

Executive summary

Sound forest management policy depends on an understanding of how forest ecosystems are affected by neighboring regions. In this study, a natural experiment induced by policy differences across a state border were used to explore this issue. Data were obtained from plots within an old-growth forest in the upper peninsula of Michigan and from an adjacent managed (second- and third-growth) forest in Wisconsin. Observed variables included measurements pertaining to canopy structure (canopy height and number of windfalls); canopy heterogeneity (categories of trees by vertical extent); and species diversity (number and species of each tree and sapling observed). The questions of interest were: As we move from wild to managed land, which forest outcomes change? Where in relation to the border do they change? And how quickly, over space, does change occur? In these analyses, a two-stage modeling process was used to fit piecewise linear functions to the observed data. In the first stage, the best-fitting changepoint was selected from twelve regions in the center of the transects, six on either side of the border. Once a changepoint had been selected, piecewise functions were constructed that paired stable values toward the ends of each transect with linear increases toward the middle. Canopy height increased fairly rapidly from managed to wild land, rising over a horizontal distance of 50 meters and stabilizing 100 meters into the Michigan side. Windfall featured a changepoint centered 300 meters into the Michigan side and a fairly gradual slope, indicating a large edge effect. The edge effect on canopy heterogeneity coincided with the border and displayed a fairly quick rise over a 100-meter span. Finally, species heterogeneity showed a “reverse edge effect,” with values stable on the wild side up to the border, and decreasing only moderately fast on the Wisconsin side. Limitations and ideas for future directions are discussed.

Background

In public policy, opportunities for controlled experiments on a scale with real-world relevance are rare. However, political boundaries can create natural experiments that allow researchers to observe the effects of contrasting policies within regions which, if not for human intervention, would presumably be homogenous. In forest management, such a natural experiment can be found along the border between Wisconsin and Michigan’s upper peninsula. On the Michigan side of this border lies an old-growth forest, never harvested for timber; on the Wisconsin side, a variety of land management practices have obtained, and timber harvesting has occurred. The effect of state policy can be understood as the difference between the forest stands on either side of this border. Given the rarity and isolations of forest populations like the Michigan forest, it is important to understand these differences. What are the effects of land management on forest ecology? Where in relation to the border do the forest characteristics begin to change? And how quickly does that change occur?

This research program is ambitious, with five substantive hypotheses, and a few dozen potential outcomes, under consideration. Due to time and space constraints, I have chosen to focus on subsets of the data with certain properties: the data pertain to trees, rather than other classes of plants, making it reasonable to use similar modeling strategies throughout; the measures of interest can be quantified usefully at the plot level, rather than at the super-plot (as with importance) or sub-plot level; and exploratory graphical analyses of the outcomes suggested that interesting edge effects might be present and did not suggest major interpretive complications.

In the analyses that follow, I focus on three of the five major research questions:

- Q2) Do canopy structure and/or stand characteristics differ from wilderness to adjacent managed lands?
- Q3) Does canopy heterogeneity or spatial heterogeneity differ from wilderness to adjacent managed lands?
- Q5) Is there a difference in species diversity as a function of distance from the wilderness border and where does this occur?

Methods

Data collection

Data were collected from a region spanning the Wisconsin-Michigan border. On the Michigan side of the border, the region is part of the Sylvania wilderness, an old-growth forest from which timber has never been harvested, and in which many human activities that would disturb the forest are prohibited. On the Wisconsin side, the land has been privately managed, and timber harvesting occurs.

Inside the region of interest, eight transects were constructed perpendicular to the border, 500 meters in extent on each side of the border. Transects were selected such that the land sampled featured neither man-made structures nor substantial surface water (e.g., lakes or bogs). Within each transect, ten meter by ten meter plots were constructed every fifty meters, starting at the 500-meter mark; ten of these plots were in Wisconsin and ten in Michigan. No data were collected at the border (distance = 0).

Primary predictor

The major predictor of interest in this dataset is distance from the border. In order to differentiate the Michigan from the Wisconsin side of the border, for all plots on the Wisconsin side, I multiplied the edge distance by -1. Accordingly, my measure of distance ranged from -500 (furthest into Wisconsin) to 500 meters (furthest into Michigan), in 50-meter intervals.

