The Spatial Structure of Endowments, Trade, and Inequality: Evidence from the Global Climate*

Jonathan I. Dingel
Chicago Booth & NBER

Solomon Hsiang
UC Berkeley & NBER

Kyle C. Meng
UC Santa Barbara & NBER

September 2017
PRELIMINARY AND INCOMPLETE

Abstract

This paper shows that the spatial structure of endowments influences global inequality through trade. In standard trade models, an increase in the spatial correlation of productivities amplifies welfare dispersion by increasing the correlation between productivity and the gains from trade. To empirically examine this prediction, we study how global agricultural trade responds to exogenous changes in the spatial correlation of agricultural productivity driven by the El Niño-Southern Oscillation, a naturally occurring climatic phenomenon. As predicted, higher spatial correlation in cereal yields increases the correlation between productivity and the gains from trade. In a forecasting application, climate-change projections for 2099 that incorporate this general-equilibrium effect exhibit 50% greater global welfare inequality than partial-equilibrium projections, with higher welfare losses in most African, South American, and many Asian countries.

*We thank Costas Arkolakis, Mark Cane, Chris Costello, Pedro DiNezio, Matt Kahn, Ben Olken, Esteban Rossi-Hansberg, and Veronica Rappaport for valuable discussions and suggestions. Thanks to seminar participants at numerous institutions and conferences for helpful comments. Kevin Dano provided excellent research assistance. Dingel thanks the Kathryn and Grant Swick Faculty Research Fund at the University of Chicago Booth School of Business for financial support. jdingel@chicagobooth.edu, shsiang@berkeley.edu, kmeng@bren.ucsb.edu.

1 Introduction

Natural endowments shape economic activity. Past research has linked local economic outcomes to local endowments. However, natural endowments are rarely confined to specific economic or political jurisdictions. Mineral deposits cross national borders, coastlines stretch across continents, and climate zones encompass much of a hemisphere. This paper addresses how these broad spatial structures in endowments influence the distribution of local economic outcomes.

Why might the spatial structure of endowments matter? For many economic phenomena, the object of interest is a change in the distribution of productivities among economies that trade. For example, consider anthropogenic climate change. As the planet warms, the spatial pattern of productivities will shift. A full account of climate change's impacts must characterize both local productivity effects and how the new spatial structure of productivities alters patterns of economic exchange in general equilibrium. In other words, it is important not only that Ethiopia and Kenya will be warmer but that both will be warmer at the same time. Analogous concerns arise in settings in which the spatial structure of productivities is altered due to the relocation of existing endowments (e.g. migrating human capital or wildlife stocks), the discovery of new uses for them (e.g. solar and wind resources), or the discovery of new endowments (e.g. fossil fuel deposits). In each case, partial-equilibrium forecasts of economic outcomes that are determined in general equilibrium and depend on the spatial structure of productivities may be biased.

In this paper, we examine how the *spatial correlation* of endowments influences the distribution of economic outcomes through their interaction via the global trading system. We show that in standard models of international trade with distance-related trade costs, dispersion in welfare across locations depends on the spatial correlation of productivity. Greater spatial correlation increases welfare inequality by increasing the correlation between productivity and the gains from trade. When endowments are spatially uncorrelated, high- and low-productivity locations are evenly distributed across space. So too are the gains from trade as all locations are similarly near both more and less productive trading partners. Conversely, when endowments are spatially correlated, there are clusters of high- and low-productivity locations. High-productivity locations gain more from trade because they are near high-productivity trading partners, while low-productivity locations gain less from trade because they are surrounded by similarly low-productivity locations. Thus, simply rearranging productivities such that they have higher spatial correlation increases inequality of welfare throughout the trading system.

To empirically investigate this general-equilibrium prediction, we study a natural experiment over the course of a half century of global agricultural trade. Causal inference is challenging when the prediction of interest concerns a system-wide feature of a trading network. In our context, we cannot experimentally rearrange the planet's natural endowments. Our solution is to exploit an annual climatic phenomenon that approximates the ideal experiment in the agricultural sector by exogenously altering the spatial correlation of agricultural productivities. In particular, historical annual variation in the El Niño-Southern Oscillation (ENSO) abruptly causes a spatial reshuffling of temperature, and thus agricultural productivity, around the planet, creating large, spatially con-

tiguous regions with similar productivities. We use a sufficient-statistics approach to infer the gains from trade from observed expenditure shares (Arkolakis, Costinot and Rodríguez-Clare, 2012). As predicted, we find that greater spatial correlation of productivity increases the correlation between productivity and the gains from trade. A one-standard-deviation increase in the spatial correlation of agricultural productivities increases the dispersion of welfare attributable to agricultural consumption by 4%.

We show how to incorporate this general-equilibrium mechanism when forecasting the economic impacts of anthropogenic climate change. Climate change is expected to induce a large spatial restructuring of productivities across the planet, resulting in both local and general-equilibrium effects. However, most projections of future climate impacts based on historic variation have omitted general-equilibrium considerations by predicting local outcomes as a function of only local conditions. More sophisticated forecasts can also incorporate system-wide variables that govern local outcomes and depend on the global climate. In our application, including the change in the spatial structure of productivities has an important consequence for our forecast of global inequality. Specifically, climate impact projections that incorporate the projected change in the spatial correlation of agricultural productivity due to climate change predict a 50% higher increase in the inequality of welfare by the end of the 21st century. Projections that omit this general-equilibrium effect considerably understate the climate-driven welfare losses for most countries in Africa and South America, and many in Asia.

This paper relates to the ongoing dialogue, stretching back centuries, as to the possible contribution that environmental and geographic endowments may have in determining the long-run well-being of societies (Sachs and Warner, 1997; Easterly and Levine, 2003). Persistent correlations between geographic endowments and economic outcomes are often remarkable (Hornbeck, 2012), with prior work articulating numerous potential channels of influence, such as differing availability of local inputs leading to differing local productivities (Nordhaus, 2006; Bleakley, 2007) as well as differing local conditions producing differing local institutions (Nunn and Puga, 2012) or differing patterns of local investment (Burchfield et al., 2006). Our analysis proceeds to explore the influence of *non-local* geographic endowments, occurring at both neighboring and distant locations, in the determination of local outcomes – an effect that depends critically on the spatial relationship between local and non-local endowments. In short, we focus on the general-equilibrium role of geography.

Our paper thus belongs to a large literature in international trade and economic geography studying local economic consequences of the geographic distribution of economic activity. Prior research has demonstrated that better market access, which is the sum of other locations' economic activity discounted by bilateral trade costs, is associated with higher per capita income and higher land values (Head and Mayer, 2004; Redding and Venables, 2004; Allen and Arkolakis, 2014; Donaldson and Hornbeck, 2016). These studies infer these benefits from observed changes in bilateral trade costs or cross-sectional variation in access to endogenous economic activity. Prior examinations of heterogeneity in the gains from trade have overwhelmingly focused on the role of country

size rather than geography (e.g. Costinot and Rodríguez-Clare 2014). By contrast, we exploit exogenous variation in the global distribution of agricultural productivities to study how the distribution of welfare outcomes depends on the spatial arrangement of economic primitives. The most closely related study of global agricultural trade is Costinot, Donaldson and Smith (2016), who examine the trade and welfare consequences of predicted shifts in comparative advantage across different crops due to climate change. Since their theoretical framework incorporates a standard gravity model of trade flows, their results in part reflect the spatial structure of productivity shocks embedded in the agronomic forecasts they employ. We contribute to the understanding of these predicted changes by empirically examining the role of spatial correlation using historical variation in agricultural productivities.

Finally, this paper speaks to the growing empirical literature examining how anthropogenic climate change may affect inequality. Prior research employing estimates from historical local temperature variation has forecast considerable dispersion in various economic outcomes under projected climate change, both across countries (Dell, Jones and Olken, 2012; Burke, Hsiang and Miguel, 2015) and within countries (Burgess et al., 2013; Houser et al., 2015). This paper shows that traditional partial-equilibrium empirical methods for projecting climate impacts may be understating changes in welfare inequality by omitting general-equilibrium effects related to the influence of climate change on the spatial structure of productivities. Thus, our paper brings the reduced-form climate impacts literature conceptually closer to recent macroeconomic analyses exploring the spatial distribution of economic activity under climate change (Brock, Engström and Xepapadeas, 2014; Desmet and Rossi-Hansberg, 2015; Krusell and Smith, 2016).

The rest of the paper is structured as follows. Section 2 provides our theoretical framework and prediction about the spatial correlation of productivity. Section 3 introduces the El Niño-Southern Oscillation and shows how it affects the spatial correlation of agricultural productivity across the world. Section 4 uses ENSO-driven variation in agricultural productivity to empirically validate our prediction. Section 5 considers implications for projecting global impacts under anthropogenic climate change. Section 6 concludes.

2 Theoretical framework

This section introduces a theoretical framework that shows how the spatial correlation of productivities affects welfare inequality and guides our empirical investigation of this prediction.

In any trade equilibrium, a country's welfare can be stated as the sum of its welfare under autarky and its gains from trade. A country's welfare under autarky depends only on its own productivity. Its gains from trade depend on the entire distribution of productivities across the trading network. The variance of welfare across countries is the variance of this sum. It therefore depends on not only the variance in productivity and gains from trade but also the covariance between these two components.

We investigate how the spatial correlation of productivities influences the covariance between a

country's productivity and its gains from trade. This requires an observable outcome that identifies the gains from trade. Across a broad class of models, a country's gains from trade are revealed by the share of its expenditure devoted to its own output. The less a country spends on its own output, the larger its gains from trade. In autarky, all expenditure is on its own output. In the trade equilibrium, this expenditure share, when combined with the "trade elasticity" governing how consumers substitute across consumption sources, summarizes the welfare gain from exchange with other locations (Arkolakis, Costinot and Rodríguez-Clare, 2012).

Section 2.1 establishes that within this class of trade models our object of interest is the covariance between a country's productivity and its domestic share of expenditure. Section 2.2 illustrates how this covariance depends on the spatial correlation of the productivity distribution and shows how to identify this *ceteris paribus* prediction in empirical settings. Section 2.3 describes conditions under which examining one sector in isolation is informative about welfare dispersion in a multi-sector world. Details and derivations are available in Appendix A.

2.1 Sufficient statistics for welfare dispersion

We consider a general economic environment within the class of models characterized by Arkolakis, Costinot and Rodríguez-Clare (2012), in which the gains from trade can be inferred from the domestic share of expenditure. We assume perfect competition in the main text, while Appendix A.1 covers the case of monopolistic competition. The world economy is comprised of $j = 1, \ldots, N$ countries.

Preferences. Individuals in country j have Dixit-Stiglitz preferences with a constant elasticity of substitution σ over goods indexed by $\omega \in \Omega$. The accompanying price index is

$$P_j = \left(\int_{\omega \in \Omega} p_j(\omega)^{1-\sigma} d\omega \right)^{1/(1-\sigma)}.$$

Production. There is one factor of production, and each country j inelastically supplies L_j units of that factor, which earns wage w_j. A country's income is therefore $Y_j = w_j L_j$. The production technology exhibits constant returns to scale and is employed by perfectly competitive firms. The cost of producing in country j depends on productivity A_j. This parameter's microeconomic meaning is model-specific: A_j governs the cost of producing j's good in the Armington model and the location parameter of j's cost distribution for a continuum of goods in the Eaton and Kortum (2002) model.

Trade costs. There are iceberg trade costs, such that selling one unit of a good to j from i requires $\tau_{ij} \geq 1$ units, with $\tau_{ii} = 1$. By the no-arbitrage condition, $p_j(\omega) \leq \tau_{ij} p_i(\omega)$.

Gravity equation. Denote sales from i to j by X_{ij} and j's total expenditure by $X_j \equiv \sum_{i=1}^{N} X_{ij}$.

The share of expenditure by j on goods from i takes the form of a gravity equation:[1]

$$\lambda_{ij} = \frac{X_{ij}}{X_j} = \frac{\chi_i \left(\tau_{ij} w_i\right)^{-\epsilon}}{\sum_{l=1}^{N} \chi_l \left(\tau_{lj} w_l\right)^{-\epsilon}} = \frac{\chi_i \left(\tau_{ij} w_i\right)^{-\epsilon}}{\Phi_j},$$

where χ_i is a function of A_i and other structural parameters that are not trade costs, ϵ is the "trade elasticity", and $\Phi_j \equiv \sum_{l=1}^{N} \chi_l \left(\tau_{lj} w_l\right)^{-\epsilon}$ is the "inward multilateral resistance" term (Head and Mayer, 2014). Φ_j is a (decreasing) transformation of j's price index that summarizes consumers' access to goods from every source.

Equilibrium. In equilibrium, labor-market clearing, goods-market clearing, and budget constraints are satisfied such that total income $Y_i = w_i L_i$ equals total expenditure X_i. Thus, an equilibrium is a set of incomes $\{Y_i\}_{i=1}^{N}$ such that

$$Y_i = \sum_{j=1}^{N} \lambda_{ij} Y_j.$$

In this environment, the results of Arkolakis, Costinot and Rodríguez-Clare (2012) imply that real consumption per capita is

$$\ln\left(C_i/L_i\right) = \ln A_i + \gamma - \frac{1}{\epsilon} \ln \lambda_{ii}, \tag{1}$$

where γ is a constant determined by structural parameters that are not productivity. The former term, $\ln A_i + \gamma$, is per capita welfare in autarky.

In the absence of trade, a country's welfare is independent of other countries' conditions and depends only on its own productivity. The latter term, $-\frac{1}{\epsilon} \ln \lambda_{ii}$, is a sufficient statistic for the gains from trade relative to autarky. It is a country's expenditure share on its own goods, mediated by the trade elasticity ϵ that governs how bilateral expenditures respond to changes in bilateral trade costs. With this standard equilibrium expression for welfare in hand, we can consider how dispersion in $\ln\left(C/L\right)$ across countries depends on the spatial distribution of productivities.

From equation (1), variance in welfare across countries is governed by the variance of productivity, the covariance of exogenous productivities and endogenous domestic shares of expenditure, and the variance of those shares.

$$var\left(\ln\left(C_i/L_i\right)\right) = var\left(\ln A_i\right) - \frac{2}{\epsilon} cov\left(\ln A_i, \ln \lambda_{ii}\right) + \frac{1}{\epsilon^2} var\left(\ln \lambda_{ii}\right) \tag{2}$$

To examine the role of spatial correlation, consider two productivity distributions – a correlated state c and an uncorrelated state u – in which the unconditional variance in productivities is identical, $var(\ln A_{i,c}) = var(\ln A_{i,u})$. Under this assumption, the difference in welfare dispersion

[1] Appendix C.1 shows that the gravity model fits trade in cereals well.

between the correlated and uncorrelated states is

$$var\left(\ln\left(C_{i,c}/L_i\right)\right) - var\left(\ln\left(C_{i,u}/L_i\right)\right) = -\frac{2}{\epsilon}\left[cov\left(\ln A_{i,c}, \ln\lambda_{ii,c}\right) - cov\left(\ln A_{i,u}, \ln\lambda_{ii,u}\right)\right] + \frac{1}{\epsilon^2}\left[var\left(\ln\lambda_{ii,c}\right) - var\left(\ln\lambda_{ii,u}\right)\right]. \quad (3)$$

The latter term should make only a second-order contribution to the difference in welfare dispersion, since $\frac{1}{\epsilon^2}$ is an order of magnitude smaller than $\frac{2}{\epsilon}$ for empirically relevant values of the trade elasticity.[2]

The first-order difference in welfare dispersion is governed by the covariance of the exogenous productivities and the endogenous domestic shares of expenditure. All else equal, this covariance will typically be positive: since a more productive country sells more to every consumer, a more productive country purchases more from itself.

Our concern here, however, is how this covariance changes with the degree of spatial correlation in productivities. We will estimate this relationship empirically, but we first illustrate why we expect that $cov\left(\ln A_{i,c}, \ln\lambda_{ii,c}\right) < cov\left(\ln A_{i,u}, \ln\lambda_{ii,u}\right)$ and thus that $var\left(\ln\left(C_{i,c}/L_i\right)\right) > var\left(\ln\left(C_{i,u}/L_i\right)\right)$.

2.2 Spatial correlation and the productivity-openness covariance

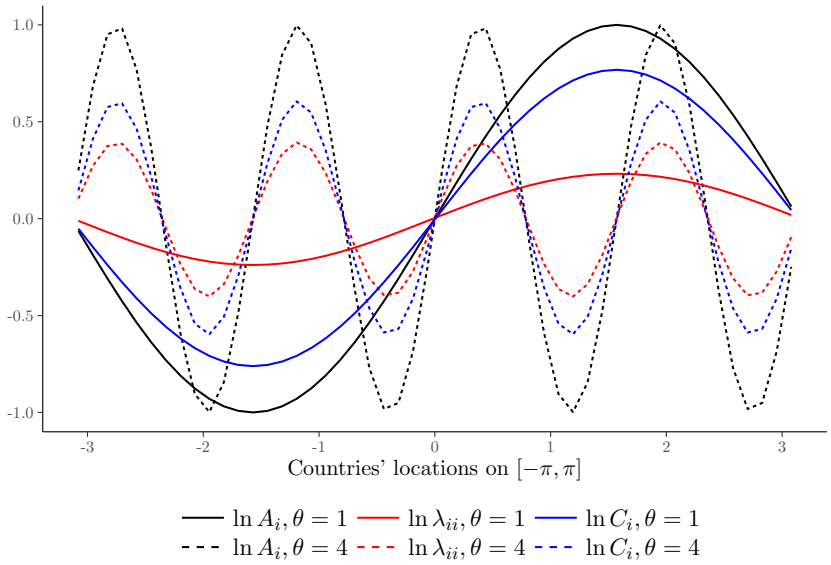
This section illustrates how the spatial correlation of productivity may influence the covariance of autarky welfare and gains from trade, and thus the variance of welfare, per equation (3). We start by examining the role of the spatial correlation of productivity when countries are perfectly symmetric except for productivity differences. This stylized example allows us to clearly illustrate our economic mechanism in isolation. We then examine how to identify this *ceteris paribus* prediction in asymmetric environments in which countries differ by other, potentially confounding, determinants of equilibrium trade flows. These more realistic examples inform how we empirically investigate our prediction.

2.2.1 Stylized example

Our ideal, stylized environment has productivity follow a sine wave over a one-dimensional space. There are N locations evenly spaced on the unit circle. These locations have equal population sizes, $L_i = 1\ \forall i$. Trade costs depend only on distance: the log trade cost between two locations is proportionate to the log distance between them. The trade elasticity is $\epsilon = 1$. Productivity $\ln A_i$ follows a sine-wave distribution, with an integer frequency of θ over the circle's circumference.

[2] Typical estimates of the aggregate trade elasticity are between 4 and 8. Caliendo and Parro (2015) estimate that the trade elasticity for agricultural goods is between 8 and 16. Provided that $var\left(\ln\lambda_{ii,c}\right) - var\left(\ln\lambda_{ii,u}\right)$ is the same order of magnitude or smaller than $cov\left(\ln A_{i,c}, \ln\lambda_{ii,c}\right) - cov\left(\ln A_{i,u}, \ln\lambda_{ii,u}\right)$, this means that the second term on the right side of equation (3) is an order of magnitude smaller than the the first term. Generally, $var\left(\ln\lambda_{ii}\right)$ is the same order of magnitude as $cov\left(\ln A_i, \ln\lambda_{ii}\right)$. For example, if trade costs are symmetric ($\tau_{ij} = \tau_{ji}$), Appendix A.1.3 shows that $var(\ln\lambda_{ii}) = \frac{\epsilon}{\epsilon+1}cov(\ln A_i, \ln\lambda_{ii}) - \frac{1+2\epsilon}{1+\epsilon}cov(\ln\Phi_i, \ln\lambda_{ii})$. Heuristically, the latter term is of smaller magnitude, since Φ_i is a price-index term that is a weighted sum of all other countries' prices, making $var(\ln\lambda_{ii})$ the same order of magnitude as $cov(\ln A_i, \ln\lambda_{ii})$. This means that the the second term on the right side of equation (3) is second-order.

Figure 1: Ideal environment: Circular geography with productivity sine wave



NOTES: This figure depicts an economy with a circular geography and a productivity distribution that follows a sine wave with frequency θ. There are $N = 50$ locations evenly spaced on the unit circle. Bilateral log trade costs are proportionate to the log length of the arc between two points on the circle. The (demeaned) distributions of productivities, equilibrium domestic shares of expenditure, and welfare are depicted for the cases of $\theta = 1$ and $\theta = 4$. See Appendix A.2 for parameterization details.

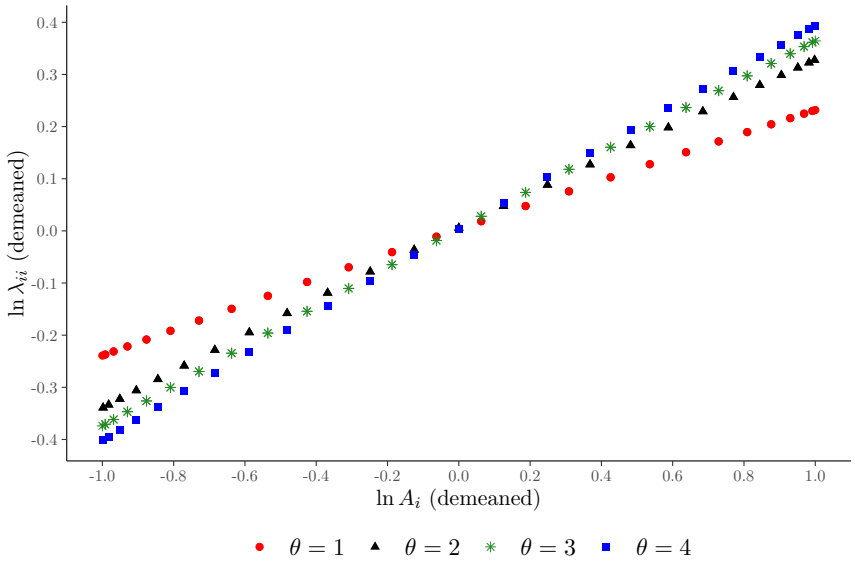
This functional form has two convenient properties. First, the spatial correlation of productivity is governed by the frequency θ: lower frequencies exhibit greater spatial correlation. Second, the mean, variance, skewness, and kurtosis of the productivity distribution are independent of the frequency. Thus, we can explore the effect of spatial correlation by varying θ alone.[3]

Figure 1 depicts the spatial distributions of productivities ($\ln A_i$), domestic shares of expenditure ($\ln \lambda_{ii}$), and welfare ($\ln C_i$) in this circular economy for the cases in which the sine wave has frequencies of $\theta = 1$ and $\theta = 4$. It is clear that the spatial correlation of productivity is greater in the $\theta = 1$ case, as location "zero" divides the circle into two contiguous regions with above-average and below-average productivity. In the $\theta = 4$ case, lower spatial correlation means that locations of opposing productivity are more closely situated. In the case of higher spatial correlation, $\theta = 1$, the equilibrium domestic share of expenditure series less closely follows the productivity series, as evident by the larger vertical gap between them. Thus, the smaller amplitude of the $\ln \lambda_{ii}$ series when productivity is more spatially correlated is accompanied by a lower value of $cov(\ln A_i, \ln \lambda_{ii})$. As a result, the amplitude of the welfare series is greater in the $\theta = 1$ case. Welfare dispersion is higher when productivity is more spatially correlated.[4]

[3] Details of the parameters underlying Figure 1 are in Appendix A.2. By Theorem 1 of Allen, Arkolakis and Takahashi (2017), the equilibrium solution depicted for each parameter value is unique.

[4] In this particular parameterization, average welfare is slightly lower when productivity is more spatially correlated. See Table A.1 for details. More generally, average welfare is not systematically related to the spatial correlation

Figure 2: One-sector sine-wave economy: $cov(\ln \lambda_{ii}, \ln A_i)$



NOTES: This figure depicts the $\ln \lambda_{ii}$–$\ln A_i$ relationship for an economy with a circular geography and a productivity distribution that follows a sine wave with frequency θ. Geographic locations and trade costs as in Figure 1. See Appendix A.2 for parameterization details.

Figure 2 depicts the expenditure-productivity relationship in our sine-wave example for more values of the sine-wave frequency, θ. The scatter plot reveals an almostly perfectly linear relationship between $\ln \lambda_{ii}$ and $\ln A_i$. The slope of this relationship, which is proportionate to $cov(\ln A_i, \ln \lambda_{ii})$, systematically varies with the frequency of the sine wave. At lower frequencies, when the productivities are more spatially correlated, a location's domestic share of expenditure is less responsive to its own productivity level. Intuitively, when productivities are more spatially correlated, an increase in a location's own productivity level is accompanied by higher productivity levels in its nearby trading partners, such that its gains from trade are larger than they would otherwise be. Thus, the spatial correlation of productivity governs the variance of welfare across countries through its effect on the covariance of $\ln \lambda_{ii}$ and $\ln A_i$.

