DISTRIBUTIONAL IMPLICATIONS OF CLIMATE CHANGE IN RURAL INDIA: A GENERAL EQUILIBRIUM APPROACH

HANAN G. JACOBY, MARIANO RABASSA, AND EMMANUEL SKOUFIAS

We develop a general equilibrium framework, based on a specific-factors trade model, to quantify the medium-term household welfare impacts of global warming in rural India. Using an hedonic approach grounded in the theory combined with detailed microdata, we estimate that three decades of warming will reduce agricultural productivity in the range of 7%–13%, with the arid northwest of India especially hard hit. Our analysis shows that the proportional welfare cost of climate change is likely to be both modest and evenly distributed across percentiles of the per capita income distribution, but this latter conclusion emerges only when the flexibility of rural wages is taken into account.

Key words: Agriculture, global warming, hedonic regression, trade.

JEL codes: Q54, Q17.

Despite uncertainties as to how climate change will unfold over the coming decades, what does seem clear is that the brunt of its economic impact will fall on low-latitude agriculture, the dominant source of livelihood for the world's poor. Partly based on this observation, the Stern Review (Stern 2006) concludes that global warming will exacerbate the poverty of nations. Yet, next to nothing is known about the implications for poverty within nations; in particular, microlevel evidence on the potential distributional consequences of climate change is lacking. Identifying the biggest losers from warming is essential for targeting policies aimed at ameliorating these negative impacts or promoting adaptation.

At first blush, there is good reason to expect that the costs of climate change will not be borne equally across income groups, even within rural areas of developing countries. After all, ownership of the most climate-vulnerable asset, farmland, is usually concentrated among the least poor of the rural population. Yet, insofar as the agricultural productivity decline is global, climate change will also raise world food prices, which, given trade openness, could benefit farmers; in particular, those that are net sellers of food. More subtly, both productivity and food price shocks will have repercussions for rural wages, often the sole source of income for the poor, and these changes, in turn, will affect other prices in the economy. In sum, it is not obvious who will lose and who will gain on balance in this adjustment process.

This article develops a general equilibrium comparative statics framework for quantitatively assessing how these sometimes conflicting forces will play out in rural India over the medium term. We view global warming as hitting the economy with an exogenous agricultural productivity shock as well as with a food price shock, the latter also exogenous given that India is, to a first approximation, a price-taker in international agricultural markets.[1] Our model yields precise predictions for how factor prices, that is, wages and land rents, will respond to these shocks and what

[1] At present, India is only a significant player in the world rice market, although its 25% share is probably not high enough to exert much monopoly power.

changes in household welfare will ensue.[2] The first part of the article thus estimates the climate sensitivity of agricultural productivity in rural India in a manner consistent with the general equilibrium framework. With these estimates and forecasts of future food price changes in hand, the second part of the article predicts the welfare consequences of climate change for rural Indian households at each point along the distribution of current welfare, as measured (conventionally) by per capita household expenditures. Aside from its sheer size and importance as an agricultural producer, India provides an ideal setting for the application of this methodology, our focus being on rural areas, home to about three-quarters of the country's huge population. With a wide range of agro-ecological zones, there is considerable variation in agricultural productivity and climate across the subcontinent. And Indian agriculture is expected to be especially hard-hit by warming temperatures in the coming decades, though not necessarily uniformly (World Bank 2009). While agriculture now contributes only about a fifth of GDP, it still absorbs almost 70% of the rural labor force; a figure that, on the whole, has declined remarkably slowly over the past two decades (Lanjouw and Murgai 2009). Productivity growth in agriculture has also been the major driver of rural poverty reduction, largely through increases in agricultural wages (Datt and Ravallion 1998; Eswaran et al. 2009; Lanjouw and Murgai 2009). Data for India are relatively abundant as well, our main source being the nationally representative National Sample Surveys (NSS). The 2002–2003 NSS round collects detailed farm-level information and the 2004–2005 round includes both a huge household expenditure survey and a parallel labor force survey.

Past empirical literature focuses on predicting the impact of a changing climate on national or global agricultural product, largely ignoring the question of how the costs will be distributed across the relevant population.[3] Agronomists base crop damage assessments on temperature responses taken from controlled agricultural experiments.

More sophisticated crop modelling (e.g., Parry et al. 1999) also accounts for certain adaptations to climate change, such as shifts in planting dates, increased irrigation, and changes in crop varieties. However, even this limited set of adaptations is assumed a priori, rather than necessarily reflecting how farmers would behave; economists prefer to use data that reveal actual behavior. The hedonic, or so-called Ricardian, approach pioneered by Mendelsohn, Nordhaus, and Shaw (1994) infers the impact of climate change on future agricultural productivity from the present-day cross-sectional relationship between climate and land values.[4]

We adopt a variant of the hedonic approach in this article, one that emerges from our general equilibrium model of factor price determination. The model, in particular, makes two key assumptions: (1) the rural Indian economy consists of many separate (district-level) labor markets across which labor does not migrate; (2) within a labor market, workers are perfectly mobile across production sectors,[5] but capital (land, in agriculture) is not.[6] Specificity of capital and perfect intersectoral mobility of labor are the hallmarks of the specific factors trade model (e.g., Jones 1971), whereas the geographic immobility of labor across rural India has been noted by, among others, Topalova (2007). Given this model-setting, land and labor prices each only partially capture variation in agricultural productivity; in other

---

[2] In a similar spirit, Fullerton and Heutel (2007) use an analytical (closed economy) general equilibrium model to study the incidence of environmental taxes in the US.

[3] Hertel, Burke, and Lobell (2010) takes a step in this direction by considering, at a fairly aggregate level, the poverty impacts of climate-induced agricultural productivity change in 15 developing countries using computable general equilibrium (CGE) modelling.

[4] Deschênes and Greenstone (2007) criticize the hedonic approach on the grounds that, in a cross-section, climate may be correlated with other variables that influence land values (we attempt to address this critique below by exploiting within-state variation). They propose, as an alternative, using panel data on farm revenue to estimate the response of agricultural productivity to year-to-year weather fluctuations (see also Guiteras 2009, for a similar approach applied to India). However, since there is little scope for adaptation to short-run weather, damage estimates based on shocks can only provide an upper bound on the long-run productivity impact of climate change; there is no guarantee, moreover, that this bound will be very informative (see also Fisher et al. 2012).

[5] Jacoby (2013) cannot reject perfect intersectoral mobility of manual labor in rural India over just a five-year horizon using data on wages and local food price shocks from 2004 to 2009.

[6] Alternatives to the assumption of perfectly immobile capital are less palatable. At the opposite extreme, if all sector-specific capital were freely convertible over the horizon in question, then factor price changes would depend only on differences in factor intensity across sectors. In this case, predicted wage responses to climate change, for example, would be *extremely* sensitive to sectoral differences in estimated factor intensities (and nothing else). In between these extremes falls a dynamic model with adjustment cost functions allowing gradual mobility of different types of capital. However, even if it proved analytically tractable, such a model would be difficult to credibly parameterize.

words, land values alone are insufficient for hedonic analysis.

Model and data together suggest that, at least for India, rural wage adjustment will be the key mechanism for redistributing the potentially substantial costs of climate change from (wealthier) landowners to the rest of the rural economy. Indeed, we show that, in proportional terms, the household welfare impacts of a combined fall in agricultural productivity and a (possible) rise in world food prices induced by global warming will be roughly the same across income groups. This conclusion holds whether geographical labor mobility *in the future* is restricted to rural areas of the same district or is allowed across rural *and* urban areas of much broader regions.[7]

Our predictions of climate-change impacts in rural India are subject to three broad caveats, worth addressing at the outset. First, our analysis is only valid for a horizon over which capital (and land) is by and large fixed in its present uses; that is to say, we abstract from *warming-induced* shifts in capital across economic sectors. Since the process of converting agricultural land to non-agricultural uses, in particular, is likely to be especially slow in India for a variety of institutional and political-economy reasons, we would argue that this horizon is apposite.

Second, first-order welfare analysis, which ignores substitution effects, may not be accurate for nonmarginal changes. Based on 30-year climate projections, we estimate a total factor productivity (TFP) decline in Indian agriculture on the order of 10% and we consider, as a worst-case scenario, an agricultural price shock of around 25%. Whether these constitute "nonmarginal" changes is a matter of judgment; Banks, Blundell, and Lewbel (1996) find that first-order approximations of the welfare effects of price changes on the order of 20% work reasonably well (i.e., with a relative error of around 10%).

Third, we do not take into account the future dynamics of river basin flows and groundwater due to, among other things,

changes in Himalayan glacier melt, nor do we consider the effects of changes in precipitation and the frequency of extreme weather. Although potentially important, the long-run evolution of each of these features is extremely uncertain and their productivity impacts—with the possible exception of those of precipitation[8]—are hard to quantify. Our focus is on the direct impacts of higher temperatures, the most predictable consequence of rising atmospheric CO_2 levels. A final agronomic caveat concerns carbon fertilization effects. While elevated CO_2 levels tend to increase plant growth under certain controlled conditions, how significant this boost will be in reality remains unclear (e.g., Parry et al. 2004; Cline 2007). For this reason, we zero out carbon fertilization effects throughout our analysis; in this sense, our productivity impacts can be viewed as worst-case scenarios of global warming as far as agriculture is concerned.

General Equilibrium Framework

Consider each district as a separate economy with three sectors: agriculture (A) and manufacturing (M), both of which produce tradable goods, and services (S), which produces a nontradable. Output Y_i in each sector $i = A, M, S$ is produced with specific capital K_i (land, in the case of agriculture), manual labor L_i, and a tradable intermediate input I_i (e.g., fertilizer in agriculture), using sector-specific production functions $Y_i = \Theta_i F_i(L_i, I_i, K_i)$, where Θ_i is sector-specific TFP. Since we are abstracting from changes in TFP in manufacturing and services, we set $\Theta_M = \Theta_S = 1$ and $\Theta_A = \Theta$ for convenience.

In India, as in most developing countries, agricultural production largely takes place on household farms using family and hired labor. Moreover, in a given year, these farms typically produce several crops on the same land (contemporaneously via multicropping and/or sequentially in multiple cropping seasons) with largely the same workers

---

[7] Timmins (2007) specifies an econometric model of labor mobility as a function of climate identified off of permanent migration decisions of household heads in Brazil. The limitation of this approach, for our purposes, is twofold: First, climate is assumed to have only amenity, and no productive, value, and, second, household asset endowments are largely irrelevant. In our view, such a model, especially applied to a rural setting where landownership differences are salient, would not yield a realistic distribution of economic costs of climate change.

[8] Quoting from the 2009 World Bank report "...the magnitude and precise timing of [changes in monsoon patterns] is unknown, as global circulation models lack accuracy at finer spatial resolutions and there remain large uncertainties in projecting local changes in climate" (p. 67). See also Sanghi, Mendelsohn, and Dinar (1998) and the review of Gornall et al. (2010) for a similar conclusion. We discuss the effect of precipitation on agricultural productivity further below.

and intermediate inputs. Hence, following, for example, Strauss (1986), we treat the representative farm in a given district as a multiproduct firm that chooses among a fixed set of c crops $\{Y_1, \ldots, Y_c\}$ to grow, transforming output between crops according to the function $Y_A = G(Y_1, \ldots, Y_c)$, where G is assumed to be homogeneous of degree one. To account for the huge agroclimatic variation across India, one should think of the set of feasible crops as varying across districts.

