Supplementary Materials for

Reckoning wheat yield trends

Marena Lin and Peter Huybers
Department of Earth and Planetary Sciences
Harvard University
20 Oxford St., Cambridge MA, 02138

Contents

1 Replicating interannual autocorrelation with phase randomization

Wheat yield data is temporally autocorrelated as plant growth is influenced by features such as cultivar selection, fertilization schedules, and planting practices as well as climate variability that may show persistence from year to year. In supplemental figure 1, we provide the example of French wheat yield with both the rising-plateau (red) and linear (blue) trends. The increasing trend in the data is perhaps the most salient example of autocorrelation. The variation about the trend also shows autocorrelation, wherein a deviation above or below the trend is likely to be followed by other deviations in the same direction. Interannual autocorrelation between yield data points means that deviations are not identically and independently distributed, resulting in a lower number of degrees of freedom than the number of data points in the set.

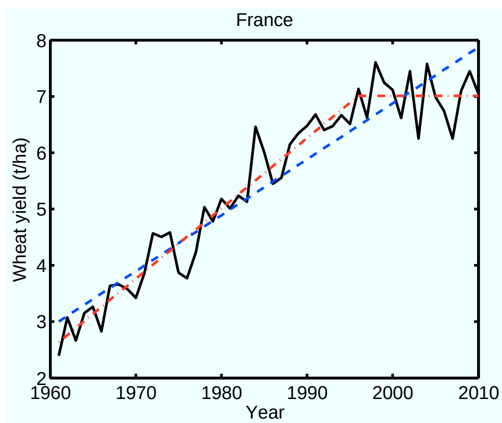

Figure 1: National French wheat yield (tons/hectare); rising-plateau model (red), linear model (blue)

As described in the main text, phase randomization can replicate autocorrelation in surrogate data sets. If autocorrelation were not accounted for, the surrogate residuals would be composed of white noise so that points would be independent of one another. Surrogate data sets with white noise residuals result in narrower null distributions of δJ because autocorrelated points are more likely to deviate systematically in the same direction than a series of independent points, leading to a higher likelihood of the surrogate data being better fit by a rising-plateau model, as opposed to the purely linear model.

To illustrate the difference in δJ distributions when autocorrelation is, and is not, accounted for, we reran our analysis using white noise for Colombian wheat yield. In the white noise case, the observed δJ appears to lead to being able to reject the null hypothesis with extremely high confidence (p-value=0.0001, supplemental figure 2a), whereas accounting for autocorrelation indicates that the result is much less significant (p-value=0.022, supplemental figure 2b). Similarly, the power of the test is inflated when using white noise (power=0.77, supplemental figure 2a) relative to the case of accounting for autocorrelation (power=0.45, supplemental figure 2b). Although both tests show the observed leveling in Colombian wheat yields to be significant, this example also illustrates how not accounting for autocorrelation could readily lead to incorrect rejection of the null hypothesis for a dataset whose results were more ambiguous.

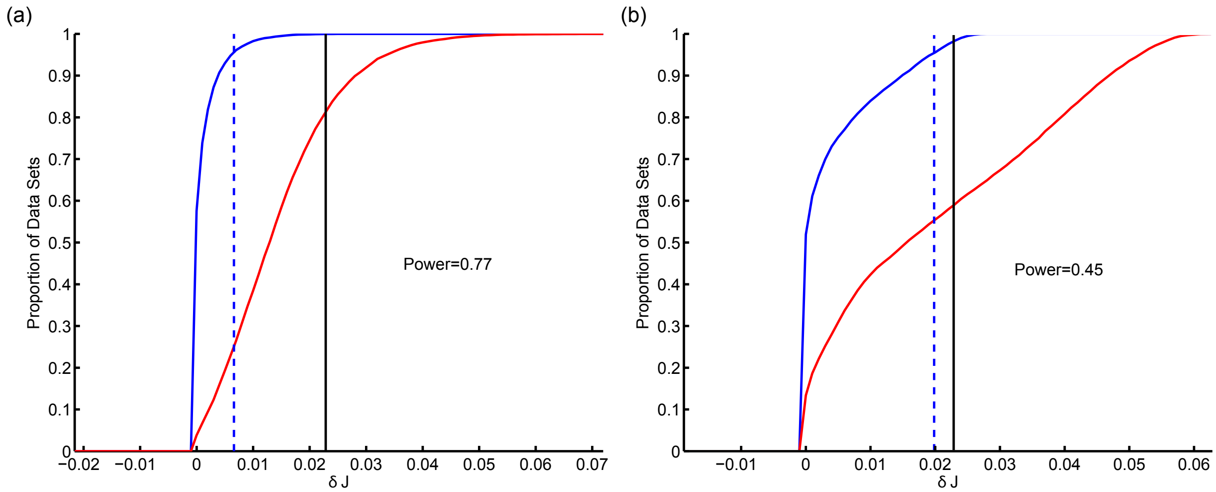

Figure 2: The cumulative sum distributions of δJ for the null (blue) and the alternative (red) hypotheses for Colombian wheat yield. The observed δJ for the original data is marked by the black line. Two cases are shown where the residuals around the model fit are modeled using (a) white noise and (b) autocorrelated variations.

2 Differences in methodology with Hafner (2003) and Brisson (2010)

Hafner (2003) found that poorer nations showed slowing growth of wheat yields, and there are several possibilities for the difference in findings. The first has to do with the difference in the time interval over which the analysis was conducted. However, rerunning our analysis using the shorter 1961-2001 time interval consistent with Hafner (2003)'s study results in only nine nations leveling at the 80% confidence level—Albania, China, Finland, Greece, Italy, Japan, Poland, Spain, and Turkey. With the exception of China, these countries also show leveling wheat yields in our 1961-2010 analysis at 80% confidence. Moreover, these nations are not consistent with the finding of poorer nations showing signs of slowed growth.

A second possibility for the difference in findings is that the functional form fitted by Hafner (2003) is sensitive to the presence of quadratic curvature in the yield data, as opposed to a leveling, though it is expected that a linear trend followed by leveling will generally also project onto a quadratic decline in yields.

Finally, the interpretation that we favor is that the distinction in results arises from differences in methodology. Hafner (2003) determined that countries showed declining yield growth if a quadratic model provides "a better fit over [a] linear model" in the least squares sense. This method is biased toward accepting the quadratic model because at worst, that model will reduce to the linear model when the quadratic term is set to zero, but will often be able to fit the data better by having a non-zero quadratic term. Our study tests for a statistical distinction between the fits of two different models, accounting for autocorrelation and the differing degrees of freedom between the two models, and seeks to find not only which function provides for the better fit, but whether the rising-plateau model provides for a fit that is better in a statistically significant sense.

The study by Brisson et al. (2010) differed from Hafner (2003) by using the models that we adopt, statistically testing whether the leveling yield model is a better fit, and accounting for the difference in degrees of freedom between the two models. However, this test assumes that all deviations in yield from the trend are independently drawn from the same normal distribution, leading to an upwardly biased estimate of the number of degrees of freedom in the data, a related positive bias in the magnitude of the F-statistic, and a related bias toward a small p-value.

3 Classifying data as homoskedastic or heteroskedastic

Simple linear regression is appropriate when variance is constant, or homoskedastic. However, heteroskedasticity appears common in wheat yield data, with variance increasing toward the present. Such behavior is intuitive as many processes will influence yield in proportion to its mean value, which also generally increases with time. In order to test a data set for heteroskedasticity, we fit a simple linear regression to the squared residuals, e_i^2 (see Eq. 1 of the main article). The values of the e_i are themselves temporally autocorrelated (as described in supplemental materials section 1) and we average consecutive seven-year intervals to reduce this dependence. More sophisticated techniques are possible, but this approach appears adequate and prevents building up a complicated hierarchy of methods. If simple linear regression of these data points gives a fitted slope that does not include a slope of zero within its 95% confidence interval, then we classify the time-series as heteroskedastic, and otherwise treat it as homoskedastic.