2.2.2 Asymmetric environments

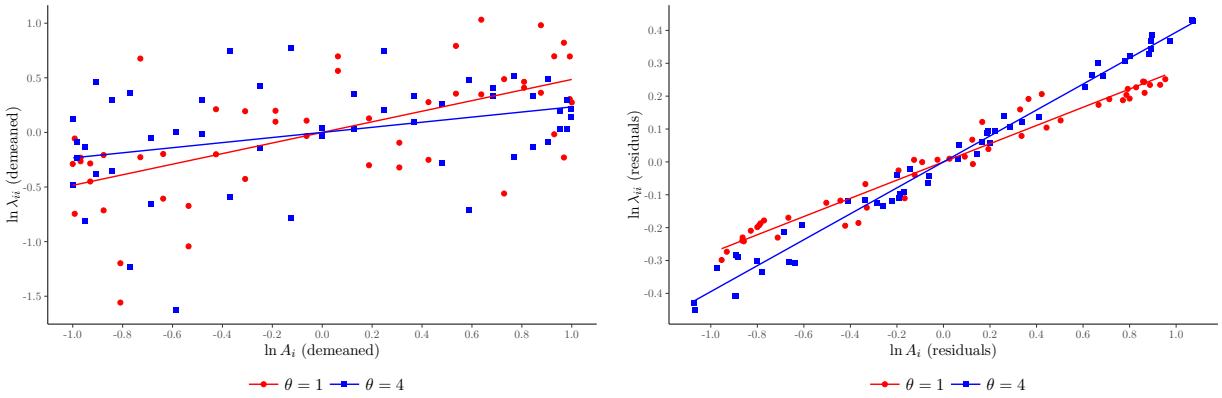
Our stylized example in Section 2.2.1 demonstrates the consequences of spatial correlation of productivity for global welfare inequality in an ideal environment that holds fixed all other economic elements. Compared to the ideal environment, any empirical setting is more complicated because (1) there are other determinants of equilibrium trade flows, (2) productivities do not follow a sine wave, and (3) countries are not evenly spaced on a circular geography. In this section, we discuss how to identify our *ceteris paribus* prediction about the impact of spatial correlation when such

of productivity.

asymmetries are present.

Economic characteristics other than productivity that influence domestic shares of expenditure complicate bivariate plots like Figure 2. A simple example is heterogeneity in country size L_i: all else equal, larger economies have a larger domestic share of expenditure. Variation in size orthogonal to productivity simply adds noise to the bivariate plot. However, variation in size correlated with productivity also introduces omitted variable bias. This can be empirically addressed by examining the covariance of the domestic share of expenditure and productivity conditional on size. More generally, any time-invariant country characteristics that influence the domestic share of expenditure and might be correlated with productivity can be absorbed by country fixed effects. We illustrate this in Figure 3, which depicts the relationship between $\ln \lambda_{ii}$ and $\ln A_i$ in an environment that features, like the previous section, a circular geography and sine-wave productivity, and, unlike the previous section, heterogeneous country sizes. In particular, country size $\ln L_i$ is positively correlated with productivity $\ln A_i$ in the $\theta = 1$ state. The left panel depicts the covariance of $\ln \lambda_{ii}$ and $\ln A_i$ for the frequencies $\theta = 1$ and $\theta = 4$. The right panel depicts these covariances conditional on country fixed effects. While our *ceteris paribus* prediction is not evident in the left panel due to omitted variable bias, the right panel shows that the covariance of $\ln \lambda_{ii}$ and $\ln A_i$ is lower when θ is lower, controlling for heterogeneous sizes.

Figure 3: Circular geography with heterogeneous sizes and productivity sine wave



Unconditional relationship

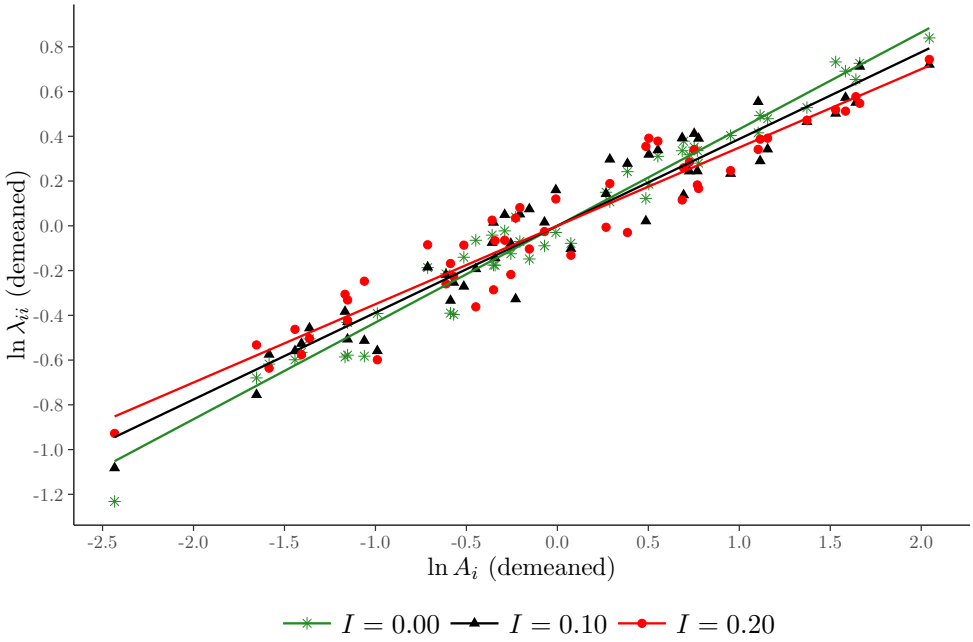
Relationship conditional on fixed effects

NOTES: This figure depicts the λ_{ii}–A_i relationship for an economy with a circular geography and a productivity distribution that follows a sine wave with frequency θ. Geographic locations and trade costs as in Figure 1. Country sizes L_i are positively correlated with A_i in the $\theta = 1$ state. See Appendix A.2 for parameterization details.

Since actual productivity does not exhibit a sine-wave distribution, we cannot conduct our empirical examination in terms of the frequency θ. Instead, we employ Moran's I, a standard measure of global spatial correlation for any characteristic arbitrarily distributed over a geography

endowed with a distance metric.[5] To illustrate our prediction with this more general measure, consider again the circular geography with equal-sized countries. We generate random productivity distributions that differ only in their spatial correlation as measured by Moran's I.[6] Figure 4 plots the expenditure-productivity relationship against Moran's I for three productivity draws with different levels of spatial correlation. There is a clear negative relationship: as Moran's I increases, the equilibrium domestic share of expenditure is less responsive to domestic productivity.

Figure 4: Circular geography with equal-sized countries and arbitrary productivities



NOTES: This figure depicts the λ_{ii}-A_i relationship for an economy with a circular geography and three selected randomly generated productivity distributions. Geography, equal-sized locations, and trade costs as in Figure 1. See Appendix A.2 for parameterization details.

Employing Moran's I to measure spatial correlation and introducing country fixed effects are also jointly sufficient to address the fact that countries are not evenly spaced on a circular geography. As mentioned above, the Moran's I statistic can be computed for any geography endowed with a distance metric. In asymmetric geographies, some countries are more "remote" from economic activity, regardless of the distribution of productivities. All else equal, remote countries exhibit a higher domestic share of expenditure. This variation is absorbed by country fixed effects.

Figure 5 plots the expenditure-productivity relationship against Moran's I for an economy with

[5] Moran's I is a commonly used measure of global spatial correlation. Formally, Moran's I for $\ln A_{it}$ in year t is

$$I_t = \frac{N}{\sum_i \sum_{j \neq i} w_{ij}} \frac{\sum_i \sum_{j \neq i} w_{ij} \left(\ln A_{it} - \overline{\ln A_t}\right) \left(\ln A_{jt} - \overline{\ln A_t}\right)}{\sum_\ell \left(\ln A_{\ell t} - \overline{\ln A_t}\right)^2},$$

where w_{ij} is the inverse distance between locations i and j and $\overline{\ln A_t}$ is the cross-sectional average in year t.

[6] See Appendix A.2 for details.

randomly located countries on a two-dimensional space and random assignments of productivity levels that differ only in their spatial correlation.[7] The vertical axis is the coefficient from a linear regression of the domestic share of expenditure on own productivity,

$$\ln \lambda_{iit} = \beta_t \ln A_{it} + \pi_i^I + \pi_t^T + \epsilon_{it}, \tag{4}$$

for country i at "time" t, where each t denotes an equilibrium associated with a different productivity distribution. The country fixed effects π_i^I control for differences in countries' time-invariant determinants of the domestic share of expenditure, such as remoteness. The "year" fixed effects π_t^T control for differences in the average domestic shares of expenditure across different spatial distributions of productivity. The year-specific (distribution-specific) slope coefficient β_t captures the resulting expenditure-productivity relationship for each spatial distribution. There is a clear negative relationship between the responsiveness of domestic share of expenditure to own productivity and the spatial correlation of productivities. While the relationship is not literally monotonic, the elasticity of λ_{iit} with respect to A_{it} is systematically lower at higher levels of Moran's I. This generalizes the patterns in our stylized, one-dimensional case to more realistic geographies.

2.3 Multiple-sector case

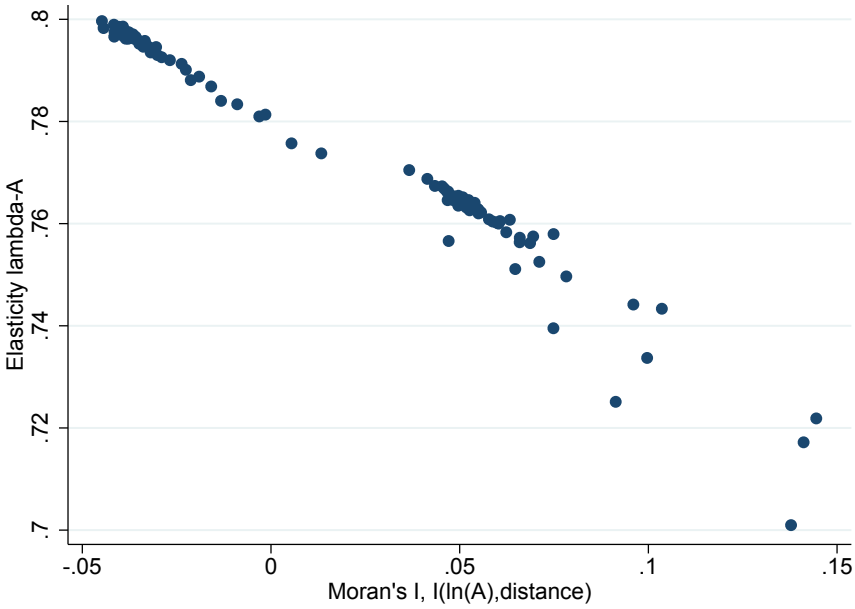
This section extends our predictions to an economy with multiple sectors. An empirical investigation of the previous section's prediction requires exogenous variation in the global productivity distribution. Our empirical strategy exploits such variation in the agricultural sector, which constitutes a small share of global trade. This section demonstrates that our prediction about the spatial correlation of productivity and the productivity-expenditure relationship holds for each sector in a multi-sector gravity model of trade. Moreover, the multi-sector welfare expression is analogous to the single-sector case, summing over sectors using their expenditure shares. Greater spatial correlation in a single sector implies greater dispersion in welfare across countries, provided that sectors' productivities are not perfectly anti-correlated. Thus, the previous section's predictions about trade in one sector can be empirically investigated in a multi-sector world.

Real consumption per capita in the multi-sector version of our economic environment is akin to equation (1), summing over sectors using their expenditure shares. We assume that consumers have Cobb-Douglas preferences over sectors $s = 1, \ldots, S$ with expenditure shares $\alpha_{js} \geq 0$ and CES preferences over varieties within sectors. As shown in Appendix A.3, real consumption per capita in this environment is

$$\ln \left(C_j / L_j \right) = \sum_{s=1}^{S} \alpha_{js} \left(\ln A_{js} + \gamma_s - \frac{1}{\epsilon_s} \ln \lambda_{jjs} \right), \tag{5}$$

[7] We generate a two-dimensional geography by randomly drawing N locations' coordinates from a normal distribution. The bilateral log trade costs between these N locations are proportionate to the log straight-line distances between them. We fix the first, second, and higher-order moments of the productivity distribution while varying its spatial correlation. Appendix A.2 details our procedure generating the geography and productivity distributions.

Figure 5: Random-geography economy



NOTES: This figure depicts how $cov(\ln \lambda_{ii}, \ln A_i)$ varies with the spatial correlation of productivity, as measured by Moran's I, over a randomly generated geography. We vary the spatial distribution of productivity while holding the first and second moments fixed. The elasticity of λ_{ii} with respect to A_i is $\hat{\beta}_t$ estimated from $\ln \lambda_{iit} = \beta_t \ln A_{it} + \pi_i^I + \pi_t^T + \epsilon_{it}$, as described in the text. See Appendix A.2 for parameterization details.

where γ_s depends on sectoral structural parameters other than productivity and ϵ_s is the sectoral trade elasticity. The first two terms, $\sum_{s=1}^{S} \alpha_{js} (\ln A_{js} + \gamma_s)$, are per capita welfare under autarky, and the final term, $-\sum_{s=1}^{S} \frac{\alpha_{js}}{\epsilon_s} \ln \lambda_{jjs}$, summarizes the gains from trade. Dispersion in per capita welfare across locations thus depends on the exogenous variation in productivities A_{js} and the endogenous variation in domestic shares of expenditure λ_{jjs}. We again compare two productivity distributions, correlated c and uncorrelated u, that have the same unconditional variance of productivity, $var\left(\sum_{s=1}^{S} \alpha_{js} \ln A_{js,c}\right) = var\left(\sum_{s=1}^{S} \alpha_{js} \ln A_{js,u}\right)$. For the sake of expositional brevity, assume also that expenditures shares are common across countries, $\alpha_{js} = \alpha_s \ \forall j$. We can then obtain the multi-sector analogue of equation (3):

$$\begin{aligned} var\left(\ln\left(C_{j,c}/L_j\right)\right) - var\left(\ln\left(C_{j,u}/L_j\right)\right) = 2\sum_{s=1}^{S}\sum_{s'=1}^{S} \frac{\alpha_s \alpha_{s'}}{\epsilon_{s'}} \Big\{ cov\left(\ln A_{js,u}, \ln \lambda_{jjs',u}\right) - cov\left(\ln A_{js,c}, \ln \lambda_{jjs',c}\right) \Big\} \\ - \sum_{s=1}^{S}\sum_{s'=1}^{S} \frac{\alpha_s \alpha_{s'}}{\epsilon_s \epsilon_{s'}} \Big\{ cov\left(\ln \lambda_{jjs,u}, \ln \lambda_{jjs',u}\right) - cov\left(\ln \lambda_{jjs,c}, \ln \lambda_{jjs',c}\right) \Big\} \end{aligned} \tag{6}$$

Just as in the single-sector case, for typical values of the sectoral trade elasticities, $2\frac{\alpha_s \alpha_{s'}}{\epsilon_{s'}}$ is an

order of magnitude larger than $\frac{\alpha_s \alpha_{s'}}{\epsilon_s \epsilon_{s'}}$. Thus, the difference in welfare dispersion is governed by the $cov\left(\ln A_{js}, \ln \lambda_{jjs'}\right)$ terms, provided that the block of $cov\left(\ln A_{js}, \ln \lambda_{jjs'}\right)$ terms is the same order of magnitude as the block of $cov\left(\ln \lambda_{jjs}, \ln \lambda_{jjs'}\right)$ terms.

Under what circumstances is studying differences in $cov\left(\ln A_{js}, \ln \lambda_{jjs}\right)$ in one sector alone informative about welfare dispersion? For simplicity, consider the two-sector case and three possible relationships between the two sectors: perfectly correlated productivities, perfectly anti-correlated productivities, and orthogonal productivities. In the perfectly correlated case, there is little scope for adjustment across sectors, and thus outcomes are similar to those obtained in the one-sector environment.[8] If sectoral productivities are perfectly anti-correlated, then outcomes in one sector may be exactly offset by outcomes in another, leaving welfare unchanged. That is, it is possible to construct circumstances in which the sum of covariances of productivities and domestic shares of expenditure within sectors is exactly the opposite of the sum of cross-sector covariances.[9]

In the orthogonal case of our illustrative sine-wave example, the presence of an additional sector introduces an additional margin of adjustment that does not qualitatively alter the economic mechanism. Naturally, this additional margin dampens the magnitude of welfare dispersion relative to the perfectly correlated case. To produce a multi-sector analogue of the sine-wave example presented in Section 2.2, we consider two symmetric sectors that differ only in their sine-wave frequency, so that their productivities are orthogonal by construction. The result is depicted in Figure 6, which is the multi-sector analogue of the example presented in Figure 2. The idiosyncratic, orthogonal variation in the second sector's productivity adds noise to the relationship between the domestic share of expenditure and productivity, but it does not change the basic pattern. When the first sector's productivity is more spatially correlated, the covariance of its log productivity and log domestic share of expenditure is smaller. This raises dispersion in welfare relative to the case in which the first sector's productivity is less spatially correlated.[10]

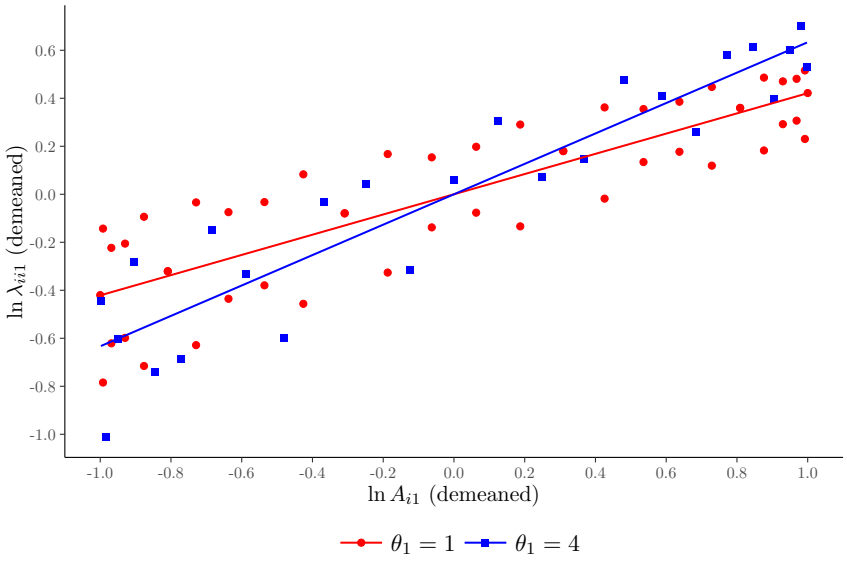
Thus, greater spatial correlation of productivity in one sector dampens the covariance of productivities and domestic expenditure shares. This is associated with greater dispersion in welfare, unless in the extreme case that there is an exact opposite pattern in the spatial distribution of productivity in other sectors. Having established that the spatial correlation of productivity may have important consequences for the distribution of welfare gains in a trading equilibrium, we next turn to an empirical setting in which there is exogenous variation in productivity.

[8] In fact, if all sectors have perfectly correlated productivities ($A_{is} \propto A_i \ \forall s$), perfectly correlated spatial linkages ($\tau_{ijs} \propto \tau_{ij} \ \forall s$) and equal trade elasticities ($\epsilon_s = \epsilon \ \forall s$), then expenditure shares are equal across sectors, $\lambda_{ijs} = \lambda_{ij}$, and the difference in welfare dispersion in equation (6) is exactly proportionate to the single-sector expression in equation (3).

[9] Consider the two-sector case with equal expenditure shares $\alpha_1 = \alpha_2 = \frac{1}{2} \ \forall j$ and equal trade elasticities $\epsilon_1 = \epsilon_2 = \epsilon$. If the two sectors' productivities are perfectly anti-correlated, such that $\ln A_{j1} + \ln A_{j2}$ is a constant, then it can be shown that $\sum_{s=1}^{S} \sum_{s'=1}^{S} \frac{\alpha_s \alpha_{s'}}{\epsilon_{s'}} cov\left(\ln A_{js}, \ln \lambda_{jjs'}\right) = 0$.

[10] See Table A.2 in Appendix A.2 for details of this example.

Figure 6: Two-sector sine-wave economy: $cov(\ln \lambda_{ii1}, \ln A_{i1})$



NOTES: This figure depicts the λ_{iis}–A_{is} relationship for $s = 1$ in a two-sector economy in which sectoral productivity follows a sine wave with frequency θ_s. The two series depicted are frequencies $\theta_1 = 1$ and $\theta_1 = 4$ for the first sector. The second sector has frequency $\theta_2 = 10$ in both cases. The two sectors have identical trade costs τ_{ijs}, expenditure shares $\alpha_{j1} = \alpha_{j2} = \frac{1}{2}$, and trade elasticities $\epsilon_1 = \epsilon_2$. The two lines are the line of best fit for each series. When productivity is more spatially correlated (when θ_1 is lower), $cov(\ln \lambda_{ii1}, \ln A_{i1})$ is lower.

3 The El Niño-Southern Oscillation

This section describes the exogenous variation that we employ to test the prediction that greater spatial correlation of productivity causes greater correlation of productivity and gains from trade. Causal empirical validation of our prediction requires a reshuffling of local productivities in a way that alters the global spatial correlation of productivities but is orthogonal to other determinants of the domestic share of expenditure. To illustrate this identification challenge, it is helpful to imagine an ideal experiment, akin to the simulations summarized in Figure 5, in which the researcher can randomly reshuffle a fixed set of productivities across the world such that the global spatial correlation of productivities is altered. In a spatially uncorrelated state of the world, productivities exhibit a more checkered-like spatial pattern. In a more spatially correlated state, productivities are rearranged such that similar productivities are spatially clustered together.

Such an ideal experiment is obviously not possible with international trade. However, we are able to approximate this idealized setting by exploiting a global climatic phenomenon known as the El Niño-Southern Oscillation (ENSO). This section summarizes the basic physics of ENSO and empirically demonstrates that it drives annual changes in the spatial correlation of agricultural productivity around the planet.

3.1 Background

The El Niño-Southern Oscillation is a naturally occurring climatic phenomenon characterized by mutuallyreinforcing circulation patterns between the atmosphere and the tropical Pacific ocean. Though the phenomenon begins in the tropics, ENSO is a major determinant of weather conditions around the planet. Indeed, at an annual frequency, ENSO is a major mode of the climate system and is often recovered as the first principal component of various local atmospheric or oceanographic variables across the planet (Sarachik and Cane, 2010).

While there is a continuum of ENSO conditions, it is simplest to describe the phenomenon as two extreme ENSO states that oscillate over a quasi-periodicity of 3-7 years. These conditions are broadly characterized by the amount of heat that is released from the tropical Pacific ocean into the atmosphere (Cane and Zebiak, 1985). In typical, "ENSO neutral" years, normal circulation patterns pushing in the westward direction maintain a pool of warm water against Indonesia and other land masses in the South Pacific. A warm or positive ENSO state, also known as an El Niño event, occurs when this circulation pattern weakens such that this pool of warm water spills eastward across a large area of the equatorial Pacific Ocean. With warm water exposed to the atmosphere over a greater area, El Niño events release more ocean heat into the atmosphere over a relatively short period. The opposite occurs during the cool or negative ENSO state, known as a La Niña event. Under La Niña, circulation patterns strengthen such that the same volume of warm water is pushed more firmly against the Indonesian landmass reducing sea-surface contact with the atmosphere and therefore reducing heat released from the ocean.

ENSO conditions in the tropical Pacific affect the spatial correlation of temperatures across the planet due to how heat travels when it is released in the tropics. Because there is almost no Coriolis effect near the equator (a result of the simple facts that the Earth is round and spins), atmospheric signals propagate rapidly throughout the tropics. Thus, the warm air initially released above the tropical Pacific Ocean propagated throughout the tropics by a transport mechanism in the atmosphere known as an "equatorial Kelvin wave" that sweeps the globe, altering climatological conditions almost simultaneously throughout the global tropics (Chiang and Sobel, 2002).[11] For this reason, it is often said that the tropical atmosphere is "teleconnected" through ENSO, as atmospheric conditions in locations distant from each other are remarkably correlated through this mechanism. Because the equatorial Kelvin wave that links climates around the equator is constrained primarily to the tropics, where the Coriolis effect is weak, the warming that prevails during El Niño events does not generally extend to higher latitudes. Instead, higher-latitude locations become slightly cooler on average because of changes to the circulation of the atmosphere.

