

Distributional Implications of Climate Change in India

Hanan Jacoby
Mariano Rabassa
Emmanuel Skoufias

The World Bank
Poverty Reduction and Economic Management Network
Poverty Reduction and Equity Unit
April 2011



Abstract

Global warming is expected to heavily impact agriculture, the dominant source of livelihood for the world's poor. Yet, little is known about the distributional implications of climate change at the sub-national level. Using a simple comparative statics framework, this paper analyzes how changes in the prices of land, labor, and food induced by modest temperature increases over the next three decades will affect household-level welfare in India. The authors predict a substantial fall in agricultural productivity, even allowing for farmer adaptation. Yet, this decline will not translate into a sharp drop in consumption for the majority of rural

households, who derive their income largely from wage employment. Overall, the welfare costs of climate change fall disproportionately on the poor. This is true in urban as well as in rural areas, but, in the latter sector only after accounting for the effects of rising world cereal prices. Adaptation appears to primarily benefit the non-poor, since they own the lion's share of agricultural land. The results suggest that poverty in India will be roughly 3–4 percentage points higher after thirty years of rising temperatures than it would have been had this warming not occurred.

This paper is a product of the Poverty Reduction and Equity Unit, Poverty Reduction and Economic Management Network. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The author may be contacted at eskoufias@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

Distributional Implications of Climate Change in India*

Hanan Jacoby

Mariano Rabassa

Emmanuel Skoufias

THE WORLD BANK

*The findings, interpretations, and conclusions of this paper are those of the authors and should not be attributed to the World Bank or its member countries. Contact Information: [hjacob@worldbank.org](mailto:hjacoby@worldbank.org); mrabassa@worldbank.org; eskoufias@worldbank.org

1 Introduction

Despite uncertainties as to how climate change will unfold in the coming decades, what seems 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, little is known about the consequences for the poverty of people within nations. There is simply a dearth of micro-level evidence on the distributional implications of climate change, evidence showing which households will suffer the greatest income losses and how far they will fall (or possibly rise) along the economic ladder.¹ Such information could be useful for targeting policies that ameliorate the negative impacts of warming or that promote adaptation to it.

There are a number of reasons to believe that the costs of climate change will not be borne equally across income groups. First, regional variation is expected in future warming as well as in the climate sensitivity of agricultural production, variation which could be correlated with average regional income. Irrigated areas, for instance, may be wealthier and, at the same time, less vulnerable to rising temperatures. Secondly, households differ in their ownership of agricultural assets, notably land, as well as in their human capital and allocation of labor. The returns on land and labor (of different types) will probably respond differently to climate change. Thirdly, households will be exposed in varying degrees to changes in food prices that are likely to occur with global warming. Urban households tend to be net buyers of food to a much greater extent than rural households (as are non-farm households within rural areas), leaving them relatively worse off when prices rise.

The aim of this paper is to quantify how these sometimes conflicting forces will play out in India over the next thirty years. Viewing global warming essentially as a change in the prices of land, labor, and food, we investigate the effects on household-level welfare using a simple comparative statics framework. Our analysis takes a step beyond conventional partial equilibrium comparative statics, though, by considering the welfare impact of these three future price changes *jointly*. The first part of the paper estimates the climate sensitivity of land and labor productivity in India. With these estimates and forecasts of future food prices changes in hand, the second part of the paper predicts the welfare consequences of climate change for households at each point along the initial welfare distribution.

Aside from its sheer size and importance as an agricultural producer, India provides an

¹Hertel et al. (2010) study the poverty impacts of climate-induced agricultural productivity change in 15 developing countries using computable general equilibrium (CGE) modeling. Our approach is distinct in its reliance on household data, which allows us to draw out more fine-grained and country-specific distributional implications.

ideal setting for the application of this methodology. With a wide range of agro-ecological zones, there is considerable variation in agricultural productivity and climate across the sub-continent. Also, Indian agriculture is expected to be especially hard-hit by warming temperatures in the coming decades, although not necessarily uniformly (World Bank, 2009). While agriculture contributes only about a fifth of GDP, productivity growth in this sector has been the major driver of poverty reduction in India, largely through increases in agricultural wages (Datt and Ravallion, 1998; Eswaran, et al. 2008; Lanjouw and Murgai, 2009). Data for India are relatively abundant as well, our main source being the nationally representative National Sample Surveys (NSS). The 2002-03 NSS round collects detailed farm-level information and the 2004-05 round includes both a huge household expenditure survey and an employment survey. These data are supplemented with a district-level panel of agricultural output and weather.

Past empirical literature focuses on predicting the impact of a changing climate on national or global agricultural product, largely ignoring the question of how damages will be distributed across the relevant population. There are three distinct approaches to making such predictions. Agronomists base crop damage assessments on temperature and, perhaps, precipitation responses taken from agricultural field trials. More sophisticated crop modeling (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 future farmers would behave. Economists prefer to use data that reflect actual farmer behavior. The hedonic or so-called Ricardian approach, pioneered by Mendelsohn, et al, (1994), infers the impact of climate change on future agricultural productivity from the present-day cross-sectional relationship between climate and land values. This allows for as much (but no more) adaptation than is revealed in the data. Predictions based on reduced-form Ricardian regressions presuppose that the technological envelope along which farmers will adapt in the future will be the same as it was in the past.² In response to perceived empirical deficiencies of such cross-sectional analyses, Deschênes and Greenstone (2007) propose using panel data on farm revenue to estimate the response of agricultural productivity to year-to-year weather fluctuations. However, since there is relatively little scope for adaptation to short-run weather, damage estimates based on shocks can only provide an upper bound on the long-run impact of climate change on agricultural productivity; there is no guarantee, moreover, that this bound will be very informative.

Our Ricardian estimates indicate that warming of about 1.25 °C expected to occur in

²The Ricardian approach also ignores the costs of adjusting along the technological frontier (see, e.g., Quiggin and Horowitz, 2003; Kelly et al., 2005).

India over the next three decades will reduce average land productivity across districts by 17 percent (on an area-weighted basis); our implementation of the weather shocks approach delivers the much higher figure of 37 percent, which suggests that farmer adaptation will have a considerable moderating influence. Results for rural wages imply that, in percentage terms, the return on labor will fall by about a third as much as the return on land in the long-run. So, although climate change will entail a significant productivity shock to the agricultural sector even accounting for adaptation, we will show that for the majority of rural households (who are not substantial land owners) the economic consequences will be muted. Finally, to obtain credible predictions of climate-induced changes in food prices, neither a Ricardian-type approach nor one based on weather shocks will do. Cereals, after all, are traded on world markets and are grown in many countries, each of which will be affected by climate change in different ways.³ We therefore base our forecast of future cereal price changes on two recent computable general equilibrium models of world demand, production, and trade, each of which takes into account the impacts of climate change on global agricultural production.

While compelling in its simplicity, comparative statics is accurate only for marginal price changes, which is one reason why we restrict our attention to the not-too-distant horizon of 30 years. If it turns out that irrigation, human capital, and other relevant investments respond sharply to the relatively modest climate change expected over this period, or that internal migration substantially increases,⁴ then our approximations may be off target, though it is far from obvious how such inaccuracy would affect our distributional findings. Climate change could also influence the processes of growth and structural transformation that are already well underway in the Indian economy. Evaluating the quantitative significance of such interactions involves multi-sector dynamic general equilibrium modeling. Although an interesting avenue for future research, such models are difficult to parameterize and not well-suited for tracing out detailed distributional implications.

On the agronomic side, 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

³While the Indian government intervenes extensively in the agricultural commodity trade in order to stabilize domestic cereal prices and is presently stockpiling enormous quantities of grain, insulating policies such as these are unlikely to be sustainable in the face of permanent world price increases.

⁴Timmins (2007) argues that climate change is nonmarginal and hence will induce significant movements of households across locations, but he considers temperature changes (in Brazil) that are about twice as large as those we assume for India. Eschewing hedonic methods, he specifies an econometric model of location choice as a function of climate identified off of permanent migration decisions of household heads. The limitation of this approach, for our purposes, is twofold: First, climate is assumed to have only amenity, and no productive, value, and, second, households 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.

do we consider the effects of changes in precipitation and weather variability. Although potentially important, the long-run evolution of each of these features is relatively uncertain and their productivity impacts (with the possible exception of those of precipitation) are hard to quantify.⁵ Our focus is on the direct impacts of higher temperatures, the most predictable consequence of rising atmospheric CO₂ levels. A final caveat concerns carbon fertilization effects. While elevated CO₂ 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 impacts can be viewed as worst-case scenarios of global warming.

The remainder of the paper proceeds as follows. In the next section, we lay out a simple household model and analytically derive the welfare consequences of climate change. Section 3 discusses the issues in estimating climate effects on endowment prices and section 4 reports the results of this estimation. The distributional analysis of climate change in India is presented in section 5 and the paper concludes in section 6 with a summary of the findings and a discussion of their implications for policy.

2 Framework

2.1 Basic household model

Households are assumed to have preferences over food and non-food consumption, the latter taken as the numeraire good. Real household per capita consumption is thus $c = z + px$, where z and x are, respectively, the per-capita quantities of non-food and food and p is the price of food. Denote the per-capita endowment of homogeneous farmland by a , which for urban households (as well as for many rural ones) may be zero. Farmland and the household's own labor are combined in fixed proportions, along with a variable quantity of hired labor h (per household member) to produce y units of food per capita with constant returns to scale. Restricting the household to producing only food crops is strictly for notational convenience and can be relaxed. As we will see shortly, however, the assumption of a Leontief production technology is consequential.⁶

⁵Quoting 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) for a similar conclusion. We discuss the effect of precipitation on agricultural productivity further in section 4.5.

⁶This assumption can be motivated by the observation that managing a farm of given size, including supervising the hired workers, requires the continual presence of family members, thus involving a relatively fixed time commitment (see Feder, 1985, and Eswaran and Kotwal, 1986, for models along these lines).

A real wage of w can be earned in the labor market. Let τ denote the per capita units of labor available to the household and let κ be the fraction of this labor used on the household's own farm, which, under our Leontief assumption, depends only on the fixed endowment of land a . Off-farm labor earnings are thus $w\tau(1 - \kappa)$ and total per capita income ι consists of these earnings plus net revenues from food production $py - wh$. Given CRS, $py - wh = \pi a$, where π is the net return on land, and $\iota = \pi a + w\tau(1 - \kappa)$. The maximized net return on land—i.e., once hired labor and any other variable inputs have been optimally chosen—is $\pi(p, w; \theta)$, a function only of prices and of the climate parameter θ , which enters via the production function. Note that π represents the cash flow per unit of land; in other words, it is gross of the cost of the household's own labor. The market value of land is just the expected discounted value of π over an infinite horizon. In equilibrium, the wage $w(\theta)$ and food price $p(\theta)$ may also depend on climate, which would induce indirect effects of θ on π due to the latter's dependence on input and output prices.

Now consider the compensation, in terms of units of real consumption, dc , required to make the household indifferent to a climate change of $d\theta$. It is straightforward to show that⁷

$$\begin{aligned} \frac{dc}{d\theta} &= \frac{d\iota}{d\theta} - \frac{dp}{d\theta}x \\ &= \left. \frac{d\pi}{d\theta} \right|_p a + \frac{dw}{d\theta}\tau(1 - \kappa) + \frac{dp}{d\theta}(y - x). \end{aligned} \quad (1)$$

This is the household's welfare loss (gain) from climate change, where the third term in the second line is precisely the expression derived by Deaton (1989) for the welfare effect of a food price change.⁸ The remaining terms, which do not enter Deaton's analysis, capture income effects due to the changing returns on land and labor.⁹ Though suppressed in the sequel, the additional notation on the profit function derivative $d\pi/d\theta$ indicates that we are holding the food price (but not the wage) constant.

⁷The first line uses the indirect utility function and follows from Roy's identity. The second line uses the envelope theorem.

⁸Since we will not be aggregating these welfare changes across households, we do not have to take a stand on the form of the social welfare function (i.e., which set of utility weights to apply).

⁹These terms resemble Roback's (1982) willingness-to-pay expression, but there are two key differences. First, in the hedonic literature, the price gradients reflect both the productive and amenity value of climate (i.e., θ is assumed to enter utility); equilibrium land values may be lower and wages higher in hotter areas because households must be made indifferent to living in such places. Second, the endowment component, a , does not account for the household's net position in the residential land market. We assume residential land is a small and relatively homogeneous component of the household budget, certainly a reasonable approximation in rural India. In sum, if the amenity value of climate is indeed important at the margin, then the interpretation of the price effect changes, but our welfare calculation and distributional analysis remains valid.

Note the role in our derivation of the Leontief assumption, which implies $d\kappa/d\theta = 0$.¹⁰ Although we are thus ruling out reallocations of labor from own farm to off-farm work in response to climate change for landed households, landless households can supply farm labor elastically to a given market through (temporary) migration. Heuristically, if we differentiate landless households by their mobility cost, the supply curve of agricultural labor to a specific location is essentially the cdf of the mobility cost distribution (later we extend the model to allow intersectoral as well as geographical labor mobility). Under perfect labor mobility, each such supply curve would be horizontal and $dw/d\theta$ would have to be zero. As we will see, the null of perfect labor mobility is strongly rejected in the data.

Dividing both sides of (1) through by c (or ι , as appropriate), we get

$$\frac{d \log(c)}{d\theta} = \lambda \frac{d \log(\pi)}{d\theta} + (1 - \lambda) \frac{d \log(w)}{d\theta} + \frac{p(y - x)}{c} \frac{d \log(p)}{d\theta} \quad (2)$$

where $\lambda = \pi a / (\pi a + w\tau(1 - \kappa))$ represents the proportion of income derived from owned land and $p(y - x)/c$ is what Deaton (1989) terms the *net consumption ratio*. The potentially heterogeneous welfare consequences of climate change can now be clearly seen. To the extent that the returns to land are disproportionately affected by rising temperature, high λ households will suffer more. On the other hand, if the price of food also rises, these same households reap a benefit insofar as they are net food producers. Urban households, by contrast, typically have $\lambda = 0$ and are net consumers of food, whereas urban wages are likely to be less sensitive than rural wages to changes in climate.

2.2 Extensions

The basic story can be enriched along two dimensions. First, suppose that land is now of two types, irrigated (I) and nonirrigated (N), such that $a_I + a_N = a$, with irrigated being the more profitable; i.e., $\pi_I > \pi_N$. With constant returns, the only way to get households to hold both types of land at the same time is to posit an increasing private annualized (per capita) cost of irrigated land, $\zeta(a_I)$, which one may think of as a conversion cost. The optimal choice of a_I will equate the per acre return differential $\pi_I - \pi_N$ to the marginal conversion cost ζ' . The comparative statics of climate change is as before with $\frac{d\pi}{d\theta}a$ in equation (1) replaced by $\frac{d\pi_I}{d\theta}a_I + \frac{d\pi_N}{d\theta}a_N$. One does not need to know by how much households will switch into (or out of) irrigated land in response to climate change as the cost of any such compositional shift will just balance the benefit at the margin (again the envelope theorem). The formula for

¹⁰If π were redefined as net of the cost of family labor, then terms involving $d\kappa/d\theta$ would envelope away and we would not have to invoke the Leontief assumption. In this case, however, π would no longer be interpretable as cash flow off land. Consequently, the link between π and land values, central to the empirical work, would be severed.

$\frac{d \log(c)}{d \theta}$ is the same as (2) with the replacement of $\frac{d \log(\pi)}{d \theta}$ by

$$\varphi \frac{d \log(\pi_I)}{d \theta} + (1 - \varphi) \frac{d \log(\pi_N)}{d \theta} \quad (3)$$

where $\varphi = \pi_I a_I / (\pi_I a_I + \pi_N a_N)$ is the share of irrigated land in total income from land.

A second extension is to distinguish labor by sectors, wherein wages have differing climate sensitivities. Take agricultural and non-agricultural jobs, for instance. Wages in the former sector, say w_A , are presumably more dependent upon climate than wages in the latter, w_{NA} . Suppose, that entry into non-agricultural occupations requires costly human capital investment and hence that labor is imperfectly substitutable across sectors. Let e be the proportion of household wage-earners in the non-agricultural sector so that total wage income is $((1 - e)w_A + ew_{NA})\tau(1 - \kappa)$. If the household's private annualized per capita cost of education is $\psi(e)$, then the income-maximizing allocation of active household members across sectors must satisfy $\psi'(e) = (w_{NA} - w_A)\tau(1 - \kappa)$. As before, we would replace $\frac{d \log(w)}{d \theta}$ in equation (2) by

$$\sigma \frac{d \log(w_{NA})}{d \theta} + (1 - \sigma) \frac{d \log(w_A)}{d \theta} \quad (4)$$

where $\sigma = ew_{NA} / ((1 - e)w_A + ew_{NA})$ is the share of skilled labor income in total labor income.

