National Park Service
U.S. Department of the Interior

Natural Resource Stewardship and Science

Assessment of Imperfect Detection of Blister Rust in Whitebark Pine within the Greater Yellowstone Ecosystem

Natural Resource Report NPS/GRYN/NRR—2017/1457

ON THE COVER
Photograph of Erin Shanahan gathering data at a whitebark pine transect
Photograph courtesy of National Park Service

Assessment of Imperfect Detection of Blister Rust in Whitebark Pine within the Greater Yellowstone Ecosystem

Natural Resource Report NPS/GRYN/NRR—2017/1457

Wilson J. Wright[1,2], Kathryn M. Irvine[2]

[1]Montana State University
Bozeman, MT 59717

[2]U.S. Geological Survey
Northern Rocky Mountain Science Center
2327 University Way, Suite 2
Bozeman, MT 59752

Editing and Design by

Tani Hubbard

National Park Service (Greater Yellowstone Inventory and Monitoring Network)
& Northern Rockies Conservation Cooperative
Tucson, AZ 85748

June 2017

U.S. Department of the Interior
National Park Service
Natural Resource Stewardship and Science
Fort Collins, Colorado

The National Park Service, Natural Resource Stewardship and Science office in Fort Collins, Colorado, publishes a range of reports that address natural resource topics. These reports are of interest and applicability to a broad audience in the National Park Service and others in natural resource management, including scientists, conservation and environmental constituencies, and the public.

The Natural Resource Report Series is used to disseminate comprehensive information and analysis about natural resources and related topics concerning lands managed by the National Park Service. The series supports the advancement of science, informed decision-making, and the achievement of the National Park Service mission. The series also provides a forum for presenting more lengthy results that may not be accepted by publications with page limitations.

All manuscripts in the series receive the appropriate level of peer review to ensure that the information is scientifically credible, technically accurate, appropriately written for the intended audience, and designed and published in a professional manner.

This report received formal peer review by subject-matter experts who were not directly involved in the collection, analysis, or reporting of the data, and whose background and expertise put them on par technically and scientifically with the authors of the information.

Views, statements, findings, conclusions, recommendations, and data in this report do not necessarily reflect views and policies of the National Park Service, U.S. Department of the Interior. Mention of trade names or commercial products does not constitute endorsement or recommendation for use by the U.S. Government.

This report is available in digital format from the Greater Yellowstone Network website and the Natural Resource Publications Management website. To receive this report in a format that is optimized to be accessible using screen readers for the visually or cognitively impaired, please email irma@nps.gov.

Please cite this publication as:

Wright, W. J., and K. M. Irvine. 2017. Assessment of imperfect detection of blister rust in whitebark pine within the Greater Yellowstone Ecosystem. Natural Resource Report NPS/GRYN/NRR—2017/1457. National Park Service, Fort Collins, Colorado.

NPS 101/138712, 136/138712, June 2017

Contents

Figures

Tables

Abstract

We examined data on white pine blister rust (blister rust) collected during the monitoring of whitebark pine trees in the Greater Yellowstone Ecosystem (from 2004-2015). Summaries of repeat observations performed by multiple independent observers are reviewed and discussed. These summaries show variability among observers and the potential for errors being made in blister rust status. Based on this assessment, we utilized occupancy models to analyze blister rust prevalence while explicitly accounting for imperfect detection. Available covariates were used to model both the probability of a tree being infected with blister rust and the probability of an observer detecting the infection. The fitted model provided strong evidence that the probability of blister rust infection increases as tree diameter increases and decreases as site elevation increases. Most importantly, we found evidence of heterogeneity in detection probabilities related to tree size and average slope of a transect. These results suggested that detecting the presence of blister rust was more difficult in larger trees. Also, there was evidence that blister rust was easier to detect on transects located on steeper slopes.

Our model accounted for potential impacts of observer experience on blister rust detection probabilities and also showed moderate variability among the different observers in their ability to detect blister rust. Based on these model results, we suggest that multiple observer sampling continue in future field seasons in order to allow blister rust prevalence estimates to be corrected for imperfect detection. We suggest that the multiple observer effort be spread out across many transects (instead of concentrated at a few each field season) while retaining the overall proportion of trees with multiple observers around 5-20%. Estimates of prevalence are confounded with detection unless it is explicitly accounted for in an analysis and we demonstrate how an occupancy model can be used to do account for this source of observation error.

Acknowledgments

We thank Kristin Legg, Erin Shanahan, Rob Daly, and the numerous Greater Yellowstone Inventory and Monitoring Network field staff throughout the years who collected these data from high to low elevation and near and far from the trailhead. We thank both Dr. Steve Cherry (Montana State University) and Dr. Jonathan Beals-Nesmith (NPS Sierra Nevada Network) for providing their constructive and insightful comments on this report. Any use of trade, firm, or product names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

1 Introduction

A long-term interagency monitoring program of whitebark pine (*Pinus albicaulis*) in the Greater Yellowstone Ecosystem (GYE) was initiated in the summer of 2004 to assess the status of the species in this region. The main objectives of this monitoring program include estimating the proportion of trees infected with blister rust (*Cronartium ribicola*), assessing the severity of blister rust infections, estimating whitebark pine survival, and assessing whitebark pine recruitment (GYWPMWG 2011).

A potential source of error in the monitoring of blister rust is differences among observers. In order to provide data on observer variability, a proportion of transects visited each field season use multiple (2 or 3) independent observers. Previous work on these data from the initial survey years (2004 and 2005) showed that there was substantial lack of agreement between some observers (Huang 2006). In this report we first summarize the data collected from 2004 to 2015 during multiple observer transects for the whitebark pine monitoring program. We look at when there was observer agreement with presence or absence of blister rust (aecia and (or) secondary indicators), location of infection (bole or canopy) and the height of the infection on the tree (upper, middle, or bottom third). Secondly, in order to account for observer variability, we demonstrate how it is possible to explicitly model detectability of blister rust and blister rust prevalence simultaneously.

We used an occupancy model that focuses on overall tree infection (instead of location of infection - bole or canopy) and uses either method of detection (aecia or secondary indicators). Explicitly modeling location of infection and type of detection were not included in the analysis because they are not needed to estimate the proportion of trees infected with blister rust, which is a main objective of this monitoring program, but could be incorporated in future analyses in order to answer additional questions of interest. The main objective of this report is to estimate the influence of observer error on detection of blister rust in the GYE monitoring program and demonstrate how the use of multiple observers can improve the accuracy of modeling efforts to estimate the proportion of trees infected with blister rust.

2 Data

2.1 Monitoring Protocol

This section provides a brief overview of field methods, more detailed methods can be found in the interagency monitoring protocol (GYWPMWG 2011). From the mapped whitebark pine stands greater than 2.0 hectares within the GYE, 150 were randomly selected for long-term monitoring. Within these 150 stands, 176 transects (10 m x 50 m) were established. Most stands have one transect, while others have two in order to measure if there is any within stand variability. For assessing blister rust status, transects are visited once every four years, with the initial observations occurring between 2004 and 2007. All whitebark pine trees taller than 1.4 meters within each transect are marked with an aluminum tag. During each site visit, observers examine these tagged trees and record the number of blister rust cankers.

Two criteria are used for identifying blister rust cankers. The first is the presence of aecia (the sporulating body of blister rust) which is considered definitive evidence of a blister rust infection. The second criterion is a canker associated with at least three out of five secondary indicators. These secondary indicators include flagging, branch swelling, roughened bark, rodent chewing, and oozing sap. Positive identification of a blister rust canker can occur when one of these two criteria is met. Additionally, the position (bole or canopy) and height within a tree (upper, middle, or bottom third) of each tree's blister rust cankers are recorded. When a canker is identified by aecia, information is not recorded on secondary indicators regardless of whether they were present or not. Aecia detections "override" recording information on secondary indicators. For instance, if there is a single canker on a tree and the observer identifies it by seeing aecia, this observer would not record anything for secondary indicators even if they were present. A second observer could still identify the canker, but only record secondary indicators if he/she did not observe the aecia.

For the multiple observer transects, two or three field crew members independently record data for each tree on a transect. Observers work on each tree independently and record whether aecia or sufficient secondary indicators are present to identify a blister rust canker. Transects visited by multiple observers can be used to examine the variability among observers. Figure 1 shows the distribution of transects and the number of times a transect received a multiple observer survey. Multiple observer surveys of transects have ranged from zero to three, meaning some transects have never had multiple observer surveys while others had anywhere from one to three multiple observer surveys.

2.2 Data Used

Data used in the report come from the multiple observer transects collected for the interagency whitebark pine monitoring program in the GYE from 2004-2015 (NPS 2016). Data from transects without multiple observers were also examined in order to provide summaries of these transects for comparison. For each tree visited, available information includes diameter at breast height (DBH), tree height class, and percent of live canopy in the upper third of the tree. Observers recorded the number of cankers at each tree height (upper, middle, or bottom third) and location (bole or canopy). Whether cankers were identified by aecia or secondary indicators was also recorded. For this report, we summarized canker counts for each observer to an indicator variable for tree infection (at least one canker detected) or no infection (zero cankers detected) when analyzing these data. For the descriptions below we also summarize counts to indicator variables, but for some scenarios these summaries are by tree height and/or location instead of for the entire tree. Covariates measured at the transect level included elevation, aspect, slope, number of trees, and time spent hiking to the stand.

2.3 Data Summaries

During each field season (2004-2015), between 3 and 18 transects were observed by multiple observers. A total of 88 transects from 78 different stands have been included in these surveys, with some transects being visited by multiple observers for multiple years. The total number of observations varies by year and ranges from 32 to 1,074 (Table 1). These observations create 2,201 observation histories, where one observation history is the multiple records from different observers of a single tree during a given year. These observation histories are composed of 4,889 different observations where each observation is the record