Primary outcomes

Under the constraints described above, the following variables were selected as primary outcomes.

Q2: Canopy height (meters). A rangefinder was used to measure distance and angle of elevation. From these numbers, canopy height was calculated trigonometrically.

Q2: Windfall (count). In each plot, the number of downed trees of substantial size was tallied.

Q3: Canopy class. Within each plot, every tree of a certain minimum size was categorized as belonging to one of four classes: upper canopy (80% or more of its crown vertically exposed); lower canopy (some of its crown exposed, but less than 80%); upper understory (crown completely covered, but reaching the canopy foliage); and lower understory (crown entirely beneath the canopy). To understand the heterogeneity of the canopy within a given plot, the Shannon-Weiner diversity index was calculated on these canopy classes ($SWDI_{canopy}$). This index is equal to $-\sum_i p_i \log(p_i)$, where p_i represents the proportion of individuals within a plot belonging to the i^{th} class, $i=1, \dots, 4$.

Q5: Species diversity. All plants observed in a plot were categorized according to species, and the Shannon-Weiner diversity index was calculated for each plot based on these classes ($SWDI_{species}$). These analyses are confined to tree species, so herbs and shrubs were not considered. Further, because an accurate count of seedlings could not be obtained at the plot level, only trees and saplings were included in $SWDI_{species}$.

Two other candidate variables pertaining to stand characteristics (Q2) were rejected based on graphical exploration: canopy cover, because its distribution with respect to distance was quite heteroskedastic, with managed lands evincing much more variability in cover; and snags, because no clear relationship with distance was evident.

Data exploration

Distributions and bivariate relationships of the four outcomes are displayed in Figure 1. Looking at the diagonal, we see that canopy height and $SWDI_{species}$ are approximately normally distributed; windfall approximates a Poisson, as expected; and $SWDI_{canopy}$ has a multimodal distribution, with low canopy heterogeneity overrepresented. The outcomes appear to be relatively independent.

To get a feel for the types of edge effects I might observe, along with their functional form, I constructed a series of graphs plotting various outcomes against edge distance. Rather

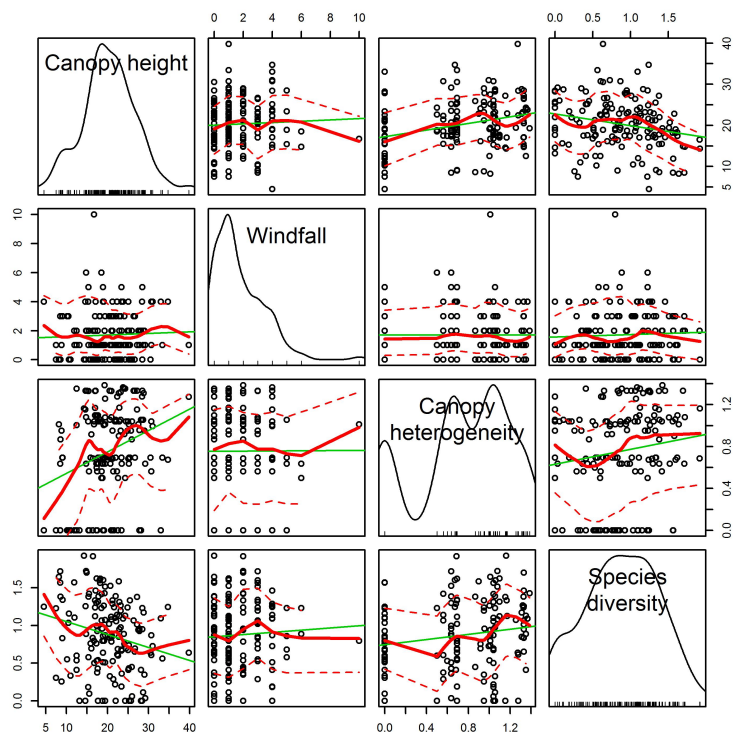


Figure 1. Scatterplot showing univariate (density) and bivariate distributions of four outcomes of interest.