For yield timeseries with heteroskedastic noise, we assume that variance scales linearly with time, t_i, and divide Eq. 1 through by $\sqrt{t_i}$, such that linear regression minimizes the homoskedastic term $\sum_{i=1}^{n} \frac{e_i^2}{t_i}$ (Judge et al. 2001). Note that for purposes of scaling variance, we initiate t at 1 and increment the value by one per year. The linear and rising-plateau models are fit to each yield data set by applying analytically derived estimators for the maximum-likelihood for the slope and intercept. In the case of the rising-plateau model, these values are determined as a function of t_p, the point of inflection, and we take the value of t_p that leads to the smallest sum of squared residuals. Depending on whether the data set has constant or increasing variance, $\sum_{i=1}^{n} e_i^2$ or $\sum_{i=1}^{n} \frac{e_i^2}{t_i}$ is being minimized in fitting the model, respectively. Accordingly, δJ is the difference in the minimized quantities, resulting in generally smaller values of δJ for data sets with increasing variance.

4 Singular value decomposition of wheat yields

Wheat yields are spatially correlated, and we use singular value decomposition to identify broad regions that vary together. In particular, singular value decomposition is performed on a matrix of wheat yields with rows corresponding to years and columns corresponding to counties. The first temporal mode extracted by the singular value decomposition of raw wheat yield would essentially give the mean time series of yield. In order to discern regions with distinct yield variations we subtract the mean yield across all counties on a year-by-year basis. After this subtraction, the singular value decomposition indicates regional departures from the mean trend.

4.1 U.S. county-level wheat yields

For U.S. county-level wheat yields, more than 70% of the variation is explained by the first mode (supplemental figure 3), and essentially all of the variation is explained by the first thirteen modes. Together, the geographic and temporal plots of the first mode show that wheat yields, relative to the U.S. mean, increase for a Western region; decrease in the Central region; and increase slowly for an Eastern region. These three distinct regional trends show that important regional variations are masked by taking the U.S. average yield. Interestingly, the wheat yield of the Western U.S. shows the longest fitted plateau, the greatest proportion of irrigated cereal land, and the highest wheat yields of the three regions (FAO 2012).

4.2 French departmental wheat yields

Results of a singular value decomposition of French departmental wheat yields also show regional divisions. More than 85% of the variation in departures from the mean French yield trend are explained by the first mode, and essentially all variation is explained by the first nine modes. Supplemental figure 4 shows the temporal and geographic influence of the first mode of the singular value decomposition, indicating a strong and smooth latitudinal gradient. The gradient of wheat yield in France is smoothly increasing northward, and the highest wheat yields in France are the farthest north, closer to 50°N. We choose to divide France into Northern and Southern regions at 46° N (supplemental figure 5). Northern France has a higher mean yield and more rapidly increasing yield than Southern France, though this discrepancy appears to have stopped growing at about the year 2000.

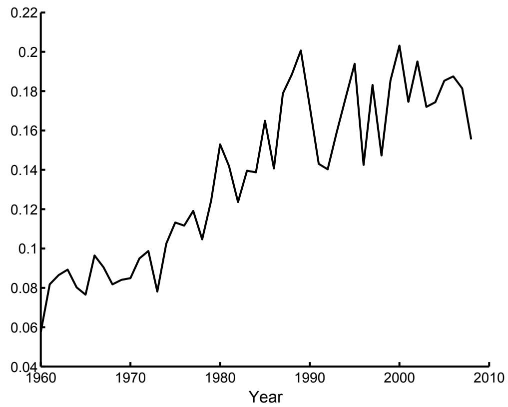

Figure 3: Temporal influence of first mode of the singular value decomposition of county-level U.S. wheat yields (first right-singular vector), with the annual mean yield removed.

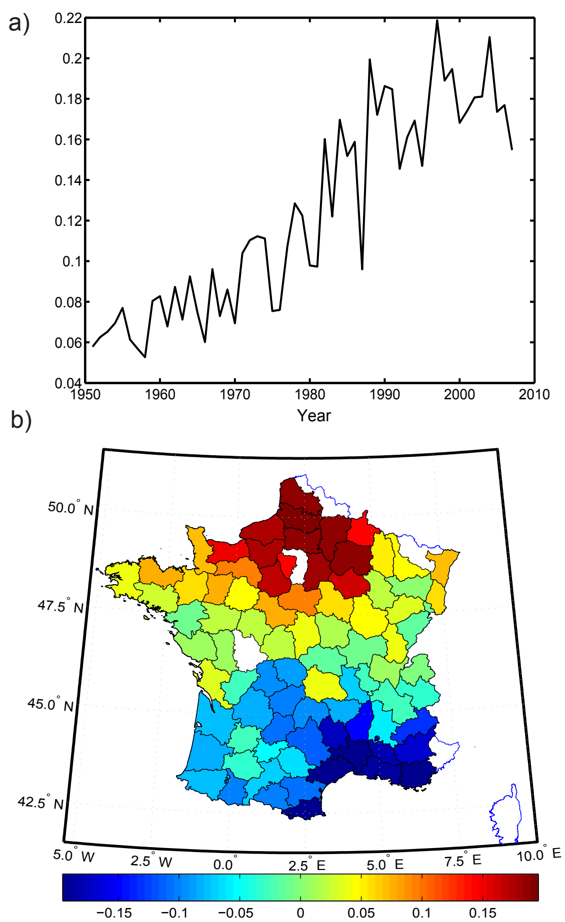

Figure 4: (a) Temporal (first right-singular eigenvector) and (b) geographic (first left-singular eigenvector) influence of first mode of the singular value decomposition of county-level French wheat yields, with the annual mean yield removed.

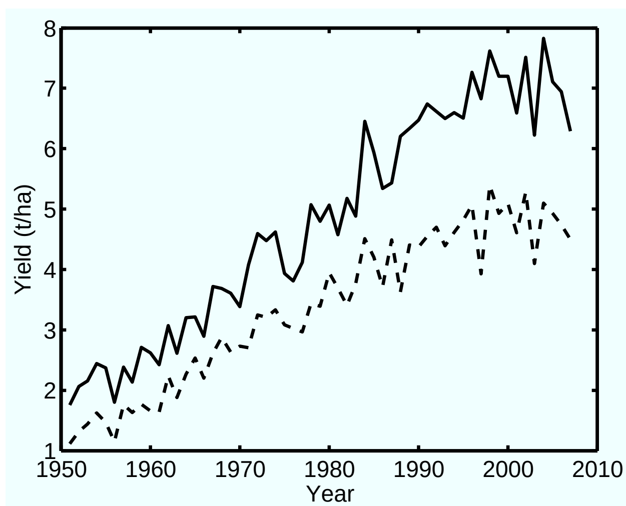

Figure 5: Regional mean wheat yields (solid line is the mean yield of the French region north of 46° N, dashed line is the French region south of 46° N)

5 Table of test values and model fits

Table 1: Values from wheat yield hypothesis test for all regions in sample. Columns from left to right are year of changepoint to plateau; observed δJ, δJ at the 95% confidence level; the power of the test, plateau yield (tons per hectare); p-value; and recent wheat yield averaged between year 2000 and most recent year of available data in tons per hectare.

Region	Plateau year	δJ	95% Level	Power	Plateau Yield	p-value	Recent Yield
Albania	1984	0.23	0.26	0.59	2.99	0.141	3.35
Austria	1990	4.32	3.76	0.58	5.05	0.002	5.03
Bangladesh	1984	2.28	2.36	0.61	2.00	0.068	2.08
China	2009	0.00	0.29	0.07	4.75	0.391	4.29
Colombia	1995	0.023	0.019	0.45	1.91	0.022	1.79
Cyprus	1995	2.13	2.56	0.36	2.02	0.130	1.78
Denmark	1995	3.66	3.38	0.56	7.19	0.033	7.23
Egypt	2004	1.08	2.30	0.26	6.32	0.155	6.37
Finland	1995	0.30	0.57	0.75	3.50	0.167	3.61
Germany	2010	1.78	1.11	0.78	7.41	0.010	7.42
Greece	1980	2.72	2.47	0.53	2.55	0.005	2.52
Hungary	1982	0.97	0.98	0.50	4.36	0.062	3.99
India	2001	0.32	0.18	0.76	2.73	0.004	2.72
Ireland	2000	6.90	5.06	0.50	8.85	0.012	8.82
Italy	1995	0.00	0.01	0.78	3.33	0.390	3.40
Japan	2003	0.492	0.485	0.24	3.81	0.048	3.81
Netherlands	1995	6.28	6.09	0.58	8.41	0.043	8.43
Norway	1982	5.65	5.05	0.56	4.35	0.010	4.32
Pakistan	2007	0.01	0.02	0.47	2.55	0.195	2.48
Poland	1988	0.077	0.087	0.59	3.62	0.092	3.78
Republic of Korea	1995	1.33	1.46	0.49	3.51	0.075	3.35
Romania	1980	0.25	0.23	0.56	2.69	0.005	2.60
Spain	2007	0.00	0.01	0.09	2.99	0.401	2.87
Sweden	1991	0.13	0.10	0.65	5.83	0.010	5.92
Switzerland	1991	6.58	5.83	0.53	5.91	0.011	5.82
Syrian Arab Republic	2004	0.01	0.06	0.16	2.28	0.422	2.39
Turkey	2000	0.17	0.49	0.66	2.14	0.256	2.25
United Kingdom	1997	4.18	3.75	0.58	7.78	0.034	7.80
Zambia	2001	3.81	3.21	0.48	6.11	0.028	6.12

