Chapter 4: Internal variability

Jeffrey Sachs, James Rising,
Tim Foreman, John Simmons, Manuel Brahm
The Earth Institute
Columbia University

October 1, 2015

Prepared for the Global Coffee Forum
Milan, Italy

Part of
The impacts of climate change on coffee: trouble brewing
http://eicoffee.net/

While long-term climate change is going to alter the landscape of coffee, many climate impacts are already occurring and represent shifts on shorter time-scales. While the most commonly known cycle in the coffee industry is the biennial production cycle, the climate has internal cycles that affect coffee production. These are produced by climate decadal variability, and are some of the most important climate dynamics for the coffee industry.

We find that the most important of these cycles is the El Niño/La Niña cycle, known as ENSO. The ENSO cycle produces vast shifts in weather climate patterns over much of the tropics. While each El Niño event is unique, they produce somewhat predictable patterns of increases and decreases in precipitation (see 1). ENSO is of particular concern today, as we enter what may be the largest El Niño event in a generation. During the last large El Niño in 1997-98, the tropics were hit by both droughts and floods, as shown in figure 2.

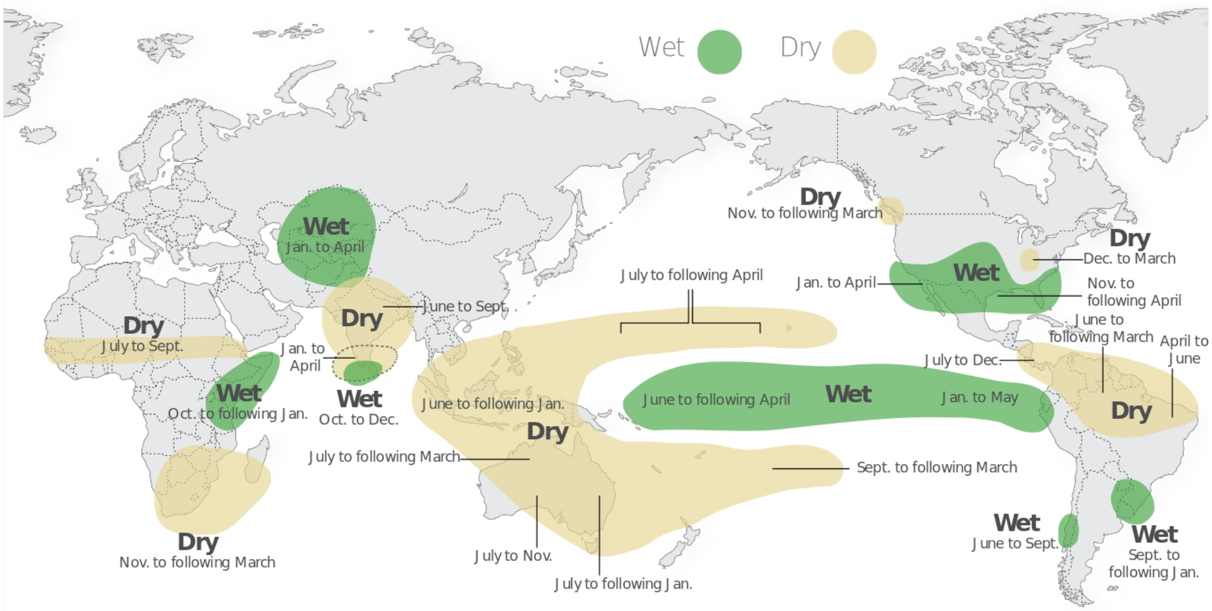

Figure 1: The typical precipitation impacts of an El Niño event, from **?** and **?** as modified by **?**.

It also coincided with infectious outbreaks in Africa (Epstein, 1999), megafires in Indonesia (Page et al., 2002), and crop failures across the tropics (Hsiang and Meng, 2015). El Niño and La Niña events can often be predicted before their impacts are felt, and knowing what to expect can make a big difference in the outcomes. In fact, it can be argued there should be fewer disasters during El Niño and La Niña because the increased predicatibilty of the seasonal climate during these events can be used to increase our preparedness for adverse climate impacts (**?**).

The ENSO cycle is not the only multi-year climate cycle. We study the effects of five climate cycles and their effects on coffee production. The results are important beyond the direct quality of our predictions, for a pragmatic reason and a theoretical reason. Pragmatically, predictions that are based on globally available climate metrics can help farmers living in areas that do not have readily available or reliable weather forecasts. For these farmers, the information that a season will be wetter or drier than normal could be more important than the particular pattern of rainfall events. Determining the predicted impacts of climate indices, using our model, relies only on globally available metrics.

This dependence on global metrics also has a theoretical benefit. Complicated feedback systems can exist

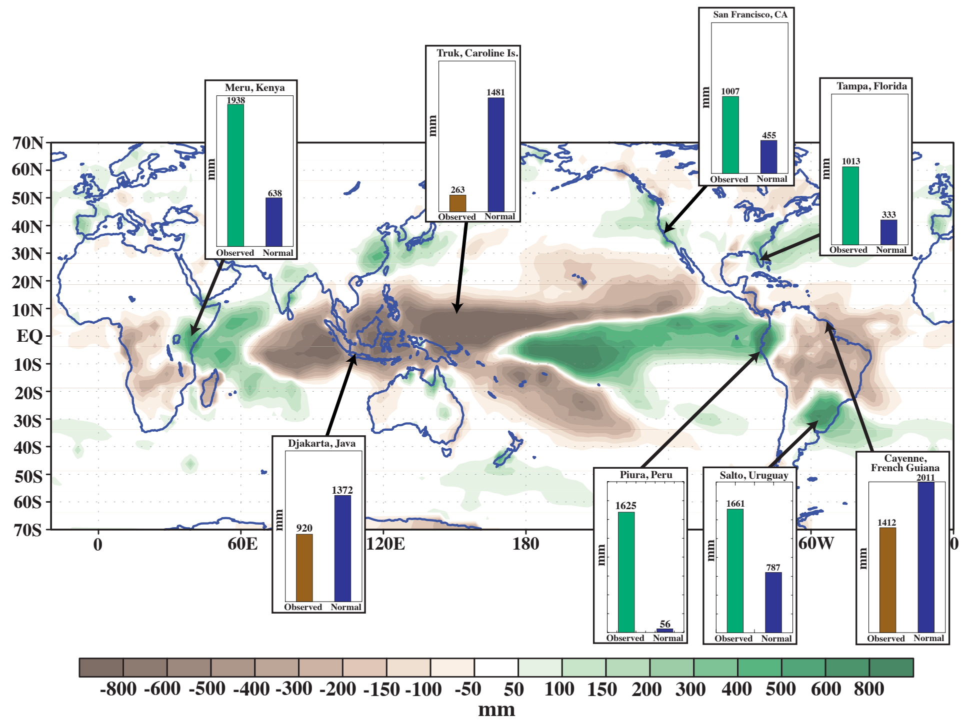

Figure 2: Rainfall anomalies from November 1997 to April 1998. Reproduced from Bell et al. (1999).

between agriculture in an area and the weather that area experiences; for example, deforestation can lead to decreases in precipitation. These feedbacks make it difficult to distinguish the effect of climate from the effect of farmers responding to, anticipating, or influencing climate. This problem is eliminated for models that use distant indices, as is the case with all of the multi-year climate cycles.

1 Impacts of El Niño

The ENSO cycle is both a concern and an opportunity for coffee production. Many coffee-producing countries are significantly impacted by these events, with changes in temperature, precipitation, and blossoming conditions. Furthermore, since El Niño and La Niña are global, they can produce large impacts on the coffee market. Like all climate events, El Niño affects both coffee plants and their associated farming communities in a way that is difficult to disentangle.

The scientific understanding of the ENSO cycle continues to evolve. Here we do a quantitative analysis of ENSO on recorded coffee prices and yields, although this is only a part of the picture. El Niño events can cause severe storms that increase erosion, producing a long-term effect that is only reflected tangentially in our data. El Niños can also affect a coffee farming operation by affecting the welfare of its farmers.

Our first analysis uses a consensus categorization of years into El Niño, La Niña, and ENSO-neutral years, as an indicator for studying impacts. These years and the shape of the NINO 3.4 indicator that corresponds to them is shown in figure 3.

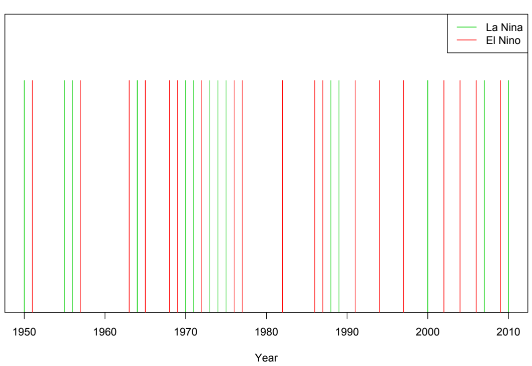

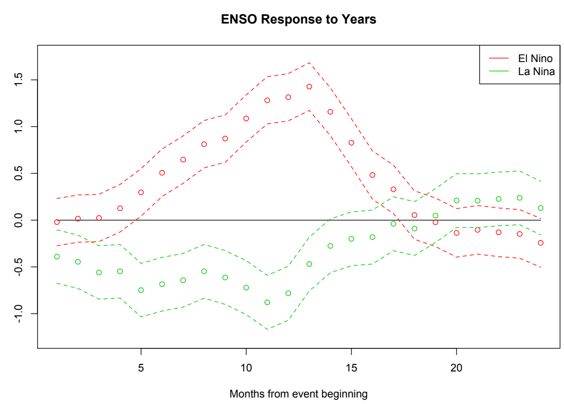

Figure 3: **Left:** The years categorized as La Niña years and El Niño years. **Right:** The estimated 24-month impulse response of the NINO 3.4 indicator to each of the three ENSO year types.