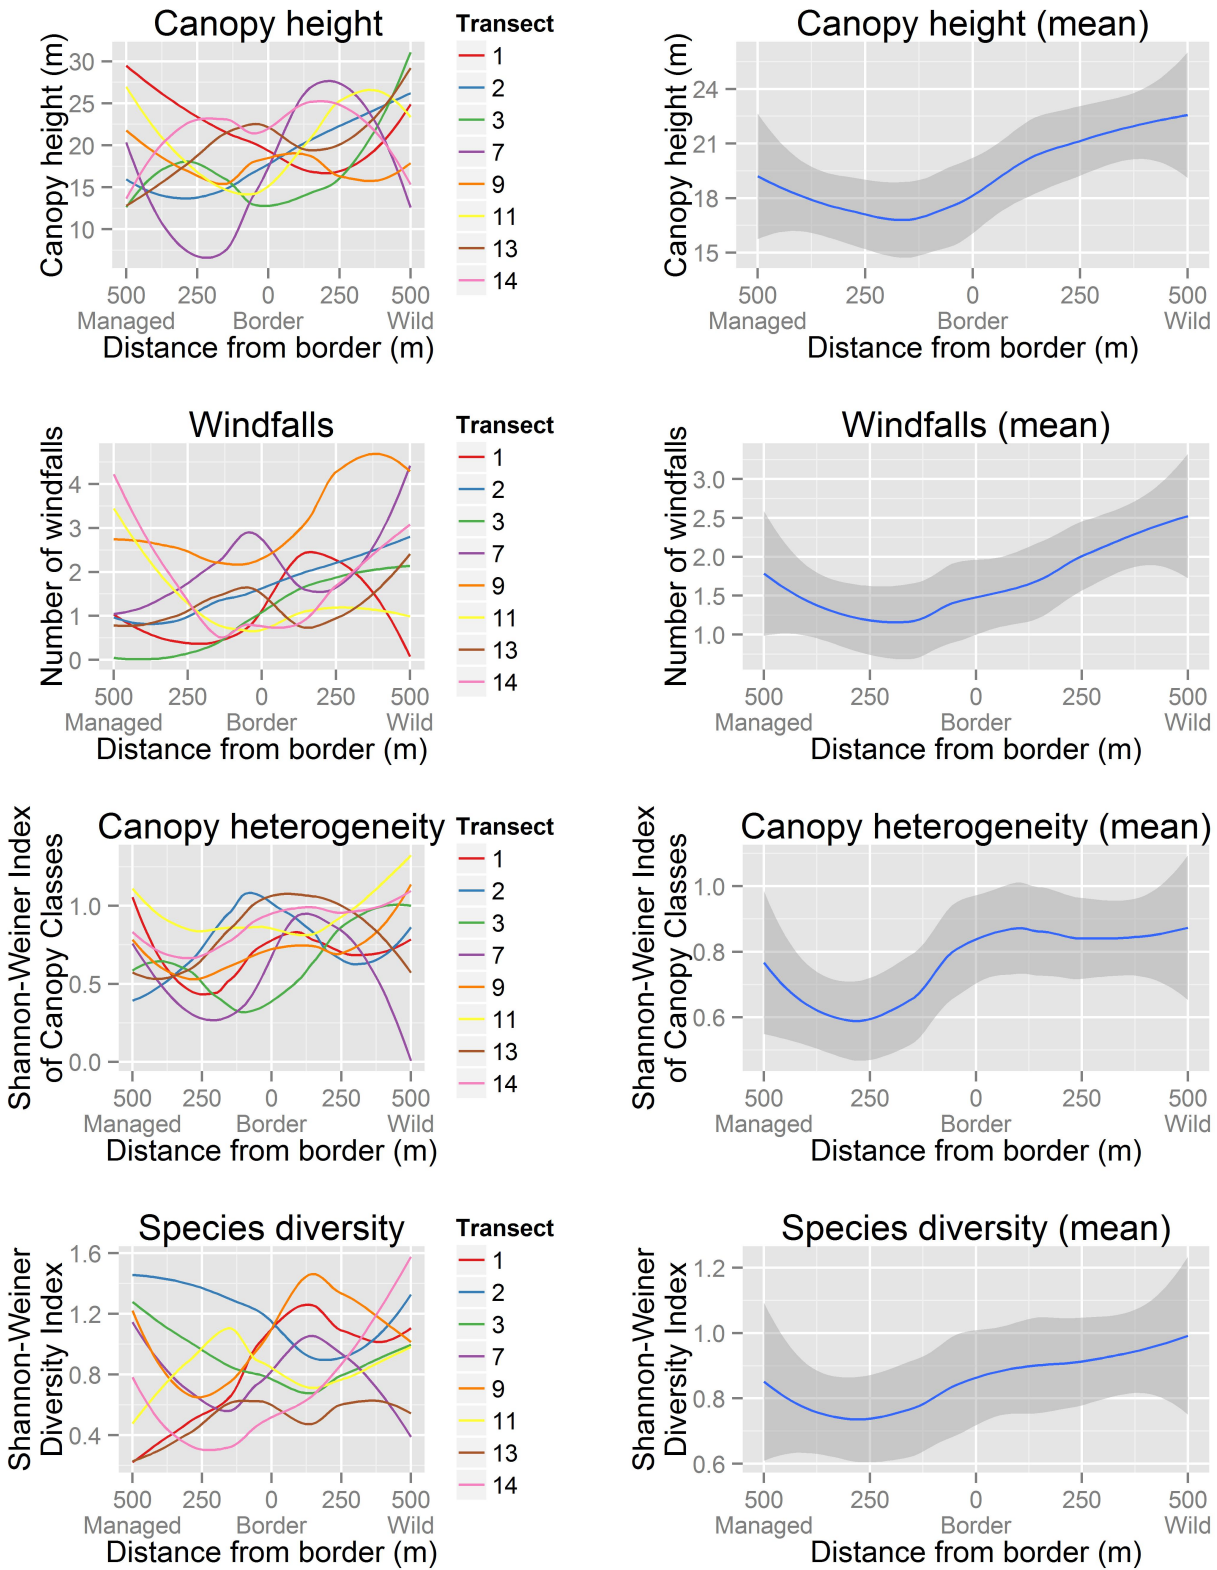


Figure 2. Using loess regression to explore functional form of relationship between plot distance and four outcomes. Left, curves for individual transects; right, curves for mean across transects at each distance.

than plotting individual datapoints, which were fairly noisy, I used loess regression to fit freeform curves to the data. Graphs from these explorations are displayed in Figure 2. In this figure, each horizontal strata represents a single outcome, ranging from canopy height (top) to species diversity (bottom). Within a given strata, the graph on the left fits curves to individual transects, whereas the graph on the right fits a curve to the mean value. Loess regression curves are too complex, and too prone to erratic behavior at the margins, to be a good edifice for a scientific theory. However, these exploratory graphs left me with three overall impressions that influenced my choice of models. First, several outcomes of interest evince similar patterns when plotted against distance. Second, these patterns are nonlinear in nature, and also unlikely to be approximated well by a simple polynomial function. And third, despite the pattern evident in the means, individual transects vary tremendously, and thus the data should not be treated as homogenous.

Statistical methods

For each of these outcomes, the primary question of interest is twofold: 1) Given that a particular outcome changes, where does the change from managed to wild land occur, and 2) What does it look like when it occurs? I operationalized the two parts of this question in two steps: first by finding the “change point,” i.e. the best dividing line between “wild” and managed”; and second by testing several possible rates of change surrounding this point.

