Appendix A: Supplemental examples of multispatial CCM analyses for: Spatial "convergent cross mapping" to detect causal relationships from short time-series

Adam Clark *et al.*
Department of Ecology, Evolution, and Behavior
University of Minnesota, Twin Cities
St. Paul, MN 55108, USA
adam.tclark@gmail.com

December 23, 2014

Abstract

This Appendix contains detailed information for the real world examples provided in the text, along with some additional examples. The first section describes how to interpret results and diagnostic statistics (**Diagnostics and interpretation**). The second section includes information about the data and methods for each real world example (**Data and methods**). The third section presents the real world examples and interprets relevant results and diagnostic statistics (**Real world examples**). The fourth and final section shows simulation methods and results for a system with bi-directional causality (**Bi-directional causality**).

1 Diagnostics and interpretation

Here, we show how to interpret the results from the same simulated system described in the text in Figure 1, and in the example files for the R package multispatialCCM. In this system, the dynamics of process X influence the

dynamics of process Y, but the reverse is not true. Before we can commence with running the actual CCM test, we first need to find the best embedding dimension for the system (as described in step (1) in the methods section in the main manuscript), and determine whether the system has sufficient information for analysis (step (2) in the methods section) (Figure A1a and A1b respectively). Once the diagnostic requirements have been met, we can move on to implementing the bootstrapped CCM algorithm (steps 3-4 in the methods section of the main manuscript). In order to interpret the results, we also need to ensure that we have completed sufficient bootstrapped iterations to generate stable estimates of the relationship between ρ and library length L (Figures A1c-d).

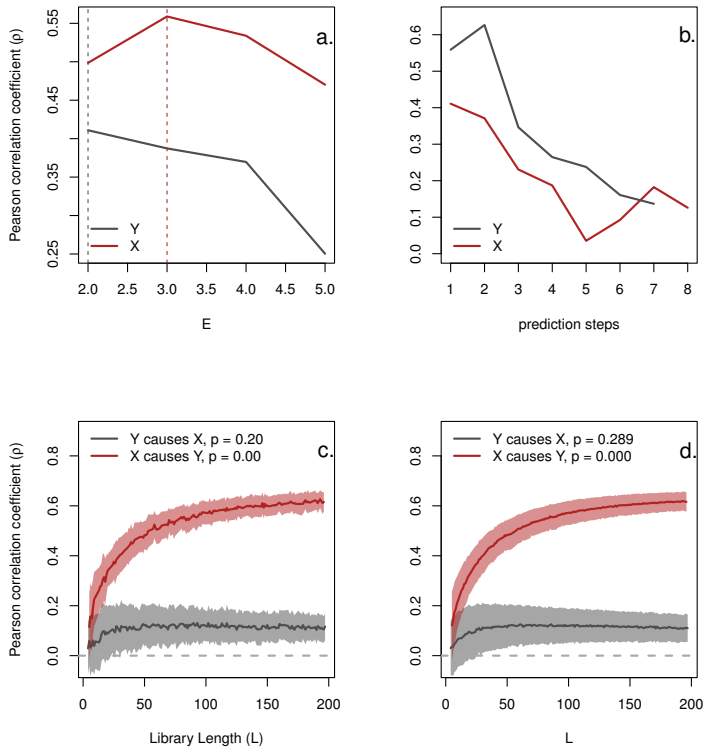

Figure A1: Results from the same simulated example from Figure 1a in the main text. Pearson Correlation coefficient ρ describes the ability of the algorithm to predict system dynamics. Shows test for best embedding dimension E (a), fit vs. prediction time step (b), and CCM results based on 100 bootstrapped iterations (c) and 1000 bootstrapped iterations (d). Solid lines show mean, and shaded region shows +/- 1 standard deviation. See main text for details.

Embedding dimension When choosing the best embedding dimension, we are looking for the dimension corresponding to the highest predictive power for one time step into the future. Ideally, this will either be at the lowest dimension tested (as for process X here, with $E = 3$, Figure A1a), or at an intermediate dimension representing a "hump" in predictive abilities between higher and lower dimensionalities (as for process Y here, with $E = 2$, Figure A1a). However, it is also important not to over-fit the model, and in

some cases it may be prudent to choose a smaller embedding dimension that has moderately lower predictive power than a higher dimensional model. We will discuss this further in the "Real world examples" section.

Prediction steps We conduct this test for two reasons. First, we wish to ensure that we have reasonable predictive ability for short time steps, or else there is likely not enough information (or too much stochastic noise) present in the data for CCM to work. Second, the predictive power should drop as the length of the prediction time step increases. Both of these conditions are met for the simulated data (Figure A1b). If this is not the case, then the system is probably purely linear, and CCM should not be used (see manuscript for other suggestions). Particularly bad is a system where predictive ability decreases, and then increases with prediction time step (a "U-shaped" ρ vs. time plot), as this suggests either some form of periodicity in system dynamics, or more likely, that data are not uniformly sampled, such that only a few observations are being predicted for longer prediction intervals. We will discuss this further in the "Real world examples" section. It should be noted that this is a "necessary but insufficient" test for non-linearity, since a linear system dominated by stochastic noise could still show the same response. In cases dominated by stochastic noise, CCM should (correctly) return "no causal link in either direction", since increasing information about the system will not increase predictive power.

Bootstrap iterations We increase the number of bootstrapped iterations for two reasons. First, it increases the precision of p-value estimates, as the lower detection limit is determined by 1/(number of iterations). This is not so important for the tests that we present here, as we only need to show $p<0.05$ for significance tests. Second, and more importantly, we need to increase iteration number in order to reduce Monte Carlo stochasticity (i.e. significantly different results from the same simulation parameters). The "best practice" for testing this is to increase the number of iterations until the mean and standard deviation estimates for ρ stabilize. Here, we show results for 100 and 1000 bootstraps (Figures A1c and A1d respectively). Though 1000 replicates provides a smoother estimate, there is no significant change in the distribution of rho estimates between the two iterations, so we chose to use only 100 iterations for our simulated performance tests in the manuscript. For real data, more iterations are generally required.