Inspection of the last two equations reveals that when the marginal effect of climate on log endowment prices are identical across sectors, the overall expression for the percentage change in consumption reverts to formula (2). In general, however,

$$\lambda = \frac{\pi_I a_I + \pi_N a_N}{\pi_I a_I + \pi_N a_N + ((1 - e)w_A + ew_{NA})\tau(1 - \kappa)}. \quad (5)$$

3 Estimating Changes in Endowment Prices

Since land is immobile, cross-sectional variation in its market value π reflects differences in its productivity, whether due to climate θ or other agronomic factors; this is the simple logic behind the Ricardian approach. To use Ricardian regressions as forecasting equations, however, one has to assume in addition that the menu of technologies available in the coming decades will be the same as that which exists in the current cross-section. Labor presents an additional wrinkle, as already alluded to. Insofar as workers are mobile across sectors or regions, wage differentials reflect variation in both supply and demand. In order to use the empirical relationship between w and θ to forecast the effects of climate change on wages, one must also assume that the elasticity of supply of labor to each local labor market will remain

the same. But this assumption may not be tenable if temperatures increase across the board and inter-regional labor migration is high. Since such migration flows may respond very little to a nationwide warming that does not alter relative wages across regions, local labor supplies in the future may be less elastic than they have been in the past. Consequently, the present cross-sectional climate-wage differential could understate the wage response to a generalized warming. On the other hand, since labor mobility within India has historically been quite low (see, e.g., World Bank, 2010), this underestimation may not be too serious.

3.1 Irrigation in the Ricardian framework

Within the Ricardian literature, the treatment of irrigation has generated considerable debate (see Cline, 2007, for an overview). Should one control for irrigation in a regression of farm profits or land values on climate, and how? Schlenker et al. (2005) answer the first question in the affirmative; irrigated land is not only more productive than rainfed land, but also differs in its responsiveness to climate variation. However, the coefficient on θ in a Ricardian regression in which irrigation infrastructure is also included as a control asks a specific question: What is the impact of climate change on farm productivity holding the stock of irrigation fixed? This question is relevant only if the stock of irrigation is indeed expected to remain fixed. Otherwise, an expansion of irrigation infrastructure must be taken into account in forecasting the impact of climate change on future agricultural productivity.¹¹

Since our objective is merely to estimate the separate climate sensitivities of irrigated and unirrigated land, there is no need to model the future trajectory of irrigation. To see why this is so, while laying out the empirical methodology, consider a regression based on a large sample of individual plots of land p located in different cross-sectional units (e.g., districts) i of the form

$$\log(\pi_{pi}) = \alpha\theta_i + \beta\theta_i \times I_{pi} + \gamma I_{pi} + \delta Z_i + \mu_i + \zeta_{pi} \quad (6)$$

where $I_{pi} = 1$ if plot p is irrigated (0 otherwise), and μ_i and ζ_{pi} are error terms representing, respectively, common geographic-level and idiosyncratic plot-level unobservables. Local climate, θ_i , affects land productivity directly and interactively with irrigation; for simplicity, other agronomic factors, Z_i , are assumed to enter only directly. Notice that the climate sensitivity of overall land productivity, $E [d \log(\pi) / d\theta] = \alpha + \beta \Pr(I = 1) + (\beta\theta + \gamma) d \Pr(I = 1) / d\theta$,

¹¹Kurukulasuriya and Mendelsohn (2007) treat irrigation as an endogenous variable in a Ricardian framework, using irrigation potential (water flow) as an instrument. If the goal is simply to estimate the overall effect of climate on farm productivity, however, structural modeling has no advantage over a reduced-form that controls for irrigation potential directly.

involves the forecast of future changes in irrigation infrastructure, $d\Pr(I = 1)/d\theta$, whereas the quantities in equation (3), $d\log(\pi_N)/d\theta = E[d\log(\pi)/d\theta|I = 0] = \alpha$ and $d\log(\pi_I)/d\theta = E[d\log(\pi)/d\theta|I = 1] = \alpha + \beta$, do not.

Equation (6) can be estimated in two stages. In the first stage, a fixed effect regression removes all variation ($\alpha\theta_i + \delta Z_i + \mu_i$) common to plots within the same geographical unit i . Under the assumption that $E[I_{pi}\zeta_{pi}] = 0$, this regression identifies β (and γ).¹² The second stage takes the estimate of the geographic fixed effect, $\alpha\theta_i + \widehat{\delta Z_i} + \mu_i$, and regresses it on θ_i and Z_i . This procedure recovers α provided that $E[\theta_i\mu_i] = 0$ and $E[Z_i\mu_i] = 0$. These are the standard identifying assumptions implicit in cross-sectional Ricardian estimation, to be discussed further in the next subsection.

3.2 Panel versus cross-sectional data

Deschênes and Greenstone (2007) advocate a panel data fixed effects estimator over the cross-sectional Ricardian regression, arguing that geographical fixed effects control for all local characteristics that may be correlated with climate. With data on (net) farm revenues per acre, denoted by R_{it} for cross-sectional unit i in year t , the regression model, suppressing any irrigation interactions, is

$$R_{it} = \alpha\theta_i + \eta\xi_{it} + \delta Z_i + \omega_i + v_{it}, \quad (7)$$

where ξ_{it} is local weather realized in year t , and ω_i and v_{it} are permanent and transitory error components, respectively. Departing from the Ricardian assumption, ω_i may now be correlated with θ_i (and with Z_i).

Since, in the long run, farmers can make a myriad of adjustments to changes in climate that would be infeasible or uneconomical in the short-run, we expect $\alpha < \eta$. Purging ω_i by fixed effects, and thus exploiting year-to-year weather variation, provides an unbiased estimate of η ; this estimate can then be taken as an upper bound for α , the parameter of true policy interest. Whether this upper bound is informative, however, depends on how large η actually is, since a large η would be as consistent with a large α as with a small one.¹³ Put another way, η represents a worst-case scenario in which virtually all avenues

¹²The identifying assumption is violated if, within the geographical unit, irrigation is more likely to be put in place on land of high (unobserved) fertility. Since placement of irrigation largely depends on the location of rivers or aquifers, as the case may be, this omitted variable bias is likely of second order importance.

¹³Applying this methodology to U.S. agriculture, Deschênes and Greenstone find that η is a relatively small number, implying that α must, in turn, be close to zero. After correcting significant errors in the original data, Fisher et al. (2009) challenge this finding as well as the conclusion that cross-sectional Ricardian estimates lack robustness. As we discuss below, Guiteras (2009) applies the weather shocks approach to panel data from Indian districts.

for adaptation to climate change have been foreclosed.¹⁴

4 Climate Impacts on Prices

4.1 Climate data

The NSS household data that we describe below were collected using clustered sampling, but the precise location of primary sampling units has not been released. Thus, the district is the lowest level of geographic disaggregation at which we can link household and climate variables. By 2003, India had 576 districts, excluding those in Lakshadweep, Andaman and Nicobar Islands, and in the mountainous northern state of Jammu/Kashmir, with an average land area of around 5000 km².

We use monthly temperature readings from 370 weather stations throughout India with at least 20 years of records during 1951-1980. These data are available from the India Meteorological Department (IMD), which also provides a high resolution daily gridded (on 1° latitude by 1° longitude cells) rainfall dataset covering 1951-2003 based on daily records from more than 1800 weather stations.¹⁵ Normal precipitation during the 1960-2000 period for a district are interpolated from the 296 cells covering India using the proportion of the district's area in each cell as weights. District-level temperature normals are an average of records from the three nearest stations to the district's geographical center, weighted by the inverse of the squared distance to the district centroid.

Figure 1 illustrates the tremendous variation in climatic conditions across India, ranging from the desert-like western Rajasthan to the eastern foothills of the Himalayas, one of the wettest regions on earth. Most 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 depends to a much greater extent on irrigation.

4.2 Land values and climate

Our cross-sectional Ricardian analysis uses reported land values collected from a farm household survey as a measure of π . In principle, the value of land should equal the expected

¹⁴Ideally, one would want a land productivity panel spanning many decades over which *climate* has actually changed, and changed to different extents across regions, while the menu of technologies has remained relatively stable. Such data are, at present, unrealistic.

¹⁵Since the IMD release does not extend past 2003, weather data for our panel analysis are taken from the Center for Climate Research, University of Delaware. Though available at a higher grid resolution (0.5° × 0.5° cells), these data are interpolated using much fewer weather stations than those of IMD.

present discounted cash flow from its optimal use. Since land values are not influenced by realized weather or prices at the time of the survey, they should more accurately reflect long-term climate and other agronomic characteristics than measures of current land productivity such as yields or net revenues. While plot values are self-assessed rather than based on records of actual market transactions, farmers are generally knowledgeable about the value of this their primary asset. Indeed, evidence from elsewhere in south Asia (Jacoby, 2000) indicates that self-assessed plot values accurately capture rental income 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.

Data on the value of more than a hundred thousand individual parcels of cropland nationwide are available in the NSS 59th round, along with area and irrigation status (see Table 1 for descriptive statistics at the district level).¹⁶ Figure 2 shows how median land values per hectare (adjusted for regional cost-of-living differences) vary over the Indian subcontinent. The highest valued cropland lies across the Indo-Gangetic plain, the most intensively farmed zone in the country, whereas land values are lowest in the far-western desert region.¹⁷ Comparison of the two maps in Figure 2 shows a distinct positive relationship between land values and irrigation, one strongly confirmed in the micro-data to which we now turn.

In Table 2, we report alternative specifications of equation (6) using different sets of explanatory variables. Since our data do not identify the main source of irrigation for a plot (i.e., canal or groundwater), we can only estimate an average return to irrigation. Model I, at any rate, assumes that this return is constant across districts. The resulting estimates (first column) indicates that, on average, an irrigated plot is 37% more productive than a rainfed plot within the same district. For the second stage, we regress the 542 district fixed effects on district-specific determinants of land productivity, weighting by the number of plot observations per district. Standard errors are calculated using Conley’s (1999) spatial GMM estimator with bandwidth set at 8 degrees lat./long. (≈ 800 km).

The second stage of Model I does not control for any geographic variables apart from climate, specifically average annual temperature and precipitation (see Table 1 for descriptive statistics). In addition, we allow weather variability, in the form of the coefficient of vari-

¹⁶Aside 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 percent 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 village common areas.

¹⁷Districts with high land values correspond fairly closely to those with high median net revenues per hectare in the 2002-03 crop year based on farm-level data from the 59th round of the NSS (The district-level correlation coefficient is 0.48). As noted, however, net revenues are affected by weather realizations and in 2002-03 India experienced an extremely poor monsoon, though rainfall in the northeast was close to normal.

ation of annual rainfall, to influence land valuations.¹⁸ Including a more detailed array of temperature and precipitation variables (e.g., monthly averages along with quadratic terms in each), as is common practice in the Ricardian literature, has the potential disadvantage of multicollinearity, leading to lack of interpretability and robustness (see Kaufman, 1998, for a critique along these lines). The results in the second column imply that a 1°C increase in average annual temperature would lead to nearly a 10 percent reduction in land values in India. Neither precipitation levels nor volatility thereof appear to be strongly associated with land productivity.

Model II allows climate to influence the return to irrigation and includes other geographic controls in the second stage. Evidence reported in column three of Table 2 suggests that the return to irrigation is lower in wetter districts, which seems plausible. Temperature, however, does not interact to a significant extent with irrigation according to our results. The inclusion of controls for topography, elevation, soil texture and type,¹⁹ as well as distance to nearest large city, a proxy for transport cost to markets, boosts the adjusted R² in the second stage from 0.124 to 0.508. Most of the associated coefficient estimates also make economic sense, with flatter terrain, medium soils, and closer proximity to markets all raising land values (Jacoby, 2000, reports the latter finding as well). Although a far richer model, the resulting point estimate of the temperature effect changes surprisingly little, rising in magnitude by only about a third to around -13 percent. This pattern suggests, albeit informally, that any further omitted variable bias is not likely to be too serious (see Altonji, Edler, and Taber, 2005).²⁰

In other (unreported) specifications, we include quadratic terms in temperature and rainfall and we also separate climate variables into aggregates over the two main growing seasons, kharif and rabi, none of which leads to appreciable improvements in fit. What is found, though, to be statistically significant and positive in the second stage is the interaction be-

¹⁸Variability could influence land values through at least three mechanisms: (1) rainfall affects contemporaneous cash flow from land in a nonlinear fashion so that expected cash flow is a function of the rainfall variance as well as the mean; (2) in making investments in agricultural assets, risk-averse farmers trade off high expected returns for safety so that productivity per acre is lower where weather variability is greater (see Rosenzweig and Binswanger, 1993, for some evidence from India); and, (3) under risk-aversion, land prices exceed the expected present discounted cash flow from land by a risk premium that is increasing in weather variability.

¹⁹See Table 1. Data on soils come from the Food and Agriculture Organization’s Soil Map of the World (SMW), which provides detailed data for each soil mapping unit, defined as a geographical area with similar soil classes. Dummies are created for the five dominant soil types in India (acrisols, cambisols, luvisols, nitosols, and vertisols). In addition, there are three textural classes (coarse, medium, and fine), which reflect the relative proportions of clay, silt, and sand in the soil, and three slopes classes (level to gently undulating, rolling to hilly, and steeply dissected to mountainous).

²⁰Sanghi, Mendelsohn, and Dinar (1998), applying a cross-sectional Ricardian approach to pooled district-level panel data from India for the period 1966-86, find a comparable temperature effect on net farm revenues of -9 percent.

tween annual temperature and precipitation (see Model III in Table 2). Figure 3 displays the resulting distribution of temperature effects on land values, based on a 1°C increase, evaluated at the sample proportion of plots irrigated in each district. The wide range of impacts mirrors, of course, the huge spatial variation in rainfall. At the left tail of precipitation distribution are districts in the arid northwest of India (in the states of Rajasthan, Haryana, and Punjab), with temperature-productivity effects exceeding –18 percent in magnitude. At the right tail, a mere seven very wet districts are predicted to slightly benefit from higher temperatures, although extrapolations this far out along the rainfall distribution should probably be taken with a pinch of salt.

4.3 Agricultural output and weather shocks

Panel data covering the production of the 20 most important crops in terms of cultivated area are available for nearly 500 districts from 1999 to 2005 (though not always continuously).²¹ We aggregate crop-wise output using state-level prices averaged over the period 2001-03 and divide by area cultivated with these crops to get gross revenue per hectare. Annual cultivation cost data are not available at the district level, so we cannot calculate net revenue, the empirical counterpart of π . However, since the sensitivity of gross revenue to bad weather shocks, in general, *overestimates* the corresponding weather sensitivity of net revenue,²² our results should still provide an upper bound on long-run damages. Table 3 reports descriptive statistics for the panel data.

How much intertemporal weather variation is available to identify productivity impacts of climate change? Consider the residuals from regressions of annual average temperature and total annual rainfall on district dummies and either (1) a linear year trend or (2) year dummies. Figure 4 plots temperature against precipitation residuals for both specifications of year effects and shows a strong negative correlation; wetter than average years in a particular district, not surprisingly, tend to be cooler than average as well. Also, while most intertemporal rainfall variation falls within a generous plus or minus 500 mm/year band around the mean, temperature variation is relatively confined. With year dummies removed, almost all residual variation in temperature, 99 percent to be precise, lies inside a $\pm 0.5^\circ\text{C}$ range, whereas around 84 percent lies within $\pm 0.25^\circ\text{C}$. Residuals are a bit more dispersed for the more restrictive linear time trend specification, with 97.5 (75) percent lying

²¹The crops are rice, wheat, jowar, bajra, cotton, maize, groundnut, sugarcane, urad, ragi, barley, gram, arhar, moong, rapeseed/mustard, potato, coconut, sunflower, soyabean, and sesamum. These constitute around 90% of cultivated area over this period.

²²Using our earlier notation, we have $py/a = \pi + wh/a$, and by the envelope theorem, $d(py/a)/d\xi = d\pi/d\xi + (dw/d\xi)h/a$. If ξ is, say, temperature, and a positive temperature shock depresses production and wages, then $d(py/a)/d\xi < d\pi/d\xi$.

in the wide (narrow) band. Clearly, then, the productivity impact of temperature shocks greater than half a degree must be inferred by extrapolation out of our sample. Yet, most climate change scenarios for India over the next 30 years predict two and a half times this much warming.

With this caveat in mind, we proceed to the panel estimates in Table 4. The first two columns report annual weather shock impacts on gross revenue per hectare controlling for, respectively, a linear year trend (quadratic and cubic terms were not significant) and for year dummies. There is no compelling reason to choose one of these specifications over the other a priori, since macro shocks affecting land productivity are unlikely to be correlated with weather (recall that, by construction, our gross revenue measure is invariant to intertemporal crop price movements). Be that as it may, the estimates show that a 1°C increase in annual temperature reduces gross productivity per hectare by 24 percent with a year trend and by 31 percent with year dummies.

Working with year dummies, we next disaggregate temperature and rainfall into their seasonal averages (or totals in the case of rainfall). Since kharif and rabi season temperature shocks are much less intercorrelated than the corresponding seasonal temperature normals, it should be easier to distinguish seasonal temperature effects with panel data than with cross-sectional data. In the event, we find that higher kharif temperatures reduce annual productivity by nearly twice as much as higher rabi temperatures (column 3 of Table 4). Together with the fact that the value of kharif output is typically double that of rabi output, this result implies that the temperature sensitivity of kharif and rabi production are roughly equal.²³ Our evidence also shows that rabi precipitation, relatively minor as it is, has no significant productivity impact, hence its exclusion from all subsequent specifications. Quadratic terms, shown in column 4, are also not significant for seasonal temperature, but there are diminishing marginal returns to kharif season precipitation.

