public Policy, University of Chicago*

[2] *Department of Ecology & Evolution University of Chicago*

[3] *Center for Economic and Policy Research, Paris, France.*

[4] *National Bureau of Economic Research, Cambridge, MA.*

[5] *Czech University of Life Sciences*

March 25, 2026

Abstract

Climate change will affect human well-being not only through direct impacts, but also through changes in ecosystem functioning that alter the flow of local amenities and disamenities. We study one such climate–ecosystem–well-being chain in U.S. forests, focusing on the weather-dependent interaction between an invasive defoliator, the spongy moth, and a fungal pathogen that normally suppresses it. Using fine-scaled defoliation maps, gridded weather, and a long-run panel of spongy moth presence, we show that warm winters and warm, dry springs—especially when they persist over multiple years—substantially increase the likelihood and spatial extent of defoliation outbreaks. We then link these ecological shocks to nearly five decades of geocoded housing transactions and estimate that exposure to defoliation lowers nearby sale prices by at least two percent, with cumulative losses of up to about six percent for properties exposed in several consecutive years. Negative price effects remain detectable for homes located up to six kilometers from defoliated areas and are roughly twice as large when defoliation is a new shock in the local housing market. The stakes of this mechanism will grow over time, as the spongy moth currently occupies only about one-third of its climatically suitable habitat and is projected to expand into many more forests and communities. Unlike direct weather–well-being channels, this class of weather–ecosystem–well-being impacts offers additional scope for mitigation and adaptation, because managers can intervene to stabilize natural enemy interactions or to substitute for declines in ecosystem functioning. Our findings show that climate-driven disruptions to natural enemy interactions can generate sizable, spatially diffuse losses in housing wealth and highlight ecosystem management as a distinct margin for climate adaptation.

*Acknowledgments: We thank Pete Rodrigue for excellent research assistance. This work is supported by the Griffin Incubator Energy Challenge at the University of Chicago, as well as by a grant from the National Science Foundation (#2043796). Corresponding author E-mail: eyalfrank@uchicago.edu. All errors are our own.

†Harris School of Public Policy, University of Chicago, Keller Center, 1307 E 60th St. Chicago, IL 60637. Email: sbanjara@uchicago.edu

‡Department of Ecology & Evolution, University of Chicago, 1101 East 57th St. Chicago, IL 60637

§Harris School of Public Policy, University of Chicago, Keller Center, 1307 E 60th St. Chicago, IL 60637. Email: eyalfrank@uchicago.edu

¶Czech University of Life Sciences

1 Introduction

Climate change will affect human well-being not only through direct impacts such as heat stress, storms, and sea level rise, but also through changes in how ecosystems function and the flow of amenities and disamenities they generate. Stable natural environments support a continuous stream of local benefits—such as shade, scenic quality, and recreational opportunities—while buffering people from disamenities like pests, allergens, and disease vectors. A growing body of work has quantified the social costs of climate change along the direct channels of the weather–well-being relationship (Dell et al. 2014; Hsiang and Kopp 2018; Kolstad and Moore 2020), but we know comparatively little about the economic consequences of climate-driven disruptions to the ecosystem interactions that regulate these amenities.[1] These ecosystem-mediated impacts matter because they can amplify or attenuate the welfare costs of extreme weather events. They also open additional margins for mitigation and adaptation through ecosystem management, not just through emissions reductions or built-environment investments. Understanding how changes in environmental conditions disrupt amenity-regulating ecosystem interactions is therefore highly relevant, particularly because sustaining ecosystem functioning at a stable level can be costly.

In this paper, we study a concrete ecosystem channel through which climate can affect local amenities: the weather-dependent interaction between an invasive forest pest, the spongy moth (*Lymantria dispar*) and a fungal pathogen that normally suppresses it. The spongy moth is often described as one of the most damaging invasive forest pests in North America (Liebhold et al. 1992; Michael and Csóka 2007). Weather conditions, especially warm and dry springs, limit the ability of the pathogen to control moth populations, allowing them to increase by several orders of magnitude and triggering large-scale defoliation events (Kyle et al. 2022; Liu et al. 2025). During these outbreaks, caterpillars consume the foliage of hardwood trees, defoliating millions of forest acres, increasing the risk of tree mortality,

[1] Climate change is already affecting species distribution and ecosystem functioning, and is projected to further shift ranges and alter community composition (Field et al. 2007; Lawler et al. 2009; Berdugo et al. 2020).

and sharply degrading the bundle of forest-related amenities enjoyed by nearby households. These visible, spatially concentrated ecological shocks provide a natural setting to study how disruptions to climate-sensitive natural enemy interactions affect ecosystem functioning and, in turn, how these changes are capitalized in the housing market.[2]

Our analysis centers on a common climate–ecosystem pathway—weather-driven disruptions to natural enemy interactions—and traces its consequences from increased defoliation risk to the capitalization of those risks in nearby property markets. Accordingly, we study both the ecological response of defoliation to weather and the economic response of housing markets to defoliation. First, we ask how weather conditions affect the probability and severity of spongy moth defoliation at scale, building on plot-level experiments and process-based projections to provide large-sample evidence on a climate–ecosystem link. Second, we ask how these defoliation shocks, interpreted as disruptions to amenity-regulating ecosystem functioning, are capitalized into nearby property values. By combining multiple decades of fine-scaled data on defoliation attributed to spongy moth, gridded weather, and geocoded housing transactions across the insect's expanding range, we document economically meaningful effects along both margins, providing quantitative evidence on a specific climate–ecosystem–well-being chain and on the economic stakes of maintaining ecosystem functioning through costly management and adaptation efforts.

We find that weather systematically pushes the system toward outbreaks: warm winters and warm, dry springs, particularly when they persist over multiple years, are associated with substantial increases in the area defoliated by spongy moth, and these climate-driven disruptions to natural enemy interactions both expand the spatial footprint of forest defoliation and impose economically meaningful losses on homeowners living in affected areas. Linking these ecological shocks to geocoded housing transactions, we show that defoliation exposure leads to a reduction of at least two percent in nearby sale prices, with cumulative losses of up to roughly six percent for properties exposed in several consecutive years, and

[2]Natural enemies are broadly defined as herbivores, predators, parasites, or pathogens that play an important role in the population dynamics of their hosts or prey (Hajek and Eilenberg 2018).

that negative effects on prices remain detectable for homes located up to six kilometers from defoliated areas. Moreover, effect sizes are nearly twice as large when defoliation represents a new shock in the local housing market, rather than a repeat occurrence, underscoring the salience of innovations in ecosystem disruptions. Taken together, these results imply that climate-driven disruptions to natural enemy interactions can both expand the spatial footprint of forest defoliation and impose economically meaningful losses on homeowners living in affected areas.

Spongy moth populations have long exhibited large, hard-to-predict outbreaks (Michael and Csóka 2007), but the emergence of the fungal pathogen *Entomophaga maimaiga*, and management efforts by the U.S. Forest Service show that the risks posed by this climate-sensitive system are partly controllable (Hajek 1997). Historically, outbreaks have defoliated millions of acres in a single year, yet biological control by the fungal pathogen and coordinated monitoring and treatment programs have substantially reduced the frequency and severity of these events across much of the insect's current range (Sharov et al. 2002; Coleman et al. 2020). Our results therefore point to specific ecosystem-level adaptation margins: avoiding management actions that exacerbate local spring drought stress and further weaken the fungal pathogen, and concentrating surveillance and large-scale suppression in drought-affected areas where the pathogen is least effective, in order to limit both the ecological footprint of defoliation and the associated losses in nearby property values.

Our findings close a key gap in the climate–economy literature by providing well-identified evidence on how climate-driven disruptions to ecosystem functioning—specifically, weather-sensitive natural enemy interactions—translate into economic losses in housing markets (Dell et al. 2014; Hsiang and Kopp 2018; Kolstad and Moore 2020; Field et al. 2007; Lawler et al. 2009; Berdugo et al. 2020). We do so by bridging three areas of research. First, we complement work that uses weather variation to estimate the social costs of climate change by focusing on an ecosystem-mediated channel through natural enemies, rather than solely on direct weather–outcome relationships (Dell et al. 2014; Hsiang and Kopp 2018; Kol-

stad and Moore 2020; Letourneau et al. 2009; Colwell et al. 2012; Kroeker and Sanford 2022; Hajek and Eilenberg 2018; Reilly et al. 2014; Kyle et al. 2020; Liu et al. 2025). Second, we contribute to an emerging literature on how biological features of the environment affect human well-being—through natural enemies, apex predators, and forests—by valuing a climate-sensitive disruption to forest ecosystem functioning in a residential context (Frank 2024; Frank and Sudarshan 2024; Raynor et al. 2021; Lu et al. 2026; Xing et al. 2023; Latifovic et al. 2024). Third, we build on studies that use invasive species to identify environmental damages, both in economic history and in modern forest ecosystems, by showing that weather shocks which destabilize natural enemies generate localized property value losses and thereby expand the set of mechanisms through which climate change imposes economic costs (Banerjee et al. 2010; O'Rourke 1994; Lange et al. 2009; Jones and McDermott 2018; Druckenmiller 2020; Leuschner et al. 1996; Lu et al. 2026; Bigsby et al. 2014; Tuthill et al. 1984; Xing et al. 2023).

Previous work on the damages from defoliation has been limited by data availability at a sufficiently fine-scaled resolution. As a result, studies have broadly generalized damage function parameters and calculated back-of-the-envelope estimates for the damages or potential benefits from slowing the spread of the moth. Considering several scenarios and parameter values, Leuschner et al. (1996) evaluate potential benefits from timber operations, recreational use of forests, and residential property value, and arrive at a value of $3.8 billion under the high efficacy scenario. Bigsby et al. (2014) simulate the costs of potential outbreaks for the city of Baltimore, estimating residential property damages of $63 million in a high outbreak scenario. Other work has focused on the disamentiy from the hair on the body of the caterpillar, which are airborne and cause skin and respiratory irritation (Tuthill et al. 1984).

In what follows, we review the history of the invasion of *L. dispar*, its population dynamics, and interactions with the fungal pathogen, *E. maimaiga*; the data sources we used in the analysis; the methods and results we obtain for the (i) weather-dependent natural enemy

interaction, and (ii) the impacts of defoliation on property values; finally, we conclude.

2 Spongy Moth Invasion, Expansion & Dynamics of Cyclical Tree Defoliation

Spongy moth is an invasive insect species that is considered to be one of the most damaging invasive species in North America (Liebhold et al. 1992; Michael and Csóka 2007).[3] Since its introduction to North America in 1869, there have been considerable eradication and control efforts that have tried to reduce its negative economic impacts and damages to forest ecosystems (Liebhold et al. 1989). During the period of 1869 to 2019, there have been three historical defoliation events: 1981, 1990, and 2006, with 1981 being by far the largest recorded defoliation event in the US. See Figure 1a for the total amount of defoliated forest area by year, along with key regime changes in the population dynamics of the moth.

Damages from the spongy moth might be higher in the future because it is still expanding its range. In recent years, the expansion rate has stabilized at about three km per year, down from its peak expansion rate of 20 km per year between 1966 and 1990 (periodic expansion rates are shown at the bottom of Figure 1a). Since its invasion, the moth has expanded its range into more than six-hundred counties by 2019. Figures 1b-1c map the expansion, moving to the south and to the west from the northeast, as well as summarize the number of newly classified moth-established counties.

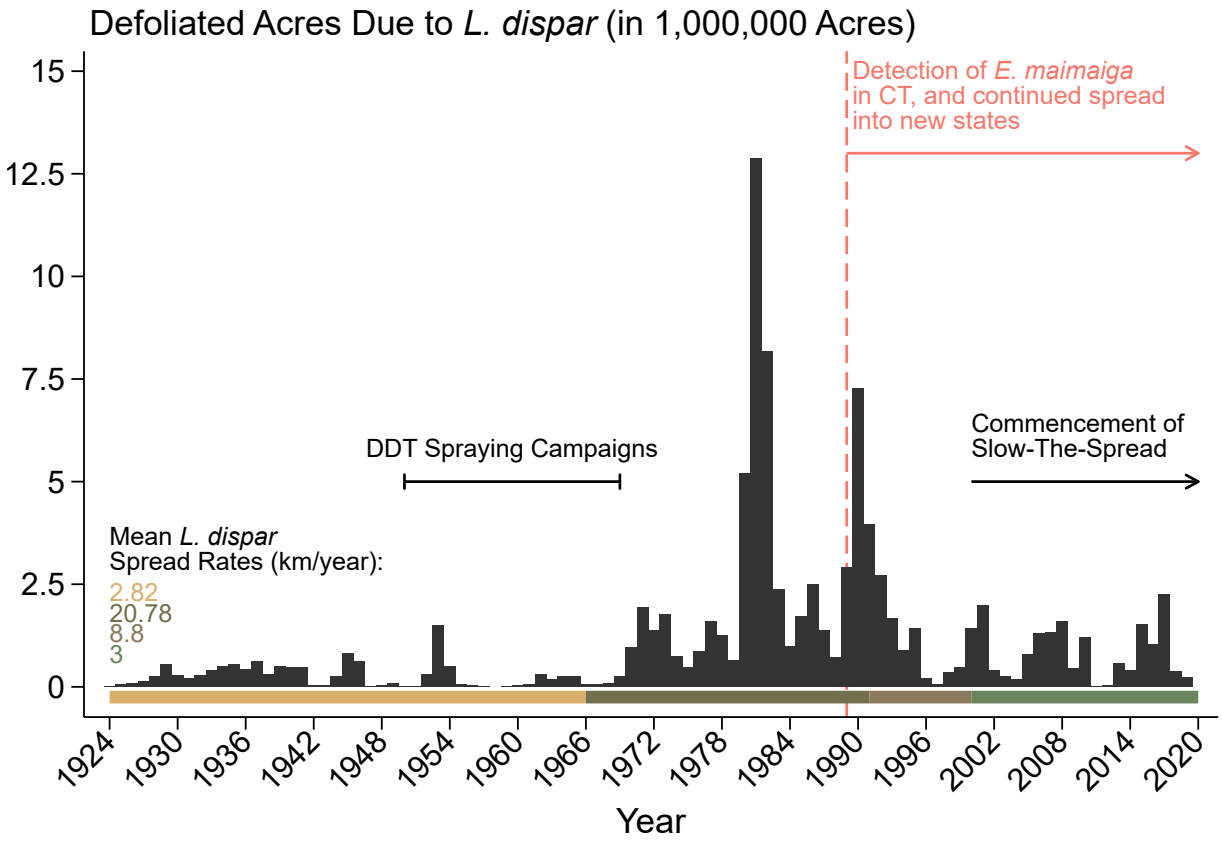
The reason why moth defoliation did not utterly devastate hardwood forests in North America is that outbreaks were reliably terminated by a fatal, species-specific baculovirus, which caused over ninty-five percent mortality (Woods and Elkinton 1987).[4] Since 1989, however, the role of the baculovirus in controlling moth populations has been almost entirely superseded by the species-specific fungal pathogen, *Entomophaga maimaiga* (Hajek 1997).

[3]The previous common name for the species, the "gypsy moth," is no longer used as of July, 2021 because the Entomological Society of America recognizes it derogatory towards the Romani people. For more details, see: https://www.nytimes.com/2021/07/09/science/gypsy-moth-romani-entomological-society.html. Accessed: December 20, 2025.

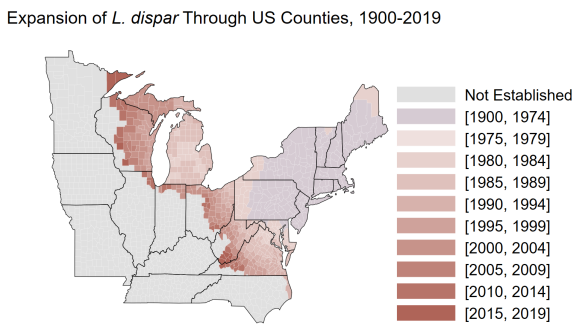
[4]Baculoviruses are a family of viruses that are known to infect insects.

Figure 1: Defoliation by the Spongy Moth & Range Expansion Over Time

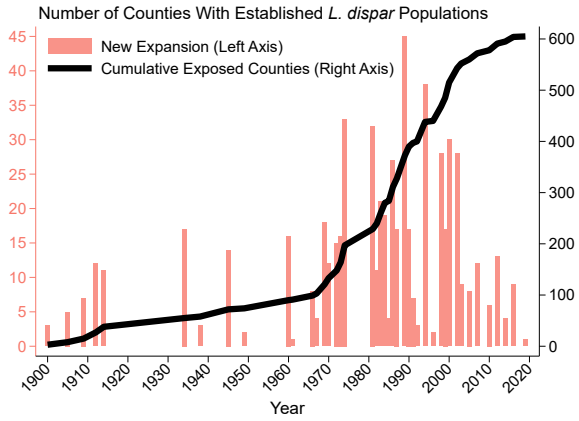
(a) National Level Defoliation



(b) Map of Range Expansion



(c) Moth-Established Counties



Notes: Summarizing in Panel (a): the area damaged by the spongy moth, and key regime changes; and in Panels (b) and (c), the expansion throughout the Northeastern United States.

