Long-term Consequences of Short-term Precipitation Shocks: Evidence From Brazilian Migrant Households

Valerie A. Mueller[a*] Daniel E. Osgood[b]

March 23, 2007

Abstract

We test for long-term damages from precipitation shocks in rural Brazil by observing wages of households that have permanently migrated from rural to urban areas. We find that large short-term precipitation shocks damage migrants' long-term wages. We offer an analytical explanation for this outcome: credit-constrained households may be willing to accept lower wages in urban areas following the depletion of their productive assets during an adverse shock. Our empirical evidence indicates further exploration is warranted on the valuation of what mechanisms affected by natural disasters cause these long-term household losses experienced in Brazil.

JEL Classification: O1, Q1, R2
Key words: Climate, shocks, agriculture, household income, migrants

1 Introduction

Droughts and floods are prevalent in low-income countries, affecting the livelihoods of several rural and urban households. An inability to effectively reduce shock-induced income loss can lead to severe long-term consequences for households in developing countries. For example,

[a]Center on Globalization and Sustainable Development, The Earth Institute at Columbia University, 150 Hogan Hall, 2910 Broadway, New York, NY 10025, US. Email: vam2105@columbia.edu.
[b]International Institute for Climate and Society, The Earth Institute at Columbia University, P.O. Box 1000, Palisades, NY 10964, US. Email: deo@iri.columbia.edu.
[*]Corresponding author. Email: vam2105@columbia.edu. Phone: +1 212-854-4312. Fax:+1 212-854-6309.

We thank Christopher Barrett, Alexander Pfaff, Cristian Pop-Eleches, Mark Rosenzweig, Glenn Sheriff, Christopher Timmins, Eric Verhoogen, and participants of the Yale Environmental Economics and Columbia Development Economics seminars for their helpful input.

precipitation shocks may directly destroy assets necessary for agriculture. Alternatively, households defer their future earning capacity in agriculture, when liquidating a substantive fraction of their assets to smooth income. Without formal insurance arrangements, savings, or credit markets, rural households in developing countries are particularly vulnerable to the income variability induced by precipitation anomalies, leaving migration as one of the few viable risk mitigation options.

By observing the income of Brazilian rural households having migrated to urban areas permanently, we detect long-term negative impacts of large precipitation shocks. Our conceptual model provides one explanation for why this may occur, although the identification of this effect is limited given the cross-sectional nature of our dataset. Once a household loses its earning potential in agriculture, low urban wages that were previously unacceptable may become attractive. Following a shock, households may prefer urban employment if they are unable to protect or replace the assets necessary to utilize their skills in agriculture. Thus, evidence of long-term climate-related damages to agricultural productivity is revealed through the wages of households that permanently migrate out of agriculture. Our findings suggest that households' inefficacy to manage risk may lead to an increase in urban poverty.

A vast literature exists measuring the ability of households to smooth consumption when facing shocks to productivity and income (Zeldes, 1989; Deaton, 1991, 1992; Paxson, 1992; Alderman and Paxson, 1994; Townsend, 1994). Recent studies have shown that households are unable to completely smooth consumption, even when informal insurance arrangements are available to the household (Jalan and Ravallion, 1999; Kazianga and Udry, 2004). Our paper takes this discussion further by assessing the long-term consequences, if any, of incomplete consumption smoothing. The approach developed in this paper is similar in spirit to the body of literature focusing on investment decisions under production and income risk in a credit-constrained economy.

Few studies offer explanations to the degenerative nature of shocks on long-term income.

Eswaran and Kotwal (1990) claim individuals are risk averse to an uneven consumption profile over time. Thus, credit-constrained households are more likely to accept lower income for lower income risk. They conclude the disparities in endowed wealth cause poor households to fall in a "low level equilibrium trap".

Rosenzweig and Binswanger (1993) argue the presence of covariate risk is an additional obstacle to individual household consumption smoothing. Borrowing resources are less available in Indian villages facing the same weather risk. The paucity of credit institutions leads low-income households to invest in less risky portfolios which reflects their limitations in smoothing consumption. Moreover, conservative portfolio choices among low-income households perpetuate the inequitable distribution of income, as risky portfolios on average yield higher returns.

Rosenzweig and Wolpin (1993) posit Indian farmers underinvest in their productive assets. In their example, bullocks generate wealth to farmers in their contribution to production, and also in terms of their sales value when facing substantial crop losses. The underinvestment in bullocks has severe implications for poor farmers who during a shock face output and income losses.

Santos and Barrett (2006) investigate the relationship between climate shocks and wealth of Ethiopian pastoralists to detect the source of poverty traps. They attribute the extant poverty to pastoralists' inabilities to protect assets from adverse shocks. A critical determinant of long-term wealth implications of shocks is herder ability. They conclude that this identification can be used to target the recipients of future asset transfers.

In our work, we use the 1995 Brazilian national household survey to measure the income impact of short-term climate shocks. Specifically, we exploit information on households who migrated from predominantly rural states to cities within the past nine years of the survey to observe the impact of past shock-induced productivity losses on current household income. In a household income regression, we include variables on ex ante risk, the mean and variance of the known distribution of precipitation in a migrant's place of origin, *and* ex post risk, the precipitation shock variables, to differentiate between households that use migration as a coping strategy, and credit-constrained households forced to permanently migrate perhaps due to the severity of the shock on their agricultural production and assets.

Rural workers are accustomed to migrating in some areas of Brazil during dry spells (Orlove and Tosteson, 1999; Finan and Nelson, 2001), and for good reason. Migration may be an effective climate adaptation mechanism for Brazil, due to the availability of transportation services in rural areas and the strong spatial variability of climate. Addressing the latter, there is dramatic geographical variation in the climate anomalies due to the El Niño Southern Oscillation (ENSO).[1] Separate stochastic processes in the north and south (Grimm et al., 1998; Hastenrath, 2000) tend to generate seasonal precipitation responses to ENSO that are strongly negatively related between the northern and southern parts of Brazil (Moura and Shukla, 1981; Nogues-Paegle et al., 2002). On average, when ENSO drives a wet year in the one region of Brazil, it drives a dry year in the other. Thus, rural households benefit from migration by moving to a rural area experiencing a favorable precipitation shock (e.g., to sugar cane plantations in São Paulo) or to an urban area during a dry period.