We see these impacts in the prices of Arabica and Robusta beans in El Niño years, relative to ENSO-neutral years, as shown in figure 4. In expectation, from the beginning of an El Niño year, prices climb for about 15 months, before beginning to decline. At their peak, prices are over 30% higher than they are predicted to have been in the absence of the El Niño event. We do not see a similar effect for La Niña events. These effects are similar in duration and form to those found by Ubilava (2012).

Computing ENSO impacts

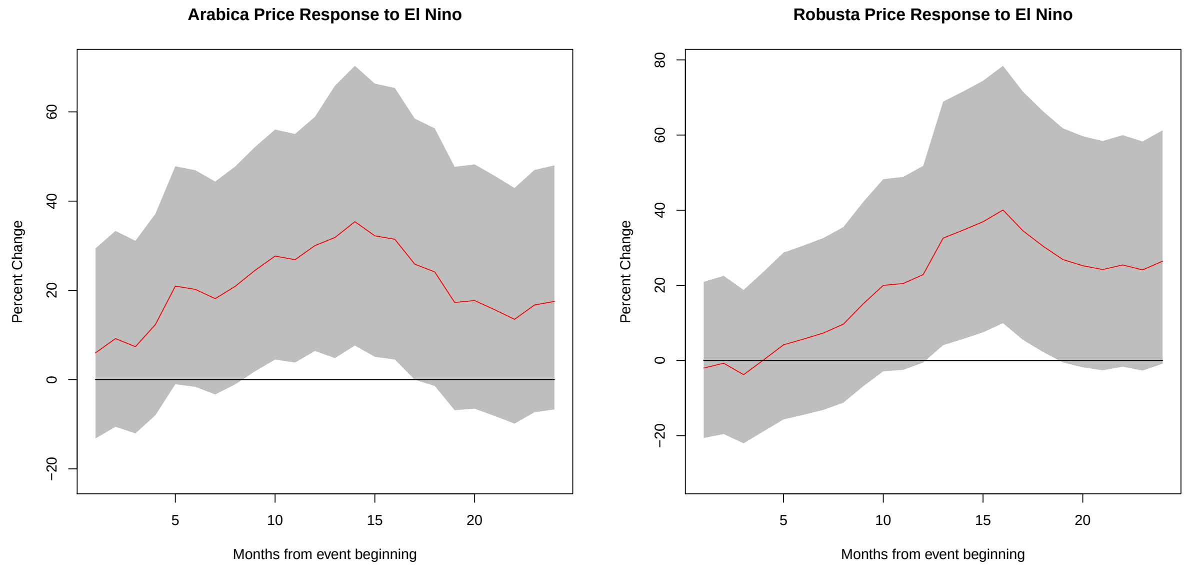

Figure 4: The graphs above show how international coffee prices respond to an El Niño event. Both Arabica and Robusta prices show increase of 20-40% over the course of the event, with potentially long-lasting impacts.

We estimate the impacts of El Niño and La Niña by estimating an "impulse response", which accounts for the multiple overlapping effects of different ENSO years and the monthly climatology of the NINO 3.4 signal.

$$y_t = \alpha + \sum_{Y=Year(t)-N/12+1}^{Year(t)} \sum_{M=1}^{N} \beta^{Class(Y)}_{12(Year(t)-Y)+M} + \gamma \sum_{s=1}^{24} \frac{y_{t-s}}{24} + \mu_{Month(t)}$$

$Year(t)$ is the year for time t and $Month(t)$ is the month for time t; $Class(Y)$ is the class of ENSO event that happened in year Y (El Niño and La Niña). N is the number of months to include in the impulse responses.

Here, the β^k_m variables describe impulse responses of length N for each class of ENSO event.

Little research has been done on the effects of the El Niño/La Niña cycle on coffee yields. Villegas et al. (2012) find that in Colombia the location of the Inter-Tropical Convergence Zone is a more important factor affecting yields, but global estimates of these effects do not appear to be available. We consider the impacts of El Niño and La Niña years globally and for each country.

An impact analogous to the price change is evident in the yields across countries. In El Niño years, yields decrease on average by 100 Hg/Ha, against an average of 6800 Hg/Ha in recent years, after accounting for long-term trajectories in yields. While this is only a drop of 1.5%, the average hides larger effects in specific regions and variability between El Niño years. No effect is seen globally in La Niña years.

Looking at individual regions, using country-level data, only a few countries appear to have large impacts from El Niño, after accounting for each country's long-term evolution in yield and production. The French

Polynesia, Gabon, Polynesia, and Thailand all show significant impacts in yields, although the directions of the impacts differ. Mauritius, Papua New Guinea, and Sri Lanka have significant impacts in total production, with large decreases for Papua New Guinea and Sri Lanka. Because of the lack of information on coffee planting areas, yields are calculated with respect to harvested areas. As a result, the countries that show impacts on production but not on yield probably reflect El Niño impacts that are hidden by selective harvesting decisions.

Regions with observed El Niño / La Niña impacts

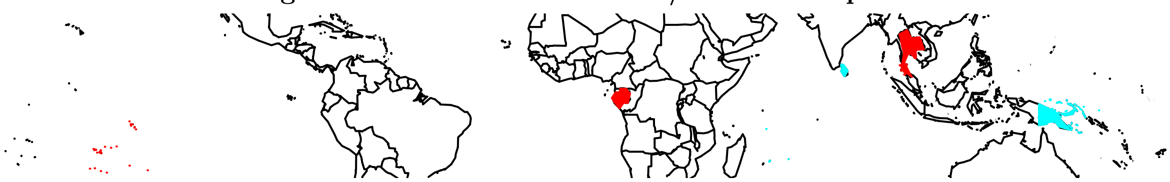

Yields	El Niño	El Niño + 1	La Niña	La Niña + 1
French Polynesia	-650 Hg/Ha	N.S.	N.S.	N.S.
Gabon	+940 Hg/Ha	N.S.	N.S.	N.S.
Polynesia	-650 Hg/Ha	N.S.	N.S.	-600 Hg/Ha
Thailand	N.S.	N.S.	N.S.	+1400 Hg/Ha
Productions	**El Niño**	**El Niño + 1**	**La Niña**	**La Niña + 1**
Mauritius	+17 MT	+13 MT	+18 MT	+17 MT
Papua New Guinea	-7,600 MT	N.S.	N.S.	N.S.
Sri Lanka	N.S.	N.S.	-2,000 MT	N.S.

Figure 5: Regions where country-level yield and production are teleconnected with El Niño and La Niña events. The location of the regions is shown at top, with yields affected by ENSO in blue and production affected by ENSO in red. In the table, values are the predicted change in yields or production in El Niño years, the year after an El Niño year, and the same for La Niña years. Entries with "N.S." show no statistically significant change at a 10% level.

2 Coherent movements

The relatively weak statistical relationship found between El Niño and country-specific yields is not uncommon among agricultural crops, but it drove an interest in our group into dissecting more clearly the relationship between global climate signals and country production. We collected five oceanic signals to explore this further, as shown in figure 6.[1]

A principal component analysis identifies regions of coherent marginal changes, across multiple time-series. This technique can be used to better understand patterns in large datasets, like the one describing country coffee production.

For each year, monthly values of the climate signals (delayed 6 months, to capture their impacts on coffee flowering) and country yields (detrended with locfit and normalized) are included. The first principal component represents the largest coherent movement of change, followed by the second component, and so on. Between the first three components, over 50% of the variation in yields can be described. The share of each of these components by year is shown in Appendix .2. Each of the components and what it suggests about the relationship between climate and yields is described below.

[1] NINO 3.4, NAO, SOI, PDO from NOAA Climate Prediction Center (CPC) (2015), unsmoothed AMO from Enfield et al. (2001).