The fungal pathogen dispersed across through most of the moth's range by 1996, and was found throughout all the states that have active moth populations by 2000 (Hajek et al. 2021).

The life cycle of the insect has four stages. Starting with an egg mass that is usually located either on the ground, rocks, or the lower end of trees. The eggs hatch into the second stage, known as the larva stage (commonly referred to as caterpillars). This is the stage where the insect inflicts its damage on trees as the larva crawl up the tree, and feed on the leafy material. After the larva has accumulated the required nutrients, it progresses to construct a puparium—a hard exterior structure (the cocoon)—where it completes the transformation from a pupa to the moth. Only male moths are capable of flying, and they fly looking for a female moth to reproduce and deposit a new egg mass, starting the cycle all over again. The fact that only male moths can fly limits the spatial extent to which different populations are connected across space, and limits the rate of expansion. However, during large outbreak events, egg masses get transported outside of the moth-infested area into new areas, allowing the range to expand considerably over short periods of time (Michael and Csóka 2007).

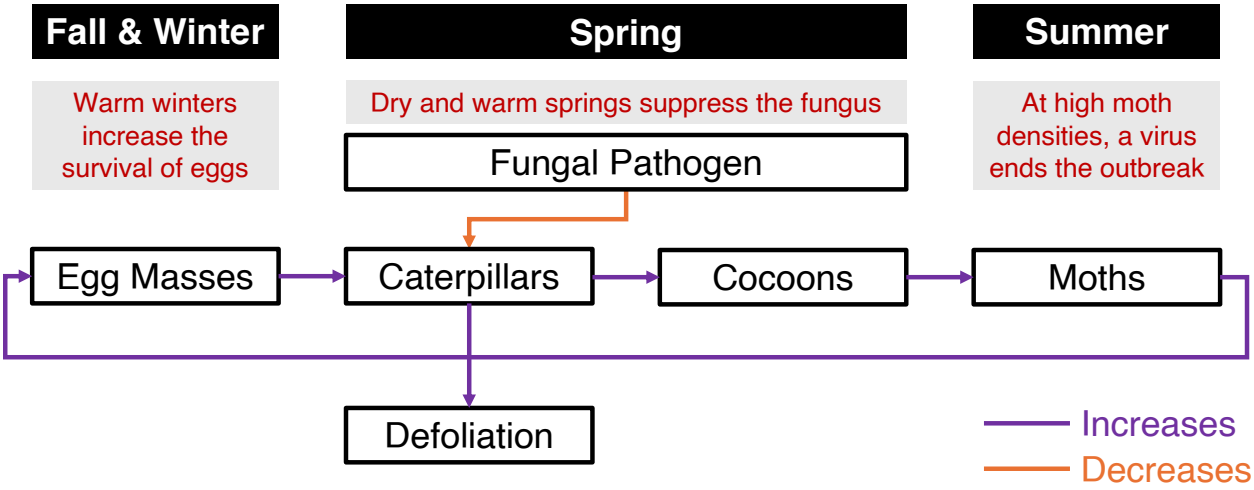
Across North America, Europe, and Asia, spongy moth populations exhibit "eruptive" dynamics that make it hard to accurately predict their population dynamics. However, recent studies have documented that population dynamics follow an eight- to twelve-year oscillatory cycle. Both ecosystem interactions and weather conditions have been suggested as mechanisms that drive these oscillatory patterns (Johnson et al. 2005). Similar to other insect populations, across large spatial scales, the moth is hypothesized to follow dual equilibria dynamics (Campbell and Sloan 1977; Liebhold et al. 2000). The dual equilibria dynamics model is a model with two stable equilibria: a high and a low population level state. Even if the population is in the low population equilibrium value, with a sufficiently large perturbation to population-density dependent control, the population of the insect can increase above the low population equilibrium and enter the attraction basin of the high population equilibrium.

There are several mechanisms that connect weather with moth population dynamics. At low moth population densities, acorn production by oak trees regulates the population size of small mammals that are predators of the insect in its larva and pupa life stages. Either frost conditions in the spring, or drought conditions in the summer, reduce the amount of acorns, and of the small mammals that depend on them during the winter. Because these small mammals are generalist predators, and the moth is not their favorite food source, they are not an important source of biological control above low insect densities (2000). Another weather factor that acts to regulate the forest pest is winter temperatures. Extremely low temperatures, especially if they persist for several days, result in large egg-mass die-offs (Liebhold et al. 1992; Ward et al. 2020). Consequently, a spring that follows a winter with one or more cold spells will see fewer larva hatching from deposited eggs. We summarize how the spongy moth, *Entomophaga maimaiga*, and weather in Figure 2.

Prior to the emergence of the fungal pathogen, when the insect reached high population densities, a virus that specifically targets the spongy moth would spread quickly and decimate the population—effectively ending the outbreak (Dwyer and Elkinton 1993; Dwyer et al. 2000). However, since the first detection on the fungal pathogen in 1989, the pathogen now out-competes the virus, and infects the larva at a greater rate than the virus. While the fungal pathogen delivers a form of biological control that requires low levels of effort by forest managers, the fungus is known to be susceptible to warm and dry weather conditions during the spring that limit the production of spores—the main mechanism for its dispersion (Reilly et al. 2014).

The ongoing cycles of spongy moth populations and subsequent defoliation events remain difficult to predict. Despite these large uncertainties regarding its population dynamics, there is broad agreement that we will see additional expansion in the range of the forest pest (Kyle et al. 2022). Currently, the forest pest is found in one-third of its suitable habitat, suggesting that many more forest acres will become exposed to future defoliation events (Sharov et al. 2002).

Figure 2: Schematic Relationship Between Forest Defoliator, Fungal Pathogen & Weather



Notes: We summarize the life cycle of the spongy moth (the forest defoliator), across three main periods: fall and winter, spring, and summer. Defoliation occurs in the spring, during the caterpillar stage. This is also the stage when the fungal pathogen (*Entomophaga maimaiga*), can produce spores that parasite and kill the caterpillar—preventing insect outbreaks with large defoliation levels. Drought conditions in the spring suppress the ability of the fungal pathogen to control the moth. Multiple seasons with drought conditions can result in a large outbreak. Warm winters further accelerate the outbreak cycle as more egg masses survive from summer to spring.

3 Defoliation, Weather & Housing Prices Data

The goal of the empirical analysis is to estimate the economic impacts of defoliation by the moth in the Northeastern US.[5] Our analysis primarily uses a panel database at the county level, spanning 1975 to 2019 on defoliation, formal classification moth establishment by the US Forest Service (USFS), and monthly-level weather conditions. In addition, we use geo-coded transaction-level data on residential property values. We provide summary statistics for the main variables in Table 1. In the following, we provide a brief explanation of the data used, with additional details included in the Online Appendix.

[5]Our sample consists of the following states: Connecticut, Delaware, Illinois, Indiana, Iowa, Kentucky, Maine, Maryland, Massachusetts, Michigan, Minnesota, Missouri, New Hampshire, New Jersey, New York, North Carolina, Ohio, Pennsylvania, Rhode Island, Tennessee, Vermont, Virginia, West Virginia, and Wisconsin.

3.1 Defoliation & Forest Cover Data

We use defoliation maps produced by the USFS. During the summer, the USFS conducts aerial surveys of forests to evaluate the degree, if any, of damage by forest pests. We aggregate the defoliation to the county level. In order to construct a defoliation rate, we use the data from National Land Cover Database (NLCD), produced by the US Geological Survey (USGS) (USGS 2019). We define the defoliation rate as the defoliated area relative to the area classified as primary forest area in the NLCD. In some years, the defoliation rate is greater than one because areas that are not classified as primary forest by the USGS algorithm still get flagged as experiencing defoliation by the USFS. Despite the rate being occasionally larger than one, we find the normalization important to capture differences across counties in forest land.

3.2 Spongy Moth Trap Data

The cooperative national Slow the Spread (STS) program and related state partners collect data on spongy moth population levels. This population surveillance is carried out by deploying pheromone-baited delta and milk-carton traps in systematic grids along the expanding invasion frontier to detect newly established populations and quantify male flight activity.(Slow the Spread Foundation, Inc. 2024; USDA APHIS 2012; Wisconsin Department of Agriculture, Trade and Consumer Protection 2026) Traps are typically placed at regular spacings on host trees using GPS-guided grids (for example, one trap every 2–3 km in many STS detection areas), with higher densities in suspected or known infestation areas, and lower densities in background detection grids (USDA APHIS 2012; North Carolina Department of Agriculture and Consumer Services 2023). Delta traps are small cardboard triangle traps used over broad areas where low male densities are expected, whereas larger milk-carton traps are deployed in selected high-risk or higher-density areas and can capture many more individuals. Counts from these traps can be interpreted as noisy but informative

measures of local adult male abundance (Liebhold et al. 2020; Minnesota Department of Agriculture 2024a). In the data, each trap has a unique site identifier and coordinates. Traps enter and exit the sample over time as the STS grid shifts westward and southward, creating a rolling, spatially non-stationary panel in which early cohorts of traps may never overlap geographically with late cohorts (Slow the Spread Foundation, Inc. 2024; USDA APHIS 2012). Each data entry reports the site ID, coordinates, along with the reported number of individual moths caught, for the year and month in which the data were collected from that site.

3.3 Weather Data

We use data from the PRISM database on gridded daily weather from 1975 to 2019 (PRISM Climate Group, Oregon State University 2014). We use data on the minimum, maximum, and mean temperatures, as well as data on precipitation levels, and aggregate the gridded daily data to a monthly-county-level data set. We use the weather data to construct standardized deviations from average weather patterns across the 1975 to 2019 time period.

In addition, we also use data from the U.S. Drought Monitor on the Drought Severity and Coverage Index (DSCI). The DSCI ranges from values of 0 to 500. A value of 0 indicates that there are no drought conditions in the area, while a value of 500 indicates that the entire area is experiencing the most severe drought conditions. The DSCI is a weighted average (by area) of five different variables, each one capturing the percent of area that is experiencing drought conditions on a scale of 0 (least severe) to 4 (most severe). The data are available at the county-quarter level for the time period of 2000 to 2019.

3.4 Residential Property Transaction Data

We use data from the CoreLogic Tax and Deed data on residential property transactions from 1975 to 2019. The data include the year of transaction, as well as the coordinates

Table 1.
Summary Statistics

Variable	Mean	SD	Min	Max	N
Panel A. County-Level Data (Unweighted)					
First Moth Year	1980.3	24.9	1900	2019	27,093
Forst Area (sq. km)	0.62	0.778	0	12.5	66,781
Defoliation Rate	16.3	621.1	0	123,606.1	66,780
Burn Rate	1.2	29.6	0	3,783.9	53,424
Panel B. County-Level Data (Weighted)					
Defoliation Rate	12.3	82	0	123,606.1	66,780
Burn Rate	0.6	10	0	3,783.9	53,424
Panel C. Property Transaction Data					
log(Sale Price)	11.53	1.4	0	24.7	79,398,437
Sale Price	242,706.1	16×10^6	1	5×10^{10}	79,398,437
Direct Defoliation Exposure$_t$	0.0041	0.064	0	1	79,398,437
Direct Defoliation Exposure$_{t-1}$	0.0044	0.067	0	1	79,271,601

Notes: Summary statistics for the main variables used in the analysis. See the main text and Data Appendix for more details.

of the property, which allow us to assign both contemporaneous and lagged exposure to defoliation. For each property transaction, we create a set of dummy variables that indicate whether the coordinates of the property intersect with an area considered to be defoliated. In addition, we measure the nearest distance of defoliated area to the property up to 25 km. We repeat this process for the year of the transaction, and up to five years prior to the sale of the property.

4 Evidence for Weather-Dependent Defoliation Patterns

Experiments at a local scale of have demonstrated that warm and dry weather conditions during the spring adversely affect the fungal pathogen that constrains outbreaks of the moth

(Reilly et al. 2014; Kyle et al. 2020). Here we test whether this relationship between weather, the fungal pathogen, and the defoliating insect replicate at larger scales. Explicitly, we exploit the sudden and widespread emergence of *Entomophaga maimaiga* (hereafter, the fungal pathogen) throughout the Northeast. As of 1989 there were no observations suggesting that the fungus is present and well-established in US forests, but by 1996, it has established across the Northeast, and by 2000 it was found throughout the range of the moth (see Section 2 for more details).

4.1 Descriptive Evidence Using Weather & Defoliation Data

To flexibly represent the contemporaneous relationships between weather and defoliation, we construct standardized weather variable for the minimum winter temperature (January to March), mean spring temperature (April to June), and total spring precipitation per-squared km. We use the full range of years, 1975 to 2019, to calculate the standardized deviations. In Figure 3, we plot four panels that summarize the distributions of two of the weather variables, at a time, and the defoliation rate, along with the joint density of the two weather variables. To account for the differences in the area that is exposed to defoliation, we weight the observations by the forest area in the county.

We interpret the pattern in Figure 3a as evidence in favor of warm springs leading to higher defoliation rates, as positive deviations in the mean spring temperature (warmer weather) exhibit higher defoliation. However, contemporaneous spring precipitation does not exhibit the predicted effect. Higher defoliation rates are observed for both positive and negative deviations (wet and dry conditions), and the highest defoliation is observed during a warm spring that also experienced, on average, a three standard deviations higher levels of precipitation.

The distribution of the minimum winter temperature deviations to that of mean spring temperature deviations, in Figure 3b, suggests there is a bi-modal distribution. There is one cluster of high defoliation rates during warm winters and warm springs, however, this is also

a high defoliation cluster during warm springs and cold winters. The latter cluster suggests that a warm spring allows the moth to compensate for lost eggs during a cold winter as long as the fungal pathogen is limited in its capacity to parasitize the caterpillars due to the warm spring conditions. However, as the joint density in Figure 3d shows, the combination of a cold winter and a dry spring is much less likely relative to a warm winter *and* a warm spring.

4.2 Estimation Results for Defoliation & Drought Conditions

Our goal is to test whether there is evidence for a weather signal in defoliation. As discussed in Section 2, the key weather effects that can contribute to defoliation are: warm winters that allow more eggs to survive to spring, and warm and dry springs that reduce the capacity of the fungal pathogen to parasitize the caterpillars of the forest pest. To test the extent to which weather conditions explain the variation in summer defoliation, we first define a stable sample of counties where both the moth and fungal pathogen are present. Specifically, we use data from 2000 to 2019 for all the counties that had an official classification by the USFS as being in the moth-infested zone, and exclude counties with either a later classification or no classification for moth presence. We use 2000 as the cutoff because by then the fungal pathogen was detected and well-established throughout most of the range of the moth.