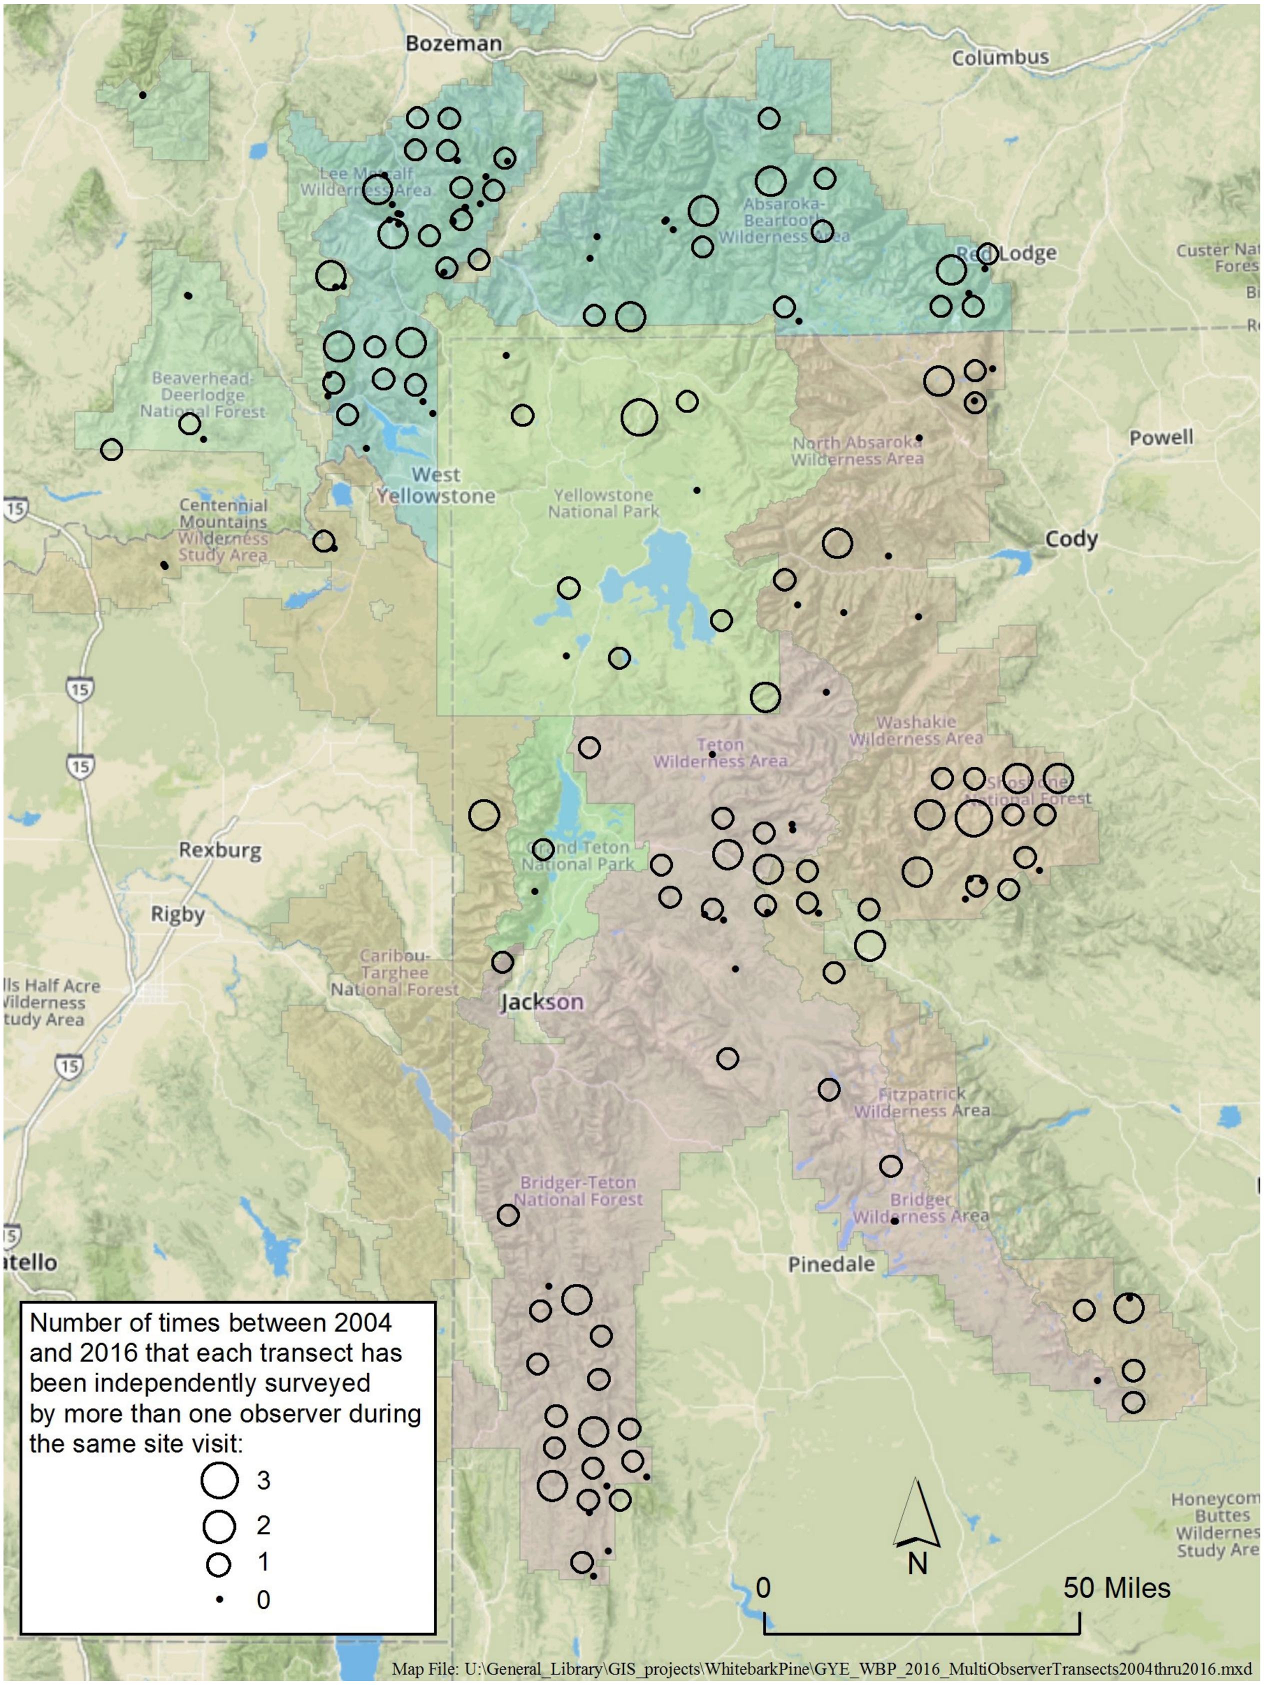

Figure 1. Distribution of Greater Yellowstone whitebark pine multiple observer transects from 2004-2015. Approximately 44 transects (a panel) are visited annually and during 2004 to 2015 each panel was visited three times to survey for blister rust infection on tagged trees.

Table 1. Tree observations on multiple observer transects and single observer transects by year.

Year	Multiple observer transect observations*	Single observer transect observations	Multiple observer total number of transects	Single observer total number of transects
2004	191	841	6	39
2005	1074	2366	18	58
2006	457	634	9	31
2007	32	276	3	12
2008	130	947	7	35
2009	307	952	11	33
2010	232	1156	10	32
2011	122	1395	9	35
2012	352	825	13	29
2013	720	2105	13	68
2014	230	1132	7	32
2015	1042	617	4	38

*Each observation is the record from one observer for a particular tree in a given year.

Table 2. Total number of trees examined by each observer and the range of survey years included. Only includes the observations from the multiple observer transects.

Observer #	Total observations (number of trees examined)	First multiple observer survey year	Last multiple observer survey year
1	9	2008	2008
2	636	2014	2015
3	164	2006	2006
4	59	2010	2010
5	170	2006	2006
6	237	2007	2011
7	267	2011	2013
8	441	2004	2005
9	64	2004	2004
10	101	2009	2010
11	44	2012	2012
12	39	2011	2011
13	76	2012	2012
14	11	2013	2013
15	13	2012	2012
16	174	2008	2011
17	268	2013	2013
18	655	2012	2015
19	318	2005	2005
20	747	2004	2015
21	173	2007	2009
22	223	2012	2013

from one observer for a particular tree in a given year. There are an additional 11,457 live tree observations from single observer records between 2004-2015. These observations come from a total of 175 transects from all 150 stands. One of the 176 transects has had multiple observer surveys conducted every site visit and therefore is not included in the single observer summaries. The total number of observations from single observer transects per field season varies between 276 and 2,366 trees (Table 1).

In total, twenty-two different observers have participated in the multiple observer transect surveys. Different observers examined a different number of trees on these surveys, ranging from 9 to 747. Some observers participated in multiple observer transects during a single summer while others observed transects during multiple summers. Table 2 shows the total number of trees examined by each observer for the multiple observer transects.

Information collected for each tree includes diameter at breast height (DBH) in centimeters, height class (1: ≤5 m, 2: >5 m and ≤10 m, 3: >10 m), and the percent of live canopy in the upper third. The DBH and height class are recorded once for each tree when aluminum tags are attached to the tree and not recorded multiple times by each observer. The live canopy volume is recorded by each observer independently during each visit. For tree records with multiple observers, the upper live canopy percent values were found by averaging across all observers for that tree. Some entries had no data and were dropped to compute this summary statistic. Tables 3 and 4 summarize these tree measurements from the multiple observer and single observer transects. Ideally, the characteristics of the trees visited by multiple observers are similar to the trees only visited by a single observer.

Table 3. Counts for trees in each height class on multiple observer transects and single observer transects.

Class	Multiple observer transects count	Single observer transects count
1: (≤5 m)	1243	6093
2: (5 m≤ H ≤10 m)	615	3907
3: (>10 m)	343	1457

Certain characteristics measured at the transect level could also be important to consider when examining detection of blister rust cankers. These include a topography measure (sin(slope) ∗ cos(aspect)) of the transect, number of trees examined on the transect, and the time required to hike to the transect. Summary statistics on these characteristics for the unique transects visited by multiple observers and by a single observer are shown in Table 5. Again, we hope to see similar transect level characteristics between the single and multiple observer transects.

Of the total 4,889 observations made on the 88 multiple observer transects, 1,342 were recorded as having at least one canker from a blister rust infection (either from aecia or 3 to 5 secondary indicators). Of these 1,342, there were 678 observations recorded of a tree having at least one canker with aecia (51% of all observations of blister rust). For the single observer transects, overall there were 3,254 total observations of infection present (identified using either criteria) and of those there were 1,797 total observations of aecia (55% of all observations of blister rust).

Table 4. Summary statistics for DBH and upper live canopy percent on multiple observer transects and single observer transects.

Summary statistic	Multiple observer transects		Single observer transects	
	DBH	Upper live canopy percent	DBH	Upper live canopy percent
Min.	1.00	0.00	1.00	0.00
1st Q	2.50	72.50	3.00	80.00
Median	7.50	87.50	8.00	90.00
Mean	11.03	75.98	10.99	77.62
3rd Q	16.50	92.50	16.00	95.00
Max.	95.00	97.67	126.50	100.00

Table 5. Summary statistics for slope at the transect, number of live trees examined, and hiking time to the transect.

	Multiple observer transects			Single observer transects		
Summary statistic	Slope (degrees) at the transect	Number of live trees examined	Hiking time (hours) to transect	Slope (degrees) at the transect	Number of live trees examined	Hiking time (hours) to transect
Min.	2.00	1.00	1.50	1.00	1.00	1.50
1st Q	11.00	6.00	4.00	12.00	5.00	4.00
Median	18.00	12.00	6.00	20.00	14.00	6.00
Mean	19.01	20.01	6.182	19.66	25.92	6.36
3rd Q	26.00	23.00	8.00	27.00	34.00	8.00
Max.	38.00	207.00	14.50	45.00	220.00	15.50

2.4 Summaries of Observer Agreement

For the total of 2,201 observation histories from multiple observer transects, the number of times observers agreed that a tree was infected or not infected can be counted. Some trees on transects that were designated for multiple observers still only had a single record due to time constraints or individual observers becoming ill partway through a survey. These records are not included in the tables below because the agreement summaries require more than one observation at the tree. The occupancy model (next sections) analyzes all tree records from both the single and multiple observer transects. Tables 6 and 7 show the number of observers present (two or three) and the number of observers who identified a blister rust infection (ranging from zero to the total number of observers present) and how many times each scenario occurred for these observation histories.