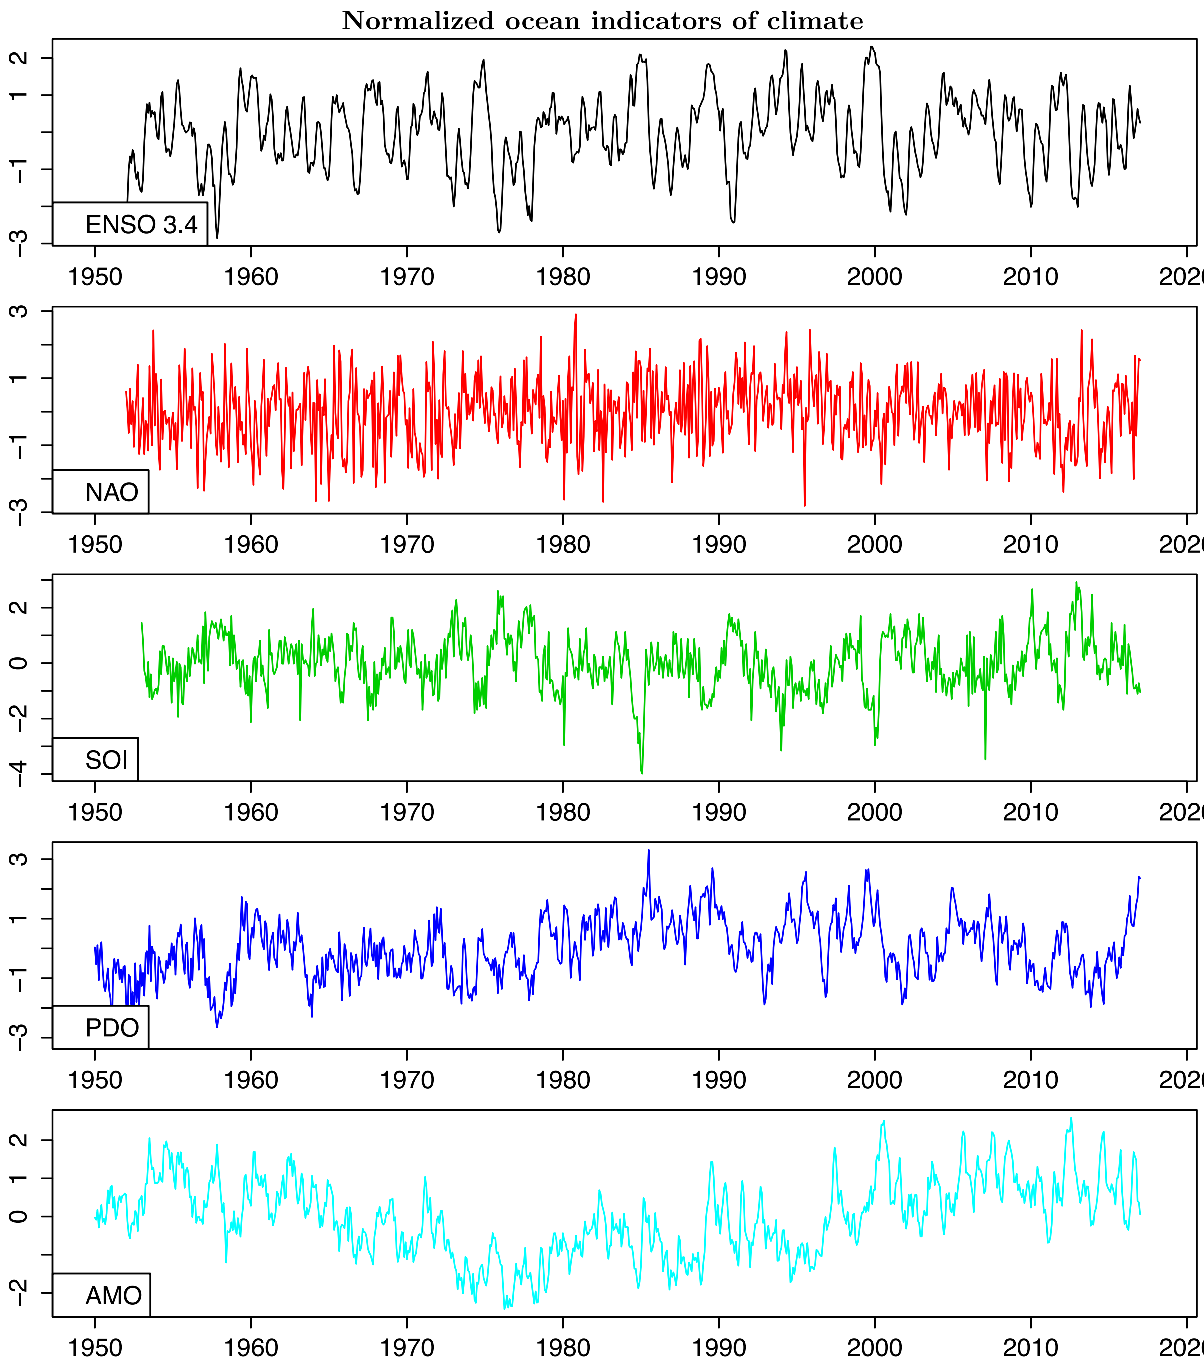

Figure 6: Normalized indicators used to study global and regional climate, sampled monthly. Each of these shows wide variability, but different periodicities. The interactions between these different signals can explain impacts in ways that individual signals cannot.

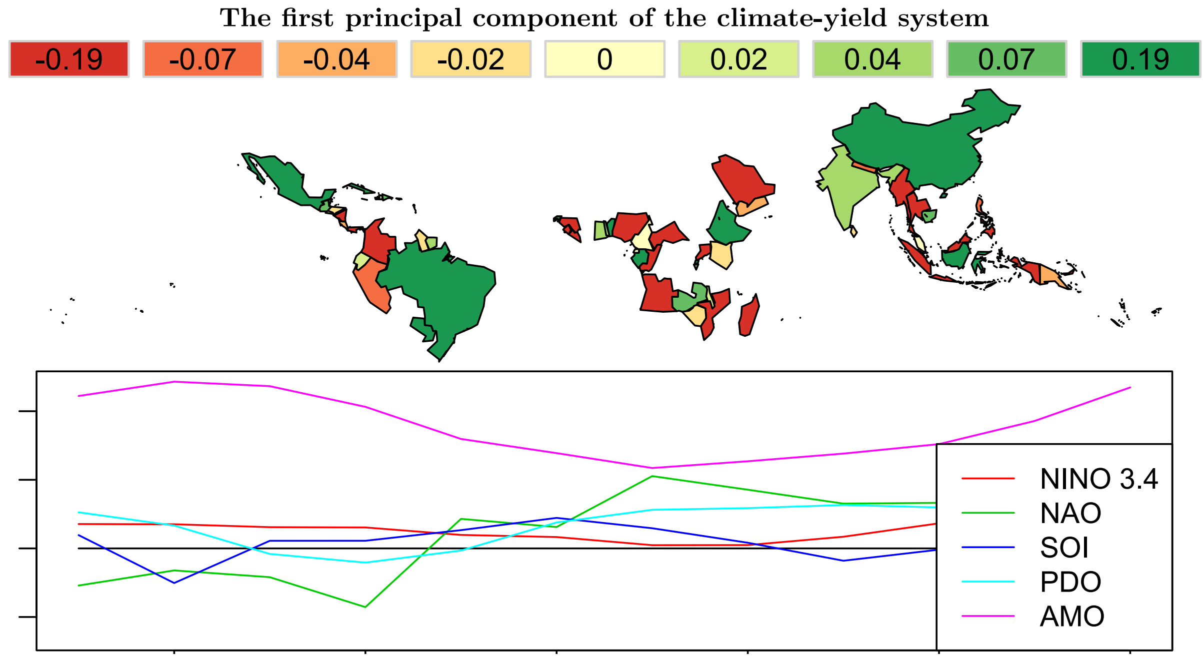

Figure 7: Spatial and temporal representation of the first principal component of the climate-yield system. Colors in the map represent increases (green) and decreases (red), and the plot below shows the climate signals across the year (delayed 6 months, so month 1 is July and month 12 is June). Explanation in the text.

The first principal component describes how yields have shifted on average over the past 50 years. Brazil, Mexico, and China have seen some the largest increases in yield, while Thailand, Myanmar and many countries in Africa have experienced the largest decreases. Most climate signals have not shown any trend, except for the Atlantic multidecadal oscillation (AMO) which is currently much higher than it was in the 1960s. As a result, all of the climate signals in the lower graph are near zero, except for AMO.

The second and third principal components are dominated by ENSO (the El Niño/La Niña cycle), represented by the NINO 3.4 index and the Southern oscillation index, which is known to be strongly correlated with ENSO but with an opposite sign. PC 2 is represented in the data when NINO 3.4 is high (El Niño) and the Pacific decadal oscillation (PDO) is also high, and its effects are reversed when these signals are both opposite in the direction of their anomalies. The largest effect of this combination, as shown in the map, is that Brazil, Paraguay, and Papua New Guinea have decreases in yields while India sees increases. This suggests that yields in these regions will often move in opposite directions, during many El Niño and La Niña years.

Observations with low values of PC 2 occur before 1975 and after 2000, while those with high values of PC 2 occur mostly in the 1980s and early 1990s. This may be driven by the slow oscillation of PDO. Since only one such cycle has occurred, it is difficult to distinguish the effects of the climate signals from socioeconomic effects, although most of this was be removed by the flexible trend used in the preprocessing step.