2 Data and methods

All data and detailed methods for the real world examples that we present in the main text and below are available on the Cedar Creek LTER data webpage at
`www.cbs.umn.edu/explore/field-stations/cedarcreek/research/data`.
Here, we give brief descriptions of each dataset that we use, and discuss existing findings and knowledge about the direction of causality for those systems. The experiment numbers that we reference can be used to look up each data set in the Cedar Creek webpage. For convenience, we also include ".csv" files that pre-collate relevant data tables and insert gaps in the time series between plot observations. These files can be used without further modification to run the analyses that we present here.

E001: Nitrogen addition experiment This experiment was established in 1982 to test how long-term nitrogen addition would alter plant species composition in an old field community. For full methods, see [8]. We analyze data from three fields (A, B, and C) and four nitrogen treatment levels, which we have grouped into "low" (0 $g\mathrm{NH_4NO_3}yr^{-1}$, with and without micronutrients), and "high" (25 and 40 $g\mathrm{NH_4NO_3}yr^{-1}$) treatments. Each field contained 12 plots in each treatment category. Above-ground species-level biomass in one field (C) was sampled annually from 1982-2011, while the others were sampled annually from 1982-2004. Because of some missing data, we retained 11 "plots" with 30 sequential measurements each, and 24 "plots" with 23 measurements each for the analyses presented here.

We focus on the dynamics of two of the most common species in the experiment, the cool-season, early-successional grass *Elymus (Agropyron) repens*, and the warm-season, late-successional grass *Schizachyrium scoparium*. In the absence of nitrogen fertilization, *A. repens* typically arrives early in succession, but is replaced by *S. scoparium* in the first 30 years of succession (see Figure 1 in [9]), because *S. scoparium* is a superior competitor for nitrogen [16]. However, in the presence of high levels of nitrogen addition, *A. repens* tends to re-colonize plots and out-compete *S. scoparium*, either for light or through indirect effects of leaf litter accumulation [3]. Fields A, B, and C were abandoned around 1968, 1957, and 1934 respectively. Thus, in "low" addition plots, we would expect to find that *A. repens* should be

disappearing or absent in all fields, and that its dynamics should be forced by those of *S. scoparium*. Conversely, in "high" addition plots, we expect that *S. scoparium* should be declining through time, and that its dynamics should be forced by *A. repens*.

Considerations for E001 analysis

- System is non-stationary (Figures A4, A6)
- Accurate prediction requires high embedding dimensions (Figures A3, A5)
- "U-shaped" prediction strength vs. interval relationship (Figures A3, A5)
- Conclusion: CCM should not be applied to these data, results are suspect

E026: Competition plots on a soil gradient This experiment was established in 1986 to test resource reduction and competitive outcome for monoculture and two-species prairie plant mixtures grown across a soil nitrogen fertility gradient. Full methods are available in [16]. Here, we focus on two soil fertility levels ("high" and "high + NH_4NO_3 fertilizer") in monoculture plots of the cool-season, early-successional grass *Agrostis scabra*. Aboveground biomass data for both species and leaf litter was collected annually from subsets of plots between 1986 and 1993. For *A. scabra* grown in "high" treatments, this included 12 plots with 8 sequential samples, 6 with 7, 3 with 4, and 9 with 3. For "high + NH_4NO_3 fertilizer" treatments, this included 4 plots with 8 sequential observations, 2 with 7, 1 with 4, and 3 with 3.

When grown on rich soils, *A. scabra* produces copious amounts of leaf litter. Because plots in this experiment were never burned, populations grown on very fertile soils (or soils with added fertilizer) show strong oscillations, possibly following chaotic dynamics, and ultimately crash, likely because they smother themselves in their own leaf litter [14]. Thus, we expect to find that *A. scabra* forces litter dynamics in less rich soils, but that it should also be forced by litter dynamics for very fertile soils.

Considerations for E026 analysis

- Different soil treatments follow different dynamics (Figure A7)
- Suspect ρ vs. prediction steps relationship for leaf litter (Figure A9a)
- Insufficient data for the case with added fertilizer (Figure A8b)
- Conclusion: Analysis in Figure A8a is likely legitimate, analysis in Figure A8b requires more data

E054: Plant biomass in old fields E054 is a subset of the long-term observational E014 study of old field succession at Cedar Creek. Full methods for E014 can be found in [2]. E054 includes annual species-level above-ground biomass samples taken from 1988-2011, taken at four subplots in each of 15 fields. Fields were abandoned between 1927 and 1998. Mean successional dynamics follow those described in [9]. Because of incomplete sampling and staggered abandonment of fields across years, we include data for 45 plots with 24 sequential samples, 2 with 23, 3 with 22, 4 with 20, 1 with 15, 4 with 11, 1 with 8, and 2 with 3.

Though there are many potential ecological questions to test in this system, we use these data to test a rather simple hypothesis about precipitation. Using total summer annual precipitation data (June - August), we test the extent to which *A. repens* and *S. scoparium* are forced by (or force) precipitation patterns at Cedar Creek. Because *A. repens* is a cool-season drought-intolerant species, its dynamics should strongly depend on water availability, and it should therefore be forced by precipitation. *S. scoparium* is a warm-season drought-tolerant species, and should therefore be more resistant to water stress, and its dynamics should not be as strongly forced by precipitation [11]. In both cases, precipitation should be an exogenous process, and its dynamics should not be forced by those of either species.

Considerations for E054 analysis

- Embedding dimension for both *A. repens* and *S. scoparium* should be reduced below the "best-fitting" E to preserve sample size and avoid over-fitting (Figure A10a-b)
- Conclusion: Both analyses in Figure A11 are likely legitimate

E120: "Big Biodiversity" experiment This experiment is the longest-running randomized test for the effects of plant diversity on ecosystem functions. Plots were established in 1994 and planted with 1, 2, 4, 8, or 16 species, and have since then been sampled annually for above-ground plant biomass. Full methods are described in [12]. The most well-known result from the experiment is that planted species number strongly, positively influences above-ground biomass production. However, because the diversity treatments are fixed, rather than dynamical variables, they do not lend themselves to CCM analysis. Instead, we focus on three other published results from the experiment: soil nitrate effects on invasion by non-planted species, biomass effects on soil nitrate, and biomass effects on insect abundance and diversity.