Table 6 shows that, for aecia, total agreement is 92.0% and 85.9% for two observer and three observer histories respectively. These values were calculated by summing the observation histories where either none or all observers detected aecia and dividing by the total number of observation histories. However, for only observation histories where at least one observer identified aecia, all observers identified aecia 57.3% and 27.8% of the time for two observer and three observer histories respectively. This calculation assumes that there are no false positives and that if one observer detected aecia, the others could potentially miss it. For these calculations, the secondary agreement percentages use the total number of trees where all observers identified aecia as the numerator and the number of trees where at least one observer identified aecia as the denominator. In this way, we are focusing more on examining the probability of detecting cankers.

For overall infection (Table 7), total agreement is 92.0% and 80.3% for two observer and three observer histories respectively. Only examining histories where at least one observer detected infection, total agreement is 75.6% and 42.0% for two and three observer histories respectively.

Observation histories (after 2004) and observer agreement can further be broken down by location in the tree. First, detections in the bole versus detections in the canopy were compared. Agreement by

Table 6. Number of observers who identified a tree as having aecia. Total observers present is two or three and the number of observers who identified a blister rust infection ranges from zero to the number of observers present.

Total observers	Observers identifying aecia			
	Zero	One	Two	Three
Two	1378*	135	181*	--
Three	400*	40	30	27*

*Indicates total agreement by all observers present which can occur when all observers classify a tree as not infected or all observers classify a tree as infected.

Table 7. Number of observers who identified a tree as being infected (aecia or secondary indicators).

Total observers	Observers identifying overall infection			
	Zero	One	Two	Three
Two	1140*	135	419*	--
Three	328*	44	54	71*

*Indicates total agreement by all observers present which can occur when all observers classify a tree as not infected or all observers classify a tree as infected.

Table 8. Number of observers who identified a tree as being infected based on aecia only by infection location.

Total observers	Observers identifying aecia in the bole				Observers identifying aecia in the canopy			
	Zero	One	Two	Three	Zero	One	Two	Three
Two	1507*	93	93*	--	1506*	105	82*	--
Three	419*	7	4	4*	357*	30	26	21*

these positions for aecia cankers can be seen in Table 8. This table shows similar values for total agreement on aecia — in the bole, 94.5% and 97.5% for two and three observers respectively and in the canopy 93.8% and 87.1% for two and three observers respectively. However, there is less agreement when only trees having at least one observer identify aecia are examined. Calculated this way, agreement is 50.0% and 26.7% in the bole for two and three observers respectively and 43.9% and 27.3% in the canopy for two and three observers respectively.

We can also examine observer agreement based on the overall infection status (aecia or secondary indicators) by infection location. These summaries are shown in Table 9. Based on these observations, total agreement on infection in the bole is 91.4% and 93.5% for two and three observers respectively and 91.1% and 79.5% in the canopy. Agreement for trees as being infected is 54.5% and 20.0% for two and three observers in the bole and 56.5% and 34.1% in the canopy.

Agreement among observers can also be compared for the six locations in a tree since all cankers are assigned to a height within the appropriate position. Table 10 shows agreement based on these locations. This table shows an overall decrease in agreement when examining the detection of cankers by location. There are far fewer observation histories in

Table 9. Number of observers who identified a tree as being infected (aecia or secondary indicators) by infection location.

Total observers	Observers identifying infection in the bole				Observers identifying infection in the canopy			
	Zero	One	Two	Three	Zero	One	Two	Three
Two	1375*	145	173*	--	1346*	151	196*	--
Three	399*	16	12	7*	299*	38	51	46*

*Indicates total agreement by all observers present which can occur when all observers classify a tree as not infected or all observers classify a tree as infected.

Table 10. Number of observers who identified a tree as being infected (aecia or secondary indicators) by infection location in the canopy or bole and by height (upper, mid, bottom).

Location (number of observers)	Observers identifying infection (aecia)				Observers identifying infection (indicators)			
	Zero	One	Two	Three	Zero	One	Two	Three
Upper Canopy (Two)	1609	59	25	--	1557	78	58	--
Upper Canopy (Three)	404	16	10	4	385	19	13	17
Mid Canopy (Two)	1593	77	23	--	1541	99	53	--
Mid Canopy (Three)	383	25	18	8	384	25	18	7
Bot Canopy (Two)	1617	52	24	--	1616	56	21	--
Bot Canopy (Three)	415	12	4	3	413	11	10	0
Upper Bole (Two)	1622	46	25	--	1564	83	46	--
Upper Bole (Three)	426	3	3	2	413	9	6	6
Mid Bole (Two)	1590	67	36	--	1557	96	40	--
Mid Bole (Three)	428	4	0	2	424	5	2	3
Bot Bole (Two)	1642	30	21	--	1627	48	18	--
Bot Bole (Three)	432	2	0	0	431	3	0	0

which all observers present agree on the location of a canker compared to instances where a portion of observers identified a canker in a particular location. Essentially, the values in the "Two" columns for two observers and "Three" columns for three observers are much less than the other columns greater than zero. In particular, compared to the previous tables, the agreement of observers appears to have declined considerably. This illustrates that an additional component of observer variability is in the estimation of canker location, in particular when height is assigned. The purpose of assigning canker locations is to get an idea of severity of infection (bole infection is believed to be more lethal for instance), but there may not be enough reliability in these data to incorporate canker height when describing the location of an infection.

2.4.1 Highlights of Observer Agreement Summaries

When examining the observer agreement summaries, there are a few key pieces of information that should be noted. Many of these summaries, such as those below, assume that there are no false positives from observers. The following bullets show percentages for total agreement when at least one positive identification for a tree's history occurred.

1. When examining the entire tree:
    a. All observers agree on aecia 57.3% and 27.8% of the time when 2 or 3 observers are present.
    b. All observers agree on infection 75.6% and 42.0% of the time when 2 or 3 observers are present.
2. When examining each tree by location (bole or canopy):
    a. For aecia, all observers agree 50.0% and 26.7% when located on the bole and 43.9% and 27.3% when located in the canopy, when 2 or 3 observers are present.
    b. For secondary indicators, all observers agree 43.9% and 31.0% when located on the bole and 45.1% and 38.3% when located in the canopy, when 2 or 3 observers are present.
    c. For overall infection, all observers agree 54.5% and 20.0% when located on the bole and 56.5% and 34.1% when located in the canopy, when 2 or 3 observers are present.
3. When including height of infection in addition to location, agreement decreases further. This indicates another potentially large source of observer variability.

Overall, these observer agreement summaries support the conclusions of earlier work on the multiple observer transects because at times there appears to be a substantial lack of agreement between observers (Huang 2006). The remainder of this report develops a model to account for imperfect detection error.

3 Model

The main reason differences in blister rust detections can occur between observers is because of observer error. In the context of blister rust monitoring, this can happen when an observer identifies a truly healthy tree as infected or identifies a truly infected tree as healthy. In this monitoring program, the more likely error to be made is the second; cankers are missed by observers and the tree is classified as healthy. These models will assume that this is the only error that is made and that there are no false positives in identifying blister rust cankers. Currently, the detection errors are not being accounted for when estimating blister rust prevalence. The multiple observer transect data can be used to model blister rust detection and incorporate it in the estimates for blister rust prevalence. The next section of this report will analyze these data using a model which explicitly accounts for the detection process.

Imperfect detection of blister rust results in a scenario analogous to that of occupancy modeling in studies on wildlife populations. In occupancy studies, the state (occupied or unoccupied) of each site in a region is of interest. However, even if a site is occupied it is possible to incorrectly observe the site as unoccupied because imperfect detection could result in not observing the species when the site is visited. In order to estimate and account for detection probabilities, each site is visited multiple times. Unbiased estimates for the proportion of all sites occupied can then be found by incorporating detection probabilities (MacKenzie et al. 2006).

In the whitebark pine monitoring program the multiple observer transects can be used to estimate detection probabilities of blister rust. Relating this scenario to the occupancy framework, each tree can be thought of as a "site" where the blister rust state (infected or not infected) is of interest. The multiple observer records can be viewed as repeated visits to each tree in order to determine the detection probability of blister rust when a tree is truly infected. These detection probabilities can then be accounted for in order to acquire more accurate estimates of the proportion of trees infected with blister rust in the GYE – a main objective of the monitoring program. In an occupancy modeling framework, there are two main parameters that are being modeled:

1. ψ is the overall probability of a tree being infected.
2. p is the probability of detecting the infection (either method), given the tree is infected.

Certain covariates are expected to influence a tree's probability of infection as well as the probability of an observer detecting an infection. For the probability of infection, previous work has shown that larger trees are more likely to be infected than smaller trees (Campbell and Antos 2000) and our data support this. Tree diameter at breast height (DBH) will be used to account for this variation. The stand characteristics of elevation, topography (sin(slope) * cos(aspect)), and number of trees were also included in the model for infection probability. The probability of detection will be modeled with covariates that are believed to be potentially important. These include size of tree (DBH), time of year (date), slope of the transect (slope), hiking time to the transect (hike), and number of trees observed on the transect (trees). Also included was an indicator variable for whether it was an observer's first field season or not in order to account for potential improvement from one year to the next.

There is an additional level of variation to account for in these data because trees are grouped together by stands. Stand random effects ($b^t_{stand_{ti}}$) were included in the model for infection (ψ) in order to account for this dependence among trees. Random effects for observer ($a_{obs_{tijk}}$) were also included in the detection model. The random effects are assumed to be normally distributed with a mean of zero and a unique variance for each distribution. The variances (σ^2_{stand} and σ^2_{obs}) will be of interest when making inferences about these random effects. Observer was included as a random effect because of the large number of observers used throughout the study and because this allows inferences to be made about the population of future observers. Including these additional covariates produces the following model specification for each parameter:

$$1.\ logit\ (\Psi_{tij}) = \beta^t_0 + \beta^t_1 DBH_{tij} + \beta^t_2 \boldsymbol{elev}_{ti} + \beta^t_3 \boldsymbol{topo}_{ti} + \beta^t_4\ trees_{ti} + b^t_{stand_{ti}}$$

$$2.\ logit\ (p_{tijk}) = \alpha_0 + \alpha_1 DBH_{tij} + \alpha_2 date_{ti} + \alpha_3 slope_{ti} + \alpha_4 \boldsymbol{hike}_{ti} + \alpha_5\ trees_{ti} + \alpha_6\ trees^2_{ti} + \alpha_7 exp_{tijk} + a_{obs_{tijk}}$$

Subscripts refer to t for time period, i for transect, tree j within transect i, and an individual record k for tree j within transect i. Note that estimates of the regression parameters associated with blister rust prevalence differ by year (denoted by t superscripts on each β) while the regression parameters associated with blister rust detection do not. This is explained in more detail below. We fit this model using a Bayesian approach with weakly informative priors for all parameters. For the standard deviations of the random effect distributions (σ_{obs} and σ^{t}_{stand}) half-Cauchy(0, 2.5) priors were assumed. For all of the remaining parameters (i.e., the occupancy and detection coefficients) we used Normal (0, 2.5) prior distributions to fit this model. These priors place more density on values closer to zero while still having a large enough variance to include likely slope coefficient values associated with centered and scaled covariates on the logit scale.