Region	Plateau year	δJ	95% Level	Power	Plateau Yield	p-value	Recent Yield
Western U.S.	1993	1.86	1.68	0.60	4.39	0.024	4.42
Central U.S.	2003	0.00	0.01	0.11	2.60	0.416	2.62
Northern France	1997	3.87	1.50	0.72	7.03	0.000	7.05
Southern France	1995	0.06	0.03	0.68	4.81	0.000	4.87
France	1996	6.13	4.55	0.56	7.01	0.002	6.96
Algeria	—	0.00	—	—	—	—	1.32
Argentina	—	0.00	—	—	—	—	2.46
Bolivia	—	0.00	—	—	—	—	1.10
Brazil	—	0.00	—	—	—	—	2.06
Canada	—	0.00	—	—	—	—	2.47
Chile	—	0.00	—	—	—	—	4.49
Lebanon	—	0.00	—	—	—	—	2.82
Nepal	—	0.00	—	—	—	—	2.02
New Zealand	—	0.00	—	—	—	—	7.60
Peru	—	0.00	—	—	—	—	1.35
Saudi Arabia	—	0.00	—	—	—	—	5.42
South Africa	—	0.00	—	—	—	—	2.59
United States	—	0.00	—	—	—	—	2.82
Uruguay	—	0.00	—	—	—	—	2.66
Tunisia	—	0.00	—	—	—	—	1.70
Eastern U.S.	—	0.00	—	—	—	—	4.06

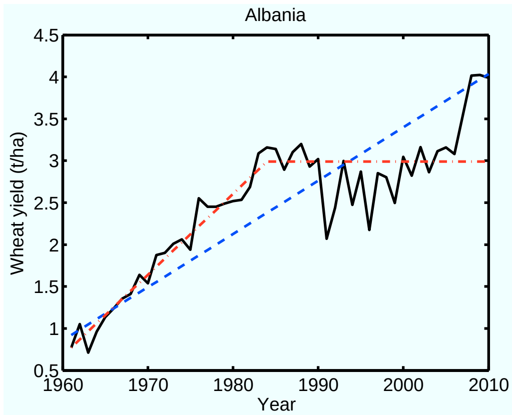

(a)

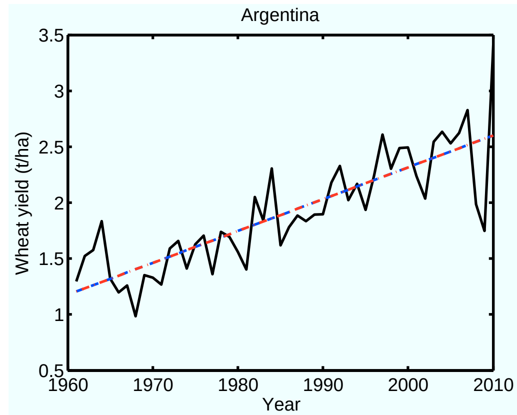

(b)

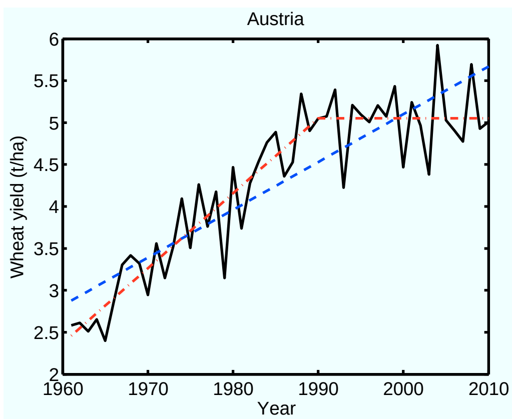

(c)

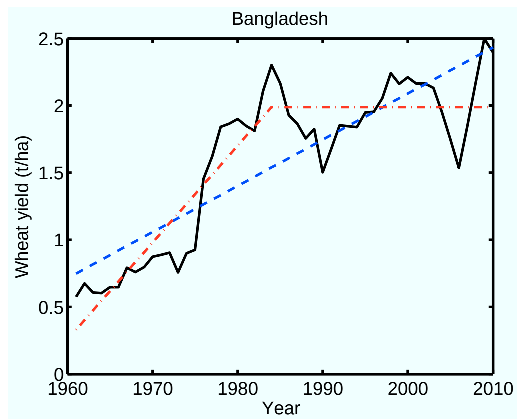

(d)

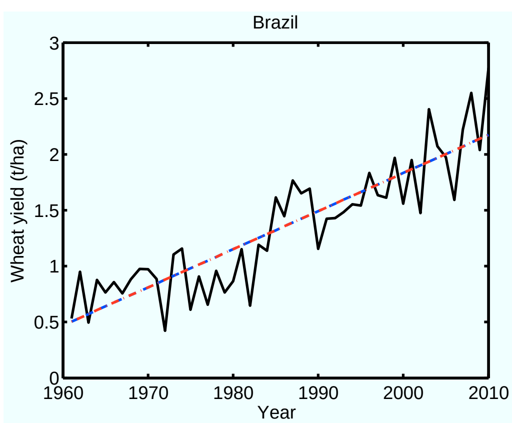

(e)

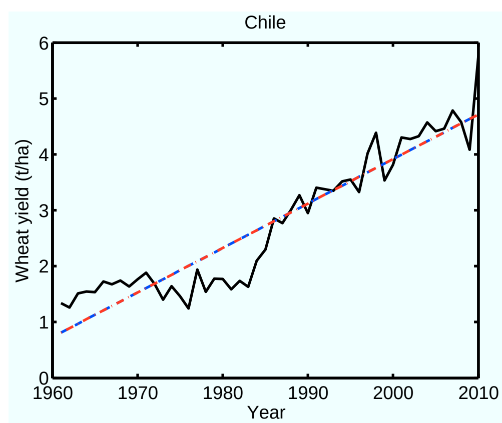

(f)

Figure 6: Wheat yield in tons per hectare (black solid line); rising-plateau model (red dot-dashed line); linear fit (blue dashed line)