A number of studies in E120 have found significant increases in insect diversity as a function of increased planted species richness [4, 5]. A posited cause of this is that increased plant diversity increases above-ground biomass, which in turn increases the foraging space and habitat structure available to insect species. Interestingly, results do not agree on the effects of above-ground biomass on insect abundance. In one case [5], diversity was found to have no significant influence on insect abundance, whereas another study [4] found significant effects of both planted diversity and above-ground biomass on insect abundance. Here, we test for the causal relationships between above-ground plant biomass and both insect species richness and abundance. Across all diversity levels, we included 162 plots with 5 observations each, 136 with 4, and 24 with 3.

Though soil nitrate was not sampled as frequently as above-ground biomass, most plots were measured for 7 sequential years between 1996 and 2002. Empirical and theoretical results show that soil nitrate levels should be reduced in high diversity mixtures compared to low diversity mixtures because of more complete utilization of niche space [12, 13]. Additionally, increased biomass is associated with decreases in soil nitrate levels, even in monocultures [15]. Consequently, we expect that above-ground biomass dynamics should influence soil nitrate dynamics. However, because species are hypothesized to maintain soil nitrate at a relatively constant level regardless of their own biomass, classical models of resource competition [6, 7] suggest that above-ground biomass should not be influenced by soil nitrate dynamics. Across all

diversity levels, we included 132 plots with 7 sequential observations, 4 with 5, and 41 with 4.

High diversity has long been associated with decreases in invasion success. Though there is much debate about this relationship in natural systems, decreased invasion by non-planted species as a function of increased planted species richness has been described by a number of studies in E120 [1, 4]. A posited mechanism for this is soil nitrate: increased diversity leads to decreased soil nitrate, which in turn reduces invader success [10]. Based on diagnostic plots of system dynamics (details in the "Real world examples" section), we combined planted diversity treatments into "low diversity" (1-2 species), "intermediate diversity" (4-8 species), and "high diversity" (16 species) for this analysis. The "low diversity" treatments had 7 sequential samples in 57 plots, 5 in 4, and 4 in 11. The "intermediate diversity" treatments had 7 in 45, and 4 in 13. The "high diversity" treatments had 7 in 27, and 4 in 7.

Considerations for E120 analysis

- Time series for soil nitrate and insects are too short to test prediction power vs. interval
- Only some combinations of diversity treatments combine to form tractable manifolds for multispatial CCM analysis
- Analyses in Figure A12b and A14a show decrease in predictive ability with library length, suggesting that plots combined in the analysis may be too dissimilar
- Conclusion: Analyses in Figures A12a and A14b are likely legitimate. Analyses in Figure A12b and Figure A14a are suspect.

3 Walk-through of real world examples

For the following examples, we used 1,000 bootstrapped iterations of the multispatial CCM algorithm, after testing for both 100 and 1,000 iterations and finding no significant difference in the mean or distribution of ρ. In each case, we walk through the diagnostic statistics that we used to validate the

CCM analysis, and discuss whether or not the results of the analyses are to be believed. We hope that this section demonstrates that naively applying the multispatial CCM algorithm without checking for continuity of manifolds across plots, proper embedding dimensions, and predictive abilities will rarely yield meaningful results. Just as with any other analysis technique, making sure to meet the assumptions of the method is half the battle.

E001: Nitrogen addition experiment Based solely on the relationship between L and ρ, results for this CCM test suggest that *S. scoparium* dynamics force *A. repens* dynamics in non-fertilized plots (Figure A2a), while neither process forces the other in fertilized plots (Figure A2b). However, both of these analyses have a number of problems that we can pick up with a few diagnostic tests, and both of these results should probably not be believed.

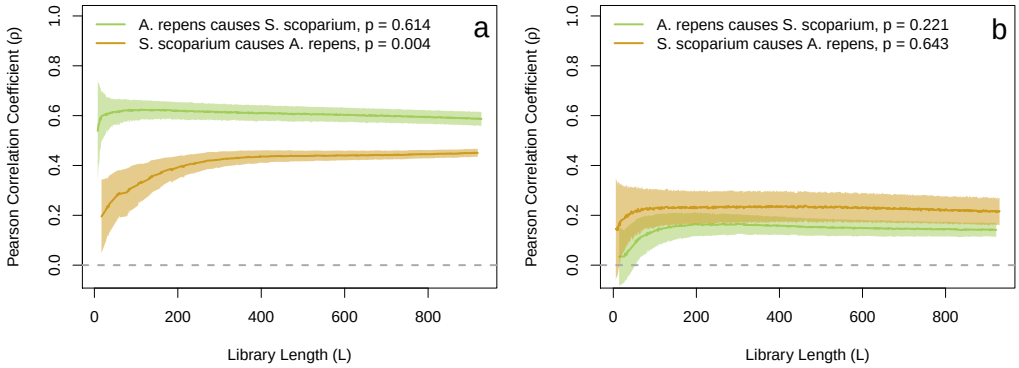

Figure A2: Test of causal forcing between *Agropyron repens* and *Schizachyrium scoparium* above-ground biomass dynamics in (a) unfertilized control plots, and (b) fertilized plots receiving 25 or 40 $g\mathrm{NH_4NO_3}yr^{-1}$.

First, let us consider the diagnostic plots for the non-fertilized case (Figure A3). For *A. repens*, the embedding dimension plot looks alright, with a relatively high predictive power achieved with 6 embedding dimensions. *S. scoparium* is somewhat more problematic because such a high embedding dimension is required to achieve moderate predictive power ($E_{max} = 14$). Even more problematic is the increase in predictive power with prediction interval. Both plots suggest that we can better predict plot dynamics 15 years into the future than we can 10 years into the future, which suggests either cyclical

patterns in the data, or a temporal trend in the data. Separating data by field (not shown) does not alleviate this problem.