Change points: To find the best change point, I constructed a series of indicator variables corresponding to various possible change points. Each of these took the value 0 for edge distances up until a particular hypothesized change point, and 1 thereafter. We can

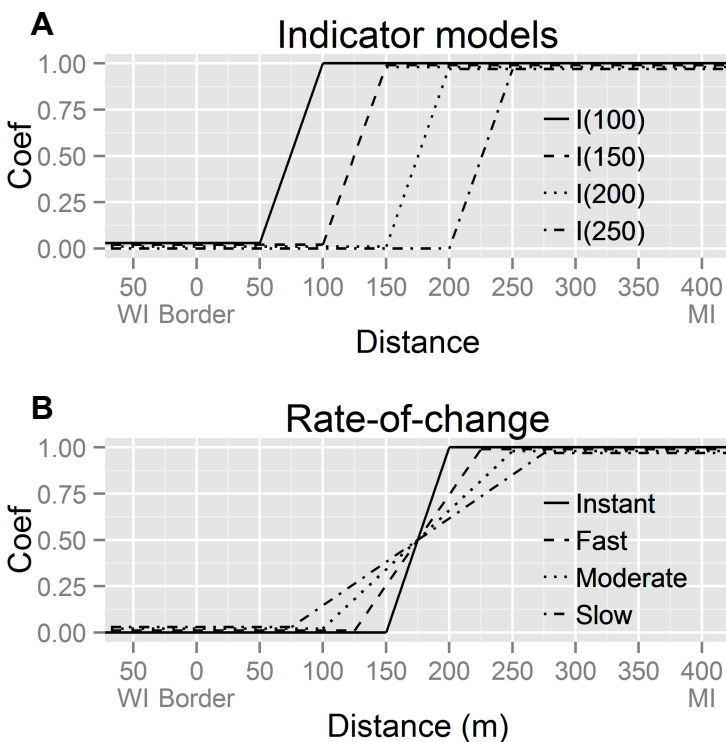


Figure 3. Schematic illustrating contrast coefficients for indicator models (A) and rate-of-change models (B).

represent this function for a given distance as $I(\text{distance})$, with locations on the Michigan side represented as positive distances. For instance, imagine we want to test the hypothesis that the change point lies on the Michigan side, between the 50-meter and 100-meter plots. We would represent this function as $I(d \geq 100)$: it takes the value 1 at the plot 100 meters into Michigan and plots further into Michigan, and takes the value 0 elsewhere. Similarly, if the changepoint were located 150 meters into Wisconsin, we would represent that as $I(-150)$. Based on the researcher’s comment that previous studies suggest edge effects may persist up to around 200 meters from a border, I constructed indicator

variables corresponding to possible change points ranging from 300 meters on the Wisconsin side to 300 meters on the Michigan side. A schematic illustrating four such variables can be found in Figure 3A. I tested the fit of each of these changepoints in twelve separate models, and selected the model with the lowest AIC value (or pseudo-AIC value, in the case of generalized linear mixed models).

Rate of change: The indicator variable model implicitly assumes a nearly instant change from managed to wild, which is unrealistic; change seems more likely to occur gradually. Once I determined the best-fit changepoint for a given outcome, I created additional predictor variables to model linear change in a piecewise fashion. For each of these, the value starts at zero at the plot 500 meters into Wisconsin; begins a linear rise at 50, 75, or 100 meters from the changepoint; and flattens out at 1 symmetrically on the opposite side. I refer to these as “fast”, “moderate”, and “slow” rates of change, respectively. A schematic of one such set of rate-of-change models, assuming a changepoint between 150 and 200 meters on the Michigan side, is illustrated in Figure 3B. Where possible, four models were compared – instant, fast, moderate, and slow; however, because no data were collected at the border, when the best-fit changepoint was at the border, only fast, moderate, and slow rates could be compared.

Within-transect correlations: The structure of the study, where plots within the same transect are typically closer in space to each other than are plots in different transects, strongly suggests a modeling strategy that takes these relationships into account. There are many such strategies, but mixed-effects modeling is the one with which I am most familiar, and so I have chosen that approach in these analyses. For simplicity, only random intercepts were considered: that is, transects are assumed to vary in their overall levels of each outcome of interest, but not in the types of edge effects those outcomes display. This may be an unrealistic assumption, but to evaluate it properly, data from more than eight transects would be needed.

Distribution of outcome variables: Some of the outcomes of interest, in particular the ones where multiple measurements contribute to the outcome (e.g. the Shannon-Weiner diversity indices), are approximately normally distributed, and for these, standard linear mixed models are appropriate. However, other outcomes, such as windfall, are counts, and will be better-modeled using generalized linear mixed models, which can cope with the restrictions they entail (e.g., predicted counts should not be below zero).

