Correction (29 July 2011): Several minor changes were made.

www.sciencemag.org/cgi/content/full/science.1204531/DC1

Supporting Online Material for

Climate Trends and Global Crop Production Since 1980

David B. Lobell,* Wolfram Schlenker, Justin Costa-Roberts

*To whom correspondence should be addressed. E-mail: dlobell@stanford.edu

Published 5 May 2011 on *Science* Express
DOI: 10.1126/science.1204531

This PDF file includes:

Materials and Methods
Figs. S1 to S11
References

Materials and Methods

Materials

We used country-level FAO data on production quantity, yield (quantity per area), and harvested area for maize, rice, soy, and wheat (*S1*). Since the gridded weather data is available on a different geographic and temporal (monthly versus annual) scale than crop data, aggregation was necessary. To do this we use the following sources

(i)The crop calendars compiled by Sacks, et al. (*S2*) give us the growing season of each crop for each 5min grid. We use the average start and end date of the growing season of all 5min grids that fall within the coarser weather grid to establish the growing season.

(ii) The agricultural maps of Monfreda. et al. (*S3*) give the fraction of each 5min grid that is devoted to grow each of the four crops in 2000. We sum the cropland area of all 5min grids that fall within the coarser climate grid to obtain the total amount grown in the cell. Since the yield data are on a country-level, we weight all climate grids in a county by the crop-specific acreages to construct area-weighted averages of all climate grids.

For these relevant growing seasons and agricultural areas, we then extracted monthly precipitation and average temperature data from the University of Delaware Terrestrial Air Temperature and Precipitation dataset (http://climate.geog.udel.edu/), and derived the growing season average precipitation and average temperature data. Finally, we aggregated the area-weighted weather data up to the national level. The weather for a country differs by crop because we average over different areas and potentially different growing seasons. An alternate climate dataset spanning 1961-2002 (CRU TS2.1 (*S4*)) was also tested as a sensitivity analysis (see Figure S6 below), but the regression fit is very comparable.

Maize, rice, and wheat all have second seasons. In the case of maize, we used only the main season because the second season contributes so little to overall production. In rice, there are several major producers with large second-season harvests, and for these we averaged the two seasons' weather data using the production weights in Matthews et al. (*S5*). Finally, in the case of wheat, we used the dominant production season for each country and use the weather data for the 120 days prior to harvest. We do so for two reasons: The distinction between winter and spring wheat was somewhat arbitrary in the crop calendar data, and because wheat is dormant and insensitive to weather through most of the winter season.

Methods

To properly identify weather effects in these panel datasets requires that “noise” from other factors are accounted for. As long as these other variables are uncorrelated to weather, they will not bias the results, but could impact the efficiency of our estimator. The main issue is the gradual trend in yields due to technology changes. Detrending is necessary to ensure that the regression is not biased from common trends in the data (i.e. a positive yield trend during a period of warming does not imply that warming is beneficial, since other factors drive the yield trend). In general, correlating two trending variables often leads to correlations that are not causal. Instead, our regression has state-specific quadratic time trends. Hence the model estimates for weather effects rely on year-to-year variations in weather and yields, and not common trends.

In some cases, countries exhibited highly nonlinear trends, with sudden jumps in the data, which means that a quadratic time trend would poorly fit the data. Rather than risk bias in our results, these countries were simply omitted from the analysis, based on tests of the 10-year lag autocorrelation of residuals from a linear trend. We also omitted the time series of a given crop and country if: (a) there were three or more instances when the FAO yield data remained exactly the same for two years in a row or (b) the cropland area did not exceed a threshold value of 10,000 hectares. All of the above criteria (sudden jumps, no change, or low area) are suggestive of noisy data and only affected very minor producers. Since the current analysis focuses on global production, omitting these small producers has little effect.

Responses to weather can depend on limitations from other factors, and given the differences in growing practices and conditions between countries, it makes sense to group nations with similar agricultural attributes before estimating the response of yield to the weather variables. For example, higher fertilizer rates can allow yields to fluctuate much more in response to weather (*S6*). Our preferred method is a four-group split corresponding to average yield quartiles, to test for heterogeneous climate sensitivities. To assess the sensitivity to the grouping method, we also ran the model on a five-group specification with low and high-yielding (below or above the median) low latitude countries, low and high-yielding high latitude countries, and North African/Middle Eastern countries separate to address the concern that including heavily irrigated countries with rainfed ones might obscure the responses for both. Finally, we ran a global pooled model, with no grouping at all. The results are qualitatively similar (See Figure S8 below).

We estimate a model of log yields -- log(Yi,t) -- as yields follow a log-normal distribution and yield variance stays comparable in relative, but not absolute terms. Transforming to log yields results in a more normally distributed variable. A log specification is also preferable because it assumes that a given change in temperature results in the same percent impact. Different countries have vastly different average yields that can vary by an order of magnitude. Assuming that temperature fluctuations have the same absolute effect in each country would imply that the same temperature fluctuation that causes a 10% change in one country leads to a 100% change in yield in another country with an average yield that is an order of magnitude lower. The model is of the form

$$\mathrm{Log}(Y_{i,t}) = c_i + d_{1i} * \mathrm{year} + d_{2i} * \mathrm{year}^2 + \beta \cdot X_{i,t} + \varepsilon_{i,t} \qquad (1)$$

where c_i is a country fixed effect, d_{1i} is the country-specific linear time trend, d_{2i} is the country-specific quadratic time trend, β is a vector of coefficients and X is a vector of variables (T, T^2, P, P^2). Following the literature we use quadratic terms for T and P to account for the fact that crops perform best at moderate T and P, and are harmed by extreme cold, hot, dry, or wet conditions. In a sensitivity check below (Fig. S11), we include separate variables for minimum and maximum temperatures. Regressions that weighted each country by its crop area were also run, with very similar results.

The fixed effects model has the advantage of accounting for time-invariant country differences (e.g., soil quality), thereby removing omitted variable bias if these time-invariant difference are correlated with the weather variables of interest. Country-specific time trends capture varying rates of progress among countries. Robustness checks indicated that results were nearly identical when including a single coefficient for year-squared rather than by country, or omitting year-squared altogether (see below Figure S10). This reflects the fact that log yield trends are very linear for most crops and countries included in the analysis.