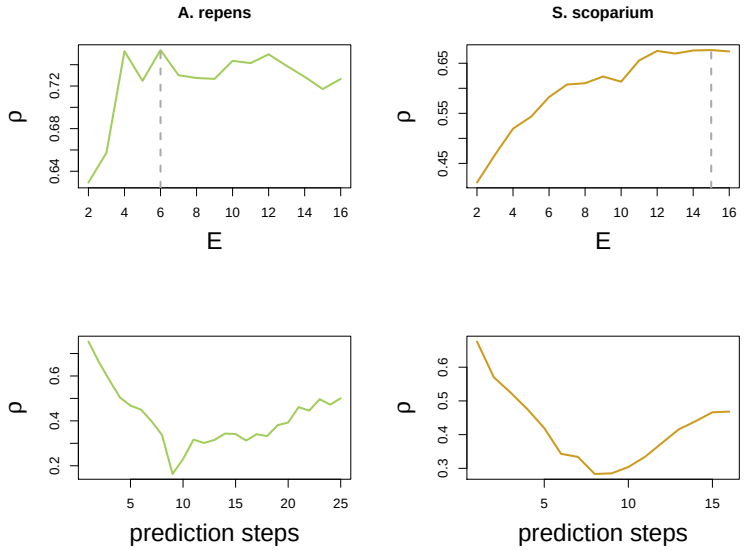

Figure A3: Diagnostic plots for test in Figure A2a (unfertilized control plots), showing relationship between predictive power ρ and embedding dimension E, or length of prediction interval.

To investigate, we can plot the dynamics of each process in two lagged dimensions (i.e. *A. repens* population this year, vs. next year, Figure A4).

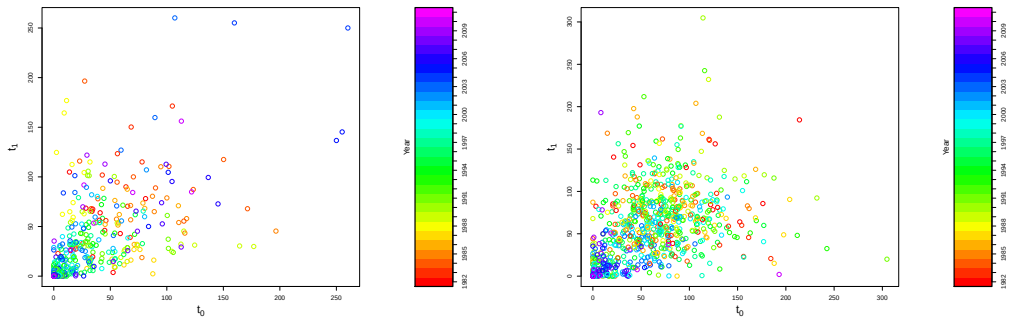

(a) *A. repens* (b) *S. scoparium*

Figure A4: Lagged population dynamics for unfertilized plots in two dimensions.

This reveals a potential cause of the problem – both dynamics have a slight temporal trend. *A. repens* is abundant only early and late in the time series (in the early 1980's, and again in the 2000's). Similarly, *S. scoparium* is relatively rare late in the time series (after about 2000). Consequently, dynamics are much easier to predict across long intervals, as long intervals always predict population sizes late in the time series. The problem is even more apparent in the fertilized plots. Here, we find similar patterns in the diagnostic plots (Figure A5), and an even more rapid decline in *S. scoparium* (Figure A6).

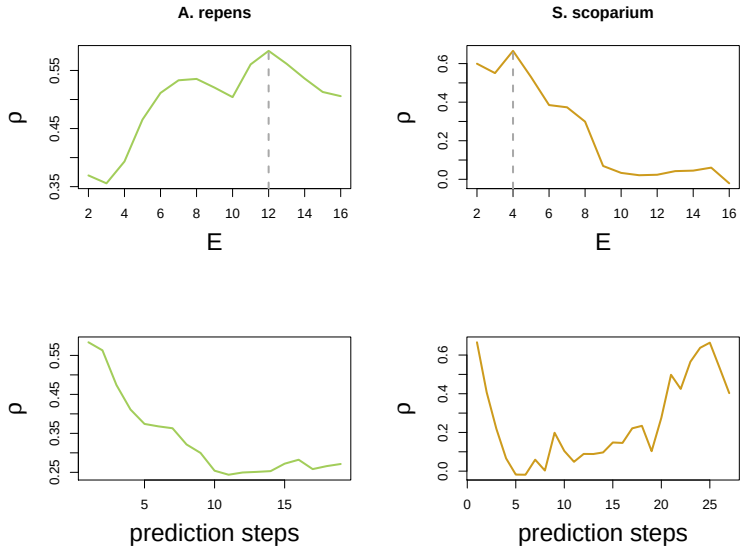

Figure A5: Diagnostic plots for test in Figure A2b (fertilized plots), showing relationship between predictive power ρ and embedding dimension E, or length of prediction interval.

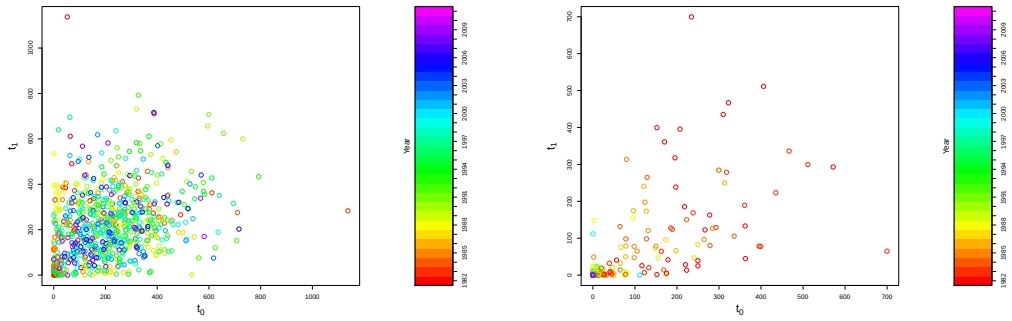

(a) *A. repens* (b) *S. scoparium*

Figure A6: Lagged population dynamics for fertilized plots in two dimensions.

In cases where the time trend is caused by an exogenous variable (e.g. climate change), this problem might be alleviated by detrending the data.