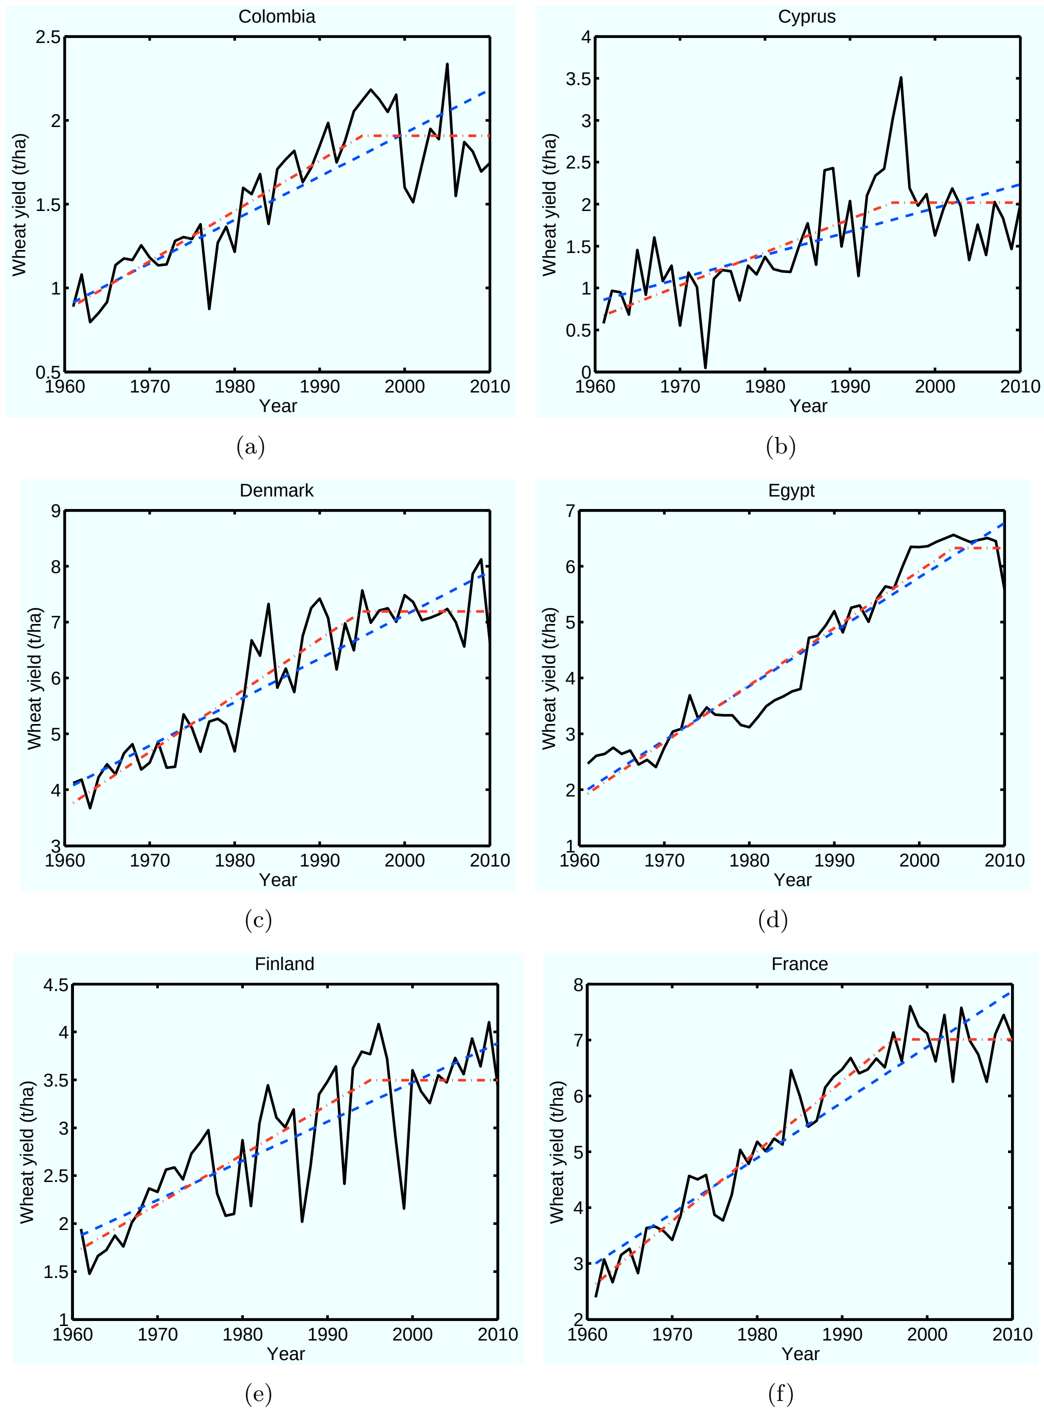

Figure 7: Wheat yield in tons per hectare (black solid line); rising-plateau model (red dot-dashed line); linear fit (blue dashed line)

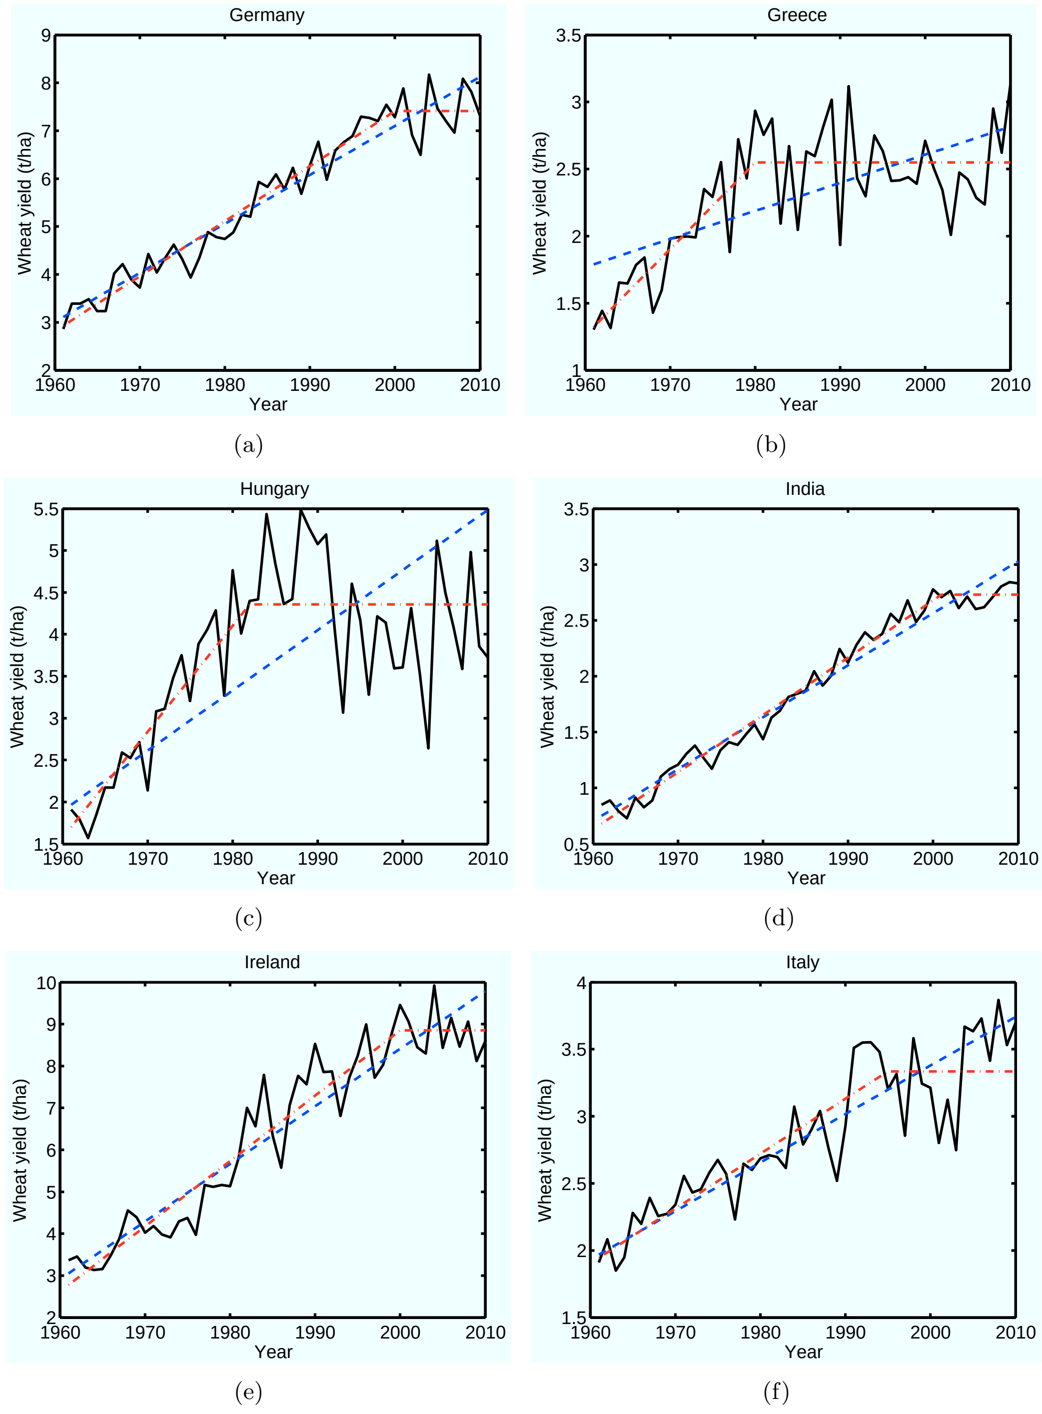

Figure 8: Wheat yield in tons per hectare (black solid line); rising-plateau model (red dot-dashed line); linear fit (blue dashed line)