We focus on drought conditions in the spring by using the Drought Severity and Coverage Index (DSCI) for April, May, and June. See Section 3 for more details on the DSCI. In Figure 4, we plot the regression coefficients from a specification that estimates how defoliation changes with more severe spring drought conditions:

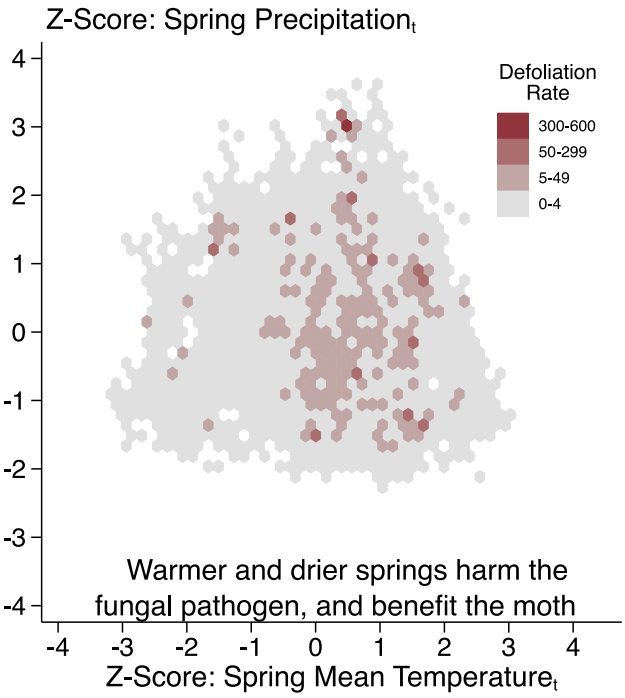
$$(\text{Defoliation})_{cst} = \sum_{k} \beta_k D^k_{cst} + \lambda_c + \delta_{st} + \varepsilon_{cst} \tag{1}$$

Where the outcome is the defoliation rate for county c, in state s, in year t. In the Online Appendix, we also report versions of the regressions that use the share of the county's area

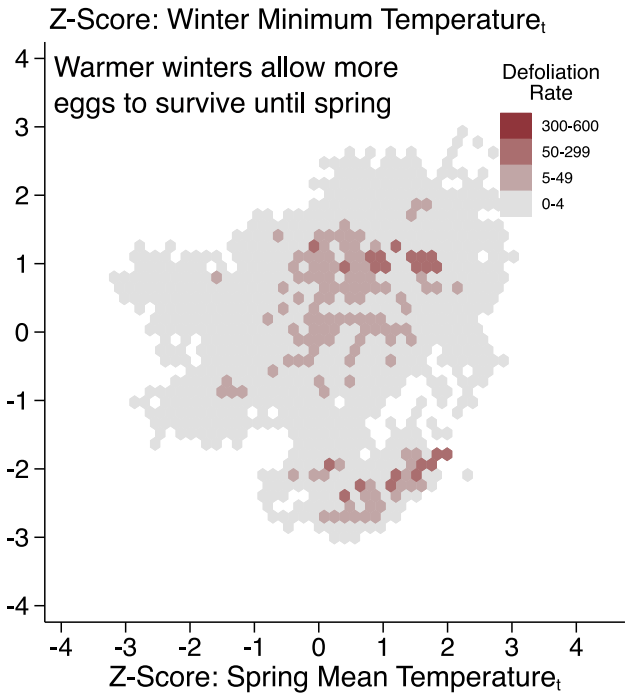
Figure 3: Weather and Defoliation Distributions

Spring Precip. VS. Spring Temp. | *Winter Temp. VS. Spring Temp.*

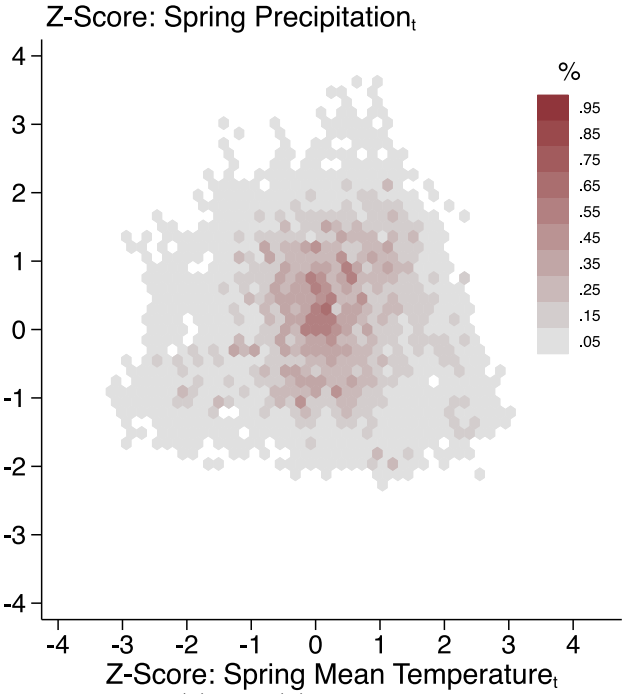
(a) Defoliation Rate



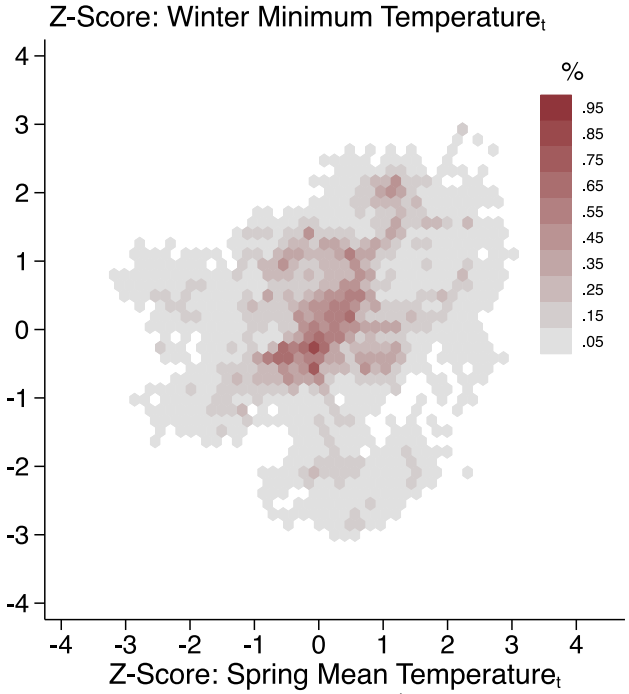
(b) Defoliation Rate



(c) Joint Density



(d) Joint Density



Notes: Panels (a) and (b) summarize the relationship between the weather deviation (Z scores calculated over the full range of the data, 1975 to 2019) and the defoliation rate. Panels (c) and (d) summarize the joint density of the weather deviations. Sample is defined as all counties classified as moth-established by 1996, with data from 1996 to 2019. Observations are weighed by forest acres at the county level.

that is defoliated, instead of the defoliation rate (which is normalized by the forested area in the county). We split the DSCI by equally-sized intervals of 25 points, with a top-coded category of above 200.[6] We also report results where we split the DSCI into deciles, conditional on having a non-zero DSCI value. Using either the equally-sized interval bins, or the deciles, in the above specification, no drought conditions in the spring serve as the omitted category, and we allow the regression to flexibly recover the relationship between drought conditions and defoliation. To account for lagged effects of spring conditions, we also calculate the mean DSCI for the current year, and up to two lagged years. When we include lagged drought values, we first calculate the mean and then split the non-zero values into bins or deciles.

To account for time-invariant characteristics at the county level, and to flexibly control for time trends, we include county fixed effects, λ_c, as well as state-by-year fixed effects, δ_{st}. We report either unweighted results, or weight observations by the number of acres that are classified as highly susceptible to defoliation by the spongy moth. Any unobserved heterogeneity is captured by the error term, ε_{cst}. We cluster the standard errors at the county level.

Higher drought conditions are monotonically linked to higher defoliation, until drought conditions become extreme, at which point defoliation declines to baseline levels. In Figure 4, we plot the results from the specification in Equation (1), using the spring drought conditions moving average for the current and lagged years. Across either the drought deciles or drought bins, we observe a similar pattern. First, at low levels of drought conditions, there is no meaningful difference in defoliation relative to the omitted category of no drought conditions during spring. Second, as drought conditions intensify, defoliation increases, and we can reject the null hypothesis that there is no difference in defoliation rates between high drought conditions relative to no drought conditions in the spring at the five percent significance level.

[6]In other words, D^1, is a dummy variable that is equal to one if the DSCI is strictly greater than zero, but smaller or equal to 25, D^2 is a dummy variable that is equal to one if the DSCI is strictly greater than 25, but smaller or equal to 50, and so on. D^9 is a dummy variable that is equal to one if the DSCI is strictly greater than 200.

Finally, when drought conditions become extreme, we observe a decline in the defoliation rate, and a return to baseline levels. This highlights that the fungal pathogen is sensitive to medium to high spring drought conditions, but that the spongy moth itself also suffers from extreme droughts. In other words, there is a range of environmental conditions that make the habitat unsuitable for the fungal pathogen, yet the moth can still perform well under those conditions.

Equally-sized interval bins (equal width) preserve the natural ordering and spacing of the underlying continuous variable, making it easier to interpret trends across the distribution and maintain meaningful distances between categories in regression models; however, they can lead to empty or sparse bins in the tails if the data is skewed, reducing statistical power and precision for those estimates. Deciles (equal-frequency bins), by contrast, ensure each category has roughly the same number of observations for more stable coefficient estimates across all bins, which is advantageous for skewed data, but they distort the x-axis spacing since bin widths vary, complicating visual interpretation of magnitudes and potentially masking non-linear patterns tied to the original scale.[7] Put differently, the first eight coefficients in Figure 4a roughly cover the same range of DSCI values as the first three coefficients in Figure 4b. We report more detailed estimation results for a combination of moving averages as well as using the share of the county's area that is defoliated, instead of the defoliation rate in Online Appendix Tables A1-A4.

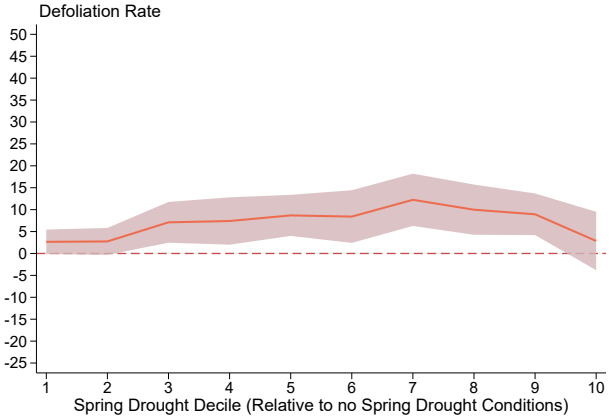
4.3 Estimation Results for Spongy Moth Trap Counts & Drought Conditions

An intermediate, but necessary, step in drought conditions resulting in large defoliation events is that the spongy moth population levels increase in drought-affected areas. Spongy moth population levels data are available along the habitat expansion border, instead of throughout the entire range that the moth inhabits. See Section 3 for more details on spongy

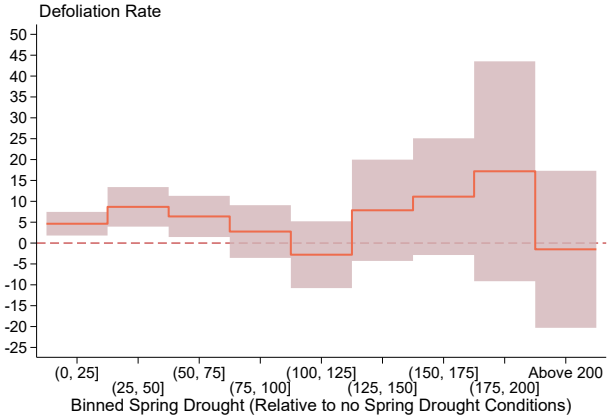
[7]In the data used in Figure 4, the decile cutoffs are: 3.85, 7.7, 11.8, 19.23, 24.7, 32.9, 42.5, 54.3, 79.9, and 247.3; while the densities in the bins are: 40.1, 20.3, 9.8, 3.9, 1.76, 1.7, 1.4, 0.6, and 0.17 percent.

Figure 4: The Effects of Spring Drought Conditions on Defoliation

(a) Spring Drought Deciles



(b) Spring Drought Bins



Notes: Estimation coefficients and 95% confidence intervals from the specification in Equation (1). We report results using the two-period moving average of spring drought conditions (using t and $t-1$). In Panel (a), we report results the split the mean DSCI values into deciles. In Panel (b), we report results that split the DSCI values into equally-sized interval bins. The omitted category in each regression is having no spring drought conditions in the county (i.e., spring DSCI = 0).

moth population surveillance. This section describes our empirical strategy for relating spongy moth trap catches to drought conditions. Our outcome is a measure of adult male moth activity at each trap, i, placed in county, c, reporting the annual, t, count of moths captured in delta or milk-carton pheromone traps.[8]

Three main challenges complicate using the spongy moth trap catch data in an analysis that links moth abundance to weather. First, an important feature of the spongy moth trapping network is the use of two distinct pheromone trap designs: small cardboard delta traps and larger milk-carton trap (Wisconsin Department of Agriculture, Trade and Consumer Protection 2026; Minnesota Department of Agriculture 2024b; USDA APHIS 2023). Delta traps are standard in broad detection grids and are deployed where expected moth densities are low, holding on the order of a few dozen males, whereas milk-carton traps are deployed in locations where populations are known or expected to be higher and can hold up to roughly 1,500 moths, aided by an internal insecticide-laced strip (Wisconsin Department

[8]On trap design and interpretation of catches, see Liebhold et al. (2020) and the USDA APHIS spongy program manual: https://www.aphis.usda.gov/sites/default/files/spongy-moth-program-manual.pdf. Accessed: March 10, 2026.

of Agriculture, Trade and Consumer Protection 2026; USDA APHIS 2012; 2014; Liebhold et al. 2020). Because the type of trap is systematically related to expected abundance and local infestation status, it is essential to allow the effect of drought on trap catches to differ between low- and high-density traps. Second, traps enter and exit the sample as the spongy moth infestation border expands, resulting in a rolling, spatially non-stationary panel. To partially restore spatial comparability across cohorts and control for broad-scale landscape heterogeneity (such as ecoregion, forest composition, and background infestation risk), we assign traps to 0.25° grid cells, and include grid-cell fixed effects in our regression models, in addition to trap and state-by-year, and state-by-month fixed effects (USDA APHIS 2012; Wisconsin Department of Agriculture, Trade and Consumer Protection 2026).[9] Third, of the nearly 1.93 million trap records, spanning 2001 to 2019, approximately 80 percent report zero catch, yet the 90th percentile of the reported catch level is 1 for low-density traps and 53 for high-density traps—reflecting a highly skewed distribution. In our analysis, we compare two main approaches to account for the large fraction of zeros and the skewed distribution: (i) a transformation-based linear fixed effects model using within-trap z-scores, and (ii) a generalized linear model for counts estimated via Poisson pseudo-maximum likelihood (PPML).

Standardized Trap Data.— A common way to homogenize highly skewed ecological counts is to transform them before applying linear models.[10] In our first approach, we convert raw counts within-trap standardized scores, z_{ict}. Meaning, we calculate z-scores separately for each trap. We then estimate a panel fixed effects model using the following regression specification, clustering standard errors at the county level to account for correlation induced

[9]Traps are located 1 to 2 km apart, and 0.25° correspond to approximately 25-30 km.
[10]For a discussion of transformations and similarity measures in ecology, see Keil et al. (2019).

by shared DSCI and unobserved county-level shocks:

$$z_{ict} = \sum_k (\beta_k^{(z)} D_{c(i)t}^k + \pi_k^{(z)} D_{c(i)t}^k \times HD_i) + \lambda_i + \delta_{st} + \gamma_{sm} + \theta_{g(i)} + \varepsilon_{ict}. \tag{2}$$

The coefficients of interest are $\beta_k^{(z)}$ and $\pi_k^{(z)}$, which capture the effect of drought conditions on the standardized moth trap outcome, across all traps, and the interaction of the DSCI bin with a dummy variable for a high-density trap, HD_i. Drought conditions are summarized as equal-interval bins by the county-year DSCI, as discussed in Sections 3 and 4.2. For completeness, and comparability with other specifications we introduce below, we include trap fixed effects, λ_i, to account for trap-site time invariant effects, even though the within-trap standardization already de-means each unit, which is equivalent to including a trap fixed effect. We flexibly control for time trends and seasonality by including state-by-year and state-by-month fixed effects, δ_{st} and γ_{sm}. To account for trap composition over time, as we explain in more detail below, we include grid-cell fixed effects, $\theta_{g(i)}$.

This specification has several advantages. First, standardizing expresses outcomes on a dimensionless scale where coefficients measure changes in units of historical variability. Second, z-scores handle zeros mechanically without requiring ad hoc adjustments such as adding a constant before taking logs. Third, the linear fixed effects framework is computationally convenient and easily accommodates multiple high-dimensional fixed effects.

Within-trap z-scores standardize each trap to its own historical mean and variance, which is helpful when studying how unusual is a specific year for a specific site. However, within-trap z-scores make cross-site and cross–trap-type comparisons hard: a one-unit change in the z-score can correspond to very different changes in raw counts across traps, and because the standard deviation is estimated from the full time series, the scaling itself partially absorbs the very drought variation we are interested in capturing.

Using z-scores involves a few more limitations. First, z-scores are not directly interpretable in terms of absolute moth abundance. Keil et al. (2019) elaborate precisely on

the interpretational challenges of standardized indices across heterogeneous units. Second, when many observations are zero—as is typical for delta traps in low-density or newly invaded areas—the standard deviation can be small and a change from zero to one moth can translate into a very large change in z_{ict}, driven more by discreteness and sampling noise than by a smooth underlying process. Finally, the linear-Gaussian approximation underlying the specification in Equation (2) is not formally consistent with the discrete, non-negative count nature of the data and can perform poorly in the presence of strong skewness or heteroskedasticity. O'Hara and Kotze (2010) provide a critique on the practice of forcing count data into Gaussian frameworks via transformations.

Station Entry, Exit & Spatial Non-Stationarity.— An additional challenge is that traps do not form a static spatial panel. Under the national Slow the Spread (STS) program, trap grids are redeployed annually along the moving invasion frontier, leading to systematic entry and exit of traps in space and time. Early cohorts of traps may be located in counties and forest types that never overlap with the locations of late cohorts, generating spatial non-stationarity in the sample. When trap cohorts follow the frontier into new areas, there may remain residual spatial correlation between drought conditions, background infestation risk, and unobserved landscape characteristics that are not fully captured by unit and time fixed effects.

In the within-trap z-score approach, trap fixed effects λ_i soak up time-invariant differences in mean catch at each trap location, but they do not, by themselves, ensure comparability across non-overlapping cohorts. If late cohorts are deployed in systematically different landscapes (e.g., more southerly, drier, or with different host composition) than early cohorts, the distribution of z_{ict} may drift over time, complicating the interpretation of the coefficients as a single, time-invariant effect of drought.