The blister rust observations are divided into 3 different time periods – T0 was 2004-2007, T1 was 2008-2011, and T2 was 2012-2015. Estimates of prevalence within each time period are of interest (since each site is visited once within a time period). Because of this, the parameters associated with prevalence (ψ) will be uniquely estimated using the observations from each time period separately (parameters noted with the t superscript). However, the estimates of the detection parameters will utilize observations across all time periods. Separate analyses indicated that there was no difference in detection parameters between T0 and T2. The model for T1 failed to converge and there appears to be insufficient data to estimate the detection parameters using only observations from T1. Since we would not expect the detection process to differ by year (after accounting for the covariates described above) and this was seen in the analyses for T0 and T2, pooling observations to estimate detection seems reasonable. This approach still allows for separate estimates of prevalence for each time period.

4 Analysis

The following analysis was performed for all observations from years 2004-2015 except for visits to transects which were not part of that panel. In some years a transect was not surveyed with its designated panel due to safety, weather conditions, or some other unplanned circumstance and as a result was often visited the following year with another panel. All of the covariates used were standardized before performing the analysis (centered and scaled). This model was fit in a Bayesian framework using Stan (mc-stan.org) called from Program R using the rstan package. Traceplots were examined to ensure the model had converged sufficiently (shown in Appendix A) and all $\hat{R}$ values were less than 1.05 which also indicated convergence.

4.1 Infection

Overall infection rate (ψ) was modeled with tree DBH, transect elevation, transect topography (sin(slope) $*$ cos(aspect)), and number of trees per transect as explanatory variables. Overall adjustments to stand blister rust prevalence were allowed by including the random effects for stand. Posterior means for the prevalence parameters and associated 95% posterior intervals are shown in Table 11 for each time period.

These estimates show that across all time periods, the odds of being infected increases as DBH increases after accounting for stand differences and other covariates. The relationship between the probability of infection and DBH appears strongest of the covariates examined here. Based on this model fit, during T1, an increase in tree DBH of 10.8 cm. is associated with an increase in the odds of infection by a factor of 4.85 with an associated 95% posterior interval for this factor from 3.94 to 6.05 after accounting for the other covariates in the model. The other slope coefficients in the model can be interpreted similarly, but the figures shown below provide a more intuitive interpretation of the model results on the probability scale. The estimated variance component of the stand random effect is quite large in each time period. This suggests there is large site variability in the overall probability of a tree being infected.

How the probability of infection (proportion of trees infected) changes over time is of particular interest in the GYE. Examining the estimated occupancy intercept for each time period describes the estimated probability of infection for the overall average tree size (DBH= 11.19 cm) at the average stand. The posterior means for these quantities and their associated 95% posterior intervals can be displayed for each time period and compared to the reported estimates of the proportion of trees infected (Shanahan et al., 2017; Figure 2). For each time period, accounting for imperfect detection with the occupancy model used here yields higher estimated probabilities of infection for the average tree compared to the ratio estimators. The posterior intervals for each time period are wider than their associated confidence intervals from the ratio estimators and there is some overlap between the two for each time period. Another way to interpret these model results is to examine estimated probabilities of infection versus

Table 11. Overall infection rate (ψ) modeled with tree DBH, transect elevation, transect topography (sin(slope) $*$ cos(aspect)), and transect number of trees as explanatory variables. Posterior means for the prevalence parameters and associated 95% posterior intervals are shown.

	T0			T1			T2		
Parameter	Mean	2.5%	97.5%	Mean	2.5%	97.5%	Mean	2.5%	97.5%
β_0	-1.15	-1.63	-0.70	-0.74	-1.20	-0.30	-0.76	-1.21	-0.29
β_1(DBH)	1.24	1.06	1.43	1.58	1.37	1.80	1.70	1.46	1.95
β_2(elev)	-0.79	-1.13	-0.45	-0.85	-1.19	-0.53	-0.51	-0.83	-0.20
β_3(topo)	0.13	-0.17	0.42	0.19	-0.09	0.47	0.10	-0.21	0.41
β_4(trees)	0.20	-0.31	0.67	0.21	-0.27	0.71	0.17	-0.33	0.67
σ_{stand}	1.87	1.53	2.27	1.71	1.38	2.12	1.59	1.28	1.95

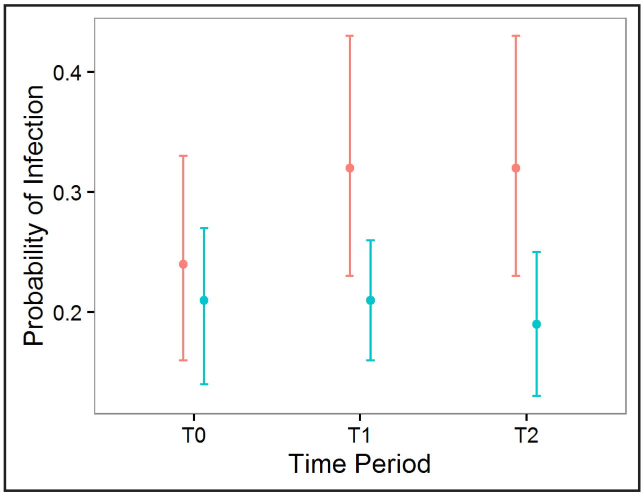

Figure 2. Estimated probabilities of infection with the associated intervals using the occupancy model (red) and the previously reported ratio estimator (blue; Shanahan et al. 2017) for each time period. For the occupancy model, these estimates are based on the intercept for each time period. Both of these statistical approaches are estimates for prevalence of blister rust infection in the GYE, in other words these are not adjusted to a probability of a particular tree having blister rust as related to that tree's characteristics and habitat, but rather are generalized population-level estimates for the GYE.

each covariate (Figures 3-5) with the other covariates at their average value. These plots show each relationship for the average stand (bold line) and the associated 95% posterior bands with the different colors distinguishing each time period (TO-blue, T1-red, T2-green).

Infection probability is estimated to increase as DBH increases in all three time periods (Figure 3). Overall, the relationships are similar across the time periods with substantial overlap between the 95% posterior bands. There is some evidence for the probability of infection being lower in T0 than the other time periods for trees with larger diameters (where the bands do not overlap in Figure 3). Recall that the majority of the trees observed (75%) have a DBH less than 16 cm. Also note that the estimates of σ_b^2 from each time period indicate substantial variability in the baseline infection rate among the different stands. The amount of variability in each time period appears approximately equal.

Based on this model, infection probability decreases as elevation increases (Figure 4). Again, we see similar trends between elevation and the probability of infection across the time periods. There was weak evidence for a meaningful relationship between

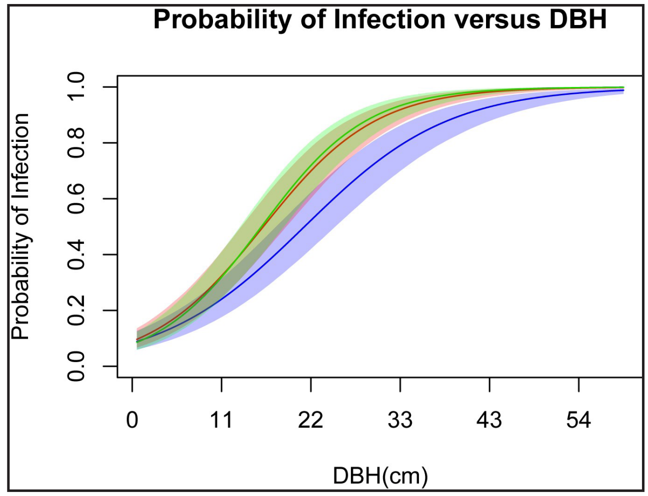

Figure 3. Estimated probability of a tree being infected with blister rust versus DBH. For each time period (blue for T0, red for T1, and green for T2) the posterior mean (bold line) and 95% posterior band are shown for this relationship. An average stand effect and all other covariates at their average values are assumed. The probability of infection is estimated to increase as DBH increases in all three time periods.

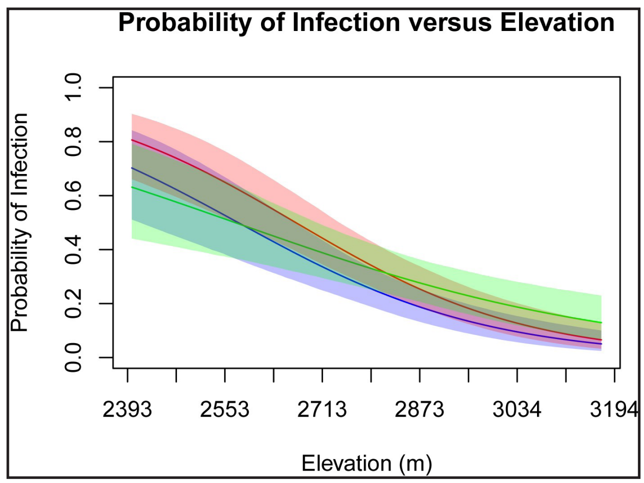

Figure 4. Estimated probability of a tree being infected with blister rust versus stand elevation. For each time period (blue for T0, red for T1, and green for T2) the posterior mean (bold line) and 95% posterior band are show for this relationship. An average stand effect and all other covariates at their average values are assumed. The probability of infection is estimated to decrease as stand elevation increases based on this model.

infection probability and the remaining covariates (topography and number of trees; Figure 5) for these time periods. However, for large numbers of trees there is a lot of uncertainty associated with this relationship because there are not as many transects with a large number of trees.

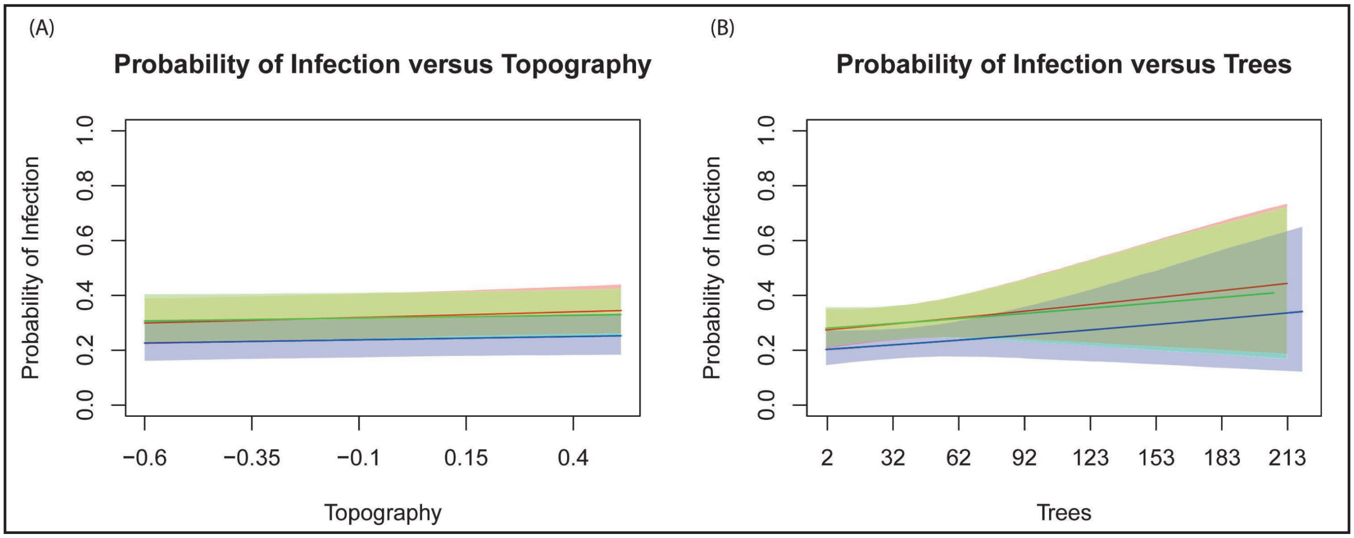

Figure 5. Estimated probability of a tree being infected with blister rust versus stand topography (A) and stand number of trees (B). For each time period (blue for T0, red for T1, and green for T2) the posterior means (bold lines) and 95% posterior bands are shown for these relationships. An average stand effect and all other covariates at their average values are assumed for each plot. Neither stand topography or stand number of trees appears to be associated with the probability of a tree being affected with blister rust.