However, in this case the dynamics are likely driven by simple interactions between the variables that we are testing. For *A. repens*, the temporal trend means that we have very low sampling density across the manifold (i.e. we have many samples along the time trend, but we do not have a lot of samples for any particular location along the time trend). For *S. scoparium*, the systems both collapse to a population size near zero, leaving minimal meaningful dynamics for the algorithm to test. Ultimately, our diagnostics suggest that these data should not be analyzed using CCM.

E026: Competition plots on a soil gradient In order to properly analyze this system, we first need to determine which plots can be included together as part of a single manifold. To do this, we again plot the dynamics of *A. scabra* in two lagged dimensions (Figure A7).

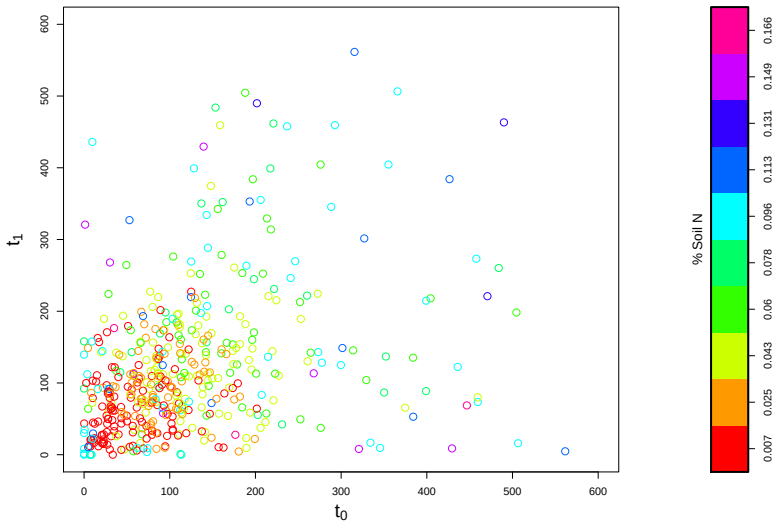

Figure A7: Lagged population dynamics *Agrostis scabra* in two dimensions.

This reveals substantially different dynamics among plots depending on total soil nitrogen, suggesting that we should analyze each soil mixture separately. Furthermore, because plots differ greatly in their above-ground biomass, we do not standardize the time series, as this could distort the manifold.

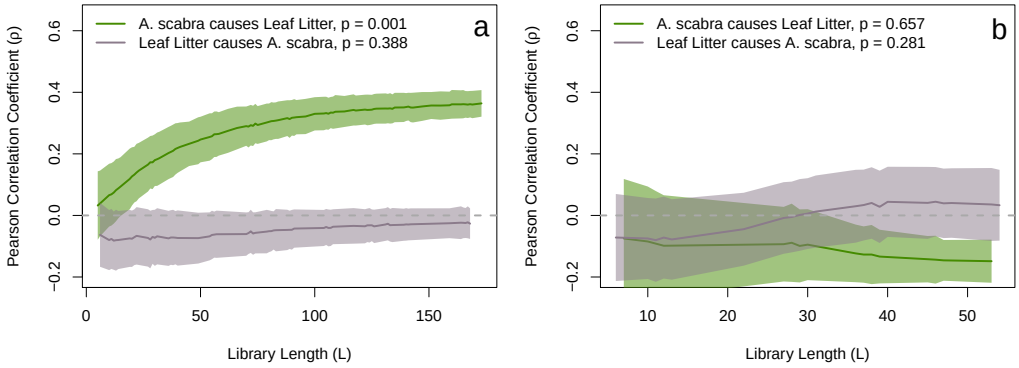

Figure A8: Test of causal forcing between *Agrostis scabra* above-ground biomass and Leaf Litter biomass dynamics in (a) fertile soil, and (b) fertile soil with added NH_4NO_3.

Based on the same diagnostic tests as we used above, we find two nitrogen treatments with sufficiently similar dynamics among plots for us to apply mutlispatial CCM. These tests suggest that in fertile soil, litter dynamics are forced by *A. scabra* dynamics (Figure A8a), but not the other way around. When additional fertilizer is added to plots, the causal direction appears to reverse, with *A. scabra* becoming forced by litter (Figure A8b). However, this second signal is not significant, likely because there are fewer plots with added fertilizer (note the shorter L).

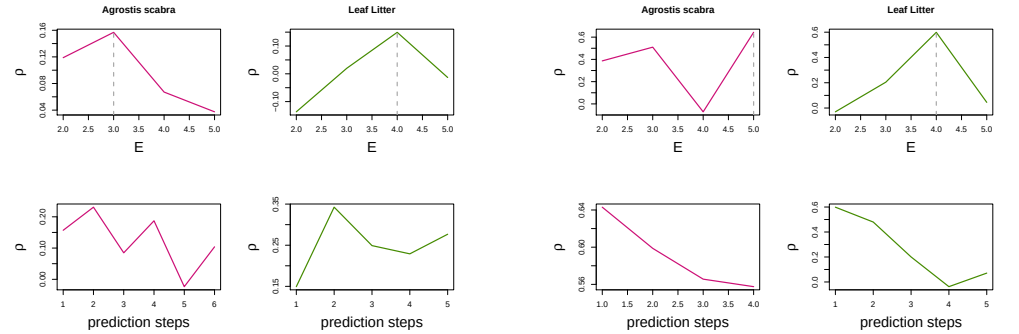

(a) fertile soil (b) fertile soil + fertilizer

Figure A9: Diagnostic plots for Figure A8.

Note that in the diagnostic plot for leaf litter in fertile soil (Figure A9a),

ρ follows a somewhat ambiguous trend as a function of prediction distance. Similarly, it is not clear that we have found the "best" embedding dimension for *A. scabra* in fertile soils with added fertilizer (Figure A9b), as the best fit occurs at the highest embedding dimension that we can test. The predictive power is also somewhat low across all diagnostic tests. While these are not ideal results for the diagnostic tests, they are not as egregious as the results for E001, and the resulting CCM tests show sensible patterns.