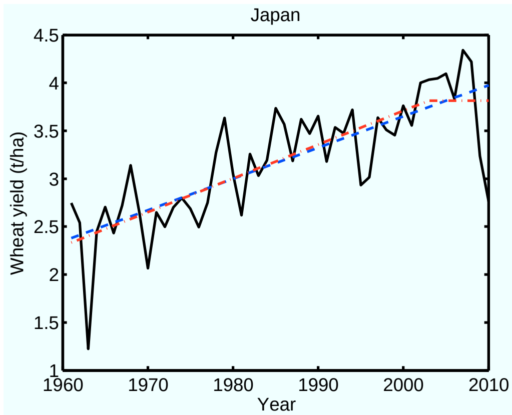

(a)

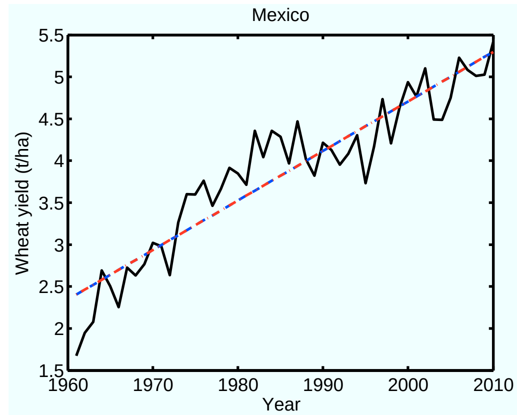

(b)

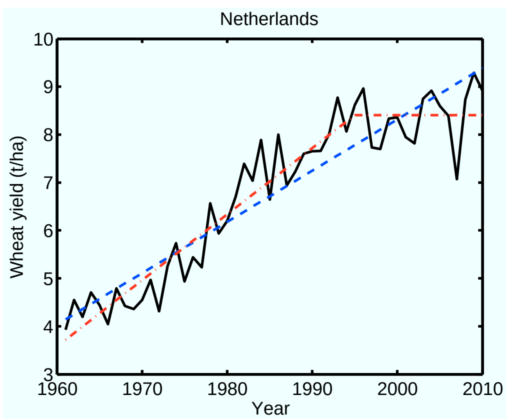

(c)

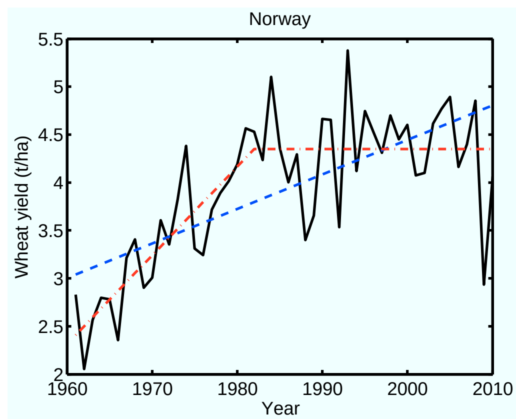

(d)

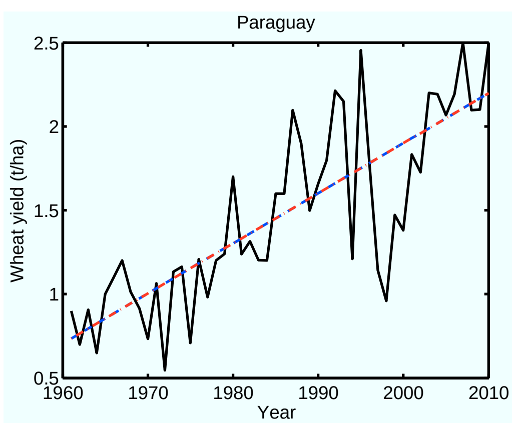

(e)

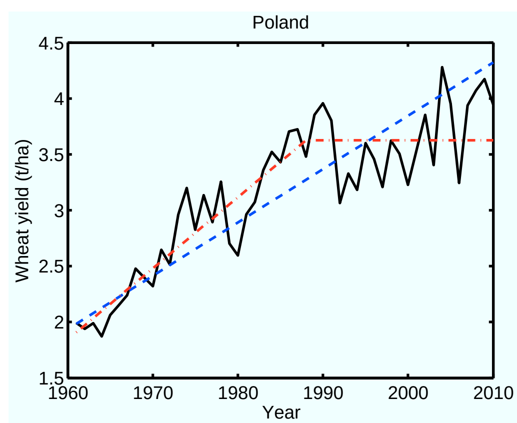

(f)

Figure 9: Wheat yield in tons per hectare (black solid line); rising-plateau model (red dot-dashed line); linear fit (blue dashed line)