The impact of climate trends was calculated in a four-step process. First we introduce some notation. Let

$T^*_{i,t}$ = predicted temperature for country i at year t, based on linear fit for 1980-2008

$Td_{i,t}$ = detrended temperature = $T_{i,t} - T^*_{i,t} + T^*_{i,1980}$

With detrended precipitation similarly defined, we then use the regression model (F(T,P)) to compute

(i) F(T,P) = predicted yields with observed weather
(ii) F(Td,P) = predicted yields with detrended temperature and observed precipitation
(iii) F(T,Pd) = predicted yields with observed temperature and detrended precipitation
(iv) F(Td,Pd) = predicted yields with detrended temperature and detrended precipitation

We compute the trends of the differences (i) – (ii), (i) – (iii), and (i)-(iv) to quantify the yield effect of trends in T, P, and both T and P, respectively.

Estimates of the 5th to 95th percentile confidence interval were obtained in all cases by bootstrap resampling, where the historical data were resampled and a new regression model was fit to the data. In each case, 500 bootstrap samples were used, so that e.g. the 5th percentile corresponds to the 25th lowest value.

Finally, to estimate price effects we first aggregate the output of all four crops based on the caloric content of the crop. Following (*S7*), we use a demand elasticity of -0.05 that implies a doubling of prices reduces global demand by roughly 5 percent, and a supply elasticity of 0.12 that implies that price doubling increases supply by 12 percent, primarily through an expansion of the growing area. The estimated price impact is simply the reduction in calories produced divided by the difference between the supply and demand elasticity. Since the price response depends on the inverse of a parameter that is normally distributed, the ratio will be non-normally distributed. We hence use 1 million bootstrap simulations from the joint distribution of the two parameters and take the average value to obtain the predicted price increase (*S7*). The estimated changes in crop production excluding and including CO_2 fertilization (columns (4) and (6) of Table 1, respectively) translate into average predicted commodity price increases of 18.9% and 6.4%.

Panel Regressions and Adaptation

It is often wrongly assumed that empirical regressions of historical yields and weather data provide predictions of climate response that assume no farmer adaptation. This is based on the perception that statistical models rely only on year-to-year variations in weather, which the farmer cannot anticipate in advance, whereas climate change represents gradual change that the farmer can adjust to. As a result, projections from statistical models are often viewed similarly to simulation results from process-based models where the management is held fixed over time (i.e. the "dumb farmer" approach).

In reality, there are many types of statistical approaches. More detailed descriptions are provided in (*S8*) but a brief description is given here. At one extreme, there are time series models that examine for a single location (e.g., a field or a country) the relationship over time between weather and yields. As such, these do ignore any adaptations other than those that farmers make on a year-to-year basis, such as applying extra irrigation during especially hot and dry periods, or planting earlier in a warmer spring. At the other extreme are cross-section models, which rely solely on comparing locations across space. These models should include all adaptations that are observed when moving across different climates, such as farmers using different varieties and sowing dates in hotter locations. Comparisons between time series and cross-section approaches can therefore be a useful bound on how effective adaptation has been historically in maintaining yields of a given crop in different climates. For example, Schlenker and Roberts (*S9*) compare time series and cross-section models for U.S. maize yield response to extreme heat, and find very similar sensitivities implying little adaptation.

In between these two extremes are panel models, which combine time series from different locations into one analysis. Intuitively, the amount of adaptation assumed in these models is somewhere in between the other two approaches, but the exact amount depends on how much the model is relying on the year-to-year variations within sites relative to variations across sites. In panel models, fixed effects (i.e. dummy variables for each location) are used to control for omitted variables that can explain differences between sites, thereby avoiding bias from omitted variables that often cast doubt on cross-sectional models. The use of fixed effects is equivalent to a joint demeaning of both the dependent and the independent variables. In other words, only deviations from the average are used in the estimation. If fixed-effects are used in a panel model with linear terms for weather variables (e.g., temperature or precipitation), then all of the information from the model comes from variations within sites. However, if fixed-effects are used in a model that has quadratic terms for weather (such as in the current study), then both within-site and across-site differences in weather are used. The quadratic term implies that the marginal effect can vary with the average weather outcome in a location, and hence differences between locations enter the identification. (The mean weather at each location is still used to identify the coefficient as demeaning a squared variable is different from the square of the demeaned variable). More explanation and examples of how panel models capture adaptations can be found in ref (*S10*).

Thus, panel regressions with quadratic weather terms overcome the main problem with cross-sectional models (omitted variable bias) while also capturing some aspects of adaptation (e.g., through changes in sowing date or variety) by allowing the marginal effect to vary with climate. However, when the analysis is limited to a single crop, panel models will not capture adaptations such as crop switching. Finally, it should be noted that because statistical models rely on observations, they cannot anticipate future innovations such as development of varieties not currently grown anywhere in the locations used to fit the model. Of course, any modeling approach will find it difficult to quantify the effects of future advances, but for an example see (*S11*). It is also possible to use statistical models with experimental trial data (*S12*), in which case the response of varieties not yet grown in farmers' fields can be evaluated.

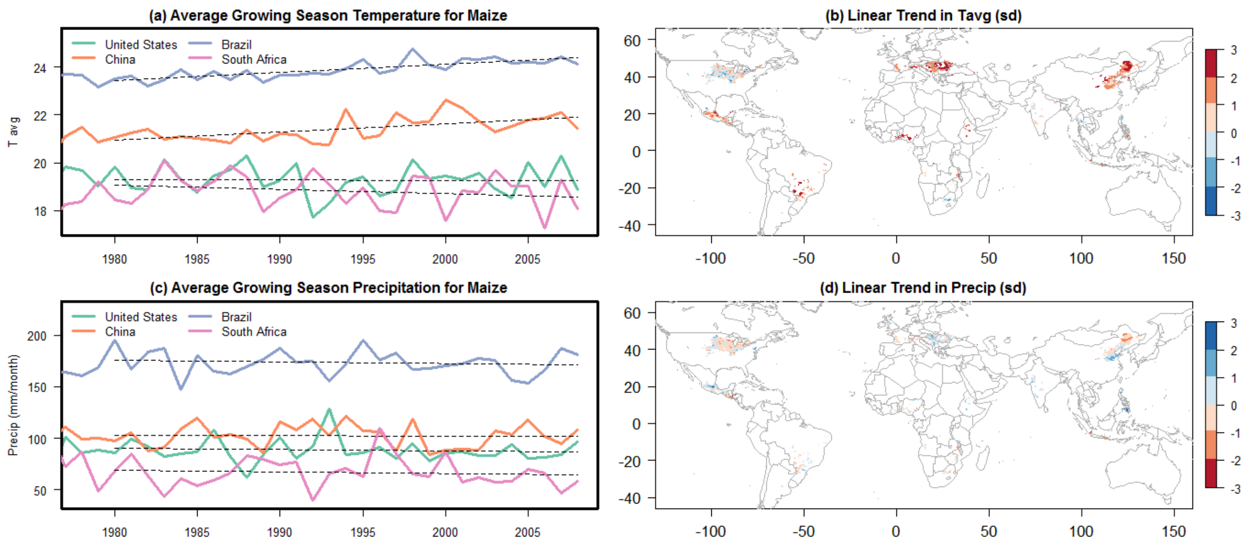

Fig. S1.