Results in column 5 of Table 4 allow for interactions of seasonal weather variables with the proportion area irrigated in the district.²⁴ Consonant with the cross-sectional findings, the

²³Denoting gross annual revenue by Y , revenue per hectare by y , and temperature by T , one can show that

$$\frac{d \log(y)/dT_K}{d \log(y)/dT_R} = \frac{Y_K}{Y_R} \frac{d \log(y_K)/dT_K}{d \log(y_R)/dT_R}$$

where subscripts denote kharif (K) and rabi (R) specific values of the relevant variable. Our estimate of Y_K/Y_R comes from the 59th round of the NSS, which gives gross revenue by season at the household level for the 2002-03 crop-year. Since the left-hand side is around two, as is Y_K/Y_R , $d \log(y_K)/dT_K$ must be close to $d \log(y_R)/dT_R$.

²⁴This variable is calculated from the 59th round of the NSS, as these data have the most complete district-wise coverage and are measured at the midpoint of the relevant period. The correlation between this irrigation ratio and a corresponding one calculated from Ministry of Agriculture data on the subset of districts for which the latter is available is about 0.8.

marginal benefits of higher (kharif) rainfall are lower in more intensively irrigated districts. Unlike the cross-sectional results, however, there is a significant and negative interaction between temperature and irrigation, at least in rabi season. This has a straightforward interpretation: Since irrigation is essential for most rabi cultivation, any impact of rabi season temperatures must be concentrated in heavily irrigated (i.e., rabi-cultivated) districts. Based on this specification, a 1°C increase in both kharif and rabi temperature would reduce output per hectare by around 29 percent for the average district. Damages would vary, however, from -21 percent in nonirrigated districts to -35 percent in fully irrigated districts. Lastly, we reconsider temperature-precipitation interactions in column 6. Though not as overwhelming as in the cross-section, the panel evidence suggests that the negative impact of higher temperatures in kharif season is, once again, attenuated in wetter districts.

Guiteras (2009) also estimates the impact of weather shocks on gross agricultural productivity in India, using district-level panel data for 1960-99. Besides the less recent production data, Guiteras' productivity index consists of just six major crops, compared to our twenty, and his sample drops most of eastern India as well. On the other hand, Guiteras uses daily gridded temperature data (not available to us) from which he constructs growing season degree-days, an arguably more relevant measure of weather than average annual or seasonal temperature. In any event, Guiteras' sophisticated econometric analysis uncovers a much smaller response to warmer temperatures than we do, even for the latter half of his sample period.²⁵ If our cross-sectional estimates based on land values are to be believed and the long-term decline of agricultural productivity due to a 1°C increase in temperature is indeed near 13 percent, then Guiteras' damage estimate of less than 11 percent is implausibly small, implying virtually no scope for adaptation. By contrast, our weather shock-based damage estimate is two and a half times that of Guiteras and is, therefore, consistent with considerable farmer adaptation to warmer temperatures in the long-run.

4.4 Rural wages and climate

The methodology for estimating the climate sensitivity of rural wages by employment sector follows closely the approach used for land values. Drawing upon the NSS employment-unemployment survey of 2004-05, the best source of labor data in India (Himanshu, 2005), we estimate the analog of equation (6) on a sample of over 36 thousand rural workers from 556 districts. We take the individual's log daily wage (adjusted for regional cost-of-living

²⁵As Guiteras points out, this was a period of increasing adoption of high-yielding varieties, which are more vulnerable to bad weather. Production data from later years would presumably exhibit even greater weather sensitivity.

differences) in the primary occupation of the last week as the dependent variable.²⁶ Aside from district fixed effects, the first-stage includes gender, a cubic polynomial in age, education dummies, and gender-education interactions. The climate interactions of interest are, in this instance, those with sector of employment. In particular, we distinguish cultivation work from non-farm tasks, since the return to the former activity is arguably more climate sensitive than the return to the latter. Almost exactly half of our sample have their main job in cultivation and these workers are on the whole substantially less educated than their counterparts engaged in non-agricultural tasks.

Table 5 reports the same sequence of specifications as in the case of land values, beginning with (Model I) no interaction terms. Holding constant observable worker attributes, we see that non-agricultural jobs pays 25 percent higher wages on average than agricultural jobs within the same district. The second stage for Model I shows that a 1°C increase in average annual temperature reduces rural wages by 5 percent.

In addition to interactions between job type and climate variables, Model II introduces other determinants of land productivity into the second stage regression.²⁷ The positive coefficient on the first-stage interaction terms indicate that non-agricultural wages are less sensitive to higher temperature than are agricultural wages, as might be expected. The difference, however, while statistically significant, is not great (1.2 percentage points for every 1°C increase in temperature). Evidently, there is sufficient substitutability among rural workers that, in the long run at least, areas of low agricultural productivity have relatively low wages outside cultivation as well.

Model III, as before, adds temperature-precipitation interactions in both the first and second stages. While we find a significant difference in the coefficients on the interaction across sectors, within each sector the interaction effects are weak, both economically and statistically. Absent conclusive evidence that the impact of temperature on wages varies with precipitation, we revert to the constant temperature effect specification of Model II in the sequel.

²⁶From the original sample of about 43 thousand observations, we drop workers in public administration/defense, education, and health/social work, as these are largely higher paying public sector jobs in which wages are not set strictly by market forces. Nevertheless, the results on this larger sample are very similar to those reported in the paper.

²⁷Among these are district-level measures of irrigation potential, both at the surface (km of perennial rivers per km²) and below it (m³ of annual replenishable groundwater per km²). Note that, in the case of land values, we controlled directly for irrigation at the plot level in the first stage, hence there was no need to include irrigation potential in the second stage.

4.5 Discussion

Taken together, our cross-sectional results imply a long-run elasticity of rural wages with respect to land productivity, $\epsilon_{w,\pi} = d\log(w)/d\log(\pi) = \frac{d\log(w)}{dT} / \frac{d\log(\pi)}{dT}$, of around 0.35. Badiani (2009) examines the response of unskilled wages (across sectors) to long-run changes in exogenous agricultural productivity using a district-level panel constructed from successive rounds of the NSS employment-unemployment survey dating back to 1983. Reassuringly, she obtains an estimated wage elasticity remarkably close to 0.35.

We do not directly estimate the impact of weather *shocks* on labor productivity, primarily because reliable, high frequency, and relatively recent data on rural wages in India is lacking (see Himanshu, 2005, for a critique of alternative sources). However, we obtain an indirect estimate using $\Delta \log(w) \approx \epsilon_{w,\pi} \frac{d\log(Y)}{dT} \Delta T$, where $\frac{d\log(Y)}{dT}$ is the output response to temperature shocks already estimated. Using this formula, a 1°C increase in temperature, for example, would reduce agricultural wages by about 10 percent in the absence of adaptation.

Finally, a word about precipitation. Except in the land value regressions, where its direct effect is insignificant, we find that higher rainfall is associated with greater agricultural productivity. Putting aside the enormous uncertainty involved in forecasting precipitation patterns over the Indian subcontinent mentioned earlier, whatever consensus has emerged points to a very modest rise in rainfall over the next three decades. Guiteras (2009), for one, uses the most recent IPCC projection of a 4 percent increase. Meanwhile, temperatures across India over this period are expected to warm by around 1.25°C on average (see below). Based on our estimates in Table 4, we predict that a 4% increase in rainfall (at the district area-weighted means of the data) would raise agricultural output per hectare by at most 1.0%, a figure dwarfed by the 37% decline in yields implied by the concomitant rise in temperature (i.e., weighted average of $\frac{d\log(Y)}{dT} \Delta T$). Results in Guiteras (2009) and Sanghi, Mendelsohn, and Dinar (1998) also support the view that climate change will overwhelmingly affect agricultural productivity in India through higher temperatures.

4.6 Cereal prices and climate

Since the main cereals produced and consumed in India, rice and wheat, are traded on both domestic and world markets, it would be nonsensical to use the cross-district correlation between current prices and climate (let alone weather) to forecast the effect of climate change on cereal prices 30 years hence. Even under the arguable proposition that, up to now, domestic grain markets have not been well integrated so that districts are for all practical purposes autarkic, such a district-level analysis would tell us little about how cereal prices would change in the future. This is because, despite short-run government attempts to

maintain price stability, cereal prices in the longer term will depend upon supply and demand, and hence climate, in the rest of the world as well as in India as a whole.

Results from two separate multi-sector computable general equilibrium models of the world economy, GTAP (Hertel, 2010) and ENVISAGE (van der Mensbrugghe, 2009), suggest that cereal prices will rise by around 10 percent by 2040 due to climate change. We should emphasize, however, that such forecasting is far from an exact science; indeed, since no standard errors are attached, the degree of prediction uncertainty cannot even be quantified. For GTAP, in particular, we are averaging price changes across a number of scenarios, including one in which the prices of all major cereals actually fall. Suffice it to say that our point estimate of a 10 percent increase is contained within a wide confidence interval.

5 Distributional Implications of Climate Change

5.1 Climate change scenario for India

We use temperature predictions from HadCM3-A1F1 (high emissions path), a commonly used climate model developed by the Hadley Centre for Climate Prediction and Research as part of the Third Assessment Report of the United Nation’s Intergovernmental Panel on Climate Change (IPPC). Our measure of warming is based on the mean monthly change in surface temperature between the baseline period 1961-1990 and the near future 2010-2039. Most other widely used climate scenarios predict very similar temperature increases in the short-run and only begin diverging significantly toward the end of the century.

Figure 5 displays the geographic distribution of projected annual temperature changes, as well as those for each agricultural season, kharif and rabi, separately. Northern India will be hardest hit by warming in the coming three decades, particularly during kharif, whereas much of the south and east will be spared the worst of the temperature increase in the principal growing season.

5.2 Putting the pieces together

With predicted temperature changes and estimated price responses in hand, it remains only to calculate the weights and the other components of the formulae of section 2 to implement the distributional analysis. This is easier said than done, however, in that households in the NSS61 expenditure survey, for whom we have consumption data, are not the same as those in the 61st round employment-unemployment survey, nor are they the same as those in the 59th round farm household survey. We must, therefore, impute several key parameters based on observable characteristics common across surveys.

- i) py/c : Computing the net (cereal) consumption ratio $p(y - x)/c$ requires the value of cereal production for each household, yet there is no farm production data in the 61st round. Using 59th round data, which has both value of production and consumption expenditure data, we estimate a tobit regression of the form $py/c = \beta'X + u$, where X includes 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 the corresponding X variables, we then impute $\widehat{py/c} = \widehat{\beta}'X$ for each household in the 61st round expenditure survey. From this imputation we subtract the *actual* value of px/c , the share of cereals in total consumption in NSS61, to obtain the household's net consumption ratio.
- ii) κ : The proportion of a household's labor allocated to its own farm can be calculated from data available in the 61st round employment-unemployment survey and predicted based on observable characteristics X using a procedure analogous to that described in (i). Imputation of a value of $\widehat{\kappa}$ to each household in the NSS61 consumption expenditure survey follows accordingly. Households owning no land (aside from their homestead) receive $\kappa = 0$ by default.
- iii) e : Whereas in our theoretical development we assumed two sectors, in practice we allow for three (rural) wage sectors: (1) agriculture (A), in which wages are affected by climate; (2) non-agricultural (NA), in which wages are also affected by climate but to a lesser degree; and (3) public sector (P), as described in fn. 26, in which wages are assumed to be unaffected by climate. Based on a sample of households with at least one wage-earner in the 61st round employment-unemployment survey and using the same set of X variables as in (i) and (ii), we predict the proportion of household wage earners whose main job is in each of these sectors and use the results to impute $\widehat{e}_A, \widehat{e}_{NA}$, and \widehat{e}_P to households in NSS61.
- iv) π and w : Weights λ , φ , and σ are functions of sector-specific land and labor prices. To obtain π_I and π_N in terms of an annualized flow, we take district median of per hectare land values for, respectively, irrigated and unirrigated land (from NSS59) multiplied by a discount rate, $\delta = 0.05$; we also try $\delta = 0.10$ as a robustness check.²⁸ For w_A, w_{NA} ,

²⁸Deschênes and Greenstone (2007) arbitrarily choose the discount rate of 0.05 to compare estimates of climate impacts on land values and on agricultural profits. Jacoby (2000), using plot-level data from Nepal, finds a median rental income to land value ratio of 0.055 for a sample of rented plots and takes this as the discount rate for a calculation similar to ours. For India, we have plot-level data from a survey of rural households in the southern state of Andhra Pradesh collected by the World Bank and the Centre for Economic and Social Studies in 2004. The survey asks landowners for the market value of each plot along

w_P we compute district medians of daily wages in each of the respective sectors. In calculating labor income, we assume 300 working days per year per active household member.

Table 6 summarizes the data for the distributional analysis, focusing on the three-quarters of the Indian population residing in rural areas in 2004-05 (urban households will be dealt with in a later subsection under a different set of assumptions). As seen in equation (2), the crucial weight variables, determining the relative importance of the different price effects, are λ and $p(y - x)/c$. When calculated using our preferred discount rate of 0.05, λ , the proportion of household income derived from owned land, averages 0.22; using the higher discount rate of 0.10 gives a mean λ of 0.29. As will become apparent, however, the choice of δ has much less of an impact on the distributional results than these averages might suggest.

Table 7 presents, first, descriptive regressions for λ calculated using the alternative discount rates. These regressions include district fixed effects, which alone account for 17 percent of the variation in λ . Adding key household characteristics—land, education, and demographics—boosts the explanatory power to over 50 percent. Recall from equation (5) that λ is a nonlinear function of household characteristics and district-level prices. Therefore, it is not surprising that much of its variation remains unexplained by this linear regression. Households with greater landownership and relatively more irrigation have much larger values of λ , just as we would expect. More educated households, by contrast, have lower values of λ due to their relatively higher imputed labor income. Households classifying themselves as members of scheduled castes, tribes, or other "backward" groups also have lower λ values, conditional on land and education. This result appears to arise out of the nonlinearity of λ : backward households tend to have lower landholdings and are more heavily concentrated in districts with lower land values. Comparing coefficient estimates across the two discount rates, we see that λ is more sensitive to variation in landholdings at $\delta = 0.10$, which it obviously should be since a higher discount rate accentuates the land component of annualized income. Otherwise, though, the two sets of coefficient estimates look quite similar.

The last column of Table 7 reports a descriptive regression for the net consumption ratio. Evidently, households with larger and better irrigated landholdings have higher production relative to consumption of cereals and, hence, a higher value of $p(y - x)/c$. More educated households also have higher net consumption ratios, but for a different reason. These households tend to be wealthier and, for this reason, devote a smaller share of their total

with its rental value, regardless of whether it was actually rented out or not. Based on a sample of nearly 5000 plots, we obtain a median rental value to market value ratio of 0.045. This result suggests that the implicit discount rate that Indian farmers use to convert their capitalized land values into annualized income flows is indeed right around 0.05.

consumption expenditures to cereals. Conversely, backward households and those with a higher proportion of dependents have larger cereal expenditure shares and, therefore, smaller values of $p(y - x)/c$.

Combining all the weights, the estimated climate sensitivity parameters, and the district-wise temperature change scenarios from the climate models, ΔT , yields household-specific predictions for the percentage change in consumption due to warming temperatures over the next three decades, $\widehat{\% \Delta c} = \widehat{\frac{d \log(c)}{dT}} \Delta T$. We work with three alternative parameter configurations: (1) *Baseline Ricardian*: Model II for land values (Table 2) and wages (Table 5); (2) *Extended Ricardian*: Model III for land values (with $T \times P$ interaction) and Model II for wages; and (3) *Zero Adaptation*: Table 4, column (6) for agricultural productivity, subsection 4.5 for agricultural wages, and 0% for non-agricultural wages. The third scenario, called zero adaptation for lack of a better term, assumes that farmers will adjust to a permanent temperature increase just as they would to a particularly warm year. For consistency, this scenario also assumes no cross-sectoral adjustments in the rural labor market; i.e., only agricultural wages decline in response to the fall in demand for cultivation labor. A decline in non-agricultural wages as well would presuppose a shift in the supply of non-agricultural workers; this can only happen, at least in our model, by increasing education, which is to say through long-run adaptation.

5.3 Household and district-level impacts of warming

Each of the components of the rural distributional analysis is displayed in Figure 6, using the Baseline Ricardian parameter scenario with $\delta = 0.05$. The upper left histogram shows the distribution of predicted percentage changes in per-capita expenditures due only to changes in the value of land; wage and cereal price effects are zeroed out for illustrative purposes. The long left tail reflects the corresponding feature of the land ownership distribution: only 35% of rural households own any land at all, apart from their homestead, and the top ten percent of landowning households possess about half of the total. Moving right, we next consider household-level climate change impacts due to wage effects alone. This distribution is also left-skewed but exhibits a comparatively narrow range of variability across households. At the lower left of Figure 6 is the distribution of predicted percentage changes in consumption due to the rise in cereal prices alone (equivalent to the rescaled distribution of the net consumption ratio $p(y - x)/c$). The median cereal price effect is close to zero, but there is an extensive right tail populated by larger farm households. If cereal prices do indeed rise by ten percent in the next thirty years, then nearly 8 percent of rural households will realize consumption gains, solely from this source, of at least 5 percent. Adding up these three

price effects of global warming, at the bottom right of Figure 6, we obtain a remarkably symmetrical distribution of $\widehat{\% \Delta c}$ across households.