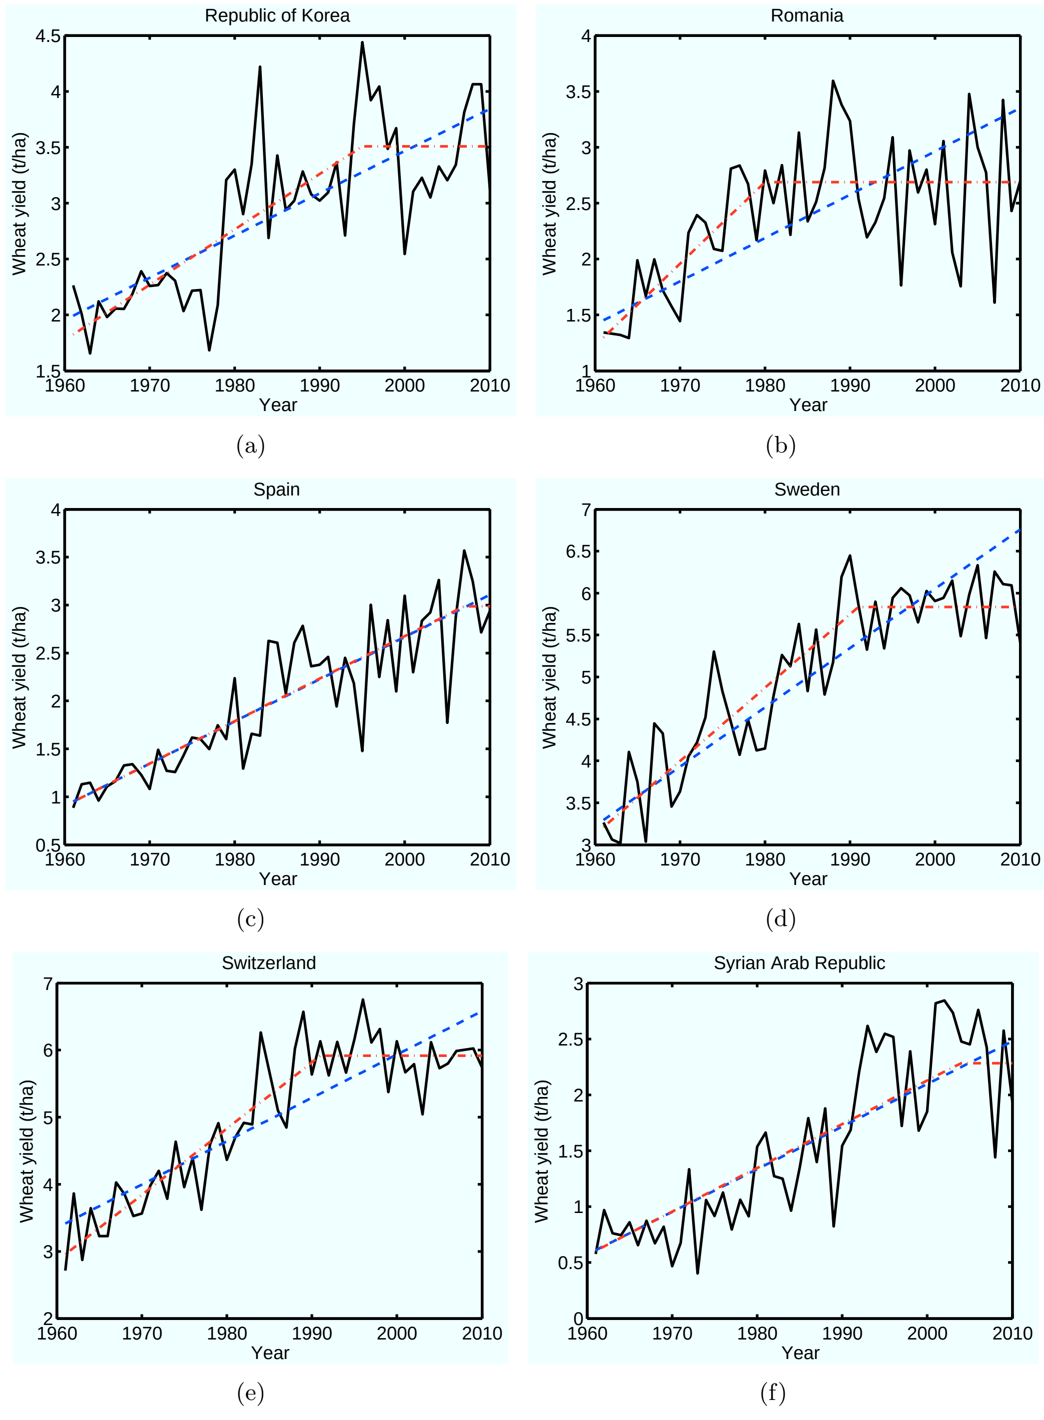

Figure 10: Wheat yield in tons per hectare (black solid line); rising-plateau model (red dot-dashed line); linear fit (blue dashed line)

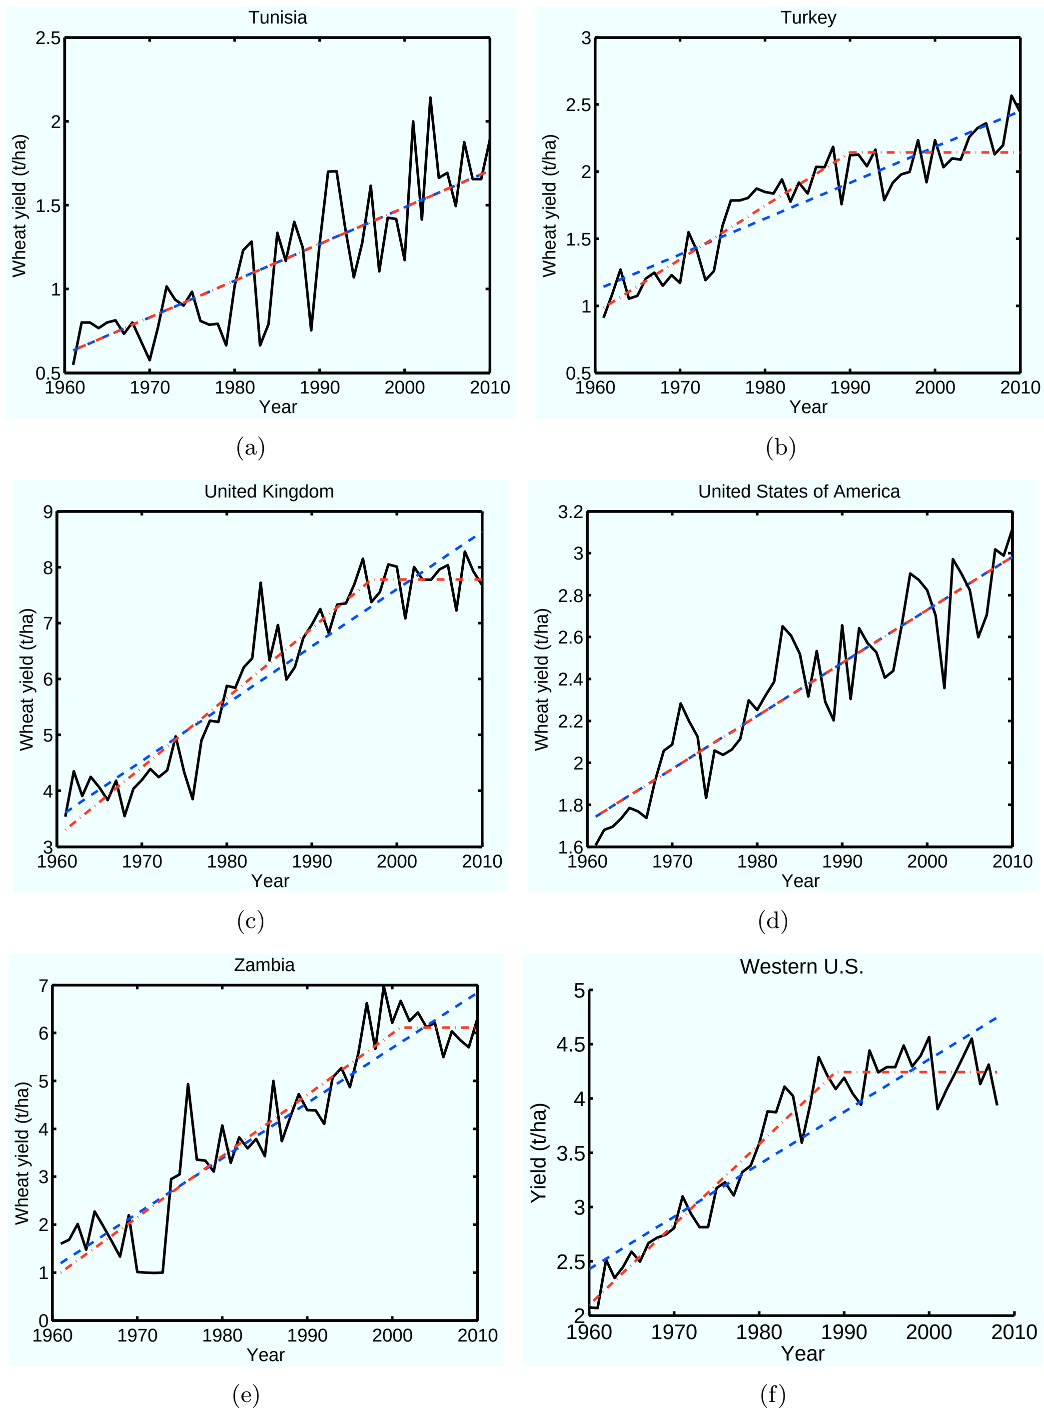

Figure 11: Wheat yield in tons per hectare (black solid line); rising-plateau model (red dot-dashed line); linear fit (blue dashed line)