Our empirical findings reveal that there may be significant damages to a household's agricultural earning potential due to large precipitation shocks. Although many households may successfully migrate to urban environments to improve their wages, we find evidence that others appear to migrate to lower wages. One characterization of the latter group is farmers who migrate as a method of last resort because of losses in their productive assets. In the next section, we develop a conceptual model to illustrate how these losses may be observed in the decline of migrants' urban wages.

[1] ENSO is the irregular warming and cooling of the Pacific Ocean's surface temperature and the above atmosphere which is a primary driving process in seasonal climate variability (IRI, 2005).

2 Conceptual Framework

We present a basic conceptual model to motivate why the incomes of permanently migrated households might reveal long-term agricultural productivity damage from large precipitation shocks. Consider a household that has the choice between two sectors, an agricultural rural sector and a non-agricultural urban sector (e.g., manufacturing). In the agricultural sector, the household could earn net income y. In the manufacturing sector, the household receives net income w. A household in the agricultural sector would not be willing to accept a low income in the manufacturing sector, if it is lower than what is available in the agricultural sector. However, if the capacity to generate income is damaged by a precipitation shock, the household would be willing to accept a lower urban income, since its agricultural alternative has worsened.

Let w^* be the manufacturing sector income at which the household is indifferent between the agricultural and manufacturing sector less transition expenses. This is the threshold above which migration occurs. If a household migrates from the rural agricultural sector to the urban manufacturing sector, the household is revealing that the urban income is preferable to the rural income. Therefore, by observing how w^* changes with respect to shocks, we can gain insight into the agricultural sector productivity impacts of the shocks.

The variable relating post migration urban incomes to the long-term impact of a climate shock is Δw^*, the change in the migration threshold income w^* with respect to a shock. To develop an expression for Δw^*, we assume that the household faces an annual optimization problem. For illustration, utility is assumed to be linear in income. Each year, the household calculates optimized income in each sector, compares the income between the sectors, and selects the sector with the highest income less transition costs. At the end of the year, the household re-evaluates its situation, taking into account any unforeseen changes that have occurred, and the annual decision problem is repeated.

Annual precipitation is reflected by the random variable $\tilde{z}_l$ with mean μ_z, variance σ_z^2, and density function $\phi(z_l)$, where l indexes locations. In the long run, the distribution of precipitation at location l is observable. For a single year, precipitation is a random draw from $\phi(z_l)$.

Consider $\boldsymbol{x}$ to be a vector of household characteristics, such as skills and endowments. The capital necessary for agricultural production, or productive assets, is k. Net agricultural income, or profits, are determined by the function $\pi(z, k, \boldsymbol{x}, \boldsymbol{p})$, which is continuous, twice differentiable, and quasi-concave in z_l and k. Let $\boldsymbol{p}$ be a vector of input and output prices. The annual agricultural sector sub problem is to perform the optimization of profit:

$$\Pi = \max_{l,k} E[\pi(\tilde{z}_l, k, \boldsymbol{x}, \boldsymbol{p})]. \tag{1}$$

The household selects optimal agricultural land location $l^*$and productive assets k^*. For the manufacturing sector, there are no decision variables. Manufacturing income is simply $w(x)$. The transition cost of changing between economies (such as migration costs or the loss of squatters rights for land or water) is c.

Given this specification, $w^*(x)$, the urban income for which a rural household is indifferent between the two sectors is

$$w^*(x) = E[\Pi(\tilde{z}_{l^*}, k^*, \boldsymbol{x}, \boldsymbol{p})] + c. \tag{2}$$

During a severe shock, some of the productive assets may be lost. This loss may be directly due to a flood or drought. It may also be because the household must liquidate some of its productive assets during the year in order to maintain a minimal level of consumption. This loss does not occur for small shocks because productive assets are not physically destroyed for small deviations from mean precipitation. In addition, the household is able to tolerate small consumption losses and mitigate marginal impacts on its productive assets using whatever limited savings it has available. Thus, $\frac{\partial k}{\partial z}|_{\mu_z} = 0$. However, as the precipitation shocks become larger, non-negligible amounts of k are lost due to the shock, so $\frac{\partial^2 k}{\partial z^2}|_{\mu_z} < 0$. Note that these expressions reflect a mechanistic destruction of k directly due to a precipitation shock, not the change in what the optimal k^* would be for different levels of precipitation.

Having completed its annual ex ante optimization, a household in the agricultural sector faces an ex-post precipitation realization z_{l0}. We define the shock as $\Delta z = \mu_z - z_{l0}$ (the deviation in precipitation from the average). How is the household in a different situation following the shock? If the shock was severe, the precipitation z_{l0} would have destroyed some of the productive assets purchased at the beginning of the year, decreasing the k^* selected in the optimization to k'. A credit-constrained household may liquidate its assets to maintain consumption and protect its productive assets. Otherwise, assets may be destroyed by the large shock directly. Credit limitations prevent the household from replenishing its productive assets for next year's optimization. Thus, the household is constrained to a level of productive assets at (or below) k'. Future agricultural income is diminished due to the reduction in productive assets. Following the shock, its potential agricultural income is $\Pi' = \Pi(z_{l^*}, k'(k^*, z_{l0}), \boldsymbol{x}, \boldsymbol{p})$. We approximate this with a Taylor series expansion:

$$\Pi' \approx \Pi(\mu_z) + \tfrac{\partial \Pi}{\partial k}|_{k^*} \tfrac{\partial k'}{\partial z}|_{\mu_z} \Delta z + \tfrac{1}{2}\left(\left(\tfrac{\partial k'}{\partial z}|_{\mu_z}\right)^2 \tfrac{\partial^2 \Pi}{\partial k^2}|_{k^*} + \tfrac{\partial \Pi}{\partial k}|_{k^*} \tfrac{\partial^2 k'}{\partial z^2}|_{\mu_z}\right) \Delta z^2 \ . \tag{3}$$