The quantities Y_j are chosen to maximize total revenue, $\sum_{j=1}^{c} P_j Y_j$, where P_j is the price of crop j, subject to the constraint that $G(Y_1, \ldots, Y_c) = Y_A$ for any given Y_A. Thus, in this set-up, production value shares $s_k = P_k Y_k / \sum_{j=1}^{c} P_j Y_j$ are determined by both agroclimatic conditions and by relative crop prices. Given the homogeneity of G, there exists a price index P_A such that $P_A Y_A = \sum_{j=1}^{c} P_j Y_j$, which upon differentiation yields

$$(1) \qquad \widehat{P}_A = \sum_j s_j \widehat{P}_j$$

where the circumflex denotes proportional changes; in other words, $\hat{x} = d \log x$. This establishes our production value share-weighted agricultural price index.

Now, we may write the return per acre in agriculture as $\Pi_A = (P_A F_A(L_A, I_A, K_A) - P_I I_A - W L_A)/K_A$, with analogous expressions for average return per unit capital in manufacturing, Π_M, and in services, Π_S, given respective output prices in these sectors, P_M and P_S. We assume that manual labor is perfectly mobile across the three sectors but its overall supply is fixed at $L = L_A + L_M + L_S$ within each district. Thus, in each district economy, there is one type of labor with a single nominal wage, W. The purpose of assuming a fixed district labor supply (i.e., no interdistrict mobility) is to allow for a more general empirical specification of the hedonic price function incorporating cross-district wage variation. In predicting the impacts of climate change, however, we are free to allow for labor markets with much broader geographic scope.

Because this is a general equilibrium framework, income effects of changes in factor prices are fully accounted for. Thus, total income y consists of the sum of value-added (revenue net input expenditures) across

sectors $i = A, M, S$

$$(2) \qquad y = \sum_i P_i Y_i - P_I I_i + E$$

with an additional exogenous component, E. Though a technical nuisance, the presence of E suits an important empirical purpose: A significant portion of household income in rural India comes from (salaried) nonmanual labor, for example, teachers, police/army, and other civil servants. Not accounting for E would severely distort the share of income derived from manual labor, especially for households in the upper tail of the income distribution. The exogeneity assumption on this income can be motivated by thinking about entry into these professions as requiring an advanced level of education (relative to unskilled labor), which cannot be acquired in the short run. We may further suppose that these salaries are paid out of a central or state government budget financed by urban taxpayers and that such labor does not contribute directly to output in any of the three rural sectors.

Given perfect competition, the condition that price equals unit cost implies

$$(3) \qquad \alpha_{LA} \widehat{W} + \alpha_{KA} \widehat{\Pi}_A = \widehat{\Theta} + \widehat{P}_A$$

where, under constant returns to scale, the input cost shares in agriculture, the α_{lA}, $l = K, L, I$, are such that $\alpha_{KA} + \alpha_{LA} + \alpha_{IA} = 1$. Similar equations hold for the other sectors as shown in appendix A, which lays out the model solution for the case of sector-specific costs shares and Cobb-Douglas production functions. However, for the sake of clarity, and because it will make no appreciable empirical difference, we henceforth assume equal input cost shares across sectors, dropping the corresponding sector subscript (i.e., $\alpha_{li} = \alpha_l$).

The relationship between wages and exogenous shocks to the economy in this case is

$$(4) \qquad \widehat{W} = \psi(\widehat{\Theta} + \widehat{P}_A)$$

$$\text{where } \psi = \frac{\beta_A}{\alpha_L + \alpha_K - R\beta_S}.$$

The $\beta_i = L_i/L$ are the sectoral labor shares and $R = \alpha_L + \alpha_K(\alpha_K + \alpha_L)(1 - E/y)$.

We can gain some intuition for the mechanics of the model by considering

the special case $\alpha_I = \beta_S = 0$; a two-input, two-sector economy (without nontradables). According to equation (4), $\psi = \beta_A$, where β_A is the share of the rural (district) labor force in agriculture. Thus, as either agricultural TFP or output prices increase, the marginal value product of labor rises in agriculture, drawing labor into farm work from the other sectors (in this case, only manufacturing) until equality of marginal value products across sectors is restored. The larger agriculture is in relation to nonagriculture, the less elastic is the supply of labor to agriculture and hence the more responsive is the wage to the exogenous shock.

If we now let $\alpha_I > 0$, then we have $\psi = \beta_A/(\alpha_L + \alpha_K) > \beta_A$. The source of this amplification effect is the increase in intermediate input use induced by higher agricultural prices, which boosts the marginal product of labor in agriculture. Because of a greater exodus of labor from manufacturing in response to agriculture's improved terms of trade, there must be an even larger wage increase than was the case in the absence of intermediates.

Finally, suppose that $\beta_S > 0$. The term $R\beta_S$ enters equation (4) after solving out for $\widehat{P}_S$, the change in the endogenous price of services. We do so, in appendix A, by equating the supply of services to the demand under the assumption of a unitary income elasticity of demand.[9] Return to the case $\alpha_I = 0$ and suppose that $E = 0$, so that there is no exogenous source of income outside of the three sectors. Here we have $\psi = \beta_A/(1 - \beta_S) > \beta_A$, which shows that the introduction of a nontradable sector also has an amplification effect. A climate-induced rise in the wage reduces the supply of services and increases the demand for services through an income effect. Both forces put upward pressure on the price of services, drawing labor into that sector. The supply curve of labor to agriculture, in turn, becomes more inelastic, making the rural wage respond even more to the exogenous shocks.

Finally, combining (3) and (4), gives

$$(5) \qquad \widehat{\Pi}_A = \frac{1}{\alpha_K}\left(1 - \alpha_L\psi\right)\left(\widehat{\Theta} + \widehat{P}_A\right).$$

The response of land values to warming-induced productivity and price shocks differs from that of wages (cf. equation 4) because of the indirect (factor price) effect of wages operating through agricultural profitability.

Households and welfare

Drilling down from the representative consumer-producer, which we invoked above to determine factor prices, consider an individual household h embedded within this rural economy. It has an endowment of farmland, K_{Ah}, that earns a per unit return of Π_A and an endowment of manual labor, L_h, returning W, which it can supply to its own farm (insofar as $K_{Ah} > 0$) or to other farms, or to firms in the services or manufacturing sectors.[10] As a concession to data availability, we will have to ignore the returns to service or manufacturing sector capital held by rural households (Jacoby (2013) shows that this omission is inconsequential).[11] Thus, income of household h is

$$(6) \qquad y_h = \Pi_A K_{Ah} + W L_h + E_h.$$

Household indirect utility is a function of income and prices, P_M, P_S, and P_j, $j = 1, \ldots, c$. Following the conventional derivation, the percent change in money-metric utility for household h, m_h, is

$$(7) \qquad \hat{m}_h = \lambda_{\Pi h}\widehat{\Pi}_A + \lambda_{Wh}\widehat{W} - \nu_{Sh}\widehat{P}_S - \sum_j \nu_{jh}\widehat{P}_j$$

where $\lambda_{\Pi h}$ and λ_{Wh} represent, respectively, the proportion of income the household derives from owned land and manual labor and ν_{jh} and ν_{Sh} are, respectively, the crop j and services shares of total consumption

---

expenditures.[12] Note that, while the last term in equation (7) is standard in the literature, the penultimate term is typically overlooked (Porto (2006) is a notable exception). Even if climate change has no impact on domestic crop prices through international trade, a decline in agricultural TFP will still lower the price of services (along with wages) and thereby raise welfare (holding income constant).

Substituting (1), (4), (5), and $\widehat{P}_S = R\widehat{W}$ into (7), leads to

$$(8) \qquad \hat{m}_h = \Omega_h \widehat{\Theta} + \sum_j (\Omega_h s_j - \nu_{jh}) \widehat{P}_j$$

where $\Omega_h = \lambda_{\Pi h}(1 - \alpha_L \psi)/\alpha_K + \psi \lambda_{Wh} - R\psi \nu_{Sh}$ is the elasticity of welfare with respect to agricultural productivity. This is the formula we implement to illustrate distributional implications of climate change. The term $\Omega_h s_j - \nu_{jh}$ is reminiscent of Deaton's (1989) well-known "net consumption ratio" (production value minus expenditures on good j divided by total consumption expenditures). Deaton's term measures the money-metric utility impact of a change in the price of good j in partial equilibrium, whereas ours measures its impact in general equilibrium, in other words, fully accounting for the induced changes in factor income. In particular, Ω_h depends on the β_j, the relative importance of each of the sectors in the relevant labor market.

Shutting down wage adjustment completely so that $\widehat{W} = 0$, we get $\widehat{\Pi}_A = (\widehat{\Theta} + \widehat{P}_A)/\alpha_K$ from equation (3), which, combined with $\widehat{P}_S = 0$ in equation (7), gives $\Omega_h|_{\widehat{W}=0} = \lambda_{\Pi h}/\alpha_K$. Since $\lambda_{\Pi h} = 0$ for households that own no land, the partial equilibrium approach has two strong distributional implications: First, landless households are unaffected by the TFP shock, and second, they are affected by the crop price shock only in proportion to the quantities of the crops they consume. Thus, accounting for general equilibrium effects operating through

---

[12] Household-specific expenditure shares that vary by income level are, strictly speaking, inconsistent with the representative consumer having Cobb-Douglas preferences (i.e., a unitary income elasticity of demand for services), an assumption we need to obtain a tractable analytical model of factor price determination. Hence, this is another concession we make to the data; expenditure shares do, in fact, vary by income level.

rural wages can, at least in principle, substantially alter the distributional story of climate change.

Agricultural productivity and climate

Our goal is to estimate the impact of climate on agricultural TFP, which allows us to predict the warming-induced productivity shock, $\widehat{\Theta}$. The regression is

$$(9) \qquad \theta_d = \gamma' Z_d + \xi_d,$$

where θ_d is *log* agricultural TFP for district d and Z_d is a vector of district-specific climatic and agro-ecological factors. The slope of this hedonic function incorporates the myriad of adaptations (crop choices, planting dates, labor intensity, etc.) that Indian farmers have made over time to local variation in climate. Provided that adaptation to climate change takes place along this same technological envelope, the marginal effect of, say, temperature from equation (9) delivers the percent change in agricultural TFP implied by any future warming. (see, e.g., Mendelsohn, Nordhaus, and Shaw 1994).

Even though we do not directly observe it, θ_d can be constructed based on equation (3), which provides the link between climate and agricultural TFP. Setting $\widehat{P}_A = 0$, since we can ignore price variation in the cross-section, and integrating both sides of that equation yields

$$(10) \qquad \alpha_L w + \alpha_K \pi = \theta + \kappa,$$

where lower-case letters denote logs. Thus, up to the constant of integration κ, log of TFP is the input cost-share weighted sum of the (log) factor prices. The reason why land values alone do not reflect agricultural productivity (i.e., $\pi \neq \theta + \kappa$) in our setting is simple: When labor is geographically immobile, spatial variation in θ induces spatial variation in wages. Since labor is necessarily an input into agricultural production, higher wages in more productive districts attenuate the increase in land values that would otherwise occur.