The top row of Table 8 presents descriptive statistics for $\widehat{\% \Delta c}$ calculated using the three different parameter configurations. Definition (1) zeros out the effects of higher cereal prices whereas definition (2) includes these effects. The population weighted average consumption decline due to warming lies between 6-11 percent, the larger figure presuming zero adaptation. Clearly, these consumption shocks, in percentage terms, are far more modest than the associated declines in land productivity. This is because the average of λ (proportion of income from land) and $\epsilon_{w,\pi}$ (elasticity of wages with respect to land productivity) are both so much smaller than unity. Comparing across definitions (1) and (2) reveals that, on average, the cereal price effect ameliorates the negative consequences of higher temperatures only very slightly.

These averages, however, conceal considerable heterogeneity. The regression coefficients in Table 8 show that, without cereal price effects, the land coefficients are negative, which means that the more land a household owns and the larger fraction of it that is irrigated, the greater the consumption losses due to climate change. But these coefficients are greatly attenuated when cereal price effects are incorporated into $\widehat{\% \Delta c}$. The reverse pattern is observed for the education coefficients. As we have seen from Table 7, both larger farmers and more educated households tend to have higher net consumption ratios and, hence, benefit more (or suffer less damage) from higher cereal prices. So, while rising food prices will have little aggregate impact on rural households, their distributional implications, as shown below, will be significant.

Before moving to this distributional analysis, we take a different cut at the data and consider the geographic pattern of impacts. Recall that in our descriptive regressions in Tables 7 and 8 we purged district fixed effects; now we turn the focus on the district means themselves. Figure 7 shows how the average λ and $p(y-x)/c$ vary across rural districts, and how they covary with each other; perhaps unsurprisingly, this spatial covariance is positive. Households in the north and west are more reliant on land for their income and they also produce greater marketed surplus, whereas toward the south and east the contrary is true on both counts. The next two figures map the district means of $\widehat{\% \Delta c}$, using definitions (1) and (2), under the two Ricardian parameter scenarios. Figure 8, the Baseline case, shows much of northern India hard-hit by the fall in agricultural productivity. But, after taking account of cereal price effects, matters improve markedly in the west and deteriorate in the east. So, the impacts of rising food prices, although washing out in the aggregate across rural areas, are not at all geographically neutral. Finally, the Extended Ricardian scenario in Figure 9, with its temperature-rainfall interaction, darkens the picture further for the arid northwest

and southern interior of India. We predict *average* consumption declines as high as 13% for some of these districts, even after incorporating the moderating influence of cereal price effects.

5.4 Climate change incidence curves

We now turn our attention to the main objective of this study, understanding the distributional implications of climate change. We do so using a climate change incidence curve (CCIC), the smoothed relationship between per capita household expenditure percentiles and predicted changes in per capita expenditures induced by increases in temperature over the next three decades ($\widehat{\% \Delta c}$). Economic growth over this time-frame is presumed to be neutral with respect to wealth ranking; a rising tide that lifts all boats, proportionally speaking. Thus, in the absence of warming, our assumption is that a household's—or, more precisely, a household *type*'s—per-capita expenditure percentile would remain unchanged.

We draw CCICs, in Figure 10, that sequentially incorporate each of the three price effects, beginning with the land price effect alone, then adding the labor price or wage effect, and finally the cereal price effect. Our initial focus is, once again, on rural areas. Panels (a)-(c) show CCICs using different versions of the Ricardian parameter estimates, though the curves look practically identical across scenarios. Doubling the discount rate used to transform land values into annualized income has no detectable impact on the incidence of climate change losses, only exacerbating the predicted consumption declines a bit. While the Extended Ricardian scenario, as we have just noted, alters the geographical distribution of impacts relative to the Baseline, it makes no discernible difference for the distribution of impacts across per-capita expenditure percentiles.

Ignoring cereal price effects, climate change appears "progressive" insofar as wealthier households suffer proportionally greater consumption losses. A household in the top percentile of the per capita expenditure distribution would experience nearly a two percentage point greater decline in consumption than a household in the bottom percentile. This progressivity is driven by the skewed land distribution and the fact that larger landowners are concentrated in the higher percentiles. By contrast, temperature-induced wage declines are relatively more costly to the poor than to the rich, mainly because the poor tend to engage in climate sensitive agricultural employment.

The picture changes dramatically after we incorporate the cereal price effect; the CCICs in panels (a)-(c) now slope upward! Indeed, climate change impacts become practically as regressive with cereal price effects as they were progressive without them. Underlying this incidence reversal is the fact that the net consumption ratio is increasing in wealth. Moving

up the per capita expenditures distribution, we find households gaining as much from the higher prices of their marketed surplus as they are losing from the lower productivity of their land.

The zero adaptation CCICs in panel (d) of Figure 10 differ somewhat from the Ricardian scenarios. First, the negative land price effect is larger in magnitude and increases more rapidly with the per capita expenditure percentile. This is a consequence of the relatively high sensitivity of agricultural productivity to temperature in the short run. Second, the wage effect has a greater distributional gradient, working even more strongly against the poor. Recall our assumption that, without adaptation, only cultivation wages are affected by rising temperatures along with the fact that India's rural poor are heavily represented in agricultural wage employment. Third, the cereal price effect, which we assume to be the same with or without adaptation, is no longer powerful enough to lead to an incidence reversal. Thus, the overall climate change impacts in the zero adaptation scenario are slightly progressive.

Comparing the pure agricultural productivity effects (land and labor) across panels (c) and (d), we see that the poorest households would be about 50 percent worse off in the absence of adaptation whereas the wealthiest households would fare almost twice as badly. In other words, adaptation appears to "favor" the non-poor. It is not, however, that the poor are ill-equipped to adapt, but only that the non-poor are receiving a bigger share of the returns to crop production, a domain in which there is a large scope for adaptation to climate change.²⁹ Of course, we ignore any costs associated with adaptation, which will presumably also be borne disproportionately by the non-poor.

Thus far we have considered the rural population in isolation. Next we bring urban households into the picture to examine the broader distributional implications of climate change in India. Around 15 percent of households classified as urban in NSS61 (representing 16 percent of the urban population) report that they own land apart from their homesteads. We treat these (presumably) peri-urban households as though they were rural, performing the same calculations as those summarized in Table 6 to obtain their $\widehat{\% \Delta c}$. For *bona fide* urban households, who have $\lambda = 0$, we assume zero wage effects. In other words, as a first approximation, we assume that urban and rural labor markets are segmented from each other.³⁰ As a consequence, these urban households will only experience the impact of higher cereal prices and, since these households are not agricultural producers, the welfare effects

²⁹Again, it comes down to $\epsilon_{w,\pi}$ being substantially less than unity. An adaptation that mitigates the decline in π due to warming by 10 percentage points, only mitigates the decline in w by around 3.5 percentage points.

³⁰This also assumes no change in the amenity value of climate for urban residents (see fn. 9), or at least that any change in the amenity value of climate is distributed equally in the population.

of higher cereal prices will be uniformly negative.

To compare per capita expenditures levels across rural and urban households, we adjust for cost-of-living difference using a new set of poverty lines produced for NSS61 by the World Bank. Panel (a) of Figure 11 highlights the rural-urban wealth disparity, showing an increasing urban representation in higher per capita expenditure percentiles. This fact largely explains the shape of the national CCIC in panel (b), which is merely a weighted average of the corresponding urban and rural curves displayed alongside. Given that rural households are made worse off relative to urban households more or less uniformly along the whole range of the wealth distribution and that the urban weight increases with wealth, the national CCIC must become increasing steep as well.

5.5 Digression: Economic growth and Engel’s Law

First-order comparative statics implies that the welfare cost of a price increase is proportional to (minus) the initial net consumption bundle, an approximation which is probably quite accurate for a change in cereal prices of just ten percent.³¹ Climate change, however, will not increase cereal prices overnight, so the question arises: What is the right "initial" net consumption bundle? So far, we have used the value of $p(y - x)/c$ observed in 2004-05. But, as a result of (distribution-neutral) economic growth,³² household incomes will also rise over the next three decades and, by Engel’s law, drive the food share, px/c , downward. Consequently, the farther off in the future the cereal price increases actually take place, the lower the (instantaneous) welfare cost is likely to be.

Suppose, then, we take the other extreme position and assume that cereal prices rise overnight, not in 2005, but only in 2040. In other words, we replace the 2005 cereal share with the predicted cereal share in 2040. Assuming that Engel’s law applies, and that the other component of the net consumption ratio, py/c , stays constant, this calculation should give a lower bound on the true welfare cost of the cereal price increase. Consider, for simplicity, the quadratic Engel curve specification

$$px/c = \beta_1 \log(c) + \beta_2 \log(c)^2 + \varphi'Z + \epsilon \tag{8}$$

used by Banks, et al. (1997). If consumption grows at an exponential rate of g for H years,

³¹Banks et al. (1996) find that even for a price increase approaching 20 percent errors in first-order welfare change estimates are tolerably small (less than 10 percent of the true welfare change based on full demand system estimation).

³²In terms of our model, where $c = z + px = \pi a + w\tau(1 - \kappa)$, assume that land and labor productivity, π and w , are rising over time at a common exponential growth-rate g , so that total per capita expenditures, c , is forced to follow suit.

then equation (8) implies that the cereal share must change by $\beta_1 gH + \beta_2 gH(gH + 2 \log(c))$ by year H . With $\beta_2 = 0$, the quadratic Engel curve collapses to the Working-Leser model, and the change in the cereal share is a constant, $\beta_1 gH$, across households. In this case, economic growth does not affect the distribution of welfare costs of higher cereal prices; every household experiences the same decline in its cereal share and hence the same fall in welfare costs.

Estimating equation (8) on the NSS61 sample of around 117 thousand households, we obtain $\widehat{\beta}_1 = -0.351$ (0.012) and $\widehat{\beta}_2 = 0.020$ (0.0009),³³ which allows us to reject Working-Leser in these data. Figure 12 displays the quadratic Engel curve along with a semi-parametric counterpart estimated by partially linear regression. The parametric model does a good job of capturing the flattening of the Engel curve at the upper tail of the per-capita expenditure distribution but overpredicts the extent to which cereal shares fall with income growth for the poorest households. Even so, adjusting the net consumption ratio to reflect the cereal shares that will prevail in 2040 has no appreciable distributional implications. Figure 13 shows the incidence of price effects by percentile in rural and urban areas assuming that $H = 35$, $g_{urban} = 1.49$, and $g_{rural} = 1.17$ (see next subsection). In both sectors, the revision of px/c leads to a small upward shift in the incidence curve. The "true" distribution of welfare costs of the 10 percent rise in cereal prices is likely to lie somewhere in between these pairs of incidence curves. Repeating this exercise using the better-fitting cubic Engel curve (see Figure 12) leads to practically identical conclusions.

5.6 Changes in poverty

Since consumption changes do not map linearly into poverty changes, predicting the increase in the poverty rate attributable to climate change involves an additional step. On top of the assumption of distributionally neutral growth underlying the CCICs, we need to take an explicit stand on the growth-rate that will prevail between our initial consumption measurement (2005) and the endpoint of our horizon (2040). In particular, the predicted poverty rate in the population given climate change, as calculated from our sample, is

$$P^{CC} = \frac{1}{pop} \sum_i I(c_i^0 \times \exp(gH) \times (1 + \widehat{\% \Delta c_i}) < \textit{poverty line}) \times \omega_i.$$

where g and H are as in the previous subsection, c_i^0 is household i 's initial per-capita expen-

³³The regression includes household demographics (size, proportion females, proportion old and young, head age and age squared), scheduled caste/tribe dummies, an urban dummy, and district fixed effects. Standard errors are adjusted for clustering on district. Separate regressions for urban and rural areas yield very similar estimates of β_1 and β_2 .

ditures, and ω_i is a household-level population weight such that $\sum \omega_i = pop$. We experiment with two alternative urban and rural growth scenarios, high and low, based on India’s actual experience during the periods 1958-1991 and 1991-2006, respectively.

Table 9 summarizes the poverty results at the national level and by sector. By 2040, economic growth alone will reduce national poverty from 44.5 to 15.7 percent in the high-growth scenario and to 27.9 percent in the low-growth scenario. With climate change, however, national poverty will be about 3-4 percentage points higher, with most of the ”new” poor coming from rural areas. Given an estimated 2040 population of around 1.5 billion, this means that warming temperatures will leave around 50 million Indians in poverty who would otherwise have risen out of it.

Finally, we can ask how much of the warming-induced poverty increase is due to a change in mean consumption versus a redistribution of consumption. This is done by recalculating P^{CC} , replacing $\widehat{\% \Delta c_i}$ with its mean value across households. As shown in Table 9, the bulk of the poverty increase is due to the fall in the mean, not the redistribution, of consumption. Consequently, price effects, which, as we have seen, contribute very marginally to the mean consumption decline, also play only a trivial role in the poverty rise.

6 Conclusions

We have estimated the welfare costs to Indian households of what might be considered moderate warming over the next three decades, and investigated the likely distribution of these costs in the population. Our main conclusions are as follows:

- The substantial fall in agricultural productivity (17 percent overall) that we predict as a result of warming will translate into a much more modest consumption decline for the majority of households. This is because these households derive the bulk of their income from wage employment and (rural) wages are estimated to fall by only a third as much as agricultural productivity.
- Climate change will have heterogeneous geographical impacts. Most strikingly, districts in the northwest of India, a region which is now comparatively well off, will suffer some of the largest percentage declines in consumption, whereas parts of eastern India, such as the very poor state of Orissa, will not fare too badly in relative terms.
- Once we account for the welfare effects of rising cereal prices, climate change impacts are regressive, falling more heavily on the poor than the rich. This is true in urban areas, where we assume that cereal price effects are the only welfare consequence of

climate change, as well as in rural areas, where the beneficial impact of higher prices to agricultural producers offsets the decline in land productivity.

- Combining rural and urban areas to consider a population of nearly 1 billion, we find that climate change is even more strongly regressive than it is in each of these sectors taken in isolation. This is simply due to the fact that the biggest losers reside in rural areas and the fraction of rural residents declines as one moves along the income distribution.
- Taking into account average income growth up to 2040, the national poverty rate will rise by 3-4 percentage points compared to the counterfactual of zero warming, which, given current population projections, will result in around 50 million more poor people than there otherwise would have been in that year.
- The benefits of adaptation will accrue disproportionately to the rich insofar as they would experience bigger consumption losses than the poor if adaptation opportunities were foreclosed—i.e., if farmers had to respond to the permanent shock as though it were a temporary one. This, again, is a consequence of the rich owning the lion's share of agricultural land.

This last finding has an important policy implication. If the objective is protecting the poor from climate change, a focus on helping farmers adapt may be misplaced, although we cannot say for certain until we know the cost of adaptation. Nevertheless, a better policy in this respect might be to reduce the share of rural income derived from climate sensitive employment, which is to say, to increase the share of income derived from non-agricultural activities. Promoting the non-farm sector, or at least removing obstacles to its development, would be one step in this direction.

In concluding, it is worth reiterating that all of the results in this paper, based as they are on a comparative statics framework, presuppose a marginal change in climate. As we move out beyond the next three decades, and India continues to warm, our modeling assumptions—e.g., no large-scale permanent migration in response to climate change—become increasingly untenable and, as a consequence, extrapolation becomes ever more precarious. Marginal, however, does not mean negligible; we have considered a 17 percent productivity shock to a sector accounting for around 70 percent of rural employment. In doing so, we have highlighted how the interplay between falling agricultural productivity and potentially rising global food prices driven by climate change shape the income distribution. Future work will extend the methodology to account for changes in the returns to land and labor that are likely to result from these food price changes in general equilibrium.

References

- [1] Altonji, Joseph G., Todd E. Elder and Christopher R. Taber. 2005. "Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools," *Journal of Political Economy*, 113(1): 151-84.
- [2] Badiani, Reena. 2009. "Factories and farms: How does economic growth impact rural incomes and education," unpublished manuscript, Yale University.
- [3] Banks, J., R. Blundell, and A. Lewbel. 1996. "Tax reform and welfare measurement: do we need demand system estimation?" *Economic Journal*, 106:1227-41.
- [4] Banks, J., R. Blundell, and A. Lewbel. 1997. "Quadratic Engel curves and consumer demand," *Review of Economics and Statistics*, LXXIX(4):527-39.
- [5] Cline, W. 2007. *Global warming and agriculture: Impact estimates by country* (Washington: Center for Global Development and Peterson Institute for International Economics).
- [6] Conley, T. 1999. "GMM estimation with cross sectional dependence," *Journal of Econometrics*, 92: 1-45.
- [7] Datt, G. and M. Ravallion. 1998. "Farm productivity and rural poverty in India," *Journal of Development Studies*, 34(4):62-85.
- [8] Datt, G. and M. Ravallion. 2011. "Has India's Economic Growth Become More Pro-Poor in the Wake of Economic Reforms," forthcoming *World Bank Economic Review*.
- [9] Deaton, A. 1989. "Rice prices and income distribution in Thailand: A non-parametric analysis," *Economic Journal*, 99:1-37.
- [10] Deschênes, O. and M. Greenstone. 2007. "The economic impacts of climate change: Evidence from agricultural profits and random fluctuations in the weather," *American Economic Review*, 97(1):354-385.
- [11] Eswaran, M. and A. Kotwal. 1986. "Access to capital and agrarian production organisation," *Economic Journal*, 96(382):482-498
- [12] Eswaran, M., A. Kotwal, B. Ramaswami, and W. Wadhwa. 2008. "How does poverty decline? Suggestive evidence from India, 1983-1999," unpublished manuscript, University of British Columbia.