Thus, in a positive ENSO year, temperature conditions around the world are essentially reshuffled such that there is a spatially contiguous area of hotter temperature across the tropics and subtropics while almost simultaneously there is a spatially contiguous area of cooler temperatures

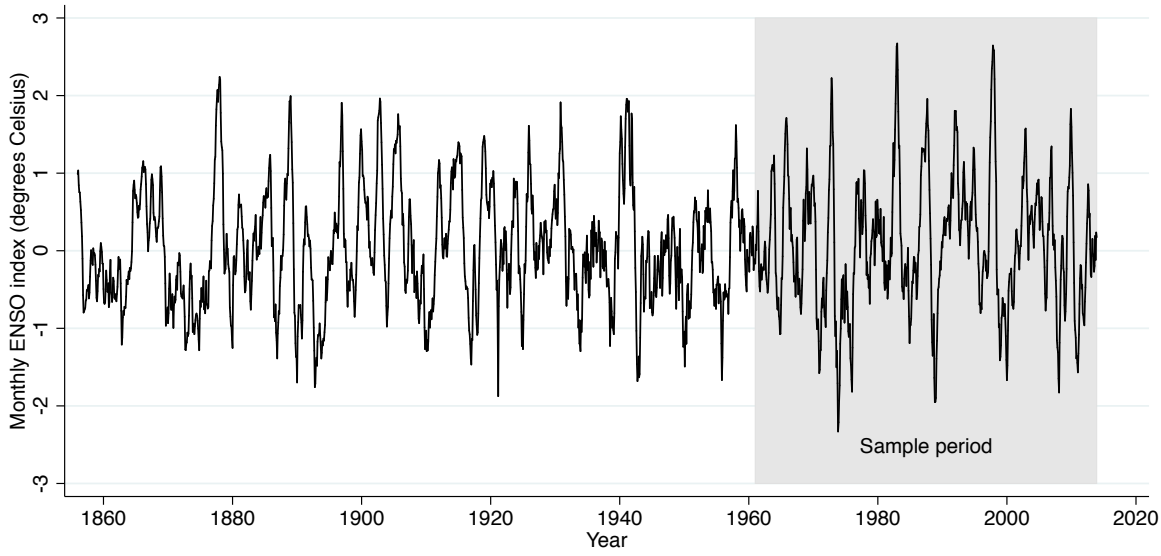
[11] This tropical phenomenon is described in Hsiang and Meng (2015). For a complete scientific treatment of ENSO physics, we refer interested readers to Sarachik and Cane (2010).

in higher-latitude locations.[12] The opposite occcurs during negative ENSO years. Less heat is released into the atmosphere and thus temperatures across the globe are less spatially organized.

ENSO conditions are typically summarized by the average sea-surface temperature over a fixed area in the tropical Pacific. In our analysis, we follow the climatology literature by employing the widely used NINO3.4 index, defined as average ocean temperature (in degrees Celsius) over a rectangular area bounded by 5°N - 5°S, 170°W - 120°W (see map in Figure E.1). Henceforth, we refer to this as the ENSO index.

Figure 7 plots the monthly ENSO index for 1856-2013, which extends further back than our sample period of 1961-2013. Observe that the index fluctuates between positive (El Niño) and negative (La Niña) values and appears stationary during this period. We turn to two important features of ENSO relevant for our empirical application: the monthly timing of ENSO events and its effects on the spatial pattern of temperatures around the planet.

Figure 7: Monthly ENSO index (1856-2013)



NOTES: Monthly ENSO index (NINO3.4) during 1856-2013. Shaded area shows sample period of analysis over 1961-2013.

3.2 Monthly timing of ENSO events

The timing of ENSO events, due to its tropical origins, are phase-shifted relative to the calendar year. An ENSO event begins appearing during April-May of a given year and lasts until the following April-May, a period known as the "tropical year." To demonstrate the within-year evolution of ENSO events, Figure E.2 isolates the 10 most positive ENSO events during our 1961-2013 sample period and plots the monthly ENSO index 12 months before and after a given December. For each ENSO event, the ENSO index begins rising around April-May of a given year and fully dissipates around June of the following year. Because the ENSO index typically peaks in December,

[12] Additionally, rainfall throughout most of this region is suppressed during positive ENSO years because the warm atmosphere prevents underlying moist surface air from rising (convecting) as readily.

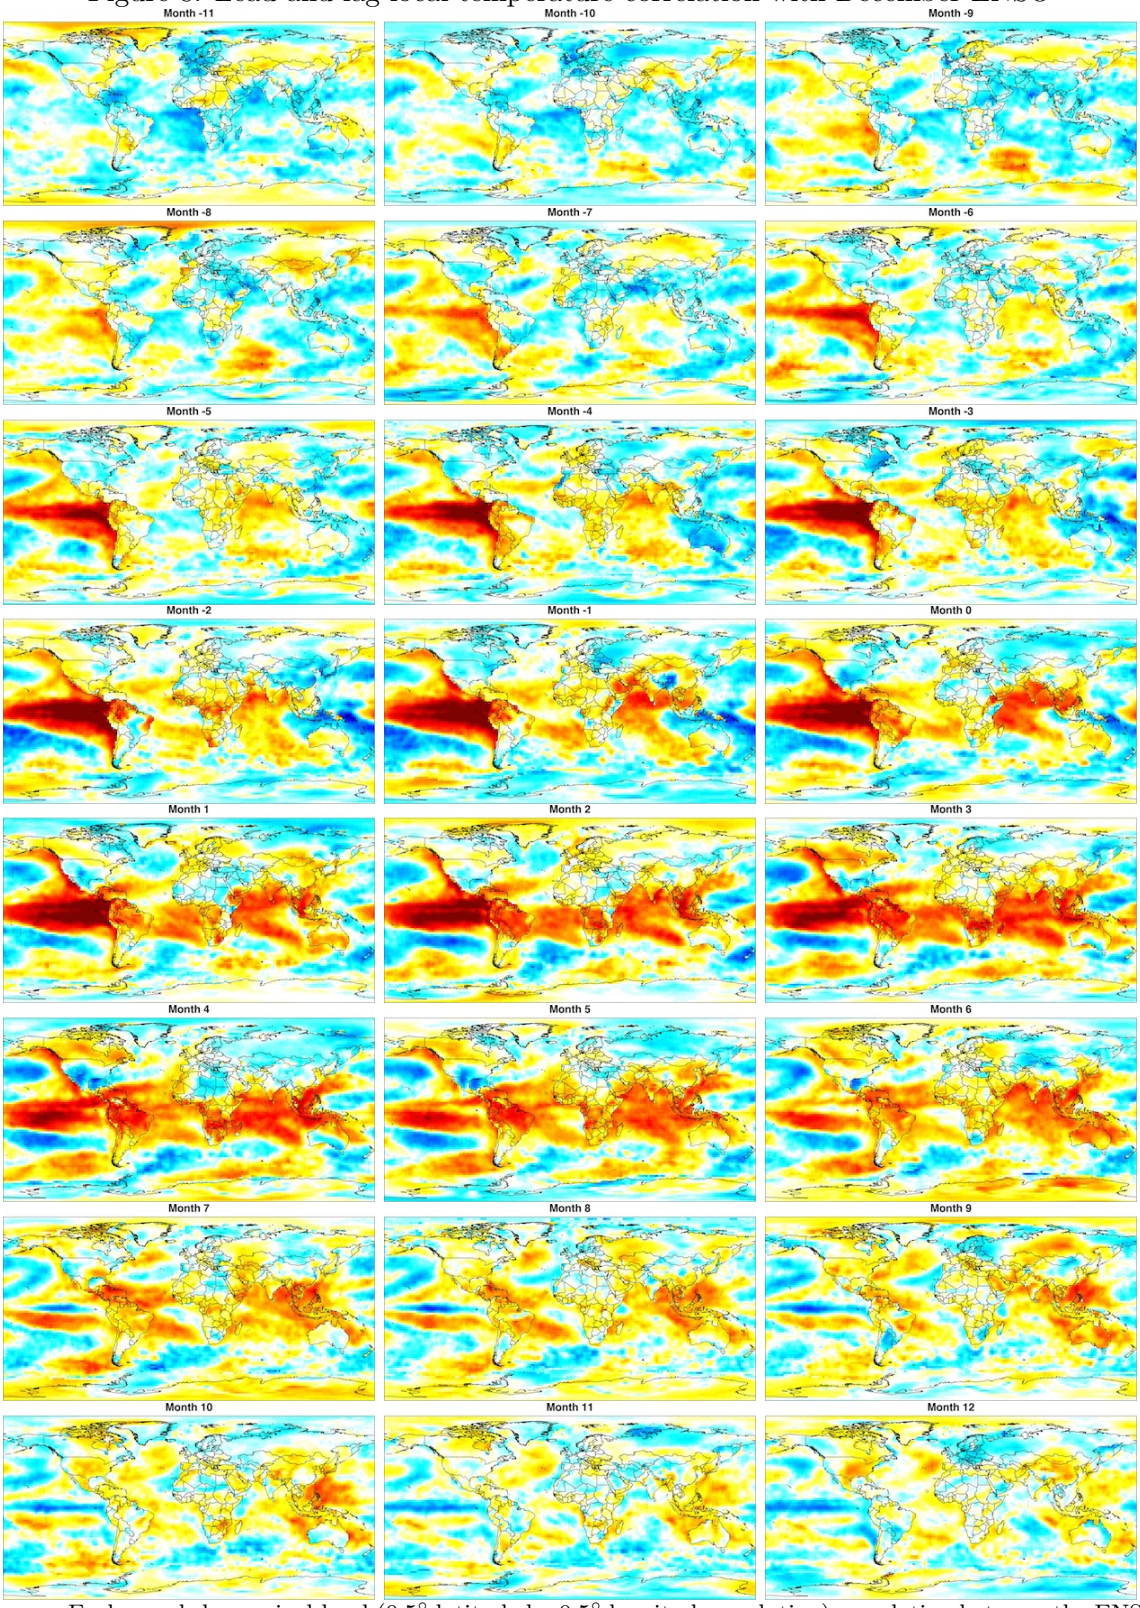
the cleanest annual measure of an ENSO event is simply the December ENSO index. This is our annualized measure of ENSO.

3.3 ENSO's spatial teleconnections

Figure 8 displays the canonical structure of warming that occurs when the ENSO index increases, month by month. Each map displays the time-series correlation of monthly temperatures for each pixel during the specified month and ENSO in month zero, defined as December. Yellow, orange, and red colors indicate locations that warm during El Niño events; blues indicate locations that cool. In May before an El Niño event (month -7), the east equatorial Pacific begins to warm. Regions throughout the tropics, both over land and the oceans, continue to warm for the next several months, peaking in the east Pacific in December (month 0) and over the rest of the tropics in March and April (months +3 and +4). The event then dissipates with little visible effect more than a year after the December peak. Higher latitudes experience some cooling through these months, though the effect is weaker. Figure 8 shows that the local impacts on temperatures around the planet from a single ENSO event straddles two calendar years. Because socio-economic data is typically indexed by calendar years, detecting all effects from an ENSO event in year t requires examining the combined effects on outcomes in year t and $t+1$.

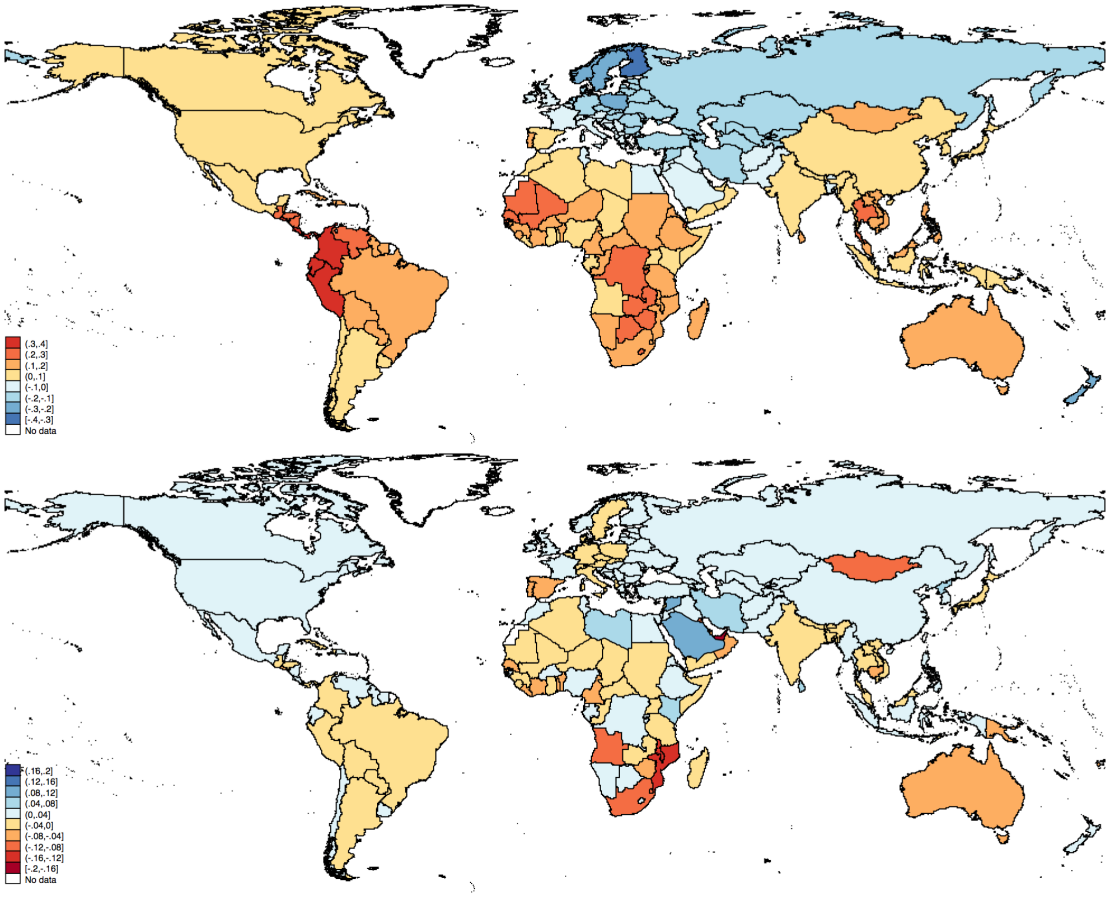
The spatial and temporal patterns shown in Figure 8 suggest that ENSO could drive global spatial correlation in cereal yields, a productivity measure that is both directly observed at the country-level with near global data coverage and has previously been documented to be sensitive to local temperature conditions (Schlenker and Roberts, 2009; Hsiang and Meng, 2015). To demonstrate this at an annual time scale, for which productivity and trade data is available, Figure 9 shows country-level combined effects of a 1-degree increase in contemporaneous and lagged December ENSO indices on crop-area weighted annual temperature (top panel) and log cereal yields (bottom panel). Each coefficient can be interpreted as the combined linear effect of a 1-degree increase in December ENSO in year t on annual temperature or productivity for that country in year t and $t+1$ after removing a country fixed effect and country-specific linear trend. ENSO fluctuations generate a clear spatial pattern in country-level temperatures and agricultural productivity. Consistent with the tropical climatic dynamics discussed above, increases in the ENSO index tend to raise temperatures and lower agricultural productivities for countries closer to the tropics. Conversely, a similar ENSO event drives cooler temperatures and increased agricultural productivities in higher-latitude countries. This latitude dependence is displayed more sharply in Figure 10, which shows country-level ENSO effects on temperature and log cereal yields as a (local polynomial) function of a country's area-weighted centroid latitude.

Figure 8: Lead and lag local temperature correlation with December ENSO



NOTES: Each panel shows pixel-level (0.5° latitude by 0.5° longitude resolution) correlation between the ENSO index in December and pixel-level monthly temperatures for 11 months before (lead) and 12 months after (lag) December. Blue shows areas with negative correlation. Red shows areas with positive correlation.

Figure 9: ENSO's spatial effects on temperature and log cereal yields



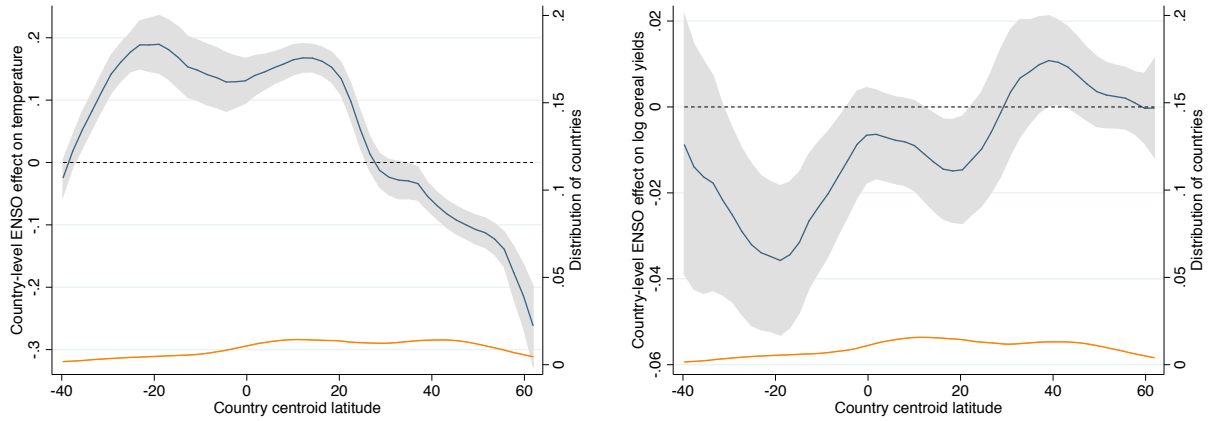
NOTES: Maps show the country-level combined linear effect of December ENSO in year t and $t-1$ on crop-area weighted temperature (top panel) and log cereal yields (bottom panel).

Figures 9 and 10 visually suggest that ENSO could drive global spatial correlation in agricultural productivity, but they do not provide a quantitative account. To quantify spatial correlation, we turn to a standard measure, the Moran's I statistic. In particular, we construct an annual Moran's I statistic for both country-level temperature and log cereal yields, with the spatial weight equal to the inverse great-circle distance between countries' area-weighted centroids.

Figure 11 plots the relationship between ENSO and the global spatial correlation in temperature (left panel) and agricultural productivity (right panel). Specifically, to capture the issue of ENSO timing in a bivariate scatter, each plot shows the relationship between Moran's I in year t on the average December ENSO index in years t and $t-1$. An increase in the ENSO index raises both the global spatial correlation in temperature and agricultural productivity. ENSO, in this simple bivariate model, explains 13% and 10% of annual variation in the global spatial correlations of temperature and agricultural productivity, respectively.

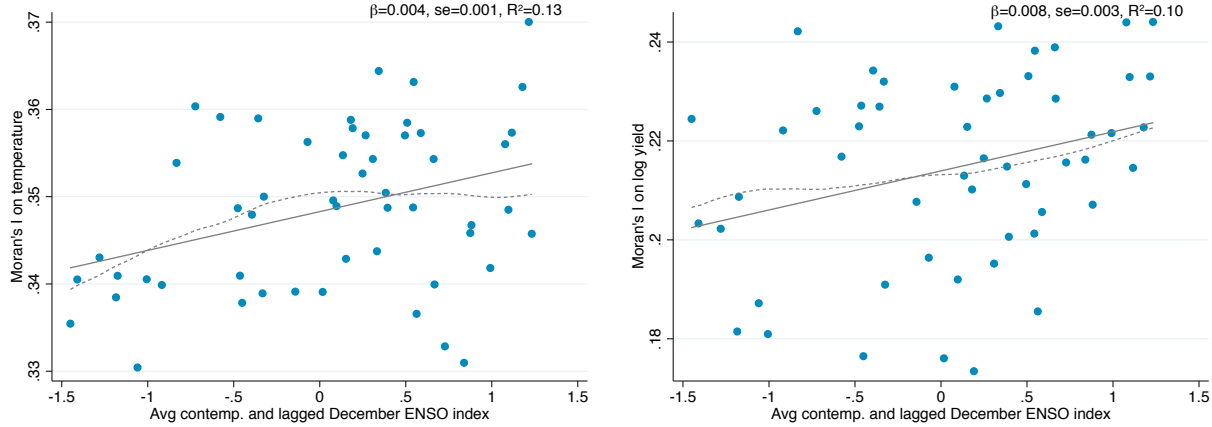
To more flexibly characterize the relationship between ENSO and the spatial correlation in agricultural productivity, Table 1 presents regressions of the annual Moran's I statistic for log cereal yields on flexible polynomial functions of December ENSO in year t and $t-1$. Two results

Figure 10: ENSO's latitude-dependent effects on temperature and log cereal yields



NOTES: Left (right) panel shows the country-level ENSO effects on crop-area weighted temperature (log cereal yields) from Figure 9 as a local polynomial function of a country's area-weighted centroid latitude. Local polynomial fit uses an Epanechnikov kernel with a rule-of-thumb bandwidth (Fan and Gijbels, 1996). 90% confidence interval shown. Orange line shows the distribution of country centroids.

Figure 11: ENSO and Moran's I for temperatures and log cereal yields



NOTES: Left (right) panel shows the relationship between Moran's I on crop-weighted country-level temperature (log cereal yields) in year t and averaged contemporaneous and lagged December ENSO. Linear fit shown as solid line. Local polynomial fit shown as dashed line.

are evident: (1) spatial correlation in cereal yields responds to the interaction of contemporaneous and lagged ENSO, and (2) there appears to be nonlinearity in the response function as captured by the precision of several quadratic ENSO coefficients. These findings inform our choice of functional form in Section 4. Finally, consistent with our assumptions in Section 2.1, ENSO does not appear to drive either the global mean or variance of agricultural productivity (see Figure E.3). That is, ENSO reshuffles agricultural productivities around the world while holding the unconditional first and second moments of agricultural productivity fixed.

Table 1: ENSO and Moran's I in log cereal yields
Outcome is Moran-I in log cereal yields

	(1)	(2)	(3)	(4)
$ENSO_t$	0.005***	0.005***	0.006***	0.004
	(0.001)	(0.001)	(0.002)	(0.003)
$ENSO_t^2$			-0.002	0.001
			(0.002)	(0.002)
$ENSO_{t-1}$	0.003**	0.004**	0.003**	0.004**
	(0.001)	(0.002)	(0.001)	(0.002)
$ENSO_{t-1}^2$			0.001	0.002**
			(0.001)	(0.001)
$ENSO_t$ x $ENSO_{t-1}$		0.004**		0.002
		(0.002)		(0.002)
$ENSO_t^2$ x $ENSO_{t-1}^2$				-0.002*
				(0.001)
Observations	53	53	53	53

NOTES: Time-series regressions of Moran's I in log cereal yields on nonlinear functions of contemporaneous and lagged December ENSO. All models include a linear time trend. Serial correlation and heteroscedasticity robust Newey-West standard errors with optimal bandwidth in parentheses (Newey and West, 1987). *** p<0.01, ** p<0.05, * p<0.1.

4 Empirical validation

The theoretical results in Section 2 suggest that the relationship between log domestic share of expenditure, $\ln \lambda_{ii}$, and log agricultural productivity, $\ln A_i$, should become flatter when the spatial correlation of productivity across the entire trading network increases. This section empirically tests this prediction using exogenous ENSO-driven changes in the global spatial correlation of productivity. We first describe our estimation strategy, then report our main finding and subject it to a series of robustness checks. Appendix B details data sources used.

4.1 Estimation strategy

For country i in year t, consider the following regression equation:

$$\ln \lambda_{iit} = \beta_0 \ln A_{it} + \beta_1 \ln A_{it} I_t + \mathbb{Z}_{it}\Pi + \mu_{it} \tag{7}$$

where λ_{iit} is country i's domestic share of expenditure in year t, A_{it} is its cereal yield, and I_t is the Moran's I statistic capturing the spatial correlation of $\ln A_{it}$ for all countries in year t. $\mathbb{Z}_{it}$ is a vector of country fixed effects, year fixed effects, and country-specific trends. Our reduced-form parameter of interest is β_1. Provided that $\widehat{\beta}_0 > 0$, $\widehat{\beta}_1 < 0$ implies that an increase in the

spatial correlation of agricultural productivity in a given year dampens the relationship between a country's domestic share of expenditure and its agricultural productivity.[13]

Through the lens of our theoretical framework in Section 2.1, an estimate of $\widehat{\beta}_1 < 0$ implies that $cov\left(\ln A_{i,c}, \ln \lambda_{ii,c}\right) < cov\left(\ln A_{i,u}, \ln \lambda_{ii,u}\right)$. When spatial correlation increases, countries benefiting from higher productivities experience greater gains from trade, as summarized by a lower domestic share of expenditure. The same increase in the spatial correlation of agricultural productivity leads countries experiencing relatively lower productivities to experience lower gains from trade. Thus, interpreted through equation (3), $\widehat{\beta}_0 > 0$ and $\widehat{\beta}_1 < 0$ would verify our prediction that trade increases the global variance in welfare as spatial correlation rises.

OLS estimation of equation (7) may be problematic to the extent that expenditure shares and productivity are simultaneously determined or if there are omitted determinants of expenditure shares that are correlated with productivity, even after conditioning on $\mathbb{Z}_{it}$. For example, demand shocks could affect expenditure shares and elicit supply responses that change average yields. Similarly, if domestic cereal output requires imported intermediate goods, then unobserved trade cost shocks may jointly affect domestic cereal yields and domestic share of expenditure. We identify β_0 and β_1 by employing instruments for $\ln A_{it}$ and $\ln A_{it}I_t$ that are plausibly uncorrelated with μ_{it}.