To compute θ_d, we regress log-wages on individual characteristics and district fixed effects and log of land values on plot characteristics and district fixed effect (FE). The first of these regressions yields estimated

district fixed effects, $\bar{w}_d$, and corresponding standard errors, $\sigma(\bar{w}_d)$, computed using the Haisken-DeNew and Schmidt (1997) procedure; similarly, the second regression gives $\bar{\pi}_d$ and $\sigma(\bar{\pi}_d)$. Combining these elements with calibrated input cost shares in agriculture by region r delivers $\bar{\theta}_d = \alpha_{Lr}\bar{w}_d + \alpha_{Kr}\bar{\pi}_d$ along with regression weights $1/(\alpha_{Lr}^2\sigma^2(\bar{w}_d) + \alpha_{Kr}^2\sigma^2(\bar{\pi}_d))$. Weighting corrects for the heteroskedasticity induced in ξ_d by the sampling variance of the two sets of estimated fixed effects underlying $\bar{\theta}_d$.[13]

Data and Estimation

Our analysis covers 465 districts in 18 geographically contiguous states of India, containing the vast bulk of its rural population and agricultural potential as well as encompassing tremendous variation in climate, ranging from the desert-like western Rajasthan to the moist eastern foothills of the Himalaya to the tropical south (see appendix figure C.1).[14]

Climate variables

Gridded weather data are available from the India Meteorological Department (IMD) on 296 $1° \times 1°$ lat./long. cells based on daily records from 395 temperature stations and more than 1800 rainfall stations (see Rajeevan et al. 2006; Srivastava, Rajeevan, and Kshirsagar 2009). Given data incompleteness and the timing of NSS surveys (see below), we interpolate normal annual temperature for the 1969–1999 period and precipitation for 1960–1999 using the proportion of the district's area in each grid-cell as weights.[15] Similarly, we compute normal annual degree days (cumulative temperature between 8 and 32°C) and harmful degree days (cumulative temperature above 34°C) for 1969–1999.

Most of India's precipitation falls during the southwestern monsoon, extending from June through September. The monsoon plays a critical role in determining the success of the kharif (summer season) harvest; rabi (winter) season production is not only less subject to temperature extremes but relies to a much greater extent on irrigation, mostly from groundwater. Indeed, in much of India, the availability of groundwater in rabi itself depends on aquifer recharge from the monsoon rains. Because of such complexities and because the timing and extent of the growing seasons are so variable across India,[16] we avoid creating monthly or seasonal temperature and precipitation variables as in other hedonic studies. Aside from problems of interpretability, a proliferation of climate variables may also lead to multicollinearity (see, e.g., Kaufman 1998).

Recent climate impact studies (e.g., Deschênes and Greenstone 2007) advocate using degree-days instead of temperature. For India, however, annual degree days and temperature normals are almost perfectly correlated across the 465 districts covered by our data, so nothing is gained by abjuring the convenience of temperature. By contrast, harmful degree days, a measure of extreme heat detrimental to plant growth, is only weakly correlated with normal annual temperature, hence we will experiment with this variable below.

Finally, in a departure from standard hedonic regressions, we allow weather variability, in the form of the standard deviation of annual rainfall, to influence agricultural productivity. Although it is difficult to incorporate uncertainty directly into our theoretical framework, land values would plausibly reflect a risk premium that is increasing in rainfall variability.

Land values

We use plot value data reported in the most recent nationally representative farm household survey of India, round 59 of the National Sample Survey (NSS59) collected

---

[13] As we discuss below, the error term ξ_d is also unlikely to be *independent and identically distributed* given spatial correlation across districts.

[14] Excluded are the peripheral states of Jammu/Kashmir in the far north and Assam and its smaller neighbors to the north and east of Bangladesh. Included states, organized into five regions, are *North*: Harayana, Himachal Pradesh, Punjab, Uttar Pradesh, and Uttaranchal; *Northwest*: Gujarat and Rajastan; *Center*: Chhattisgarh, Madhya Pradesh, Maharashtra, and Orissa; *East*: Bihar, Jharkhand, and West Bengal; *South*: Andhra Pradesh, Karnataka, Kerala, and Tamil Nadu.

[15] Although the NSS household surveys were collected using clustered sampling, the precise location of primary sampling units (villages in rural areas) has not been publicly released. Thus, the district is the lowest level of geographic disaggregation at which we can link household and climate variables.

[16] To take a well-known example, kharif season planting is timed to the onset of the summer monsoon, which varies by about a month between the southern and northern extremities of India.

in 2002–2003. Although these land valuations are self-assessed, rather than based on records of actual market transactions, farmers generally know the worth of this, their primary asset. Indeed, evidence from elsewhere in south Asia (Jacoby 2000) indicates that self-assessed plot values accurately capture the rental income flow from land. Idiosyncratic biases, moreover, would tend to wash out when averaging over a vast number of observations on individual plot values collected from tens of thousands of farmers.

From the eighteen major states, we have value, area, and irrigation status for 91,684 plots (reported by 41,703 households) across 468 districts.[17] More than half of the plots are irrigated, although the data do not identify the main source (i.e., canals or groundwater). At any rate, irrigation is most prevalent across the Indo-Gangetic plain in the north and east of the country, where surface water is most abundant; not coincidentally, this is also the region with the highest reported land values. In the first-stage regression of log plot value per hectare, we control for the plot's irrigation status along with the district fixed effects. The results show that, on average, an irrigated plot is 38.9% more valuable (with a robust standard error of 1.7 percentage points) than a rainfed plot within the same district. Given that access to irrigation is typically a *sine qua non* for dry-season cultivation in India, this magnitude does not seem unreasonable.[18]

Wages

For the second component of our agricultural TFP index, $\bar{\theta}_d$, we draw upon wage data from two rounds of the NSS Employment-Unemployment Survey (EUS), the 55th (1999–2000) and the 61st (2004–2005), which straddle the 59th round, from which we

obtain land values. In the spirit of the theoretical model, we focus on "manual labor," which constitutes nearly 80% of wage or salaried employment. Combined, the two surveys yield a sample of 74,659 rural wage earners from 496 districts in the major states of India. We take individual log daily wages in the last week as the dependent variable in our first-stage regression. Aside from district fixed effects, the first-stage model is standard (and hence coefficient estimates are not reported here), including a quadratic in age, literacy, and education dummies, an agricultural employment dummy, gender interactions with all of the previous variables, and an NSS round dummy interacted with all of the above.

Productivity results

The TFP index $\bar{\theta}_d$ is an input cost share weighted sum of the district fixed effects extracted from the land values and wage regressions just described, for which there are 465 districts in common. The cost shares are calculated from NSS59 data by five broad regions as defined in note 14 (see table 3 and appendix B for details).[19] As discussed above, hedonic regressions based on equation (9) are weighted by the inverse estimated standard error of $\bar{\theta}_d$. Since the error term, ξ_d, is likely to be spatially autocorrelated, we calculate standard errors using Conley's (1999) spatial GMM estimator with bandwidth of 8 degrees latitude-longitude (≈ 800 km).

Table 1 reports estimates of equation (9), starting with simple linear specifications in three baseline climate variables: temperature, precipitation, and the standard deviation of precipitation. All regressions include district-level agro-ecological controls: elevation, topography, soil texture, and type,[20] as well as distance to nearest large city, an indicator of transport cost to markets (see table C.1). From the estimates in column 1, we conclude that a 1°C increase in average annual temperature would lead to a nearly

---

[17] Aside from irrigated and unirrigated land devoted to seasonal crops, the only other category of agricultural land specified in the survey is "orchards and plantations (including forests)," which constitutes 4% of agricultural land area. Due to its heterogeneous nature, we exclude this category from our analysis. Privately held grassland dedicated exclusively to pasture also appears to be rare in India, as livestock are typically fed fodder or grazed in fallow fields or in village common areas.

[18] Since we have many multi-plot households in the data, we also estimate a household fixed effects model to clean out unobserved variation in land quality (a household's plots are typically located in close proximity and thus should be much more similar to each other than to other plots in the same district). The resulting return on irrigation is 37.5%, also very precisely estimated, suggesting that omitted variable bias is not important at this stage of the estimation.

---

[19] We could, in principle, estimate α_l for each district. However, given the sample size of NSS59, these estimates would tend to be quite noisy.

[20] The three latter variables are derived from the Food and Agriculture Organization's Soil Map of the World (SMW). We use dummies for five dominant soil types in India (acrisols, cambisols, luvisols, nitosols, and vertisols), three textural classes (coarse, medium, and fine) reflecting the relative proportions of clay, silt, and sand in the soil, and three slopes classes: flat (0–8% slope), undulating (8–30% slope), and hilly (>30% slope).

Table 1. Climate Effects on Agricultural Productivity across Indian Districts

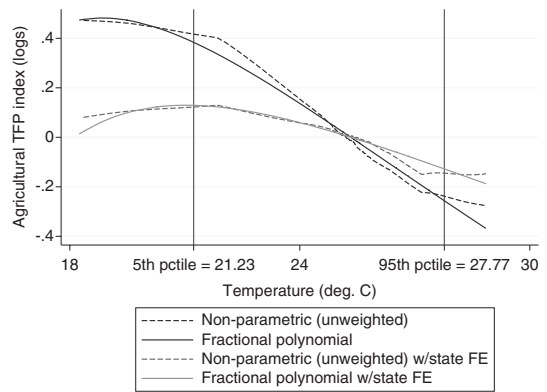|	(1)	(2)	(3)	(4)	(5)	(6)
T (temp., °C)	−0.0977***	−0.0398*				
	(0.0154)	(0.0223)				
T^{-2}			−3.835	−5.138*	−4.081	−4.890*
			(3.278)	(2.752)	(3.273)	(2.655)
$T^{-2}log(T)$			28.73***	21.12**	29.18***	20.12**
			(10.31)	(9.090)	(10.29)	(9.083)
Harmful deg. days					−0.000790	−0.000944
					(0.00117)	(0.00166)
P (precip., meters)	0.250***	0.396***	0.269***	0.423***	0.233**	0.401**
	(0.0782)	(0.144)	(0.0793)	(0.149)	(0.105)	(0.166)
Std. Dev. of P	−0.233***	−0.274***	−0.245***	−0.290***	−0.224***	−0.277***
	(0.0543)	(0.0928)	(0.0566)	(0.0966)	(0.0679)	(0.105)
State fixed effects	No	Yes	No	Yes	No	Yes
Adj. R^2	0.595	0.675	0.597	0.675	0.597	0.676

Note: Spatially robust standard errors in parentheses (***$p < .01$, **$p < .05$, *$p < .1$) based on Conley (1999) with 8 degree bandwidth ($N = 465$ districts). Dependent variable for all regressions is $\bar{\theta}_d = \alpha_{Lr}\bar{w}_d + \alpha_{Kr}\bar{\pi}_d$, all are weighted by $1/(\alpha_{Lr}^2\sigma^2(\bar{w}_d) + \alpha_{Kr}^2\sigma^2(\bar{\pi}_d))$, and all include a constant term and controls for soil type and texture, terrain, elevation, and distance to nearest major city.

10% fall in agricultural productivity in India; wetter districts and those with less variable precipitation are, meanwhile, significantly more productive.

To address the critique of Deschênes and Greenstone (2007) concerning omitted variable bias in hedonic regressions, we include state fixed effects in column 2. Relative to column 1, there is a roughly 60% attenuation of the temperature coefficient. There are two possible, not necessarily mutually exclusive, explanations for this difference: On the one hand, since state dummies (state dummies + controls) explain 77% (85%) of the variation in temperature across districts, the within-state estimator will tend to accentuate any effects of measurement error, biasing the temperature coefficient toward zero.[21] On the other hand, there may indeed be an omitted variable bias affecting the column 1 estimates. Indian states, for example, set their own agricultural policies independent of the national government, and these policies may be correlated with climate (e.g., hotter and drier states might provide greater subsidies to electricity for pumping groundwater). The true average productivity impact of 1°C higher temperatures, at any rate, probably lies somewhere in between −4% and −10%.[22]

Figure 1. Nonlinearities in Temperature

Note: Alternative estimates of the relationship between agricultural TFP (θ) and mean temperature (*T*).