Since $\frac{\partial k}{\partial z}|_{\mu_z} = 0$, the change in agricultural income due to the shock is:

$$\Delta \Pi = \frac{1}{2} \frac{\partial \Pi}{\partial k}|_{k^*} \frac{\partial^2 k'}{\partial z^2}|_{\mu_z} \Delta z^2. \tag{4}$$

Applying (4) to (2), we obtain Δw^*, the income indifference threshold due to shock Δz:[2]

$$\Delta w^* \approx \frac{1}{2}\frac{\partial \Pi}{\partial k}|_{k^*}\frac{\partial^2 k'}{\partial z^2}|_{\mu_z}\Delta z^2. \tag{5}$$

If $\frac{\partial \Pi}{\partial k}$ is positive, Δw^* is negative since $\frac{\partial^2 k}{\partial z^2}|_{\mu_z} < 0$. The household has a lower potential rural income and is willing to accept lower urban income than prior to the shock due to its loss of productive assets. It may face the lower wage indefinitely, if, for example, the household abandons the agricultural sector altogether to maintain short-term consumption following the shock. In this case, there is no mechanism to rebuild the productive assets necessary to yield an agricultural economy income sufficient to meet minimum consumptive needs. The poverty trap literature identifies this outcome during the transition between economies (for example, migration costs) or the liquidation of assets lost in the transition process (such as badly defined rights for land or water) (Azariadis and Stachurski, 2006).

It is possible that precipitation shocks could be reflected in the wages of migrants without long-term damages to agricultural productivity. If the poorest agricultural households have no capacity to smooth fluctuations in production between years, then we might observe them sorting into urban jobs due to short-term production impacts and consumption constraints. In this case, lower urban wages reflect the selection effect, and not an actual long-term loss of productivity. Because it impacts people so close to the margin, this sorting effect is likely to have some of its components observable through the first order shock term. In addition, since seasonal migration is common, a migrant who's long-term productive capability has not been damaged would be likely to eventually return to the greater agricultural income.

[2] Note that in our regression specification, we estimate the difference of log incomes, which approximates $\Delta w^*/w^*$.

3 Data

We use the 1995 Pesquisa Nacional de Amostra Domicilios (PNAD), a nationwide household survey administered by the Brazilian government annually. The PNAD questionnaire collects a variety of information regarding demographics, wages, housing, and migration for each member of the households sampled. This information is georeferenced to the municipality level, analogous to a U.S. county, using Brazilian GIS data (IBGE, 1998). The household survey data precisely reflects on the destination of a migrant, i.e., the municipality in which the migrant was surveyed. However, the information regarding previous locations of migrants is limited to the survey design. Households were asked if they migrated, how many years ago, and from what state they were born and migrated from (not the municipality).[3]

To formulate our climate variables, we use the NOAA NCEP CPC global precipitation data product. This is based on a combination of weather station observations, satellite estimates, and numerical model predictions to provide a comprehensive georeferenced grid of climate indicators at a 2.5 degree spatial resolution. The resolution is at a convenient scale for our state and municipality level census data. Monthly data begins in 1979. In Brazil, this could incorporate a great deal of weather station data, having an extensive network of 13,197 stations. Some of stations were in operation as far back as 1897.[4]

Figure 1 is a map of Brazil and decomposes the country into five regions, North (N), Northeast (NE), Southeast (SE), South (S), and Midwest (MW). Figures 2 and 3 present the long-term climate distributions faced across Brazil. They also illustrate the scale of the 2.5 degree pixel. Climate pixels data were spatially averaged using state boundaries, also consistent with the reporting of migrants' origins at the state level. Climate data mean,

[3]The survey asks individuals if they were born in the municipality in which they are currently living. Thus, the birth municipalities of individuals that reported being born in their municipality of residence is known. Additionally, the survey asks the individual if he migrated from another municipality in the current state of residence, without having him report name of the municipality.

[4]Agência Nacional de Energia Elétrica (ANEEL), http://ingrid.ldeo.columbia.edu/SOURCES/.ANEEL

variance, and spatial aggregation operations were performed through queries to the IRI Data Library web interface.[5] In Figure 2, it is clear that the Amazonian Northwest has some of the wettest areas in Brazil, and the Northeast has some of the driest areas. The Northeast also experiences substantial variation in precipitation (see Figure 3).

In the analysis, we focus on rural households that migrated to urban areas within the last nine years of the 1995 survey. Approximately 85,000 households were sampled in the 1995 survey. Data regarding which state migrants lived in previously is available for this group (not the municipality). There are several migration patterns in Brazil (e.g., rural to rural, rural to urban, urban to rural, and rural/urban to Amazon) and several motivations for the various choices of migration. For clarity, we focus on the rural to urban migration pattern. We further restrict the sample to households currently living in urban areas. Urban areas are defined as municipalities with populations greater than or equal to one hundred thousand. We also limit the migrants in the analysis to those who came from predominantly rural states. A state is declared rural if its percentage of the population living in rural areas is greater than or equal to twenty-five percent. This leaves us with 40,005 households, 2,339 of the households having migrated between one to nine years prior to the survey.[6]

Table 1 compares the demographic characteristics of migrants and non-migrants in the 95 PNAD dataset. There are few differences between these two groups. The composition of the households differs. Migrant households tend to have younger male household heads with more children than sedentary households. They also possess fewer assets, but their monthly income is slightly larger than non-migrants. One notable feature of this group is that they tend to come from places where the precipitation variation is slightly greater. This provides preliminary evidence that there may be migrants who manage income risk by moving to areas where the risk is lower.

[5] http://iridl.ldeo.columbia.edu

[6] In an effort to focus on long-term migration, we do not consider households that migrated during the year of the survey.