(a) Time series of growing season average temperature for maize in four major producing countries and (b) maps of linear trend for 1980-2008 for major growing areas. (c)-(d) same as (a)-(b) but for precipitation. Dashed lines in (a) and (c) show linear trends for 1980-2008. Units of trends in (b) and (d) are the total trend for the 29-year period (e.g. ˚C per 29 years), divided by the historical standard deviation for the 1960-2000 period. Only cells with at least 5% of area in maize are shown.

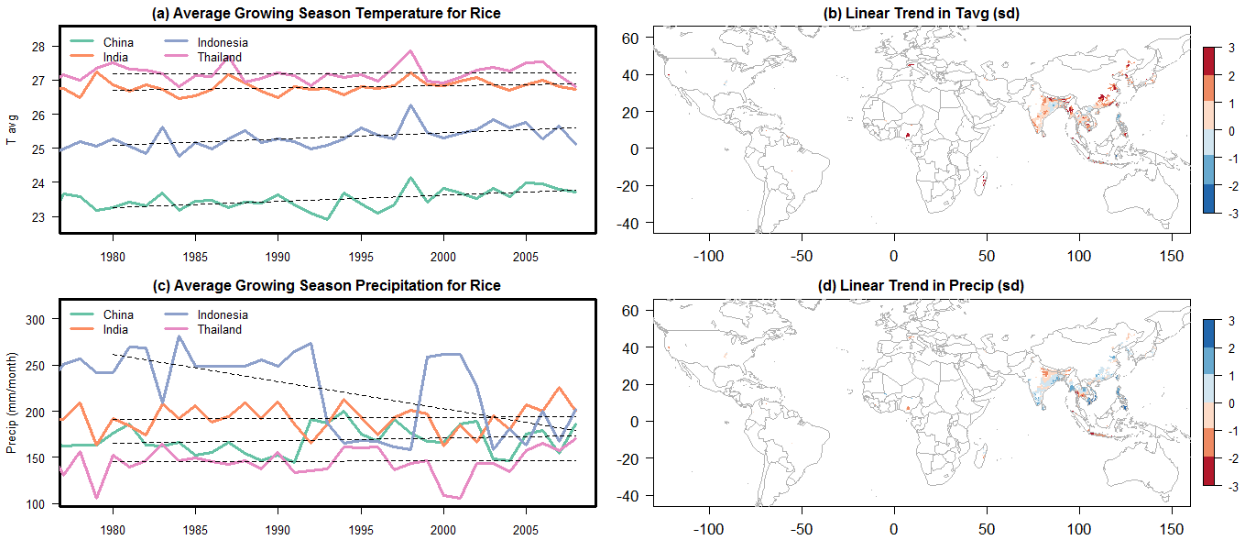

Fig. S2

(a) Time series of growing season average temperature for rice in four major producing countries and (b) maps of linear trend for 1980-2008 for major growing areas. (c)-(d) same as (a)-(b) but for precipitation. Dashed lines in (a) and (c) show linear trends for 1980-2008. Units of trends in (b) and (d) are the total trend for the 29-year period (e.g. ˚C per 29 years), divided by the historical standard deviation for the 1960-2000 period. Only cells with at least 5% of area in rice are shown.

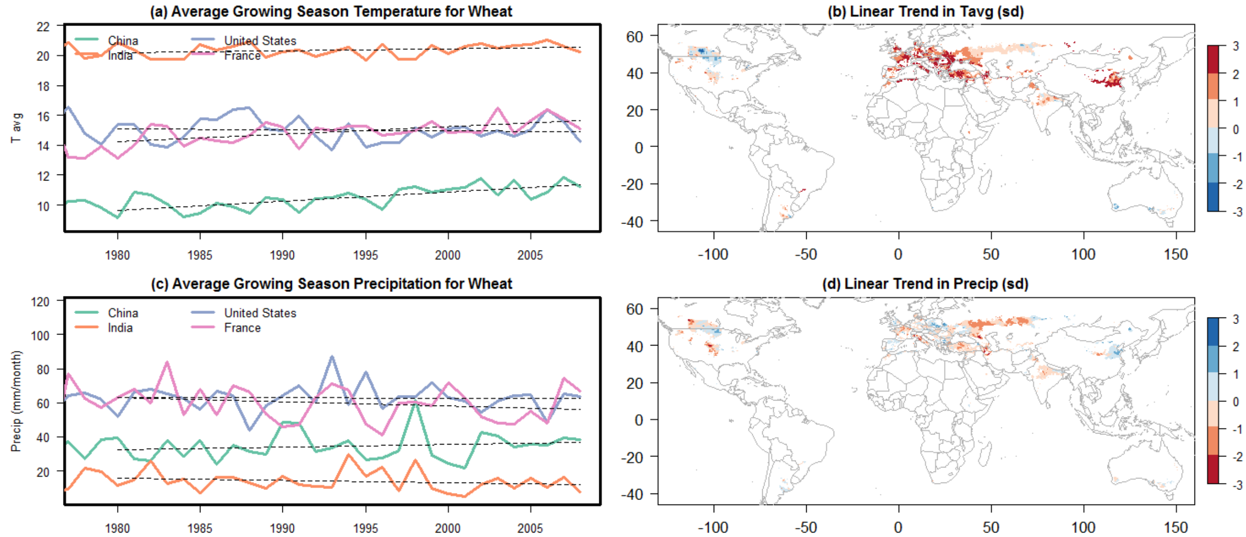

Fig. S3

(a) Time series of growing season average temperature for wheat in four major producing countries and (b) maps of linear trend for 1980-2008 for major growing areas. (c)-(d) same as (a)-(b) but for precipitation. Dashed lines in (a) and (c) show linear trends for 1980-2008. Units of trends in (b) and (d) are the total trend for the 29-year period (e.g. ˚C per 29 years), divided by the historical standard deviation for the 1960-2000 period. Only cells with at least 5% of area in wheat are shown.