Figure 1 illustrates the estimated nonlinear relationship between the agricultural TFP index, $\bar{\theta}_d$, and temperature in the corresponding district, based on (1) a partially linear nonparametric (in temperature) regression, and (2) the best-fitting fractional polynomial parametric regression (see columns 3 and 4 of table 1). The two sets of estimates track each other very closely, both indicating that over most of range of district-level temperatures (5th–95th percentiles), the

---

[21] Measurement error is not due solely to inaccuracy of temperature records and the like, but, more broadly, to the fact that average annual temperature is inevitably just a proxy for the 'true' measure of temperature affecting agricultural productivity.
[22] Guiteras (2009) estimates the effect of weather shocks on gross agricultural productivity in India from 1960–99, which he takes as an upper bound on the long-run impacts of climate

change. If we double his estimated effect of a 0.5°C temperature shock in the kharif season, we obtain a nearly 10% decline in productivity per 1°C warming. Thus, our results seem sensible using this admittedly crude metric.

**Table 2. Temperature-Precipitation Interaction Effects on Productivity**

	(1)	(2)	(3)	(4)
T (temperature, °C)	−0.152***	−0.116***	−0.191***	−0.186***
	(0.0248)	(0.0225)	(0.0407)	(0.0319)
P (precipitation, meters)	−1.017***	−1.150***	−2.705*	−4.261**
	(0.389)	(0.235)	(1.621)	(1.813)
P^2			0.505	1.043*
			(0.487)	(0.574)
$T \times P$	0.0514***	0.0640***	0.113*	0.186***
	(0.0164)	(0.0116)	(0.0655)	(0.0697)
$T \times P^2$			−0.0174	−0.0421*
			(0.0201)	(0.0230)
Std. Dev. of P	−0.234***	−0.289***	−0.260***	−0.238**
	(0.0540)	(0.0939)	(0.0654)	(0.0949)
State fixed effects	No	Yes	No	Yes
Adj. R^2	0.606	0.688	0.609	0.690

Notes: See notes to Table 1.

productivity-temperature relationship is practically linear. Figure 1 also shows a flatter gradient with state fixed effects than without, confirming the finding from the linear regression. Columns 5 and 6 add harmful degree days, our measure of daily temperature extremes, but its coefficient falls well short of attaining significance, and we drop the variable in the sequel.

Maintaining, for the sake of parsimony, linearity in temperature, we next explore interactions between temperature and precipitation with results reported in table 2. As seen in columns 1 and 2, the negative temperature impact on agricultural productivity is strongly attenuated in wetter districts, even more so after including state fixed effects. Figure 2 traces out the marginal effect of temperature as a function of annual rainfall, allowing for quadratic terms in rainfall (columns 3 and 4 of table 2). Interestingly, for districts at the 5th percentile of precipitation, the impacts of a 1°C temperature rise are quite similar whether state fixed effects are included (−12%) or excluded (−15%); by contrast, at the 95th percentile of precipitation, the marginal temperature effect is +2% with state fixed effects and −5% without. Thus, both within and between state evidence suggests that the driest districts of India, those in the arid northwest (Rajastan, Haryana, and Punjab), will suffer the greatest productivity losses from global warming. As we will see, this finding in itself has distributional implications.

Regarding precipitation, we have already noted a positive and statistically significant
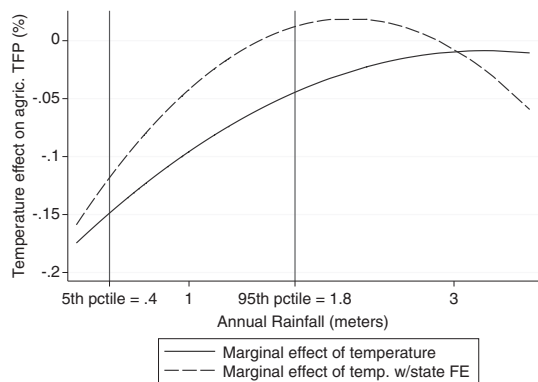
Figure 2. Marginal Effect of Temp. by Precip. Level (*P*)

Note: Marginal effects of mean temperature (T) on agricultural TFP (θ) evaluated at different values of mean precipitation (P).

productivity impact of rainfall. Putting aside the enormous uncertainty involved in forecasting precipitation patterns over the Indian subcontinent mentioned earlier, whatever consensus has emerged points to a slight rise in rainfall over the next three decades. The 2007 Intergovernmental Panel on Climate Change (IPCC) projection suggests a 4% increase, which, based on our estimates, would increase agricultural productivity by no more than about 1% on average. However, even this small positive impact would be counterbalanced, according to the findings in table 1, by the probable increase in inter-annual rainfall variability.[23]

---

[23] Results in Guiteras (2009) and Sanghi, Mendelsohn, and Dinar (1998) further support the view that climate change will

Assembling the Components

We use the mean monthly change in surface temperature between the baseline period 1961–1990 and the near future 2010–2039 as predicted by the HadCM3-A1FI (high emissions path) scenario from the Hadley Centre for Climate Prediction and Research as part of the Third Assessment Report of the United Nation's IPCC. Most other widely used climate models forecast similar temperature increases over this timeframe and only begin diverging significantly toward the end of the century. On average, India is predicted to warm by about 1.25 degrees Celsius, but temperature increases are generally greater in the North (appendix figure C.1).

Although India's government presently pursues a policy of on-again-off-again trade insulation vis á vis agricultural staples, this may not be sustainable indefinitely. If not, the question arises: How will international food prices be affected by warming? Only a calibrated general equilibrium global trade model can provide an answer. Using GTAP, a multiregion, multisector, comparative static, computable general equilibrium model of the world economy, Hertel, Burke, and Lobell (2010) forecast changes in the prices of major agricultural commodities after 30 years of climate change (2000–2030). Results are available for three scenarios based on the distribution of potential yield impacts taken from the empirical literature for each region and crop: low, medium, and high productivity, with the medium case representing the most likely set of outcomes of warming.

For consistency, any productivity scenario for India embedded in these global projections should lie roughly in the range of our own estimates of agricultural TFP change (−7% to −13%). But this is not the case for Hertel et al.'s high productivity scenario, in which yields of most crops are predicted to *rise* worldwide by 4%, leading to price *declines* on the order of 20%; therefore, we do not emphasize this scenario. While the low-productivity scenario represents the other extreme, in which global temperature change is rapid, crops are highly sensitive to warming worldwide, and CO_2 fertilization is minimal, the assumed yield declines for India of between −10% and −15% are not out of

affect agricultural productivity in India primarily through higher temperatures.

line with our own Ricardian assessment.[24] Thus, based on s_j calculated separately for each of the five regions,[25] we obtain production value share weighted price changes, $\widehat{P}_A$, ranging from −5% to −2% in the medium productivity scenario and from 18% to 27% in the low-productivity scenario (see table 3).

Our sample of just over 60,000 rural households in 18 major states is drawn from the nationally representative consumer expenditure survey (CES) collected as part of NSS61 in 2004–2005. We use these data, along with the corresponding round of the NSS-EUS, to calculate the various components of $\hat{m}_h$ as given in equation (8).

Figure 3 shows how the relevant expenditure shares are distributed across per capita expenditure percentiles in 2004–2005. Both curves in this figure, as well as all of those in figures 4, 6, 7, and 8, represent nonparametric regressions of the variable of interest, averaged across households within each of the 100 expenditure-percentile groups, on the per capita expenditure percentile itself (henceforth, simply percentile). As is typical, the share of household expenditures devoted to food crop products declines steadily with income. The percentile profile of the services (nontradable) share,[26] by contrast, is rather flat, increasing only within the top two quintiles of the per capita expenditure distribution. The remainder of household consumption expenditures consist of clothing, durables, and miscellaneous items, all of which can be considered manufactured goods, as well as some non-crop food categories (e.g., meat and dairy).

Appendix B details the steps in the construction of household income shares $\lambda_{\Pi h}$ and λ_{Wh}. As seen in table 3, the mean of $\lambda_{\Pi h}$ (λ_{Wh}) varies from 0.19 (0.48) in the North region to 0.07 (0.57) in the South (see also appendix figure C.2 for a more detailed picture); among landowning households, the corresponding figures are 0.29 (0.45) and 0.15

---

[24] Since it is not at all straightforward to associate any of Hertel et al.'s price scenarios with a particular temperature scenario, we maintain the assumptions of HadCM3-A1FI throughout.

[25] Specifically, we use farm-level data from NSS59 covering the kharif and rabi seasons of 2002–2003 to compute s_{rj} as the ratio of the (population weighted) mean production value in region r of crop j to the corresponding mean of total production value.

[26] Nontraded goods expenditure categories include: firewood and other local fuel, transport and travel, tailoring expenses, house rental and related expenses, medical treatment expenses, education expenses, remittances and gifts, recreation and leisure expenses, and taxes and fees.

Table 3. Data for Distributional Analysis by Region

|	North	Northwest	Central	East	South
Region-level					
α_L	0.331	0.304	0.258	0.317	0.260
α_I	0.264	0.325	0.258	0.250	0.238
Revenue Shares (s_j)					
Rice	0.234	0.028	0.226	0.527	0.315
Wheat	0.377	0.195	0.141	0.199	0.004
Maize	0.015	0.054	0.023	0.021	0.036
Oilseeds	0.028	0.284	0.115	0.021	0.184
Sugarcane	0.125	0.042	0.082	0.006	0.072
Cotton	0.016	0.137	0.098	0.000	0.035
Other crops	0.206	0.260	0.314	0.226	0.354
$\widehat{P}_A$: Medium Global Productivity	−3.8	−2.2	−3.2	−4.7	−3.1
$\widehat{P}_A$: Low Global Productivity	18.0	25.2	24.0	23.2	27.3
District-level					
β_A	0.691	0.755	0.805	0.722	0.688
	(0.145)	(0.103)	(0.113)	(0.143)	(0.188)
β_S	0.213	0.160	0.125	0.187	0.187
	(0.099)	(0.082)	(0.083)	(0.108)	(0.125)
$\widehat{\Theta}$: without state fixed effects	−0.162	−0.164	−0.122	−0.111	−0.121
	(0.043)	(0.046)	(0.033)	(0.028)	(0.053)
$\widehat{\Theta}$: with state fixed effects	−0.099	−0.118	−0.058	−0.042	−0.068
	(0.051)	(0.051)	(0.033)	(0.024)	(0.048)
Household-level					
λ_Π	0.193	0.145	0.117	0.107	0.070
	(0.221)	(0.172)	(0.147)	(0.161)	(0.117)
λ_W	0.479	0.554	0.567	0.524	0.573
	(0.187)	(0.183)	(0.192)	(0.182)	(0.181)
Expenditure shares (v_j)					
Rice	0.060	0.018	0.128	0.201	0.147
	(0.052)	(0.027)	(0.134)	(0.089)	(0.077)
Wheat	0.107	0.076	0.053	0.057	0.009
	(0.052)	(0.066)	(0.061)	(0.056)	(0.011)
Maize	0.002	0.015	0.004	0.003	0.000
	(0.008)	(0.043)	(0.026)	(0.014)	(0.004)
Oilseeds	0.045	0.055	0.051	0.054	0.046
	(0.020)	(0.026)	(0.022)	(0.020)	(0.021)
Sugar	0.031	0.032	0.027	0.016	0.017
	(0.017)	(0.013)	(0.016)	(0.008)	(0.012)
Other crops	0.119	0.137	0.151	0.126	0.141
	(0.042)	(0.066)	(0.064)	(0.038)	(0.053)
Nontradables (v_S)	0.182	0.169	0.170	0.162	0.178
	(0.076)	(0.072)	(0.071)	(0.067)	(0.082)

Note: Means (standard deviations) of district and household level variables, using household sampling weights in latter case. Input cost shares (α_l) and revenue shares from NSS59 (see Appendix B); crop-specific price scenarios underlying $\widehat{P}_A$ from Hertel et al. (2010); $\widehat{\Theta}$ based on col. 3 or 4 of Table 2; sectoral labor shares (β_i) from NSS61-EUS; expenditure shares from NSS61-CES (cotton share set to zero).