Our instrumental-variables strategy exploits the change in the global spatial correlation of yields as ENSO "reshuffles" agricultural productivity around the world, as described in Section 3. This means that ENSO drives both local variation in cereal yields and global variation in the spatial correlation of yields. We first estimate ENSO-driven effects on country-specific yields and global spatial correlation in yields in order to use their predicted values as instruments.[14] Specifically, we estimate:

$$\ln A_{it} = f_i(\mathbf{ENSO}_t) + \mathbb{Z}_{it}\Omega_1 + \omega_{1it} \tag{8}$$

$$I_t = g(\mathbf{ENSO}_t) + \mathbb{W}_t\Omega_2 + \omega_{2t} \tag{9}$$

where following our discussion in Section 3.2 about the timing mismatch between the "tropical year" that characterizes ENSO events and the calendar year in which most socio-economic data is indexed, $\mathbf{ENSO}_t \equiv \{ENSO_t, ENSO_{t-1}\}$. This enables the detection of a given ENSO event $\mathbf{ENSO}_t \equiv \{ENSO_t, ENSO_{t-1}\}$ across two calendar years. The contemporaneous and lagged ENSO effects on country-level cereal yields and global spatial correlation in yields are dictated by the functions $f_i()$ and $g()$.[15] Ω_1 and Ω_2 are coefficients on panel and time-series semi-parametric

[13] Observe that estimating equation (7) is a variant of equation (4), used to theoretically demonstrate our prediction under the most realistic set of modeling assumptions. Equation (7) further imposes the functional-form assumption that the $\ln \lambda_{iit} - \ln A_{it}$ relationship is linear in Moran's I, which we verify below.

[14] Using a predicted variable as an instrument variable (IV) is also known as "simulated" IV. Prominent applications can be found in health (Currie and Gruber, 1996; Cutler and Gruber, 1996; Schlenker and Walker, 2015) and public economics (Moffitt and Wilhelm, 2000; Gruber and Saez, 2002; Dahl and Lochner, 2012).

[15] Rather than separately estimating equations (8) and (9), an alternative approach would be to only estimate equation (8), obtain ENSO-predicted yields, $\widehat{f}_i(\mathbf{ENSO}_t)$, and use it to construct a predicted Moran's I statistic. Such an approach, however, implicitly assumes that $I_t = I(\{\widehat{f}_i(\mathbf{ENSO}_t)\})$ is the correct functional relationship between ENSO and the global spatial correlation in yields. Our approach, which follows Kelejian (1971) and Amemiya (1974),

controls, respectively.

Using our two instruments, $\widehat{f}_i(\mathbf{ENSO}_t)$ and $\widehat{f}_i(\mathbf{ENSO}_t)\widehat{g}_i(\mathbf{ENSO}_t)$, from equations (8) and 9, we then estimate the following two first-stage regressions, one for each of our endogenous variables:[16,17]

$$\ln A_{it} = \alpha_{11}\widehat{f}_i(\mathbf{ENSO}_t) + \alpha_{12}\widehat{f}_i(\mathbf{ENSO}_t)\widehat{g}(\mathbf{ENSO}_t) + \mathbb{Z}_{it}\Gamma_1 + \upsilon_{1it} \tag{10}$$

$$\ln A_{it}I_t = \alpha_{21}\widehat{f}_i(\mathbf{ENSO}_t) + \alpha_{22}\widehat{f}_i(\mathbf{ENSO}_t)\widehat{g}(\mathbf{ENSO}_t) + \mathbb{Z}_{it}\Gamma_2 + \upsilon_{2it} \tag{11}$$

where Γ_1 and Γ_2 are coefficients on semi-parametric controls. The second stage in our IV estimator is then equation (7). Standard errors are clustered at both the country and continent-year levels to allow arbitrary forms of serial correlation and heteroscedasticity within a country and spatial correlation and heteroscedasticity across countries within a particular continent and year.

There are several potential concerns regarding this approach. First, as with any IV estimator, our instruments should satisfy the exclusion restriction and generate sufficiently strong first-stage regressions. We conduct several indirect tests of our exclusion-restriction assumption in Section 4.3 and demonstrate that our estimates pass standard weak-instrument tests. Second, because our estimation uses a global climatic index driving both country-level agricultural productivity and the global spatial correlation in that productivity, one may be concerned that estimates of β_0 and β_1 are identified using a particular nonlinear functional form of $\mathbf{ENSO}_t$. In what follows, we show that we recover very similar estimates using a variety of functional forms for $f_i()$ and $g_i()$, such that our main result is not driven by any particular nonlinear transformation of $\mathbf{ENSO}_t$.

4.2 Main results

It is instructive to first consider the analogous OLS regression of equation (7), as shown in column 1 of Table 2. The top panel includes linear country-specific trends, while the bottom panel includes quadratic country-specific trends. Both models include country and year fixed effects. Both OLS models provide precise and positive estimates of β_0. However, the estimate of β_1 is noisy and opposite the expected sign.

To address the potential bias of OLS estimates, columns 2 through 5 of Table 2 report IV estimates of β_0 and β_1 in equation (7) by varying (i) the polynomial order of the ENSO function in $f_i()$ and $g()$ in equations (8) and (9) (across columns) and (ii) the polynomial order of time trends in equations (7)-(11) (top vs bottom panels). All models incorporate country and year fixed effects as included instruments. We show several first-stage tests. The Cragg-Donald and Kleibergen-Paap

allows for flexible estimation of $g()$ and in general will yield the more efficient IV estimator.

[16] While use of predicted regressors in an OLS setting typically yields incorrect standard errors that omit sampling variability, predicted regressors can be used as instruments for endogenous variables. This is because in an IV setting any instrumenting of an endogenous variable retains its sampling variability (see Wooldridge 2002, p.117).

[17] This IV model is exactly identified: there are two endogenous variables and two instruments. Alternatively, one could estimate $\widehat{f}_i(\mathbf{ENSO}_t)$ and $\widehat{g}(\mathbf{ENSO}_t)$ directly within equations (10) and (11). However, as will be discussed below, $f_i()$ involves the estimation of country-specific parameters and such an approach would result in a highly over-identified IV setting that is more vulnerable to weak-instrument concerns (see Angrist and Pischke 2009, p.209).

F-statistics presents the joint F-statistic for both first-stage regressions under the assumption of iid and cluster-robust errors, respectively (Cragg and Donald, 1993; Kleibergen and Paap, 2006). Because there is no clear relationship between the F-statistics and properties of the IV estimator in cases with multiple endogenous variables and non-iid standard errors, we also present the p-value from the Anderson-Rubin test against a null in which the endogenous regressors in equation (7) are jointly zero (Anderson and Rubin, 1949). This test is robust to the presence of weak instruments.

Columns 2 through 5 of Table 2 shows estimates of β_0 and β_1 that are of the expected sign, statistically significant, and not sensitive to the polynomial order of time trends. Importantly, our estimates are fairly stable across different combinations of linear and quadratic $\mathbf{ENSO_t}$ functions in $f_i()$ and $g()$ used to generate our instruments. The similar reduced-form coefficients obtained from different nonlinear transformations of $\mathbf{ENSO_t}$ shows that our results are not uniquely identified by a particular functional-form assumption. For all specifications in Table 2, our IV estimates exhibit strong first-stage regressions and second-stage estimates that are robust even if instruments were weak. To quantify these estimates in welfare terms, suppose the global spatial correlation of agricultural productivity were to increase by one standard deviation relative to the historical mean. Our IV estimates, when applied to the expression for the variance of welfare in equation (2), imply a 4% increase in the global variance of welfare under our benchmark model in column 5, panel A of Table 2 (see Appendix D.1 for details). Appendix Table F.1 replicates Table 2 with standard errors that are clustered at the country and year levels. Allowing for arbitrary forms of spatial correlation and heteroscedasticity across all countries within a given year does not qualitatively change our main result.

4.3 Additional robustness checks

This section presents several robustness checks of our main empirical result. They are designed to test both the validity of our statistical modeling assumptions and the interpretation of our results. Our benchmark model throughout is that shown in column 5, panel A of Table 2.

Indirect tests of exclusion-restriction assumption Our IV model correctly identifies β_0 and β_1 when ENSO conditions influence a country's domestic share of cereal expenditure only through its effects on local yields and the global spatial correlation in yields.[18] While it is unlikely that ENSO, as a naturally occurring climatic phenomenon, is affected by economic activity, the exclusion violation could be violated if ENSO affects domestic cereal expenditure outside of its influence on

[18] For example, Hsiang, Meng and Cane (2011) show over the same period that warmer ENSO conditions increase the likelihood of civil conflicts in the tropics. This relationship, however, need not imply an exclusion-restriction violation. Suppose ENSO increases civil conflicts in the tropics only through lowered cereal yields. In that scenario, civil conflict would serve as a "bad control" in our IV specification, potentially biasing our coefficients of interest (see Angrist and Pischke 2009, p. 64-68). Our exclusion-restriction assumption is invalid only if ENSO increases civil conflicts partially through non-agricultural channels and civil conflicts affect domestic share of expenditure by, for example, raising international trade costs relative to internal trade costs. Because the current climate-conflict literature currently supports both agricultural and non-agricultural channels (Hsiang, Burke and Miguel, 2013), the inclusion of conflict as a control would not delivera unique interpretation: it either jeopardizes a valid identification strategy or show that our instrumental-variable strategy is invalid.

Table 2: Main result: Domestic share of expenditure and spatial correlation of productivity

Outcome is log domestic share of expenditure					
	(1) OLS	(2) IV	(3) IV	(4) IV	(5) IV
			Panel A: Linear trends		
$\ln A_{it}$	0.363**	3.754***	3.785***	2.133***	2.219***
	(0.166)	(1.248)	(1.158)	(0.791)	(0.757)
$\ln A_{it} \times I_t$	0.645	-14.333**	-14.471***	-7.382**	-7.783**
	(0.757)	(5.810)	(5.445)	(3.510)	(3.328)
Pct. change in welfare variance	-0.300	6.510**	6.571***	3.389**	3.571**
	(0.353)	(2.601)	(2.409)	(1.643)	(1.567)
ENSO order for predicting $\ln A_{it}$		1	1	2	2
ENSO order for predicting I_t		1	2	1	2
Cragg-Donald F-stat		22.73	24.55	63.92	67.20
Kleibergen-Paap F-stat		29.56	17.10	15.12	16.86
Anderson-Rubin weak-id robust p-value		4.66e-07	2.72e-08	0	0
Observations	5681	5681	5681	5681	5681
			Panel B: Quadratic trends		
$\ln A_{it}$	0.406*	3.473***	3.575***	3.100**	3.091***
	(0.210)	(1.300)	(1.029)	(1.240)	(0.982)
$\ln A_{it} \times I_t$	0.501	-13.471**	-13.952***	-12.073**	-12.028***
	(0.894)	(6.210)	(4.892)	(5.722)	(4.411)
Pct. change in welfare variance	-0.233	6.127**	6.341***	5.503**	5.483***
	(0.415)	(2.711)	(2.142)	(2.597)	(2.052)
ENSO order for predicting $\ln A_{it}$		1	1	2	2
ENSO order for predicting I_t		1	2	1	2
Cragg-Donald F-stat		17.09	20.80	24.62	25.77
Kleibergen-Paap F-stat		11.21	11.03	6	11.60
Anderson-Rubin weak-id robust p-value		5.62e-05	4.63e-06	5.13e-07	3.10e-08
Observations	5681	5681	5681	5681	5681

NOTES: Estimates of equation (7). Column 1 shows OLS estimates. Columns 2-5 show IV estimates. Columns 2-5 differ in the ENSO functional form used for predicting instruments. All models incorporate country and year fixed effects as included instruments. Panel A (B) includes global or country-specific linear (quadratic) trends in equations (7)-(11). Percentage change shown in the variance of welfare for a one-standard-deviation increase in Moran's I relative to the historic mean. Standard errors clustered at country and continent-year levels in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

cereal yields. For example, as noted above, a violation would occur if ENSO were to directly affect trade costs.

As indirect tests of our exclusion-restriction assumption, Table 3 augments our benchmark model with additional controls designed to capture time-and-country-varying trade costs. Column

Table 3: Robustness: Controls for trade costs
Outcome is log domestic share of expenditure

	(1)	(2)	(3)	(4)	(5)
$\ln A_{it}$	2.219***	2.260***	2.231***	2.207***	2.209***
	(0.757)	(0.751)	(0.761)	(0.749)	(0.779)
$\ln A_{it} \times I_t$	-7.783**	-7.982**	-7.815**	-7.800**	-7.702**
	(3.328)	(3.306)	(3.338)	(3.262)	(3.453)
ln oil price × average ln λ_{ii}	No	Yes	No	No	No
ln oil price × centrality	No	No	Yes	No	No
Year FE × average ln λ_{ii}	No	No	No	Yes	No
Year FE × centrality	No	No	No	No	Yes
Cragg-Donald F-stat	67.20	67.23	67.25	69	66.23
Kleibergen-Paap F-stat	16.86	16.99	16.92	17.35	17.48
Anderson-Rubin weak-id robust p-value	0	0	0	1.46e-10	0
Observations	5681	5681	5681	5681	5681

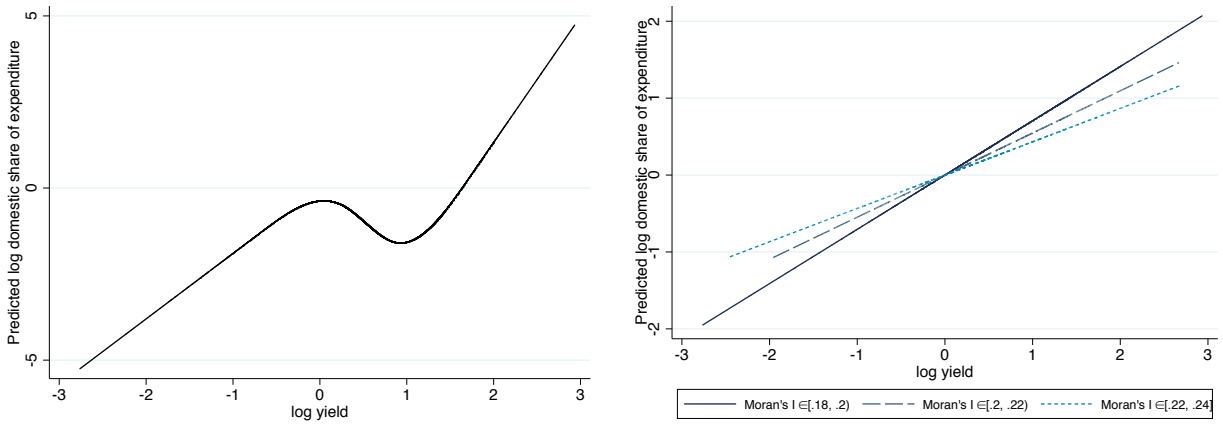
NOTES: Estimates of equation (7). All models incorporate quadratic functions of $\mathbf{ENSO}_t$ to predict $\ln A_{it}$ and I_t, linear time trends, and country and year fixed effects as included instruments. Column 1 replicates column 5, panel A of Table 2. Column 2 (3) controls for the interaction of global log oil price and cross-sectional average log domestic share of expenditure (output-weighted inverse distance averaged across all other countries). Column 4 (5) controls for the interaction of year fixed effects and cross-sectional average log domestic share of expenditure (output-weighted inverse distance averaged across all other countries). Standard errors clustered at country and continent-year levels in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

1 replicates our benchmark result. In column 2, we add a country-year-varying proxy for trade costs by interacting the global annual crude oil price with the country's average log domestic share of expenditure over our sample period. Column 3 uses an alternative measure of cross-sectional trade openness by interacting the global oil price with a country's centrality, measured as a country's average output weighted inverse distance to every other country, where output weights is the time-series averaged log agricultural output of the trading partner. In both cases, these proxies for trade costs do not meaningfully alter our estimates of β_0 and β_1. Columns 4 and 5 provide a more flexible specification by interacting year fixed effects with these two cross-sectional measures of trade openness used in columns 2 and 3. Again, our coefficients of interest are relatively unaffected by the inclusion of these controls, suggesting that unobserved shocks to trade costs are not correlated with our ENSO-driven instruments.

Nonlinearity in second-stage regression Our second-stage regression in equation (7) assumes that log domestic share of expenditure is linear in both log cereal yields and the interaction of log cereal yields and Moran's I. We now examine these functional-form assumptions. Tests of nonlinearity require breaking our main instruments into multiple variables. This potentially creates weaker first-stage regressions. To minimize this concern, we avoid jointly examining nonlinearity in both the interacted and uninteracted terms of equation (7). Rather, we first examine whether log

domestic share of expenditure is linear in log cereal yields. Conditional on that linear relationship, we then examine whether the interaction between log yield and the spatial correlation of log yields is linear. The first test is presented in the left panel of Figure 12 in which a restricted cubic spline is fitted onto log yields in a version of equation (7) without the interaction term. The relationship between log domestic share of expenditure and log yields appears roughly linear. The right panel of Figure 12 examines how this relationship varies with Moran's I by estimating separate slopes for each tercile in the empirical distribution of Moran's I. The linear relationship between log domestic share of expenditure and log yield becomes flatter as Moran's I values increase. This interaction relationship appears monotone and roughly linear and is the empirically estimated analogue to our stylized example in Figure 2.

Figure 12: Robustness: Nonlinearities in second-stage functional form



NOTES: Left panel fits a restricted cubic spline function to unconditional log yield. Spline functions for endogenous yield and exogenous ENSO-predicted yield have four knots placed along the support according to Harrell (2001). Right panel assumes log yield function is linear and discretizes Moran's I into three evenly-spaced bins with similarly discretized instruments. All models based on IV estimator corresponding to column 5, panel A of Table 2.

Time-varying parameters Table 4 examines whether β_0 and β_1 vary over our sample period. Column 1 shows our benchmark result. Column 2 shows estimates using the years 1961-1987, the first half of our sample period. Column 3 restricts the years 1988-2013, the second half of our sample period. We do not find that our results differ statistically across the two subsample periods.

Randomization placebo tests Our main source of identifying variation is global time-series fluctuations in the ENSO cycle, as shown in Figure 7. While it is plausible that in sufficiently long samples ENSO would be uncorrelated with unobserved determinants of gains from trade, spurious correlations between ENSO conditions and the error term μ_{it} could occur within a 53-year sample period. This would lead to biased IV estimates.

To examine the relevance of this concern, we conduct a placebo test by randomly reshuffling years in our panel data, breaking the time-series link between domestic shares of expenditure and ENSO-driven changes in country-level cereal yields and the global spatial correlation in yields. This allows us to obtain an empirical distribution of our reduced-form coefficients, β_0 and β_1,

Table 4: Robustness: Time-varying parameters
Outcome is log domestic share of expenditure

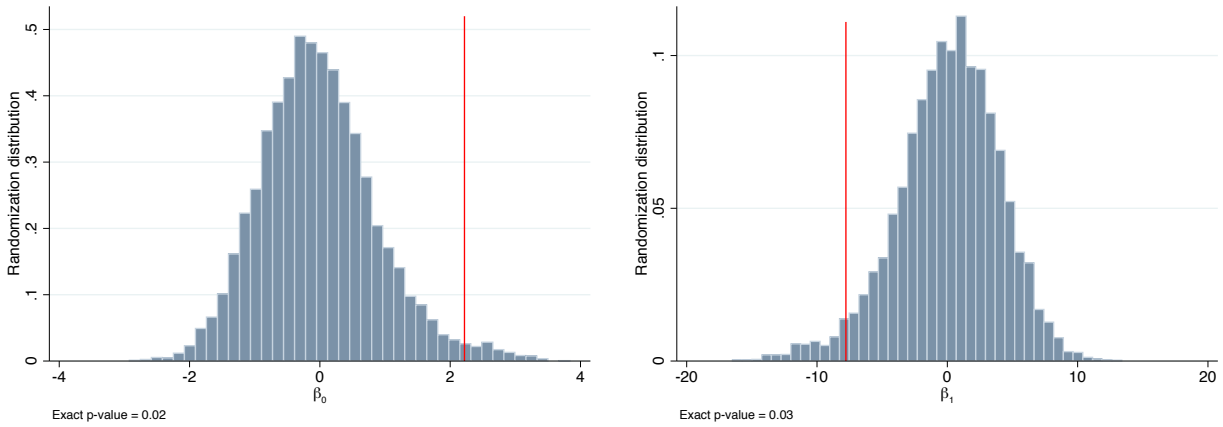
	(1)	(2)	(3)
$\ln A_{it}$	2.219***	1.633	2.030***
	(0.757)	(0.994)	(0.702)
$\ln A_{it} \times I_t$	-7.783**	-5.278	-7.448**
	(3.328)	(4.615)	(3.087)
Pct. change in welfare variance	3.850**	2.066	4.116**
	(1.693)	(1.777)	(1.779)
Cragg-Donald F-stat	67.20	32.22	21.84
Kleibergen-Paap F-stat	16.86	8.962	13.19
Anderson-Rubin weak-id robust p-value	0	3.80e-06	1.22e-05
2nd stage sample period	1961-2013	1961-1987	1988-2013
Observations	5681	2757	2919

NOTES: Estimates of equation (7). All models incorporate quadratic functions of $\mathbf{ENSO}_t$ to predict $\ln A_{it}$ and I_t, linear time trends, and country and year fixed effects as included instruments. Column 1 replicates column 5, panel A of Table 2. Column 2 restricts sample to 1961-1987. Column 3 restricts sample to 1988-2013. Standard errors clustered at country and continent-year levels in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

under placebo conditions and compute the probability of observing our benchmark estimates if years were randomly assigned.

The left and right panels of Figure 13 shows the empirical distribution for $\widehat{\beta}_0$ and $\widehat{\beta}_1$, respectively, for 10,000 randomly reshuffled years. The vertical line shows the location of our estimated $\widehat{\beta}_0$ and $\widehat{\beta}_1$ from the observed data. The exact p-values for our estimated $\widehat{\beta}_0$ and $\widehat{\beta}_1$ are 0.02 and 0.03, respectively, implying that it is highly unlikely that our main result is due to small-sample bias.

Figure 13: Robustness: Randomization placebo distributions



NOTES: Empirical distributions of β_0 (left panel) and β_1 (right panel) from 10,000 random assignment of years. Vertical lines show estimates of β_0 and β_1 from observed data using benchmark model from column 5, panel A of Table 2.

Dynamic effects There may be dynamic responses that are omitted from the static model in Section 2. For example, improvements in the forecasts of strong ENSO events since the 1980s (Chen et al., 2004) could allow the domestic share of expenditure to respond to future ENSO-driven cereal yields. Furthermore, the possibility of storing cereals could affect the cross-sectional relationship between domestic share of expenditure and agricultural productivity. Table 5 tests whether such dynamics are present. Column 1 replicates our benchmark contemporaneous-productivity specification for a sample in which t is restricted to 1962-2012, the sample period which allows for both lead and lagged yields. To examine whether agricultural trade anticipates future ENSO events, column 2 tests for the effects of lead log yields, as instrumented by ENSO in years $t+1$ and t. Future ENSO-driven yields do not influence the domestic share of expenditure. Column 3 tests for the effects of lagged log yields, as instrumented by ENSO in years $t-1$ and $t-2$. Past yields might influence the domestic share of expenditure if past productivity affects contemporary productivity through intertemporal channels such as soil depletion or if past output is stored to facilitate current consumption. Column 3 shows that past ENSO-driven yields do not influence the domestic share of expenditure. To more directly examine a potential role for stored cereals, column 4 reports our benchmark contemporaneous-productivity specification using a modified measure of the domestic share of expenditure that excludes stored cereals.[19] The estimated coefficients are similar to our benchmark estimates, suggesting that storage behavior is not responsible for the response of expenditure shares to ENSO-driven productivity variation.

5 Application: Inequality under future climate change

In this section, we demonstrate how to incorporate general-equilibrium mechanisms, such as the spatial correlation of productivity, into forecasts of economic outcomes under different productivity distributions. In particular, we illustrate how changes in the spatial correlation of productivities due to anthropogenic climate change affect future welfare inequality and quantify the bias in partial-equilibrium projections that omit spatial considerations. We begin by summarizing the standard partial-equilibrium approach to climate impact projections in Section 5.1. We then incorporate climate-driven changes in spatial correlation into this approach in Section 5.2.

5.1 Standard approach to climate impact projections

The typical objective is to project future economic impacts of climate change. For country i in year t, the economic outcome of interest y_{it} may be consumption or welfare. Because this outcome is determined in general equilibrium, it is determined by both local temperature, T_{it}, and the global spatial correlation of productivities induced by temperature, I_t. Denote that function as $h(T_{it}; I_t)$. Climate change may affect both arguments.

[19] Our standard measure of country-level domestic good consumption is output minus exports. To construct a measure that excludes storage, we further subtract the country-level difference between stored cereals in year $t-1$ and year t. This implicitly assumes that all stored cereals are domestically produced. The sample is smaller due to observations with missing storage data.