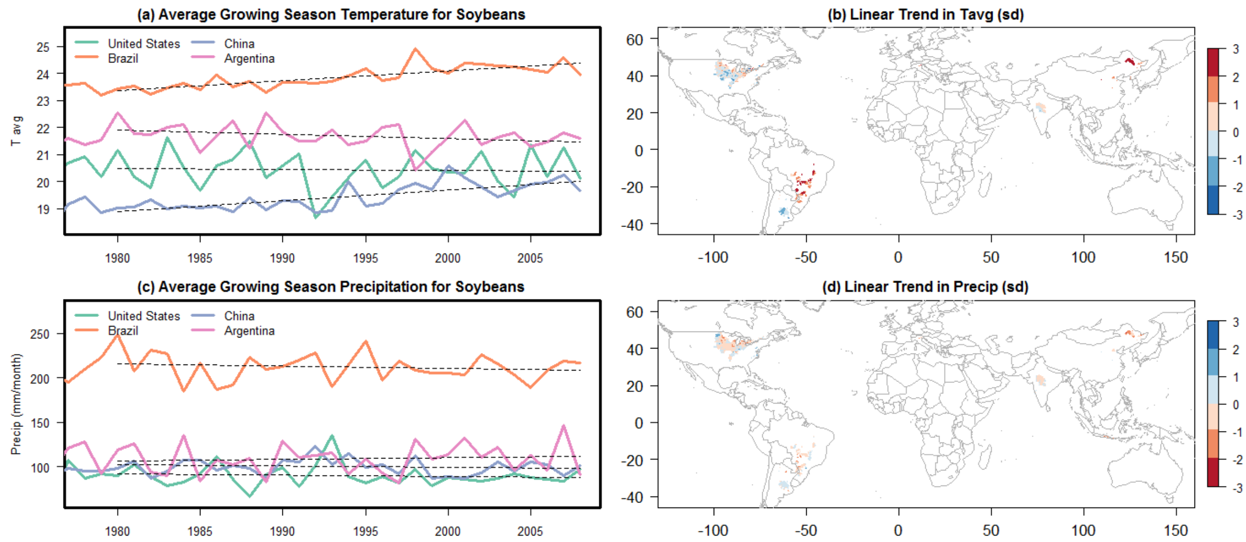

Fig. S4

(a) Time series of growing season average temperature for soybean in four major producing countries and (b) maps of linear trend for 1980-2008 for major growing areas. (c)-(d) same as (a)-(b) but for precipitation. Dashed lines in (a) and (c) show linear trends for 1980-2008. Units of trends in (b) and (d) are the total trend for the 29-year period (e.g. ˚C per 29 years), divided by the historical standard deviation for the 1960-2000 period. Only cells with at least 5% of area in soybean are shown.

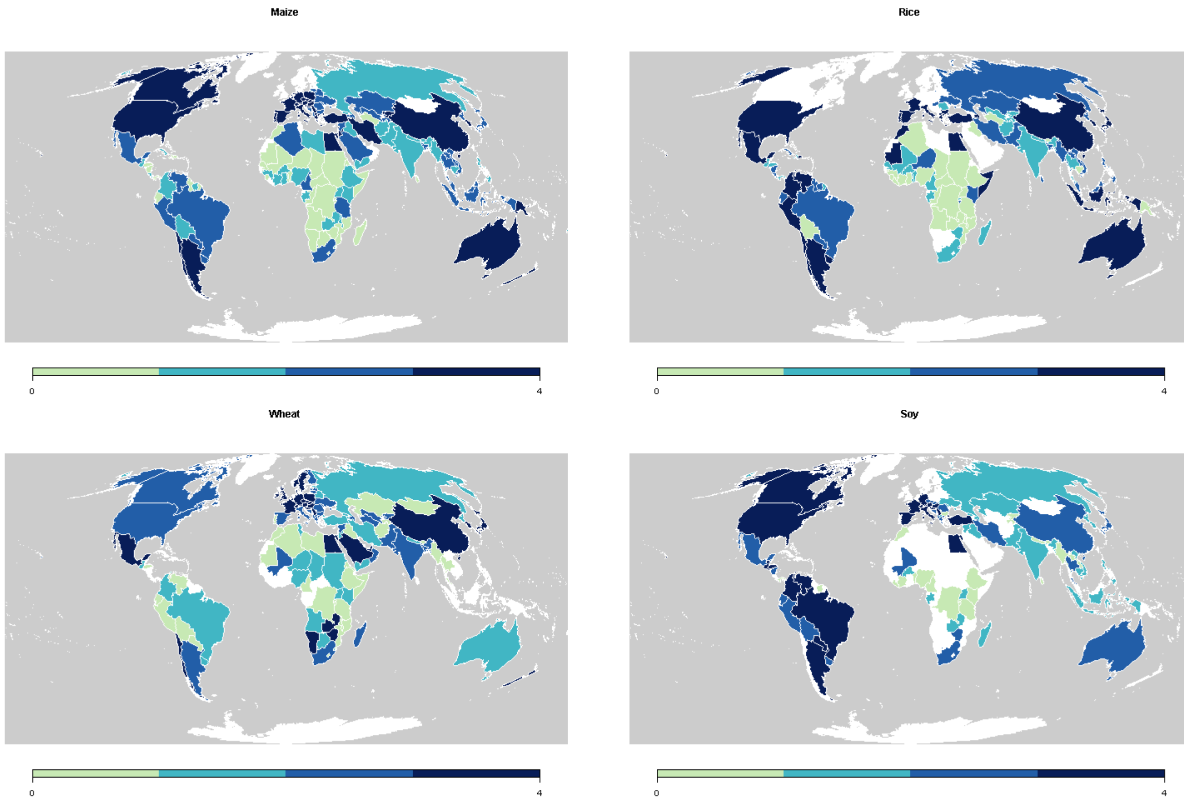

Fig. S5

Maps showing yield quartiles for each crop, which were used to define country groupings for panel regressions. A separate regression was fit to each group, in order to capture possible interactions between weather responses and management intensity (e.g. fertilizer and irrigation rates). White indicates countries where FAO does not report yield for that crop.