4.2 Detection

The estimates associated with detectability of blister rust are presented in Table 12. Based on this model, the overall detection probability for the average tree is high (posterior mean around 0.78). There is strong evidence that detection decreases as tree DBH increases (posterior mean of -0.30 with the 95% posterior interval excluding zero). The other parameter estimates show there is also strong evidence that detection of blister rust increases as the slope of the transect increases, but only moderate evidence of negative relationships between hiking time and detection (the majority of the 95% posterior intervals for these parameters exclude zero). There is also moderate evidence that observers' ability to detect blister rust improves after the first year. Figure 6 shows how the probability of detection was associated with these covariates for the average observer. The green lines show first year observers and the blue lines show second year observers. The estimated (based on posterior means) relationships between each covariate when the others are at the average value and the 95% posterior bands for each are shown.

Table 12. This table shows the model estimates associated with detectability for each of the explanatory variables used in the model as well as the random observer effects. These estimates assume detection is constant across time periods (after accounting for the other covariates, differences among observers, and observer improvement).

Parameter	Post. mean	2.5%	97.5%
α_0	1.25	0.88	1.62
α_1(DBH)	-0.30	-0.39	-0.21
α_2(date)	-0.01	-0.12	0.11
α_3(time)	-0.13	-0.30	0.04
α_4(slope)	0.33	0.23	0.44
α_5(trees)	-0.11	-0.34	0.11
α_6(trees^2)	0.32	0.17	0.46
α_7(exper)	0.27	-0.05	0.62
σ_{obs}	0.60	0.32	0.98

The model suggests a quadratic relationship between detection and the number of trees on a transect (Figure 6D), which seems surprising. There is no evidence of an association between Julian date and the probability of blister rust detection. There is evidence of some variability among observers, but this is quite less than the variability of infection rates among stands.

Based on Figure 6D, there appears to be a slight decrease in detection as the number of trees initially increases. However, for transects with a very large number of trees per transect, detection is estimated to be higher. Note that the majority of transects (75%) have less than 33 trees, however. The evidence for an increase in detection could be due to higher detection rates for transects with many trees as a result of unmodeled covariates.

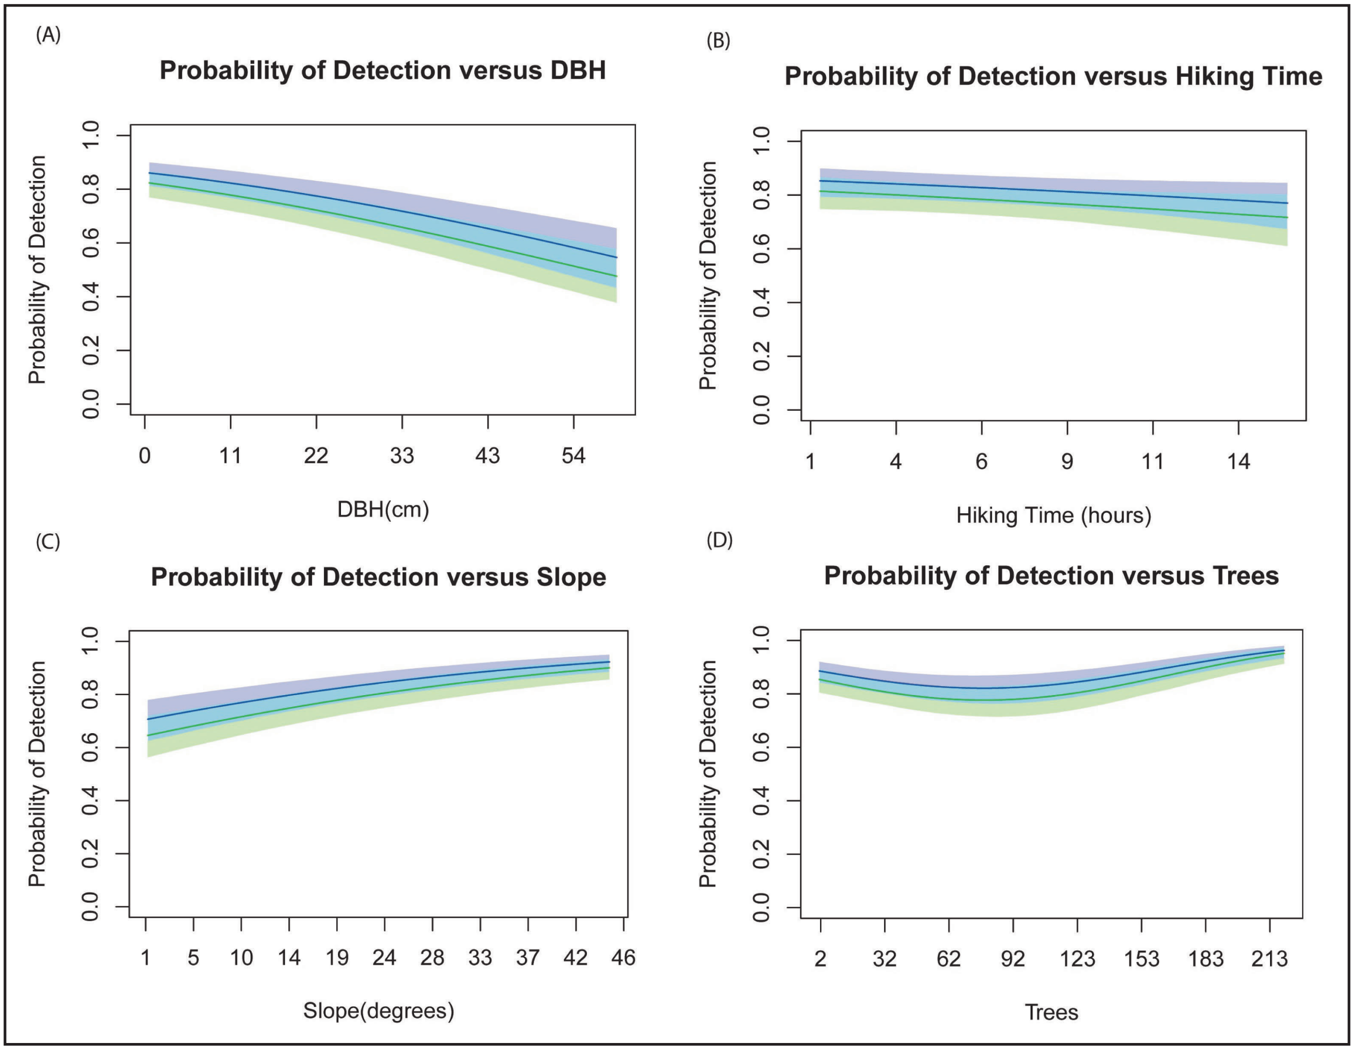

Figure 6. Estimated probability of detecting an infection versus tree DBH (A), hiking time (B), slope (C), and number of trees (D). For all of these relationships, the posterior means (bold lines) and 95% posterior bands are shown for first year observers (green) and experienced observers (blue). In each plot, all additional covariates were assumed to be at their average values.

5 Discussion

5.1 Blister Rust Detection

Estimating blister rust detection serves two main purposes. The first is understanding the detection process and how it is influenced by covariates could help improve study design. The second, and more important, is to account for imperfect detection when obtaining estimates for infection (or prevalence of blister rust in the GYE). Understanding how detection is influenced by tree and stand characteristics can aid in the training of observers. Observers should be aware that these covariates can impact detection probabilities. Being mindful of these limitations could help focus observers on ways to improve detection. However, detection of blister rust will never be perfect and there will be a need to account for it when estimating prevalence. The relationship between the number of trees on the transect and probability of detection was not expected – in particular that detection was estimated to be higher on transects with many trees. This relationship could be difficult to interpret because the model is no longer differentiating between aecia and secondary indicator methods of detection (see more in a later section) or because of additional unmodeled covariates (as mentioned above).

The second purpose of modeling blister rust detection is crucial because it allows unbiased estimates for the parameters used to address main objectives of the monitoring program. Without accounting for imperfect detection, estimates for the proportion of trees infected will likely be biased lower than the true value. While improving detection probabilities is beneficial, if modeled appropriately, detection of blister rust can be low without negative consequences on the parameter estimates of interest. In comparison to the ratio estimators, which ignore imperfect detection, the infection point estimates for each time period from the occupancy model are all higher but there is still some overlap in the associated interval estimates. This is likely due to the overall high detection rates which result in a relatively small amount of bias for prevalence estimates because bias decreases as detection probabilities increase. Note, however, that unless detection is perfect there will always be some bias in estimates of prevalence when detection is ignored.

Imperfect detection can also impact estimates of the relationships between prevalence and covariates of interest. For instance, based on the analysis above, the probability of detection decreased as tree DBH increased while the probability of infection increased with that variable. Assuming detection was perfect could suggest that there was actually no relationship between infection and DBH (or that the relationship was not as strong). MacKenzie et al. 2006 explain that "the bias results as logistic regression simply model the relationship between habitat and where the species is *found* (a combination of occupancy and detectability), not where the species *is* (occupancy)" (pg. 34). For T1 and T2, the differences between the overall prevalence estimates from the occupancy model and the ratio estimators are much larger than that of T0. This illustrates another important benefit of utilizing occupancy models to account for imperfect detection when estimating blister rust prevalence. Changes in prevalence are confounded with changes in detection when an occupancy model is not used. Even though the parameters describing detection are assumed constant over time, average detection rates may change due to changes in the size distribution of whitebark pine trees or observer turnover. The results of the occupancy model analysis indicated that both of these variables are related to detection and ignoring them could mask trends in prevalence as their distributions change over time.

Additionally, this analysis focused on overall infection classification, but agreement (and detection) could be lower if examining trees by location or using each detection method independently. This was suggested by the multiple observer agreement summaries. These additional sources of observer error were not a concern in this analysis because we aggregated observations to the entire tree level instead of by tree location. However this variability could be important in analyses that address other questions of interest concerning blister rust infection severity which depend on canker locations as well.

5.2 Comparing Single and Multiple Observer Transects

For the most part, the multiple observer transects are similar in characteristics to the single observer transects. This suggests that the multiple observer

transects should provide detection estimates which are representative of all transects in the study. Transects with a larger number of trees were included on the multiple observer transects in the most recent field season, but previously had only been included in the single observer transects. Based on the comparison of this variable, the single and multiple observer data appear to be similar to one another (and also cover the same range of values), but future field seasons should continue to visit transects with a large number of trees with multiple observers. Note that this analysis pooled observations across time periods in order to estimate detection probability parameters. Attempting to perform completely separate analyses for each time period would not have met this standard. In other words, the analyses for T0 and T1 would have had to assume that the association between detection and number of trees was unchanged for the transects with a very large number of trees even though those trees were not observed by multiple observers for those years.