The descriptive statistics for the precipitation shock and shock-squared variables are also included in Table 1. Since our model focuses on an annual production cycle, a natural definition of the shock is the deviation of the realized value of rainfall the year prior to migration from the mean precipitation level. A positive value of the shock variable implies a dry shock, and a negative value, a wet shock. The value of the shock variable depends on the year the household migrated and the place of origin. From Table 1, it appears that, on average, our sample of migrants faced normal to slightly drier-than-average climate prior to the move.

Table 2 exhibits the distribution of the origin and destination for migrant households. Two noticeable patterns exist. First, most of the migrants come from the northeastern region of Brazil which is particularly vulnerable to dry spells. Second, of those coming from the northeast, many migrate to the cities in the northeast and southeast regions.

A final consideration regards the timing of the survey and migration. Our sample of migrants traveled during the time frame of 1986 to 1994. The challenge we face in identifying the effect of climate-induced migration on differences in household income is the concurrent events of precipitation shocks and economic recessions. For example, there were five severe economic recessions in the period of 1987 to 1992 (Chauvet, 2002), and quite a few dry periods between 1983-1994, most of them in the northern region of Brazil.[7] We include region and time fixed effects in our empirical models to differentiate between the effects of recessions and precipitation shocks on household income differences.

[7] Figure 4 compares the 1980-1996 trend of precipitation (indexed from negative one to one) in the northern and southern regions of Brazil. From the figure, it is clear that the northern region has experienced quite a few dry periods during that period.

Table 1: Descriptive Statistics of Demographic and Climate Variables for Migrant and Non-migrant households

Variable	Migrants Mean	Migrants Std. Dev.	Non-Migrants Mean	Non-Migrants Std. Dev.
Head of Household Characteristics				
Education	7.88	4.67	7.60	4.53
Age	36.49	12.02	43.54	13.42
Male	0.83	0.38	0.79	0.41
Household Characteristics				
Household size	3.82	1.79	4.05	1.87
Number of members less 10 years old	0.93	1.04	0.78	0.99
Homeownership	0.49	0.50	0.74	0.44
Own Washing Machine	0.28	0.45	0.35	0.48
Monthly Income (1995 Reais)	909.59	1175.67	898.66	1298.54
Distance Migrated (km)	1194.40	752.18		
Climate				
Precipitation mean (mm/day)	3.65	1.12	3.65	0.88
Precipitation variance (mm/day)	0.51	0.21	0.41	0.19
Precipitation shock (mm/day)	0.02	0.67		
Precipitation shock squared (mm/day)	0.45	0.68		
Observations	2,339		37,666	

Table 2: Percent of Migrant Sample by Origin and Destination

Destination	Origin North	Northeast	Southeast	South	Midwest
North	5.22	4.23	0.64	0.81	0.51
Northeast	3.16	13.68	1.15	0.47	0.13
Southeast	1.97	20.69	9.49	4.49	0.81
South	0.98	1.15	0.98	7.01	0.51
Midwest	4.36	10.30	3.55	1.92	1.75

4 Econometric Model

We use the sample of households to estimate the regressions for total monthly income y per household h in their current location j. Our benchmark regression is the following:

$$\begin{aligned} \log y_{hj} &= \beta_0 + Age_{hk}\beta_1 + \log \widehat{y}_{hk}\beta_2 + \overline{R}_k\beta_3 + \sigma^2_{R_k}\beta_4 + (\overline{R}_k - \boldsymbol{\delta}^t_h \mathbf{R}^t_{hk})\beta_5 \\ & +(\overline{R}_k - \boldsymbol{\delta}^t_h \mathbf{R}^t_{hk})^2\beta_6 + \beta^d_{hj} + \beta^o_{hk} + \beta^{do}_{hjk} + \beta^t_{hk} + \varepsilon_{hjk}. \end{aligned} \tag{6}$$

We include exogenous household characteristic variables in the regression to control for idiosyncratic features of the household that may affect their current earnings. The head of household's age Age_{hk} is a proxy for the head of household's stage in the life cycle. To control for differences in labor market opportunities across locations given the observable skill set of the household, we impute a wage index $\log \widehat{y}_{hk}$. The wage index at the state of origin is imputed from a monthly wage regression that includes all workers in the survey that were actively employed the time of the survey and were at least ten years of age.[8] The estimates of the parameters and standard errors for this regression are presented in Table 8 in the Appendix.[9] The parameters are used to impute the potential wage market level for individuals at their original location. Then, the imputed wages for all employed household members are summed to calculate the control for monthly household labor market wage opportunities at its place of origin. Both the current household income $\log y_{hj}$ and the

[8] Because a fraction of actively employed individuals do not report income, we estimate the wage regression accounting for sample selectivity using Full Information Maximum Likelihood (FIML) (see Greene, 1997). The Log-likelihood ratio test indicated that we can reject the null hypothesis that the correlation coefficient is equal to zero at the one percent significance level. The variables included in the selection equation are dummy variables indicating the age category of an individual, gender, the relationship to the head of household, receipt of pension, receipt of remissions, receipt of dividends or interest from investments, and regional dummy variables. OLS regressions ignoring sample selectivity were also estimated although the results are not reported here. The R-squared was 0.52, indicating that much of the variation in wages is being captured in this regression.

[9] To simplify the presentation of the wage equation parameters, we include the average and standard deviation of the parameters across twenty-seven states. Only twenty-six dummy variables were included as state fixed effects. We omit a dummy variable for the state of Sao Paulo.

imputed wage in the state of origin $\log \widehat{y}_{hk}$ correspond to the terms expressed in (5), the income indifference thresholds of the urban non-agricultural and rural agricultural sectors. Note that this control does not adjust for unobserved characteristics of households. Rather, the imputed wage $\log \widehat{y}_{hk}$ is more effective in controlling for cross-sectional differences in wage markets between regions than for imputing the true wages that a household might expect.