(0.59). Note that about a third of income for the average household comes from exogenous nonmanual labor, or E in terms of our notation.

The final elements in the computation of Ω_h are the sectoral labor shares, the β_i, which we calculate for each district using NSS61-EUS data on the industry of workers' primary occupation. According to equation (4), β_A, β_S, and the α_l (the input cost shares) also pin down ψ, the elasticity of the rural wage with respect to food prices
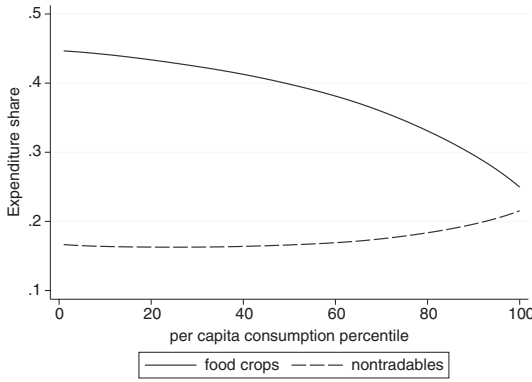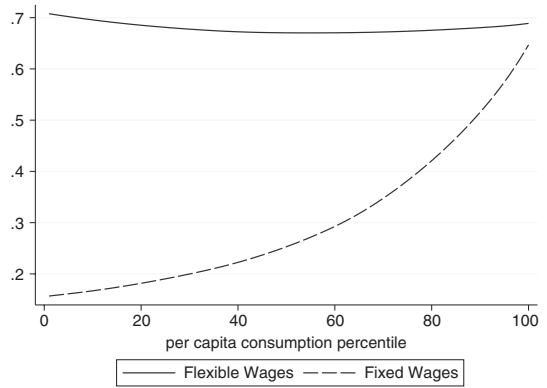
Figure 3. Distribution of Expenditure Shares

Note: Smoothed relationship between mean (within-percentile) expenditure share and per capita expenditure percentile.

Figure 4. Distribution of Welfare-Productivity Elasticity (Ω_h)

Note: Smoothed relationship between mean (within-percentile) Ω_h (elasticity of welfare w.r.t. agricultural productivity) and per capita expenditure percentile under alternative assumptions about rural wage flexibility.

(or agricultural TFP) for each district.[27] It is important to note that particularly large or small values for ψ are not preordained by the model; in particular, $\beta_S = 1 \Rightarrow \psi = 0$ whereas $\beta_A = 1 \Rightarrow \psi = \frac{1}{\alpha_K + \alpha_L} > 1$. As it happens, the elasticity is large, with a median value of 1.16 and exceeding unity for 80% of districts. The reason for this is the importance of agricultural employment in rural India; more than 90% of districts have $\beta_A > 0.5$, and more than half have $\beta_A > 0.75$. Later on, however, when we consider a broader labor market than the rural district, the relevant agricultural labor share will typically be lower.

Figure 4 displays the per capita expenditure percentile profile of Ω_h —the elasticity of welfare with respect to agricultural productivity—as well as the corresponding profile of $\Omega_h|_{\widehat{W}=0} = \lambda_{\Pi h}/\alpha_K$, the value of Ω_h that would obtain in a (hypothetical) fixed-wage economy. First, note that $\Omega_h|_{\widehat{W}=0}$ (or $\lambda_{\Pi h}$) is strongly increasing in household income (proxied by per capita expenditures), as landownership in India is concentrated among the better off. By contrast, in a flexible wage economy, the distribution of Ω_h across percentiles is remarkably flat. Three forces combine to produce this result. First, Ω_h is increasing in $\lambda_{\Pi h}$ (recall that land income is more sensitive to warming than wage income) and, as we have just seen, $\lambda_{\Pi h}$ is increasing in income. Hence, for this reason, Ω_h is increasing in income. Second, since

a higher price of services is welfare-reducing, Ω_h is decreasing in the expenditure share of services, v_{Sh}, and, as already noted, v_{Sh} is increasing in income, albeit weakly. For this reason, Ω_h is decreasing in income. Third, districts with a large share of labor deployed in agriculture tend to be poorer ones, and, as it turns out, Ω_h is increasing in β_A. This also means that Ω_h is decreasing in income. Thus, the second and third forces together counteract the first. What figure 4 shows is that the counterbalance is nearly perfect.

Finally, we examine the distribution of the predicted shock to agricultural TFP. The hedonic estimate is $\widehat{\Theta}_d = \gamma_{T_d}\Delta T_d$, where γ_{T_d} is the marginal effect of temperature in district d and ΔT_d is the predicted temperature change for the district based on HadCM3-A1FI. Thus, $\widehat{\Theta}_d$ varies across districts with the extent of predicted warming and, because we base γ_{T_d} on the last two interactive specifications in table 2, it also varies with the current precipitation level. Both of these factors, greater predicted temperature rise and lower rainfall, conspire to accentuate the agricultural TFP decline in the North and Northwest relative to the rest of India, as seen in the first two maps of figure 5. Since rural areas of northern India are relatively well off (third map of figure 5), we obtain the falling percentile profiles observed in figure 6; that is, when we assign each rural household in the NSS61-CES the estimated $\widehat{\Theta}_d$ for their district of residence. Note that these profiles are nearly identical whether we base γ_{T_d} on

---

[27] In computing R, we use an estimate of E/y at the district rather than household level, since this income ratio enters R through the aggregate demand for services.
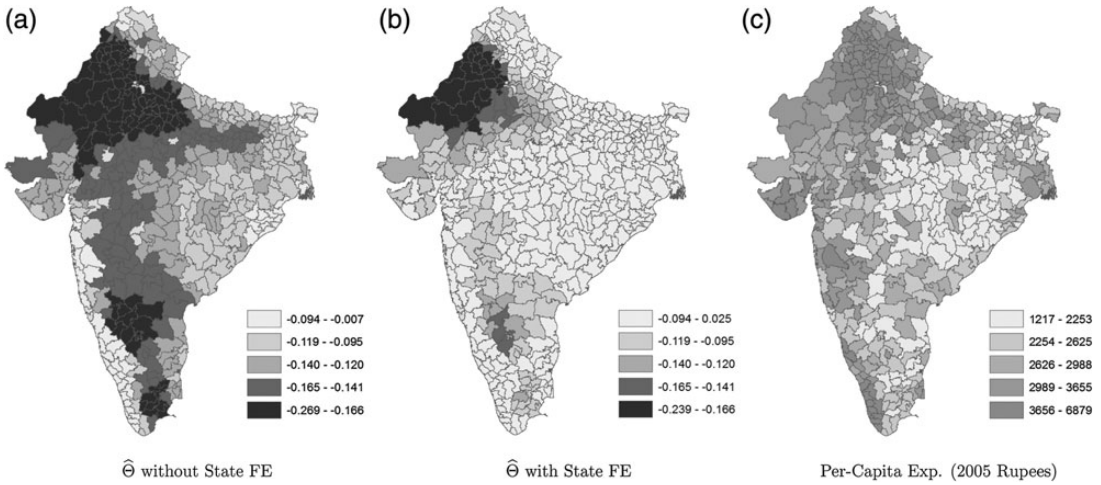
Figure 5. TFP Shock ($\widehat{\Theta}$) and Mean Per-Capita Household Expenditures across 18 Major States of India
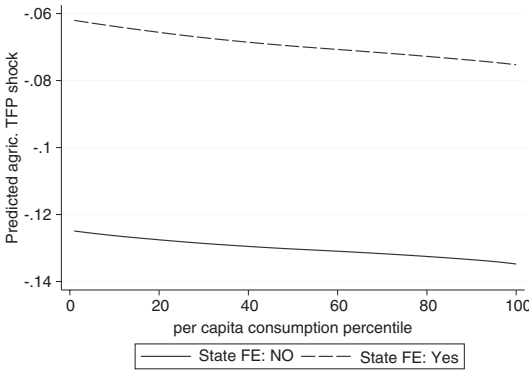
Figure 6. Distribution of $\widehat{\Theta}$

Note: Smoothed relationship between mean (within-percentile) predicted agricultural TFP shock ($\widehat{\Theta}$) and per capita expenditure percentile. Each household assigned shock corresponding to district of residence.

the hedonic estimates with state fixed effects or on those without; only the levels differ.

Distributional Implications of Climate Change

We calculate the money-metric utility change for each household from equation (8) based on the components discussed in the previous section. Figure 7 illustrates climate change incidence curves (CCIC), showing how the welfare costs of warming are distributed across per capita expenditure percentiles.[28]

The left-hand side panels depict the medium global productivity scenario for food prices and the right-hand side panels the low global productivity. Bear in mind that these alternative productivity scenarios refer to the rest of the world, not to India in particular. For India itself, we use our own TFP change predictions based on the hedonic estimates with and without state fixed effects.

Focusing on the flexible wage scenarios (panels *a* and *b*), we find that the proportional welfare impacts of climate change fall within a fairly narrow range. At no percentile would the welfare shock be much worse than around a 10% annual consumption loss (of course, this is an ongoing rather than a one-time cost).[29] Indeed, if global food prices rise substantially (panel *b*) richer households might even experience a slight welfare gain. The distribution of the shock is also practically neutral; in other words, the CCIC is essentially flat. In the low global productivity scenario, the poor gain from higher rural wages and, given the importance of wage income, this gain can offset the consumption-side loss from higher food

_____

[28] Implicit throughout is the assumption that future economic growth will be neutral with respect to current income ranking;

in other words, in the absence of warming, households of a given type would remain locked in their per capita expenditure percentile circa 2004–2005.