The third principal component also occurs when ENSO is in its El Niño state, and AMO is high or increasing. In this case, India, Peru, and southern areas in Africa show decreases, while other areas are not heavily affected. Both PC 2 and PC 3 can equally be understood in their La Niña form (and

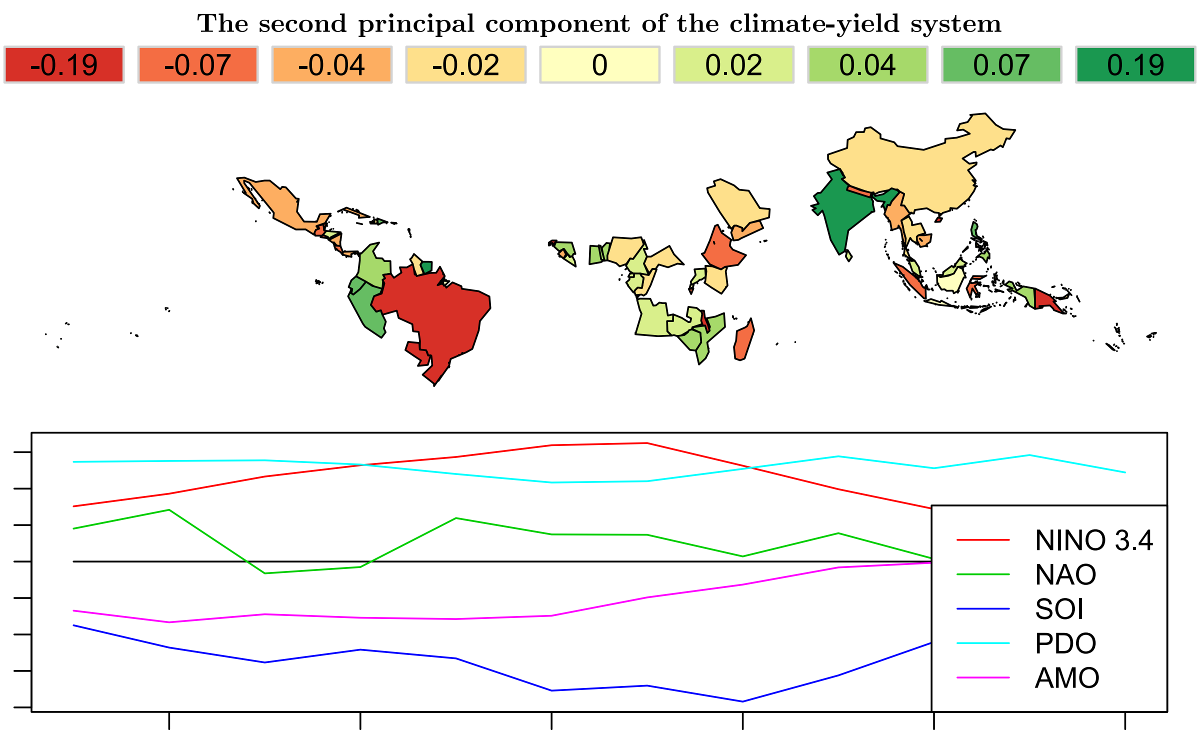

Figure 8: Spatial and temporal representation of the second principal component of the climate-yield system. Colors in the map represent increases (green) and decreases (red), and the plot below shows the climate signals across the year (delayed 6 months). Explanation in the text.

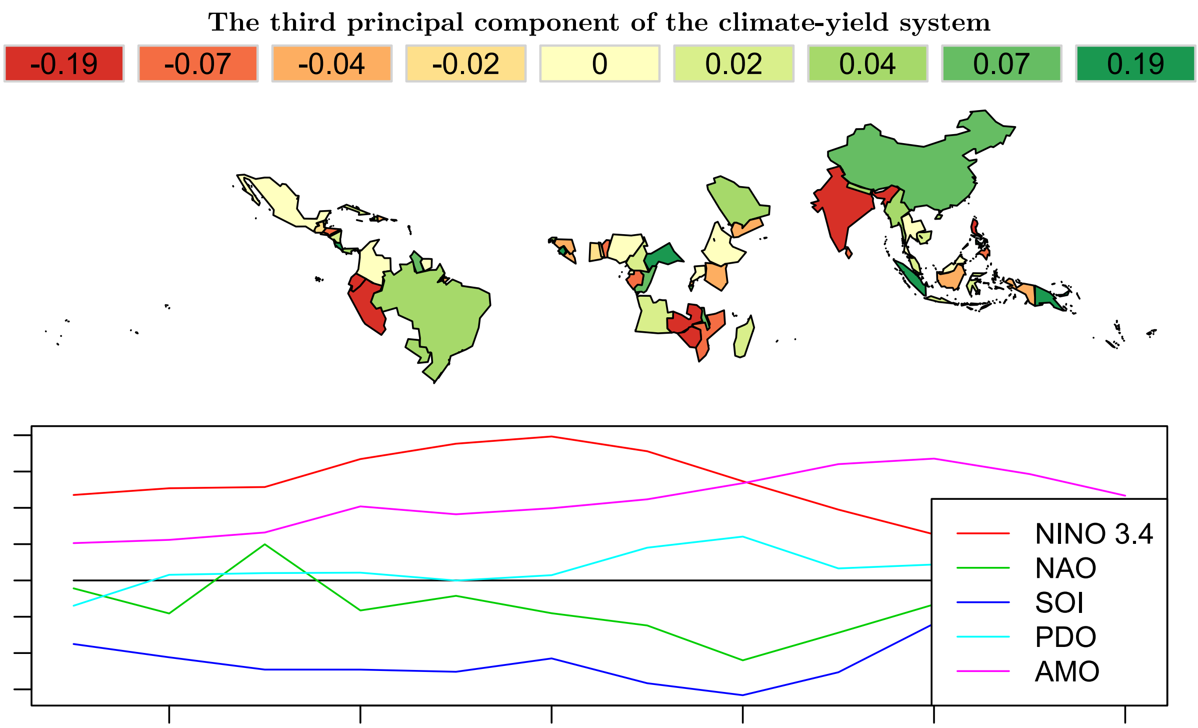

Figure 9: Spatial and temporal representation of the third principal component of the climate-yield system. Colors in the map represent increases (green) and decreases (red), and the plot below shows the climate signals across the year (delayed 6 months). Explanation in the text.

associated low values of PDO for PC 2 and low values of AMO for PC 3), which produce changes in yields in the opposite direction.

Between PCs 2 and 3, the effects of El Niño and La Niña appear across much of the globe. Because the impacts on most countries result from an interaction between the ENSO cycle and AMO or PDO, the results did not appear in the initial analysis.

3 Projection for 2015-2016

The El Niño predicted of 2015-16 is expected to have a similar magnitude to the event in 1997-98. However, there is considerable unresolved uncertainty in our analysis as to the most likely outcomes of this event.

The El Niño of 1997-98 produced catastrophic impacts in many areas, but its effect on coffee was fairly minor, coinciding with in a 9.6% drop in production. A large fraction of this global effect was due to Brazil, which had a 32% drop in yields, largely as a consequence of its biennial cycle. Excluding Brazil, the rest of global production only decreased 1.3%. The regional picture is more nuanced, with large decreases also in Oceania. The top part of figure 10 shows these results

From the analysis above, the most damaging El Niños coincide with consistently high values in PDO, such as we see today. By decomposing the existing constellation of climate signals into the three coherent groupings shown above, we project the estimates shown in figure 10. As described above, there may be large decreases in yields in Brazil and Central America. Our projection also identifies losses in India, and Southern Africa. This is at odds with the recorded values from 1997-98, which saw the most widespread losses across Indonesia and Papua New Guinea.

Inputs to the El Niño projection

Our projection is based on the most recent consensus projection of the NINO 3.4 index of ENSO, from International Research Institute for Climate and Society (2015). We try to apply reasonable values to the other indices, using the negative of NINO 3.4 for SOI, given its -0.6 correlation with NINO 3.4; a zero value for NAO, given its rapid shifts; and constant extrapolations for PDO and AMO at their most recent value, given the slow shifts in these signals.

Projecting these signals onto the principal component axes gives loadings of 1.3, 3.3, and 1.9, for the three components respectively.

To account for any spurious effect of our decision-making process, we estimate the values for 2014-15 as well, and report the difference.

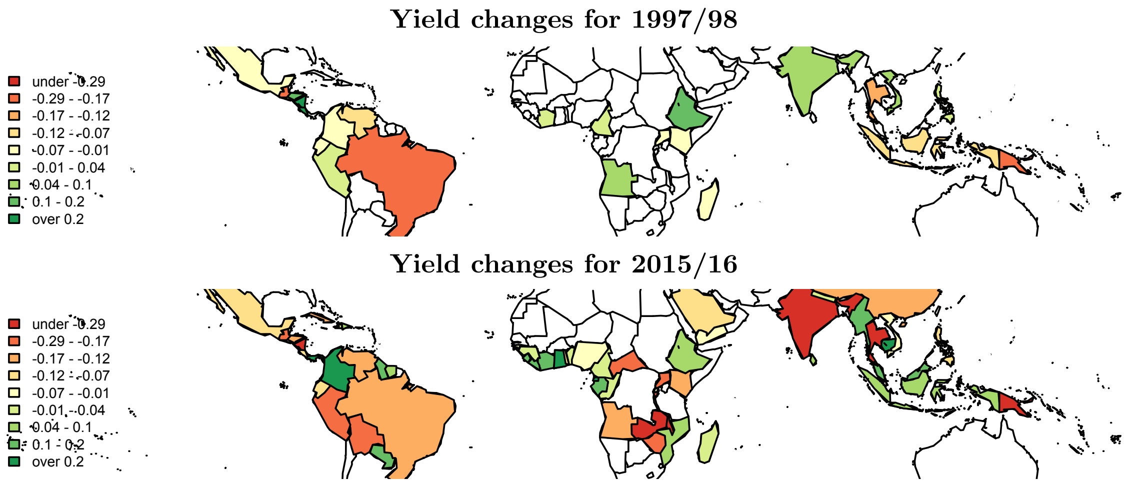

Figure 10: **Above:** Yield changes coinciding with the 1997/98 El Niño, relative to 1996/97, from International Coffee Council (1998). **Below:** Yield changes predicted for the 2015/16 El Niño, relative to 2014/15.