The climate variables in the regression control for ex ante risk on income, and elicit the losses of the constrained households characterized in the theoretical model. They are $\overline{R}_k$, $\sigma^2_{R_k}$, $(\overline{R}_k - \boldsymbol{\delta}^t_h \mathbf{R}^t_{hk})$, and $(\overline{R}_k - \boldsymbol{\delta}^t_h \mathbf{R}^t_{hk})^2$. The latter two variables are the precipitation shock and shock-squared. They depend on the year the household migrated t and the place of origin k. $\boldsymbol{\delta}^t_h$ is a row vector of year dummy variables indicating the year a household migrated, and $\mathbf{R}^t_{hk}$ is a vector of precipitation values for municipality k in a given year t. The precipitation shock variables $(\overline{R}_k - \boldsymbol{\delta}^t_h \mathbf{R}^t_{hk})$ and $(\overline{R}_k - \boldsymbol{\delta}^t_h \mathbf{R}^t_{hk})^2$ capture the lower wages a credit-constrained household is willing to accept because of the shock's impact on their production and assets. Our model recognizes that some rural households may be quite capable of self-insuring against small shocks by liquidating or consuming their assets, but their coping strategies are limited for larger shocks. After losing their productive assets, households may migrate to urban areas accepting a lower urban income because of the damage imposed on their local alternatives to agricultural production.

Regional destination β^d_{hj}, origin β^o_{hk}, and destination by origin β^{do}_{hjk} fixed effects, as well as time fixed effects β^t_{hk} capture the effect of unobservable spatial and time variables that influence current income. The spatial fixed effects are four dummy variables for the original location of each household, i.e. north, northeast, south, and midwest,[10] four dummy variables for the present location of each household, and fifteen dummy variables that interact the regional dummy variables for the original and present locations of the household.[11] Spatial

[10] A dummy variable for the southeastern region is omitted from the regression.

[11] We omit the dummy variable that interacts the northeastern original location and northeastern current location dummy variables from the model.

fixed effects account for the effect of moving costs, migration networks (Munshi, 2003), and access to city amenities (Blomquist et al., 1988) on income. Time fixed effects are eight year-specific dummy variables representing the actual year an individual migrated. These capture competing shocks that may affect income-generating activities, such as a recession (von Wachter and Bender, Forthcoming).

The last term ε_{hjk} in regression (6) represents the idiosyncratic error term. We assume that the error term is distributed normal with mean equal to zero and variance σ^2. We compute heteroskedasticity-robust standard errors for all specifications of model (6) (White, 1980).

5 Identification and Results

The results from our baseline regression are presented in Table 3. The experiment variable (the shock-squared parameter) is negative *and* significantly different from zero, supporting our hypothesis that large short-term shocks have long-term negative consequences on income. The magnitude of the shock parameter is close to zero and insignificant, implying that small wet and dry shocks have little measurable effect on long-term household income. Households may be somewhat able to cope with small unanticipated shocks so that they do not lead to long-term consequences. However, the shock-squared parameter reveals that a large shock does appear to cause long-term damage, perhaps from a loss of productive assets of the income of resource-constrained households.[12]

It is interesting to briefly discuss the parameters recovered for mean and variance of historical precipitation. Since they are included primarily as a control for observable climate information, are cross-sectional, and are recovered without careful effort to ensure identifi-

[12] To confirm shocks are in fact unanticipated, we estimate our wage regression including five additional variables for the shocks in the five years after the PNAD survey was conducted. Four out of five of these parameters were statistically insignficiant.

cation, these results are suggestive at best. Precipitation mean is negative and marginally significant, while precipitation variance is positive and significant. This may indicate that those migrating from harsher (e.g., drier, more variable) climates do not necessarily represent the poorest groups, and may actually have developed skills that are valuable in urban settings. The long-term migration process, may in general, yield greater household wages for those more likely to face climate shocks as long as they are not migrating in response to losses arising from a specific shock.

To improve identification of the shock-squared parameter, we include a battery of additional confounding variables that may bias the shock variable parameters. We include regional origin, regional destination, regional origin by destination, and time fixed effects to absorb bias that may exist on the shock variable parameters. Table 9 displays the results from regressions that exclude spatial and time fixed effects, include only origin fixed effects, include only origin and destination fixed effects, and include only origin, destination, and origin by destination fixed effects. The results from these regressions are used to compare the impact of controlling for unobservable factors such as moving costs, migration networks, and limited access to city amenities, all having a potential positive impact on income. The experiment results are fairly robust with the shock-squared parameter differing only slightly in magnitude and retaining significance across specifications. After controlling for spatial fixed effects, the magnitude of the shock-squared parameter estimates are slightly larger. The regression controlling for time-specific shocks yields a slightly smaller (in magnitude) negative and significant parameter on the shock-squared variable than in the fixed effect regressions in Table 9, although the time fixed effects are not jointly significantly different from zero in our baseline regression.

To avoid problems due to endogeneity, we used few controls for auxiliary household characteristic variables that influence income in our baseline model. As a diagnostic, we estimated additional specifications of the household income regression that include additional

Table 3: Household Income Regressions

Variable	Baseline	Excluding Imputed Wage	Excluding Precipitation Mean, Variance
Intercept	0.5372***	6.4013***	0.4179***
	(0.1681)	(0.2019)	(0.0132)
Age	-0.0066***	0.0018	-0.0066***
	(0.0012)	(0.0018)	(0.0012)
Precipitation mean	-0.0570*	-0.0607	
	(0.0346)	(0.0519)	
Precipitation variance	0.2526***	0.3850***	
	(0.0943)	(0.1387)	
Precipitation shock	-0.0004	-0.0503	-0.0032
	(0.0266)	(0.0412)	(0.0265)
Precipitation shock squared	-0.0664***	-0.0902**	-0.0492**
	(0.0245)	(0.0396)	(0.0239)
Log of imputed wage	0.9963***		0.9974***
	(0.0185)		(0.0185)
R-squared	0.59	0.05	0.59
Observations	2,339	2,339	2,339
F test: All coefficients=0	94.90***	3.65***	99.08***
F test: Shock variables=0	3.98**	2.84*	2.27*
F test: Origin Fixed effects=0	2.18*	2.57**	2.40**
F test: Destination Fixed effects=0	0.69	1.87	0.92
F test: Origin×Destination Fixed effects=0	2.21***	2.79***	2.22***
F test: Time Fixed effects=0	0.93	0.05	0.88

1 Heteroskedasticity-robust standard errors are in parenetheses. ***, **, and * indicate the parameters are significant at the 1, 5, and 10 percent critical levels.