[29] By contrast, welfare losses are generally greater than 10% in the (unreported for reasons already discussed) *high* global productivity scenario in which most agricultural commodity prices fall sharply. However, even in this case, the welfare costs of warming are quite evenly distributed across percentiles.
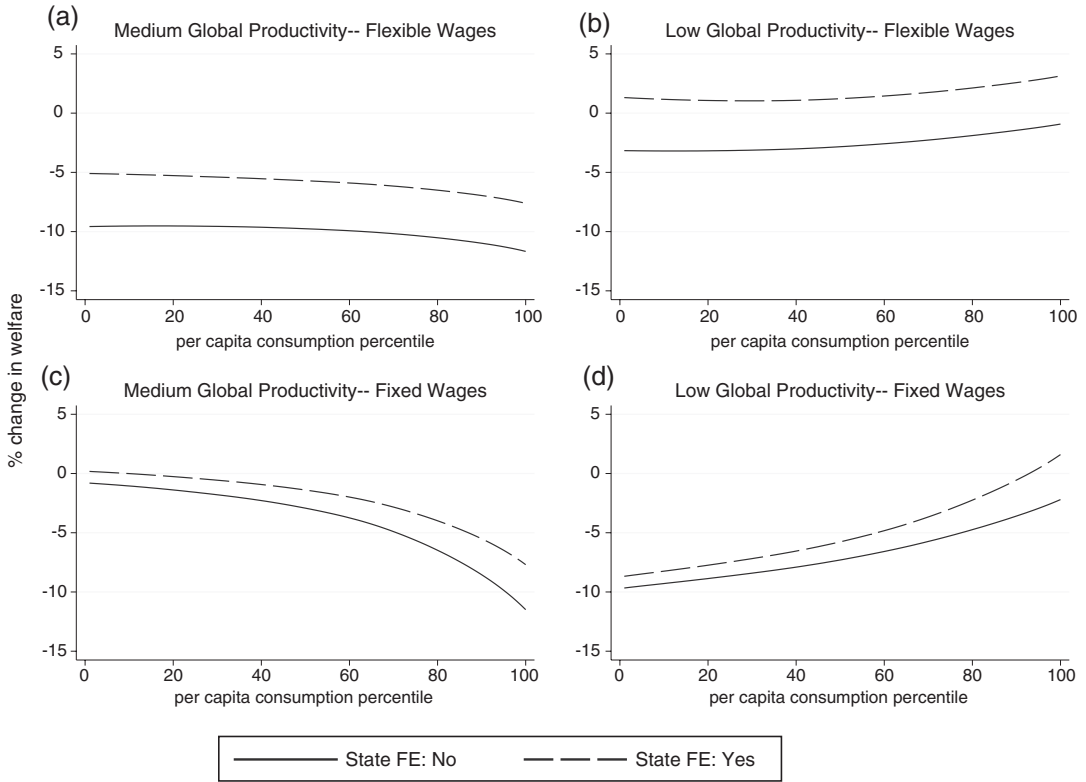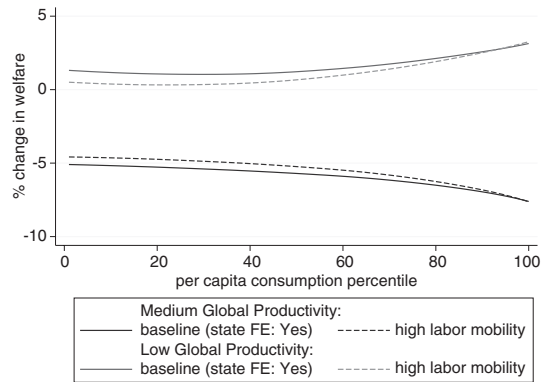
Figure 7. Climate Change Incidence Curves

Note: Smoothed relationship between mean (within-percentile) predicted welfare change ($\hat{m}_h$) and per capita expenditure percentile.

prices.[30] Our results thus imply that an ostensibly "pro-poor" policy of insulating the domestic economy from rising food prices may, over the medium run, be self-defeating (see also Jacoby 2013).[31]

It is illuminating to compare these findings to the partial equilibrium case of no wage adjustment, as shown in the lower panels of figure 7. The contrast, in terms of distributional implications, is stark. When food prices are stable or even declining (panel *c*), the poorest households are hardly touched by climate change. After all, neither the cost of
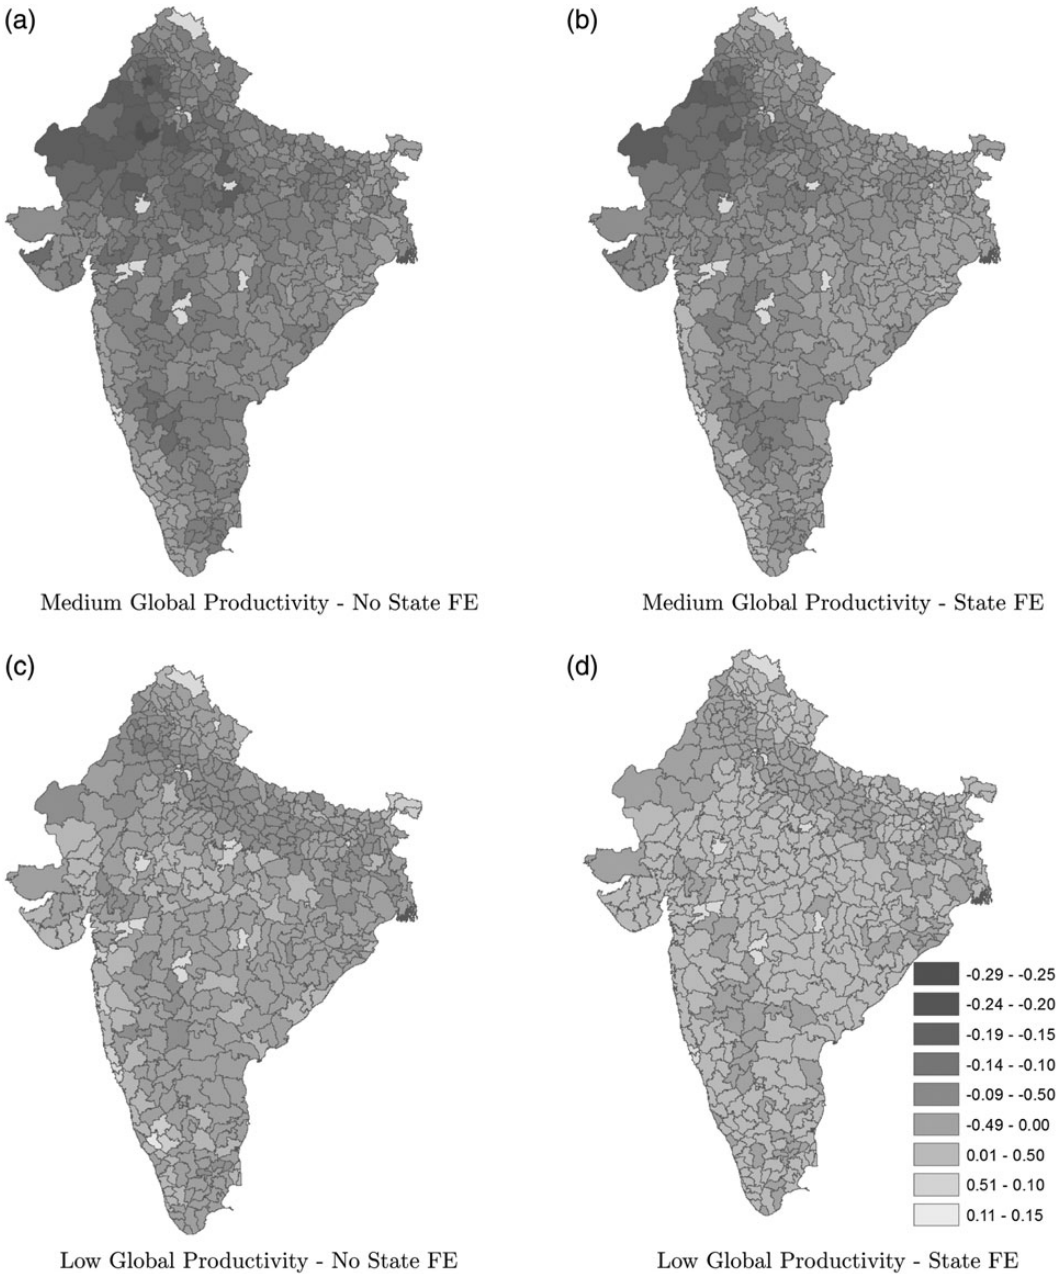

Figure 8. Robustness to Labor Mobility

Note: Smoothed relationship between mean (within-percentile) predicted welfare change ($\hat{m}_h$) and per capita expenditure percentile under alternative assumptions about geographic labor mobility.

their consumption bundle nor their income has changed much. Since better off (largely landowning) households bear the brunt of the productivity decline in this scenario, climate change looks like a steeply progressive consumption tax. When food prices rise

---

[30] Figure C.2 in the Appendix separates welfare effects due to changes in agricultural TFP from those due to changes in international commodity prices, showing how they roughly counterbalance in the low global productivity scenario. Also in that figure is a decomposition of welfare effects into consumption-side changes (due to higher prices of food and services) and income-side changes (due to higher wages and agricultural profits), which are also approximately off-setting.

[31] Of course, this ignores the urban poor who may not benefit as much from higher rural wages. A full analysis of distributional impacts in urban areas, however, would require careful modeling of rural-urban labor market linkages, which is beyond the scope of the present investigation.

(a)

Medium Global Productivity - No State FE

(b)

Medium Global Productivity - State FE

(c)

Low Global Productivity - No State FE

(d)

Low Global Productivity - State FE

■	-0.29 - -0.25
■	-0.24 - -0.20
■	-0.19 - -0.15
■	-0.14 - -0.10
■	-0.09 - -0.50
■	-0.49 - 0.00
□	0.01 - 0.50
□	0.51 - 0.10
□	0.11 - 0.15

Figure 9. Distribution of Welfare Changes ($\widehat{m}_h$) across 18 Major States of India

(panel *d*), however, the situation is reversed; climate change acts like a regressive tax. The poor lose out from higher food prices without the countervailing wage income gain, whereas the blow to the rich is softened, or even eliminated, by higher producer prices. In this counterfactual setting, trade insulation policies are indeed pro-poor.

How important for these results is the assumption that each rural district will, over the three-decade horizon, remain a hermetically sealed labor market? Figure 8 compares our baseline predicted welfare changes under each price scenario (solid curves) to those based on an alternative calculation of Ω_h, which assumes an integrated rural-urban labor market in each of the five regions (dashed curves). In other words, instead of computing β_j for each rural district, we do so for rural *and* urban areas of each

region.[32] Yet, as is evident from the figure, the differences are negligible.

Returning to the baseline flexible wage case, figure 9 shows how welfare changes are distributed across rural Indian districts. Under the medium global productivity scenario, harder hit (darker colored) districts lie primarily in the northwest and, to a lesser extent, in the southern interior. This pattern largely reflects the geographical distribution of estimated TFP changes noted in figure 5 as well as greater values of Ω_h.[33] Substantially higher food prices (coupled with trade-openness) under the low global productivity scenario noticeably alter the picture. Aside from the more modest welfare losses (larger gains) across districts, the geographic disparities are also more muted.

Conclusions

Guided by a simple general equilibrium model, we have sketched the potential distributional consequences of warming on the vast canvas that is rural India. As with all such exercises, we had to narrow our purview for the sake of empirical implementation. Nevertheless, our focus on the multifarious effects of higher temperatures operating through rural incomes and food prices has yielded several important insights. First, we estimate that the warming of about 1.25 degrees Celsius expected to occur in India over the next three decades will reduce overall agricultural productivity in the range of 7%–13%—once, that is, farmers have adapted their production methods to the new climate. However, in the arid wheat-belt of northwest India, a region now comparatively well-off, average productivity losses will be in the 10%–16% range.

Second, we have shown that rural wage adjustment is likely to be a quantitatively important part of the distributional story of climate change. Wages adjust to the productivity shock induced by warmer temperatures, which in turn lowers the price of services. Wages also respond to any rise in international agricultural commodity prices transmitted to domestic markets. Moreover,

according to our model and the parameters calculated from micro-data, these responses are elastic. On balance, wage adjustment helps the poorest households when food prices rise substantially. Indeed, if the price increases predicted by a low global productivity scenario do eventually materialize, poor (as well as rich) rural households may even stand to gain from climate change.