4 Pest growth and control

An associated project for this report explored the effects of temperature in agricultural pest outbreaks. The project focuses coffee rust fungus, La Roya, in Guatemala, an area of high coffee production and recent extreme rust outbreaks (Georgiou et al., 2014). It examined how changes in monthly temperature, the associated 'incubation period' for the fungus (*Hemileia vastatrix*), and the inclusion of a vigilant farmer can affect the outbreak size distribution over time.

Hotter conditions have supported the increase of fungus spread that is killing coffee trees in altitudes that were once free of fungus. This has even caused farmers to switch to more resistant, but lower quality strains of coffee such as Coffea Robusta. This problem impacts the export earnings of coffee-producing countries but more importantly it directly impacts the employment of hundred of thousands of coffee workers who depend on the harvesting earnings to feed their families, with little income to cover a lost season (Magrath, 2014). Additionally, as coffee is a globally traded commodity and Central America is one of the top exporters, the proliferation of coffee rust fungus also has implications for other countries around the world.

Many attempts have been made to find simple and mechanistic solutions to both understand and predict the outbreaks dynamics, but as shown by Lockwood and Lockwood (2008), it has often proven fruitless to capture the effects of weather through linear models. Though many techniques have been utilized to capture the non-linear behavior for different spatial and temporal domains, this project follows their strategy of using a spatial, agent-based model to understand the interactions of space and time.

4.1 A rust model

We make a number of simplifications to study the fungus outbreaks. First, we will only consider the fungus's interaction with the host plant, even though it has been found to utilize other plants for different stages of its growth cycle. Secondly, we assume that the only factors influencing it spread are temperature and the health of the host plant, thereby ignoring wind and rain impacts that are also known to be important (Ferreira and Boley, 1991). Similarly ignored are higher order effects from the application of fungicide, where fungicide can also impact some flora and fauna that regulate the fungus, leading to potentially unpredictable disruptions in the natural system.

The model is initialized as a two-dimensional grid of farm space, each grid cell having a certain probability of an appearance of a fungal outbreak. For each time-step, chosen to be one month after examining the reproductive cycle of coffee rust, the outbreak will begin to increase in size as a function of both its current size, the temperature, the amount of host plants available and spread to neighboring grid-cells. The temperature used in the model was obtained from surface temperature Reanalysis Data from the National Center of Atmospheric Research, spatially averaged over the area of Guatemala (without ocean cells) and temporally averaged to each month (NCAR, 2015). A time series of these average temperatures is shown in Figure 11, left. Here, one can note a relatively consistent seasonal amplitude of ˜2.5°C around a mean of ˜22°C, with a slight upwards skew.

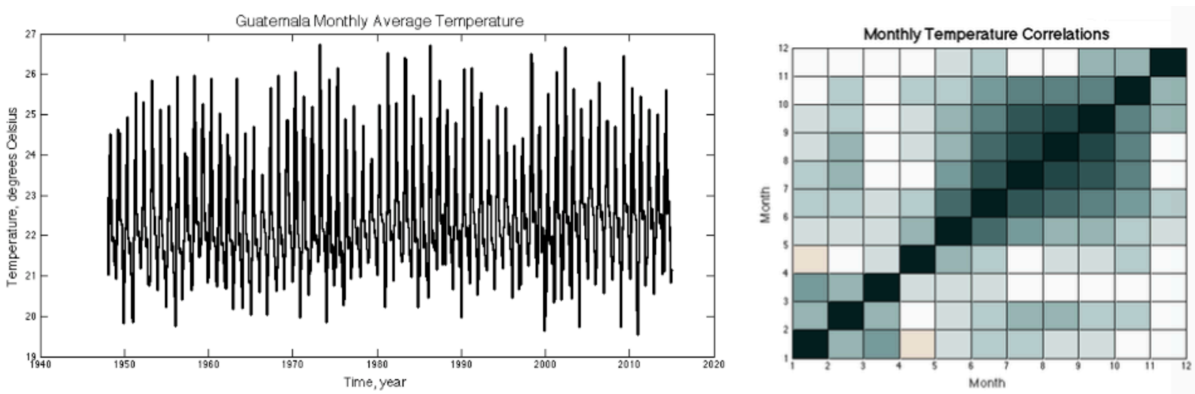

Figure 11: **Left:** Timeseries of temperatures. **Right:** Autocorrelation of monthly temperature.

The high level of correlation between months requires that the temperature selected in each monthly time step depend on the temperature in the previous month (see figure 11, right). In particular, the summer months are highly correlated, reaching correlations of 0.8 with the previous month in some cases. To account for this, the model is set up to draw a random season from the 67-year time series, employing a three month time series corresponding to an instance of summer, fall, winter or spring, depending on which is needed.

Our basic growth equation can be described by the following equation, with the basic assumption being that higher temperatures increase the growth rate of the fungus (at least at the temperatures seen in Guatemala).

$$N_{t+1} = N_t e^{r_i T_t / T_c}$$

where N is the population of a particular grid cell, r_i is the initial growth rate, T_t is the temperature at that specific time, T_c is the average temperature (over all months). Normally, the quantification of the growth rate is usually conducted with a consideration of both the daily maximum and daily minimum temperature (Magrath, 2014). However this was simplified for inclusion in our model.

When an additional population of fungus is created in subsequent time steps, it is distributed among the original and nearby grid-cells proportional to the health of the host plant in the new grid cell, the population of the source grid-cell and a multiplicative term similar to the prior growth equation. The maximum fungus population for each grid cell is 1, representing 100% infection of the host plant.

The disturbance of each grid-cell is also be subject to density-dependent pressure from predators, in this case the farmer spraying fungicide. Once a particular grid cell reaches a certain percentage of infection it is detected by the farmer. Detected, the population is decreased by a certain fraction, through the application of the fungicide. In addition, the fungus population will also decrease at a rate proportional to its current population and the relative health of the host species, independent of temperature.

$$N_{t+1} = N_t - N_t(F - 1)$$

where F is the percentage of available host plants for the fungus to grow on.

4.2 Experiments and results

To evaluate how temperature affects the spread of pests, experiments with three temperature scenarios were run. The first one uses each time step temperature (month) from the historical seasonal data for the region. The other two temperature scenarios considered global warming, one where the mean temperature was increased by 2°C and the other by 4°C. Those three temperature scenarios were run under two pest control conditions, the first one without any kind of pest control and the second one with a farmer's control by using fungicide that eliminated a fraction of the fungus when detected.

4.3 Historical temperature data without pest control

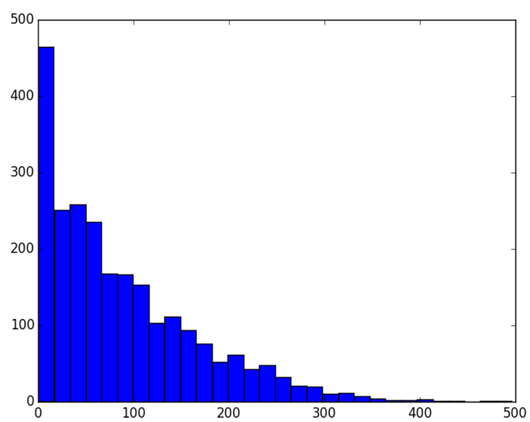

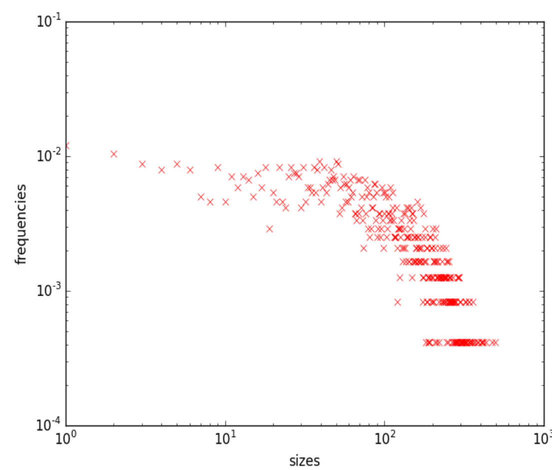

Figure 12: Historical temperature data without pest control. **Left:** Histogram of outbreak sizes. **Right:** Log-Log Plot of outbreak sizes.

We set the 'infected threshold' to about 0.3, from the simple fact that we found several instances in which leaves about 1/3 covered in coffee rust where considered 'heavily infected'. In figure 12 we have plotted the distribution of outbreak sizes, counting each heavily infected grid-cell at each time step. Please note that as we utilized a 30 by 30 grid cell, complete infection can be represented by a score of 900. Therefore, 490, the largest event, signifies that 54% of the crop is heavily infected.