2 Regional destination fixed effects, regional origin fixed effects, fixed effects interacting regional destination and regional origin dummy variables, and time fixed effects are included in all wage model specifications.

controls. One regression included eight dummy variables indicating the number of household members in a given age category, and the head of household's educational attainment. Another regression included the distances between the origin and destination states to proxy the impact of moving costs explicitly on household income. The results are presented in Table 10. We include these only as diagnostics, since many of the household regressors are endogenous, and we lack suitable instruments for these variables.[13] Note the distance and distance-squared variable parameters are not statistically significant upon controlling for spatial fixed effects, perhaps because the spatial fixed effects act as an exogenous control for moving costs. Comparing the results across models, the specification of household characteristics does not impact our experiment results. In other words, the order of magnitude and significance of the shock-squared parameter is robust across these model specifications.

Recent attention has been placed on the potential self-selection of migrants into more favorable labor markets (Chiswick, 1999). Our theoretical model, however, characterizes a potential disfavorable selection problem. If our sample of migrants reflect low-income households that sort themselves into urban low-paid jobs irrespective of any long-term damage from shocks, then we must address that selection problem. There is some preliminary evidence to suggest that the migration patterns we observe is not simply the sorting of poor households. For example, the descriptive statistics for migrants do not support such a story. Instead of being poorer and less educated than non-migrants, the mean education and income of migrants is slightly higher. In addition, the mean and variance terms in the benchmark regression imply that migrants from drier, more variable climates may face higher wages in urban environments, which is the opposite of what might be expected if the wage loss from a shock was driven by poor households sorting without long-term damage. Because these pieces of evidence are merely suggestive of the selection processes that may be occurring, we

[13] For example, Angrist and Krueger (1999) provide an in depth discussion of the endogeneity of education variables.

perform more formal corrections for potential bias from sample selection.

We test whether our shock variables are identifying the result of low-income households sorting into urban jobs due to short-term production impacts and consumption constraints rather than long-term losses of productivity. The possible bias induced by the former would exist if omitted variables associated with wealth are correlated with the shock. Suppose poorly educated households are responding to the negative shocks by migrating. In this case, we would expect our parameter of interest to be negatively biased. To demonstrate the potential bias induced by factors related to wealth, we estimate our wage regression replacing education, self-employment, and ownership of a washing machine as the dependent variable. The regression results reveal the shock variables are uncorrelated with factors related to wealth, such as education and washing machine ownership (see Table 5).[14] The shock-squared coefficient in the self-employed regression is positive suggesting a potential positive bias on the shock-squared parameter in the wage regression rather than the anticipated negative bias.

Realizing dry and wet shocks may have distinct impacts on household incomes, we estimated a model that includes five climate dummy variables to determine if the sign of the shock influences the empirical results: very dry, dry, normal, wet, and very wet.[15] The dummy variable that was omitted from the regression was classified as normal climate. Normal climate was considered to be one standard deviation of the shock variable centered around the mean. The dummy variables dry and very dry were given a value of unity if the value of the shock variable fell between 0.5 and 1 standard deviation above the mean, and greater than 1 standard deviation above the mean, respectively. The dummy variables wet

[14] Two households in the original sample of migrants were excluded from the education regression because they did not provide information on the head of household's education level. Two hundred seventeen households in the original sample of migrants were excluded from the self-employment regression because they did not provide information on the head of household's occupational status.

[15] Of the sample of migrants, 19.20, 8.38, 11.76, and 16.08 percent fell into the very wet, wet, dry, and very dry categories.

and very wet were given a value of unity if the value of the shock variable fell between -0.5 and -1 standard deviation below the mean, and less than 1 standard deviation below the mean, respectively. Table 6 displays the results from this regression. When stratifying the shock variable into this manner, only the dry dummy variable parameter was significantly different from zero. Since the sign and order of magnitude of the impact of a dry shock is the same as that of the shock-squared term in previous regressions, it is likely that the parameter for the shock-squared term in the benchmark regression is largely driven by this category of shocks.

Our final model specification attempts to differentiate between the short-term and long-term impacts of precipitations shocks. The model includes variables that interact the shock variables, and dummy variables indicating whether the household migrated 1 to 4 years ago, and 5 to 9 years ago. The regression results are presented in Table 7.[16] The only interaction that is significant is the shock squared×Migrate 5 to 9 years ago, indicating that long-term (5-9 year) process does indeed appear to be driving our findings. This supports our hypothesis that extreme climate events impose long-term consequences on household welfare.

6 Conclusion

Large precipitation shocks can damage households' productive capacities in rural areas. It is our contention that households lose a number of assets during large shocks, preventing them from continuing working in the agricultural sector. Since many of these households are credit-constrained, they are unable to weather the shock by borrowing or depleting savings. The wage where households are indifferent between working in the rural agricultural and

[16] Diagnostic regressions were also performed that included interactions between the historical climate and shock variables (such as between the shock squared and the variance parameter). None of these interactions were significant.

urban non-agricultural sectors, decreases, perhaps to meet a minimum level of consumption.

Using the Brazilian census data, some of the households facing these constraints are observed. Our empirical evidence suggests that these households lose long-term income in the agricultural sector because of short-term severe precipitation shocks. Although migrants face lower long-term income in urban areas, staying in their original locations may have led to an even worse outcome than urban migration. The observed decline in income may reveal the loss of worthwhile alternatives as opposed to a damage from migration itself.

We use current household economic and demographic information, and information regarding their places of origin to identify these losses. The results provide insight to understanding the long-term damage to household welfare from large precipitation shocks in rural areas. Household data describing socioeconomic backgrounds before and after the shock would help improve the identification of climate-induced losses and allow for a direct test of the hypothesis posed in our conceptual model. Future work will also involve better understanding the risk management schemes of these farmers to observe the contents necessary for effective risk-reducing portfolio strategies in areas with variable climate.