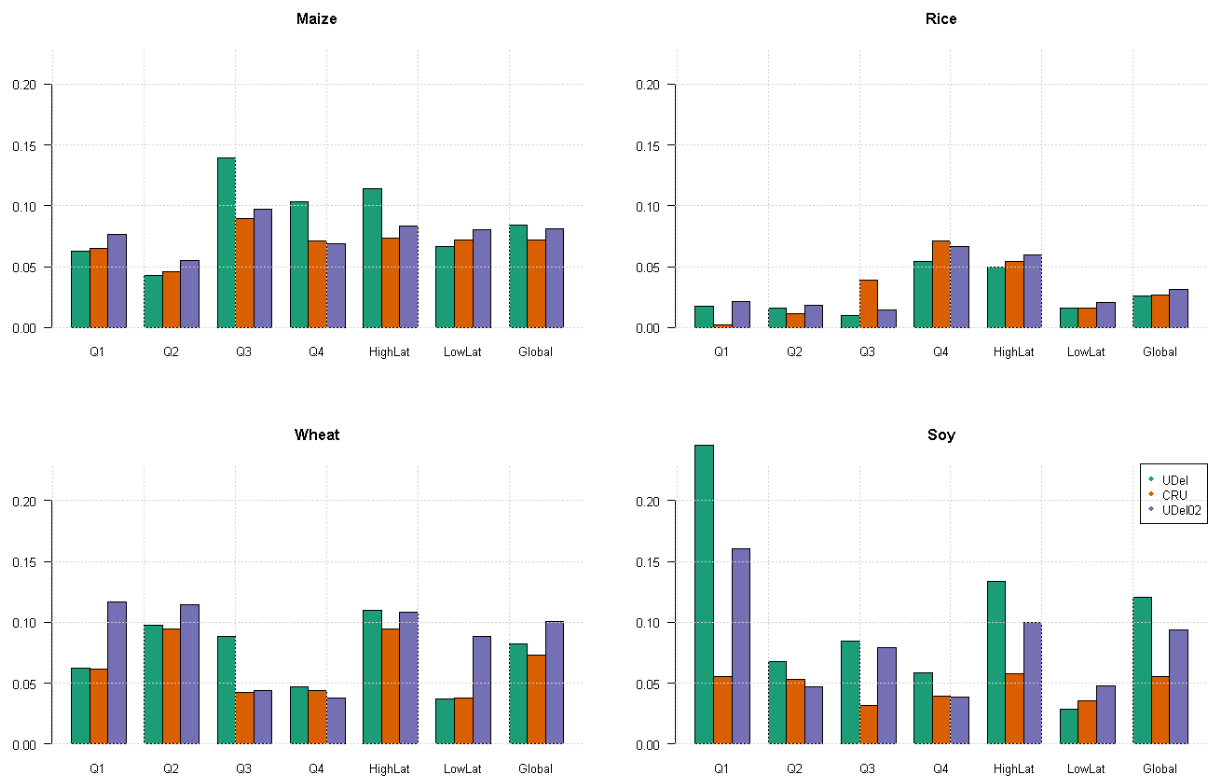

Fig. S6

Summary of regression fits, showing the amount of additional variance explained by T and P compared to a fixed effects model with only countries and time trends. Values show the percent of remaining variance explained by T and P, for different groups of countries (e.g., Q1=quartile 1, HighLat=high latitude countries, Global = all countries). Different colors represent models fit with three climate datasets (full University of Delaware dataset, 1960-2008 (UDel); Climate Research Unit TS2.1 Dataset, 1960-2002 (CRU), and University of Delaware dataset over same time period as CRU, 1960-2002 (UDel02).

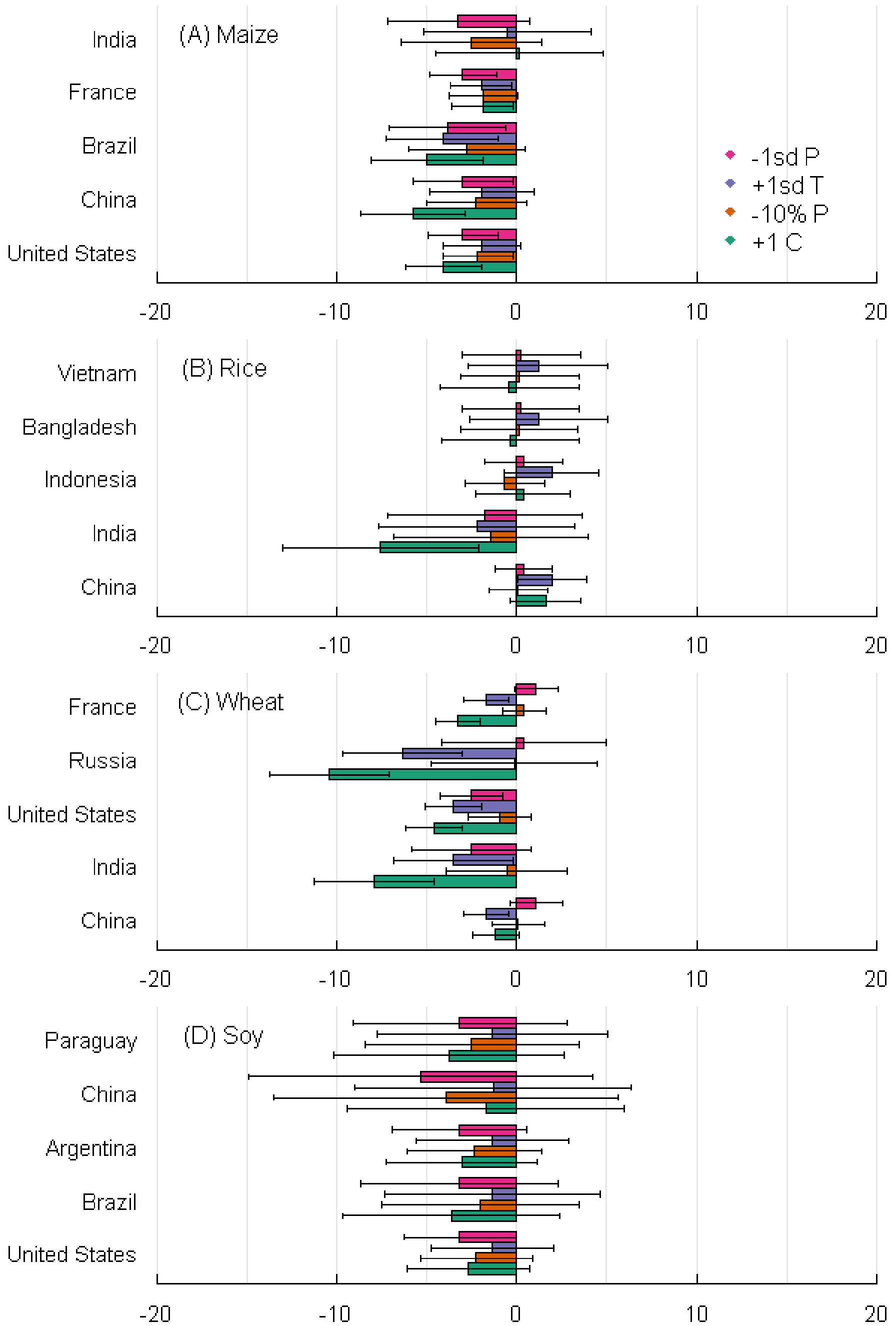

Fig. S7

Summary of regression model coefficients, showing estimated sensitivity to changes in T and P for major producers of each crop. Estimated yield changes (in %) are shown both for +1˚C and -10% P, as well as a 1σ change in each.