- [13] Feder, G. 1985. "The relation between farm size and farm productivity : The role of family labor, supervision and credit constraints," *Journal of Development Economics*, 18(2-3): 297-313.
- [14] Fisher, A., M. Hanemann, M. Roberts, and W. Schlenker. 2009. "Climate change and agriculture reconsidered," unpublished manuscript, University of California at Berkeley.
- [15] Guiteras, R. 2009. "The impact of climate change on Indian agriculture," unpublished manuscript, University of Maryland.
- [16] Hertel, T., M. Burke, and D. Lobell. 2010. "The poverty implications of climate-induced crop yield changes by 2030," GTAP Working Paper No. 59, Purdue University.
- [17] Himanshu. 2005. "Wages in rural India: sources, trends and comparability," *Indian Journal of Labour Economics*, 48(2).
- [18] Jacoby, H. 2000. "Access to markets and the benefits of rural roads," *Economic Journal*, 110(465):713-737.
- [19] Kaufman, R. 1998. "The impact of climate change on US agriculture: a response to Mendelsohn et al. (1994)," *Ecological Economics*, 26:113-119.
- [20] Kelly, D., C. Kolstad, and G. Mitchell. 2005. "Adjustment costs from environmental change," *Journal of Environmental Economics and Management*, 50:468-95.
- [21] Kurukulasuriya and Mendelsohn. 2007. "Modeling endogenous irrigation: The impact of climate change on farmers in Africa," Policy Research Working Paper. The World Bank, Washington, DC.
- [22] Lanjouw, P. and R. Murgai. 2009. "Poverty decline, agricultural wages, and nonfarm employment in rural India: 1983-2004" *Agricultural Economics*, 40:243-263.
- [23] 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.
- [24] 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.
- [25] 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.

- [26] Quiggin, J. and J. Horowitz. 2003. "Costs of adjustment to climate change," *The Australian Journal of Agricultural and Resource Economics*, 47(4):429-46.
- [27] Roback, J. 1982. "Wages, rents, and the quality of life," *Journal of Political Economy*, 90(6):1257-1278.
- [28] Rosenzweig, M. and H. Binswanger. 1993. "Wealth, weather risk and the composition and profitability of agricultural investments," *Economic Journal*, 103:56-78.
- [29] Sanghi, A., R. Mendelsohn, and A. Dinar. 1998. "The climate sensitivity of Indian agriculture." Chapt. 4 in *Measuring the Impact of Climate Change on Indian Agriculture*, The World Bank, Washington, DC.
- [30] Schlenker, W., W. M. Hanemann, and A. Fisher. 2005. "Will U.S. agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach," *American Economic Review*, 95(1): 395-406.
- [31] Stern, N. 2006. *The economics of climate change: The Stern Review*. London: H.M. treasury.
- [32] 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.
- [33] Van der Mensbrugge, D. 2009. "The environmental impact and sustainability applied general equilibrium (ENVISAGE) model," The World Bank, Washington, DC.
- [34] World Bank. 2009. *South Asia: Climate change strategy*. South Asia region, The World Bank, Washington, DC.
- [35] World Bank. 2010. *India's employment challenge: Creating jobs, helping workers*. Oxford: Oxford University Press.

Table 1: DATA FOR RICARDIAN ANALYSIS

Variable	Mean	Std. dev.	Min.	Max.	Source
Mean annual temperature ($^{\circ}\text{C}$)	24.8	2.9	13.7	29.2	s1
Mean annual precipitation (mm)	1296	706	175	4350	s2
CV of annual precipitation	1.42	0.29	0.84	2.52	s2
Acrisols soils	0.087	0.282	0.000	1.000	s3
Cambisols soils	0.214	0.411	0.000	1.000	s3
Luvisols soils	0.280	0.450	0.000	1.000	s3
Nitisols soils	0.089	0.284	0.000	1.000	s3
Vertisols soils	0.157	0.364	0.000	1.000	s3
Hilly	0.101	0.302	0.000	1.000	s3
Undulating	0.417	0.494	0.000	1.000	s3
Fine textured soils	0.194	0.396	0.000	1.000	s3
Medium textured soils	0.677	0.468	0.000	1.000	s3
Elevation 250-1500m	0.406	0.412	0.000	1.000	s4
Elevation > 1500m	0.040	0.157	0.000	1.000	s4
Groundwater (m^3/km^2)	19.8	15.4	0.1	103.8	s5
River length (km/km^2)	0.088	0.052	0.000	0.349	s6
Km to nearest city of > 1 million	221.6	180.5	3.7	1011.6	s7
Proportion irrigated plots	0.521	0.319	0.000	1.000	s8
Plot value ('000 Rs. per hectare)	307.7	303.7	13.2	2250.7	s8
Wage (Rs. per day)	75.0	47.4	20.4	358.0	s9
Proportion non-ag. employment	0.537	0.263	0.000	1.000	s9

Notes: Summary statistics for s1-s8 taken over 542 districts, data for s9 over 565 districts. Data sources are s1: 1951-1980 (monthly) weather station data, India Meteorological Department; s2: 1960-2000 (daily) 1° gridded data, India Meteorological Department; s3: 1992 Soil Map of the World, FAO; s4: SRTM 90m Digital Elevation Data, CGIAR-CSI; s5: 2004 Dynamic Ground Water Resources of India, Ministry of Water Resources; s6: 1997 Digital Chart of the World, ESRI; s7: 2001 Census of India; s8: NSS59 Schedule 18.2; s9: NSS61 Schedule 10.

Table 2: LAND VALUES AND CLIMATE

1st stage: Var. × Irrigated plot 2nd stage: Variable	Model I		Model II		Model III	
	First	Second	First	Second	First	Second
Constant	0.367*** (0.017)	2.822*** (0.902)	0.886*** (0.230)	4.619*** (0.826)	0.279 (0.440)	6.511*** (1.052)
Mean annual temp. (T)		-0.096*** (0.023)	-0.009 (0.007)	-0.130*** (0.023)	0.017 (0.018)	-0.218*** (0.039)
Mean annual precip. (P)		-0.396 (0.247)	-0.206*** (0.042)	-0.129 (0.113)	0.297 (0.307)	-1.951*** (0.694)
CV of annual precip.		-0.000 (0.004)	-0.057 (0.067)	-0.003* (0.002)	-0.071 (0.068)	-0.002 (0.002)
$T \times P$					-0.021* (0.013)	0.077*** (0.028)
Hilly				-0.248 (0.190)		-0.333 (0.203)
Undulating				-0.350*** (0.063)		-0.316*** (0.056)
Elevation 250-1500m				-0.413*** (0.091)		-0.396*** (0.092)
Elevation > 1500m				-0.526 (0.389)		-0.329 (0.456)
Fine textured soils				0.021 (0.150)		0.047 (0.144)
Medium textured soils				0.182** (0.085)		0.229*** (0.079)
Acrisols				0.682*** (0.203)		0.668*** (0.207)
Cambisols				-0.168 (0.134)		-0.140 (0.133)
Luvisols				-0.150* (0.089)		-0.129 (0.085)
Nitisols				0.198* (0.115)		0.146 (0.120)
Vertisols				-0.089 (0.172)		-0.064 (0.167)
Km to nearest city of > 1 mil.				-0.003*** (0.000)		-0.003*** (0.000)
Adjusted R^2		0.124		0.508		0.522

Notes: Standard errors in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$), adjusted for district-level clustering in first stage and for spatial correlation in second stage using Conley (1999) with 8 degree radius. Dependent variable is the log of plot value per hectare. Independent variables are interacted with irrigated dummy plot in first stage district fixed-effects estimation (sample size = 101,021 plots) and enter directly in second stage estimation (sample size = 542 districts). Second stage regressions weighted by number of plots in district. Omitted categories: topography (level to gently undulating), elevation (< 250 m), soil texture (coarse), soil type (other).

Table 3: DATA FOR PANEL ANALYSIS

Variable	Mean (std. dev.) by Year						
	1999	2000	2001	2002	2003	2004	2005
TEMPERATURE °C							
Annual	25.3 (2.5)	25.5 (2.1)	25.5 (2.5)	25.6 (2.7)	25.6 (2.5)	25.5 (2.2)	25.9 (1.8)
Kharif (June-Oct.)	27.7 (2.6)	27.8 (2.3)	27.7 (2.5)	28.1 (2.7)	27.8 (2.5)	27.5 (2.2)	28.2 (1.8)
Rabi (Nov.-March)	20.9 (3.6)	21.4 (3.2)	21.5 (3.5)	21.1 (3.7)	21.4 (3.4)	21.8 (3.4)	22.1 (2.8)
PRECIPITATION MM							
Annual	1162 (690)	1098 (627)	1189 (699)	1014 (666)	1138 (662)	1200 (790)	1329 (722)
Kharif (June-Oct.)	943 (539)	905 (504)	980 (560)	810 (538)	983 (557)	934 (579)	1076 (616)
Rabi (Nov.-March)	79 (122)	56 (87)	85 (101)	99 (97)	63 (78)	106 (113)	130 (184)
PRODUCTION							
Gross revenue (Rs./ha)	1140 (688)	1071 (704)	1124 (621)	984 (674)	1167 (732)	1070 (783)	1097 (725)
Proportion area irrigated	0.55 (0.32)	0.56 (0.32)	0.55 (0.32)	0.55 (0.32)	0.55 (0.32)	0.49 (0.30)	0.53 (0.31)
Number of districts	489	481	483	486	494	404	257

Notes: Weather variables derived from (monthly) gridded data, Center for Climatic Research, University of Delaware; gross revenue from the Directorate of Economics and Statistics, Ministry of Agriculture; irrigation from NSS59.

Table 4: GROSS AGRICULTURAL PRODUCTIVITY AND WEATHER SHOCKS

	(1)	(2)	(3)	(4)	(5)	(6)
Annual temperature	-0.242*** (0.029)	-0.308*** (0.037)				
Kharif temperature (T_K)			-0.206*** (0.026)	-0.438* (0.238)	-0.208*** (0.045)	-0.280*** (0.059)
T_K^2				0.005 (0.004)		
$T_K \times$ prop. area irrig.					0.020 (0.059)	0.097 (0.085)
Rabi temperature (T_R)			-0.117*** (0.019)	0.033 (0.129)	0.011 (0.029)	0.007 (0.030)
T_R^2				-0.004 (0.003)		
$T_R \times$ prop. area irrig.					-0.178*** (0.046)	-0.173*** (0.047)
Annual precipitation	0.402*** (0.049)	0.285*** (0.048)				
Kharif precip. (P_K)			0.155*** (0.027)	0.346*** (0.061)	0.786*** (0.141)	-0.404 (0.569)
P_K^2				-0.039*** (0.013)	-0.087*** (0.030)	-0.076*** (0.029)
$P_K \times$ prop. area irrig.					-0.840*** (0.218)	0.197 (1.197)
$P_K^2 \times$ prop. area irrig.					0.089* (0.050)	0.081 (0.054)
$T_K \times P_K$						0.042** (0.021)
$T_K \times P_K \times$ prop. area irrig.						-0.038 (0.040)
Rabi precip. (P_R)			0.043 (0.078)			

Notes: Standard errors in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$), adjusted for district-level clustering. Dependent variable is the log gross revenue per hectare. Sample size is 3094 district-year observations, falling to 3025 when irrigation variables are included in cols. (5) and (6). All specifications include district fixed effects. Column (1) includes a year trend and columns (2)-(6) include year dummies.

Table 5: WAGES AND CLIMATE

1st stage: Var. \times non-ag. empl.	Model I		Model II		Model III	
2nd stage: Variable	First	Second	First	Second	First	Second
Constant	0.253*** (0.014)	1.006*** (0.321)	0.003 (0.160)	1.182*** (0.265)	-0.686** (0.294)	1.601*** (0.383)
Mean annual temp. (T)		-0.049*** (0.014)	0.012** (0.005)	-0.049*** (0.010)	0.040*** (0.011)	-0.065*** (0.015)
Mean annual precip. (P)		0.137* (0.080)	-0.055** (0.028)	0.069 (0.051)	0.513** (0.221)	-0.294 (0.294)
CV of annual precip.		0.001 (0.002)	-0.000 (0.001)	0.001 (0.001)	-0.000 (0.001)	0.001 (0.001)
$T \times P$					-0.023** (0.009)	0.015 (0.013)
Hilly				0.308*** (0.046)		0.319*** (0.043)
Undulating				0.098*** (0.035)		0.098*** (0.038)
Elevation 250-1500m				-0.253*** (0.044)		-0.253*** (0.043)
Elevation $>$ 1500m				0.073 (0.115)		0.070 (0.122)
Fine textured soils				-0.000 (0.056)		0.000 (0.055)
Medium textured soils				0.096*** (0.036)		0.096*** (0.035)
Acrisols				0.097 (0.079)		0.105 (0.082)
Cambisols				-0.205*** (0.055)		-0.206*** (0.054)
Luvisols				-0.147*** (0.053)		-0.149*** (0.052)
Nitisols				0.037 (0.082)		0.039 (0.081)
Vertisols				-0.091 (0.085)		-0.092 (0.086)
Km to nearest city of $>$ 1 mil.				-0.000* (0.000)		-0.000* (0.000)
River length (km/km^2)				0.591* (0.322)		0.583* (0.317)
Groundwater (m^3/km^2)				-0.000 (0.001)		-0.000 (0.001)
Adjusted R^2		0.277		0.568		0.564

Notes: Standard errors in parentheses (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$), adjusted for district-level clustering in first stage and for spatial correlation in second stage using Conley (1999) with 8 degree bandwidth. Dependent variable is the log wage. Independent variables are interacted with non-agricultural employment dummy in first stage district fixed-effects estimation (sample size = 36,201 individuals) and enter directly in second stage estimation (sample size = 556 districts). First-stage regressions also include a cubic polynomial in age, gender dummy, education/literacy dummies, and gender-education interactions. Second stage regressions weighted by number of sample individuals in district. Omitted categories: topography (level to gently undulating), elevation ($<$ 250 m), soil texture (coarse), soil type (other).

Table 6: DESCRIPTIVE STATISTICS FOR RURAL DISTRIBUTIONAL ANALYSIS

	Mean (std. dev.)	Label	Source
VARIABLES			
π_I	278.6 (240.58)	median value ('000 Rs. per ha.) for irrigated plots	NSS59 Sch18.2
π_N	204.9 (207.11)	median value ('000 Rs. per ha.) for unirrigated plots	NSS59 Sch18.2
a_I	0.078 (0.215)	per capita landholdings of irrigated land	NSS61 Sch1
a_N	0.083 (0.243)	per capita landholdings of unirrigated land	NSS61 Sch1
w_A	46.59 (23.59)	median daily wage for cultivation activities	NSS61 Sch10
w_{NA}	63.60 (25.11)	median daily wage for nonfarm activities	NSS61 Sch10
w_P	186.46 (82.18)	median daily wage for public sector activities	NSS61 Sch10
\hat{e}_A	0.35 (0.238)	imputed proportion of cultivation workers	NSS61 Sch10
\hat{e}_{NA}	0.42 (0.189)	imputed proportion of nonfarm workers	NSS61 Sch10
\hat{e}_P	0.22 (0.196)	imputed proportion of public sector workers	NSS61 Sch10
$\hat{\kappa}$	0.46 (0.374)	imputed proportion of on-farm hours worked	NSS61 Sch10
τ	0.56 (0.213)	number of active hh members (age 15-65) per hh member	NSS61 Sch1
\hat{v}_{land}	4.88 (14.74)	imputed income ('000 Rs.) from owned land ($\delta = .05$)	
\hat{v}_{labor}	21.93 (24.20)	imputed income ('000 Rs.) derived from wage labor	
WEIGHTS			
φ	0.56 (0.448)	proportion of land income from irrigated land	
σ_A	0.24 (0.202)	proportion of wage income from agricultural employment	
σ_{NA}	0.35 (0.185)	proportion of wage income from non-agricultural employment	
σ_P	0.41 (0.277)	proportion of wage income from public sector employment	
$\lambda_{\delta=.05}$	0.22 (0.268)	proportion of total income from owned land	
$\lambda_{\delta=.10}$	0.29 (0.314)	proportion of total income from owned land	
$\frac{p(y-x)}{c}$	0.034 (0.276)	net consumption ratio (partially imputed)	

Note: Sample size is 76,213 rural households from NSS61 Schedule 1, representing a rural population of 723,709,342 people.