We suspect that this could be part of the reason the separate analyses did not work for the T1 time period. This emphasizes the importance of collecting adequate data on the multiple observer transects in order to estimate detection and be able to assume that those estimates of detection are applicable to the single observer transects. Achieving this within each time period would be advisable for subsequent seasons. While the assumption that the detection parameters are constant across time periods seems justified (and separate estimates for T0 and T2 indicated this), that may not be the case for subsequent years. It may be important to be able to assess this more thoroughly in future analyses.

5.3 Possible Model Extensions

Including the location of the infection (in the canopy, bole, or both) and different detection probabilities for aecia and secondary indicators are ways the above model could be extended. This more complicated approach could provide a more realistic description of the detection process as well as providing more detailed information about blister rust prevalence and where it is occurring within trees. While these aspects of the model may be of interest to address some research questions, the model used does not need to include these additions if the primary interest is making inference about the rate of overall blister rust infection. Here we focused on the simpler model because it is more likely that this version will ultimately be used when estimating the proportion of trees infected with blister rust.

However, not including the different detection methods could make the interpretation of the parameters associated with detection more difficult. This is because now the coefficients are describing the relationships between overall detection (either method) and each variable. These relationships could be different for detecting aecia and detecting the secondary indicators. Based on the current analysis, it is impossible to describe how these covariates are related to aecia or secondary indicators individually. For instance, it is expected that time of year would primarily impact the probability of detecting aecia but not influence the probability of detecting the secondary indicators. While this model could make interpretation of the detection coefficients more difficult, these additional model components would not change the estimates for infection or how these are interpreted.

Another way the occupancy model could be expanded is to include stand random effects for the probability of detection. There could be additional sources of heterogeneity in the detection probabilities not accounted for by in our model. This includes additional covariates (weather related or other stand characteristics) as well as the possibility of detection within a stand being related to the overall infection rate within the stand. In other words, the stand random effects for occupancy included in the model would be correlated with the stand random effects for detection. We did not explore the possibility for these effects but they could be incorporated into future analyses.

Finally, our approach assumed the only errors being made by observers were due to imperfect detections resulting in false negatives when recording blister rust detection for a tree. For these data this assumption appeared reasonable, but we acknowledge that incorrectly recording a healthy tree as infected is possible (false positives). This would most likely be due to identifying the necessary secondary indicators as a result of other factors but incorrectly attributing them to be the result of a blister rust infection. Models have been developed to also account for false positives in an occupancy framework. This could be an important extension because even low false positive rates have been shown to bias estimates

(Royle and Link 2006). More work would be needed to explore whether false positives are a concern for these data or not.

5.4 Preliminary Ideas for Multiple Observer Sampling

In order to help provide guidance on the multiple observer sampling protocol, we discuss some general study design recommendations from MacKenzie et al. (2006). These recommendations focus on the allocation of survey effort in occupancy studies and typically focus on scenarios of constant occupancy (ψ) and detection (p) probabilities. Since these assumptions are not realistic and the blister rust monitoring program is quite different from typical occupancy studies, some of these recommendations may not be applicable or logistically feasible. However, they will provide some general guidelines for the multiple observer transects that can help inform the sampling protocol. These recommendations generally attempt to minimize the standard error associated with the estimates of ψ (which is done by appropriately estimating detection). The authors acknowledge that the "optimal" study designs they present will not be possible for many studies.

MacKenzie et al. (2006) generally recommend the use of a "standard design" where each site (which is a tree in this monitoring protocol) is visited the same number of times for the replicate surveys (i.e., same number of multiple observers). Based on this design, Table 6.1 in MacKenzie et al. (2006) shows the optimal number of revisits to each site. The optimal number ranges from a low of 2 revisits up to 34. Once a reasonably high probability of detection is reached ($p>0.4$), only two or three revisits are needed for an optimal design. MacKenzie et al. later say that despite their table 6.1, they generally "suggest that at least three surveys per site be conducted" (pg. 180) because typically covariates for detection are of interest (where estimation requires more repeat surveys). Again note that this is assuming all sites are visited an equal number of times. Even though this isn't feasible for the blister rust monitoring, taking into consideration the general recommendation of at least three observers per tree (on multiple observer transects) could be beneficial. Based on the current data, all multiple observer transects only had two observers per tree after T0. This could be another reason the separate analysis of T1 resulted in estimation issues – there were not enough repeat visits to a tree to reliably estimate how detection was associated with the various covariates of interest. Again, three observers per tree on the multiple observer transects may not be feasible given the logistical constraints of these surveys, but is recommended by MacKenzie et al. (2006).

Since blister rust monitoring protocol utilizes a portion of trees which are observed only once and another portion which are observed by multiple observers, this is classified as a "double sampling" design. Unless the probability of detection is very high (over 0.8), this design is not optimal (i.e., better to have every site revisited) based on the criteria described by the authors. They state that the reason the double sampling design is not optimal (usually) is that there is more uncertainty associated with whether a tree is infected or not compared to actually estimating parameters associated with detection. This information means that overall detection is fairly high for blister rust (i.e., the double sampling design is reasonable). Additionally, the standard errors from this monitoring program will be small because of the large number of trees being examined.

Another thing to note is that estimation of detection probabilities only uses the multiple observer trees where at least one person identified infection. In other words, multiple observer transects where infection is very low will not contribute much information to estimate detection despite the increased survey effort. Our previous suggestions on the multiple observer sampling protocol focused on sampling so that the range of potential detection covariates was covered. This should still be a priority, but may be done more efficiently than by randomly selecting entire transects to be observed by multiple observers. It would be completely acceptable if a portion of a transect is examined by more than one observer and the remaining trees have single observers. This means that instead of selecting entire transects for the multiple observer protocol, we only need to select a portion of trees within transects (assuming transects are always visited by at least two observers at once). This will still provide data to estimate detection and its relationship with covariates. Since most of the covariates vary at the stand level, if this is done for every transect (or most) then the range of detection covariates will be adequately covered without relying on random selection of transects. Another option could be to stratify based on these covariates and then sample based on this stratification in order

to ensure the multiple observer transects cover the ranges of each covariate.

From the past survey years, typically 5-20% of the trees observed are examined by multiple observers. Spreading this same level of effort across all the transects could be a more efficient protocol than intensely surveying a few transects with multiple observers each year. This systematic approach would guarantee that the ranges of covariates (mentioned above) are represented in the multiple observer data. This protocol could be performed by utilizing multiple observers per tree on the first (or last/ middle; could randomize this) few trees on each transect. How many trees are included in the multiple observer set depends on what is feasible based on the other logistical constraints of the field protocol (somewhere in the 5-20% range would match the effort from previous years).

5.5 Conclusion

Imperfect detection is a common component of ecology research and is evident in blister rust monitoring in whitebark pine. In order to address the main objectives of the monitoring program, this imperfect detection should be accounted for and modeled appropriately. This report shows how this could be done to obtain estimates for blister rust prevalence in trees. The apparent imperfect detection of blister rust can be statistically adjusted by using the methods explored here. Although proper training of observers is paramount for long-term monitoring, detection of blister rust will never be perfect. The next steps for the analysis of these data include working to incorporate detection of blister rust into the overall estimate of the proportion of trees infected with blister rust.

6 Literature Cited

Campbell, E. M. and J. A. Antos. 2000. Distribution and severity of white pine blister rust and mountain pine beetle on whitebark pine in British Columbia. Canadian Journal of Forestry Research 30:1051-1059.

Greater Yellowstone Whitebark Pine Monitoring Working Group (GYWPMWG). 2011. Interagency whitebark pine monitoring protocol for the Greater Yellowstone Ecosystem, version 1.1. Greater Yellowstone Coordinating Committee, Bozeman, Montana.

Huang, M. 2006. A statistical analysis of observer variability in the identification of blister rust infection occurring in white-bark pine monitoring. Department of Mathematical Sciences, Montana State University, Bozeman, Montana.

MacKenzie, D. I., J. D. Nichols, J. A. Royle, K. H. Pollock, L. L. Bailey, and J. E. Hines. 2006. Occupancy estimation and modeling: Inferring patterns and dynamics of species occurrence. Academic Press, New York.

National Park Service. 2016. Multiple observer and single observer data provided to Kathi Irvine and Wilson Wright, Northern Rocky Mountain Science Center, U.S. Geological Survey. Interagency Whitebark Pine Long-Term Monitoring Database - accessed January 29, 2016. NPS Greater Yellowstone Network, Bozeman, Montana.

National Park Service (NPS). 2016. Data compiled from master database for full surveys 2004 through 2015 provided to Kathi Irvine and Wilson Wright, USGS Northern Rocky Mountain Science Center. Interagency Whitebark Pine Long-Term Monitoring Database. Accessed January 29, 2016. NPS Greater Yellowstone Network, Bozeman, Montana.

Shanahan, E., K. Legg, and R. Daley. 2017. Status of whitebark pine in the Greater Yellowstone Ecosystem: A step-trend analysis with comparisons from 2004 to 2015. Natural Resource Report NPS/GRYN/NRR—2017/1445. National Park Service, Fort Collins, Colorado.

Royle, J. A. and W. A. Link. 2006. Generalized site occupancy models allowing for false positives and false negative errors. Ecology, 87: 835-841.

Appendix A

Traceplots for the parameters of interest are shown on the following pages. All of these suggested adequate convergence. For this analysis, 4 chains were used (with random initial values for each). Each chain had a 500 iteration burn-in followed by 1,000 iterations which were saved for inference.

The estimated densities associated with each chain for all the parameters were also examined (not shown). These suggested well behaved posterior distributions for all parameters and that each chain indicated similar posterior distributions for the parameters. A cursory examination of residual plots (not shown) for occupancy state was also used to evaluate whether there was evidence of unmodeled patterns with these covariates. These plots indicated no patterns and suggest an adequate fit of this model based on these assessments.