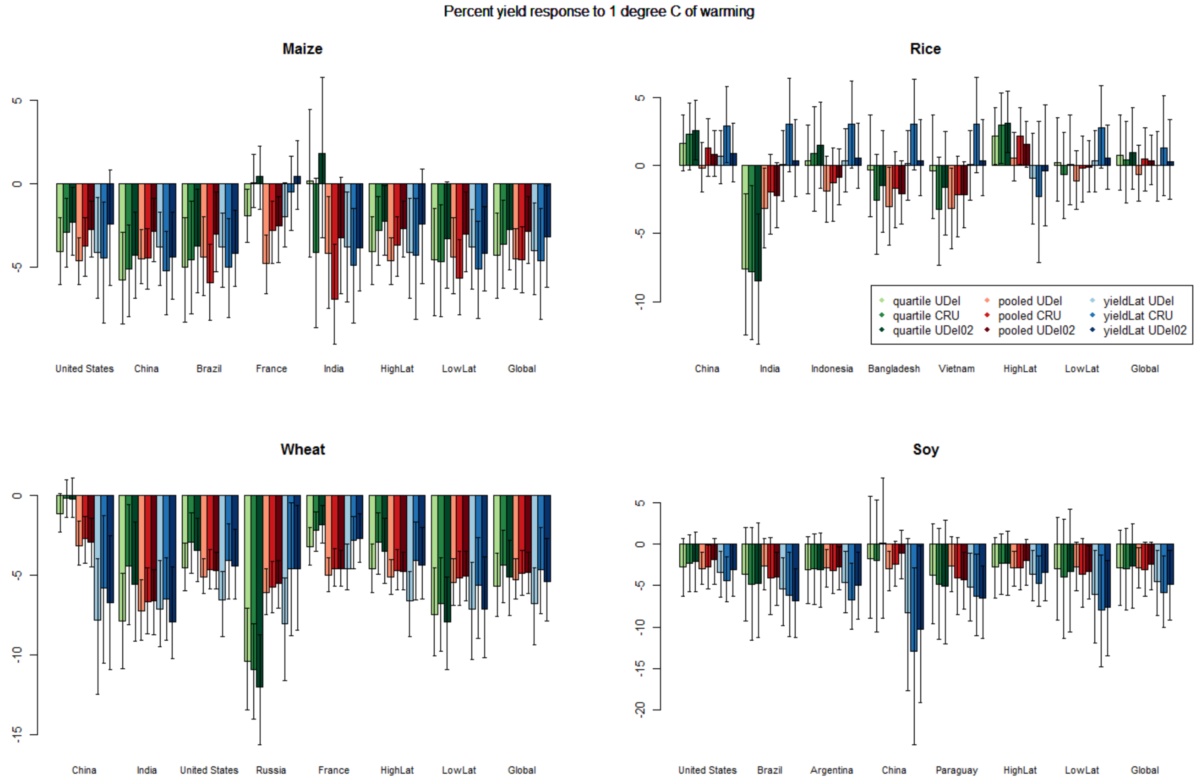

Fig. S8

Summary of sensitivity of regression results to different country groupings or climate datasets. Estimated yield response to 1˚C warming are shown for top five producers of each crop as well as aggregate totals for high and low latitude countries and global aggregates. Countries without estimates were omitted because of highly nonlinear yield trends. Results are shown for three groupings (grouping used in main paper by yield quartile (quartile), grouping all countries into one global panel (pooled), and grouping countries first by latitude range and then by whether yield was above or below median yield for given latitude group (yieldLat)) and for three climate datasets (full University of Delaware dataset, 1960-2008 (UDel); Climate Research Unit TS2.1 Dataset, 1960-2002 (CRU), and University of Delaware dataset over same time period as CRU, 1960-2002 (UDel02). All models contain a quadratic time trend by country.