E054: Plant biomass in old fields This test shows some of the most easily interpretable causal results that we present here. They also show the importance of choosing a sensible embedding dimension. Though it is usually prudent to choose E that maximizes predictive ability, for both *A. repens* and *S. scoparium* we chose somewhat smaller embedding dimensions with predictive powers that are comparable to, but slightly smaller than, the "best fitting" E (Figure A10, using $E = 2$ rather than $E = 7$ and $E = 4$ rather than $E = 10$ respectively). We do this both to prevent over-fitting the model, and to retain a longer time series, as increasing E necessarily reduces the maximum library length that we can test.

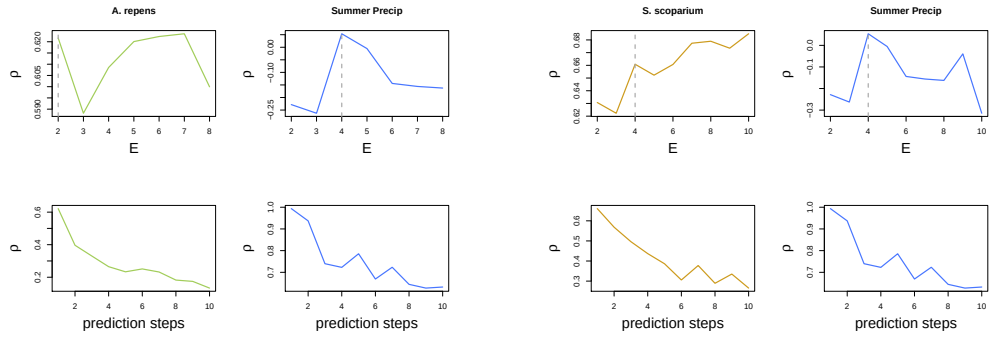

(a) *A. repens* (b) *S. scoparium*

Figure A10: Diagnostic plots for Figure A11.

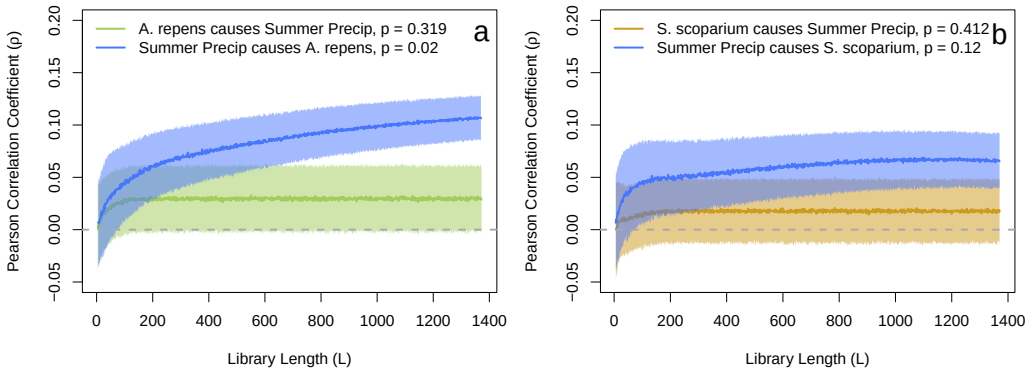

Figure A11: Test of causal forcing between total summer precipitation and population dynamics for (a) *Agropyron repens* and (b) *Schizachyrium scoparium*.

Sensibly, we find that plant population dynamics do not influence precipitation dynamics for either species. For *A. repens*, the more drought-sensitive of the two species, we find significant causal forcing by annual summer precipitation, as expected (Figure A11a). Note that this is the test presented in Figure 1c in the main text. For the more drought-tolerant *S. scoparium*, we find a trend that suggests forcing by summer precipitation, but it is not significant (Figure A11b). This is consistent with a weaker forcing effect, which would take somewhat more data to detect as significant.

E120: "Big Biodiversity" experiment Our results for insect abundance and richness partially match those for existing studies. For richness, we find a clear trend across all diversity treatments showing that above-ground plant biomass influences insect abundance dynamics, but not the other way around (Figure A12a). Insect richness and plant biomass do not appear to be significantly causally related, though the ρ vs. L relationship is somewhat strange (Figure A12b).

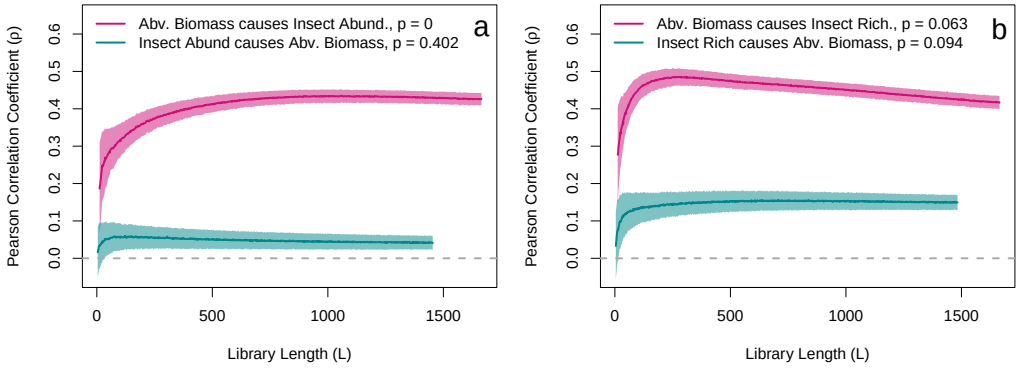

Figure A12: Test of causal forcing between total above-ground biomass and (a) insect abundance and (b) insect species richness.

Diagnostics are somewhat less useful in this case, because we don't have sufficiently long time series within each plot to show how predictive power varies with prediction interval for insect abundance or richness (Figure A13). Since we only have five sequential observations and an embedding dimension of 2, we can only predict three time steps into the future. While diagnostics for above-ground biomass seem okay, we cannot really determine whether the insect dynamics are appropriate for CCM analysis.

The CCM plot itself (Figure A12b) offers some possible information. Because ρ increases, but then decreases, with library length, our analysis shows that increasing the amount of information we have about the system (by adding more plots) decreases our ability to predict its dynamics. This suggests that the plots that we have combined for the analysis are not all well-predicted by the manifold we have constructed. However, separating plots by diversity treatment does not yield meaningful results either (analyses not shown). This suggests that insect richness dynamics differ among plots for some other reason. In the case of insect abundance, however, the CCM plot shows a more or less monotonic increase with library length, suggesting that the plots can be well-described by our estimated manifold. Thus, our results for insect abundance are probably believable, whereas our results for richness are likely confounded.