Software: Graphs were constructed in R using ggplot2, with loess smoothing for exploratory graphs. Statistical models were evaluated using SAS 9.3 PROC MIXED (for normal linear mixed models) and PROC GLIMMIX (for generalized linear mixed models, assuming a Poisson distribution).

Results

Stand characteristics: Canopy height

Across all plots, the mean canopy height was 19.4 meters (SD = 7.0). The distribution was approximately normal. To estimate the importance of transect as a factor in these data, we

can use the intraclass correlation, a measure which quantifies the relative importance of a single source of variance compared to the total variance in a dataset on a scale from 0 (low importance) to 1 (high importance). For canopy height, this value was quite small, ICC = 0.034.

Changepoint location. The AIC values for the twelve indicator models ranged from 1063.1 to 1076.4, with the minimum value occurring between 50 and 100 meters on the Michigan side of the border (corresponding to an indicator function of $I(d \geq 100)$).

Rate of change. Four piecewise linear functions representing different rates of change were considered: instant, fast, moderate, and slow. Of these, the instant rate of change fit best (AIC = 1063.1), and the edge effect in the final model was statistically significant (Wald $t(151) = 3.83$, $p = .0002$). The estimated intercept and slope (SE) were 17.6 (0.8) and 4.0 (1.1), respectively. In other words, among plots on the Wisconsin side of the change region, average canopy height was 17.6 meters; this rose over a distance of 50 meters to an average of 21.6 on the Michigan side, or an increase of approximately 0.8 meters in height per decimeter in distance from the change point.

Model diagnostics. Residual plots show nothing that should concern us; residuals are normally distributed, and no relationship with predicted value is evident.

Stand characteristics: Windfall

Across all plots, the median number of windfalls observed was 1 (range: 0 to 10). Because count data are nonnegative, linear mixed models, which permit predicted values below zero, are inappropriate; hence, generalized linear mixed models in the Poisson family were used instead. The estimated covariance¹ attributable to transect was 0.12.

Changepoint location. The pseudo-AIC values for the twelve indicator models ranged from 434.16 to 442.74, with the minimum value occurring at the innermost location examined on the Michigan side ($I(d \geq 300)$).

Rate of change. Instant, fast, moderate, and slow models were compared. The moderate rate of change fit best (pseudo-AIC = 433.24). The edge effect in the final model was statistically significant (Wald $t(151) = 4.14$, $p < .0001$). The estimated intercept and slope (SE) were 0.32 (0.15) and 0.55 (0.13), respectively. On the original scale, this corresponds to an estimated mean windfall count of 1.38 on the Wisconsin side of the change region and an estimated mean count of 2.41 on the Michigan side, with a moderate exponential rise over the 150 meters between. These estimates correspond reasonably well to the raw data, in which we see that among plots on the Wisconsin and Michigan sides of the change region, the median windfall counts were 1 (range: 0-6) and 2 (range: 0-10).

¹ Because Poisson regression does not estimate a separate parameter for error variance, I did not calculate an intraclass correlation, which requires it.

Model diagnostics. A quantile-quantile plot for the windfall model is shown in Figure 4. This plot compares the observed versus expected distribution of the model residuals. The curvature of this plot, in combination with the preponderance of zeroes (indicated by the strange clustering in the lower left quadrant), suggest that these data may be overdispersed compared to a Poisson distribution, i.e. there is more variability in observed windfall than can be accounted for by this model. Implications of this are discussed in the limitations section below.




Figure 4. Windfall model: quantile-quantile residual plot showing overdispersion.

Canopy heterogeneity

Across all plots, the mean $SWDI_{\text{canopy}}$ was 0.76 (SD = 0.44). The distribution did not show extreme deviation from normal, in that skewness and kurtosis values were moderate; however, a pronounced spike was observed at $SWDI_{\text{canopy}} = 0$. Very little effect of transect was observed (ICC = 0.016).

Changepoint location. The AIC values for the twelve indicator models considered ranged from 182.9 to 189.7, with the minimum value occurring on the border ($I(d \geq 50)$).