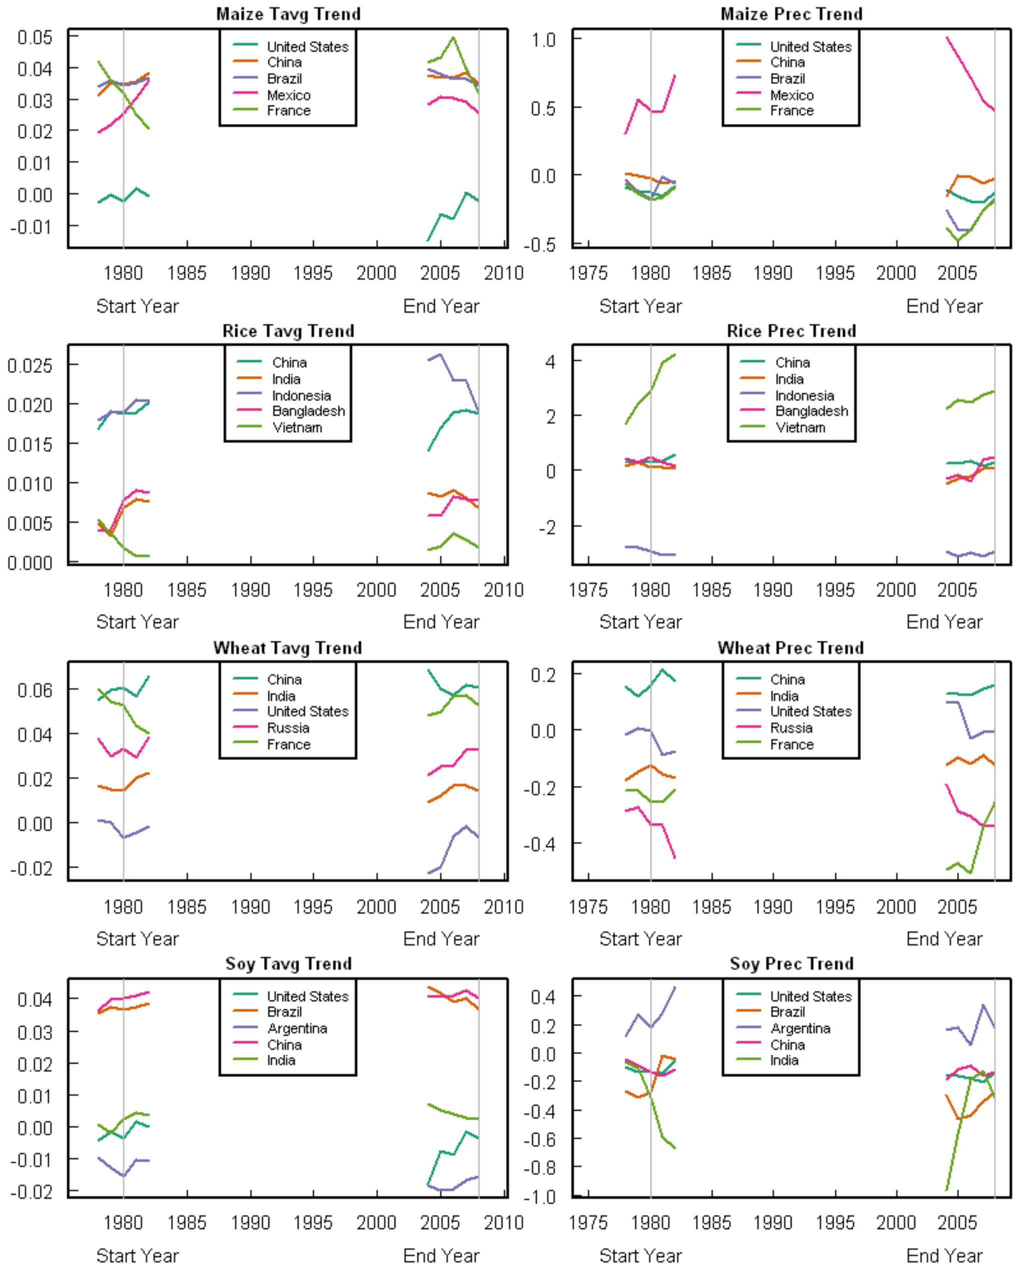

Fig. S9

Sensitivity of T and P trends to different start and end years, for top five producers of each crop. The vertical lines show the start and end years used in this study (1980 and 2008). There were no systematic effects of moving the start or end year, with some trends increasing and others decreasing. Changes were generally smaller than the differences between countries, with some exceptions.

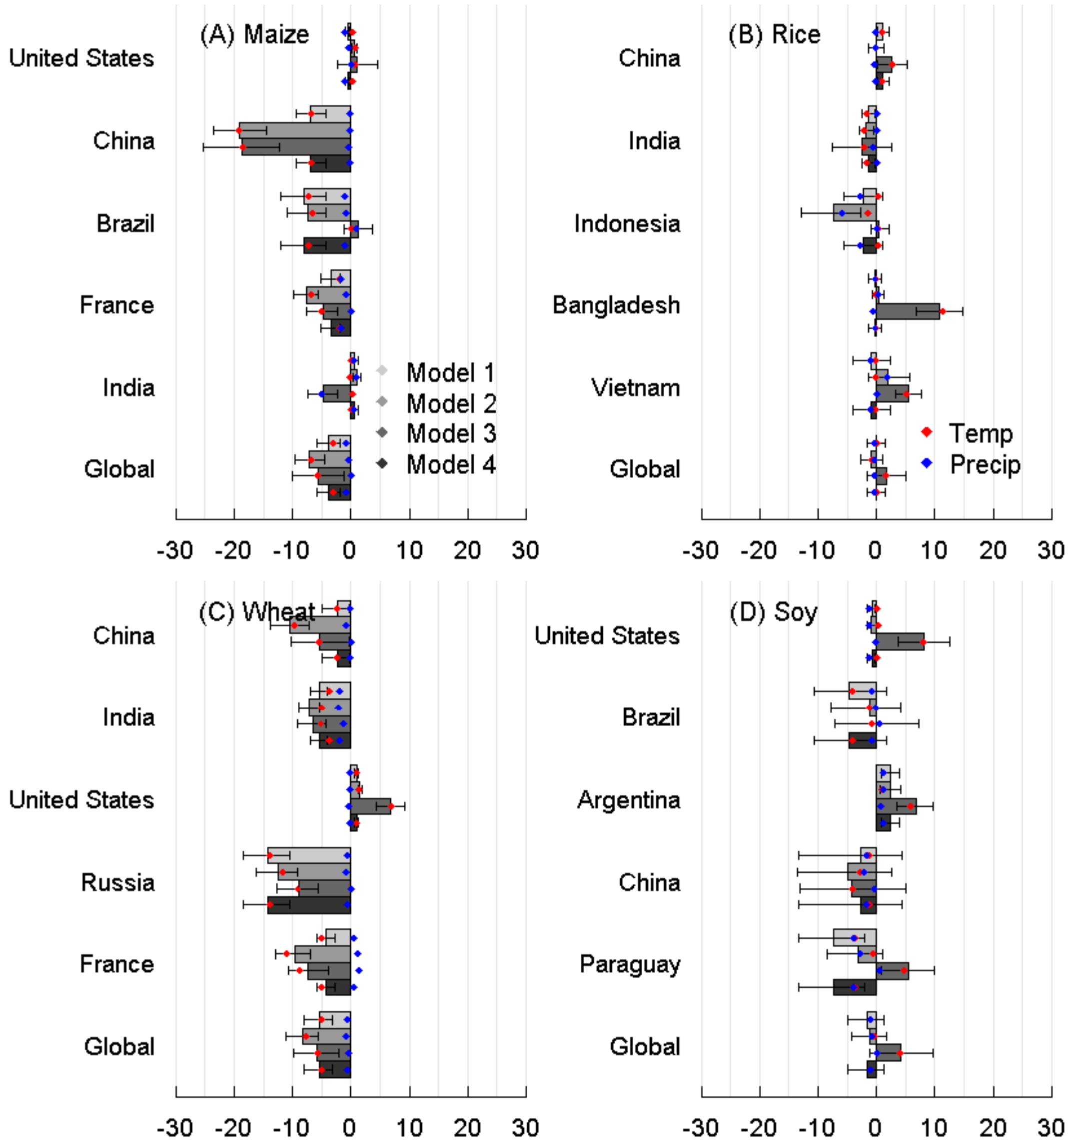

Fig. S10

Same as Figure 3 in main paper but showing results for different model specifications. Model 1 = quadratic time trends by country (main model used in paper), Model 2 = only linear time trends by country, Model 3 = linear time trend by country, pooled quadratic trend, and Tmin and Tmax modeled separately instead of using Tavg, Model 4 = linear time trend by country, pooled quadratic trend, and Tavg. Model 3 uses CRU climate data to 2002, whereas others use UDel climate data to 2008. CRU impacts are scaled to 39 years (i.e. multiplied by 39/33) to match time scale of for UDel based results.