Table 7: DESCRIPTIVE REGRESSIONS FOR RURAL HOUSEHOLDS

	$\lambda_{\delta=.05}$	$\lambda_{\delta=.10}$	$\frac{p(y-x)}{c}$
Land (0, 2] ha. per capita	0.281*** (0.007)	0.386*** (0.008)	0.252*** (0.006)
Land > 2 ha. per capita	0.738*** (0.011)	0.794*** (0.010)	0.547*** (0.023)
Proportion irrigated land	0.089*** (0.008)	0.096*** (0.009)	0.187*** (0.006)
Primary complete	-0.014** (0.006)	-0.012* (0.006)	0.041*** (0.006)
Secondary incomplete	-0.054*** (0.006)	-0.055*** (0.006)	0.107*** (0.006)
Secondary complete and above	-0.168*** (0.009)	-0.184*** (0.010)	0.180*** (0.010)
Backward caste/tribe	-0.058*** (0.005)	-0.061*** (0.005)	-0.047*** (0.004)
Dependency ratio	0.013** (0.005)	0.012** (0.006)	-0.029*** (0.005)
Proportion of males	-0.024*** (0.007)	-0.024*** (0.008)	0.005 (0.006)
Adj. R^2	0.505	0.578	0.621

Notes: Standard errors in parentheses (** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$), adjusted for district-level clustering. Number of observations is 76213 households. Regressions are population weighted and include district fixed effects. Omitted categories are: land (landless), education (less than complete primary).

Table 8: DESCRIPTIVE REGRESSIONS FOR RURAL HOUSEHOLDS: $\widehat{\% \Delta C}$

Definition	Baseline Ricardian		Extended Ricardian		Zero Adaptation	
	(1)	(2)	(1)	(2)	(1)	(2)
Mean $\widehat{\% \Delta C}$ (std. dev.)	-0.0637 (0.0407)	-0.0602 (0.0297)	-0.0678 (0.0487)	-0.0644 (0.0356)	-0.1127 (0.1040)	-0.1093 (0.0866)
REGRESSION COEFFS.						
Land (0, 2] ha.	-0.0322*** (0.0009)	-0.0070*** (0.0008)	-0.0371*** (0.0015)	-0.0119*** (0.0011)	-0.0856*** (0.0025)	-0.0605*** (0.0022)
Land > 2 ha.	-0.1007*** (0.0018)	-0.0461*** (0.0023)	-0.1239*** (0.0032)	-0.0693*** (0.0033)	-0.2487*** (0.0052)	-0.1942*** (0.0049)
Prop. irrigated land	-0.0187*** (0.0012)	0.0001 (0.0010)	-0.0205*** (0.0016)	-0.0018 (0.0014)	-0.0474*** (0.0034)	-0.0286*** (0.0032)
Primary completed	0.0122*** (0.0009)	0.0163*** (0.0008)	0.0128*** (0.0011)	0.0169*** (0.0009)	0.0276*** (0.0022)	0.0317*** (0.0019)
Secondary incomplete	0.0339*** (0.0011)	0.0446*** (0.0011)	0.0353*** (0.0012)	0.0461*** (0.0012)	0.0670*** (0.0024)	0.0777*** (0.0022)
Secondary complete	0.0722*** (0.0016)	0.0902*** (0.0015)	0.0761*** (0.0018)	0.0942*** (0.0016)	0.1220*** (0.0038)	0.1400*** (0.0036)
Backward caste/tribe	0.0079*** (0.0008)	0.0031*** (0.0006)	0.0099*** (0.0010)	0.0052*** (0.0007)	0.0181*** (0.0022)	0.0133*** (0.0019)
Dependency ratio	0.0057*** (0.0008)	0.0028*** (0.0006)	0.0060*** (0.0009)	0.0032*** (0.0008)	0.0113*** (0.0020)	0.0084*** (0.0018)
Proportion of males	-0.0056*** (0.0010)	-0.0051*** (0.0008)	-0.0050*** (0.0012)	-0.0045*** (0.0010)	0.0068*** (0.0026)	0.0073*** (0.0023)
Adj R^2	0.544	0.522	0.548	0.503	0.532	0.501

Notes: Standard errors in parentheses ($***p < 0.01$, $**p < 0.05$, $*p < 0.1$), adjusted for district-level clustering. Number of observations is 76213 households. Descriptive statistics and regressions are population weighted. All regressions include district fixed effects. Definition (1) of $\widehat{\% \Delta C}$ zeros out cereal price effect; definition (2) includes cereal price effect. Omitted categories are: land (landless), education (less than complete primary).

Table 9: PREDICTED POVERTY RATES

	2004/05	2040			
		High growth		Low growth	
		(1)	(2)	(1)	(2)
RURAL					
P	48.8	13.9	13.9	29.9	29.9
P^{CC}	48.8	18.0	18.3	35.9	35.8
$P^{CC} _{\widehat{mean\% \Delta C}}$	48.8	18.0	17.8	35.7	35.7
URBAN					
P	31.1	5.2	5.2	14.6	14.6
P^{CC}	31.1	5.3	5.8	14.7	15.7
$P^{CC} _{\widehat{mean\% \Delta C}}$	31.1	5.3	5.6	14.8	15.5
ALL					
P	44.5	11.8	11.8	26.2	26.2
P^{CC}	44.5	14.9	15.3	30.8	31.0
$P^{CC} _{\widehat{mean\% \Delta C}}$	44.5	14.9	14.8	30.7	30.6

Notes: Annual growth rates in mean consumption derived from several NSS rounds are drawn from Datt and Ravallion (2009). Low growth refers to 1958-1991 (0.58 rural and 0.79 urban), and high growth rates to 1991-2006 (1.17 rural and 1.49 urban). Definition (1) zeros out cereal price effect; definition (2) includes cereal price effect.