Table 4: Household Income Regressions Accounting for Migrant Selectivity

Variable	Heckman	OLS
Intercept	0.4399***	0.5309***
	(0.1777)	(0.1682)
Age	-0.0075***	-0.0067***
	(0.0012)	(0.0012)
Precipitation mean	-0.0600*	-0.0565*
	(0.0344)	(0.0346)
Precipitation variance	0.2607***	0.2542***
	(0.0937)	(0.0944)
Precipitation shock	-0.0009	0.0005
	(0.0264)	(0.0266)
Precipitation shock squared	-0.0651***	-0.0660***
	(0.0224)	(0.0245)
Log of imputed wage	0.9980***	0.9978***
	(0.0184)	(0.0185)
Sigma	0.6339	
	(0.0134)	
Rho	0.0969	
	(0.0639)	
R-squared/Value of Log-likelihood function	-9633.12	0.59
Observations	35,090	2,337
F/Wald test: All coefficients=0	3502.89***	94.82***
F/Wald test: Shock variables=0	7.69**	3.96**
F/Wald test: Origin Fixed effects=0	8.95*	2.01*
F/Wald test: Destination Fixed effects=0	3.43	0.70
F/Wald test: Origin×Destination Fixed effects=0	31.80***	2.20***
F/Wald test: Time Fixed effects=0	8.12	0.95
Wald test: Rho=0	2.27	

1 Heteroskedasticity-robust standard errors are in parentheses. ***, **, and * indicate the parameters are significant at the 1, 5, and 10 percent critical levels.

2 Regional destination fixed effects, regional origin fixed effects, fixed effects interacting regional destination and regional origin dummy variables, and time fixed effects are included in all wage model specifications.

3 The selection equation in the Heckman model includes an intercept, and variables that interact four head of household age categorical variables (greater than 25, 26 to 40, 41 to 55, and greater than 55 years old) and four regional dummy variables representing the head of household's birth place (north, northeast, south, and midwest). Fourteen of the seventeen coefficients in the regression were statistically significant at the 5 percent critical level. The selection predicts 91 percent of the observations correctly, where the threshold used to determine that value was 0.25 (see Greene, pp. 892-893). Approximately, 7 percent of the sample of households migrated in the last nine years of the survey

4 The OLS regression is the same as the baseline regression, except it includes two less migrant households, as they do not report information on the head of household's age or birth place. The results are presented to assess the difference between the results from the OLS and Heckman specifications.

Table 5: Regressions for Detecting Selection of Rural Migrants into Lower Paying Jobs

Variable	Education	Self-Employed	Own Washing Machine
Intercept	-14.694***	0.156	-1.087***
	(0.893)	(0.115)	(0.096)
Age	-0.078***	0.006***	0.006
	(0.009)	(0.001)	(0.019)
Precipitation mean	0.023	0.018	-0.013
	(0.181)	(0.023)	(0.056)
Precipitation variance	1.214**	-0.007	-0.132
	(0.500)	(0.065)	(0.056)
Precipitation shock	0.007	0.017	-0.009
	(0.149)	(0.019)	(0.018)
Precipitation shock squared	0.049	0.035**	0.009
	(0.122)	(0.018)	(0.015)
Log of imputed wage	3.948***	-0.038***	0.208***
	(0.120)	(0.013)	(0.012)
R squared	0.47	0.04	0.22
Observations	2,337	2,122	2,337

1 Heteroskedasticity-robust standard errors are in parentheses. ***, **, and * indicate the parameters are significant at the 1, 5, and 10 percent critical levels.

2 Regional destination fixed effects, regional origin fixed effects, fixed effects interacting regional destination and regional origin dummy variables, and time fixed effects are included in all wage model specifications.

Table 6: Household Income Regression Including Very Dry, Dry, Wet, and Very Wet Dummy Variables

Variable	
Intercept	0.5013***
	(0.0168)
Age	-0.0066***
	(0.0012)
Precipitation mean	-0.4217
	(0.0346)
Precipitation variance	0.2236**
	(0.0939)
Wet	0.0100
	(0.0518)
Very wet	-0.0584
	(0.0443)
Dry	-0.0917**
	(0.0460)
Very dry	-0.0408
	(0.0451)
Log of imputed wage	0.9964***
	(0.0185)
R-squared	0.59
Observations	2,339
F/Wald test: All coefficients=0	91.36***
F/Wald test: Shock dummy variables=0	1.39
F/Wald test: Origin Fixed effects=0	1.97*
F/Wald test: Destination Fixed effects=0	0.74
F/Wald test: Origin×Destination Fixed effects=0	2.12***
F/Wald test: Time Fixed effects=0	0.87

1 Heteroskedasticity-robust standard errors are in parentheses. ***, **, and * indicate the parameters are significant at the 1, 5, and 10 percent critical levels.

2 Regional destination fixed effects, regional origin fixed effects, fixed effects interacting regional destination and regional origin dummy variables, and time fixed effects are included in all wage model specifications.

3 Five climate dummy variables were generated to determine if the sign of the shock influenced the empirical results. The dummy variable that was omitted from the regression was classified as normal climate. Normal climate was considered to be one standard deviation of the shock variable centered around the mean. The dummy variables dry and very dry were given a value of unity if the value of the shock variable fell between 0.5 and 1 standard deviation above the mean, and greater than 1 standard deviation above the mean, respectively. The dummy variables wet and very wet were given a value of unity if the value of the shock variable fell between -0.5 and -1 standard deviation below the mean, and less than 1 standard deviation below the mean, respectively.