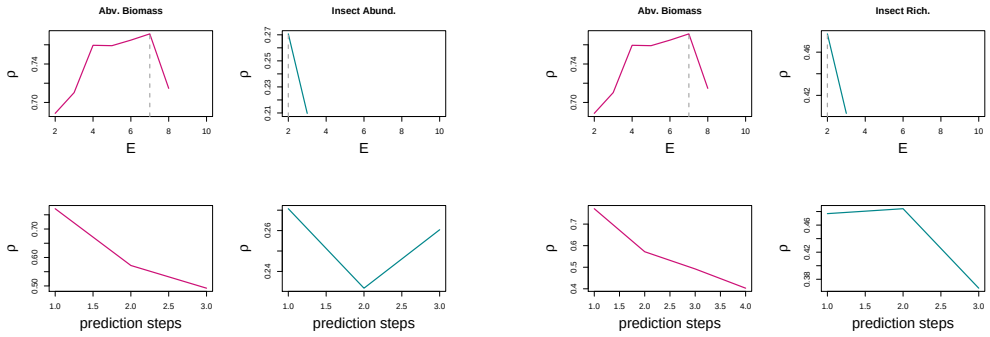

(a) Insect abundance

(b) Insect richness

Figure A13: Diagnostic plots for Figure A12.

Next, we consider causal relationships between soil nitrate dynamics and above-ground biomass (Figure A14a). There is a clear signal that above-ground biomass forces nitrate dynamics, matching our expectations. However, results for the effects of soil nitrate on above-ground biomass are somewhat more ambiguous. The p-value suggests that soil nitrate does not force biomass dynamics, but there is again an increase in ρ, followed by a decrease. Since the nitrate time series is relatively short (at most 7 sequential observations per plot) while the best embedding dimension is rather large ($E = 5$, or potentially higher), we cannot glean much information from the relationship between predictive power and number of prediction steps. Separating plots by diversity treatment does not provide a clearer pattern either (analysis not shown). Possibly, species differ sufficiently in their effect on soil nitrate that plots cannot be combined using our method.

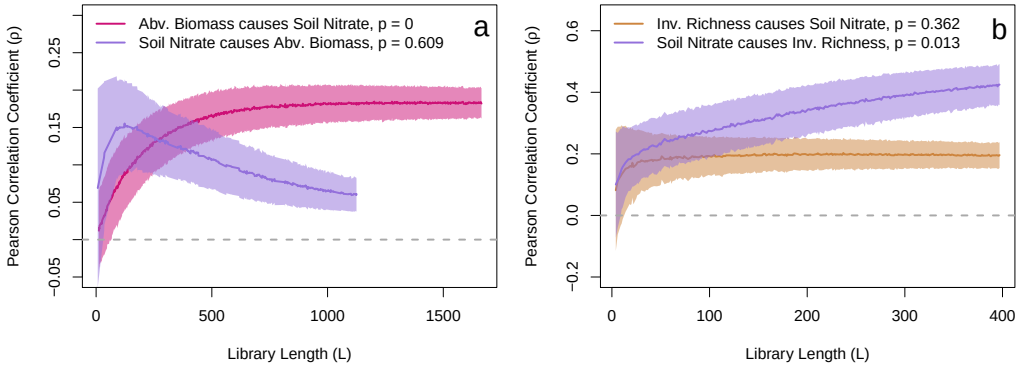

Figure A14: Test of causal forcing between soil nitrate and (a) total biomass or (b) species richness of invading (non-planted) plant species.

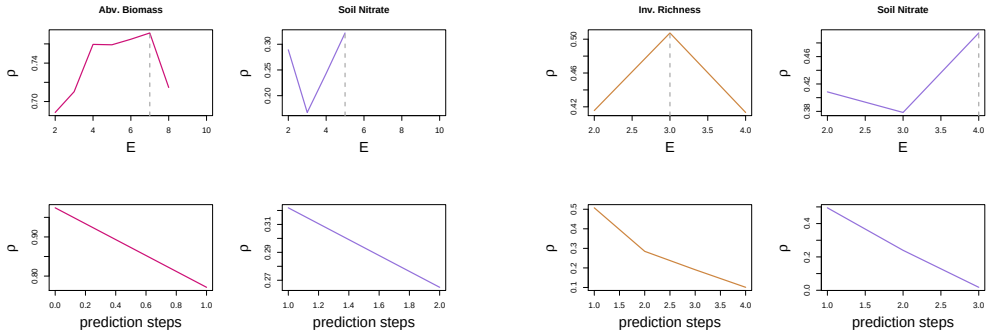

(a) *Above-ground biomass* (b) *Invader richness*

Figure A15: Diagnostic plots for Figure A14.

Results for the relationship between invader richness and soil nitrate are somewhat clearer. Here, we are able to find a subset of diversity treatments that appeared to share a single low-dimensional manifold (4 and 8 species treatments). While predictive power vs. prediction steps is still unclear in this case because of short time series (Figure A15b), the CCM analysis itself shows relatively clear signals (Figure A14b). In this case, the results suggest that invader richness does not force nitrate dynamics, whereas nitrate dynamics do influence plant invader species richness, matching our expectations from previous research. Note that this is the same as the analysis presented in Figure 1b in the main text.

4 Bi-directional causality

In order to test algorithm performance when bi-directional forcing exists, we repeated the simulations described in the manuscript, except with $\beta = 1.25$ (i.e. effect of process Y on process X, which was 0 for all other analyses). We include results from 160 simulations of the parameter ranges discussed in the main manuscript. Results show similar performance as reported in the uni-directional case for both directions of forcing (Figure A16).