To address this, we augment the model with coarse spatial grid-cell fixed effects. Using trap coordinates, we assign each trap to a 0.25°-equivalent grid cell (approximately 25–30

km on a side) and include grid-cell indicators $\theta_{g(i)}$, where $g(i)$ denotes the spatial grid cell containing trap i. These spatial fixed effects help in two ways. First, they capture persistent landscape-level heterogeneity (e.g., ecoregions, soil types, forest composition) that may be correlated both with drought exposure and with moth population potential, thereby reducing omitted-variable bias when trap cohorts move into new areas.[11] Second, by anchoring traps within coarse spatial units that are observed over multiple years and across cohorts, they partially restore a notion of spatial overlap, making DSCI effects less sensitive to the particular sequence of frontier expansion. While grid-cell fixed effects are nested within trap fixed effects for traps that never move, the presence of multiple traps per cell and the movement of the frontier imply that they capture variation at a broader spatial scale than individual trap identifiers.

Preserve Count Data Structure Using Poisson Pseudo-ML.— To respect the discrete and often zero-inflated nature of trap counts, we instead model y_{ict} directly as a count outcome using a generalized linear model with a log link and Poisson pseudo-maximum likelihood (PPML).[12] The conditional mean is specified as

$$\mathbb{E}[y_{ict} \mid D^k_{c(i)t}, HD, \lambda_i, \delta_{st}, \gamma_{sm}, \theta_{g(i)}] = \exp\Big(\sum_k (\beta^{(P)}_k D^k_{c(i)t} + \pi^{(P)}_k D^k_{c(i)t} \times HD_i) + \lambda_i + \delta_{st} + \gamma_{sm} + \theta_{g(i)}\Big) \tag{3}$$

We estimate the regression specification in Equation (3) via PPML with multiple levels of fixed effects, as in Equation (2).

This approach offers several advantages. First, zeros are handled naturally: there is no need to add arbitrary constants or discard zero observations, and the log link works at the level of the conditional mean, not as a transformation of the dependent variable.[13]

[11] For the role of spatial fixed effects and random effects in ecological models, see Lee and Napier (2004).

[12] On PPML with high-dimensional fixed effects, see Correia et al. (2020). For related applications to insect counts and abundance estimation, see Hilborn and Mangel (2011).

[13] PPML is consistent under fairly general forms of overdispersion and heteroskedasticity; see Wooldridge

Second, coefficients on the DSCI bin indicators have a clear multiplicative interpretation: $\beta_k^{(P)}$ is the log incidence rate ratio for bin k relative to the omitted bin, so $\exp(\beta_k^{(P)}) - 1$ is the percentage change in expected moth count associated with that drought category, conditional on fixed effects. Third, the count nature and skewness of the data are explicitly acknowledged, aligning better with the biology of rare events and occasional local outbreaks. Fourth, modern PPML implementations allow for high-dimensional fixed effects without the need to estimate a large number of dummy coefficients directly. A potential concern with Poisson models is overdispersion relative to the Poisson variance assumption. However, PPML remains consistent for the conditional mean even under overdispersion, and robust variance estimators account for misspecification of the variance function.[14]

In summary, the within-trap z-score linear fixed effects model provides a convenient way to standardize heterogeneous trap series and to express results in standard deviation units, but at the cost of losing a direct link to absolute abundance and of imposing an approximate Gaussian framework on discrete, zero-heavy data. The PPML count model, by contrast, respects the count nature of the data, handles zeros naturally, and yields interpretable multiplicative effects of drought. Both approaches allow to flexibly incorporate multiple levels of fixed effects, including coarse spatial grid-cell effects that help address spatial non-stationarity due to the moving invasion frontier. Given these considerations, our preferred specification is the PPML approach in Equation (3), with count data as the dependent variable and high-dimensional fixed effects capturing trap-, time-, and space-specific heterogeneity.

Estimation Results for Drought & Moth Trap Catch.— The abundance of spongy moths caught in traps increases with higher levels of drought conditions in the county,

(1999) and Correia et al. (2020).

[14] See the discussion in Correia et al. (2020). In settings with extreme overdispersion or structural zeros, negative binomial or hurdle models can be considered, but they are less straightforward to combine with many fixed effects.

but only in the high-density traps (the milk cartons). In Figure 5, we report estimation results using the OLS and PPML estimators for the within-trap and unadjusted count data, respectively. For each, we report results from using either contemporaneous DSCI values, and up to two lags when calculating the mean DSCI values for the county, before constructing the DSCI bins.[15]

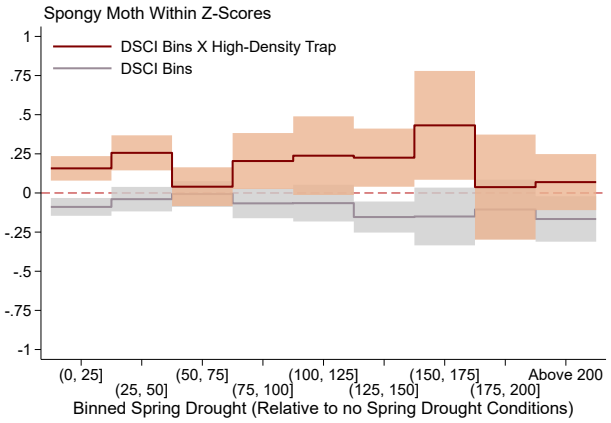
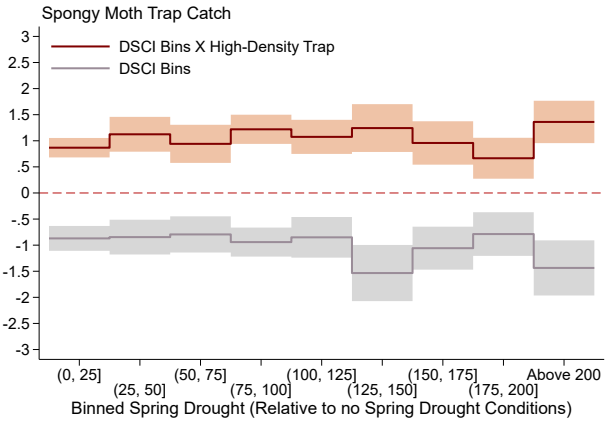
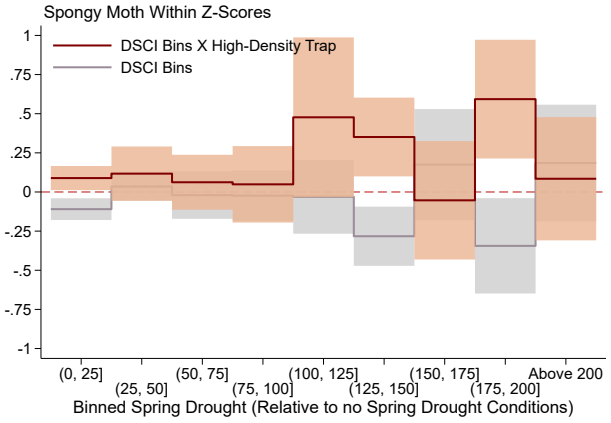
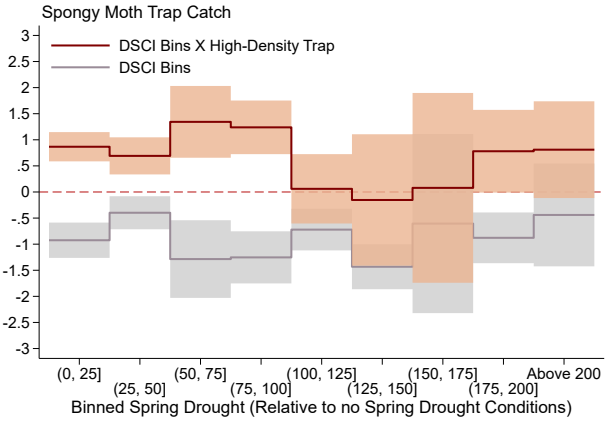
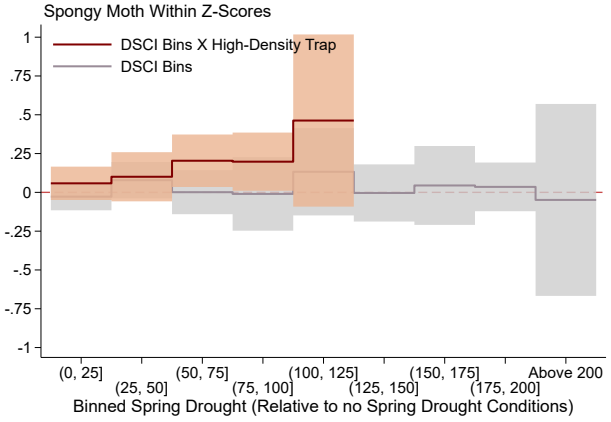
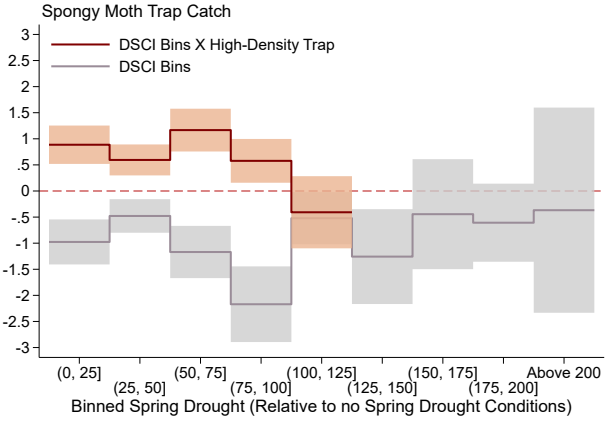
The non-interacted coefficients, on average, we recover no change in moth trap catch as a function of DSCI bins (gray line). For the PPML coefficients, if anything, we observed lower total catch during drought conditions. For high-density traps, however, we estimate positive coefficients that link higher moth trap catch with higher drought severity (maroon line). For example, for the within-trap z-score, using contemporaneous DSCI values only (Figure 5a), moth catch is higher by approximately 0.16 standard deviations.[16] For the unadjusted catch levels, using contemporaneous DSCI values only (Figure 5b), moth catch is approximately 183 percent larger when the county experiences drought conditions, relative to no drought conditions.[17] The exact magnitude and precision change across the two estimators, and the moving averages used to construct the DSCI bins, yet the fundamental pattern holds: more moths are caught in traps that are designed to capture high-moth densities when drought conditions are higher, relative to periods without any drought conditions.

[15] When using two lags of DSCI values, some high-density traps have no support for mean DSCI values above 125, and those bins are not estimated.

[16] Calculated using the mean for the linear combination of the first four coefficients of DSCI values ranging from 0 to 100, capturing 46.5 percent of the non-zero DSCI values in the moth trap sample.

[17] Calculated using the mean for the linear combination of the first four coefficients of DSCI values ranging from 0 to 100, capturing 46.5 percent of the non-zero DSCI values in the moth trap sample.

Figure 5: The Effects of Spring Drought Conditions on Spongy Moth Trap Catch

Within-Trap Z-Score (OLS)	Unadjusted Count Data (PPML)
(a) t	(b) t
	
(c) $t,\ t-1$	(d) $t,\ t-1$
	
(e) $t,\ t-1,\ t-2$	(f) $t,\ t-1,\ t-2$
	

Notes: Estimation results and 95 percent confidence intervals from estimating the specifications in Equations (2) and (3). Each plot reports results from one regression. The first row includes data using contemporaneous DSCI values, the second row uses the mean of contemporaneous and lagged DSCI values, and the third row uses contemporaneous and two years of lagged DSCI values. All regressions include trap ID, state-by-year, state-by-month, and grid-cell fixed effects.

5 Impacts of Defoliation on Property Values

Large defoliation events can reduce the amenity value of living near a forest. Not only is the view transformed from a green forest to a defoliated forest, but a moth outbreak event also includes two caterpillar-related disamenities: hairs from the caterpillars, and their excrement. The caterpillars grow long thin hairs that get carried away by wind, up to a mile, and can cause either skin irritation and rashes (Etkind et al. 1982; Tuthill et al. 1984), or respiratory difficulties if inhaled (Coleman et al. 2020). As the caterpillar digests large amounts of leafy material, they release their excrement, known as frass, creating sanitary and odor nuisances for homeowners (Bigsby et al. 2014; Coleman et al. 2020).

We use geocoded property transaction data, which we match with defoliation exposure that intersect with the location of the property, or is located in its vicinity. Our goal is to estimate the effect of defoliation events on residential property values, using the sale price of houses that sold during years with or without high defoliation levels. The analysis here assumes that defoliation acts as a sudden and unexpected shock to amenity values, lowering the WTP of buyers for houses that have recently been exposed to the effects of defoliation.

A moth defoliation event might only have a local effect in proximity to where the defoliation occurred. Some of the reduction in value might be transitory, as values return to their baseline levels after the defoliation event ends, while some of the devaluation might be permanent if potential homeowners know the area is in the moth infestation zone. The uncertain dynamics that lead to an outbreak provide as-good-as random variation, with respect to local housing markets, in the exposure to the disamenity shocks potentially caused by moth defoliation. This allows us to estimate a two-way-fixed-effects model that recovers the effect of nearby defoliation on the sale price using the following specification:

$$\log(\text{Sale Price})_{pvcsym} = \beta(\text{Defoliation Dummy})_{py} + \lambda_v + \delta_{sy} + \gamma_{sm} + \varepsilon_{vcsym} \tag{4}$$

Where the defoliation dummy is equal to one when the property is sold and is exposed to defoliation. For each property p, we define a set of dummy variables that measure exposure to defoliation in the year of the transaction y, as well as in years prior to the transaction. Exposure to defoliation could be due to direct overlap with, or due to adjacency to, a defoliated area. To reflect this, we define dummies that are equal to one if the centroid of the property intersects with a defoliated area, in either the year of the transaction, or up to a five-year lag. However, as there could be measurement error in both the centroid location, and the exact defoliation location, we measure the distance of the property to the defoliated area, and construct flexible distance bins from defoliation.

Each property, located in county subdivision v, county c, and state s, might be affected by real-estate dynamics operating at the neighborhood level, the state, and as part of real-business cycles. To account for these effects, we include county subdivision fixed effects, λ_v, which control for any neighborhood time-invariant characteristics. We allow for flexible time trends using state-by-year fixed effects, δ_{sy}, which capture within-state dynamics that would affect housing as well as national-level business cycles. We also allow for different seasonality by including state-by-month fixed effects, γ_{sm}. Finally, all unobserved heterogeneity is accounted for by the error term, ε_{vcsym}, which we cluster at the county subdivision level.

Effects of Direct Defoliation Exposure Across Samples.— Property values drop by more than two percent when defoliation occurs in the year they are sold. In Table 2, we report the estimation results from Equation (4) using a variety of sample restrictions. In the full sample, when we impose no other restrictions, we precisely estimate that a property whose location intersects with the defoliation map in the year it sells for, sells for 3.7 percent less (column 1). We observe virtually no change in the magnitude or precision of this effect if we restrict only to counties that ever experience defoliation in our sample (column 2), focus on rural areas by including only non-metro counties (column 3), or include up to and including small metro counties (column 4). If we restrict the sample only to ZCTAs that

experienced a defoliation event during our sample (column 5), we observe a smaller, but precisely estimated, decline of 2.3 percent in the sale price. Alternatively, if we restrict the sample of properties to be those that are within at least 5 km of a forest area (column 6), we recover a similar sample size and effect size as in the full sample (column 1). We interpret these results as evidence for the negative effect that defoliation has on property values. In Online Appendix Table A6, we demonstrate these results hold when flexibly accounting for observable property characteristics. In subsequent analysis, we also recover similar magnitudes of negative prices effects when we use a repeated sales sample.

Table 2.
The Effect of Direct Defoliation Exposure on Property Values

	Full Sample	County With Defoliation	Non-Metro Counties	Non-Large Metro Counties	ZCTAs With Defoliation	Within 5 km to Forest
	(1)	(2)	(3)	(4)	(5)	(6)
Defoliation	-0.037	-0.037	-0.040	-0.040	-0.023	-0.036
	(0.016)	(0.015)	(0.020)	(0.014)	(0.010)	(0.014)
R^2	0.320	0.308	0.281	0.300	0.278	0.310
n$\times 10^6$	79.39	48.96	8.8	28.44	24.55	70.58
Clusters	22,208	9,865	13,857	18,739	5,193	20,387

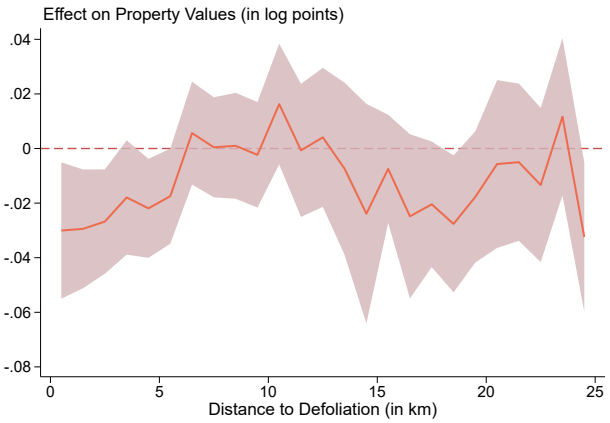
Notes: Estimation results from the specification in Equation (4). Each regression includes county sub-division, state-by-year, and state-by-month fixed effects. Each column reports results from using different sample restrictions. In column (1), we report results from the full sample. In column (2), we restrict the sample to counties that experienced defoliation in our sample. In column (3), we restrict to non-metro counties, using the 2023 Rural-Urban Continuum Codes 2024. In column (4), we add to the sample in (3) by including small metro counties. In column (5), we restrict results ZCTAs that experienced defoliation in our sample. In column (6), we restrict results to properties that are within 5 km of a forest. Standard errors are clustered at the county sub-division level.