The Log-Log result, though resembling a power law from 100 on, demonstrates some inconsistent behavior in smaller events. Indeed, in the histogram shelves can be observed where lower values have relatively the same probability. With no warming, the most common event is between $0 - 17$, or between 0 and 2% of the crop (though this refers to a single month, not a harvest cycle). The largest event, at 54% of the crop is actually lower than the actual 70% loss in Guatemala in 2012, though there was not an indication of how this figure was calculated. Nevertheless, the model predicts scenarios where the fungus infects over 30% (300 sites) of the field.

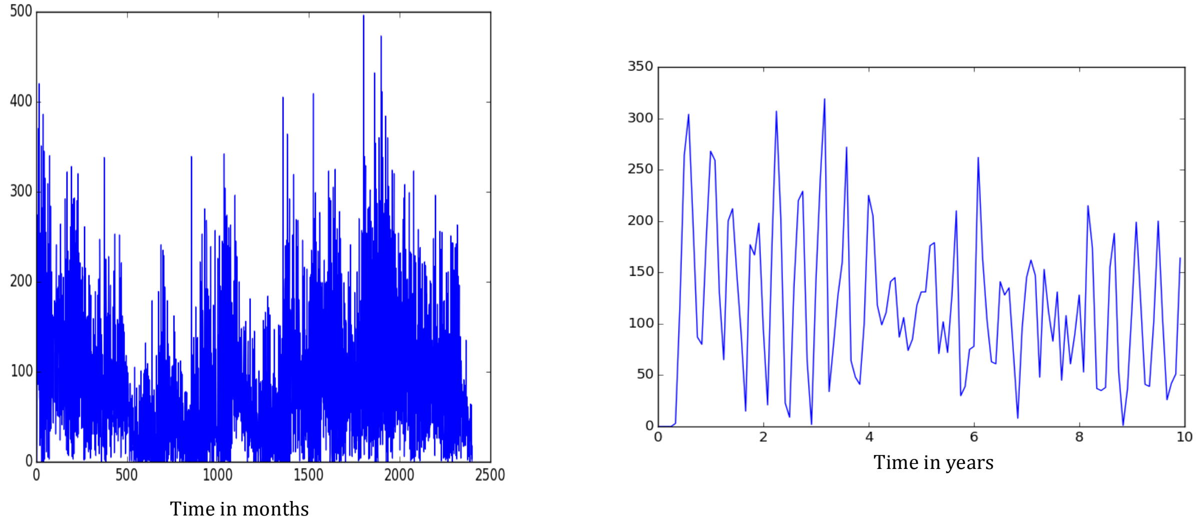

Figure 13: **Left:** Time series of outbreak sizes by month. **Right:** Time series of outbreak sizes by year.

Figure 13 shows the time series of the previous figure in months. It is clear that the system can be entrenched within certain domain (i.e., large or small events) for many years. As the behavior between 1750 and 2000, a 21-year period, consistently shows some of the highest outbreaks, while still containing intervening low events. This is possibly because the growth of fungus and plant, which are both tied to temperature, sometimes became more synchronous. However, this cannot be determined outright, and thus for future study we might want to run for more time steps, to understand more fully the nature of this large- scale periodicity.

On smaller scales, it can also be noted that many of the largest events come directly after a period of relative calm, as the host plant has had a chance to regain health and provide much more nourishment to the attacking fungus. This small scale rebound, can be seen with more detail. Though the rapid up and down movement can be shown on a scale of a few months, there always seems to be a larger periodicity on the scale of a few years; however, the randomness in the system makes it difficult to conclude anything concrete.

4.4 Historical temperature data with pest control

Next, we implemented the farmer control, where obvious pest presence would immediately be sprayed with a fungicide and reduced to a fraction of its value in the next month. This fungicide and the necessary training to use it correctly may currently be absent within the poorer farms in the area. Thus, this allows us to see how implementation might change the situation in the future. In Figure 14, the distribution has a much smaller mean and median than the without the pest control measures. We note a stronger power law relationship, though like the last iteration, it is slightly concave down. This suggests that moderate events are marginally more likely than they would have been. However, the histogram might be slightly skewed by a larger prevalence of zero events. Additionally, one can note a larger spread than the previous model run for lower probability events.

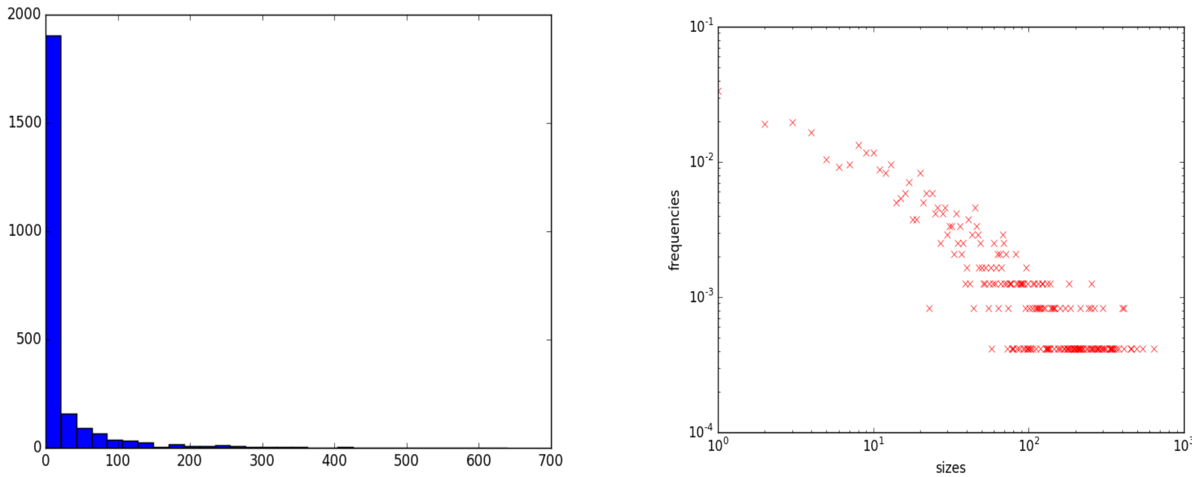

Figure 14: Historical temperature data with pest control. **Left:** Histogram of outbreak sizes. **Right:** Log-Log Plot of outbreak sizes.

Unexpectedly, the largest event, 643, signifying about 77% of the crop, is much higher than the previous iteration. This suggests that though fungicide keeps the fungus levels low for the average month, the healthy status of the host plant will make it so the correct temperature conditions or perturbation can cause a huge event, even before the farmer can react (here at a 1 month lag). This is very reminiscent of real world pest control experiences, where application can have unforeseen consequences, such as diminishing the population of a pest predator, and thus upsetting the natural structure of the system and allowing a pest to flourish later (Modern Farmer, 2014). However, catastrophic losses at a few points do not offset the considerable gains shown across the histogram.

Nevertheless, this line of thinking is further corroborated by figure 15, where decent periods of little activity are punctuated by huge events, a common feature in nonlinear spatial systems. However, when zoomed in to a period of 10 years, one can note the similarity between the control and non- control scenarios, where the lower bound in the control situation (within inter-month cycles) is replaced with 0.

4.5 2°C global warming temperature data without pest control

For the next run, we linearly increased the temperature of each month by 2 degrees, in order to represent possible regional warming over the next century. Increasing the temperatures to above normal, and thus often increasing the ability of the fungus to reproduce, causes the histogram of outbreak sizes to shift rightward. The log-log plot, while showing linear behavior for the right tail of the distribution, mimics this change. In an average month 10% of the crop is considered heavily infected, with the tail hitting about 75% of the crop as a maximum value. The time series (Figure 17) shows very few events with absolutely no fungus though the behavior in terms of both large-scale and small-scale periodicities does not seem to be drastically different.

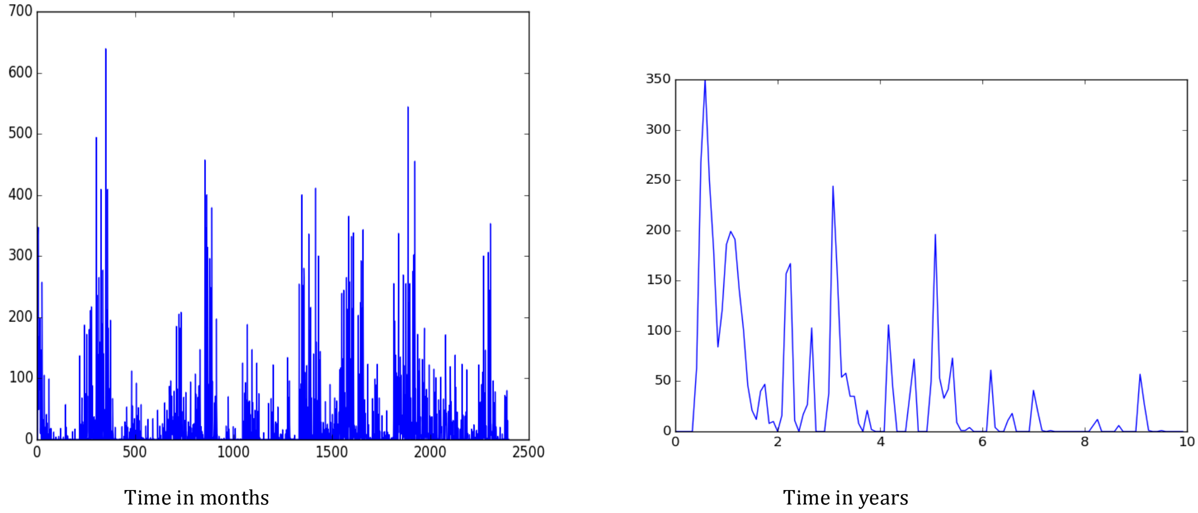

Figure 15: **Left:** Time series of outbreak sizes by month. **Right:** Time series of outbreak sizes by year.