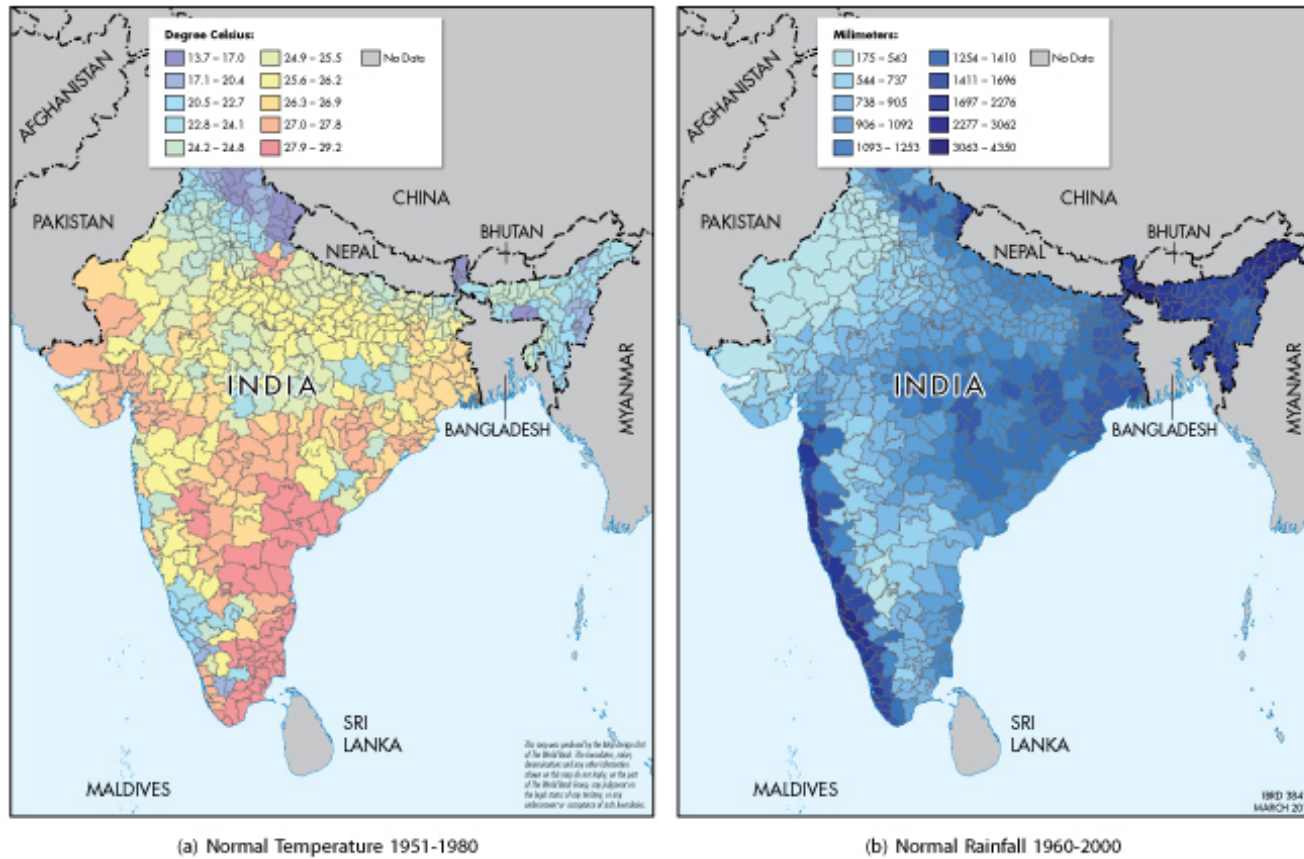


FIGURE 1 —CLIMATE OF INDIA

Note: Excluding Andaman & Nicobar Islands, Lakshadweep, and Jammu & Kashmir

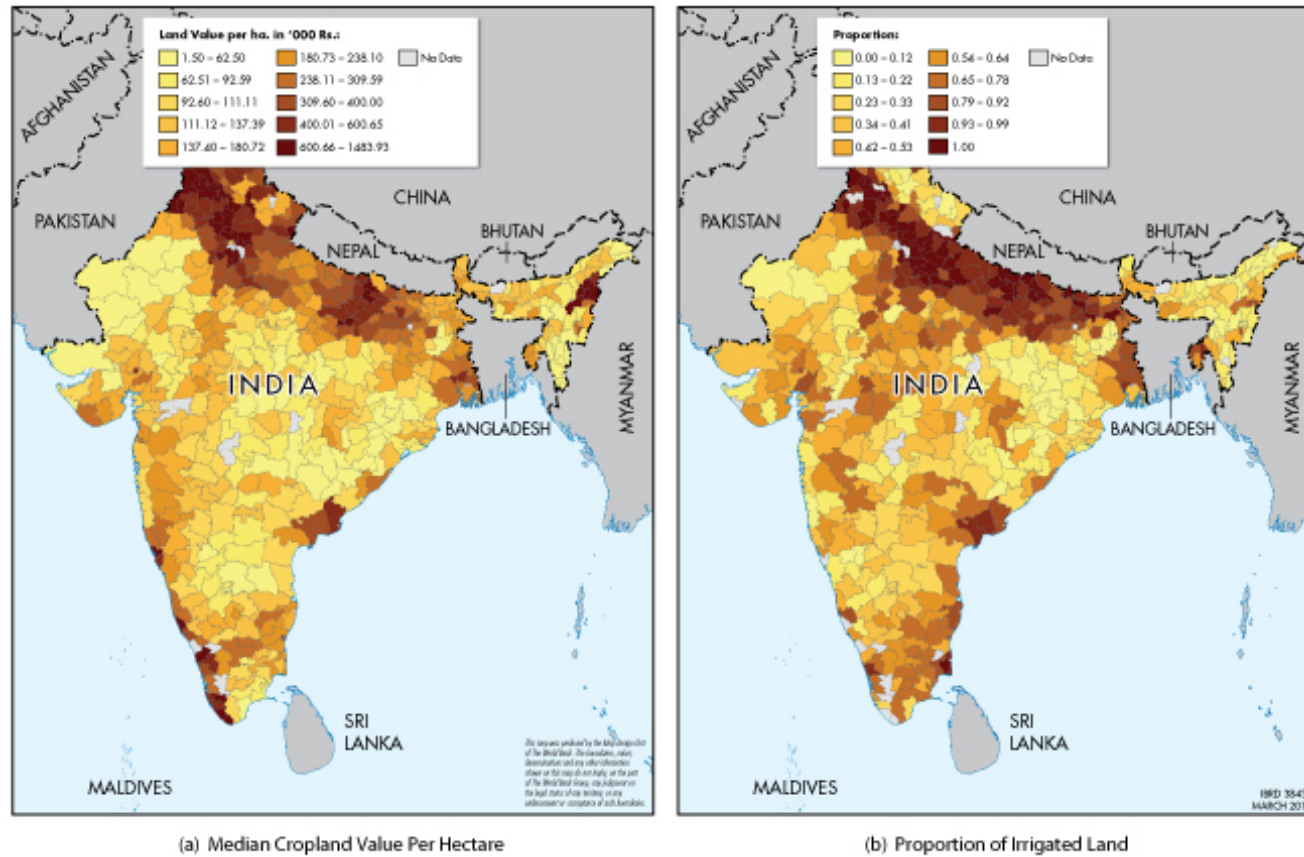


FIGURE 2 —LAND VALUES AND IRRIGATION

Note: Excluding Andaman & Nicobar Islands, Lakshadweep, and Jammu & Kashmir

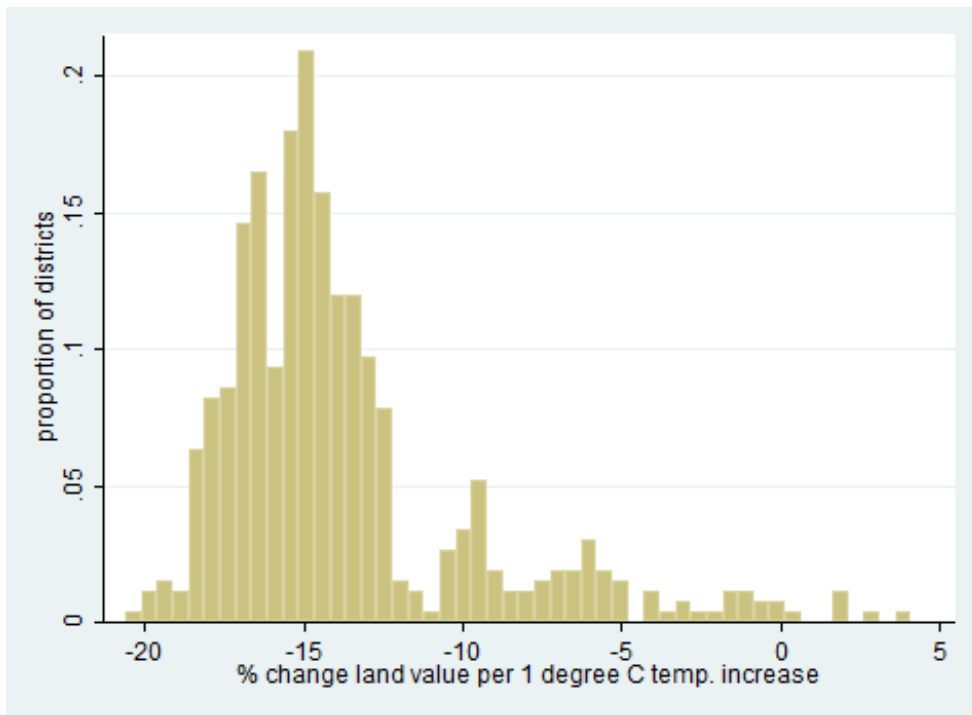


FIGURE 3 — DISTRIBUTION OF TEMPERATURE IMPACTS ON LAND VALUES

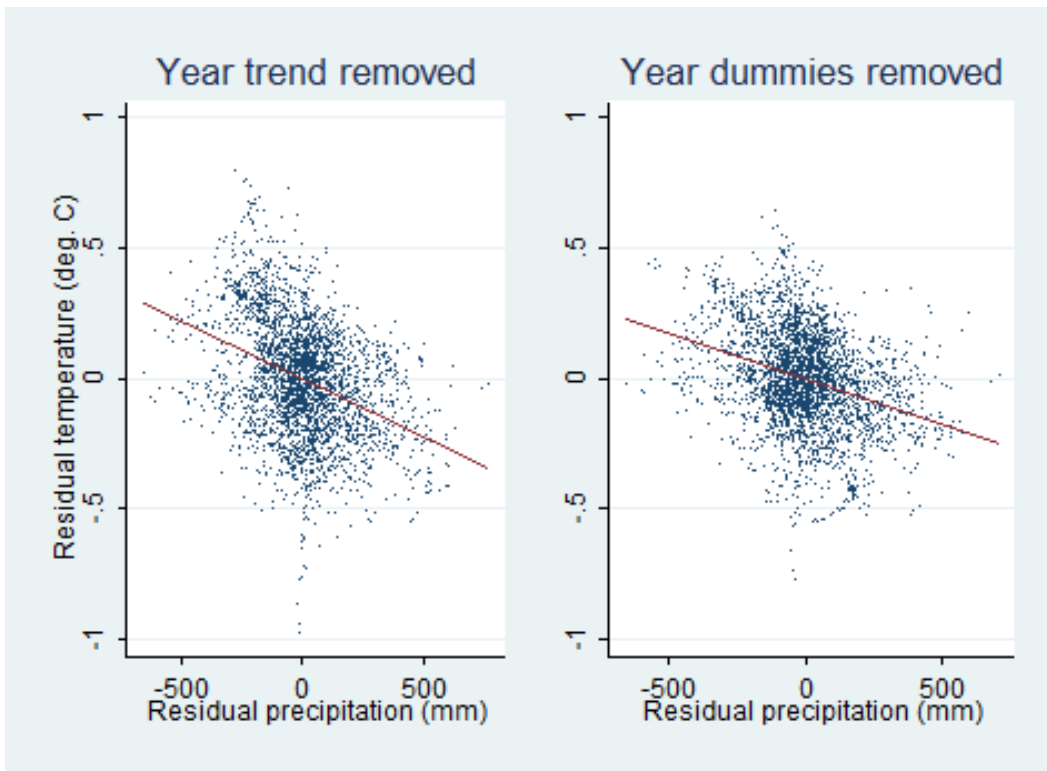


FIGURE 4 —RESIDUAL WEATHER VARIATION: INDIAN DISTRICTS 1999-2005

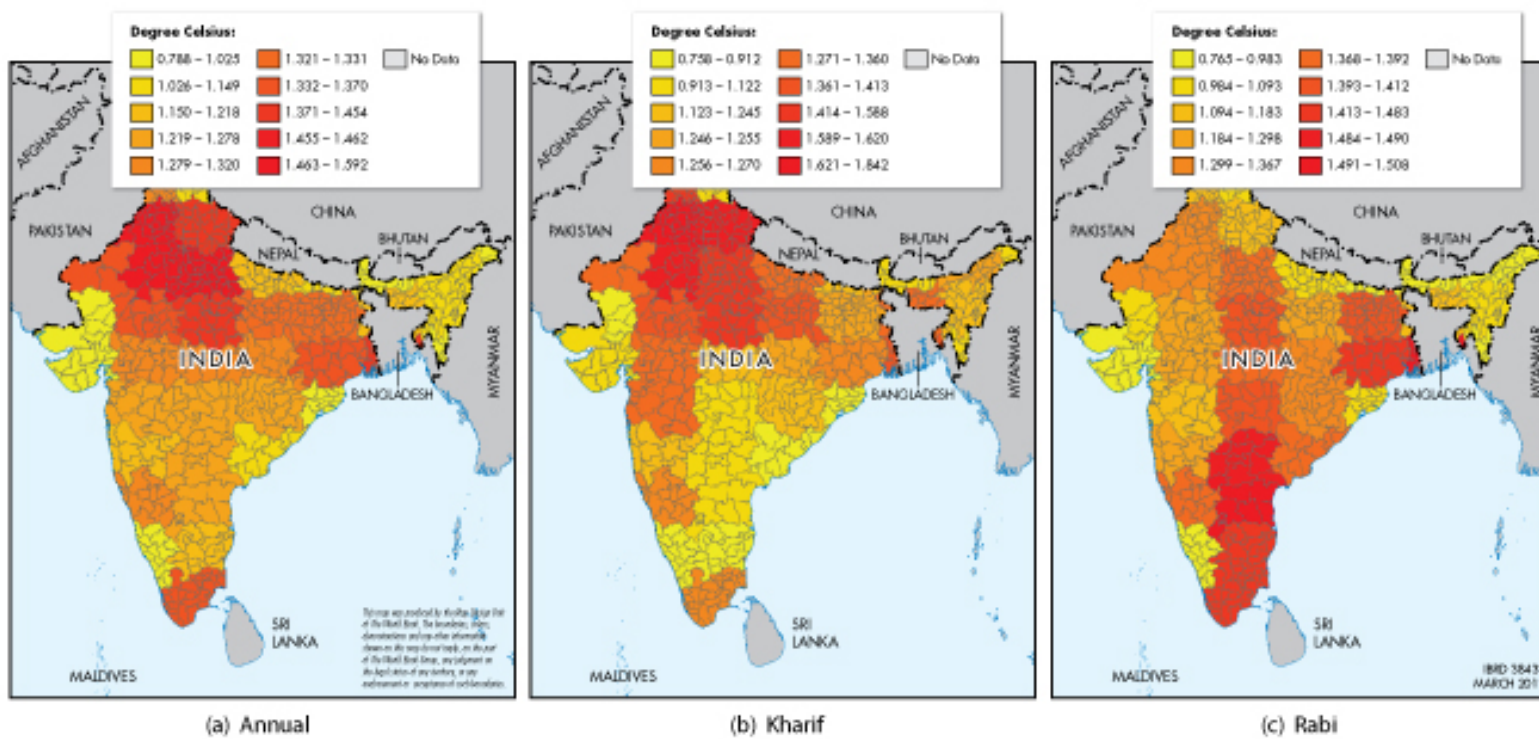


FIGURE 5 — PREDICTED CHANGES IN TEMPERATURE - HADCM3 A1F 2040

Note: Excluding Andaman & Nicobar Islands, Lakshadweep, and Jammu & Kashmir

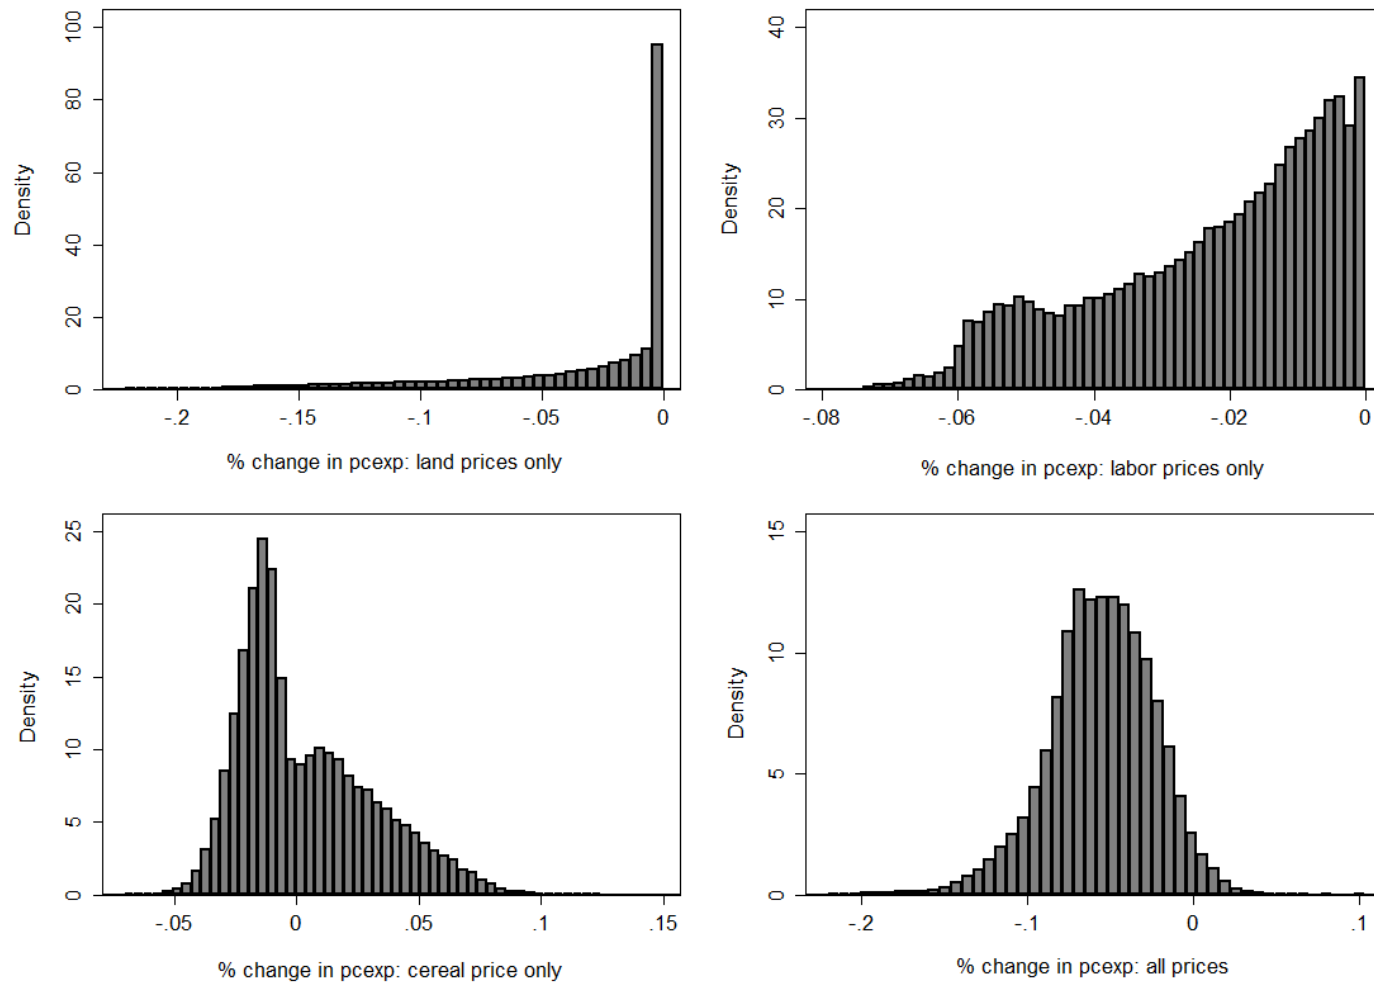


FIGURE 6 —DISTRIBUTION OF TEMPERATURE IMPACTS ON PER CAPITA CONSUMPTION

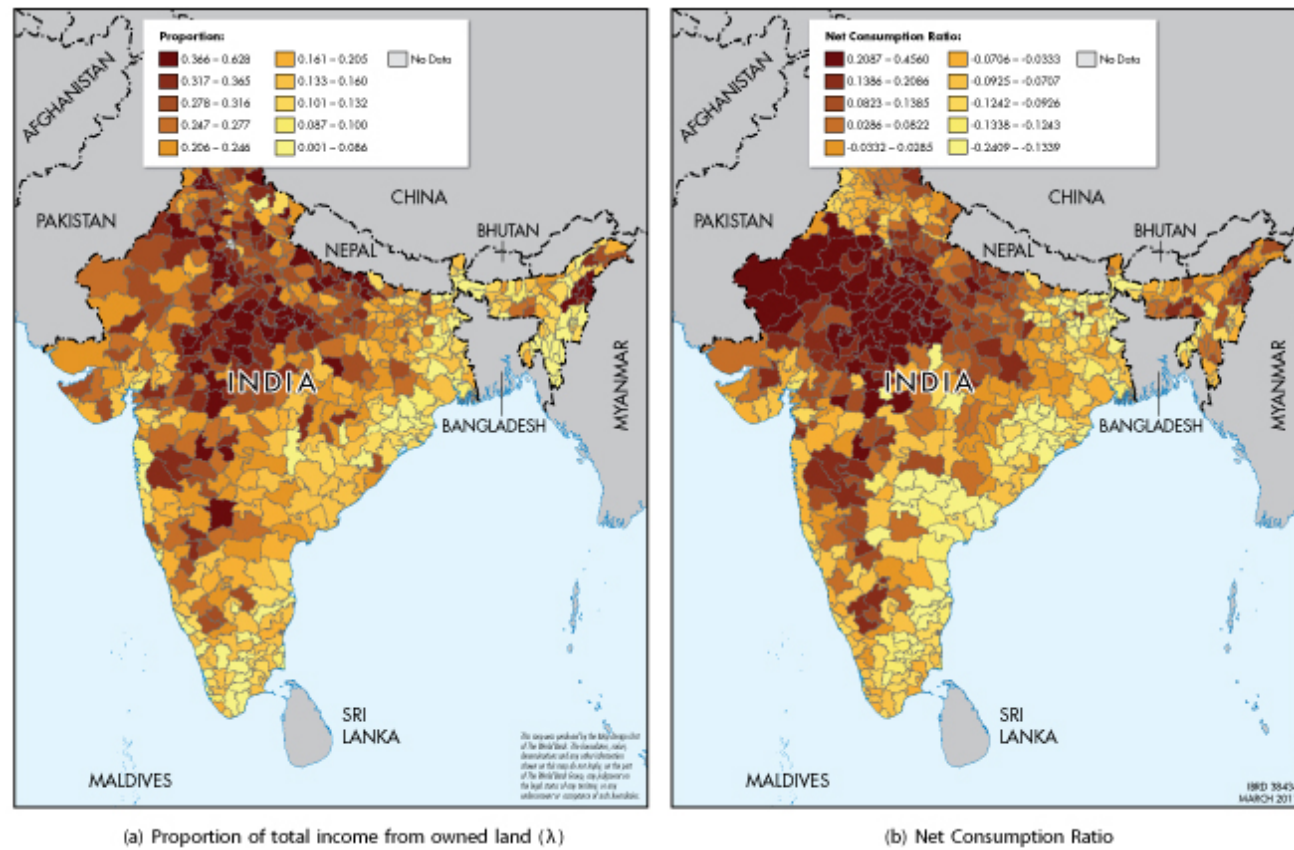


FIGURE 7 —WEIGHTS FOR DISTRIBUTIONAL ANALYSIS (RURAL HOUSEHOLDS)

Note: Excluding Andaman & Nicobar Islands, Lakshadweep, and Jammu & Kashmir

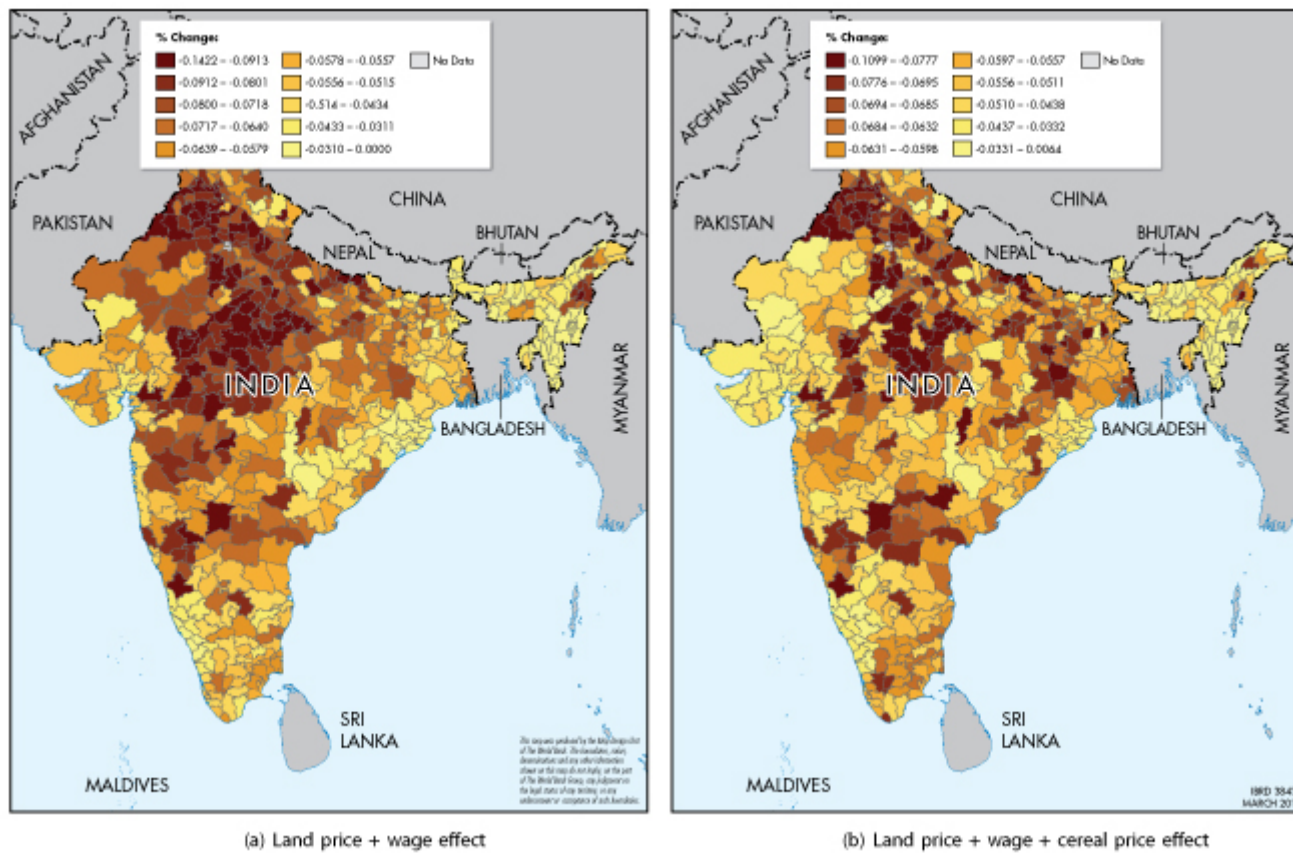


FIGURE 8 —PERCENTAGE CHANGE IN PER CAPITA CONSUMPTION: BASELINE RICARDIAN (RURAL HOUSEHOLDS)