Heterogeneity With Respect to Distance from Defoliation Exposure.— Defoliation exposure suppresses property values up to six km from the defoliated area. In Figure 6, we plot estimation results that use the distance to the nearest defoliated area. We run a regression of the logged sale price on a set of distance bin dummies, up to 25 km. The omitted category in these regressions are properties that are more than 25 km away from

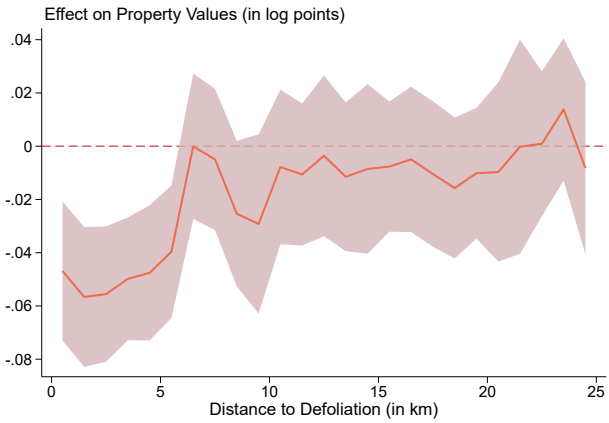
any defoliated area. We report results for two samples: counties that ever experienced defoliation, and non-large metro counties. We find that in both cases, the effect is concentrated in the distance bins that are close to defoliated areas, up to six km. The effect is larger, and decays more clearly, in the sample of non-large metro areas (Figure 6b). We use this finding to focus our attention in subsequent analysis on properties that are within six km of defoliation.

Figure 6: Effects on Property Values as a Function of Distance to Defoliation

(a) Counties With Defoliation



(b) Non-Large Metro Counties



Notes: We report coefficients and 95% confidence intervals from a regression of the logged sale price on distance bin dummies from defoliation. The omitted category are properties that are further than 25 km from any defoliated area in the year they are sold. In Panel (a), we use the sample of counties that ever had defoliation exposure during our sample period. In Panel (b), we use the sample that excludes large metro areas. Each regression includes county subdivision, state-by-year, and state-by-month fixed effects. Standard errors are clustered at the county subdivision.

Heterogeneity With Respect to Timing of Defoliation Exposure.— Contemporaneous exposure to defoliation has the largest negative effect, however, defoliation exposure in previous years can also have a negative effect on the sale price. In Table 3, we report estimates on how lagged exposure to defoliation affects the sale price of a property. We include lags from one year up to five years before the property sells. We either define defoliation exposure as the intersection between a property and the defoliated area, or, based on the results in Figure 6, as the property being up to six km from the defoliated area. We find that

contemporaneous exposure has the largest effect, but lags up to three years can still have a negative, meaningful, and precisely estimated effect on the sale price of the property. Once we include defoliation exposure lags for four and five years, we no longer find those have a meaningful or precisely estimated effect. In other words, as most spongy moth outbreaks last around three years, if we use the plausible scenario that the house is experiencing exposure to defoliation in the year of sale, and each of the previous two years as well, then the mean decline in the sale price is 5.9 percent (column 2).[18]

Decomposing the lagged defoliation exposure by season reveals that defoliation has the largest negative effect on sale price when the house is sold in a season when either spongy moth outbreak or defoliation are observed. We interacts contemporaneous and lagged defoliation exposure with seasonal dummies. For simplicity, we define winter as January, February, and March, and assign the remaining seasons to the remaining quarters of the year. In Table 4, we report that the negative effect of defoliation exposure increases from 3.7 percent (Table 2, column 2) to 5.5 or 4.2 percent for sales that take place during the spring and summer, when outbreaks and defoliation occur (Table 4, column 1). The effect for sales that happen in the fall, after defoliation has happened, is a 3.3 percent decline in sale price (Table 4, column 1), reflecting that even though most tree in our sample area shed their leaves in the fall, there is still a negative effect following exposure to defoliation. These estimates are precisely estimated, allowing us to reject the null hypothesis of no change in sale price after defoliation exposure at the one percent level. Exposure to defoliation suppresses sale prices even if it occurred in one year or two years before the house sold (Table 4, columns 2 and 3); however, the magnitude of the effects declines with longer lags. The negative effect on sale price does not appear for houses that sell in the winter (Table 4, columns 1-3), regardless of the lag structure. Combined, these results demonstrate that defoliation exposure has the largest negative impacts on the value of the property when defoliation is easy to observe

[18]There are 684,040 properties in the sample, reflecting 1.4 percent of the sample, that experience defoliation within six km in their sale year, as well as the two preceding years—amounting to three consecutive years of defoliation exposure.

Table 3.
Effects of Lagged Defoliation Exposure on Property Values

Panel A. Intersection With Defoliated Areas

	(1)	(2)	(3)	(4)	(5)
Defoliation_{t}	-0.034	-0.033	-0.032	-0.031	-0.029
	(0.014)	(0.014)	(0.014)	(0.015)	(0.015)
$\text{Defoliation}_{t\text{-}1}$	-0.025	-0.022	-0.022	-0.019	-0.017
	(0.014)	(0.013)	(0.013)	(0.013)	(0.014)
$\text{Defoliation}_{t\text{-}2}$		-0.026	-0.023	-0.024	-0.021
		(0.014)	(0.013)	(0.013)	(0.013)
$\text{Defoliation}_{t\text{-}3}$			-0.019	-0.017	-0.017
			(0.013)	(0.012)	(0.012)
$\text{Defoliation}_{t\text{-}4}$				0.000	0.003
				(0.012)	(0.011)
$\text{Defoliation}_{t\text{-}5}$					-0.009
					(0.010)
R^2	0.307	0.306	0.305	0.303	0.302

Panel B. Distance to Defoliated Areas $\leq$6 km

	(1)	(2)	(3)	(4)	(5)
Defoliation_{t}	-0.020	-0.019	-0.019	-0.020	-0.020
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
$\text{Defoliation}_{t\text{-}1}$	-0.019	-0.014	-0.013	-0.013	-0.013
	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
$\text{Defoliation}_{t\text{-}2}$		-0.026	-0.023	-0.023	-0.023
		(0.005)	(0.004)	(0.004)	(0.004)
$\text{Defoliation}_{t\text{-}3}$			-0.014	-0.012	-0.012
			(0.006)	(0.005)	(0.005)
$\text{Defoliation}_{t\text{-}4}$				-0.011	-0.011
				(0.009)	(0.007)
$\text{Defoliation}_{t\text{-}5}$					0.000
					(0.010)
R^2	0.308	0.308	0.308	0.308	0.308
n$\times 10^6$	48.96	48.96	48.96	48.96	48.96
Clusters	9,865	9,865	9,865	9,865	9,865

Notes: Estimation results from modifying the specification in Equation (4). We add lagged defoliation exposure, for either the intersection with defoliated areas (in Panel A), or with defoliation occurring up to six km from the property. Sample consists of properties in counties that ever had defoliation during our sample period. Each regression includes county sub-division, state-by-year, and state-by-month fixed effects. Standard errors are clustered at the county sub-division level.

or predict, but almost no effect in winter months, when defoliation events might not be as salient for buyers.

Using Repeated Sales to Account for Additional Property Characteristics.— To rule out that properties that are sold during defoliation events are different in terms of their characteristics, we restrict the sample to repeated sales, and include a property fixed effect. Specifically, we run a modified regression of the specification in Equation (4) where we add both a dummy for each sale order, as well as the property fixed effect. For each property, we use its unique identifier in the CoreLogic data to first count the number of sales in our sample. We then count how many defoliation events occurred within six km. We focus on the properties that sold at least once with defoliation exposure, and at least once without defoliation exposure. Such that each property was sold in both states. We exclude properties that sold more than 10 times during out sample period, which only accounts for 0.25 percent of the sample.

Exposure to defoliation has a negative effect on property values even when we account for property characteristics by using within-property variation. In Table 5, we first report results for the baseline specification, but restrict the sample to the repeated sales sample, as described above. We find a smaller estimated decline of 1.3 percent in property values (column 1), which we can reject the null hypothesis of no change in property values at the significance level of 10 percent. Then, we report results from the augmented specification, described above, for properties that sold exactly twice: once during defoliation, when it occurred within six km, and once when it did not. It is important to note that for some properties, defoliation exposure could happen at the first sale, and not in the second, while for other properties, the reverse is true. We report a precisely estimated decline of 3.1 percent in this repeated sales sample (column 2), similar to the results reported in the non-repeated sales sample (Table 2). Next, we extend the number of repeated sales to three, four, and five sales. We find negative effects for all samples, but of smaller magnitude (columns 3, 4,

Table 4.
Effects of Defoliation Exposure on Property Values by Season

	(1)	(2)	(3)
Winter_{t}	-0.006	-0.007	-0.007
	(0.019)	(0.018)	(0.018)
Spring_{t}	-0.055	-0.049	-0.048
	(0.017)	(0.015)	(0.015)
Summer_{t}	-0.042	-0.039	-0.038
	(0.016)	(0.015)	(0.015)
Fall_{t}	-0.033	-0.033	-0.033
	(0.016)	(0.015)	(0.015)
Winter_{t-1}		-0.004	-0.006
		(0.017)	(0.015)
Spring_{t-1}		-0.050	-0.044
		(0.015)	(0.014)
Summer_{t-1}		-0.030	-0.026
		(0.014)	(0.013)
Fall_{t-1}		-0.008	-0.007
		(0.018)	(0.016)
Winter_{t-2}			0.004
			(0.017)
Spring_{t-2}			-0.048
			(0.015)
Summer_{t-2}			-0.035
			(0.013)
Fall_{t-2}			-0.016
			(0.017)
R^2	0.306	0.306	0.306

Notes: Estimation results from modifying the specification in Equation (4). We interact the dummy for defoliation exposure with seasonal dummies, and lags, up to two years. Sample consists of properties in counties that ever had defoliation during our sample period. Each regression includes county sub-division, state-by-year, and state-by-month fixed effects. Standard errors are clustered at the county sub-division level.

and 5). Finally, we report the estimation results for all the repeated sales in our sample, which include properties that sold as many as 10 times during our sample period. We find a decline in the sale price of 1.4 percent (column 6); however, we can only reject the null hypothesis of no change in sale price at the significance level of 10 percent.

Table 5.
The Effect of Defoliation Using Repeated Property Sales

Repeated Sales:	≤10	=2	≤3	≤4	≤5	≤10
	(1)	(2)	(3)	(4)	(5)	(6)
Defoliation	-0.013	-0.031	-0.019	-0.015	-0.014	-0.014
	(0.008)	(0.008)	(0.007)	(0.007)	(0.007)	(0.008)
Sale Order FEs		X	X	X	X	X
Property FEs		X	X	X	X	X
R^2	0.256	0.716	0.644	0.601	0.575	0.546
n$\times 10^6$	14.69	3.08	7.02	10.26	12.39	14.69
Clusters	6,723	6,583	6,678	6,710	6,718	6,723

Notes: Estimation results from modifying the specification in Equation (4). Each regression includes county sub-division, state-by-year, and state-by-month fixed effects. The regression in column (1) uses the baseline specification in Equation (4), but restrict the sample to the repeated sales sample for properties that are in counties that ever had defoliation occur during out sample period. Regression in columns (2) to (6), add sale order and property fixed effects. Exposure to defoliation is defined as defoilation within six km in the year the property sold. Standard errors are clustered at the county sub-division level.

Exploring the Effect of First-Time Defoliation Exposure.— As the spongy moth expands in range and outbreaks occur over time, some areas will experience a defoliation event for the first time. The effect of a "defoliation shock," in the sense that not only was defoliation not expected, but the fact that defoliation could occur was also not salient. We test for whether defoliation exposure that is new in a housing market has a differential effect relative to any defoliation exposure. Explicitly, we define for each county subdivision the first year we observe a defoliation event. Then, we define a new dummy variable that is equal

to one if the property is sold in the year of a defoliation shock. We estimate regressions that modify the specification in Equation (4) by adding an interaction term for defoliation exposure and a defoliation shock.

Defoliation exposure that occurs for the first time in a county subdivision meaningfully reduces the sale price beyond the suppressing effect on price of defoliation alone. In Table 6, we report estimation results for the samples where counties ever had a defoliation event, and all non-large metro counties. Overall, defoliation exposure continues to have a negative effect on the sale price of a property of two to four percent. For sales that happen when the defoliation occurs for the first time in the county subdivision, the sale price is 2.5 to 5.3 percent lower, depending on wether we estimate direct defoliation exposure, or exposure to defoliation within six km of the property.

Dynamic Hedonic Models and Time-Varying Environmental Disamenities.— In the analysis we have presented up to here, we have abstracted away from considering whether households are forward-looking when purchasing houses, and how their expectation of the change to the forest amenity will affect their willingness to pay (WTP). In this section, we report a diagnostic measure that allows us to determine whether our estimated negative effects of defoliation on sale price are an underestimate or an overestimate, and by what order of magnitude. To do so, we follow Bishop and Murphy (2019) who develop a framework for valuing *time-varying* local amenities in a hedonic model when households are forward-looking and face moving costs. Their key insight is that the traditional static Rosen (1974) hedonic model can yield biased estimates of WTP when the attribute of interest evolves over time, because current prices reflect both current amenity levels and expectations about their future path. Instead of estimating a fully dynamic structural model—which requires data on the housing choices of households over time— they derive a simpler forward-looking hedonic model and characterize when static and dynamic approaches deliver similar or different implicit prices for a given time-varying attribute.

Table 6.
The Effect of Defoliation Exposure Shocks on Property Values

	County With Defoliation		Non-Large Metro	
	(1)	(2)	(3)	(4)
Direct Defoliation	-0.025		-0.031	
	(0.012)		(0.015)	
Direct Defoliation×Defoliation Shock	-0.053		-0.042	
	(0.044)		(0.022)	
Defoliation≤6 km		-0.020		-0.041
		(0.007)		(0.009)
Defoliation≤6 km×Defoliation Shock		-0.036		-0.025
		(0.011)		(0.014)
R^2	0.308	0.308	0.300	0.300
n×10^6	48.96	48.96	28.44	28.44
Clusters	9,865	9,865	18,739	18,739

Notes: Estimation results from a modification to the specification in Equation (4). We interact the defoliation exposure dummy with a dummy varialbe for whether the property sells in the first year that the county sub-division experiences a defoliation event. Each regression includes county sub-division, state-by-year, and state-by-month fixed effects. In columns (2) and (3) , we restrict the sample to counties that experienced defoliation in our sample. In column (3) and (4), we restrict to non-metro and small-metro counties, using the 2023 Rural-Urban Continuum Codes 2024. Standard errors are clustered at the county sub-division level.

The standard model assumes that households are myopic: they choose a location by trading off current housing and neighborhood characteristics against current price, so the marginal price of an attribute is interpreted directly as marginal WTP for that attribute in the current period. Bishop and Murphy (2019) instead allow households to be forward-looking and to internalize the discounted stream of future amenity levels and future moving decisions, leading to a forward-looking housing price function $\tilde{r}(\bar{x})$ that depends on the expected *path* of the amenity, not just its current level x. They show that when the amenity is evolving over time, the static hedonic price function $r(x)$ will generally differ from the forward-looking price function $\tilde{r}(\bar{x})$, so the static model can either under- or over-estimate WTP depending on the transition properties of the amenity.

For the spongy moth setting, these issues arise because defoliation is not a permanent one-time shock but an evolving environmental disamenity whose severity, spatial extent, and perceived future risk change over time. Remote-sensing evidence shows that canopy loss and productivity can be substantial during outbreak years but that key forest indicators, such as leaf-area-index and gross primary productivity, often partially or largely recover in subsequent years, consistent with non-permanent damage (Latifovic et al. 2024). In such a context, current house prices in defoliated areas are likely to reflect not only the current level of the disamenity, but also expectations about how quickly the forest will recover, how likely repeat outbreaks are, and how costly it is to move away from affected locations.

The central object in the Bishop and Murphy (2019) analysis is the *transition process* for the amenity: how current levels map into expected future levels. They emphasize two key dimensions: (i) the overall trend in the amenity over a horizon of T years (e.g., improving, deteriorating, or trendless), and (ii) whether the cross-sectional distribution of the amenity across locations is *mean-reverting*, *mean-diverging*, or neither. When the amenity is trending uniformly across all locations (for example, air quality that improves at a similar rate everywhere), the static and forward-looking hedonic models can yield similar estimates of marginal WTP, even though the amenity is time-varying. In contrast, when locations with

low current amenity experience faster improvements (mean reversion) or slower improvements (mean divergence) than high-amenity locations, static and forward-looking models generally produce different price functions and implicit prices.