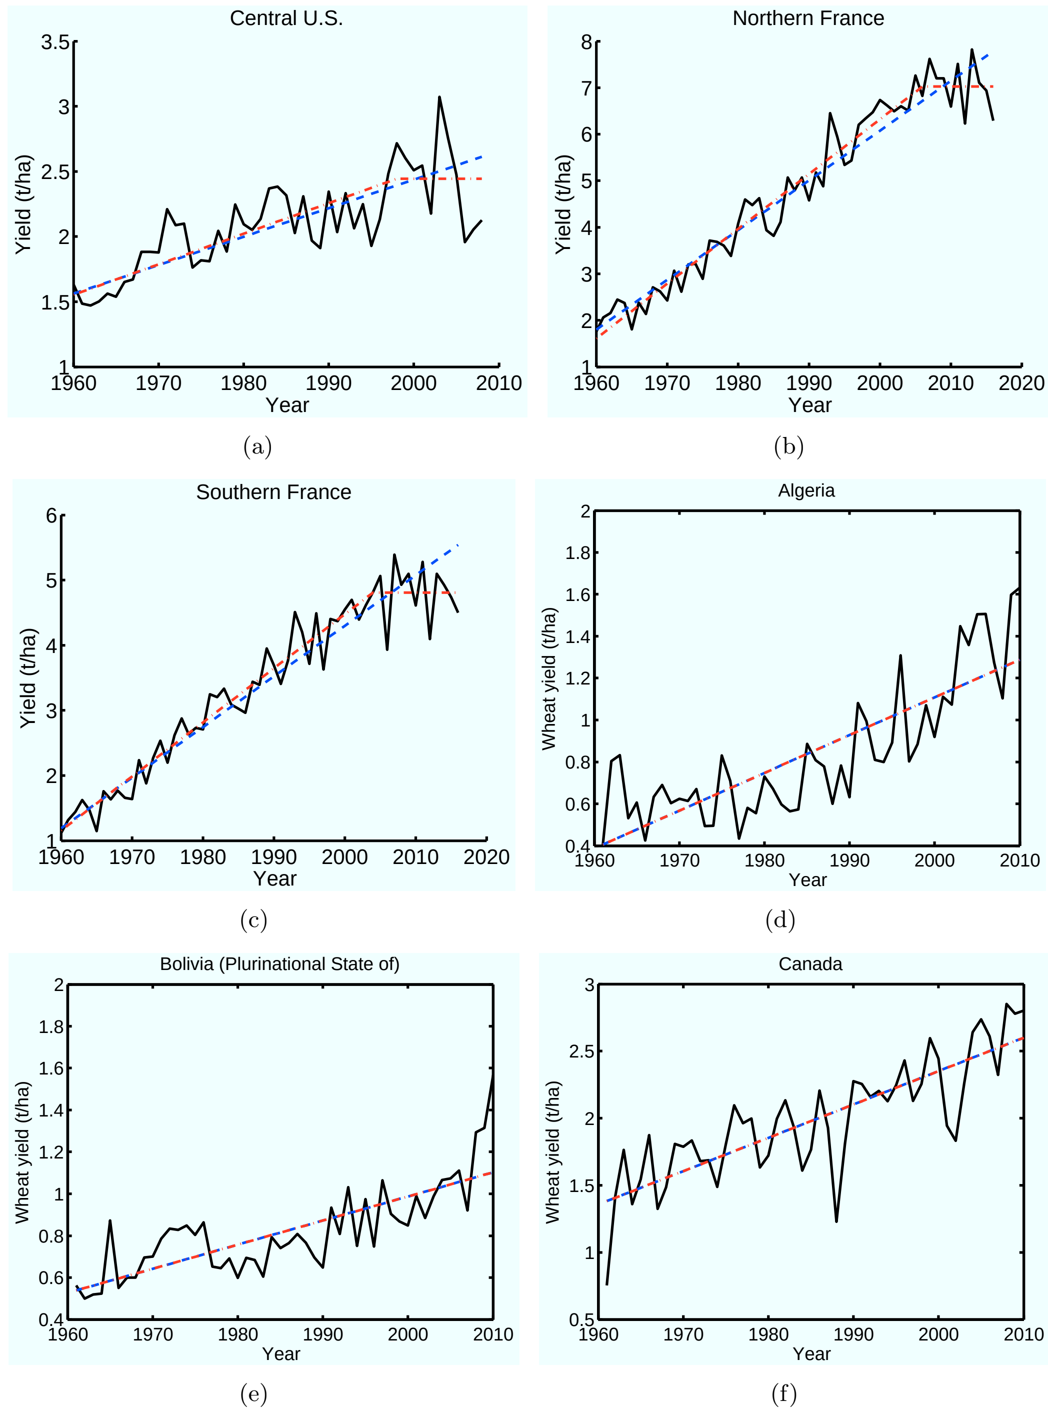

Figure 12: Wheat yield in tons per hectare (black solid line); rising-plateau model (red dot-dashed line); linear fit (blue dashed line)

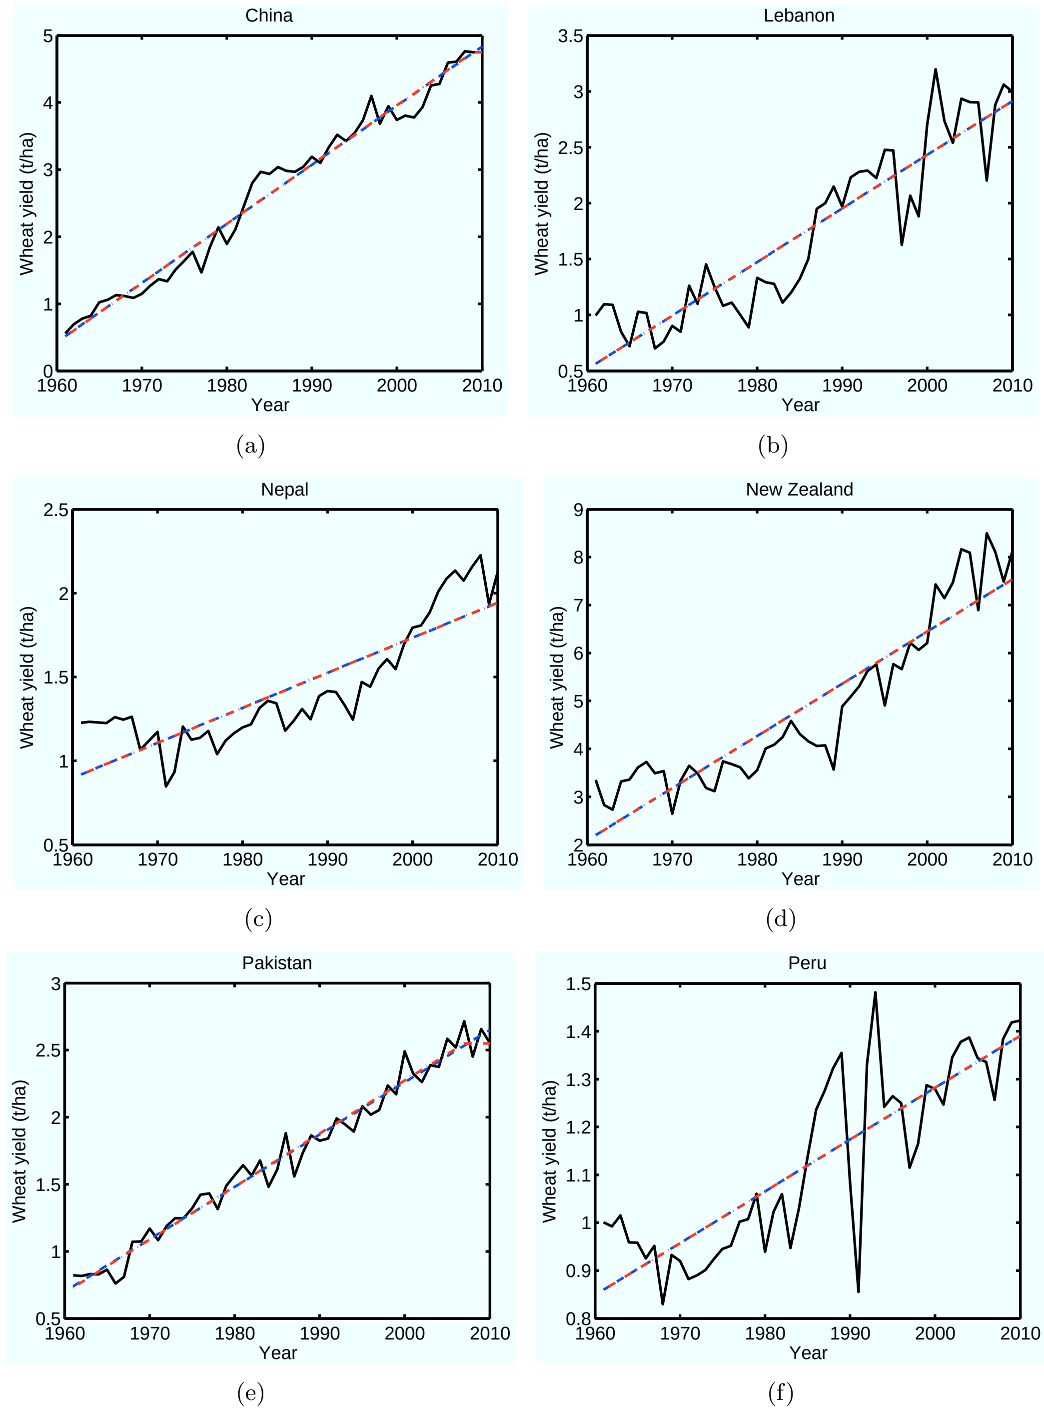

Figure 13: Wheat yield in tons per hectare (black solid line); rising-plateau model (red dot-dashed line); linear fit (blue dashed line)