Table 5: Robustness: Dynamic effects
Outcome is log domestic share of expenditure

	(1)	(2)	(3)	(4)
$\ln A_{it}$	2.178***			2.096**
	(0.770)			(1.035)
$\ln A_{it} \times I_t$	-7.588**			-8.289*
	(3.363)			(4.634)
$\ln A_{it+1}$		0.726		
		(1.180)		
$\ln A_{it+1} \times I_{t+1}$		-2.447		
		(5.423)		
$\ln A_{it-1}$			-0.290	
			(0.911)	
$\ln A_{it-1} \times I_{t-1}$			0.480	
			(4.234)	
Remove storage?	No	No	No	Yes
Cragg-Donald F-stat	65.71	60.99	48.15	50.50
Kleibergen-Paap F-stat	17.53	11.55	11.26	10.25
Anderson-Rubin weak-id robust p-value	0	0.0768	0.282	0.00121
2nd stage sample period	1962-2012	1962-2012	1962-2012	1961-2013
Observations	5458	5457	5456	5367

NOTES: Estimates of equation (7). Column 1 reproduces benchmark model using log yields in year t instrumented by December ENSO in years t and $t-1$. Column 2 uses log yields in year $t+1$, instrumented by December ENSO conditions in years $t+1$ and t. Column 3 uses log yields in year $t-1$, instrumented by December ENSO conditions in years $t-1$ and $t-2$. Column 4 uses log yields in year t to examine effects on a measure of domestic share of expenditure that excludes cereal storage. Sample period for 2nd stage equation is 1962-2012 for columns 1-3 and 1961-2013 for column 4. All models incorporate country fixed effects, year fixed effects, and country-specific trends as included instruments. Standard errors clustered at country and continent-year levels in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

To illustrate how climate impacts are typically estimated, define $t \in [\underline{t}, \bar{t}]$ as the historical estimation period. It is common practice to first estimate the effect of historic temperature T_{it}^H on y_{it} using the following specification:

$$y_{it} = h(T_{it}^H; I_t^H) + \mathbb{X}_{it}\Pi + \nu_{it} \quad \text{for } t \in [\underline{t}, \bar{t}] \tag{12}$$

where $\mathbb{X}_{it}$ is a vector of controls. The nonlinear response function $h()$ is typically estimated using variation in local temperature T_{it}^H and implicitly depends upon the historical global spatial correlation of temperature-induced productivities, I_t^H. In a second step, a future outcome is projected by evaluating the estimated response function at a future temperature T_{it}^F. In particular, one is typically interested in predicting future outcome levels driven only by temperature changes while holding everything else fixed to the last estimation year. That projection is

$$\widehat{y}_{it} = \widehat{h}(T_{it}^F; I_t^H) + \mathbb{X}_{i\bar{t}}\widehat{\Pi} + \widehat{\nu_{i\bar{t}}} \quad \text{for } t \in [\bar{t}+1, \mathcal{T}]. \tag{13}$$

This approach assumes that climate change affects local outcome y only through local temperature T_{it}.[20] That is, using historical local temperature variation alone to estimate $h()$ implicitly assumes that climate change does not affect I_t.[21] However, as the next section shows, climate change also alters I_t. This is consequential since a country's welfare depends on its productivity both directly (autarky welfare) and indirectly through any correlation between a country's productivity and its gains from trade, as governed by the spatial correlation of productivities, I_t. Because the standard approach does not directly estimate the effect of changes in I_t, any changes in the spatial correlation of productivity due to climate change would be omitted from projections using equation (13).

5.2 Incorporating spatial correlation into climate impact projections

We now illustrate how climate projections are altered by incorporating general-equilibrium effects from climate-driven changes in the spatial correlation of productivities. Before doing so, it is important to highlight two elements of our exercise so that it not be interpreted as a literal prediction of future impacts. First, with the exception of changes in local temperature due to climate change over the 21st century, we hold everything else fixed at recent historical values. Second, we apply estimates based on past exogenous annual changes in agricultural productivity to future long-term productivity changes due to climate change. Both practices are standard in econometric estimates of climate impacts. In particular, the latter practice suggests that our welfare projections based on historical annual fluctuations may be an upper bound for future climate impacts when economies have a longer time horizon to adapt to expected climate change (Deschênes and Greenstone, 2007; Hsiang, 2016).

Using the same sample period as in Section 4, we first estimate a global nonlinear log cereal yield response function on country-year-level crop-area-weighted temperature. Specifically, for the period 1961-2013, we estimate:

$$\ln A_{it} = k(T_{it}) + \mathbb{X}_{it}\Psi + \nu_{it} \tag{14}$$

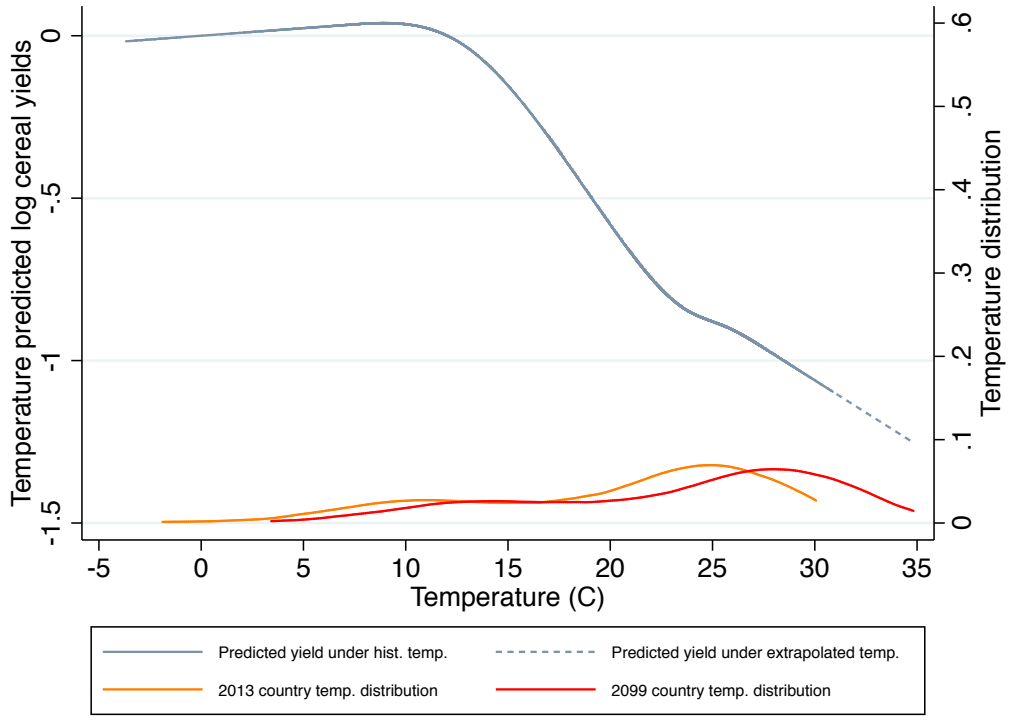
where $k()$ is a restricted cubic spline and the set of controls in $\mathbb{X}_{it}$ include quadratic precipitation terms, country fixed effects, year fixed effects, and country-specific quadratic trends.[22] Figure 14 shows the predicted cubic spline response function, $k()$, using six knots, as well as the cross-country distribution of crop-area-weighted temperatures observed in 2013, the end of our estimation period, and under a business-as-usual (Representative Concentration Pathway 8.5) climate scenario

[20] In practice, the literature often focuses on the change in outcomes attributable to the increase from $T^H_{i\bar{t}}$ to T^F_{it} and reports $\Delta\widehat{y}_{it} = h(T^F_{it}; I^H_t) - h(T^H_{i\bar{t}}; I^H_t)$.

[21] Including control variables $\mathbb{X}_{it}$ in equation (12), as commonly done in the literature, may further omit historical general-equilibrium effects. For example, if the econometrician includes regional fixed effects so that $\widehat{h}()$ is estimated using only within-region temperature variation, the identifying variation exhibits spatial correlation at a more localized scale and thus may not be applicable to a global climate change projection even if I_t is unchanged.

[22] Note that by modeling temperature effects on log cereal yield rather than an economic outcome as shown in equation (12), we are isolating the direct productivity effects of local temperature.

Figure 14: Estimated temperature response function for log cereal yields

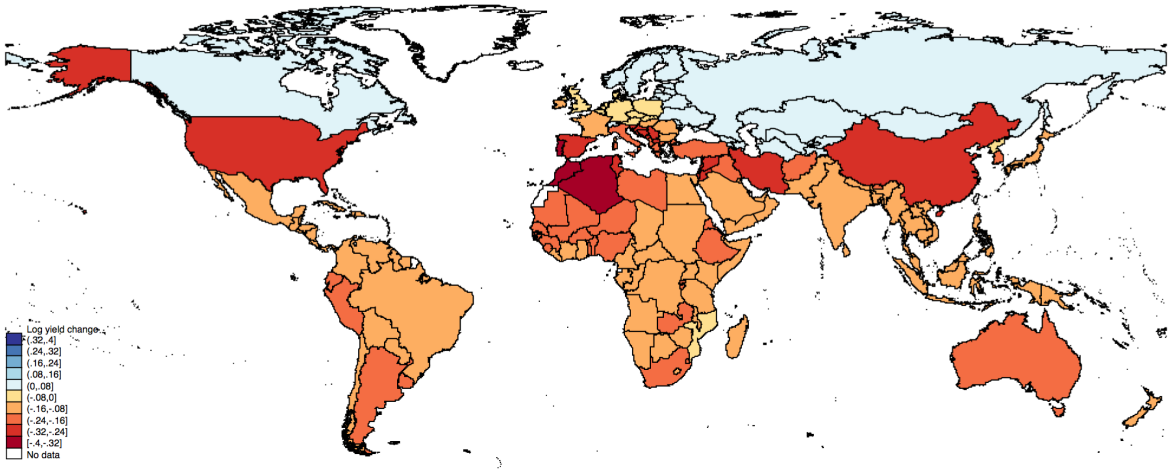


NOTES: Gray solid line shows $k()$ in equation (14), the predicted relationship between crop-area-weighted country-level temperature and log cereal yields, estimated during 1961-2013. Restricted cubic spline estimated using six knots placed along the temperature support according to Harrell (2001). Gray dashed line shows additional predicted log cereal yields using extrapolated temperature in 2099. Estimated model corresponds to column 4 of Table F.2. Orange line shows the distribution of observed country-level temperature in 2013. Red line shows country-level projected temperature in 2099 from CMIP5 ensemble mean under a business-as-usual (RCP 8.5) scenario.

in 2099 using the Coupled Model Intercomparison Project (CMIP5) multi-model ensemble mean climate projection.[23] Column 4 of Table F.2 shows that the coefficients of this functional form are jointly statistically significant with an optimal temperature around 9°C. Table F.2 also indicates that the predicted optimal temperature is relatively insensitive to the number of knots in the spline function and to the omission of precipitation controls. Figure 15 maps the projected change in country-level log cereal yields under climate change between 2013-2099, $\widehat{\ln A}_{i,2099} - \ln A_{i,2013}$. Projections of temperature-driven log cereal yields, $\widehat{\ln A}_{it}$, for each projection year come from the estimated coefficients in equation (14) together with local temperature projections from the CMIP5 multi-model ensemble mean under a business-as-usual scenario. Consistent with the nonlinear yield response function shown in Figure 14, climate change lowers agricultural productivities across much of the world with the exception of a few upper-latitude countries that experience modest gains.

[23] The Coupled Model Intercomparison Project is a coordinated effort by the climate-science community to harmonize model runs across various climate models. The average climate projection across CMIP models is known as the multi-model ensemble mean. CMIP5 was used to inform the Fifth Assessment Report of the Intergovernmental Panel on Climate Change (see Taylor, Stouffer and Meehl (2012) for details).

Figure 15: Change in log cereal yields under climate change (2013-2099)



NOTES: Map shows the projected change in country-level log cereal yields under climate change between 2013-2099. Projected change in log yields obtained after estimating equation (14) with six knots as shown in column 4 of Table F.2. Climate projections from CMIP5 ensemble mean under a business-as-usual (RCP 8.5) scenario.

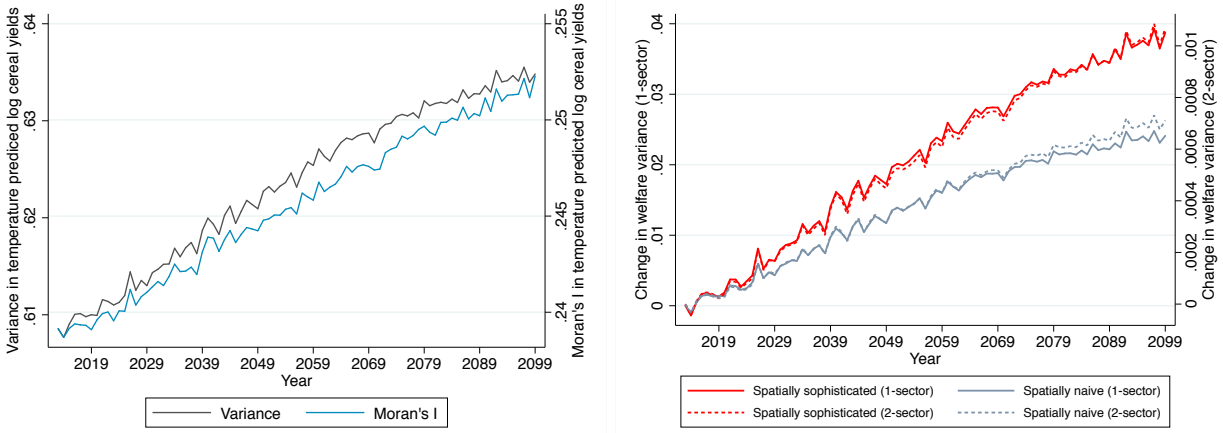
Climate change alters two important moments of the cross-country agricultural productivity distribution.[24] First, the variance of agricultural productivity increases. Because climate change raises temperatures across the globe, it shifts countries from the relatively flat portion of the yield response function depicted by the gray line in Figure 14 to the declining portion. This increases the cross-country variance, depicted by the gray line in the left panel of 16.

Second, as temperatures increase across the planet, the spatial correlation of agricultural productivities increases. This is due to the non-monotone yield response function shown in Figure 14 and the fact that surface temperatures are generally decreasing in distance to the equator. Absent climate change, mid-latitude locations experience the productivity-maximizing temperature. Locations closer to and farther from the equator both generally exhibit lower yields. As climate change increases temperatures across the globe, latitudes closer to the poles now experience the productivity-maximizing temperature, with less productive locations becoming more bunched around the equator. This bifurcation of global agriculture into high-productivity poles and a low-productivity band around the equator increases the spatial correlation of agricultural productivity. The resulting increase in Moran's I for agricultural productivity under climate change is indicated by the blue line in the left panel of Figure 16.

The change in the spatial correlation of agricultural productivity due to climate change is consequential for global inequality. To calculate the change in the variance of welfare under climate

[24] Because of surface albedo changes due to polar ice melt, there is higher warming in the poles relative to lower latitude areas. This implies that the country-specific shift in temperature due to climate change is not spatially uniform. However, this latitude-dependent gradient is relatively small compared to the average global temperature change and thus plays a second-order role in determining changes in the spatial correlation of productivity. For example, the projected temperature change over the 21st century for Gabon, the country located closest to the equator, is 3.4°C. For Finland, the country located closest to the north pole, the projected temperature change is 2.9°C.

Figure 16: Variance in climate-driven welfare over the 21st century



NOTES: Left panel shows the change in the global variance (gray line) and Moran's I (blue line) of log cereal yields over the 21st century under climate change. Right panel shows the change in the variance of welfare omitting (gray line) and including changes in the spatial correlation of log cereal yields over the same period. Solid lines under one-sector (agriculture) model. Dash lines under two-sector (agriculture and non-agriculture) model. Climate projections from CMIP5 ensemble mean under a business-as-usual (RCP 8.5) scenario.

change, we combine projected local agricultural productivity, $\widehat{\ln A_{it}}$, and global spatial correlation of productivity, $\widehat{I}_t$, with our benchmark reduced-form estimates $\widehat{\beta}_0$ and $\widehat{\beta}_1$ in Section 4 to construct future $\widehat{\ln \lambda_{iit}}$ using equation (7). This allows us to construct each component of the welfare variance in equation (2): $var(\widehat{\ln A_{it}})$, $cov(\widehat{\ln A_{it}}, \widehat{\ln \lambda_{iit}})$, and $var(\widehat{\ln \lambda_{iit}})$. To examine the importance of incorporating changes in spatial correlation, we construct a spatially "naive" projection in which the spatial correlation of temperature-predicted agricultural productivities is fixed at the 2013 value and a spatially "sophisticated" projection in which we allow the spatial correlation of productivities to be affected by climate change (see Appendix D.2 and D.3 for details).[25]

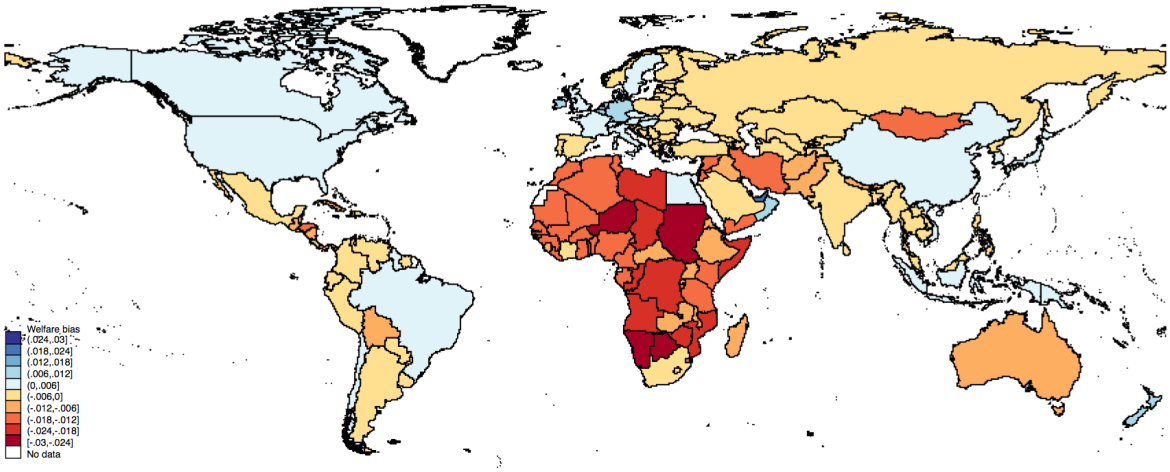
The projected change in the variance of welfare over the 21st century is shown in the right panel of Figure 16. The gray line shows the spatially naive scenario and the red line shows the spatially sophisticated scenario. The solid and dashed lines show the projected change in welfare variance under a one-sector (agriculture) and two-sector model (agriculture and non-agriculture), respectively. In a one-sector model, projections that incorporate changes in the spatial correlation of agricultural productivity due to climate change predict a 61% greater increase in welfare inequality between 2013-2099 compared to standard projections that omit changes in spatial correlation. In a two-sector model where there is adjustment in the non-agricultural sector, we predict a 50% greater increase in welfare inequality.

How do the projections differ across countries? High-productivity countries have higher welfare in the spatially sophisticated projection than the spatially naive projection. This is due to the role of spatial correlation in mediating the relationship between productivity and gains from trade. Climate

[25] We choose to fix the spatial correlation of productivity to the 2013 value in order to ensure that our spatially naive and sophisticated projections have the same starting value in 2013. Observe that is understate the bias from the standard approach to climate impact projection discussed in Section 5.1 which implicitly holds the spatial correlation of productivity fixed to the historical mean in the estimation period and not the last year of the estimation period.

change increases the spatial correlation of productivity in the spatially sophisticated projection, making the gains from trade higher in more productive countries and lower in less productive countries relative to the spatially naive projection in which spatial correlation is held fixed. Figure 17 maps the difference between the two projections for the one-sector model.[26] Most countries in Africa, South America, and many in Asia have lower gains from trade in the sophisticated projection, while parts of Europe and North American have higher gains from trade. In the simplest terms, a key feature of climate change is that it makes Ethiopia and Kenya less productive at the same time.

Figure 17: Bias in welfare change from spatially naive projection (2013-2099)

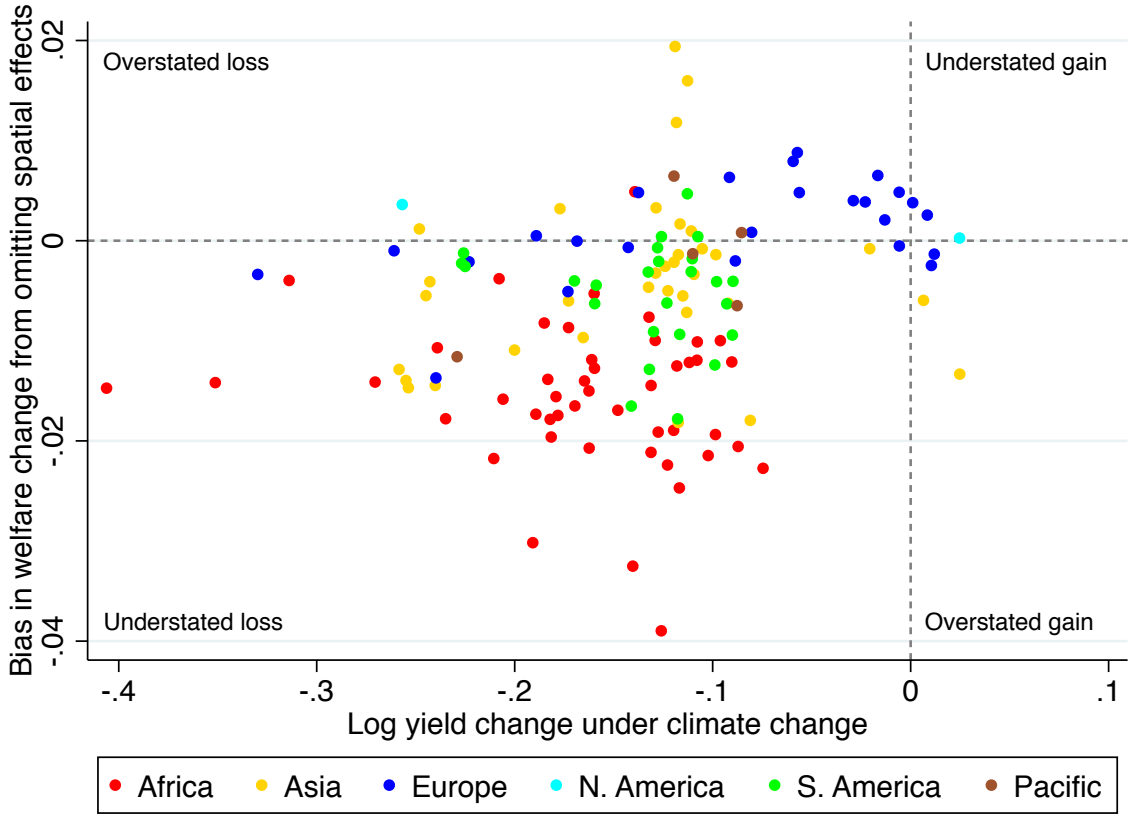


NOTES: Map shows the difference in projected 2013-2099 welfare changes between a spatially sophisticated projection (Figure D.2) and a spatially naive projection (Figure D.1). Climate projections from CMIP5 ensemble mean under a business-as-usual (RCP 8.5) scenario.

Figure 18 depicts climate-change forecasts for each country in terms of the predicted change in agricultural productivity from Figure 15 and the projection bias due to omitting the change in spatial correlation from Figure 17. The vast majority of countries suffer a decrease in agricultural productivity as a result of increased temperature. As in Figure 17, the strongest contrast in Figure 18 is between the projection biases of Africa and Europe.[27] While there is substantial heterogeneity in the productivity declines across countries within each continent, most economies in Africa, South America, and many in Asia experience welfare losses that would be considerably understated by the spatially naive projection method. By contrast, the welfare projections for most European economies are altered modestly by incorporating the increase in spatial correlation of productivity attributable to climate change, and most have welfare losses that are overstated by the spatially naive projection. Among the few economies with increased productivity under climate change, the projection bias is close to zero.

[26] See equation (D.18) at the end of Appendix D.4.
[27] See Figure E.4 for the same plot with country identifiers.

Figure 18: Country-level projection bias against change in agricultural productivity (2013-2099)



NOTES: Scatter shows country-level welfare bias from spatially naive projection (from Figure 17) plotted against change in log cereal yields under climate change (from Figure 15). Countries color-coded by continent.

6 Conclusion

This paper studies how the spatial arrangement of the global productivity distribution shapes the distribution of welfare across the globe. In particular, we predict that, holding the unconditional moments of the productivity distribution fixed, greater spatial correlation of productivity increases dispersion in welfare. The increase in spatial correlation causes high-productivity locations to enjoy larger gains from trade than they would if productivity were spatially uncorrelated. To empirically investigate this prediction, we exploit the El Niño-Southern Oscillation, a naturally occurring global climatic phenomenon, to examine how the spatial correlation of agricultural productivity governs the response of the domestic share of expenditure, a sufficient statistic for the gains from trade in a broad class of trade models, to local productivity. Using data from the past five decades of agricultural trade, we find that high-productivity locations enjoy larger gains from trade when agricultural productivity is more spatially correlated, as predicted.

The interplay between local productivity and the gains from trade is potentially important in many domains. Many economically valuable endowments – such as human capital, energy resources,

minerals, and wildlife stocks – tend to exhibit substantial spatial correlation. Furthermore, the spatial correlation of these endowments may change following the relocation of existing endowments, the discovery of new uses for existing endowments, or the discovery of new endowments. This paper provides a framework for analyzing the general-equilibrium welfare consequences of such system-wide changes to the global trade network. In settings in which welfare differences are arbitraged away by mobile factors of production, we expect that greater spatial correlation of productivity would make population density, rather than welfare per capita, more unequal across locations. Our application in this paper demonstrates that incorporating changes in the spatial correlation of productivities alters predictions about global welfare inequality when forecasting the consequences of anthropogenic climate change.