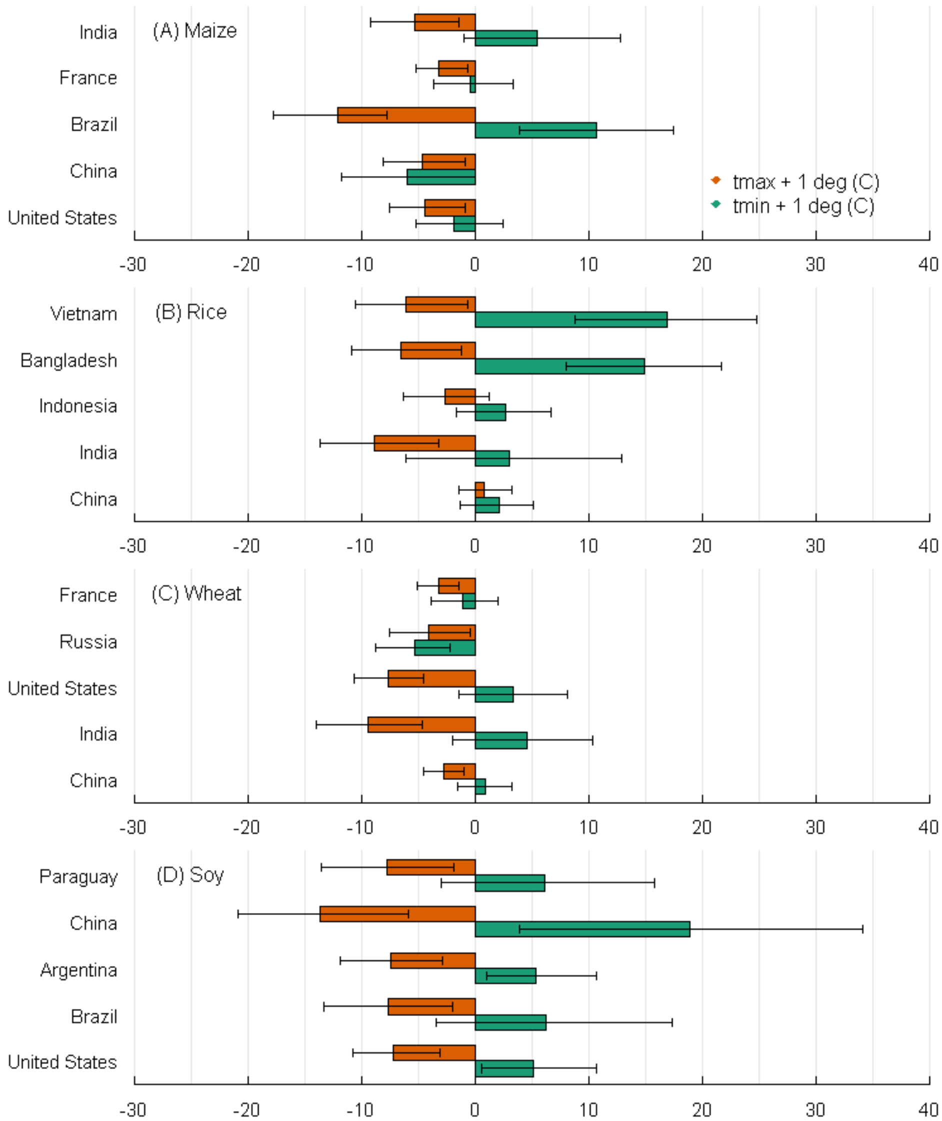

Fig. S11

The inferred importance of maximum (Tmax) vs. minimum (Tmin) temperature from a model that includes both separately. Bars show inferred sensitivity of yields to a 1°C rise in Tmax or Tmin, with error bars showing 5-95% confidence interval. Large error bars reflect the fact that correlation between the two are high in most regions ($r > 0.7$), making it difficult to uniquely identify the effect of each. The overall greater importance of Tmax is consistent with previous studies that emphasize the importance of extreme heat (*S9*, *13*). This sensitivity analysis used the CRU dataset because UDel contains only average temperature. We emphasize that covariance between Tmax and Tmin make interpretation of any single coefficient difficult (i.e. a positive coefficient on Tmin does not necessarily mean that yields are helped by increased Tmin).

References

S1. FAO, *Food and Agriculture Organization of the United Nations (FAO), FAO Statistical Databases*. http://faostat.fao.org

S2. W. J. Sacks, D. Deryng, J. A. Foley, N. Ramankutty, *Global Ecology and Biogeography* **in press**, (2010).

S3. C. Monfreda, N. Ramankutty, J. A. Foley, *Global Biogeochem. Cycles* **22**, (2008).

S4. T. D. Mitchell, P. D. Jones, *International Journal of Climatology* **25**, 693 (May, 2005).

S5. R. B. Matthews, M. J. Kropff, D. Bachelet, H. H. Van Laar, *Modeling the Impact of Climate Change on Rice Production in Asia*. (CABI, Los Banos, Philippines, 1995).

S6. W. Schlenker, D. B. Lobell, *Environmental Research Letters*, 014010 (8pp) (2010).

S7. M. Roberts, W. Schlenker, in *NBER Working Paper No. 15921*. (2010).

S8. D. Lobell, M. Burke, *Climate Change and Food Security: Adapting Agriculture to a Warmer World*. (Springer Verlag, 2009).

S9. W. Schlenker, M. J. Roberts, *Proceedings of the National Academy of Sciences* **106**, 15594 (2009).

S10. W. Schlenker, *Working Paper "Inter-annual weather variation and crop yields"*, (2006).

S11. A. J. Challinor, E. S. Simelton, E. D. G. Fraser, D. Hemming, M. Collins, *Environmental Research Letters* **5**, 034012 (2010).

S12. D. B. Lobell, M. Banziger, C. Magorokosho, B. Vivek, *Nature Clim. Change* **1**, 42 (2011).

S13. D. B. Lobell, *Agricultural and Forest Meteorology* **145**, 229 (2007).