Rate of change. Because the changepoint chosen coincided with the border, only fast, moderate, and slow models could be compared. The fast rate of change fit best (AIC = 182.9), and the edge effect in the final model was statistically significant (Wald $t(142) = 2.93$, $p = .0039$). The estimated intercept and slope (SE) were 0.65 (0.053) and 0.20 (0.069), respectively. This indicates that among plots on the Wisconsin side of the change region, average heterogeneity was 0.65; this value rose over a distance of 100 meters to an average of 0.86, or an increase of approximately 0.02 “heterogeneity units” per decimeter.

Model diagnostics. Residual plots were examined. Subject to the constraints of the dataset, these residuals are reasonable; as in Figure 3, some clustering in the lower left was due to numerous plots in which only one canopy class was observed. Based on the observed edge effect, one might expect to see less canopy homogeneity on the Michigan side, but this was not the case (Wisconsin: 23%; Michigan: 16%; $\chi^2(1) = 1.15$, $p = .28$).

Diversity of species: trees and saplings

Across all plots, tree and sapling exemplars of 23 distinct species were observed. The mean $SWDI_{\text{species}}$ was 0.85 (SD = 0.48). The intraclass correlation was small but not ignorably so, ICC = 0.15.

Changepoint location. The AIC values for the twelve indicator models ranged from 212.5 to 217.9, with the minimum value occurring between 150 and 100 meters on the Wisconsin side of the border ($d \geq -100$).

Rate of change. Instant, fast, moderate, and slow models were compared. The moderate rate of change fit best (AIC = 212.3). The edge effect in the final model was statistically significant (Wald $t(151) = 2.54$, $p = .012$). Notably, and unlike other outcomes considered in these analyses, the “edge” in question was on the Wisconsin side of the border. The intercept and slope (SE) were 0.74 (0.087) and 0.19 (0.075), respectively. In practical terms, among plots on the Wisconsin side of the change region, average heterogeneity was 0.74; this rose over a distance of 150 meters to an average of 0.93, or an increase of approximately 0.013 “heterogeneity units” per decimeter.

Model diagnostics. Residual plots were examined. No notable model defects emerged, with residuals closely approximating a normal distribution.

Discussion

Summary

In these analyses, a two-stage modeling process was used to fit piecewise linear functions mapping four forest stand characteristics – canopy height, windfall, canopy heterogeneity, and species diversity – to distance from a significant boundary representing a difference in land management. In the first modeling stage, the best-fitting changepoint was selected from twelve regions near the border, six on either side. Once a changepoint had been selected, piecewise functions were constructed that paired stable, dissimilar values toward the ends of each transect with linear change in the region between.

For each outcome, these piecewise functions explained a portion of the variation: in other words, all four outcomes differed from wild to managed land, with locations and rates of change that varied from outcome to outcome. Canopy height increased fairly rapidly from managed to wild land, starting at 17.6 m on the managed side and rising, over a horizontal distance of 50 meters, to 21.6 m on the wild side, 100 meters from the border. Windfall appeared to be affected more dramatically, with a very interior changepoint centered 300 meters into the Michigan side and a moderate linear increase in $\log(\text{windfall})$ spanning 150 meters. The edge effect on canopy heterogeneity coincided with the border and displayed a fairly fast rise from 0.65 to 0.86 over a 100-meter horizontal span. Finally, for species heterogeneity, the observed edge effect was centered around a point 100 meters on the Wisconsin side of the border, with values rising over 150 horizontal meters from 0.74, on the managed side, to 0.93 close to the border.

The estimates for the windfall model are somewhat provocative, in that they suggest an edge effect that ranges fairly far into the wild side of the border. Because the best fit occurred at the innermost changepoint I tested, it is actually possible that the true model would locate a changepoint even farther into the wilderness. To fit a better model with a more interior changepoint, however, more interior data on the Michigan side would be

needed. Although potential problems with this estimate are discussed below, if we were to take this estimate at face value, these results could have policy implications.