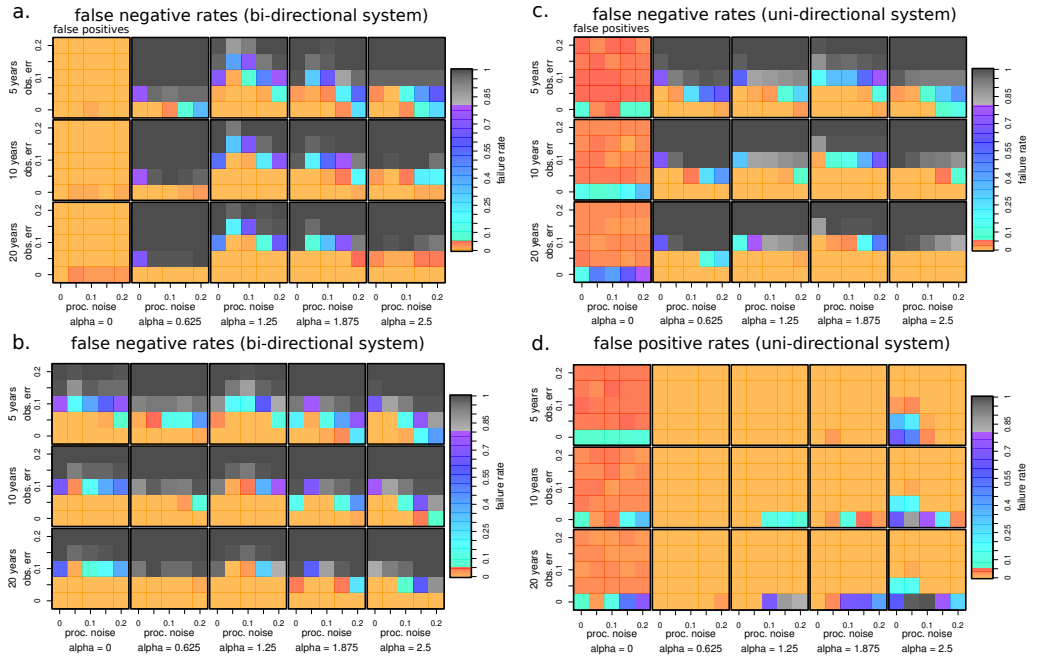

Figure A16: Multispatial CCM performance for bi-directional causal simulations for forcing of Y by X (a), and forcing of X by Y (b). We also include performance from the uni-directional test for forcing of Y by X (c), and forcing of X by Y (d) in the main text (where Y does not influece X) for comparison (Figure 2 in the main text). See Figure 2 caption in the main text for details.

References

[1] J. E. Fargione and D. Tilman. Diversity decreases invasion via both sampling and complementarity effects. *Ecology Letters*, 8(6), 2005. Times Cited: 105.

[2] R. S. Inouye, N. J. Huntly, D. Tilman, J. R. Tester, M. Stillwell, and K. C. Zinnel. OLD-FIELD SUCCESSION ON a MINNESOTA SAND PLAIN. *Ecology*, 68(1):12–26, 1987. Times Cited: 206.

[3] F Isbell, D Tilman, S Polasky, S Binder, and P Hawthorne. Low biodiversity state persists two decades after cessation of nutrient enrichment. *Ecology Letteres*, 2013.

[4] J.M.H. Knops, D. Tilman, N.M. Haddad, S. Naeem, C.E. Mitchell, J. Haarstad, M.E. Ritchie, K.M. Howe, P.B. Reich, E. Siemann, and J. Groth. Effects of plant species richness on invasion dynamics, disease outbreaks, insect abundances and diversity. *Ecology Letters*, 2(5):286–293, September 1999.

[5] Evan Siemann, David Tilman, John Haarstad, and Mark Ritchie. Experimental tests of the dependence of arthropod diversity on plant diversity. *The American Naturalist*, 152(5):738–750, 1998.

[6] D. Tilman. RESOURCE COMPETITION BETWEEN PLANKTONIC ALGAE - EXPERIMENTAL AND THEORETICAL APPROACH. *Ecology*, 58(2):338–348, 1977. Times Cited: 538.

[7] D Tilman. *Resource Competition and Community Structure*. Princeton University Press, Princeton, NJ, 1982.

[8] D Tilman. Secondary succession and the pattern of plant dominance along experimental nitrogen gradients. *Ecological Monographs*, 57(3):189–214, 1987. Cited References Count:40|ECOLOGICAL SOC AMER|2010 MASSACHUSETTS AVE, NW, STE 400, WASHINGTON, DC 20036.

[9] D Tilman. Competition and biodiversity in spatially structured habitats. *Ecology*, 75(1):2–16, 1994. Cited References Count:77|ECOLOGICAL SOC AMER|2010 MASSACHUSETTS AVE, NW, STE 400, WASHINGTON, DC 20036.

[10] D Tilman. Niche tradeoffs, neutrality, and community structure: A stochastic theory of resource competition, invasion, and community assembly. *Proceedings of the National Academy of Sciences of the United States of America*, 101(30):10854–10861, 2004. Cited References Count:51|NATL ACAD SCIENCES|2101 CONSTITUTION AVE NW, WASHINGTON, DC 20418 USA.

[11] D. Tilman and John A. Downing. Biodiversity and stability in grasslands. *Nature*, 367(6461):363–365, 1994.

[12] D. Tilman, J. Knops, D. Wedin, P. Reich, M. Ritchie, and E. Siemann. The influence of functional diversity and composition on ecosystem processes. *Science*, 277(5330):1300–1302, 1997. Tilman, D Knops, J Wedin, D Reich, P Ritchie, M Siemann, E.

[13] D. Tilman, C. L. Lehman, and K. T. Thomson. Plant diversity and ecosystem productivity: Theoretical considerations. *Proceedings of the National Academy of Sciences of the United States of America*, 94(5):1857–1861, 1997. Times Cited: 467.

[14] D. Tilman and D. Wedin. OSCILLATIONS AND CHAOS IN THE DYNAMICS OF a PERENNIAL GRASS. *Nature*, 353(6345):653–655, 1991. Times Cited: 105.

[15] David Tilman and David Wedin. Plant traits and resource reduction for five grasses growing on a nitrogen gradient. *Ecology*, 72(2):685, April 1991.

[16] David Wedin and David Tilman. Competition among grasses along a nitrogen gradient: Initial conditions and mechanisms of competition. *Ecological Monographs*, 63(2):199, May 1993.