References

Allen, Treb, and Costas Arkolakis. 2014. "Trade and the Topography of the Spatial Economy." *The Quarterly Journal of Economics*, 129(3): 1085–1140.

Allen, Treb, Costas Arkolakis, and Yuta Takahashi. 2017. "Universal Gravity."

Amemiya, Takeshi. 1974. "The nonlinear two-stage least-squares estimator." *Journal of Econometrics*, 2(2): 105–110.

Anderson, James E., and Eric van Wincoop. 2003. "Gravity with Gravitas: A Solution to the Border Puzzle." *American Economic Review*, 93(1): 170–192.

Anderson, T. W., and Herman Rubin. 1949. "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations." *The Annals of Mathematical Statistics*, 20(1): 46–63.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton University Press.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. "New Trade Models, Same Old Gains?" *American Economic Review*, 102(1): 94–130.

Bleakley, Hoyt. 2007. "Disease and development: evidence from hookworm eradication in the American South." *The Quarterly Journal of Economics*, 122(1): 73–117.

Brock, William, Gustav Engström, and Anastasios Xepapadeas. 2014. "Spatial climate-economic models in the design of optimal climate policies across locations." *European Economic Review*, 69: 78 – 103. Sustainability and Climate Change: From Theory to Pragmatic Policy.

Burchfield, Marcy, Henry G Overman, Diego Puga, and Matthew A Turner. 2006. "Causes of sprawl: A portrait from space." *The Quarterly Journal of Economics*, 121(2): 587–633.

Burgess, Robin, Olivier Deschenes, Dave Donaldson, and Michael Greenstone. 2013. "The Unequal Effects of Weather and Climate Change:Evidence from Mortality in India." *mimeo.*

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel. 2015. "Global non-linear effect of temperature on economic production." *Nature*, 527(7577): 235–239.

Caliendo, Lorenzo, and Fernando Parro. 2015. "Estimates of the Trade and Welfare Effects of NAFTA." *The Review of Economic Studies*, 82(1): 1–44.

Cane, Mark A., and Stephen E. Zebiak. 1985. "A Theory for El Niño and the Southern Oscillation." *Science*, 228(4703): 1085–1087.

Chen, Dake, Mark A. Cane, Alexey Kaplan, Stephen E. Zebiak, and Daji Huang. 2004. "Predictability of El Nino over the past 148 years." *Nature*, 428(6984): 733—736.

Chiang, John C.H., and Adam H. Sobel. 2002. "Tropical Tropospheric Temperature Variations Caused by ENSO and Their Influence on the Remote Tropical Climate." *Journal of Climate*, 15(18): 2616–2631.

Costinot, Arnaud, and Andrés Rodríguez-Clare. 2014. "Trade Theory with Numbers: Quantifying the Consequences of Globalization." Vol. 4 of *Handbook of International Economics*, Chapter 4, 197–261. Elsevier.

Costinot, Arnaud, Dave Donaldson, and Cory Smith. 2016. "Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World." *Journal of Political Economy*, 124(1): 205–248.

Cragg, John G., and Stephen G. Donald. 1993. "Testing Identifiability and Specification in Instrumental Variable Models." *Econometric Theory*, 9(2): 222–240.

Currie, Janet, and Jonathan Gruber. 1996. "Health Insurance Eligibility, Utilization of Medical Care, and Child Health." *The Quarterly Journal of Economics*, 111(2): 431–466.

Cutler, David M., and Jonathan Gruber. 1996. "Does Public Insurance Crowd out Private Insurance?" *The Quarterly Journal of Economics*, 111(2): 391–430.

Dahl, Gordon B., and Lance Lochner. 2012. "The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit." *American Economic Review*, 102(5): 1927–56.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *American Economic Journal: Macroeconomics*, 4(3): 66–95.

Deschênes, Olivier, and Michael Greenstone. 2007. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather." *American Economic Review*, 97(1): 354–385.

Desmet, Klaus, and Esteban Rossi-Hansberg. 2015. "On the spatial economic impact of global warming." *Journal of Urban Economics*, 88: 16 – 37.

Donaldson, Dave, and Richard Hornbeck. 2016. "Railroads and American Economic Growth: A "Market Access" Approach." *The Quarterly Journal of Economics*, 131(2): 799–858.

Easterly, William, and Ross Levine. 2003. "Tropics, germs, and crops: how endowments influence economic development." *Journal of monetary economics*, 50(1): 3–39.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica*, 70(5): 1741–1779.

Fan, Jianqing, and Irene Gijbels. 1996. *Local Polynomial Modelling and Its Applications.* Chapman & Hall.

Gruber, Jon, and Emmanuel Saez. 2002. "The elasticity of taxable income: evidence and implications." *Journal of Public Economics*, 84(1): 1 – 32.

Harrell, Frank. 2001. *Regression modeling strategies: with applications to linear models, logistic and ordinal regression, and survival analysis.* Springer.

Head, Keith, and Thierry Mayer. 2004. "The empirics of agglomeration and trade." , ed. J. V. Henderson and J. F. Thisse Vol. 4 of *Handbook of Regional and Urban Economics*, Chapter 59, 2609–2669. Elsevier.

Head, Keith, and Thierry Mayer. 2014. "Gravity Equations: Workhorse,Toolkit, and Cookbook." Vol. 4 of *Handbook of International Economics*, Chapter 3, 131–195. Elsevier.

Hornbeck, Richard. 2012. "Nature versus Nurture: The Environment's Persistent Influence through the Modernization of American Agriculture." *American Economic Review Papers and Proceedings*, 102(3): 245–249.

Houser, T., R. Kopp, S.M. Hsiang, M. Delgado, A.S. Jina, K. Larsen, M. Mastrandrea, S. Mohanand R. Muir-Wood, D.J. Rasmussen, J. Rising, and P. Wilson. 2015. *Economic risks of climate change: an American prospectus.* Columbia University Press.

Hsiang, Solomon. 2016. "Climate Econometrics." *Annual Review of Resource Economics*, 8(1): 43–75.

Hsiang, Solomon M., and Kyle C. Meng. 2015. "Tropical Economics." *American Economic Review: Papers and Proceedings*, 105(5): 257–61.

Hsiang, Solomon M, Kyle C Meng, and Mark A Cane. 2011. "Civil conflicts are associated with the global climate." *Nature*, 476(7361): 438–441.

Hsiang, Solomon M., Marshall Burke, and Edward Miguel. 2013. "Quantifying the Influence of Climate on Human Conflict." *Science*, 341(6151).

Kaplan, A., M. Cane, Y. Kushnir, A. Clement, M. Blumenthal, and B. Rajagopalan. 1998. "Analyses of global sea surface temperature 1856-1991." *Journal of Geophysical Research*, 103: 18,567–18,589.

Kelejian, Harry H. 1971. "Two-Stage Least Squares and Econometric Systems Linear in Parameters but Nonlinear in the Endogenous Variables." *Journal of the American Statistical Association*, 66(334): 373–374.

Kleibergen, Frank, and Richard Paap. 2006. "Generalized reduced rank tests using the singular value decomposition." *Journal of Econometrics*, 133(1): 97 – 126.

Krugman, Paul. 1980. "Scale Economies, Product Differentiation, and the Pattern of Trade." *American Economic Review*, 70(5): 950–59.

Krusell, Per, and Anthony A. Smith. 2016. "Climate Change around the World." *mimeo.*

Legates, D. R., and C. J. Willmott. 1990*a*. "Mean Seasonal and Spatial Variability Global Surface Air Temperature." *Theoretical and Applied Climatology*, 41: 11–21.

Legates, D. R., and C. J. Willmott. 1990*b*. "Mean Seasonal and Spatial Variability in Gauge-Corrected, Global Precipitation." *International Journal of Climatology*, 10: 111–127.

Moffitt, Robert, and Mark Wilhelm. 2000. "Taxation and Labor Supply Decisions of the Affluent." In *Does Atlas Shrug? Economic Consequences of Taxing the Rich.* , ed. Joel Slemrod, 193–234. Russell Sage Foundation.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–708.

Nordhaus, William D. 2006. "Geography and macroeconomics: New data and new findings." *Proceedings of the National Academy of Sciences of the United States of America*, 103(10): 3510–3517.

Nunn, Nathan, and Diego Puga. 2012. "Ruggedness: The blessing of bad geography in Africa." *Review of Economics and Statistics*, 94(1): 20–36.

Ramankutty, Navin, Amato T. Evan, Chad Monfreda, and Jonathan A. Foley. 2008. "Farming the planet: 1. Geographic distribution of global agricultural lands in the year 2000." *Global Biogeochemical Cycles*, 22.

Redding, Stephen, and Anthony J. Venables. 2004. "Economic geography and international inequality." *Journal of International Economics*, 62(1): 53–82.

Sachs, Jeffrey D, and Andrew M Warner. 1997. "Fundamental sources of long-run growth." *The American economic review*, 87(2): 184–188.

Sarachik, Edward S., and Mark A. Cane. 2010. *The El Niño-Southern Oscillation Phenomenon.* Cambridge University Press.

Schlenker, Wolfram, and Michael J Roberts. 2009. "Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change." *Proceedings from the National Academy of Science*, 106(37): 15594–8.

Schlenker, Wolfram, and W. Reed Walker. 2015. "Airports, Air Pollution, and Contemporaneous Health." *The Review of Economic Studies.*

Simonovska, Ina, and Michael E. Waugh. 2014. "The elasticity of trade: Estimates and evidence." *Journal of International Economics*, 92(1): 34–50.

Taylor, Karl E., Ronald J. Stouffer, and Gerald A. Meehl. 2012. "An Overview of CMIP5 and the Experiment Design." *Bulletin of the American Meteorological Society*, 93(4): 485–498.

Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross Section and Panel Data.* MIT Press.

Appendix

A Theory appendix

A.1 General economic environment

This appendix section provides details of perfect-competition results presented in Section 2 and extends them to the Krugman (1980) model of monopolistic competition.

A.1.1 Perfect-competition microfoundations

Production function. In the Armington model, each country i produces a distinct variety using a linear production technology such that one unit of labor yields A_i units of output. Under perfect competition, its price is w_i/A_i and thus the Dixit-Stigliz price index is $P_j = \left(\sum_{i=1}^{N} \left(w_i \tau_{ij}/A_i\right)^{1-\sigma}\right)^{1/(1-\sigma)}$. In the Eaton and Kortum (2002) model, there is a continuum of varieties, and each country's efficiency in producing them follows a Fréchet distribution with location parameter T_i and dispersion parameter ϑ.

Gravity equation. Written in terms of expenditure shares, the gravity equation is

$$\lambda_{ij} = \frac{\chi_i \left(\tau_{ij} w_i\right)^{-\epsilon}}{\sum_{l=1}^{N} \chi_l \left(\tau_{lj} w_l\right)^{-\epsilon}}$$

In the Armington model, $\epsilon = \sigma - 1$ and $\chi_i = A_i^{\epsilon}$. In the Eaton-Kortum model without intermediate inputs, $\epsilon = \vartheta$ and $\chi_i = T_i$. Thus, the equilibrium trade flows associated with a productivity distribution $\{A_i\}_i$ and trade elasticity ϵ in the Armington model are equal to the equilibrium trade flows for an Eaton-Kortum model in which efficiency distributions have location parameters $T_i = A_i^{1/\epsilon}$.

Welfare. Equation (1) is an immediate consequence of the main result in Arkolakis, Costinot and Rodríguez-Clare (2012). They show that, in a broad class of models, the gains from trade relative to autarky are equal to $-\frac{1}{\epsilon} \ln \lambda_{jj}$. In our theoretical environment, autarky welfare is equal to $\ln A_j + \gamma$, as implied by equation (1). γ is a function of structural parameters. In the Armington model with symmetric preferences, γ is zero. In the Eaton and Kortum (2002) model, $\gamma = \left[\Gamma\left(1 + \frac{1-\sigma}{\theta}\right)\right]^{1/(1-\sigma)}$ where Γ is the gamma function.

A.1.2 Krugman (1980) microfoundations

We now discuss the case of monopolistic competition with homogeneous firms, in which the measure of varieties available in equilibrium is endogenously determined. Consider a many-country version of the Krugman (1980) model in which fixed costs f_j and marginal costs c_j may vary across countries. Denote the equilibrium number of homogeneous firms producing in country j by n_j.

Gravity equation. In the free-entry equilibrium, $\epsilon = \sigma - 1$, $n_i = L_i/\left(\sigma f_i\right)$, and $\chi_i = n_i c_i^{-\epsilon}$.

Welfare. In this setting, real per capita consumption is

$$\ln\left(C_j/L_j\right) = \frac{1}{\epsilon}\ln n_j - \ln c_j + \ln\left(\frac{\sigma-1}{\sigma}\right) - \frac{1}{\epsilon}\ln\lambda_{jj}$$

If population size L_j and fixed costs f_j are country-invariant ($n_j = n\ \forall j$) and we interpret productivity as shifting marginal costs, $A_j = c_j^{-1}$, then this is equivalent to equation (1) with $\gamma = \frac{1}{\epsilon}\ln n + \ln\left(\frac{\sigma-1}{\sigma}\right)$. If population size L_j and marginal costs c_j are country-invariant and we interpret productivity as shifting fixed costs, $A_j = f_j^{-1/\epsilon}$, then this is equivalent to equation (1) with $\gamma = \frac{1}{\epsilon}\ln\left(L/\sigma\right) + \ln\left(\frac{\sigma-1}{\sigma c}\right)$.

In the case of countries with heterogeneous population sizes, equation (1) must be extended to replace γ with a location-specific γ_j that depends on structural parameters other than productivity. For example, if productivity shifts marginal costs, $A_j = c_j^{-1}$, then we obtain

$$\ln\left(\frac{C_j}{L_j}\right) = \ln A_j + \gamma_j - \frac{1}{\epsilon}\ln\lambda_{jj}, \tag{1'}$$

where $\gamma_j = \frac{1}{\epsilon}\ln n_j + \ln\left(\frac{\sigma-1}{\sigma}\right)$. If we assume that $cov(\ln A_i^u, \gamma_i) = cov(\ln A_i^c, \gamma_i)$, we obtain an extension of equation (3):

$$\begin{aligned} var\left(\ln C_{i,c}/L_i\right) - var\left(\ln C_{i,u}/L_i\right) = &-\frac{2}{\epsilon}\left[cov\left(\ln A_{i,c}, \ln\lambda_{ii,c}\right) - cov\left(\ln A_{i,u}, \ln\lambda_{ii,u}\right)\right] \\ &-\frac{2}{\epsilon}\left[cov\left(\gamma_i, \ln\lambda_{ii,c}\right) - cov\left(\gamma_i, \ln\lambda_{ii,u}\right)\right] \\ &+\frac{1}{\epsilon^2}\left[var\left(\ln\lambda_{ii,c}\right) - var\left(\ln\lambda_{ii,u}\right)\right]. \end{aligned} \tag{3'}$$

A.1.3 The case of symmetric trade costs

Starting from the equilibrium system of equations and using the assumption of symmetric trade costs ($\tau_{ij} = \tau_{ji}$):

$$\begin{aligned} Y_i = w_i L_i &= \sum_j \left(\frac{w_i}{A_i}\right)^{-\epsilon} \tau_{ij}^{-\epsilon}\frac{w_j L_j}{\Phi_j} = \left(\frac{w_i}{A_i}\right)^{-\epsilon}\Omega_i \\ \Rightarrow \frac{w_i}{A_i} &= \left(\frac{\Omega_i}{A_i L_i}\right)^{\frac{1}{\epsilon+1}} \\ \Rightarrow \Phi_i &= \sum_j \tau_{ji}^{-\epsilon}\left(\frac{w_j}{A_j}\right)^{-\epsilon} = \sum_j \tau_{ji}^{-\epsilon}\left(A_j L_j/\Omega_j\right)^{\frac{\epsilon}{\epsilon+1}} \\ &= \sum_j \tau_{ji}^{-\epsilon}\left(A_j L_j/\Phi_j\right)^{\frac{\epsilon}{\epsilon+1}} \end{aligned}$$

The last equality exploits the fact that we can normalize incomes such that $\Phi_i = \Omega_i$ when trade is balanced and $\tau_{ij}^{-\epsilon}$ is symmetric, as established in Anderson and van Wincoop (2003) and Head and Mayer (2014).

Since $\lambda_{ii} = \frac{(w_i/A_i)^{-\epsilon}}{\Phi_i}$, the above expression for w_i/A_i and the assumption of symmetric trade costs yield the result that

$$\lambda_{ii} = \frac{(\Omega_i/(A_i L_i))^{\frac{-\epsilon}{\epsilon+1}}}{\Phi_i} = (A_i L_i)^{\frac{\epsilon}{\epsilon+1}} \Phi_i^{-\frac{2\epsilon+1}{\epsilon+1}}$$

Thus, in the case in which $L_i = L \; \forall i$, $var(\ln \lambda_{ii}) = \frac{\epsilon}{\epsilon+1} cov(\ln A_i, \ln \lambda_{ii}) - \frac{1+2\epsilon}{1+\epsilon} cov(\ln \Phi_i, \ln \lambda_{ii})$.

A.2 Illustrative examples

This appendix section details the construction and parameterization of the illustrative examples presented in Section 2.2.

A.2.1 Circular geography

Figures 1 and 2 present an illustrative example in which we parameterize A_i and τ_{ij} on a circular geography. As stated in the main text, $N = 50$, $\epsilon = 1$, and $L_i = 1 \; \forall i$. Countries have locations given by $l_i = \frac{\pi}{N}(2i - 1 - N)$ for $i = 1, \ldots, N$. Productivity $\ln A_i$ has a mean value of 10 and follows a sinewave with amplitude 1 and frequency θ.[28] Bilateral trade costs are given by $\ln \tau_{ij} = .8 \ln \|l_i - l_j\|$, where $\|l_i - l_j\|$ is the distance between locations i and j on the circle.

Figures 1 and 2 depict demeaned distributions. Table A.1 reports the means and variances of countries' welfare per capita under autarky and trade.

Table A.1: Outcomes for one-sector sine-wave economy

Frequency of $\ln A$ sine wave (θ)	1	2	3	4
Autarky welfare ($\ln A$) mean	10	10	10	10
Autarky welfare ($\ln A$) variance	0.510204	0.510204	0.510204	0.510204
Trading-equilbrium welfare ($\ln C/L$) mean	12.2654	12.2769	12.2807	12.2836
Trading-equilbrium welfare ($\ln C/L$) variance	0.298203	0.226274	0.203006	0.184882

Figure 3 introduces heterogeneous sizes to the parameterization used in figures 1 and 2. Size $\ln L_i$ is the sum of a sinewave with mean value 10, amplitude 1, and frequency $\theta_L = 1$ and Gaussian noise $\sim \mathcal{N}(0, 1)$.

Figure 4 depicts the expenditure-productivity relationship in a circular geography with equal-sized countries for three selected productivity vectors. These productivity distributions differ only in their spatial correlation. They were generated by shuffling a vector $\ln A_0 \sim N(0, 1)$.

Figure 6 depicts a two-sector sine-wave economy with two symmetric sectors that differ only in their sine-wave frequency. Table A.2 reports the means and variances of countries' welfare per

[28] The standard deviation of a sine wave is proportional to its amplitude and independent of its frequency. This is true for both the function and our N-point discretization.

Table A.2: Outcomes for two-sector sine-wave economy

Frequency of $\ln A_1$ sine wave (θ_1)	1	2	3	4
Autarky welfare ($\frac{1}{2}\ln A_1 + \frac{1}{2}\ln A_2$) mean	10	10	10	10
Autarky welfare ($\frac{1}{2}\ln A_1 + \frac{1}{2}\ln A_2$) variance	0.255102	0.255102	0.255102	0.255102
Trading-equilbrium welfare ($\ln C/L$) mean	12.3610	12.3699	12.3721	12.3736
Trading-equilbrium welfare ($\ln C/L$) variance	0.114255	0.097649	0.092189	0.087935

capita under autarky and trade. Compared to Table A.1, the variance of autarky welfare is lower because autarky welfare is the simple average of two sectors' (orthogonal) productivities with the same mean and variance. The mean trading-equilibrium welfare is higher (the gains from trade are larger) in the multi-sector case due to gains from specialization according to comparative advantage. This additional margin of adjustment also dampens the degree to which greater spatial correlation of productivity in one of the two sectors affects the variance of welfare in the trading equilibrium, but our main prediction still obtains in this multi-sector setting with orthogonal productivities.

A.2.2 Randomly generated geography

In Figure 5, we examine how the relationship between productivity and domestic share of expenditure depends on the spatial correlation of productivity in a randomly generated geography. We draw $N = 50$ locations' coordinates on a plane $l_j \in \mathbb{R}^2$ from a standard normal distribution. We generate bilateral, distance-related trade costs using $\ln \tau_{ij} = \tau \ln(1 + d_{ij})$, where distance $d_{ij} = \|l_i - l_j\|$ is the Euclidean norm on $\mathbb{R}^2$ and τ is a positive scaling factor. countries are of equal size, $L_j = 1\ \forall j$. The trade elasticity $\epsilon = 4$. The spatial weight matrix W we use to compute Moran's I is constructed using inverse distances, $w_{ij} = \frac{1}{d_{ij}}$ with $w_{ii} = 0$ by convention.

To construct M productivity distributions that have identical first and second moments and different spatial covariances, we employ a procedure that makes use of the eigenvectors of a transformation of the spatial weight matrix. Let I denote the identity matrix and consider $J = ee^\intercal$, where e is the constant vector of ones. The object of interest is the matrix $\Pi = CWC$ where C corresponds to the centering matrix: $C = I - \frac{1}{N}J$ and W is our spatial weight matrix. It can be shown that the upper and lower bounds of Moran's I are given by: $\lambda_{max}\frac{N}{W_0}$ and $\lambda_{min}\frac{N}{W_0}$ where λ_{max} and λ_{min} denote the largest and the smallest eigenvalues of Π respectively and $W_0 = \sum_{i=1}^{N}\sum_{j=1}^{N} W_{ij}$. More generally, Moran's I will be equal to $\lambda_i\frac{N}{W_0}$ when evaluated at the ith eigenvector. It can also be shown that the non-constant eigenvectors of Π are centered and have identical second moments.

With these eigenvectors in hand, we apply a translation $T_c : v \to v + c$ with $c \in \mathbb{R}^2_+$ to the non-constant eigenvectors of Π. We employ linear combinations of the strictly positive vectors produced by this transformation to produce a set of M productivity distributions denoted $(A_1, \ldots, A_M)$. By construction, these productivity distributions have identical first and second moments but vary in their spatial covariance as measured by Moran's I.

A.3 Multi-sector case

This appendix section provides details of the multi-sector economic environment summarized in Section 2.3.

Preferences. Individuals in country j have preferences that are Cobb-Douglas over sectors $s = 1, \ldots, S$ and constant elasticity of substitution (CES) within sectors. Thus, the relevant price indices are

$$P_j = \prod_{s=1}^{S} P_{js}^{\alpha_{js}} \text{ and } P_{js} = \left(\int_{\omega_s \in \Omega_s} p_j(\omega_s)^{1-\sigma_s} d\omega_s \right)^{1/(1-\sigma_s)},$$

where $\alpha_{js} \geq 0$ are expenditure shares ($\sum_{s=1}^{S} \alpha_{js} = 1$) and σ_s are sectoral elasticities of substitution across varietes.

Production. Productivity in country j in sector s is A_{js}.

Trade costs. Selling one unit to j from i in sector s requires producing $\tau_{ijs} \geq 1$ units, with $\tau_{iis} = 1$.

Gravity equation. Denote sales from i to j in sector s by X_{ijs} and j's total expenditure by $X_j \equiv \sum_{i=1}^{N} \sum_{s=1}^{S} X_{ijs}$. Across sectors, Cobb-Douglas preferences cause optimizing consumers to spend α_{js} of their total expenditure in sector s, $X_{js} = \alpha_{js} X_j$. Within each sector, CES preferences result in the share of expenditure by j on goods from i in sector s taking the form of a gravity equation:

$$\lambda_{ijs} = \frac{X_{ijs}}{X_{js}} = \frac{\chi_{ijs} \left(\tau_{ijs} w_i\right)^{-\epsilon_s}}{\sum_{l=1}^{N} \chi_{ljs} \left(\tau_{ljs} w_l\right)^{-\epsilon_s}} = \frac{\chi_{ijs} \left(\tau_{ijs} w_i\right)^{-\epsilon_s}}{\Phi_{js}}.$$

Equilibrium. In a competitive equilibrium, labor-market clearing, goods-market clearing, and budget constraints are satisfied such that total income $Y_i = w_i L_i$ and sectoral income $Y_{is} = w_i L_{is}$ satisfy $Y_{is} = \sum_{j=1}^{N} X_{ijs}$, $Y_i = \sum_{s=1}^{S} Y_{is}$, and $X_{is} = \alpha_{is} Y_i$ for all countries. The equilibrium system of equations is

$$Y_{is} = \sum_{j=1}^{N} \lambda_{ijs} \alpha_{js} \sum_{s'=1}^{S} Y_{js'}.$$

In this environment, real consumption per capita is

$$\ln\left(\frac{C_j}{L_j}\right) = \sum_{s=1}^{S} \alpha_{js} \left(\ln A_{js} + \gamma_s - \frac{1}{\epsilon_s} \ln \lambda_{jjs} \right).$$

The first two terms, $\sum_{s=1}^{S} \alpha_{js} \left(\ln A_{js} + \gamma_s\right)$, are per capita welfare under autarky, and the final term, $-\sum_{s=1}^{S} \frac{\alpha_{js}}{\epsilon_s} \ln \lambda_{jjs}$, summarizes the gains from trade.