Bishop and Murphy define mean reversion as a case where the expected future amenity level $\bar{x}(x_{i,t})$ is a linear contraction of the current level,

$$\bar{x}(x_{i,t}) = \phi + \gamma x_{i,t}, \quad \gamma < 1, \tag{5}$$

This implies that low-amenity locations are expected to *catch up* to high-amenity locations, while mean divergence would instead feature $\gamma > 1$ and widening cross-sectional gaps. Under mean reversion, the forward-looking hedonic price function is everywhere steeper than the static price function, because a given difference in current amenity implies a smaller difference in expected future amenity; thus, the static model tends to *understate* the marginal WTP for the amenity.

Application to Spongy Moth Defoliation.— In the spongy moth defoliation setting, the relevant amenity is a bundle of forest-related attributes: mature canopy cover, shade, privacy, and associated ecological services around the residence.For at least for some forest types and outbreak regimes, the amenity loss from defoliation can be *partially mean reverting*: canopy and aesthetic quality recover over time, and severely defoliated locations may improve faster than lightly defoliated ones as the outbreak subsides. In other words, during an outbreak, locations are more different than each other than in the short few year after an outbreak, when forests are largely able to recover from the defoliation.

Mapping this to Bishop and Murphy's framework, mean reversion in spongy moth damage would arise if the expected future forest condition around a house is a contraction of the current loss: heavily defoliated areas today are expected to recover relatively more (in absolute or percentage terms) than lightly affected areas. This could occur because of biological

recovery (trees refoliating and regenerating after the outbreak), targeted management (e.g., suppression efforts or replanting in severely affected stands), or changes in host composition that reduce future defoliation risk. In that case, the cross-sectional dispersion in forest amenity would shrink over time, satisfying a mean-reversion condition analogous to $\gamma < 1$ in Bishop and Murphy's notation. Under this scenario, the static hedonic regression of prices on contemporaneous defoliation would tend to *understate* households' forward-looking WTP to avoid defoliation, because it ignores that a marginal increase in current amenity (less defoliation) also implies a more favorable future path.

By contrast, defoliation would be *mean-diverging* if heavily defoliated areas are expected to stay worse or deteriorate further relative to lightly defoliated areas. This could occur if repeated outbreaks cluster in the same landscapes, if tree mortality is much higher in some stands, or if interacting stressors (e.g., drought, pests, and disease) lead to long-term forest decline in already heavily impacted neighborhoods. In such a case, the cross-sectional gaps in forest amenity widen over time (analogous to $\gamma > 1$), and the static hedonic model would tend to *overstate* the forward-looking marginal WTP to avoid an additional unit of defoliation.

In the spirit of Bishop and Murphy (2019), the evolution of spongy moth damage can be represented with an autoregressive process in which current defoliation in county c and year t, $\text{Defoliation}_{c,t}$, depends on its lagged value, $\text{Defoliation}_{c,t-1}$, and a time trend, estimated using one regression, with a coefficient per each commuting zone, m, according to an AR(1) model:[19]

$$\text{Defoliation}_{c,t} = \rho_1 \, \text{Defoliation}_{c,t-1} \times \kappa_m + \rho_2 t \times \kappa_m + \lambda_c + \varepsilon_{c,t} \tag{6}$$

where ρ_1 captures the persistence of defoliation over time. In their framework, forward-looking households care about the discounted average of future amenity levels, so the ex-

[19] We use a sample of 578 counties that had spongy moth declared established by 2010—to allow for enough time to estimate the AR1 process. The 578 counties reside in 104 commuting zones.

pected average future defoliation around a property can be written as

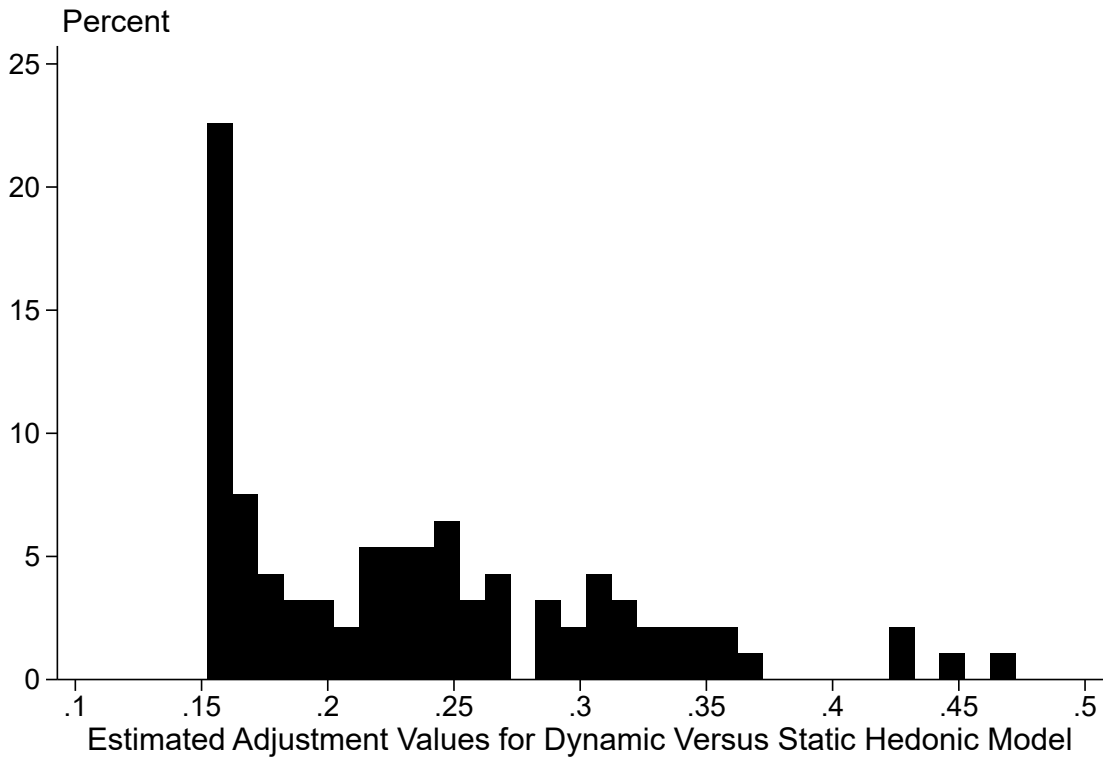
$$\bar{x}(\text{Defoliation}_{c,t}) = \phi + \gamma\, \text{Defoliation}_{c,t} \tag{7}$$

where γ summarizes how current conditions map into the discounted stream of future conditions.

$$\gamma = \frac{\sum_{s=1}^{T} \beta^{s-1} \rho_1^{s-1}}{\sum_{s=1}^{T} \beta^{s-1}} \tag{8}$$

Following the choice of parameters in Bishop and Murphy (2019) as a benchmark, we use a relatively high discount factor, $\beta = 0.95$, and a finite horizon $T = 7$, representing the mean turnover of households in residential properties, the period over which households value local environmental conditions. The results in Figure 7 capture $\hat{\gamma}$ values, estimated at the commuting zone level by pooling panel data from multiple counties, that are well below one—indicating strong mean reversion in defoliation, at the commuting zone level. In this framework, the static coefficients reported in the analysis can be interpreted as underestimating the full suppressing effect of defoliation exposure on housing prices.

Figure 7: Estimation of Dynamic Versus Static Adjustment Values



Notes: Distribution of estimated γ values in Equation 8 using the AR1 process in Equation (6). Values below one indicate that the static hedonic model is likely underestimating the effects of defoliation on the sale price relative to a fully dynamic model of housing location choice. See text for more details.

6 Conclusions

This paper documents two empirical findings on spongy moth defoliation events at a large scale, across the Northeastern US, using county-level and detailed property transaction data. We corroborate previous plot-level and lab experiments that documented that the fungal pathogen that plays a large role in preventing a spongy moth outbreak fails during warm and dry spring conditions. This raises a concern that future weather conditions will result in a more favorable environment for the spongy moth as a result of climate change. We document a reduction in the sale price of residential properties of at least two percent. The negative effect on the sale price can compound if defoliation occurs around the property for several consecutive years, and is more severe when defoliation is unprecedented in the local housing market. These findings highlight that climate change impacts can operate through an ecosystem functioning channel. The policy implication of these results is that over time, more housing capital will risk devaluation as the spongy moth expands its range, and climate change makes large outbreaks of the moth more likely.

References

Banerjee, Abhijit, Esther Duflo, Gilles Postel-Vinay, and Tim Watts. 2010. "Long-Run Health Impacts of Income Shocks: Wine and Phylloxera in Nineteenth-Century France." *The Review of Economics and Statistics* 92 (4): 714–728.

Berdugo, Miguel, Manuel Delgado-Baquerizo, Santiago Soliveres, Rocı o Hernández-Clemente, Yanchuang Zhao, Juan J Gaitán, Nicolas Gross, et al. 2020. "Global ecosystem thresholds driven by aridity." *Science* 367 (6479): 787–790.

Bigsby, Kevin M, Mark J Ambrose, Patrick C Tobin, and Erin O Sills. 2014. "The cost of gypsy moth sex in the city." *Urban Forestry & Urban Greening* 13 (3): 459–468.

Bishop, Kelly C, and Alvin D Murphy. 2019. "Valuing time-varying attributes using the hedonic model: When is a dynamic approach necessary?" *The Review of Economics and Statistics* 101 (1): 134–145.

Campbell, Robert W, and Ronald J Sloan. 1977. "Natural Regulation of Innocuous Gypsy Moth Populations." *Environmental Entomology* 6 (2): 315–322.

Coleman, Tom W, Laurel J Haavik, Chris Foelker, and Andrew M Liebhold. 2020. "U.S. Department of Agriculture • Forest Service." *Forest Insect & Disease Leaflet* 162:1–20.

Colwell, Robert K, Robert R Dunn, and Nyeema C Harris. 2012. "Coextinction and Persistence of Dependent Species in a Changing World." *Annual Review of Ecology, Evolution, and Systematics* 43:183–203.

Correia, Sergio, Paulo Guimarães, and Thomas Zylkin. 2020. "Fast Poisson Estimation with High-Dimensional Fixed Effects." *The Stata Journal* 20 (1): 95–115. https://scorreia.com/software/ppmlhdfe/.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken. 2014. "What Do We Learn from the Weather? The New Climate-Economy Literature." *Journal of Economic Literature* 52 (3): 740–798.

Druckenmiller, Hannah. 2020. "Estimating an economic and social value for healthy forests: Evidence from tree mortality in the American West." *Working Paper.*

Dwyer, Greg, Jonathan Dushoff, Joseph S Elkinton, and Simon A Levin. 2000. "Pathogen-Driven Outbreaks in Forest Defoliators Revisited: Building Models from Experimental Data." *The American Naturalist* 156 (2): 105–120.

Dwyer, Greg, and Joseph S Elkinton. 1993. "Using Simple Models to Predict Virus Epizootics in Gypsy Moth Populations." *The Journal of Animal Ecology* 62 (1): 1–11.

Etkind, P H, T M Odell, A T Canada, S K Shama, A M Finn, and R Tuthill. 1982. "The gypsy moth caterpillar: a significant new occupational and public health problem." *Journal of occupational medicine.: official publication of the Industrial Medical Association* 24 (9): 659–662.

Field, Christopher B, David B Lobell, Halton A Peters, and Nona R Chiariello. 2007. "Feedbacks of Terrestrial Ecosystems to Climate Change." *Annual Review of Environment and Resources* 32 (1): 1–29.

Frank, Eyal G. 2024. "The Economic Impacts of Ecosystem Disruptions: Costs From Substituting Biological Pest Control." *Science* 385:eadg0344.

Frank, Eyal G, and Anant Sudarshan. 2024. "The Social Costs of Keystone Species Collapse: Evidence From The Decline of Vultures in India." *American Economic Review* 114 (10): 3007–3040.

Hajek, Ann E. 1997. "Fungal and Viral Epizootics in Gypsy Moth (Lepidoptera: Lymantriidae) Populations in Central New York." *Biological control: theory and applications in pest management* 10 (1): 58–68.

Hajek, Ann E, Andrea L Diss-Torrance, Nathan W Siegert, and Andrew M Liebhold. 2021. "Inoculative Releases and Natural Spread of the Fungal Pathogen Entomophaga maimaiga (Entomophthorales: Entomophthoraceae) into U.S. Populations of Gypsy Moth, Lymantria dispar (Lepidoptera: Erebidae)." *Environmental entomology* 50 (5): 1007–1015.

Hajek, Ann E, and Jørgen Eilenberg. 2018. *Natural Enemies: An Introduction to Biological Control.* Cambridge University Press.

Hilborn, Ray, and Marc Mangel. 2011. "Estimating population abundance from sparse trap counts: an introduction with applications to insects." *Ecological Applications* 21 (8): 2849–2860. https://doi.org/10.1890/10-2117.1.

Hsiang, Solomon, and Robert E Kopp. 2018. "An Economist's Guide to Climate Change Science." *The Journal of Economic Perspectives* 32 (4): 3–32.

Johnson, Derek M, Andrew M Liebhold, Ottar N Bjornstad, and Michael L Mcmanus. 2005. "Circumpolar variation in periodicity and synchrony among gypsy moth populations." *The Journal of animal ecology* 74 (5): 882–892.

Jones, Benjamin A, and Shana M McDermott. 2018. "Health Impacts of Invasive Species Through an Altered Natural Environment: Assessing Air Pollution Sinks as a Causal Pathway." *Environmental & Resource Economics* 71 (1): 22–43.

Keil, Petr, Walter Jetz, and Katrin Böhning-Gaese. 2019. "Z-scores unite pairwise indices of ecological similarity and beta diversity." *Ecosphere* 10 (4): e02616. https://doi.org/10.1002/ecs2.2616.

Kolstad, Charles D, and Frances C Moore. 2020. "Estimating the Economic Impacts of Climate Change Using Weather Observations." *Review of Environmental Economics and Policy* 14 (1): 1–24.

Kroeker, Kristy J, and Eric Sanford. 2022. "Ecological Leverage Points: Species Interactions Amplify the Physiological Effects of Global Environmental Change in the Ocean." *Annual review of marine science* 14:75–103.

Kyle, C H, J Liu, J Wang, R Kotamarthi, and G Dwyer. 2022. "Eco-climate models predict severe effects of climate change on a bio-control agent of an invasive forest pest." *Working Paper.*

Kyle, Colin H, Jiawei Liu, Molly E Gallagher, Vanja Dukic, and Greg Dwyer. 2020. "Stochasticity and Infectious Disease Dynamics: Density and Weather Effects on a Fungal Insect Pathogen." *The American Naturalist* 195 (3): 504–523.

Lange, Fabian, Alan L Olmstead, and Paul W Rhode. 2009. "The Impact of the Boll Weevil, 1892–1932." *The journal of economic history* 69 (3): 685–718.

Latifovic, Rasim, M. Altaf Arain, Alemu Gonsamo, and M. and others Hussain. 2024. "Assessment of spongy moth infestation impacts on forest productivity and carbon loss using Sentinel-2 satellite remote sensing and eddy covariance flux data." *Carbon Balance and Management* 19 (1): 1–22. https://doi.org/10.1186/s13021-024-00242-0. https://pmc.ncbi.nlm.nih.gov/articles/PMC11093731/.

Lawler, Joshua J, Sarah L Shafer, Denis White, Peter Kareiva, Edwin P Maurer, Andrew R Blaustein, and Patrick J Bartlein. 2009. "Projected climate-induced faunal change in the Western Hemisphere." *Ecology* 90 (3): 588–597.

Lee, J., and J. Napier. 2004. "The Spatial Random Effects and the Spatial Fixed Effects Model: Theory and Practical Implications." *Journal of Agricultural, Biological, and Environmental Statistics* 9 (1): 1–17.

Letourneau, Deborah K, Julie A Jedlicka, Sara G Bothwell, and Carlo R Moreno. 2009. "Effects of Natural Enemy Biodiversity on the Suppression of Arthropod Herbivores in Terrestrial Ecosystems." *Annual review of ecology, evolution, and systematics* 40 (1): 573–592.

Leuschner, William A, John A Young, Sally A Waldon, and F William Ravlin. 1996. "Potential Benefits of Slowing the Gypsy Moth's Spread." *Southern Journal of Applied Forestry* 20 (2): 65–73.

Liebhold, A, J Elkinton, D Williams, and R-M Muzika. 2000. "What causes outbreaks of the gypsy moth in North America?" *Population Ecology* 42 (3): 257–266.

Liebhold, Andrew, Victor Mastro, and Paul W Schaefer. 1989. "Learning from the legacy of Leopold Trouvelot." *Bulletin of the ESA* 35 (2): 20–22.

Liebhold, Andrew M, Joel A Halverson, and Gregory A Elmes. 1992. "Gypsy Moth Invasion in North America: A Quantitative Analysis." *Journal of biogeography* 19 (5): 513–520.