Third, whatever the outlook for global food prices, the welfare gains and losses from climate change will be distributed fairly evenly across income groups in rural India. One might think that wealthier landowners would lose out as long as the agricultural prices they face are stable but would gain (in relative terms) if these prices rise enough. However, the robust rural wage response overturns this superficial intuition. For the same reason, government intervention that aims to insulate poor rural consumers from higher food prices, but which has the side-effect of insulating producers, is likely to be counter-productive.

References

Banks, J., R. Blundell, and A. Lewbel. 1996. Tax Reform and Welfare Measurement: Do We Need Demand System Estimation? *Economic Journal* 106:1227–41.

Cline, W. 2007. *Global Warming and Agriculture: Impact Estimates by Country*. Washington: Center for Global Development and Peterson Institute for International Economics.

Conley, T. 1999. GMM Estimation with Cross Sectional Dependence. *Journal of Econometrics* 92:1–45.

Corden, W. M., and J. P. Neary. 1982. Booming Sector and De-Industrialisation in a Small Open Economy. *Economic Journal* 92:825–48.

Datt, G., and M. Ravallion. 1998. Farm Productivity and Rural Poverty in India. *Journal of Development Studies* 34 (4): 62–85.

Deaton, A. 1989. Rice Prices and Income Distribution in Thailand: A Non-Parametric Analysis. *Economic Journal* 99:1–37.

Deschênes, O. and M. Greenstone. 2007. The Economic Impacts of Climate Change: Evidence from Agricultural Profits and Random Fluctuations in the Weather.

---

[32] The share of employment in agricultural in these broader labor markets ranges from 58–68%.

[33] Figure C.3 in the Appendix displays the geographic pattern of both $\lambda_{\Pi h}$ and Ω_h, showing that households in the northwest derive a relatively large fraction of income from land (high $\lambda_{\Pi h}$).

American Economic Review 97 (1): 354–85.

Eswaran, M., A. Kotwal, B. Ramaswami, and W. Wadhwa. 2009. Sectoral labour Flows and agricultural Wages in india, 1983–2004: Has Growth trickled Down? *Economic and Political Weekly*, 46–55.

Fisher, A., M. Hanemann, M. Roberts, and W. Schlenker. 2012. The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment. *American Economic Review* 102 (7): 3749–60.

Fullerton, D., and G. Heutel. 2007. The General Equilibrium Incidence of Environmental Taxes. *Journal of Public Economics* 91 (3): 571–91.

Gornall, J., R. Betts, E. Burke, R. Clark, J. Camp, K. Willet, and A. Wiltshire. 2010. Implications of Climate Change for Agricultural Productivity in the Early Twenty-First Century. *Philosophical Transactions of the Royal Society B* 365:2973–80.

Guiteras, R. 2009. The Impact of Climate Change on Indian Agriculture, unpublished manuscript, University of Maryland.

Haisken-DeNew, J., and C. Schmidt. 1997. Interindustry and Interregion Differentials: Mechanics and Interpretation. *Review of Economics and Statistics* 79 (3): 516–21.

Hertel, T., M. Burke, and D. Lobell. 2010. The Poverty Implications of Climate-Induced Crop Yield Changes by 2030. *Global Environmental Change* 20 (4): 577–85.

Jacoby, H. 2000. Access to Markets and the Benefits of Rural Roads. *Economic Journal* 110 (465): 713–37.

———. 2013. Food Prices, Wages, and Welfare in Rural India. Policy Research Working Paper no. 6412 (April), World Bank, Washington, DC.

Jones, R. 1971. A Three-Factor Model in Theory, Trade, and History. In *Trade, the Balance of Payments, and Growth*, ed. J. Bhagwati. Amsterdam: North-Holland.

Kaufman, R. 1998. The Impact of Climate Change On US Agriculture: A Response To Mendelsohn et al. (1994). *Ecological Economics* 26:113–19.

Kovak, B. 2011. Local Labor Market Effects of Trade Policy: Evidence From Brazilian

Liberalization, unpublished manuscript, Carnegie Mellon University.

Lanjouw, P., and R. Murgai. 2009. Poverty Decline, Agricultural Wages, and Non-farm Employment in Rural India: 1983–2004. *Agricultural Economics* 40:243–63.

Marchand, B. 2012. Tariff Pass-Through and the Distributional Effects of Trade Liberalization. *Journal of Development Economics* 99:265–81.

Mendelsohn, R., W. Nordhaus, and D. Shaw. 1994. The Impact of Global Warming on Agriculture: A Ricardian Analysis. *American Economic Review* 84 (4): 753–71.

Narayanan, B., A. Aguiar, and R. McDougall, Editors. 2012. *Global Trade, Assistance, and Production: The GTAP 8 Data Base*, Center for Global Trade Analysis, Purdue, IN: Purdue University.

Parry, M., C. Rosenzweig, A. Iglesias, G. Fischer, and M. Livermore. 1999. Climate Change and World Food Security: A New Assessment. *Global Environmental Change* 9:51–67.

Parry, M., C. Rosenzweig, A. Iglesias, M. Livermore, and G. Fischer. 2004. Effects of Climate Change on Global Food Production under SRES Emissions and Socio-Economic Scenarios. *Global Environmental Change* 14:53–67.

Porto, G. 2006. Using Survey Data to Assess the Distributional Effects of Trade Policy. *Journal of International Economics* 70:140–60.

Rajeevan, M., J. Bhate, J. Kale, and B. Lal. 2006. High Resolution Daily Gridded Rainfall Data for the Indian Region: Analysis of Break and Active Monsoon Spells. *Current Science* 91 (3): 296–306.

Sanghi, A., R. Mendelsohn, and A. Dinar. 1998. The Climate Sensitivity of Indian Agriculture. In *Measuring the Impact of Climate Change on Indian Agriculture*, chap. 4. Washington, DC: The World Bank.

Srivastava, A., M. Rajeevan, and S. Kshirsagar. 2009. Development of a High Resolution Daily Griddeds Temperature Data Set (1969–2005) for the Indian Region. *Atmospheric Science Letters*. 10 (4): 249–54.

Stern, N. 2006. *The Economics of Climate Change: The Stern Review*. London: H.M. treasury.

Strauss, J. 1984. Joint Determination of Food Consumption and Production in Rural Sierra Leone: Estimates of a Household-Firm Model. *Journal of Development Economics* 14 (1): 77–103.

Timmins, C. 2007. If You Cannot Take the Heat, Get Out of the Cerrado...Recovering the Equilibrium Amenity Cost of Nonmarginal Climate Change in Brazil. *Journal of Regional Science* 47 (1): 1–25.

Topalova, P. 2007. Trade Liberalization, Poverty, and Inequality: Evidence from Indian Districts. In *Globalization and Poverty*, ed. A. Harrison, Chicago: University of Chicago Press.

World Bank. 2009. *South Asia: Climate Change Strategy*. South Asia region. Washington, DC: World Bank.

Appendix

A Model Solution

We assume Cobb-Douglas production functions with input cost shares $\alpha_{Li} + \alpha_{Ii} + \alpha_{Ki} = 1$ in each sector $i = A, M, S$.[34] The first step is to solve the following system of four equations

$$(A.1) \qquad \alpha_{LA}\widehat{W} + \alpha_{KA}\widehat{\Pi}_A = \widehat{\Theta} + \widehat{P}_A$$

$$\alpha_{LM}\widehat{W} + \alpha_{KM}\widehat{\Pi}_M = 0$$

$$\alpha_{LS}\widehat{W} + \alpha_{KS}\widehat{\Pi}_S = \widehat{P}_S$$

$$\beta_A\widehat{\Pi}_A + \beta_M\widehat{\Pi}_M + \beta_S\widehat{\Pi}_S = \widehat{W}$$

for $\widehat{W}$ and the $\widehat{\Pi}_i$ (recall, $\widehat{P}_M = \widehat{P}_I = 0$ by assumption). The first three equations are the sectoral price-equals-unit-cost conditions, whereas the last equation is derived from the labor constraint (which implies $\sum_i \beta_i \widehat{L}_i = 0$) and the fact that $\widehat{L}_i = \widehat{\Pi}_i - \widehat{W}$ in the Cobb-Douglas case. Solving the equation system (A.1) leads to

$$(A.2) \quad \widehat{W} = (\beta_A/\alpha_{KA}D)(\widehat{\Theta} + \widehat{P}_A)$$
$$+ (\beta_S\alpha_{KS}/D)\widehat{P}_S,$$

---

[34] Elasticity formulae can also be obtained using more general production functions (see e.g., Corden and Neary 1982; Kovak 2011), but these do not lend themselves to empirical implementation in our setting. For example, elasticities of substitution between inputs are hard to estimate; the Cobb-Douglas restricts all of these to unity.

where $D = 1 + \sum_i \beta_i\alpha_{Li}/\alpha_{Ki}$.

Next, we need to solve for $\widehat{P}_S$ by equating changes in service sector supply $\hat{Y}_S$ and demand $\hat{X}_S$. As discussed in the text, we assume a Marshallian aggregate demand function for services of the form $X_S = \eta y/P_S$, where η is the share parameter (i.e., Cobb-Douglas preferences). Differentiating, we obtain

$$(A.3) \quad \hat{X}_S = \widehat{y} - \widehat{P}_S$$
$$= (1 - E/y)(\omega_A\widehat{P}_A + \omega_S\widehat{P}_S) - \widehat{P}_S$$

where $\omega_j = (1 + \alpha_{Kj}/\alpha_{Lj})\beta_j/(1 + \sum_i \beta_i\alpha_{Ki}/\alpha_{Li})$. Note that the second line uses aggregate income equation (2) and the fact that with Cobb-Douglas production functions, sectoral labor shares are equivalent to sectoral *value-added* shares (another advantage of our functional form assumption from the point of view of empirical implementation).

On the supply side, from the services production function and the specificity of capital, we have

$$(A.4) \quad \hat{Y}_S = \alpha_{LS}\widehat{L}_S + \alpha_{IS}\hat{I}_S.$$

Meanwhile, the condition that input prices equal respective marginal value products delivers $\widehat{W} = \widehat{P}_S + \widehat{F}_{L_S} = \widehat{P}_S - \widehat{L}_S + \hat{Y}_S$ and $\widehat{P}_S = -\widehat{F}_{I_S} = \hat{I}_S - \hat{Y}_S$, where the second equality in each case follows from the total differentiation of the marginal product functions F_{L_S} and F_{I_S}. Solving these two equations, after first substituting out $\hat{I}_S$ from the second using (A.4), yields

$$(A.5) \quad \hat{Y}_S = \frac{\alpha_{LS} + \alpha_{IS}}{\alpha_{KS}}\widehat{P}_S - \frac{\alpha_{LS}}{\alpha_{KS}}\widehat{W}.$$

Substituting equation (A.2) into (A.5), equating the result to (A.3), and solving gives

$$(A.6)$$

$$\widehat{P}_S = \frac{\alpha_{KS}(1 - E/y)\omega_A D + \alpha_{LS}\beta_A/\alpha_{KA}}{D(1 - \alpha_{KS}(1 - E/y)\omega_S) - \alpha_{LS}\beta_S/\alpha_{KS}}$$
$$(\widehat{\Theta} + \widehat{P}_A),$$

which, inserted into equation (A.2), delivers a general expression for ψ. For the case of equal input cost shares across sectors, $D = 1 + \alpha_L/\alpha_K$ and equation (A.6) simplifies to $\widehat{P}_S = (\widehat{\Theta} + \widehat{P}_A)R\beta_A/(\alpha_K + \alpha_L - R\beta_S)$ where

$R = \alpha_L + \alpha_K(\alpha_K + \alpha_L)(1 - E/y)$. Substituting this expression for $\hat{P}_S$ into (A.2) delivers the ψ in equation (4).