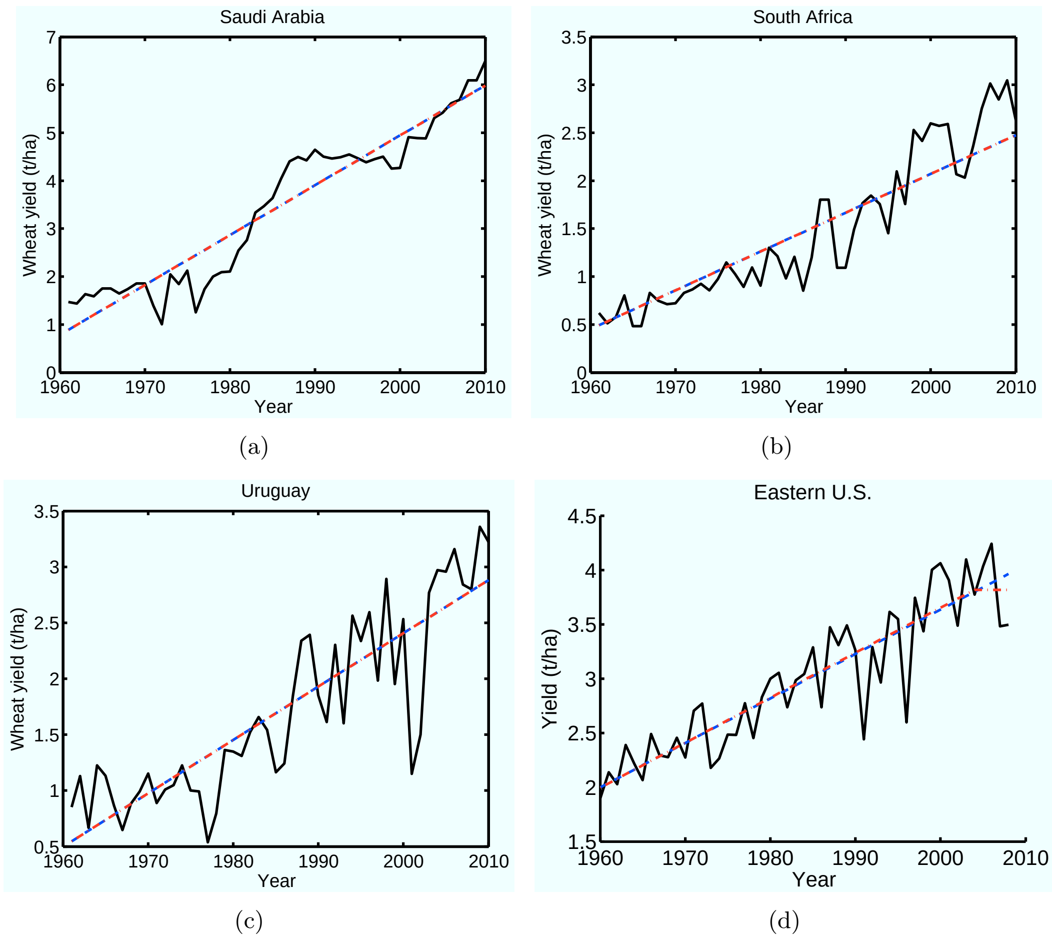

Figure 14: Wheat yield in tons per hectare (black solid line); rising-plateau model (red dot-dashed line); linear fit (blue dashed line)

6 Excluded major wheat producers

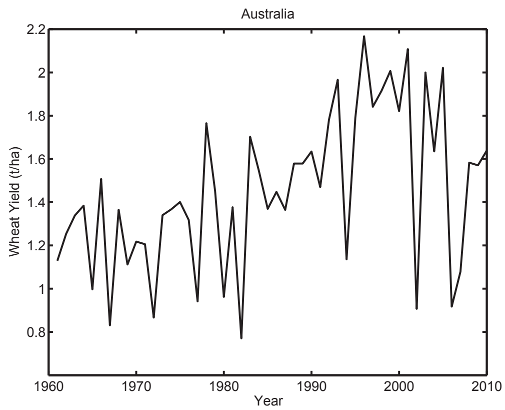

(a)

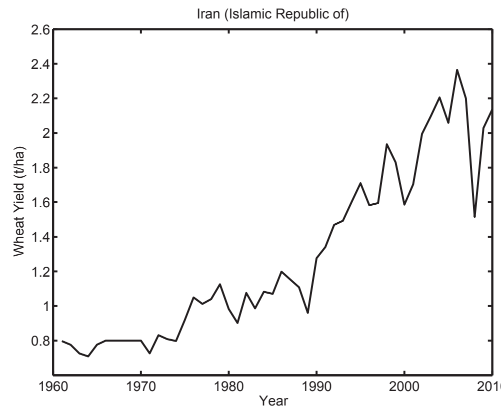

(b)

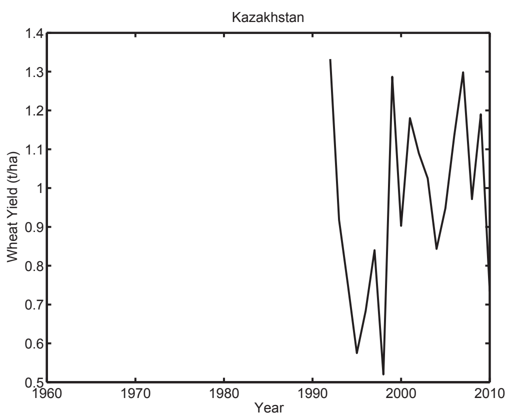

(c)

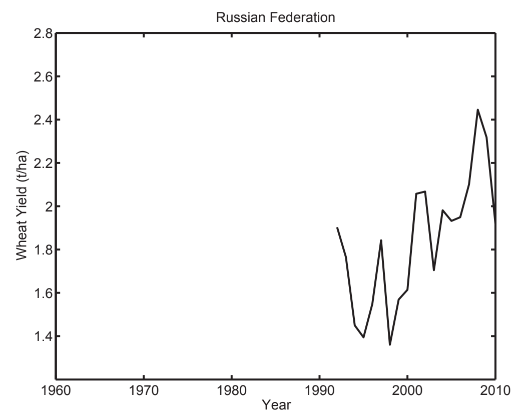

(d)

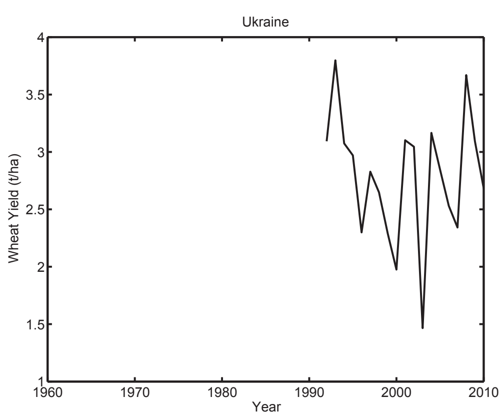

(e)

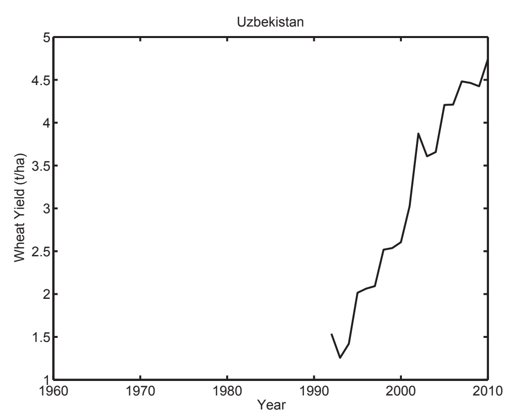

(f)

Figure 15: Yield timeseries of major wheat producers (within top twenty wheat producers in 2007) excluded for unreliable reporting or poor model fits.

References

Food and Agriculture Organization of the United Nations (FAO) 2012 *FAO Crop Production and Trade Statistics* . Available from: http://faostat.fao.org/site/339/default.aspx.

Brisson N *et al* 2010 *Field Crops Research* **119**, 201–212

Hafner S 2003 *Agric. Ecosyst. Environ.* **97**, 275–283

Judge G G, Hill R C & Griffiths W E 2001 *Undergraduate Econometrics, Second Edition* (Hoboken: John Wiley & Sons, Inc.)