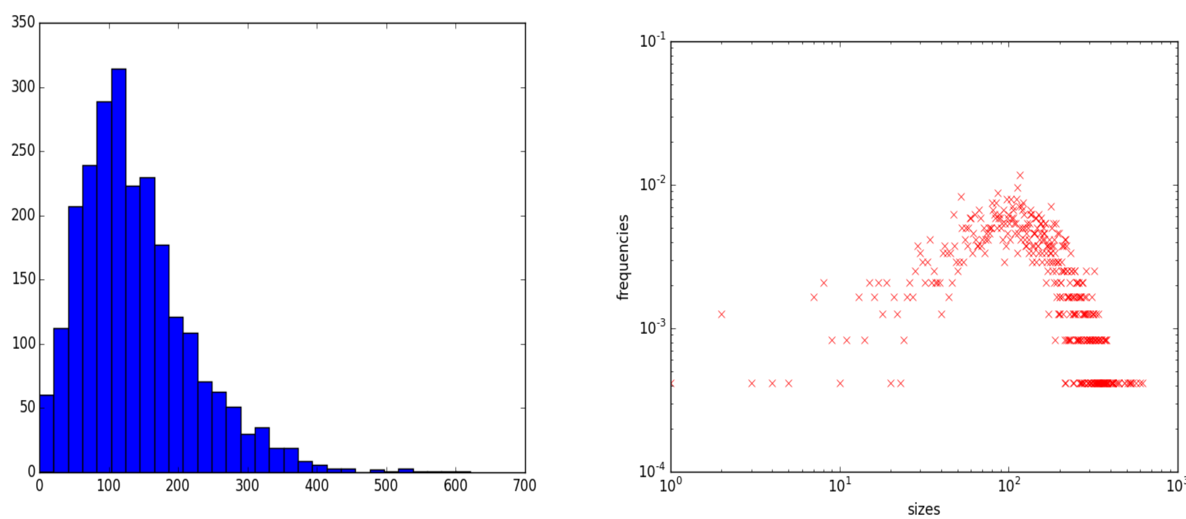

Figure 16: 2°C global warming temperature data without pest control. **Left:** Histogram of outbreak sizes. **Right:** Log-Log Plot of outbreak sizes.

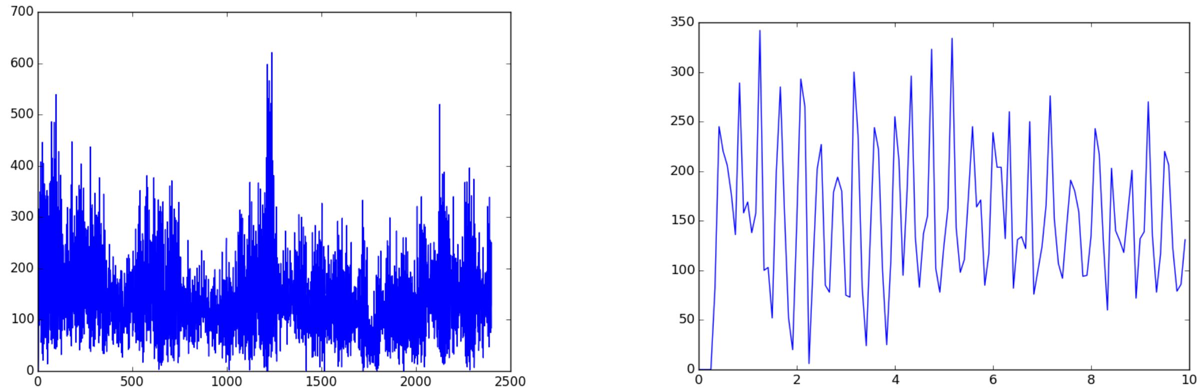

Figure 17: **Left:** Time series of outbreak sizes by month. **Right:** Time series of outbreak sizes by year.

4.6 2°C global warming temperature data with pest control

The addition of pest controls to the warmed scenario has a similar effect as we have noted in the previous iteration. The distribution begins to resemble a power law, however here with a slightly thinner tail. Nevertheless, even in the warmed environment the measures do a reasonable job of controlling the pests, with levels far below the untreated, cooler scenario. The Log-log plot, is however slightly more concave than the previous scenario with pest control.

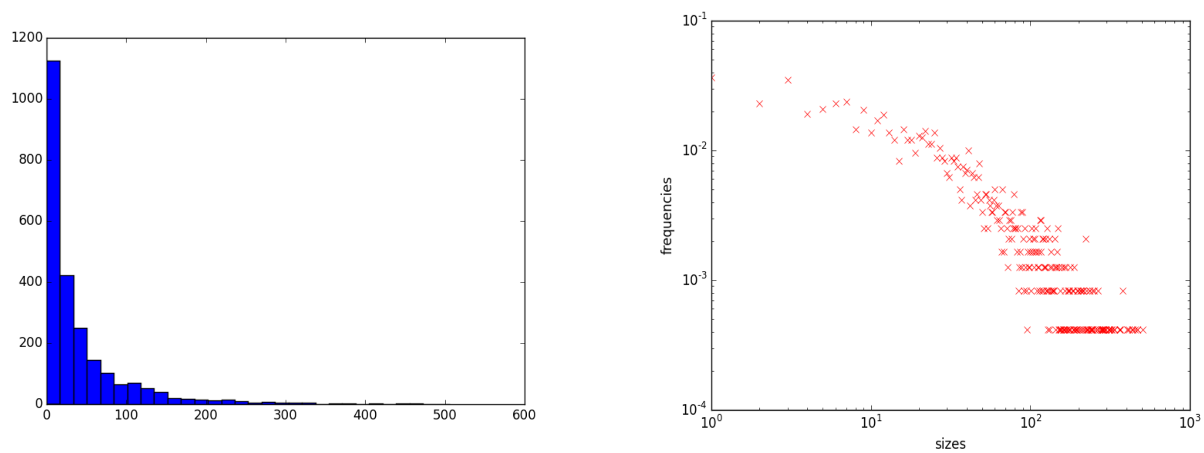

Figure 18: 2°C global warming temperature data with pest control. **Left:** Histogram of outbreak sizes. **Right:** Log-Log Plot of outbreak sizes.

Results for a warming of 4°C are included in Appendix **??**. Under these conditions and without pesticide, the health of the crop is so poor that it cannot maintain a full outbreak. Even with pesticide, it is impossible to full contain the disease.

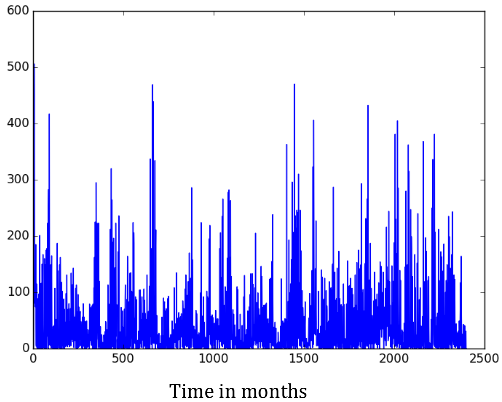

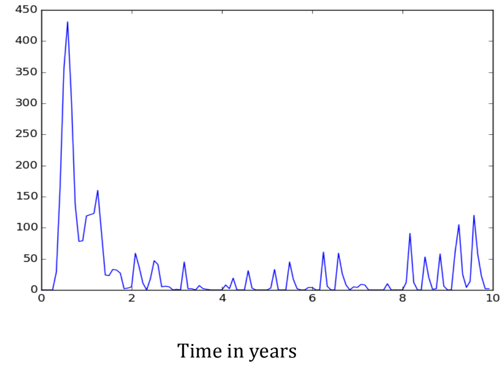

Figure 19: **Left:** Time series of outbreak sizes by month. **Right:** Time series of outbreak sizes by year.

4.7 Discussion

Throughout the drafting and modeling process, we made many other simplifications. We ignored rain and wind as possible spreading agents, instead opting for a random approach. We chose significant parameters such as our time step through very simplified observations of the fungus. We ignored the vegetation cycle (as Guatemala has a very defined wet and dry season which must affect plant growth), though this might be somewhat mitigated by the fact that we tied temperature to the growth of the host plant. We also made considerable simplifying assumptions about the qualities of fungicide application and fungus growth and spread.

Nevertheless, we believe that our results in the change of distribution are representative of what might occur in the real world, given a particular coffee field. We have noted that warming induces a rightward shift in event distribution. The subsequent health decline of the plants may inhibit huge shocks to the system, as the conditions are not ideal for a full fungus takeover. Pest control, while curtailing the infection of an average month can lead to thicker tails and larger rare events. This is possibly because plants are kept at a healthier level, an ideal condition for a quick fungus take over and a drawback of an artificially controlled environment. Additionally, pest control appears to be efficient at compensating for the increased fungus growth rates caused by warming, as even in the 4°C warmer environment, it is able to bring the distribution back to the less disastrous approximate power law, albeit with a mode higher than zero. In the future, under the extreme scenario, it is very possible that some sort of artificial control will be necessary to continue to grow coffee in this region. This will possibly bring more complications and unpredictable dynamics that we cannot comment on with such a simple model.