To assess the validity of the equal cost shares assumption, we took cost shares of value-added for Indian manufacturing and service sectors based on national accounts (from Narayan, Aguiar, and McDougall 2012) and computed ψ according to the more general formula. The results are virtually identical to those of the equal cost shares case described in the text. While the labor cost share in manufacturing is, to be sure, much lower than that in agriculture (and in services), so is the capital share. Hence, the ratio of capital to labor shares, which is what is most relevant to our calculations, is very similar across sectors.

B Computation Details

Production function parameters

To compute the labor cost share, $WL_A / \sum_j P_j Y_j$, note that $L_A = H + \ell$, where H is hired labor in agriculture and ℓ is family labor. Note also that $WL_A = WH(1 + f)$, where $f = \ell/H$ is the ratio of family to hired farm labor. For a labor market in equilibrium, f should equal the ratio of the number of agricultural laborers working on their own farm to the number working for wages on other farms. Thus, we can calculate f for each of the five regions from individual employment data in NSS61-EUS. Comparable data on hired labor expenses (for regular and casual farm workers), WH, and on total crop revenues are available at the farm-level by season from NSS59. Summing up WH across seasons and households within each region (using sampling weights) multiplying by $(1 + f)$ and dividing by a similarly computed sum of crop revenues gives the labor share for each region as reported in table 3. A similar calculation gives the intermediate input share $P_I I_A / \sum_j P_j Y_j$, where the numerator is the total expenditures on non-labor variable inputs (seed, fertilizer, pesticide, and irrigation).

Household income shares

To get household labor income, we need employment information on household members, which is only available in the NSS61-EUS. Since households in NSS61-CES are not the same as those covered by NSS61-EUS, we must impute the requisite variables based on common characteristics available in both data sets. Thus, let n_h be the proportion of economically active household members (age 15–60) in NSS61-EUS. We first run a regression of n_h on the following explanatory variables: linear splines in the amount of land the household cultivates, household demographic composition and age of head, proportion of household members with different education levels, social category of household (scheduled caste/tribe), and district dummies. Using these regression coefficients, we impute a value of n_h^* for each household in NSS61-CES (conditional on having members in the age range) based on the same set of explanatory variables, dropping the small percentage of households for whom $n_h^* \leq 0$.

Daily wage income per worker is the weighted sum of (exponentiated) state fixed effects from a log-wage regression based on NSS-EUS and controlling for sector of employment (agricultural/nonagricultural), type of labor (manual/nonmanual), and NSS survey round. Weights are the imputed proportion of household days worked (in last week) in agricultural manual labor, nonagricultural manual labor, and non-manual labor using the same imputation procedure as for n_h^*. Thus, we have adjusted state mean daily wages $\overline{W}_{s,m}$, for $m = MAL$ (*manual ag. labor*), MNL (*manual nonag. labor*), NML (*non-manual*), and the corresponding household weights $d_{h,m}^*$. Annual per capita income from manual labor is then $WL_h^* = 365 n_h^* (d_{h,MAL}^* \overline{W}_{s,MAL} + d_{h,MNL}^* \overline{W}_{s,MNL})$ and from non-manual labor is $E_h^* = 365 n_h^* d_{h,NML}^* \overline{W}_{s,NML}$.

An analogous approach is used to generate state-specific prices for irrigated and unirrigated land, whose sum, weighted by, respectively, hectares of irrigated and unirrigated land owned per household member, constitutes the total value of land per capita. To get land income per capita, $\Pi_A K_{Ah}$, we multiply this latter figure by a discount rate of 0.05 (as in, e.g., Deschênes and Greenstone 2007). Finally, we compute $\lambda_{\Pi h} = \Pi_A K_{Ah}/(\Pi_A K_{Ah} + WL_h^* + E_h^*)$ and $\lambda_{Wh} = WL_h^*/(\Pi_A K_{Ah} + WL_h^* + E_h^*)$.
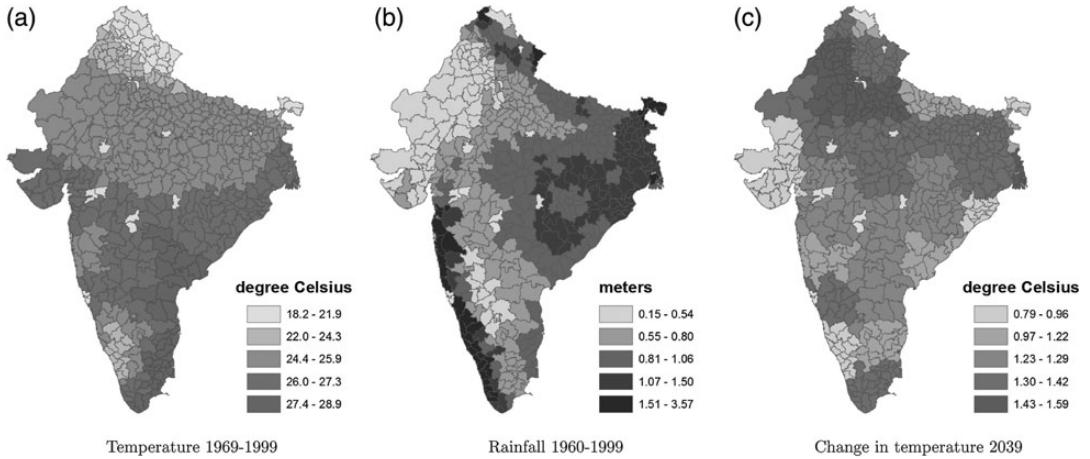
C Appendix Figures and Tables

Temperature 1969-1999   Rainfall 1960-1999   Change in temperature 2039

Figure C.1. Climate and Predicted Temperature Change across 18 Major States of India

TFP versus Price Changes
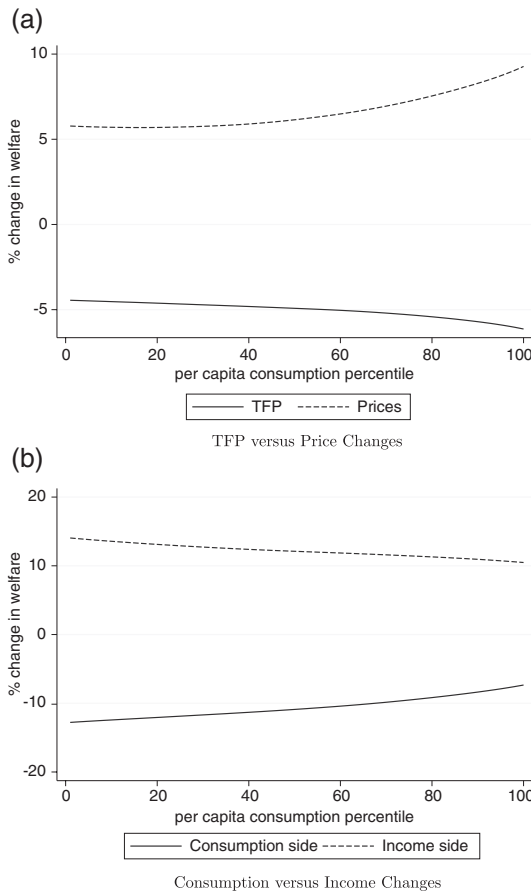

Consumption versus Income Changes

Figure C.2. Alternative Decompositions of Welfare Changes

Note: Smoothed relationship between mean (within-percentile) predicted welfare change, separated by source, and per-capita expenditure percentile. All decompositions are based on Low Global Productivity scenario and state fixed effects specification of hedonic model. Panel (a) decomposition separates welfare effects due to agricultural TFP changes $(\widehat{\Theta})$ and world price changes $(\widehat{P}_A)$. Panel (b) decomposition separates welfare effects due to changes in household consumption (including of nontradables) and changes in household income.
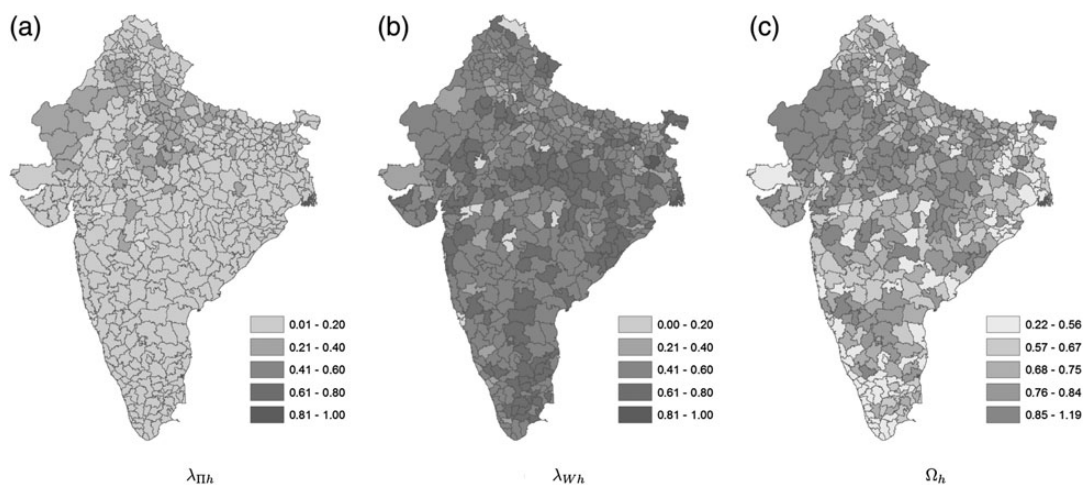
Figure C.3. Distribution of Key Parameters across 18 Major States of India

Table C.1. Data for Hedonic Regressions

Variable	Mean	Std. Dev.	Min	Max
Agricultural TFP index ($\bar{\theta}_d$)	−0.0142	0.387	−0.978	1.127
Mean annual temperature (°C)	25.33	1.800	18.24	28.85
Mean harmful degree days	13.18	16.53	0	71.47
Mean annual precipitation (meters)	0.941	0.473	0.150	3.571
Std. Dev. of mean annual precip.	1.289	0.595	0.299	4.595
Proportion area at elevation 250–1500 m	0.410	0.419	0	1
Proportion area at elevation >1500 m	0.0246	0.128	0	0.999
Dominant topography: Hilly	0.0581	0.234	0	1
Dominant topography: Undulating	0.432	0.496	0	1
Dominant soil texture: Fine	0.217	0.413	0	1
Dominant soil texture: Medium	0.634	0.482	0	1
Dominant soil type: Acrisols	0.0172	0.130	0	1
Dominant soil type: Cambisols	0.213	0.410	0	1
Dominant soil type: Luvisols	0.320	0.467	0	1
Dominant soil type: Nitosols	0.0925	0.290	0	1
Dominant soil type: Vertisols	0.181	0.385	0	1
Distance to nearest city of $>10^6$ pop. (km)	170.6	105.9	3.687	575.1

Note: Summary statistics for 465 districts. Source for city populations is the 2001 Census of India. All other sources mentioned in the text.