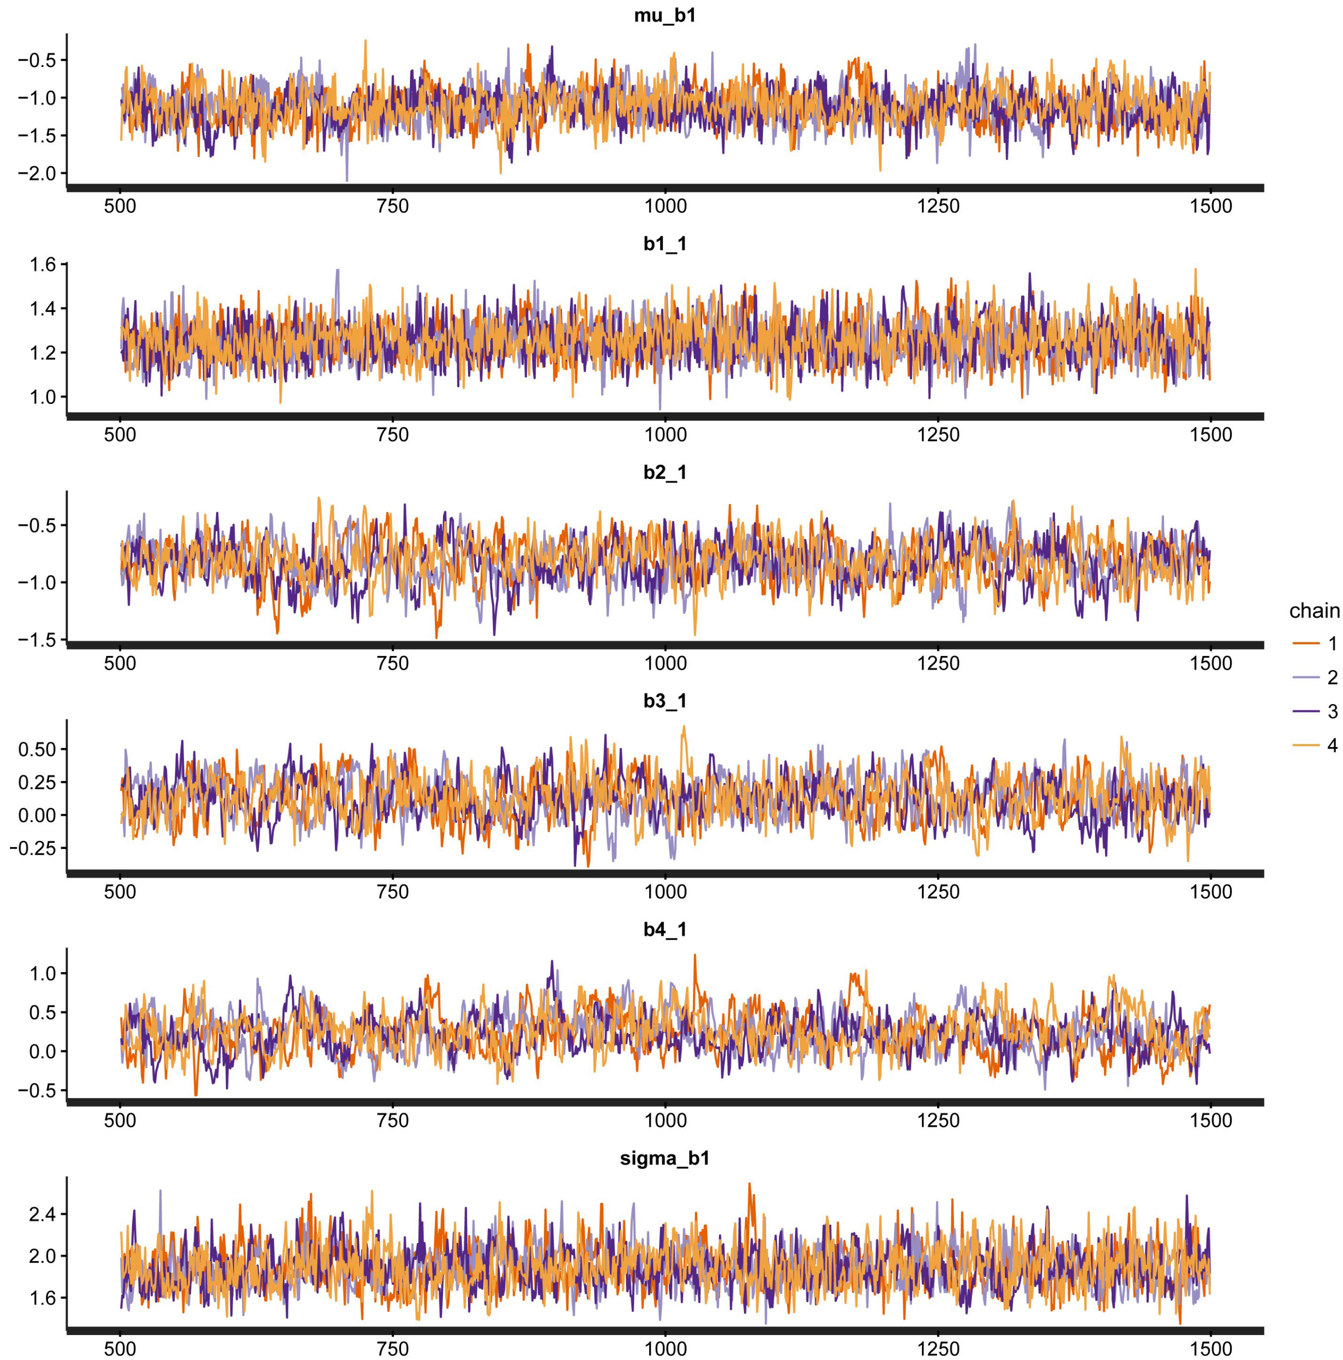

Figure A1. Traceplots of the saved iterations for the parameters describing the probability of tree infection in the T0 time period.

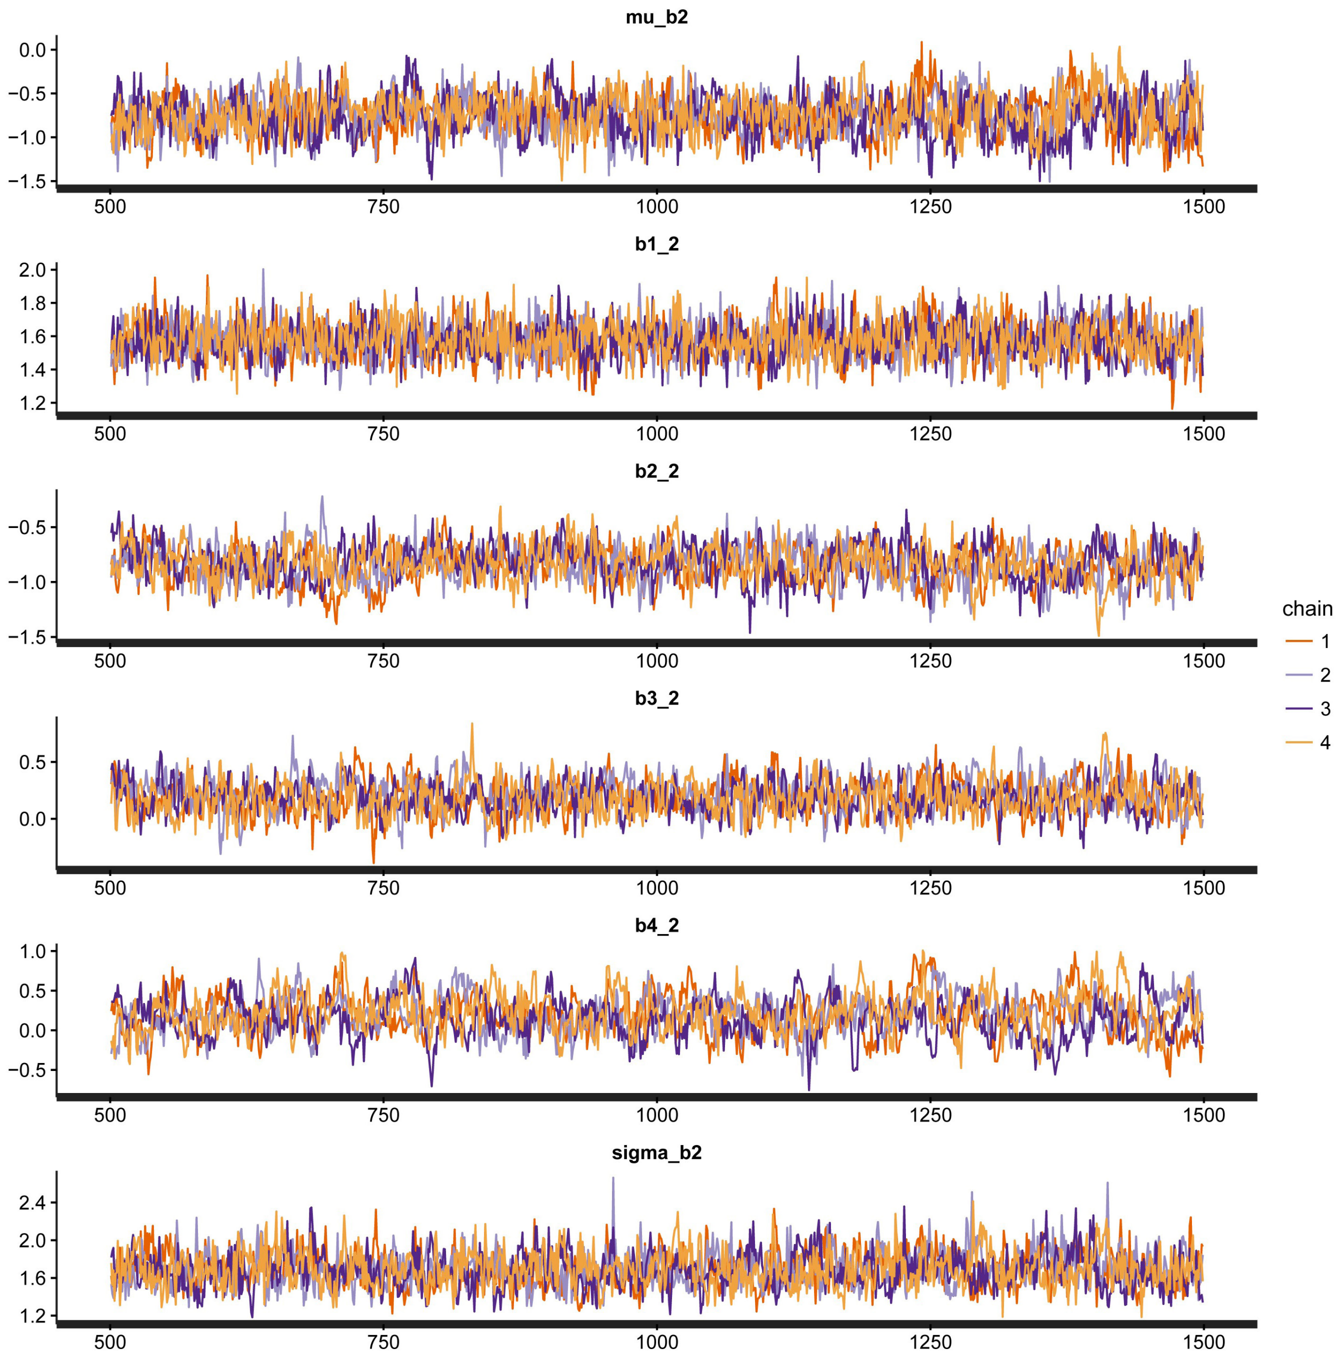

Figure A2. Traceplots of the saved iterations for the parameters describing the probability of tree infection in the T1 time period.