One interesting feature of the species diversity model I estimated is the changepoint estimate it produced, which, unlike the other changepoint estimates in these analyses, is on the managed side of the border, 100 meters into Wisconsin. I could imagine a scenario by which this might occur: animals living in the old-growth forest may carry seeds across the border in a process resembling a random walk, with trajectories inside the managed land becoming gradually less common as distance from the border increases. However, given the small range of AIC values produced by these models, it is also possible that this finding is an artifact. Ultimately, the plausibility of each of these models must be evaluated with disciplinary knowledge in mind.

Limitations and future directions

The models I have presented feature several limitations. Three of the four models have some distributional weakness, though two of those are tempered by the functional form of my models. The windfall model, as noted above, is somewhat overdispersed; there exist statistical methods for managing this, and in future work those methods should be explored. Perhaps more troubling are the diversity indices; like windfall, these have a lower bound of zero in reality, but my models do not account for this fact. Both of these limitations would be fairly worrisome in the context of a continuous linear model, as extrapolating could lead the model to predict impossible values of the outcome. Because I have defined these relationships in a piecewise fashion, however, these limitations are less problematic in light of the particular model results obtained, which do not predict impermissible values for any plot. Given edge effects much further from the center and/or observed datapoints much closer to the absolute lower bound, the modeling strategy would need to be modified.

Although I have made an effort to handle the correlational structure induced by the study design, I have given little attention in this report to the nature of that structure. In point of fact, for two of the models, the covariance parameter estimate corresponding to transect is close to zero. Because for our purposes this is basically a nuisance parameter, I did not take the time to investigate the reasons for this, but glancing at the between-transect differences illustrated in the left half of Figure 1, a likely possibility is that there is not an overall level difference between transects in these outcomes. This does not rule out other types of transect effects, such as interactions between transect and distance – i.e. an edge effect which itself changes from transect to transect, with some regions affected more than others. I felt that accounting for such relationships was likely to require data from more than eight transects, and so decided that it was beyond the scope of this project.

A larger limitation of my analyses pertains to my specifications of functional form. The models here represent a piecewise linear approach to modeling what is essentially a nonlinear relationship: homogeneity on either side of a border region of indeterminate size, paired with some gradient effect within that region. The advantages of this approach are threefold: it is simple; it is appropriate for a relatively small dataset; and when it fails, it

should fail in an obvious fashion. However, the chief disadvantage is that it is unlikely to be a particularly faithful representation of the true functional form of this relationship.

Nonlinear modeling procedures, which can cope with a variety of functional forms, may be a closer fit for the research question. Like the S-shaped logistic models which are often used to model binary data, these techniques offer a flexible way to model curvilinear relationships, sometimes with parameters that afford fairly direct interpretations pertaining to the location of a change point and the associated rate of change. In my initial explorations with this dataset, I experimented with logistic growth models of some of the outcomes. However, the results of these explorations were generally unsatisfying, with some models failing to converge and others providing implausible parameter estimates. There could be several reasons for this. The other complexities of this dataset (such as the Poisson-distributed outcomes and the potential for within-transect correlations) necessitated the use of fairly sophisticated procedures, such as PROC NLMIXED, with which I am not familiar and which may fail in nonobvious ways. On the other hand, the difficulty could be a matter of sample size. The sheer variety of functional forms on display in Figure 1 do not suggest that a simple logistic function will be a particularly good fit for all plots of land; with only eight transects to consider, it may be that my nonlinear modeling attempts merely gave the data enough rope to hang itself. If such models are to be fit, they should be chosen with the benefit of a good grounding in the relevant ecological literature.

In this report, I have not addressed what are perhaps the two most interesting research questions, those pertaining to the makeup of the communities of species on either side of the border. One way to approach these questions might involve dimension reduction. In this approach, the first stage would involve using a clustering technique of some kind, such as factor analysis, to find groups of species that tend to occur together. Each plot could then be assigned a separate score pertaining to the strength of representation of each species cluster. In the second stage, those outcomes could then be modeled using the piecewise linear predictors described above – or, if one were feeling particularly adventurous, using some nonlinear modeling technique. The advantages of this approach are twofold: assuming that species depend on each other in some way, it reduces the number of outcomes considerably; and it does so while allowing us to consider data from “unimportant” species, which may well be unimportant only as individuals, but quite important in aggregate.