Note: Excluding Andaman & Nicobar Islands, Lakshadweep, and Jammu & Kashmir

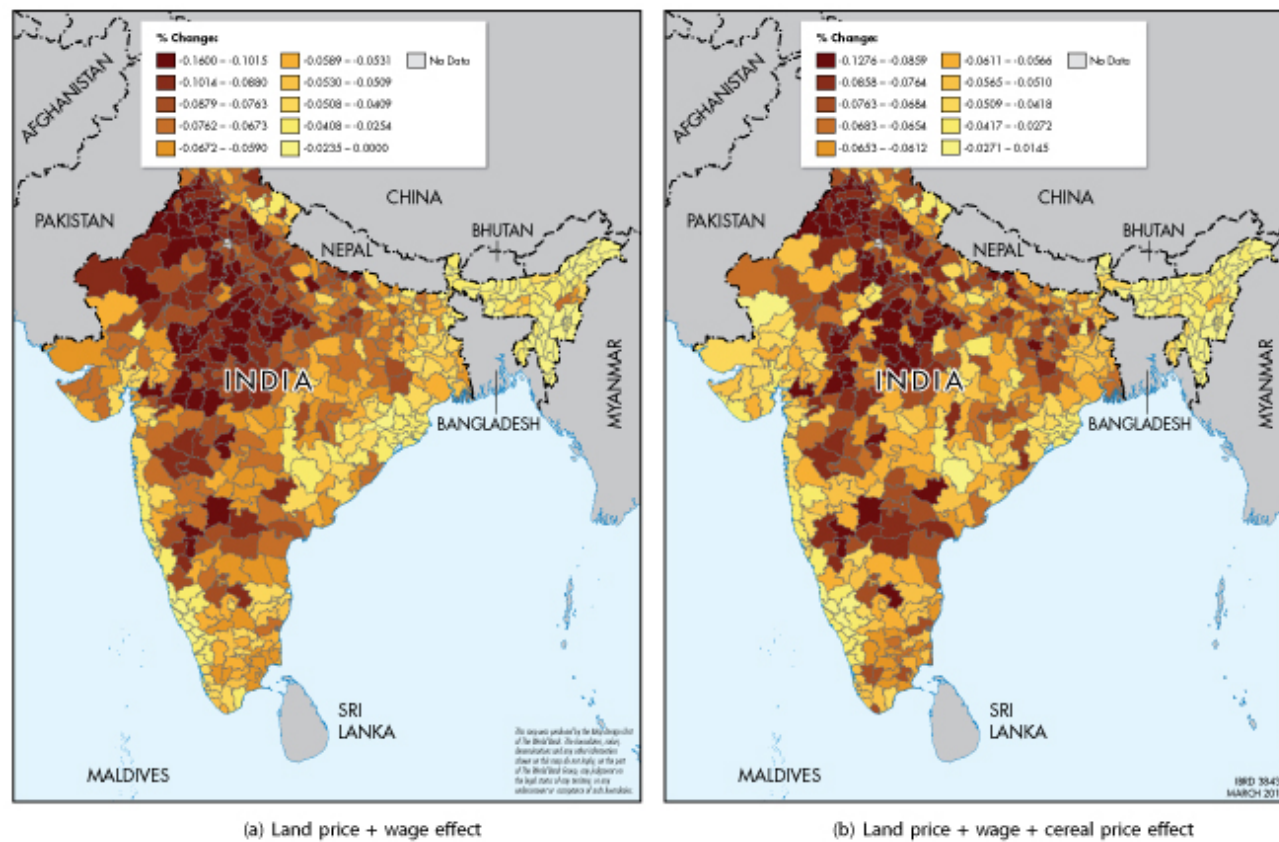


FIGURE 9 —PERCENTAGE CHANGE IN PER CAPITA CONSUMPTION: EXTENDED RICARDIAN (RURAL HOUSEHOLDS)

Note: Excluding Andaman & Nicobar Islands, Lakshadweep, and Jammu & Kashmir

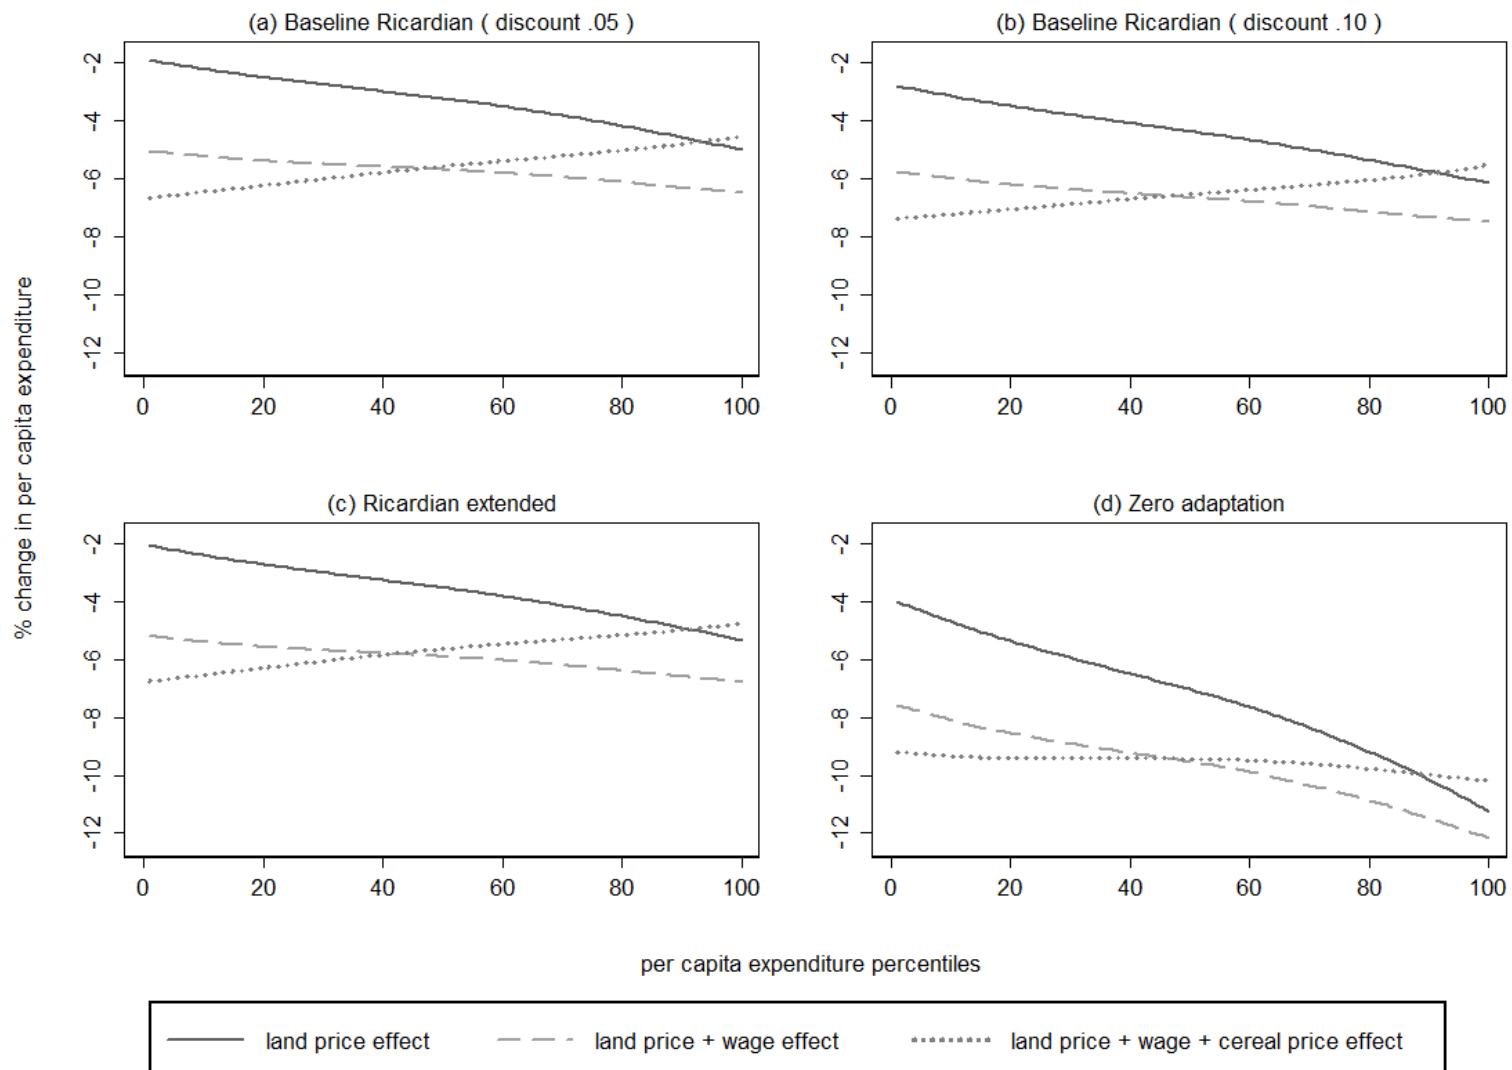
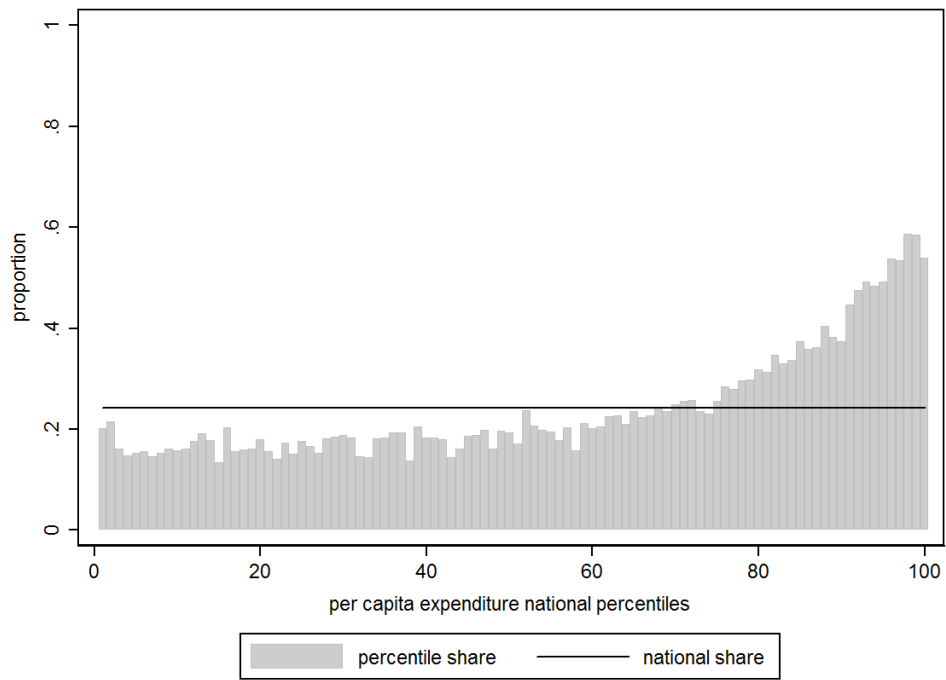
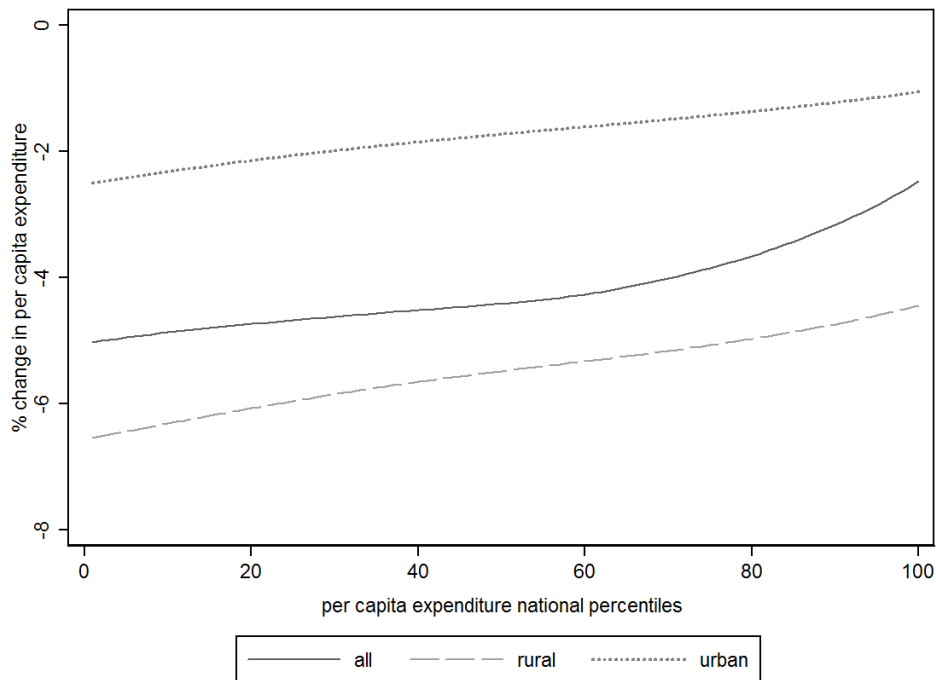


FIGURE 10 —CLIMATE CHANGE INCIDENCE CURVES FOR RURAL POPULATION



(a) Proportion of urban population



(b) Incidence curves

FIGURE 11 —CLIMATE CHANGE INCIDENCE CURVE FOR RURAL AND URBAN POPULATION

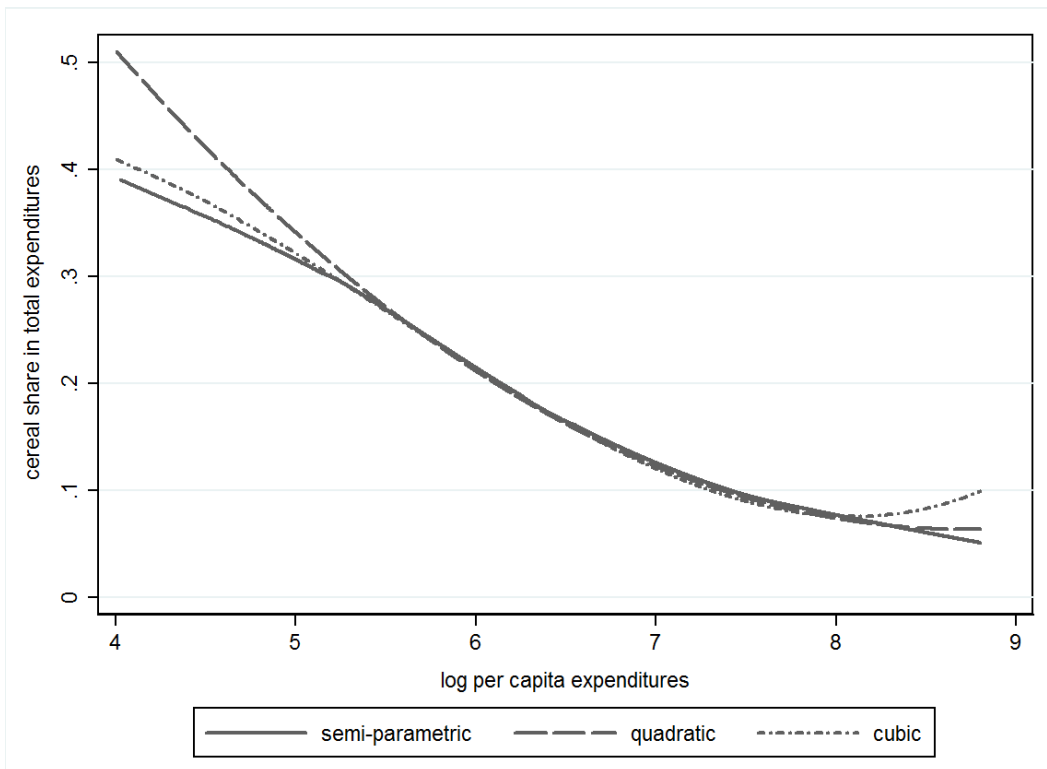


FIGURE 12 —ALL-INDIA ENDEL CURVES FOR CEREAL EXPENDITURES

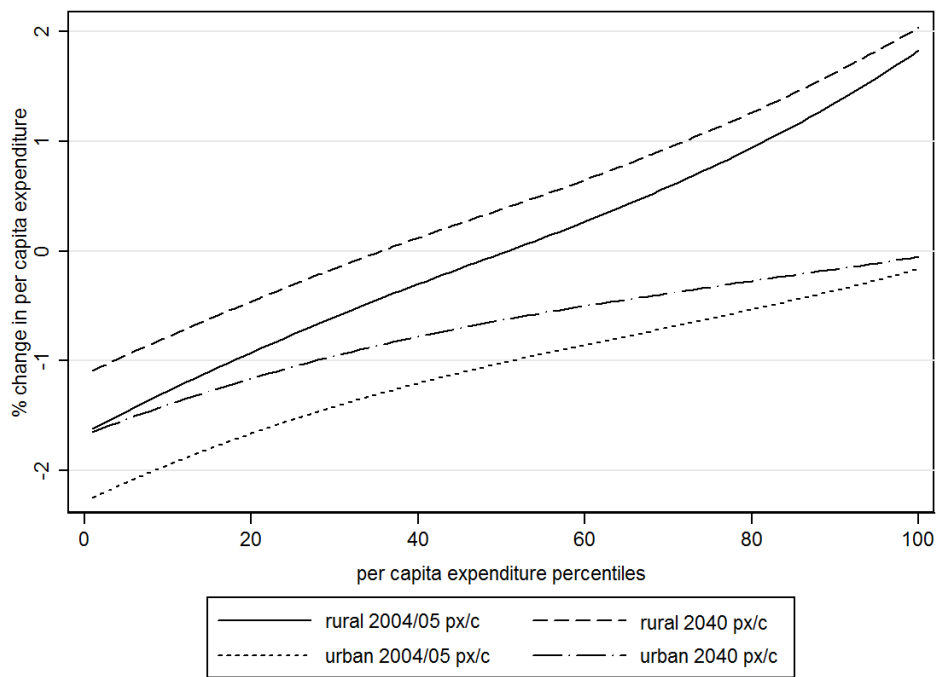


FIGURE 13 — ECONOMIC GROWTH AND ENGELS' LAW