Liebhold, Andrew M., Kim W. Thorpe, and Deborah G. McCullough. 2020. "Bounds on Absolute Gypsy Moth (*Lymantria dispar dispar*) Population Density as Derived from Counts in Single Milk Carton Traps." *Environmental Entomology* 49 (5): 1056–1067. https://doi.org/10.1093/ee/nvaa089. https://pmc.ncbi.nlm.nih.gov/articles/PMC7601369/.

Liu, Jiawei, Colin Kyle, Jiali Wang, Rao Kotamarthi, William Koval, Vanja Dukic, and Greg Dwyer. 2025. "Climate change drives reduced biocontrol of the invasive spongy moth." *Nature Climate Change* 15 (2): 210–217.

Lu, Han, Stephan Heblich, Christopher Timmins, and Yanos Zylberberg. 2026. "Cool Cities: The Value of Green Infrastructure," *American Economic Review: Insights.*

Michael, Mcmanus a, and – György Csóka. 2007. "History and Impact of Gypsy Moth in North America and Comparison to Recent Outbreaks in Europe." *Acta Silvatica et Lignaria Hungarica* 3:47–64.

Minnesota Department of Agriculture. 2024a. *MDA to Set Over 20,000 Traps to Survey for Spongy Moth.* https://www.mda.state.mn.us/mda-set-over-20000-traps-survey-spongy-moth-0. Press release, Accessed 2026-03-10.

———. 2024b. *Spongy Moth Survey.* https://www.mda.state.mn.us/mda-set-over-20000-traps-survey-spongy-moth-0. Press release on delta and milk-carton trap deployment, Accessed 2026-03-10.

North Carolina Department of Agriculture and Consumer Services. 2023. *Spongy Moth Trapping.* https://www.ncagr.gov/divisions/plant-industry/plant-protection/entomological-services/spongy-moth. Accessed 2026-03-10.

O'Hara, Robert B., and D. Johan Kotze. 2010. "Do not log-transform count data." *Methods in Ecology and Evolution* 1 (2): 118–122. https://doi.org/10.1111/j.2041-210X.2010.00021.x.

O'Rourke, Kevin. 1994. "The Economic Impact of the Famine in the Short and Long Run." *The American economic review* 84 (2): 309–313.

PRISM Climate Group, Oregon State University. 2014. *Parameter-elevation Regressions on Independent Slopes Model.*

Raynor, Jennifer L, Corbett A Grainger, and Dominic P Parker. 2021. "Wolves make roadways safer, generating large economic returns to predator conservation." *Proceedings of the National Academy of Sciences of the United States of America* 118 (22).

Reilly, James R, Ann E Hajek, Andrew M Liebhold, and Ruth Plymale. 2014. "Impact of Entomophaga maimaiga (Entomophthorales: Entomophthoraceae) on outbreak gypsy moth populations (Lepidoptera: Erebidae): the role of weather." *Environmental entomology* 43 (3): 632–641.

Rosen, Sherwin. 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *The Journal of Political Economy* 82 (1): 34–55.

Sharov, A A, D Leonard, A M Liebhold, E A Roberts, and W Dickerson. 2002. ""Slow The Spread": A National Program to Contain the Gypsy Moth." *Journal of Forestry* 100 (5): 30–36.

Slow the Spread Foundation, Inc. 2024. *Slow the Spread (STS) of the Spongy Moth Program.* https://www.slowthespread.org. Accessed 2026-03-10.

Tuthill, R W, A T Canada, K Wilcock, P H Etkind, T M O'Dell, and S K Shama. 1984. "An epidemiologic study of gypsy moth rash." *American journal of public health* 74 (8): 799–803.

U.S. Department of Agriculture, Economic Research Service. 2024. *Rural-Urban Continuum Codes.*

USDA APHIS. 2012. *Spongy Moth Program Manual.* https://www.aphis.usda.gov/sites/default/files/gypsy_moth.pdf. USDA Animal and Plant Health Inspection Service, Accessed 2026-03-10.

———. 2014. *Approved Survey Method for Lymantria dispar: Delta and Milk Carton Traps.* https://pest.ceris.purdue.edu/services/approvedmethods/sheet.php?v=1616. National Agricultural Pest Information System (NAPIS) method sheet, Accessed 2026-03-10.

———. 2023. *Spongy Moth Survey in Progress.* https://www.aphis.usda.gov/sites/default/files/doorhanger-spongy-moth-survey.pdf. Door hanger explaining delta and milk-carton traps, Accessed 2026-03-10.

USGS. 2019. *National Land Cover Database.*

Ward, Samuel F, Brian H Aukema, Songlin Fei, and Andrew M Liebhold. 2020. "Warm temperatures increase population growth of a nonnative defoliator and inhibit demographic responses by parasitoids." *Ecology,* e03156.

Wisconsin Department of Agriculture, Trade and Consumer Protection. 2026. *Spongy Moth Traps and Egg Mass Surveys.* https://datcp.wi.gov/Pages/Programs_Services/SMSurvey.aspx. Accessed 2026-03-10.

Woods, S A, and J S Elkinton. 1987. "Biomodal patterns of mortality from nuclear polyhedrosis virus in gypsy moth (Lymantria dispar) populations." *Journal of invertebrate pathology* 50 (2): 151–157.

Wooldridge, Jeffrey M. 1999. "Distribution-free estimation of some nonlinear panel data models." *Journal of Econometrics* 90 (1): 77–97.

Xing, Jianwei, Zhiren Hu, Fan Xia, Jintao Xu, and Eric Zou. 2023. "Urban Forests: Environmental Health Values and Risks." *NBER WP 31554,* Working Paper Series, no. 31554.

Online Appendix

A Additional Results

A.1 Defoliation-Weather Results Using Defoliation Share Instead of Defoliation Rate

In the main text, in Section 4, we report results using the defoliation rate as the outcome, and either deciles or bins for the DSCI values. Here, we repeat the analysis using the share of area in the county that is defoliated, and normalize it by the area of the county instead of the forested area. In Tables A1-A4, we report weighted and unweighted results, using no lags and up to two years lagged values in the mean DSCI values. Qualitatively, the results are similar. Higher drought conditions are correlated with more defoliation. When we use two lagged years to calculate the mean DSCI values, before splitting them into bins, we are left with no observations that have DSCI values above 200. That coefficient is not esimtated in those regressions (columns 3 and 6 of Tables A3 and A4

A.2 Assessing Sensitivity to Specification Changes

In the main text, we present results in Table 2 that use county subdivision, state-by-year, and state-by-month fixed effects. Here, we focus on the sample of non-large metro counties (column 4 in Table 2) and change the specification to absorb more granular spatial or temporal variation. By and large, the results appear to be insensitive to the exact combination chosen. In Table, A5, we report the baseline specification in column (1). Then, we change the spatial fixed effects and include either ZCTA fixed effects (column 2), or EPA hexagon fixed effects instead of the county subdivision fixed effects.[20] We allow the time fixed effect to be more flexible by allowing the seasonality to change, by including state-by-year-by-month

[20]EPA hexagon split the US into 40 km sized hexagons. We match each property to its hexagon id based on its coordinates.

Table A1.
The Effects of Spring Drought Conditions on Defoliation

Weighted by:	Equally-Weighted			Susceptible Acres		
Averaging t to $t-s$:	$s=0$	$s=1$	$s=2$	$s=0$	$s=1$	$s=2$
	(1)	(2)	(3)	(4)	(5)	(6)
Drought Decile 1	0.98	2.66	5.12	1.08	4.05	4.67
	(1.31)	(1.42)	(1.90)	(1.78)	(1.77)	(2.09)
Drought Decile 2	1.36	2.75	6.98	0.81	2.45	7.12
	(1.81)	(1.57)	(2.36)	(2.45)	(1.96)	(2.60)
Drought Decile 3	2.43	7.10	7.40	1.83	8.83	7.09
	(1.73)	(2.37)	(3.30)	(2.50)	(3.46)	(3.20)
Drought Decile 4	3.88	7.40	10.33	4.39	8.48	12.59
	(2.62)	(2.75)	(3.18)	(3.65)	(2.95)	(3.73)
Drought Decile 5	18.80	8.69	7.50	19.70	11.24	7.12
	(5.96)	(2.39)	(2.77)	(6.29)	(3.22)	(2.76)
Drought Decile 6	5.43	8.42	9.92	6.85	9.12	8.36
	(2.56)	(3.06)	(3.09)	(3.52)	(3.80)	(3.83)
Drought Decile 7	3.64	12.24	11.81	3.19	12.44	10.64
	(3.29)	(3.03)	(2.87)	(4.40)	(3.85)	(3.53)
Drought Decile 8	9.27	9.99	12.38	9.45	10.33	12.18
	(3.37)	(2.91)	(3.36)	(4.36)	(4.04)	(3.98)
Drought Decile 9	7.41	8.94	9.87	3.68	8.94	9.45
	(3.50)	(2.42)	(3.24)	(4.91)	(3.29)	(4.33)
Drought Decile 10	8.68	2.85	4.10	6.47	0.01	1.09
	(3.71)	(3.39)	(3.62)	(4.44)	(4.28)	(4.44)
R^2	0.477	0.476	0.476	0.526	0.525	0.525
N	9,198	9,198	9,198	9,198	9,198	9,198
Clusters	511	511	511	511	511	511

Notes: Estimation results for the specification in Equation (1). The outcome is the defoliation rate (defoliated acres relative to the mean level of forest acres in the county). We regress the defoliation rate on spring drought conditions, as captured by the Drought Severity and Coverage Index (DSCI). We split the index into deciles, conditional on non-zero values of the DSCI. The omitted category is no drought conditions (DSCI=0). Each regression includes county and state-by-year fixed effects. Sample includes all counties that had the spongy moth classified as established by 2000. Sample spans 2002 to 2019. Observations are either equally-weighted (columns 1 to 3) or weighted by the number of susceptible forest acres (columns 4 to 6). Standard errors are clustered at the county level.

Table A2.
The Effects of Spring Drought Conditions on Defoliation

Weighted by:	Equally-Weighted			Susceptible Acres		
Averaging t to $t-s$:	$s=0$	$s=1$	$s=2$	$s=0$	$s=1$	$s=2$
	(1)	(2)	(3)	(4)	(5)	(6)
Drought Decile 1	0.07	0.15	0.13	0.07	0.23	0.13
	(0.07)	(0.07)	(0.07)	(0.11)	(0.10)	(0.10)
Drought Decile 2	0.02	0.12	0.27	-0.05	0.08	0.30
	(0.09)	(0.08)	(0.09)	(0.16)	(0.12)	(0.13)
Drought Decile 3	0.11	0.31	0.17	0.09	0.45	0.17
	(0.10)	(0.13)	(0.10)	(0.17)	(0.22)	(0.14)
Drought Decile 4	0.20	0.26	0.43	0.28	0.36	0.56
	(0.14)	(0.11)	(0.13)	(0.22)	(0.16)	(0.19)
Drought Decile 5	0.65	0.47	0.19	0.85	0.63	0.19
	(0.20)	(0.14)	(0.11)	(0.28)	(0.22)	(0.14)
Drought Decile 6	0.27	0.27	0.26	0.36	0.29	0.20
	(0.15)	(0.13)	(0.14)	(0.24)	(0.18)	(0.21)
Drought Decile 7	-0.00	0.49	0.41	-0.03	0.54	0.37
	(0.15)	(0.16)	(0.14)	(0.24)	(0.22)	(0.22)
Drought Decile 8	0.37	0.47	0.46	0.39	0.51	0.44
	(0.18)	(0.18)	(0.17)	(0.26)	(0.29)	(0.26)
Drought Decile 9	0.27	0.38	0.45	0.11	0.40	0.46
	(0.21)	(0.14)	(0.18)	(0.34)	(0.23)	(0.29)
Drought Decile 10	0.50	0.05	-0.01	0.43	-0.11	-0.17
	(0.20)	(0.20)	(0.20)	(0.28)	(0.29)	(0.30)
R^2	0.460	0.460	0.460	0.503	0.503	0.503
N	9,198	9,198	9,198	9,198	9,198	9,198
Clusters	511	511	511	511	511	511

Notes: Estimation results for the specification in Equation (1). The outcome is the percent of defoliation area in the county. We regress the defoliation rate on spring drought conditions, as captured by the Drought Severity and Coverage Index (DSCI). We split the index into deciles, conditional on non-zero values of the DSCI. The omitted category is no drought conditions (DSCI=0). Each regression includes county and state-by-year fixed effects. Sample includes all counties that had the spongy moth classified as established by 2000. Sample spans 2002 to 2019. Observations are either equally-weighted (columns 1 to 3) or weighted by the number of susceptible forest acres (columns 4 to 6). Standard errors are clustered at the county level.

Table A3.
The Effects of Spring Drought Conditions on Defoliation

Weighted by:	Equally-Weighted			Susceptible Acres		
Averaging t to $t-s$:	$s=0$	$s=1$	$s=2$	$s=0$	$s=1$	$s=2$
	(1)	(2)	(3)	(4)	(5)	(6)
DSCI∈(0,25]	3.12	4.63	6.70	3.06	5.83	6.62
	(1.44)	(1.44)	(2.24)	(1.92)	(1.70)	(2.23)
DSCI∈(25,50]	6.61	8.67	9.69	7.08	8.85	7.93
	(2.33)	(2.41)	(2.49)	(2.43)	(3.19)	(3.22)
DSCI∈(50_75]	7.65	6.38	8.33	7.67	6.47	7.05
	(3.41)	(2.51)	(3.12)	(4.47)	(3.51)	(4.27)
DSCI∈(75,100]	6.69	2.76	-0.20	2.41	1.21	-3.93
	(3.55)	(3.22)	(3.73)	(4.54)	(4.33)	(4.51)
DSCI∈(100,125]	6.12	-2.79	11.88	3.68	-6.76	8.47
	(3.40)	(4.07)	(4.17)	(5.47)	(4.75)	(4.44)
DSCI∈(125,150]	10.63	7.85	10.14	8.37	4.39	4.97
	(5.66)	(6.17)	(5.81)	(4.90)	(5.98)	(5.56)
DSCI∈(150,175]	10.50	11.12	3.47	7.90	3.91	-2.53
	(3.72)	(7.12)	(6.22)	(4.15)	(6.28)	(8.44)
DSCI∈(175,200]	4.78	17.20	12.96	1.71	14.64	14.15
	(11.93)	(13.40)	(3.45)	(13.35)	(11.02)	(3.96)
DSCI Above 200	0.11	-1.49	0.00	-6.26	-9.77	0.00
	(9.22)	(9.57)	(.)	(10.74)	(11.33)	(.)
R^2	0.475	0.476	0.476	0.524	0.525	0.524
N	9,198	9,198	9,198	9,198	9,198	9,198
Clusters	511	511	511	511	511	511

Notes: Estimation results for the specification in Equation (1). The outcome is the defoliation rate (defoliated acres relative to the mean level of forest acres in the county). We regress the defoliation rate on spring drought conditions, as captured by the Drought Severity and Coverage Index (DSCI). We split the index into equally-sized interval bins, with top-coded bin for DSCI values above 200. The omitted category is no drought conditions (DSCI=0). Each regression includes county and state-by-year fixed effects. Sample includes all counties that had the spongy moth classified as established by 2000. Sample spans 2002 to 2019. Observations are either equally-weighted (columns 1 to 3) or weighted by the number of susceptible forest acres (columns 4 to 6). Standard errors are clustered at the county level.

Table A4.
The Effects of Spring Drought Conditions on Defoliation

Weighted by:	Equally-Weighted			Susceptible Acres		
Averaging t to $t-s$:	$s=0$	$s=1$	$s=2$	$s=0$	$s=1$	$s=2$
	(1)	(2)	(3)	(4)	(5)	(6)
DSCI∈(0,25]	0.12	0.22	0.21	0.13	0.30	0.22
	(0.07)	(0.06)	(0.07)	(0.11)	(0.09)	(0.10)
DSCI∈(25,50]	0.27	0.32	0.30	0.38	0.34	0.22
	(0.10)	(0.13)	(0.13)	(0.14)	(0.20)	(0.20)
DSCI∈(50_75]	0.22	0.26	0.36	0.23	0.29	0.34
	(0.17)	(0.16)	(0.18)	(0.24)	(0.26)	(0.29)
DSCI∈(75,100]	0.20	0.06	-0.23	-0.00	-0.03	-0.45
	(0.20)	(0.21)	(0.20)	(0.29)	(0.30)	(0.29)
DSCI∈(100,125]	0.34	-0.27	0.44	0.29	-0.50	0.33
	(0.23)	(0.21)	(0.22)	(0.40)	(0.30)	(0.28)
DSCI∈(125,150]	0.61	0.39	0.36	0.57	0.19	0.12
	(0.29)	(0.30)	(0.26)	(0.30)	(0.36)	(0.30)
DSCI∈(150,175]	0.53	0.40	0.07	0.42	0.19	-0.21
	(0.16)	(0.25)	(0.25)	(0.21)	(0.34)	(0.39)
DSCI∈(175,200]	0.48	0.81	0.43	0.35	0.42	0.51
	(0.67)	(0.41)	(0.14)	(0.80)	(0.44)	(0.22)
DSCI Above 200	0.14	-0.05	0.00	-0.20	-0.55	0.00
	(0.34)	(0.36)	(.)	(0.47)	(0.51)	(.)
R^2	0.459	0.460	0.460	0.502	0.503	0.503
N	9,198	9,198	9,198	9,198	9,198	9,198
Clusters	511	511	511	511	511	511

Notes: Estimation results for the specification in Equation (1). The outcome is the percent of defoliation area in the county. We regress the defoliation rate on spring drought conditions, as captured by the Drought Severity and Coverage Index (DSCI). We split the index into equally-sized interval bins, with top-coded bin for DSCI values above 200. The omitted category is no drought conditions (DSCI=0). Each regression includes county and state-by-year fixed effects. Sample includes all counties that had the spongy moth classified as established by 2000. Sample spans 2002 to 2019. Observations are either equally-weighted (columns 1 to 3) or weighted by the number of susceptible forest acres (columns 4 to 6). Standard errors are clustered at the county level.

fixed effects, which nest the state-by-year and state-by-month fixed effects (column 5). We then revert to the state-by-year and state-by-month fixed effects, and include linear time trends at the county level (column 5), or at the county subdivision level (column 6).