Dispersion in per capita welfare across countries thus depends on the exogenous variation in productivities A_{js} and the endogenous variation in domestic shares of expenditure λ_{jjs}. For the sake of expositional brevity, assume that $\sum_{s=1}^{S} \alpha_{js} \gamma_s$ is a constant. In that case, the variance of

per capita welfare is

$$var\left(\ln\left(\frac{C_j}{L_j}\right)\right) = var\left(\sum_{s=1}^{S}\alpha_{js}\ln A_{js}\right) + var\left(\sum_{s=1}^{S}\frac{\alpha_{js}}{\epsilon_s}\ln\lambda_{jjs}\right) - 2\sum_{s=1}^{S}\sum_{s'=1}^{S}\frac{\alpha_s\alpha_{s'}}{\epsilon_{s'}}cov\left(\ln A_{js},\ln\lambda_{jjs'}\right)$$

To examine the role of spatial correlation, consider two productivity distributions – a correlated state c and an uncorrelated state u. Assume that unconditional variance of the productivity distributions is the same, $var\left(\sum_{s=1}^{S}\alpha_{js}\ln A_{js}^{c}\right) = var\left(\sum_{s=1}^{S}\alpha_{js}\ln A_{js}^{u}\right)$. The difference in welfare dispersion then depends on the covariance of productivities and domestic shares of expenditure, both within and across sectors, and between domestic shares of expenditure.

$$\begin{aligned} var\left(\ln\left(C_{j,c}/L_j\right)\right) - var\left(\ln\left(C_{j,u}/L_j\right)\right) = {} & 2\sum_{s=1}^{S}\sum_{s'=1}^{S}\frac{\alpha_s\alpha_{s'}}{\epsilon_{s'}}\left\{cov\left(\ln A_{js,u},\ln\lambda_{jjs',u}\right) - cov\left(\ln A_{js,c},\ln\lambda_{jjs',c}\right)\right\} \\ & - \sum_{s=1}^{S}\sum_{s'=1}^{S}\frac{\alpha_s\alpha_{s'}}{\epsilon_s\epsilon_{s'}}\left\{cov\left(\ln\lambda_{jjs,u},\ln\lambda_{jjs',u}\right) - cov\left(\ln\lambda_{jjs,c},\ln\lambda_{jjs',c}\right)\right\} \end{aligned}$$

Just as in the single-sector case, for typical values of the sectoral trade elasticities, $2\frac{\alpha_s\alpha_{s'}}{\epsilon_{s'}}$ is an order of magnitude larger than $\frac{\alpha_s\alpha_{s'}}{\epsilon_s\epsilon_{s'}}$. Thus, the difference in welfare dispersion is governed by the $cov\left(\ln A_{js},\ln\lambda_{jjs'}\right)$ terms, provided that the block of $cov\left(\ln A_{js},\ln\lambda_{jjs'}\right)$ terms is the same order of magnitude as the block of $cov\left(\ln\lambda_{jjs},\ln\lambda_{jjs'}\right)$ terms.

B Data sources

Agricultural data Country-year cereal[29] output (in metric tons), yield (in metric tons per hectare), export quantity (in metric tons), export unit value (in current USD), and change in storage for 1961-2013 obtained from FAOstat.[30] Agriculture domestic share of expenditure constructed as the ratio of domestic expenditure over total expenditure. Agriculture domestic expenditure is the product of consumption of domestic good (i.e. output less export quantity) and export unit value. Agriculture total expenditure is the sum of domestic expenditure and import value (i.e. product of import quantity and import unit value). Bilateral value of cereal trade flows (exporter-importer-year value in current USD) for 1986-2013 obtained from FAOstat.

Non-agricultural data Country-year total output (obtained in current local currency and converted to current USD) for 1970-2013 from U.N. National Accounts Official Country Data.[31] Non-agriculture output value is the difference between total output value and agriculture output value. Country-year total import and export values (in current USD) for 1970-2013 from the World Bank's

[29] Aggregated across 8 major cereals: barley, maize, millet, oats, rice, rye, sorghum, and wheat.
[30] Available at http://www.fao.org/faostat/en/#data
[31] Available at http://data.un.org/Data.aspx?d=SNA&f=group_code%3a206

World Development Indicators (WDI).[32] Non-agriculture trade values is the difference between total trade values and agriculture trade values. Non-agriculture domestic share of expenditure constructed as the ratio of domestic expenditure over total expenditure. Non-agriculture domestic expenditure is the difference between non-agriculture output value and non-agriculture export value. Non-agriculture total expenditure is the sum of non-agriculture domestic expenditure and non-agriculture import value.

ENSO index Annual ENSO variations can be detected using different indices, with the most commonly used being equatorial Pacific sea surface temperature (SST) anomalies. Because ENSO activity peaks during December, we utilize 1960-2013 December values of the monthly Kaplan NINO3.4 index which averages SST over the area 5°N - 5°S, 170°W - 120°W (Kaplan et al., 1998).[33]

Historical temperature and precipitation Global temperature (in degrees centigrade) and precipitation (in mm/month) variable constructed from monthly gridded global weather data at a 0.5° latitude by 0.5° longitude resolution from the Center for Climatic Research at the University of Delaware (Legates and Willmott, 1990*a*,*b*). 1960-2013 monthly data was first spatially aggregated from pixel to country-level using cross-sectional crop area weights from Ramankutty et al. (2008). Annual values constructed by averaging January-December monthly values.

Projected temperature under climate change Global multi-model ensemble mean temperature temperature (in degrees centigrade) from monthly gridded global data at the 2.5° latitude by 2.5° longitude resolution from the Coupled Model Intercomparison Project version 5 (CMIP5).[34] 2014-2099 monthly data was first spatially aggregated from pixel to country-level using cross-sectional crop area weights from Ramankutty et al. (2008). Annual values values constructed by averaging monthly values.

Geography Country latitude and longitude are defined as crop area-weighted average using the global cross-sectional distribution of crop area in 2000 (Ramankutty et al., 2008). Great-circle distances between these country centroids are computed using the haversine formula.

Global oil prices Monthly West Texas Intermediate crude oil spot price obtained from the St. Louis Federal Reserve.[35] 1961-2013 annual values constructed by averaging January-December monthly values.

[32] Available at http://data.worldbank.org/data-catalog/world-development-indicators
[33] Available at http://iridl.ldeo.columbia.edu/SOURCES/.Indices/.nino/.EXTENDED/
[34] Available at https://climexp.knmi.nl/selectfield_cmip5.cgi
[35] Available at https://fred.stlouisfed.org/series/WTISPLC

C Additional empirical results

C.1 Gravity estimates for cereals trade

The theoretical model of Section 2 assumes that trade in cereals follows a gravity specification for exporter i, importer j and year t:

$$\ln X_{ijt} = -\epsilon \ln \tau_{ijt} + \ln\left(\frac{\chi_{it}}{w_{it}^{\epsilon}}\right) + \ln\left(\frac{X_{jt}}{\Phi_{jt}}\right)$$

Estimating this log-linear equation using data on bilateral trade flows in cereals yields patterns similar to those found in aggregate data (e.g., Head and Mayer 2014). Bilateral cereal trade flows are available from FAOStat starting in 1986. As is standard, we estimate in our panel setting using bilateral distance as a (time-invariant) source of variation in bilateral trade costs while employing exporter-year (it) fixed effects and importer-year jt fixed effects. Table C.1 reports the results. While the estimated distance coefficient of -1.5 differs from the coefficient of -1.0 typically estimated for aggregate trade flows, the regression exhibits the typical explanatory power, accounting for the majority of the variation in cereal trade flows. In addition, 61% of countries that trade cereals in a given year both import and export cereals that year. At the level of importer-exporter pairs, 18% of trading pairs sell cereal in both directions. Thus, cereals are far from homogeneous commodities, and international trade in cereals is well described by the gravity specification.

Table C.1: Gravity regression for international trade in cereals

Outcome is log import value	
	(1)
$\ln dist_{it}$	-1.519***
	(0.100)
R-squared	0.545
Country-level intra-industry trade share	0.614
Bilateral intra-industry trade share	0.236
Observations	59,927

NOTES: OLS estimates of gravity model for bilateral (importer-reported) trade value during 1986-2013. All models include importer-year and exporter-year fixed effects. Intra-industry trade shares are fraction of country-year and country-pair-year observations with positive exports and imports, conditional on positive exports or imports. Standard errors clustered at the importer and exporter levels in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

C.2 Non-agriculture domestic share of expenditure

In order to perform welfare projections under climate change in a multi-sector setting (see Appendix D.3), we estimate how the non-agriculture domestic share of expenditure responds to local agricultural productivity and the global spatial correlation of agricultural productivity. Specifically, we conduct the same IV procedure detailed in Section 4.2, but replace the outcome variable in equation 7 with non-agricultural domestic share of expenditure:

$$\ln \lambda_{iimt} = \phi_0 \ln A_{it} + \phi_1 \ln A_{it} I_t + \mathbb{Z}_{it} \Pi_m + \mu_{imt}$$

Table C.2 shows estimates of ϕ_0 and ϕ_1 corresponding to the agricultural domestic share of expenditure benchmark result in column 5, panel A of Table 2.

Table C.2: Non-agricultural domestic share of expenditure and spatial correlation in productivity

Outcome is log domestic share of expenditure	
	(1)
$\ln A_{it}$	-0.030
	(0.134)
$\ln A_{it} \times I_t$	0.053
	(0.562)
Cragg-Donald F-stat	23.83
Kleibergen-Paap F-stat	6.641
Anderson-Rubin weak-id robust p-value	0.784
Observations	2028

NOTES: IV estimates of equation (7) with non-agricultural domestic share of expenditure. Quadratic ENSO functional form used for predicting instruments. Model incorporates country fixed effects, year fixed effects, and country-specific linear trends as included instruments. Standard errors clustered at country and continent-year levels in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

D Welfare calculations

This section details the welfare calculations in the main text.

D.1 Calculating historical variance of welfare effect

This section details the within-sample welfare calculation discussed in Section 4.2 and shown in Table 2. Recall the expression for the variance of welfare in equation (2):

$$var\left(\ln\left(C_i/L_i\right)\right) = var\left(\ln A_i\right) - \frac{2}{\epsilon} cov\left(\ln A_i, \ln \lambda_{ii}\right) + \frac{1}{\epsilon^2} var\left(\ln \lambda_{ii}\right)$$

We employ this expression to quantify the magnitude of our reduced-form results in welfare terms. Consider the following thought experiment: suppose the spatial correlation of productivity increases from the 1961-2013 historical mean, $\bar{I} = .215$, by one standard deviation, $\sigma_I = .0192$. What is the resulting percentage change in the cross-sectional variance of welfare, holding everything else fixed? We denote these two hypothetical states as uncorrelated state u and correlated state c.

For the uncorrelated state, we define variance productivity as the average cross-sectional productivity variance during 1961-2013:

$$var(\ln A_{iu}) \equiv E_t[var_i(\ln A_{it}|t)] \tag{D.1}$$

Next, we define covariance between productivity and domestic share of expenditure during the uncorrelated state as the average cross-sectional covariance during 1961-2013:

$$cov(\ln A_{iu}, \ln \lambda_{iiu}) \equiv E_t[cov_i(\ln A_{it}, \ln \lambda_{iit}|t)] \tag{D.2}$$

We further define the variance of domestic share of expenditure during the uncorrelated states as the average variance during 1961-2013:

$$var(\ln \lambda_{iiu}) \equiv E_t[var_i(\ln \lambda_{iit}|t)] \tag{D.3}$$

Note that equations (D.1)- (D.3) can be directly computed from data since $\ln A_{it}$ and $\ln \lambda_{iit}$ are observed.

For the correlated state c, $var(\ln A_{ic})$ is also given by equation (D.1) since we assume productivity variance is unaltered by changes in spatial correlation. $cov(\ln A_{ic}, \ln \lambda_{iic})$ and $var(\lambda_{iic})$, however, have to be calculated as one does not directly observe data from a year in which only $I_t = \bar{I} + \sigma_I$ while everything else is fixed at the historical mean. To do this, first recall our reduced-form expression for $\ln \lambda_{it}$ from equation (7):

$$\ln \lambda_{iit} = \beta_0 \ln A_{it} + \beta_1 \ln A_{it} I_t + \mathbb{Z}_{it}\Pi + \mu_{it}$$

Our estimates of this equation can be employed to construct each component of equation (2) for the

correlated state. The covariance between productivity and domestic share of expenditure during the correlated state is:

$$\begin{aligned} cov(\ln A_{ic}, \ln \lambda_{iic}) \equiv (\widehat{\beta}_0 + \widehat{\beta}_1(\bar{I} + \sigma_I))E_t[var_i(\ln A_{it}|t)] \\ + E_t[cov_i(\ln A_{it}, \mathbb{Z}_{it}\widehat{\Pi}|t)] + E_t[cov_i(\ln A_{it}, \widehat{\mu}_{it}|t)] \end{aligned} \tag{D.4}$$

The variance of domestic share of expenditure during the correlated state is:

$$\begin{aligned} var(\ln \lambda_{iic}) \equiv (\widehat{\beta}_0^2 + 2\widehat{\beta}_0\widehat{\beta}_1(\bar{I} + \sigma_I) + \widehat{\beta}_1^2(\bar{I} + \sigma_I)^2)E_t[var_i(\ln A_{it}|t)] \\ + E_t[var_i(\mathbb{Z}_{it}\widehat{\Pi}|t)] + E_t[var_i(\widehat{\mu}_{it}|t)] \\ + 2(\widehat{\beta}_0 + \widehat{\beta}_1(\bar{I} + \sigma_I))E_t[cov_i(\ln A_{it}, \mathbb{Z}_{it}\widehat{\Pi}|t)] \\ + 2(\widehat{\beta}_0 + \widehat{\beta}_1(\bar{I} + \sigma_I))E_t[cov_i(\ln A_{it}, \widehat{\mu}_{it}|t)] \\ + 2E_t[cov_i(\mathbb{Z}_{it}\widehat{\Pi}, \widehat{\mu}_{it}|t)] \end{aligned} \tag{D.5}$$

Observe that every term in equations (D.4) and (D.5) is either directly observed or can be obtained from estimation of equation (7). For example, for the model estimated in column 5 of panel A in Table 2, with $\widehat{\beta}_0 = 2.219$ and $\widehat{\beta}_1 = -7.783$, we have:

$$\begin{aligned} E_t[var_i(\ln A_{it}|t)] &= .454 \\ E_t[var_i(\mathbb{Z}_{it}\widehat{\Pi}|t)] &= .816 \\ E_t[var_i(\widehat{\mu}_{it}|t)] &= .100 \\ E_t[cov_i(\ln A_{it}, \mathbb{Z}_{it}\widehat{\Pi}|t)] &= -.234 \\ E_t[cov_i(\ln A_{it}, \widehat{\mu}_{it}|t)] &= .000784 \\ E_t[cov_i(\mathbb{Z}_{it}\widehat{\Pi}, \widehat{\mu}_{it}|t)] &= -.00151 \end{aligned}$$

Applying equation (2), the percentage change in the variance of welfare in the correlated state, relative to the uncorrelated state, is

$$\frac{var(\ln(C_{ic}/L_{ic})) - var(\ln(C_{iu}/L_{iu}))}{var(\ln(C_{iu}/L_{iu}))} = \frac{var\left(\ln A_{ic}\right) - (2/\epsilon)cov\left(\ln A_{ic}, \ln \lambda_{iic}\right) + (1/\epsilon^2)var\left(\ln \lambda_{iic}\right)}{var\left(\ln A_{iu}\right) - (2/\epsilon)cov\left(\ln A_{iu}, \ln \lambda_{iiu}\right) + (1/\epsilon^2)var\left(\ln \lambda_{iiu}\right)} - 1 \tag{D.6}$$

To complete the calculation, let agricultural trade elasticity be $\epsilon = 8.59$ (Caliendo and Parro, 2015). Values from equation (D.6) are shown in Table 2.

D.2 Calculating change in variance of welfare under climate change

In Section 5.2, we calculate the percentage change in the variance of welfare between the end of our estimation period, $\bar{t} = 2013$, and the end of our projection period, $\mathcal{T} = 2099$, under climate change, holding everything else fixed. Compared with the Section 4.2 welfare calculation, there is

an added complication: climate change also changes the variance of productivity.

To begin, recall equation (14), our estimation equation for $\ln A_{it}$ during the estimation period, $t \in [\underline{t}, \bar{t}]$:

$$\ln A_{it} = k(T_{it}) + \mathbb{X}_{it}\Psi + \nu_{it}$$

Using equation (14) and our business-as-usual CMIP5 ensemble mean projected temperatures under climate change, we first compute country-year agricultural productivity under climate change during the projection period $t \in [\bar{t}+1, \mathcal{T}]$:

$$\widehat{\ln A}_{it} = \widehat{k}(T_{it}) + \mathbb{X}_{i\bar{t}}\widehat{\Psi} + \widehat{\nu}_{i\bar{t}} \tag{D.7}$$

Figure 14 shows $\widehat{k}()$ for our benchmark specification. Figure 15 shows the associated change in log agricultural productivity, $\widehat{\ln A}_{i,2099} - \ln A_{i,2013}$, for each country. In the left panel of Figure 16, the black line shows $var(\widehat{\ln A}_{it})$ while the blue line shows $\widehat{I}_t$ computed using $\widehat{\ln A}_{it}$ for each projection year.

To compute the change in variance of welfare from the end of the estimation period, $\bar{t}$, to any year during the projection period, $t \in [\bar{t}+1, \mathcal{T}]$, we difference equation 2:

$$\begin{aligned} var\left(\ln(C_{it}/L_{it})\right) - var\left(\ln(C_{i\bar{t}}/L_{i\bar{t}})\right) = & [var(\widehat{\ln A}_{it} - var(lnA_{i\bar{t}})] \\ & - (2/\epsilon)[cov(\widehat{\ln A}_{it}), \widehat{\ln \lambda}_{iit}) - cov(\ln A_{i\bar{t}}, \ln \lambda_{ii\bar{t}})] \\ & + (1/\epsilon^2)[var(\widehat{\ln \lambda}_{iit}) - var(\ln \lambda_{ii\bar{t}})] \end{aligned} \tag{D.8}$$

We consider two scenarios for obtaining future domestic share of expenditure, $\widehat{\ln \lambda}_{iit}$. In the first scenario, the projection is "naive" in the sense that the spatial correlation of productivities is fixed at its value at the end of the estimation period, $I_{\bar{t}}$, throughout the projection period. In the second spatially "sophisticated" projection, we allow climate change to alter the spatial correlation of productivity.

Spatially "naive" variance projection Holding the spatial correlation of productivity fixed, the domestic share of expenditure during the projection period, $t \in [\bar{t}+1, \mathcal{T}]$, is computed using our benchmark estimation of equation 7:

$$\widehat{\ln \lambda}^n_{iit} = (\widehat{\beta}_0 + \widehat{\beta}_1 I_{\bar{t}})\widehat{\ln A}_{it} + \mathbb{Z}_{i\bar{t}}\widehat{\Pi} + \widehat{\mu}_{i\bar{t}} \tag{D.9}$$

Equations (D.7) and (D.9) allows construction of $var(\widehat{\ln \lambda}^n_{iit})$ and $cov(\widehat{\ln A}_{it}, \widehat{\ln \lambda}^n_{iit})$ for each year in the projection period. These are then fed into equation (D.8) to compute the change in the variance of welfare since 2013 over the 21st century for the spatially naive projection. This spatially naive welfare variance projection is shown as the solid gray line in the right panel of Figure 16.

Spatially "sophisticated" variance projection Allowing the spatial correlation of productivity to vary under climate change, the domestic share of expenditure during the projection period, $t \in [\bar{t}+1, \mathcal{T}]$, is computed using our benchmark estimation of equation 7:

$$\widehat{\ln \lambda}^{s}_{iit} = (\widehat{\beta}_0 + \widehat{\beta}_1 \widehat{I}_t)\widehat{\ln A}_{it} + \mathbb{Z}_{i\bar{t}}\widehat{\Pi} + \widehat{\mu}_{i\bar{t}} \tag{D.10}$$

Equations (D.7) and (D.10) allows construction of $var(\widehat{\ln \lambda}^{s}_{iit})$ and $cov(\widehat{\ln A}_{it}, \widehat{\ln \lambda}^{s}_{iit})$ for each year in the projection period. These are then fed into equation (D.8) to compute the change in the variance of welfare since 2013 over the 21st century for the spatially sophisticated projection. This spatially sophisticated welfare variance projection is shown as the solid red line in the right panel of Figure 16.

Projection bias For the period from 2013-2099, we calculate the percentage difference in the change in welfare variance under the spatially sophisticated projection relative to the spatially naive projection:

$$\frac{var\left(\ln(C^{s}_{i,2099}/L^{s}_{i,2099})\right) - var\left(\ln(C_{i,2013}/L_{i,2013})\right)}{var\left(\ln(C^{n}_{i,2099}/L^{n}_{i,2099})\right) - var\left(\ln(C_{i,2013}/L_{i,2013})\right)} - 1 \tag{D.11}$$

For our baseline parameters, the spatially sophisticated projection predicts a 61% greater increase in welfare variance than the spatially naive projection.

D.3 Calculating change in variance of welfare under climate change: 2-sector

This section calculates welfare projection implications in a multi-sector setting. We consider two simplifications to the multi-sector variance of welfare expression in equation (6). First, there are two sectors, agriculture and non-agriculture, $s \in \{a, m\}$. The global variance of welfare in any period t is:

$$\begin{aligned} var\left(\ln\left(\frac{C_{it}}{L_{it}}\right)\right) = {} & \alpha_a^2 var(\ln A_{iat}) + \alpha_m^2 var(\ln A_{imt}) + 2\alpha_a\alpha_m cov(\ln A_{iat}, \ln A_{imt}) \\ & - \frac{2\alpha_a^2}{\epsilon_a} cov\left(\ln A_{iat}, \ln \lambda_{iiat}\right) - \frac{2\alpha_a\alpha_m}{\epsilon_m} cov(\ln A_{iat}, \ln \lambda_{iimt}) \\ & - \frac{2\alpha_m^2}{\epsilon_m} cov\left(\ln A_{imt}, \ln \lambda_{iimt}\right) - \frac{2\alpha_n\alpha_a}{\epsilon_a} cov(\ln A_{imt}, \ln \lambda_{iiat}) \\ & + \frac{\alpha_a^2}{\epsilon_a^2} var(\ln \lambda_{iiat}) + \frac{\alpha_m^2}{\epsilon_m^2} var(\ln \lambda_{iimt}) + \frac{2\alpha_a\alpha_m}{\epsilon_a\epsilon_m} cov\left(\ln \lambda_{iiat}, \ln \lambda_{iimt}\right) \end{aligned} \tag{D.12}$$

where now $\epsilon_a = 8.59$ and we add $\epsilon_m = 4.14$ (Simonovska and Waugh, 2014). The average agricultural expenditure share in our data is $\alpha_a = 0.17$ and $\alpha_m = 1 - \alpha_a = .83$.

Second, we assume that climate change does not affect the variance of non-agricultural produc-

tivity, $var(\ln A_{imt})$, or the spatial correlation of non-agricultural productivity, I_{mt}.[36] This implies that any variance or covariance term in equation (D.12) containing $\ln A_{imt}$ will be differenced away under climate change. Thus, the difference in the variance of welfare from the end of the estimation period, $\bar{t}$, to any year in the projection period, $t \in [\bar{t}+1, \mathcal{T}]$ is:

$$\begin{aligned} var\left(\ln\left(\frac{C_{it}}{L_{it}}\right)\right) - var\left(\ln\left(\frac{C_{i\bar{t}}}{L_{i\bar{t}}}\right)\right) = \alpha_a^2 \left(var(\widehat{\ln A}_{iat}) - var(\ln A_{ia\bar{t}})\right) \\ - \frac{2\alpha_a^2}{\epsilon_a}\left(cov\left(\widehat{\ln A}_{iat}, \widehat{\ln \lambda}_{iiat}\right) - cov\left(\ln A_{ia\bar{t}}, \ln \lambda_{iia\bar{t}}\right)\right) \\ - \frac{2\alpha_a\alpha_m}{\epsilon_m}\left(cov(\widehat{\ln A}_{iat}, \widehat{\ln \lambda}_{iimt}) - cov(\ln A_{ia\bar{t}}, \ln \lambda_{iim\bar{t}})\right) \\ + \frac{\alpha_a^2}{\epsilon_a^2}\left(var(\widehat{\ln \lambda}_{iiat}) - var(\ln \lambda_{iia\bar{t}})\right) \\ + \frac{\alpha_n^2}{\epsilon_m^2}\left(var(\widehat{\ln \lambda}_{iimt}) - var(\ln \lambda_{iim,\bar{t}})\right) \\ + \frac{2\alpha_a\alpha_m}{\epsilon_a\epsilon_m}\left(cov\left(\widehat{\ln \lambda}_{iiat}, \widehat{\ln \lambda}_{iimt}\right) - cov\left(\ln \lambda_{iia\bar{t}}, \ln \lambda_{iim\bar{t}}\right)\right) \end{aligned} \tag{D.13}$$

The first, second, and fourth terms in equation (D.13) are agricultural terms identical to those used in the single-sector welfare calculations. Computing the remaining three terms in equation (D.13) requires one additional estimation: the causal effect of local agricultural yield and the global spatial correlation in yield on the domestic share of expenditure for non-agricultural goods. Those additional effects, discussed in Appendix C.2, are ϕ_0 and ϕ_1 in equation (C.2) and estimated in Table C.2. Each of the non-agricultural terms in equation (D.13) depends on spatial correlation, so they will differ between the spatially "naive" and "sophisticated" climate projections.