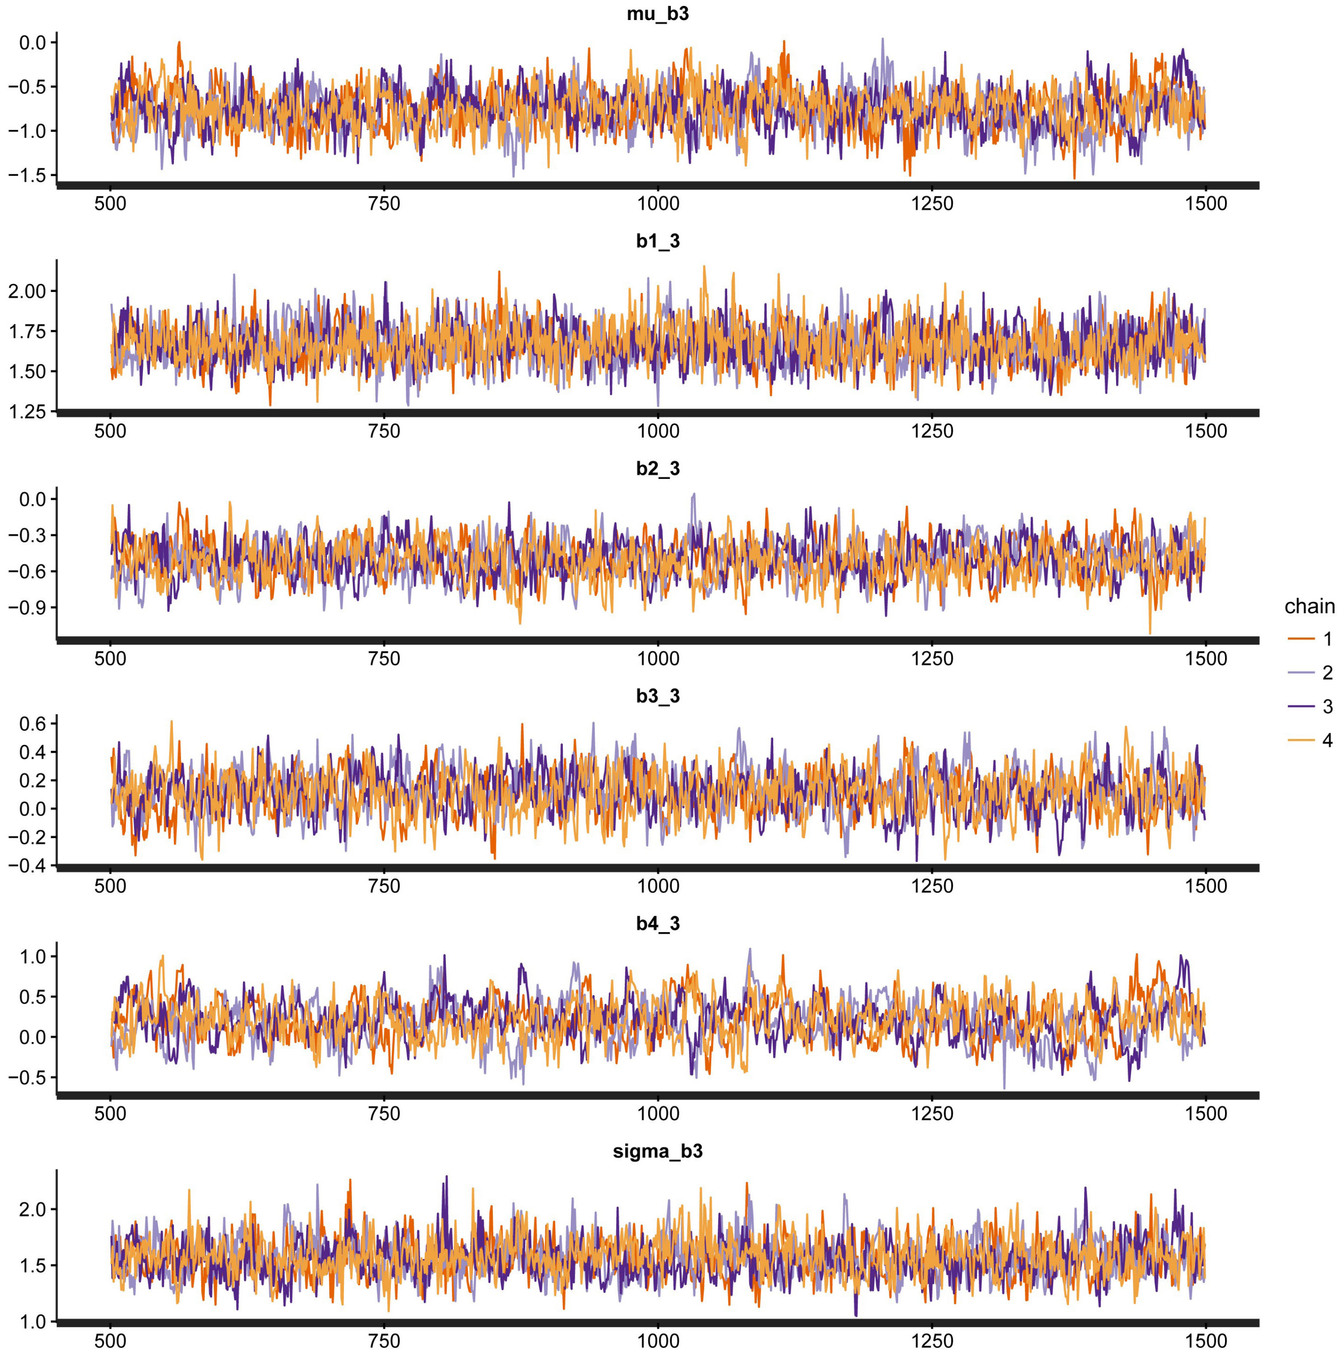

Figure A3. Traceplots of the saved iterations for the parameters describing the probability of tree infection in the T2 time period.

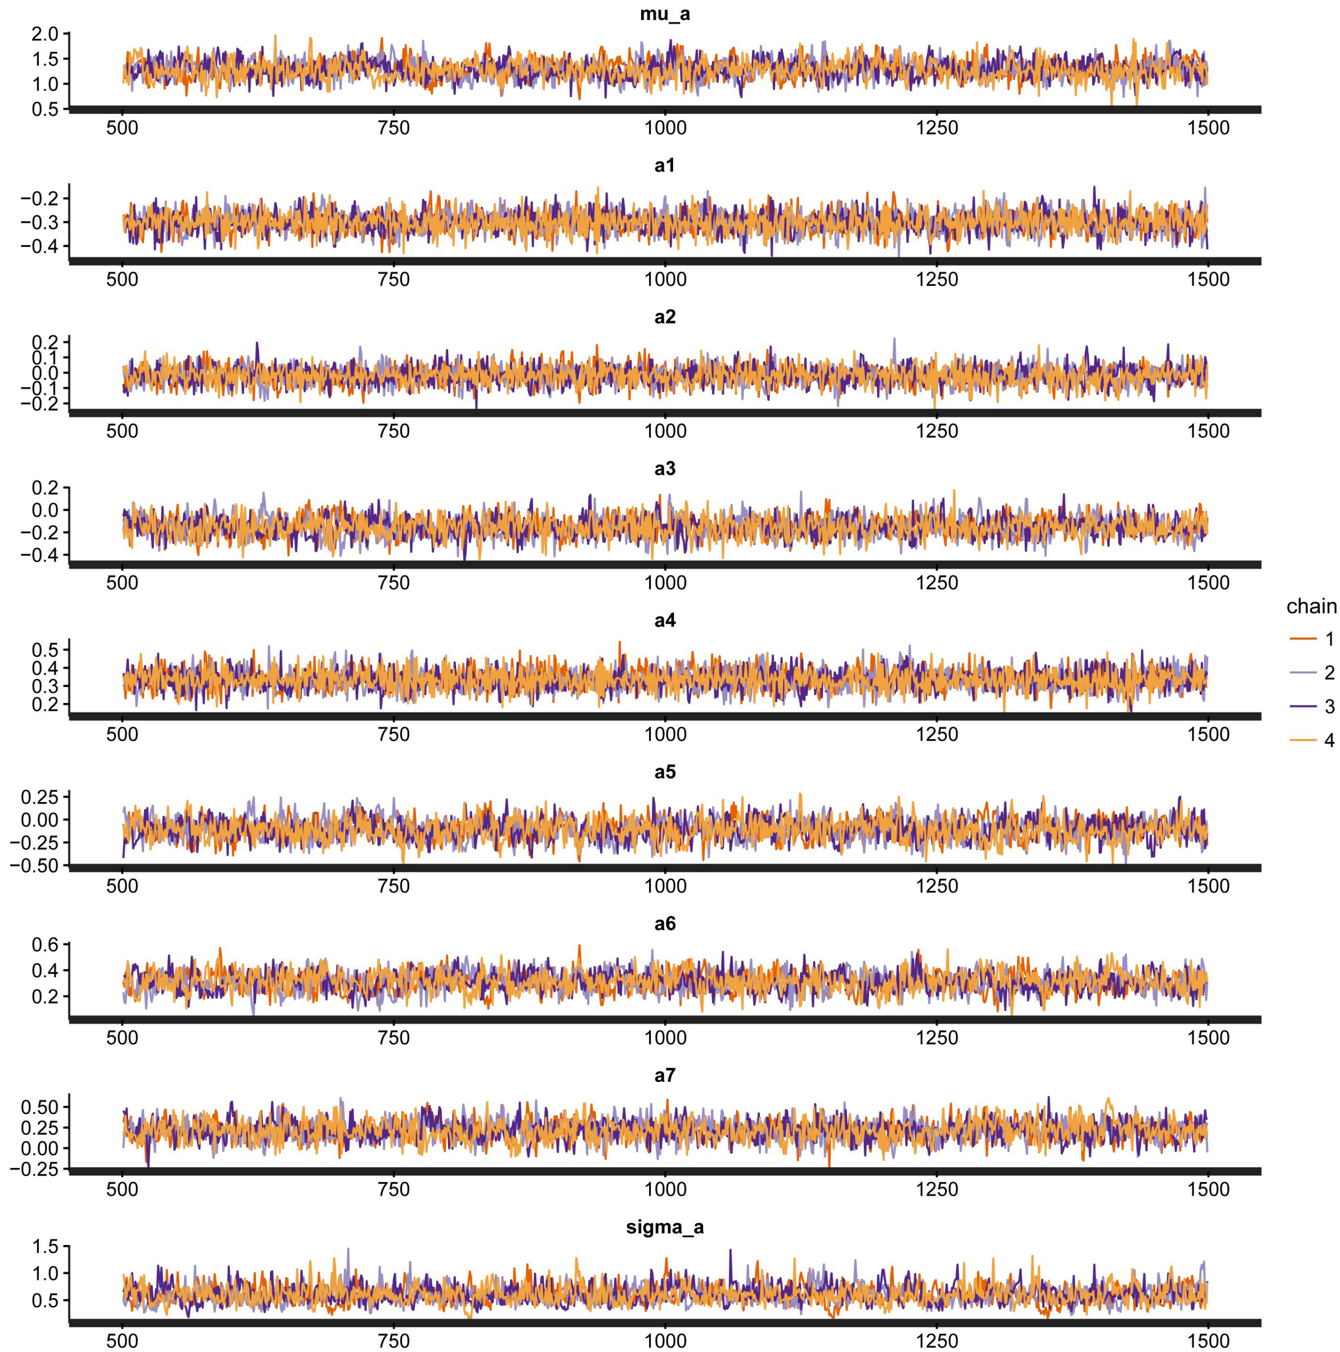

Figure A4. Traceplots of the saved iterations for the parameters describing the probability detecting a blister rust infection.

NPS 101/138712, 136/138712, June 2017

National Park Service
U.S. Department of the Interior

Natural Resource Stewardship and Science
1201 Oak Ridge Drive, Suite 150
Fort Collins, Colorado 80525

www.nature.nps.gov