Table A5.
Robustness to Specification Changes

	(1)	(2)	(3)	(4)	(5)	(6)
Defoliation	-0.040	-0.044	-0.036	-0.038	-0.040	-0.028
	(0.014)	(0.015)	(0.014)	(0.014)	(0.014)	(0.011)
County Sub-division FEs	X			X	X	X
ZCTA FEs		X				
Hexagon FEs			X			
State-by-Year FEs	X	X	X		X	X
State-by-Month FEs	X	X	X		X	X
State-by-Year-by-Month FEs				X		
County Trends					X	
County Sub-division Trends						X
R^2	0.300	0.300	0.311	0.304	0.301	0.310
n$\times 10^6$	28.44	28.44	28.44	28.44	28.44	28.44
Clusters	18,739	18,897	18,894	18,739	18,739	18,739

Notes: Estimation results from the specification in Equation (4). Each column reports results from using the same sample, which excludes large metro counties (same as in Table 2, column 4). Each column uses a different set of fixed effects. In column (1), we report results from the baseline set of fixed effects, as in Table 2. In column (2), we include ZCTA instead of county sub-division fixed effects. In column (3), we use EPA Hexagon instead of county sub-division fixed effects. In column (4), we use state-by-year-by-month instead of state-by-year and state-by-month fixed effects. In column (5), we include linear trends at the county level. In column (6), we include linear trends at the county sub-division level. Standard errors are clustered at the county sub-division level.

A.3 Estimation Results for Property Values When Including Property Characteristics as Controls

In the main text, we report estimation results across a variety of samples (Table 2), but only include location and time fixed effects to control for differences across locations, and absorb state-level pooled shocks and seasonality. Here, we repeat the estimation and include

property characteristics. We augment the specification in Equation (4) by adding fixed effects for the total number of rooms, the year built, a dummy for new construction properties, and deciles for the acres of the property. Including controls in the regression lowers the sample size, because not all characteristics are observed for all properties. In Table A6, we first report how the results change when we restrict the samples to have non-missing values for those control variables (Panel A), but do not include them in the regression, and then report the results from including the controls in the estimation (Panel B). First, comparing the results between the two panels, we find small differences in each column, highlighting that the inclusion of the controls themselves does not meaningfully affect the estimation results. When we compare to the results in Table 2, we note two differences: smaller sample sizes, and larger estimated negative effects on the sale price. For example, in the case of counties that ever-experience defoliation (column 2), we observe in the main sample (Table 2) a negative effect on sale price of 3.7 percent, estimated using 48.96 million property sales, across 9,865 county subdivisions. Restricting the sample to properties with non-missing values for the control variables, lowers the sample size to 27.18 million sales (44.5 percent reduction) across 6,163 county subdivisions (37.5 percent reduction), and recovers an estimated negative effect on sale price from defoliation exposure of 5 percent. Including the controls in the regressions recovers a negative effect on the sale price of 5.3 percent. We cannot reject that these coefficients are the same at conventional levels of statistical significance. To summarize, flexibly accounting for observable property characteristics does not absorb the effect of defoliatin exposure.

A.4 Balance Tables for Repeating Sales Samples

In the main text, we report results when estimating how defoliation exposure affects the sale price for a subsample of properties that sell more than once, and have variation in defoliation exposure. Here, we report how similar the properties are across the non-repeated sample, the repeated sample, and how similar the properties are that sell twice or more. In general,

Table A6.
The Effect of Direct Defoliation Exposure on Property Values With Controls

	Full Sample	County With Defoliation	Non-Metro Counties	Non-Large Metro Counties	ZCTAs With Defoliation	Within 5 km to Forest
Panel A. Sample With Non-Missing Property Controls						
	(1)	(2)	(3)	(4)	(5)	(6)
Defoliation	-0.051	-0.050	-0.040	-0.037	-0.028	-0.047
	(0.018)	(0.018)	(0.035)	(0.015)	(0.012)	(0.018)
R^2	0.322	0.326	0.282	0.318	0.313	0.323
n$\times 10^6$	45.19	27.18	3.91	16.46	11.88	41.09
Clusters	13,498	6,163	6,986	10,579	3,130	12,506
Panel B. Including Property Controls						
	(1)	(2)	(3)	(4)	(5)	(6)
Defoliation	-0.054	-0.053	-0.046	-0.040	-0.030	-0.049
	(0.016)	(0.016)	(0.031)	(0.015)	(0.011)	(0.016)
R^2	0.373	0.366	0.350	0.376	0.343	0.373
n$\times 10^6$	45.19	27.18	3.91	16.46	11.88	41.09
Clusters	13,498	6,163	6,986	10,579	3,130	12,506

Notes: Estimation results from the specification in Equation (4). We make the following modifications to the smaple and the specification: In Panel A, we report results for the same samples as in Table 2 but restrict those to the observations that have non-missing values for the control variables of total number of rooms, year built, new construction sales dummy, and deciles for the acres of the parcel. In Panel B, we report the results when including the controls in the regression. Standard errors are clustered at the county sub-division level.

we regress a dummy variable for being included in a specific sample on several property characteristics: the number of bedroom, the number of bathrooms, the total number of rooms, the year built, a dummy for the property being newly constructed, and the number of acres for the property's parcel.

We first define a dummy variable for having more than one sale in our sample (Table A7). We then focus on properties that have more than one sale, and define a dummy for having a variation in exposure to defoliation (Table A8). Within the repeating sales sample with variation in defoliation exposure, we focus on comparing properties with exactly two sales to properties that had more than two sales (Table A9). We continue to compare properties that sold three to five times (Table A10), and properties that sold six to 10 times (Table A11).

We recover precisely estimated coefficients for nearly all property characteristics, however, they are small in magnitude relative to the mean of the outcome, that we interpret them as precisely estimated zero differences between the types of properties being compared. The one exception is that properties that are newly constructed are meaningfully less likely to appear in the repeating sales sample, and are less likely to sell multiple times (more than two).

Table A7.
Balance Table for Repeated Sales Property

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Bedrooms	0.0003						
	(0.0000)						
Bathrooms		0.0050					
		(0.0001)					
All Rooms			0.0025				0.0033
			(0.0001)				(0.0001)
Year Built				0.0002			0.0004
				(0.0000)			(0.0000)
New Construction					-0.1074		-0.1185
					(0.0002)		(0.0003)
Acres (in 1,000s)						-0.0002	-0.0002
						(0.0001)	(0.0001)
Constant	0.8073	0.8030	0.7956	0.4556	0.8015	0.7949	-0.0887
	(0.0001)	(0.0003)	(0.0004)	(0.0029)	(0.0000)	(0.0000)	(0.0038)
Dep. Var. Mean	0.808	0.815	0.809	0.814	0.795	0.795	0.809
n$\times 10^6$	46.84	57.54	46.04	66.33	77.09	73.26	43.88

Notes: Estimation results from a simple regression of a dummy variable on a set of property characteristics. Robust standard errors in parentheses.

Table A8.
Balance Table for Repeated Sales Defoliation Sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Bedrooms	0.0007						
	(0.0000)						
Bathrooms		0.0018					
		(0.0001)					
All Rooms			-0.0003				0.0052
			(0.0000)				(0.0002)
Year Built				-0.0015			-0.0011
				(0.0000)			(0.0000)
New Construction					-0.1013		-0.0764
					(0.0002)		(0.0003)
Acres (in 1,000s)						0.0002	-0.0001
						(0.0001)	(0.0000)
Constant	0.2089	0.2143	0.2141	3.2661	0.2467	0.2378	2.3328
	(0.0001)	(0.0002)	(0.0002)	(0.0036)	(0.0001)	(0.0001)	(0.0048)
Dep. Var. Mean	0.211	0.218	0.213	0.258	0.241	0.238	0.202
n$\times 10^6$	37.87	46.87	37.25	54.02	61.28	58.23	35.48

Notes: Estimation results from a simple regression of a dummy variable on a set of property characteristics. Robust standard errors in parentheses.

Table A9.
Balance Table for Properties Two Times Only

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Bedrooms	-0.0001						
	(0.0000)						
Bathrooms		-0.0003					
		(0.0001)					
All Rooms			0.0013				0.0004
			(0.0001)				(0.0001)
Year Built				0.0001			0.0000
				(0.0000)			(0.0000)
New Construction					0.0780		0.0713
					(0.0007)		(0.0010)
Acres (in 1,000s)						-0.0003	0.0055
						(0.0003)	(0.0014)
Constant	0.2179	0.2107	0.2108	-0.0126	0.2076	0.2112	0.1722
	(0.0001)	(0.0002)	(0.0005)	(0.0062)	(0.0001)	(0.0001)	(0.0087)
Dep. Var. Mean	0.218	0.210	0.218	0.205	0.210	0.211	0.222
n$\times 10^6$	8	10.23	7.92	13.92	14.79	13.85	7.16

Notes: Estimation results from a simple regression of a dummy variable on a set of property characteristics. Robust standard errors in parentheses.

Table A10.
Balance Table for Properties Sold Three to Five Times

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Bedrooms	0.0001						
	(0.0000)						
Bathrooms		0.0058					
		(0.0003)					
All Rooms			0.0003				0.0006
			(0.0001)				(0.0001)
Year Built				0.0005			0.0003
				(0.0000)			(0.0000)
New Construction					-0.0257		-0.0283
					(0.0007)		(0.0011)
Acres (in 1,000s)						0.0008	-0.0033
						(0.0003)	(0.0015)
Constant	0.6343	0.6208	0.6330	-0.3442	0.6345	0.6321	0.0641
	(0.0002)	(0.0008)	(0.0003)	(0.0075)	(0.0001)	(0.0001)	(0.0102)
Dep. Var. Mean	0.634	0.634	0.634	0.636	0.634	0.632	0.632
n$\times 10^6$	8	10.23	7.92	13.92	14.79	13.85	7.16

Notes: Estimation results from a simple regression of a dummy variable on a set of property characteristics. Robust standard errors in parentheses.

Table A11.
Balance Table for Properties Sold Six to Ten Times

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Bedrooms	0.0000						
	(0.0000)						
Bathrooms		-0.0055					
		(0.0003)					
All Rooms			-0.0015				-0.0010
			(0.0001)				(0.0001)
Year Built				-0.0006			-0.0003
				(0.0000)			(0.0000)
New Construction					-0.0523		-0.0430
					(0.0005)		(0.0007)
Acres (in 1,000s)						-0.0005	-0.0022
						(0.0002)	(0.0008)
Constant	0.1479	0.1685	0.1562	1.3568	0.1579	0.1568	0.7637
	(0.0001)	(0.0007)	(0.0005)	(0.0059)	(0.0001)	(0.0001)	(0.0076)
Dep. Var. Mean	0.148	0.156	0.148	0.159	0.156	0.157	0.147
n$\times 10^6$	8	10.23	7.92	13.92	14.79	13.85	7.16

Notes: Estimation results from a simple regression of a dummy variable on a set of property characteristics. Robust standard errors in parentheses.

A Data Appendix

A.1 Defoliation Maps

Defoliation acts as both an outcome, when we study how weather shocks affect defoliation, and as a treatment when we study how defoliation affects housing markets. Defoliation maps are available in raster formats: one raster for each year between 1975 and 2019 inclusive. These maps are produced each year through aerial surveys carried out by the US Forest Service in the summer months. Each pixel in the raster represents an area of 2km x 2km on the Earth surface. A pixel is assigned the value of 1 if defoliated in the given year and a value of 0 otherwise. An example for the variation in the defoliation levels over time can be seen in Figure A1.

A.2 Weather

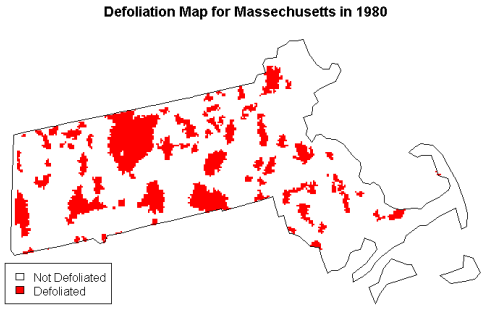
For weather, we use the PRISM Climate Data, which is available in the form of monthly rasters for the entire contiguous United States. Each pixel in the raster represents ≈4 km X 4km area on Earth. The dataset contains following weather measurements:

- Monthly total precipitation (including rain and melted snow)
- Monthly average of daily mean temperature
- Monthly average of daily minimum temperature
- Monthly average of daily maximum temperature
- Monthly average of daily mean dew point temperature

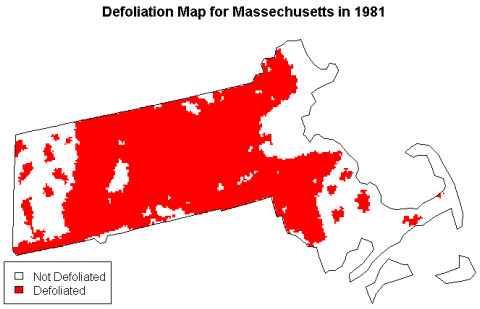
To study the impact of weather on defoliation, we need to match each defoliation pixel to a weather pixel. For this, we first convert each annual defoliation raster to a point shapefile where each point represents the centroid of a pixel. Then we overlay the shapefile with

Figure A1: Defoliation Maps for Massachusetts 1980-1983

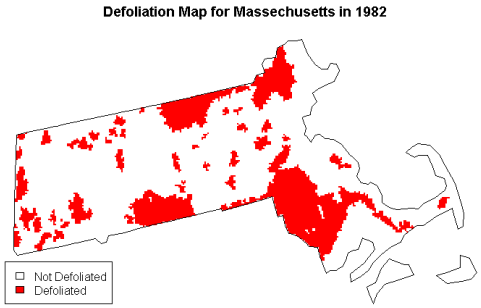
(a) 1980



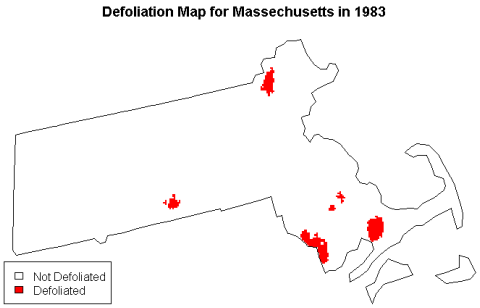
(b) 1981



(c) 1982



(d) 1983



Notes: Demonstrating the spatiotemporal variation in defoliation intensity. These maps focus on Massachusetts around the peak defoliation event that occurred in 1981.

the monthly PRISM rasters of the corresponding year in Google Earth Engine. Finally, we extract the monthly weather measurements that correspond with the location of centroids of defoliation pixels. Similarly, to study the impacts of weather on gypsy moth density, we overlay the trap shapefile and extract monthly weather data the same way we do for defoliation.

A.3 Property Transactions

Data on property transactions allow us to study how the housing market in the moth-infested region respond to gypsy moth outbreaks. We obtain geocoded property-level transaction data from CoreLogic Tax and Deed Data. Access to the data is provided by the Fama-Miller

Center for Research in Finance at the University of Chicago Booth School of Business and the Becker Friedman Institute for Economics at the University of Chicago. The dataset covers transactions between 1976 and 2019 inclusive and contains identifiers for specific property types, such as residential buildings, single-family homes, vacant land etc. After we exclude records with either a missing location, date of sale or price, non-arms-length transactions, and non-residential properties, we are left with ≈56 million observations for our study area, which we use in the analysis.

To study the impact of defoliation on the housing market, we need to create a treatment variable by matching each property transaction to defoliation events. The property that gets matched to a defoliation event is considered treated and considered untreated otherwise. We match each transaction in a given year with defoliation events in the past five years, including the present year, thereby producing 6 treatment variables. We do this to take into account the possible lag between defoliation and hypothesised treatment effect. For our analysis, we create two kinds of treatment variables:

- **Intersection:** This treatment variable simply takes the value of 1 if the centroid coordinates of a property transaction falls within a defoliated region and 0 otherwise.
- **Intersection area:** This treatment variable is defined fraction of circle of given radius centred at a given transaction that intersects with a defoliated area.