Spatially "naive" variance projection Holding the spatial correlation of productivity fixed, the non-agriculture domestic share of expenditure during the projection period, $t \in [\bar{t}+1, \mathcal{T}]$, is computed using our benchmark estimation of equation C.2:

$$\widehat{\ln \lambda}^n_{iimt} = (\widehat{\phi}_0 + \widehat{\phi}_1 I_{\bar{t}})\widehat{\ln A}_{it} + \mathbb{Z}_{i\bar{t}}\widehat{\Pi}_m + \widehat{\mu}_{im\bar{t}} \tag{D.14}$$

Equation (D.14) additionally allows construction of $cov(\widehat{\ln A}_{it}, \widehat{\ln \lambda}^n_{iimt})$, $var(\widehat{\ln \lambda}^n_{iimt})$ and $cov(\widehat{\ln \lambda}^n_{iiat}, \widehat{\ln \lambda}^n_{iimt})$, for each year in the projection period. These are then fed into equation (D.13) to compute the change in the variance of welfare since 2013 over the 21st century for the spatially naive projection. This spatially naive welfare variance projection is shown as the dashed gray line in the right panel of Figure 16.

[36] This assumption is partly necessitated by data limitations: we are not aware of a source of plausibly exogenous variation in the spatial correlation of non-agricultural productivity that could be employed to estimate the relevant parameters. This is why our paper focuses on agricultural trade.

Spatially "sophisticated" variance projection Allowing the spatial correlation of productivity to vary under climate change, the non-agriculture domestic share of expenditure during the projection period, $t \in [\bar{t}+1, \mathcal{T}]$, is computed using our benchmark estimation of equation C.2:

$$\widehat{\ln \lambda}^{s}_{iimt} = (\widehat{\phi}_0 + \widehat{\phi}_1 \widehat{I}_t)\widehat{\ln A}_{it} + \mathbb{Z}_{i\bar{t}}\widehat{\Pi}_m + \widehat{\mu}_{im\bar{t}} \tag{D.15}$$

Equation (D.15) additionally allows construction of $cov(\widehat{\ln A}_{it}, \widehat{\ln \lambda}^{s}_{iimt})$, $var(\widehat{\ln \lambda}^{s}_{iimt})$ and $cov(\widehat{\ln \lambda}^{s}_{iiat}, \widehat{\ln \lambda}^{s}_{iimt})$, for each year in the projection period. These are then fed into equation (D.13) to compute the change in the variance of welfare since 2013 over the 21st century for the spatially sophisticated projection. This spatially sophisticated welfare variance projection is shown as the dashed red line in the right panel of Figure 16.

Projection bias For the period 2013-2099, we again apply equation (D.11) to calculate the percentage difference in the change in welfare variance under the spatially sophisticated projection relative to the spatially naive projection. For our baseline parameters, the spatially sophisticated projection predicts a 50% greater increase in welfare variance than the spatially naive projection.

D.4 Calculating change in country-level welfare under climate change

From equation (1), the expression for welfare of country i in year t is

$$\ln (C_{it}/L_{it}) = \ln A_{it} + \gamma - \frac{1}{\epsilon} \ln \lambda_{iit}.$$

For the projection period, the productivity of country i, $\widehat{\ln A}_{it}$ is given by equation (D.7). Next, we calculate the difference in welfare projections for individual countries between the spatially "naive" and "sophisticated" projections.

Spatially "naive" country welfare projection Holding the spatial correlation of productivity fixed, the other component of welfare during the projection period, $t \in [\bar{t}+1, \mathcal{T}]$, is:

$$\widehat{\ln \lambda}^{n}_{iit} = (\widehat{\beta}_0 + \widehat{\beta}_1 I_{\bar{t}})\widehat{\ln A}_{it} + \widehat{\kappa} I_{\bar{t}} + \mathbb{Z}_{i,\bar{t}}\widehat{\Pi} + \widehat{\mu}_{i\bar{t}} \tag{D.16}$$

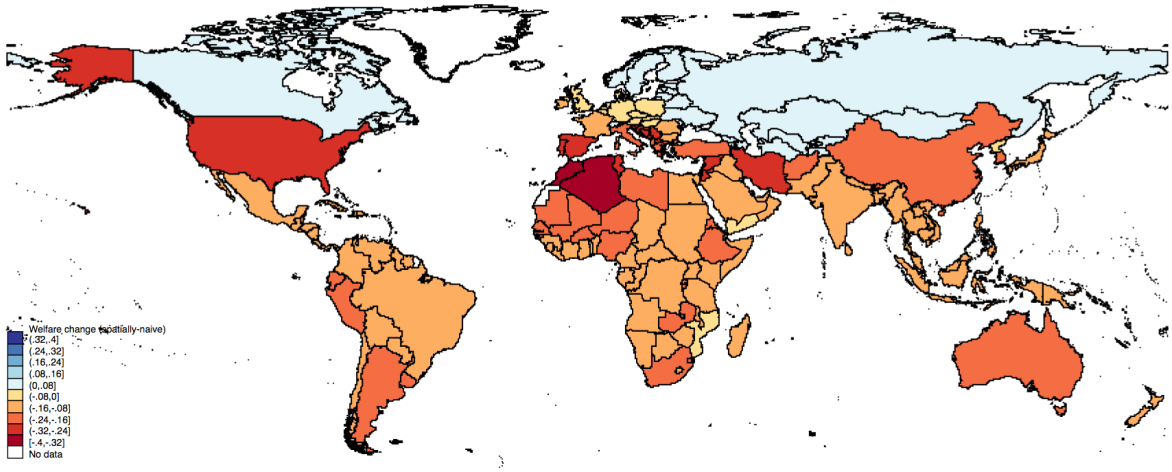
where spatial correlation affects both the average domestic share of expenditure (κ) and its relationship to domestic productivity (β_1).[37] The difference in country i welfare between end of the estimation period and the end of the projection period is:

$$\ln(C^{n}_{i\mathcal{T}}/L^{n}_{i\mathcal{T}}) - \ln(C_{i\bar{t}}/L_{i\bar{t}}) = [\widehat{\ln A}_{i\mathcal{T}} - \ln A_{i\bar{t}}] - (1/\epsilon)\left[(\widehat{\beta}_0 + \widehat{\beta}_1 I_{\bar{t}})(\widehat{\ln A}_{i\mathcal{T}} - \ln A_{i\bar{t}})\right] \tag{D.17}$$

Figure D.1 shows the change in welfare for each country under the spatially naive projection.

[37] To obtain κ, we first recover year fixed effects from equation (7). $\widehat{\kappa} = 11.17$ is the coefficient from a linear regression of year fixed effects on I_t.

Figure D.1: Welfare change under spatially naive projection (2013-2099)



NOTES: Map shows the projected 2013-2099 welfare changes from the spatially naive projection under a business-as-usual CMIP5 ensemble mean climate projection.

Spatially "sophisticated" country welfare projection Allowing the spatial correlation of productivity to vary under climate change, the other component of welfare during the projection period, $t \in [\bar{t}+1, \mathcal{T}]$ is:

$$\widehat{\ln \lambda}^{s}_{iit} = (\widehat{\beta}_0 + \widehat{\beta}_1 \widehat{I}_t)\widehat{\ln A}_{it} + \widehat{\kappa}\widehat{I}_t + \mathbb{Z}_{i,\bar{t}}\widehat{\Pi} + \widehat{\mu}_{i,\bar{t}}$$

Similarly, the difference in country i welfare between end of the estimation period and the end of the projection period is:

$$\ln(C^s_{i\mathcal{T}}/L^s_{i\mathcal{T}}) - \ln(C_{i\bar{t}}/L_{i\bar{t}}) = [\widehat{\ln A}_{i\mathcal{T}} - \ln A_{i\bar{t}}] - (1/\epsilon)\left[(\widehat{\beta}_0 + \widehat{\beta}_1 \widehat{I}_{\mathcal{T}})\widehat{\ln A}_{i\mathcal{T}} - (\widehat{\beta}_0 + \widehat{\beta}_1 I_{\bar{t}}) \ln A_{i\bar{t}} + \widehat{\kappa}(\widehat{I}_{\mathcal{T}} - I_{\bar{t}})\right]$$

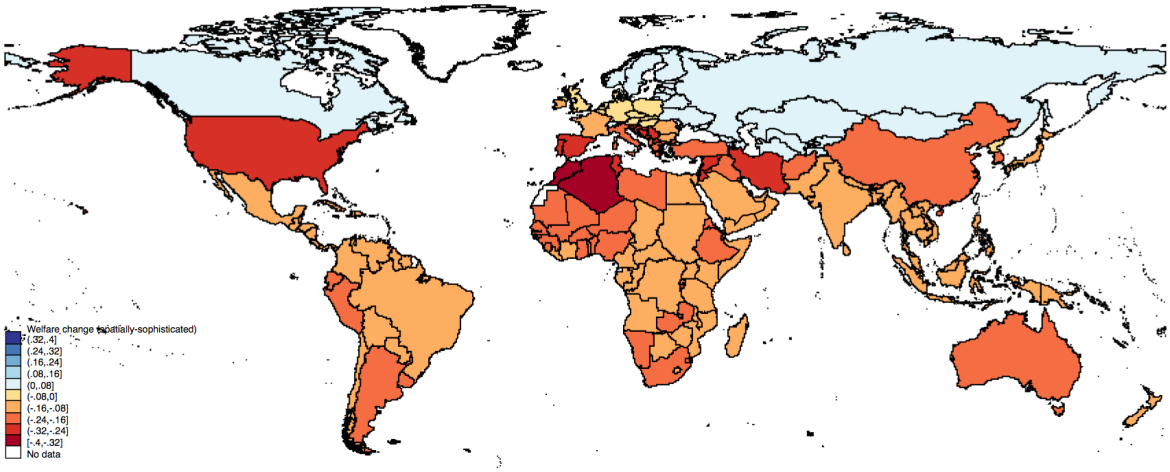
Figure D.2 shows the change in welfare for each country under the spatially sophisticated projection.

Projection bias Bias under the spatially naive projection is defined as:

$$[\ln(C^s_{i\mathcal{T}}/L^s_{i\mathcal{T}}) - \ln(C_{i\bar{t}}/L_{i\bar{t}})] - [\ln(C^n_{i\mathcal{T}}/L^n_{i\mathcal{T}}) - \ln(C_{i\bar{t}}/L_{i\bar{t}})] = -(1/\epsilon)[(\widehat{\beta}_1 \widehat{\ln A}_{i\mathcal{T}} + \widehat{\kappa})(\widehat{I}_{\mathcal{T}} - I_{\bar{t}})] \quad \text{(D.18)}$$

Figure 17 shows the country-level bias.

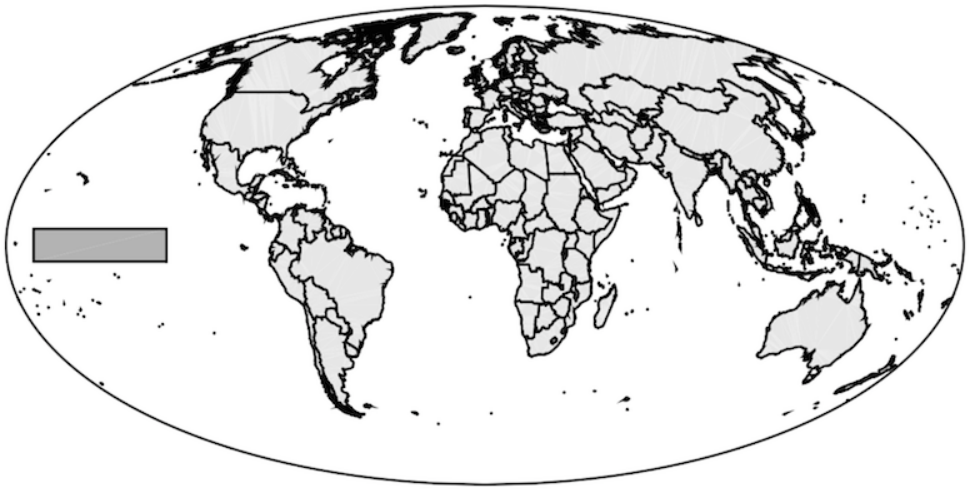
Figure D.2: Welfare change under spatially sophisticated projection (2013-2099)



NOTES: Map shows the projected 2013-2099 welfare changes from the spatially sophisticated projection under a business-as-usual CMIP5 ensemble mean climate projection.

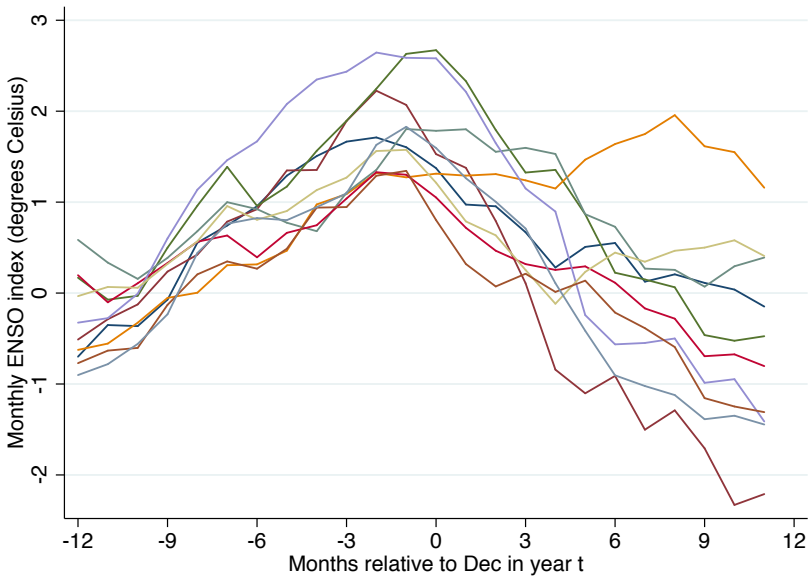
E Appendix figures

Figure E.1: Location of NINO3.4 sea-surface temperature measurements



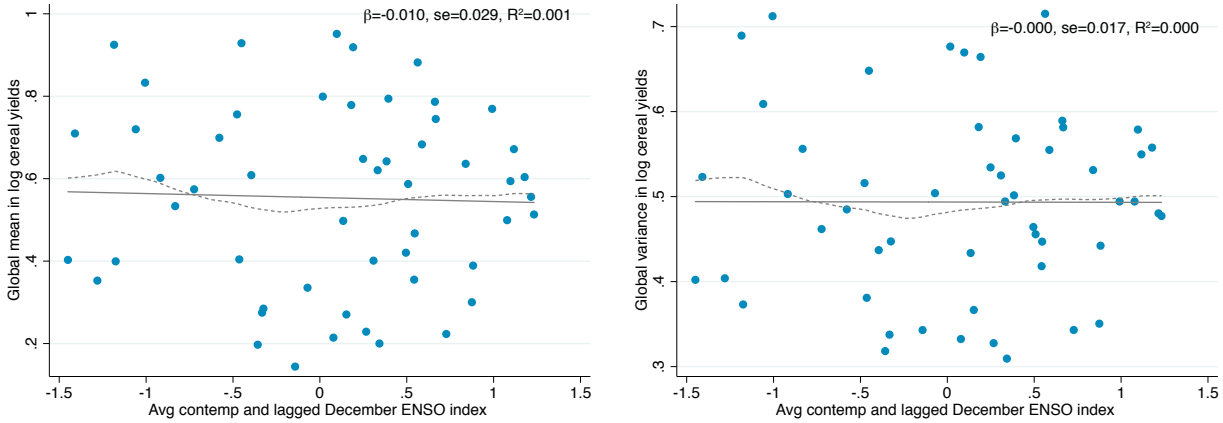
NOTES: Dark grey rectangle spanning 5°N-5°S and 170°W-120°W is the NINO3.4 region over which sea surface temperatures are averaged to compute the NINO3.4 index.

Figure E.2: Monthly ENSO index for top 10 positive events



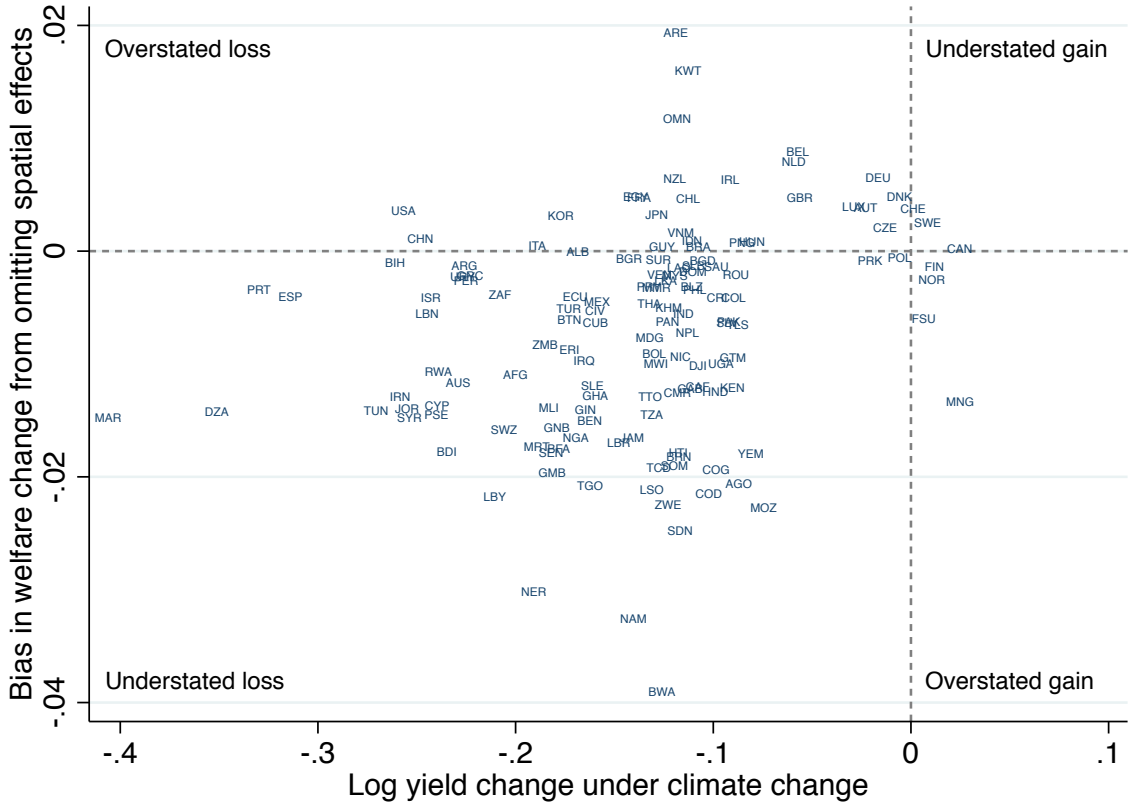
NOTES: Monthly evolution of ENSO index 12 months before and after the 10 most positive ENSO events over 1961-2013. ENSO events occur during the winters of 1965, 1972, 1982, 1986, 1991, 1994, 1997, 2002, 2006, and 2009.

Figure E.3: ENSO and the global cross-sectional mean and variance in log cereal yields



NOTES: Left (right) panel shows the relationship between the mean (variance) of cross-sectional country log cereal yields and averaged contemporaneous and lagged December ENSO. Linear fit shown as solid line. Local polynomial fit shown as dashed line.

Figure E.4: Country-level projection bias against change in agricultural productivity (2013-2099)



NOTES: Scatter shows country-level welfare bias from spatially naive projection (from Figure 17 plotted against change in log cereal yields under climate change (from Figure 15) over 2013-2099. Countries labeled by ISO3 identifier.

F Appendix tables

Table F.1: Robustness: country and year clustered standard errors

Outcome is log domestic share of expenditure					
	(1) OLS	(2) IV	(3) IV	(4) IV	(5) IV
			Panel A: Linear trends		
$\ln A_{it}$	0.363**	3.754***	3.785***	2.133***	2.219***
	(0.163)	(1.229)	(1.043)	(0.660)	(0.605)
$\ln A_{it} \times I_t$	0.645	-14.333**	-14.471***	-7.382**	-7.783***
	(0.747)	(5.743)	(5.018)	(2.935)	(2.655)
Pct. change in welfare variance	-0.300	6.510**	6.571***	3.389**	3.571***
	(0.348)	(2.558)	(2.162)	(1.372)	(1.251)
ENSO order for predicting $\ln A_{it}$		1	1	2	2
ENSO order for predicting I_t		1	2	1	2
Cragg-Donald F-stat		22.95	24.78	64.52	67.84
Kleibergen-Paap F-stat		30.55	10.03	11.16	10.23
Anderson-Rubin weak-id robust p-value		1.62e-09	0	0	0
Observations	5681	5681	5681	5681	5681
			Panel A: Linear trends		
$\ln A_{it}$	0.406*	3.473***	3.575***	3.100***	3.091***
	(0.204)	(1.240)	(0.980)	(1.063)	(0.777)
$\ln A_{it} \times I_t$	0.501	-13.471**	-13.952***	-12.073**	-12.028***
	(0.870)	(5.959)	(4.747)	(4.967)	(3.537)
Pct. change in welfare variance	-0.233	6.127**	6.341***	5.503**	5.483***
	(0.404)	(2.585)	(2.036)	(2.228)	(1.626)
ENSO order for predicting $\ln A_{it}$		1	1	2	2
ENSO order for predicting I_t		1	2	1	2
Cragg-Donald F-stat		17.25	21	24.86	26.02
Kleibergen-Paap F-stat		8.783	10.06	5.453	10.70
Observations	5681	5681	5681	5681	5681

NOTES: Estimates of equation (7). Column 1 shows OLS estimates. Columns 2-5 show IV estimates. Columns 2-5 differ in the ENSO functional form used for predicting instruments. All models incorporate country and year fixed effects as included instruments. Panel A (B) includes global or country-specific linear (quadratic) trends in equations (7)-(11). Percentage change shown in the variance of welfare for a one-standard-deviation increase in Moran's I relative to the historic mean. Standard errors clustered at country and year levels in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table F.2: Robustness: log cereal yield and local temperature

Outcome is log cereal yields

	(1)	(2)	(3)	(4)	(5)	(6)
Temperature 1st term	0.004	0.004	0.005	0.005	0.007	0.005
	(0.010)	(0.010)	(0.011)	(0.010)	(0.013)	(0.012)
Temperature 2nd term	-0.183***	-0.165***	-0.222***	-0.203**	-0.126*	-0.100
	(0.047)	(0.043)	(0.085)	(0.082)	(0.075)	(0.072)
Temperature 3rd term	0.650***	0.599***	0.418*	0.393*	0.020	-0.031
	(0.201)	(0.189)	(0.245)	(0.237)	(0.255)	(0.242)
Temperature 4th term	-1.162*	-1.100*	0.356	0.248	1.320*	1.394*
	(0.663)	(0.641)	(0.807)	(0.765)	(0.784)	(0.740)
Temperature 5th term			-2.204	-1.801	-2.895	-3.370
			(2.199)	(2.104)	(2.437)	(2.341)
Temperature 6th term					1.830	3.213
					(5.077)	(4.929)
Precipitation		0.003***		0.003***		0.003***
		(0.001)		(0.001)		(0.001)
Precipitation squared		-0.000***		-0.000***		-0.000***
		(0.000)		(0.000)		(0.000)
Number of knots	5	5	6	6	7	7
Temp. joint p-value	0.0001	0.0003	0.0002	0.0006	0.0005	0.0013
Optimal temp.	8.81	8.91	8.87	8.94	7.80	7.70
Observations	7226	7226	7226	7226	7226	7226

NOTES: Estimates of cubic spline terms for $h()$ in equation (14). The number of knots placed along the temperature support according to Harrell (2001) varies across columns. Odd (even) numbered columns exclude (include) quadratic precipitation terms. P-value from a joint significance test of temperature terms shown. Optimal temperature indicates temperature associated with highest predicted log cereal yield. Standard errors clustered at country and continent-year levels in parentheses. *** p<0.01, ** p<0.05, * p<0.1.