.1 Computing ENSO impacts

We estimate the impacts of El Niño and La Niña by estimating an "impulse response", which accounts for the multiple overlapping effects of different ENSO years and the monthly climatology of the NINO 3.4 signal.

$$y_t = \alpha + \sum_{Y=Year(t)-N/12+1}^{Year(t)} \sum_{M=1}^{N} \beta_{12(Year(t)-Y)+M}^{Class(Y)} + \gamma \sum_{s=1}^{24} \frac{y_{t-s}}{24} + \mu_{Month(t)}$$

$Year(t)$ is the year for time t and $Month(t)$ is the month for time t; $Class(Y)$ is the class of ENSO event that happened in year Y (El Niño and La Niña). N is the number of months to include in the impulse responses.

Here, the β_m^k variables describe impulse responses of length N for each class of ENSO event.

.2 Additional PCA details

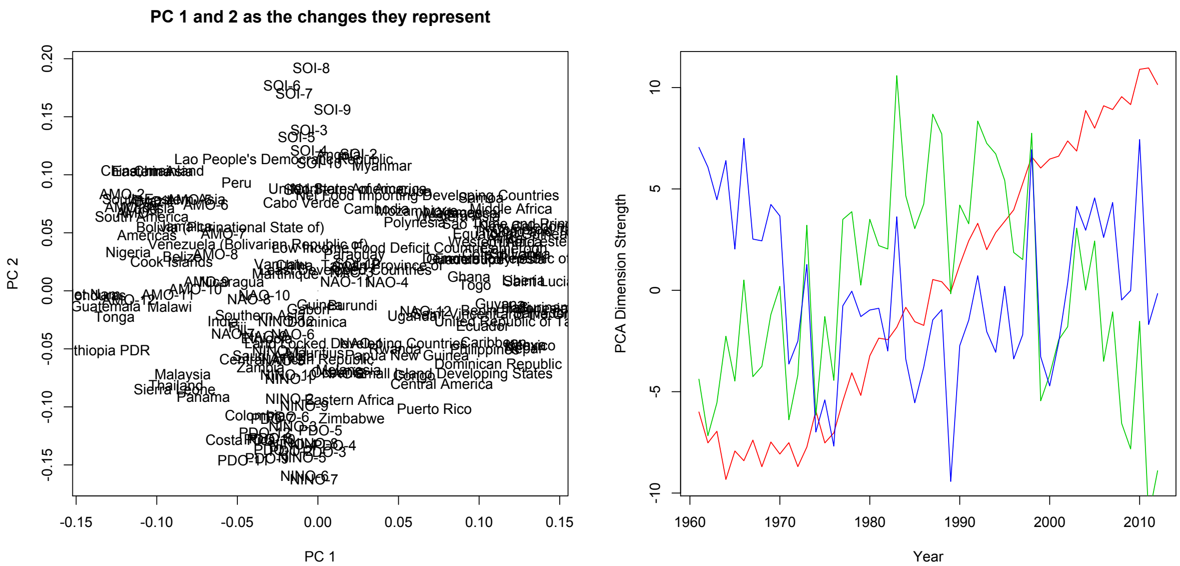

Figure 20: **Left:** The first and second principal component, in terms of the marginal effects of countries and climate signals. These are displayed more clearly in the main report. **Right:** The values of the first three principal components (PC 1 = red, PC 2 = green, PC 3 = blue) across years. As years progress, PC 1 generally increases, and PC 2 first decreases and then increases.

.3 Monthly production

Production records are generally maintained on a yearly basis, but different price information is available monthly. Different countries harvest and ship beans during different months, and this information can be used to infer the monthly production added to the global market.

We use the coffee harvested calendar from the Sweet Maria's Coffee Production Timetable, which is admittedly uncertain and subject to yearly change. However, they provide a general cycle around which actual yearly production is assumed to vary. We distribute the production for each country amongst its harvesting months, and evenly distribute throughout the year production for countries not represented in the calendar (most notably, Vietnam). We also distinguish between countries that produce Arabica and Robusta coffees, or those that produce a combination of both every year. The result in the figure above shows wide variations from month to month.

It is also interesting that the range of variation has increased significantly. The peak of production each year has increased much faster than the yearly minimum: In the 1960s, the best years produced monthly

Figure 21: Inferred monthly production for Arabica and Robusta coffee, based on the harvesting calendars of their producing countries.

peaks over 10 million bags, while the slowest months produced only 3 million bags. In the last decade, the greatest monthly production has been over 15 million bags, but the worst months have only produced 5 million bags. The situation is even starker for Arabica coffee, where the worst months in the last decade are comparable to those in the 1960s, although the best months have increased over 20%.

We can use this monthly production data to inform the coffee market model, described below.

References

Bell, G. D., Halpert, M. S., Kousky, V. E., Gelman, M. E., Ropelewski, C. F., Douglas, A. V., and Schnell, R. C. (1999). Climate assessment for 1998. *Bulletin of the American Meteorological Society*, 80(5):1040–1040.

Enfield, D. B., Mestas-Nunez, A. M., Trimble, P. J., et al. (2001). The atlantic multidecadal oscillation and its relation to rainfall and river flows in the continental u. s. *Geophysical Research Letters*, 28(10):2077–2080.

Epstein, P. R. (1999). Climate and health. *Science(Washington)*, 285(5426):347–348.

Ferreira, S. and Boley, R. (1991). Hemileia vastatrix. *Crop Knowledge Master. Available at: www. extento. hawaii. edu/kbase/crop/type/h_ vasta. htm. Acessed on*, 7(27):2012.

Georgiou, S., Jacques, A., and Imbach Bartol, P. A. (2014). An analysis of the weather and climate conditions related to the 2012 epidemic of coffee rust in guatemala. technical report.

Goddard, L. and Dilley, M. (2005). El niño: catastrophe or opportunity. *Journal of Climate*, 18(5):651–665.

Hsiang, S. M. and Meng, K. C. (2015). Tropical economics. *American Economic Review*, 105(5):257–61.

International Coffee Council (1998). The ”el niño southern oscillation event (enso)” and its impact on coffee production. Technical report. Available at `http://www.ico.org/documents/eb3657r1e.pdf`.

International Research Institute for Climate and Society (2015). IRI/CPC ENSO Predictions Plume. Available at `http://iri.columbia.edu/our-expertise/climate/forecasts/enso/current/`, accessed on Sep. 23, 2015.

International Research Institute for Climate and Society (2016). What changes in rainfall are typical during el niño? `http://iridl.ldeo.columbia.edu/maproom/IFRC/FIC/elninorain.html`. Accessed: 2016-01-25.

Lockwood, D. R. and Lockwood, J. A. (2008). Grasshopper population ecology: catastrophe, criticality, and critique. *Ecology and Society*, 13(1):34.

Magrath, J. (2014). Coffee rust fungus threatens employment collapse in central america.

Mason, S. J. and Goddard, L. (2001). Probabilistic precipitation anomalies associated with enso. *Bulletin of the American Meteorological Society*, 82(4):619–638.

Modern Farmer (2014). Battling the coffee rust: Photos of farmers fighting an epidemic. Retrieved May 16, 2015, from `http://modernfarmer.com/2014/08/battlingcoffee-growers-struggleepidemic/`.

NCAR (2015). Ncep/ncar reanalysis monthly jeans and other derived variables. Retrieved 17 Apr. 2015 from `http://www.esrl.noaa.gov/psd/data/gridded/data.ncep.reanalysis.derived.html`.

NOAA Climate Prediction Center (CPC) (2015). Climate indices: Monthly atmospheric and ocean time series. Available at `http://www.esrl.noaa.gov/psd/data/climateindices/list/`.

Page, S. E., Siegert, F., Rieley, J. O., Boehm, H.-D. V., Jaya, A., and Limin, S. (2002). The amount of carbon released from peat and forest fires in indonesia during 1997. *Nature*, 420(6911):61–65.

Ropelewski, C. F. and Halpert, M. (1989). Precipitation patterns associated with the high index phase of the southern oscillation. *Journal of climate*, 2(3):268–284.

Ubilava, D. (2012). El niño, la niña, and world coffee price dynamics. *Agricultural Economics*, 43(1):17–26.

Villegas, N., Barrientos, J. C., and Málikov, I. (2012). Relationship between ocean-atmospheric parameters and green coffee production in colombia. *Revista Colombiana de Ciencias Hortícolas*, 6(1):88–95.

Acknowledgments

We would like to thank **Walter Baethgen** at the International Research Institute for Climate and Society for his thoughtful reviews of the work here and all of his comments and suggestions.

Many thanks also to:

Andrea Illy, illycaffè S.p.A., Chairman/CEO

Mario Cerutti, Lavazza S.p.A., Corporate Relations

Belay Begashaw, Columbia Global Centers, Director

Mauricio Galindo, International Coffee Organization, Head of Operations

Alexandra Tunistra, Root Capital, Advisory Services

Amir Jina, University of Chicago, Economics

Joann de Zegher, Stanford University, Environment and Resources

Marion Dumas, Columbia University, School of International and Public Affairs