Table 7: Household Income Regression with Shock and Time Interaction Variables

Variable	
Intercept	0.5493***
	(0.1684)
Age	-0.0066***
	(0.0012)
Precipitation mean	-0.0559
	(0.0346)
Precipitation variance	0.2487***
	(0.0944)
Precipitation shock×Migrate 1 to 4 years ago	-0.0702
	(0.0553)
Precipitation shock×Migrate 5 to 9 years ago	-0.0190
	(0.0408)
Precipitation shock squared×Migrate 1 to 4 years ago	0.0485
	(0.0641)
Precipitation shock squared×Migrate 5 to 9 years ago	-0.0906***
	(0.0319)
Log of imputed wage	0.9951***
	(0.0185)
R-squared	0.59
Observations	2,339
F/Wald test: All coefficients=0	91.06***
F/Wald test: Shock variables=0	3.13***
F/Wald test: Origin Fixed effects=0	2.12*
F/Wald test: Destination Fixed effects=0	0.74
F/Wald test: Origin×Destination Fixed effects=0	2.22***
F/Wald test: Time Fixed effects=0	1.03

1 Heteroskedasticity-robust standard errors are in parentheses. ***, **, and * indicate the parameters are significant at the 1, 5, and 10 percent critical levels.

2 Regional destination fixed effects, regional origin fixed effects, fixed effects interacting regional destination and regional origin dummy variables, and time fixed effects are included in all wage model specifications.

References

Alderman, H., Paxson, C. H., 1994. Do the poor insure? a synthesis of the literature on risk and consumption in developing countries. In: Bacha, E. L. (Ed.), Economics in a Changing World. Vol. Development, Trade, and the Environment. Macillan, London, pp. 48–78.

Angrist, J. D., Krueger, A. B., 1999. Empirical strategies in labor economics. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics Vol. 3, Part. Elsevier, pp. 1277–1366.

Azariadis, C., Stachurski, J., 2006. Handbook of Economic Growth. http://emlab.berkeley.edu/users/chad/Handbook.html, Ch. Poverty traps.

Blomquist, G. C., Berger, M. C., Hoehn, J. P., March 1988. New estimates of quality of life in urban areas. American Economic Review 78 (1), 89–107.

Chauvet, M., 2002. The brazilian business and growth cycles. Revista Brasileira de Economia 56 (1), 75–106.

Chiswick, B. R., 1999. Are immigrants favorably self-selected? American Economic Review 89 (2), 181–185.

Deaton, A., 1991. Saving and liquidity constraints. Econometrica 59 (5), 1221–1248.

Deaton, A., 1992. Household saving in LDCs: Credit markets, insurance and welfare. Scandinavian Journal of Economics 94 (2), 253–273.

Eswaran, M., Kotwal, A., 1990. Implications of credit constraints for risk behavior in less developed countries. Oxford Economic Papers 42 (2), 473–482.

Finan, T. J., Nelson, D. R., 2001. Making rain, making roads, making do: Public and private adaptations to drought in ceara, northeast brazil. Climate Research 19 (2), 97–108.

Greene, W. H., 1997. Econometric Analysis, 3rd Edition. Prenctice Hall, New Jersey, U.S.

Grimm, A., Ferraz, S., Gomes, J., November 1998. Precipitation anomalies in southern brazil associated with el nino and la nina events. Journal of Climate 11 (11), 2863–2880.

Hastenrath, S., March 2000. Interannual and longer-term variability of upper air circulation in the northeast brazil-tropical atlantic sector. Journal of Geophysical Research-Atmospheres 105 (D6), 7327–7335.

Instituto Brasileiro de Geografia e Estatistica (IBGE), 1995. Pesquisa nacional por amostra de domicilios. CD Rom.

Instituto Brasileiro de Geografia e Estatistica (IBGE), 1998. Malha municipal digital do brasil.

International Research Institute for Climate Prediction (IRI), 2005. online.

Jalan, J., Ravallion, M., 1999. Are the poor less well insured? evidence on vulnerability to income risk in rural china. Journal of Development Economics 58 (1), 61–81.

Kazianga, H., Udry, C., 2004. Consumption smoothing? livestock, insurance, and drought in rural burkina faso. Economic growth center discussion paper no. 898, Yale University.

Moura, A. D., Shukla, J., December 1981. On the dynamics of droughts in northeast brazil: Observations, theory, and numerical experiments with a general circulation model. Journal of Atmospheric Sciences 38 (12), 2653–2675.

Munshi, K., May 2003. Networks in the modern economy: Mexican migrants in the u.s. labor market. Quarterly Journal of Economics 118 (2), 549–599.

Nogues-Paegle, J., Mechoso, C. R., Fu, R., Berbery, E. H., Chao, W. C., Chen, T.-C., Cook, K., Diaz, A. F., Enfield, D., Ferreira, R., Grimm, A. M., Kousky, V., Liebmann, B., Marengo, J., Mo, K., Neelin, J. D., Paegle, J., Robertson, A. W., Seth, A., Vera, C. S., Zhou, J., 2002. Progress in pan american clivar research: Understanding the south american monsoon. Meteorologica 27 (12), 1–30.

Orlove, B. S., Tosteson, J. L., 1999. The application of a seasonal to interannual climate forecasts based on el niño-southern oscillation (ENSO) events: Lessons from australia, brazil, ethiopia, peru, and zimbabwe. Workshop on Environmental Politics Working Papers WP99-3, Berkeley, http://globetrotter.berkeley.edu/EnvirPol/WP/03-Orlove.pdf.

Paxson, C. H., 1992. Using weather variability to estimate the response of savings to transitory income in thailand. American Economic Review 82 (1), 15–34.

Rosenzweig, M. R., Binswanger, H. P., 1993. Wealth, weather risk and the composition and profitability of agricultural investments. The Economic Journal 103 (416), 56–78.

Rosenzweig, M. R., Wolpin, K. I., 1993. Credit market constraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investments in bullocks in india. Journal of Political Economy 101 (21), 223–244.

Santos, P., Barrett, C., 2006. Heterogeneous wealth dynamics: On the roles of risk and ability. Working paper, Department of Applied Economics and Management, Cornell University.

Townsend, R. M., 1994. Risk and insurance in village india. Econometrica 62 (3), 539–591.

von Wachter, T., Bender, S., Forthcoming. In the right place at the wrong time-the role of firms and luck in young workers careers. American Economic Review.

White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica 48 (4), 817–830.

Zeldes, S., 1989. Consumption and liquidity constraints: An empirical investigation. Journal of Political Economy 97 (2), 305–346.

APPENDIX

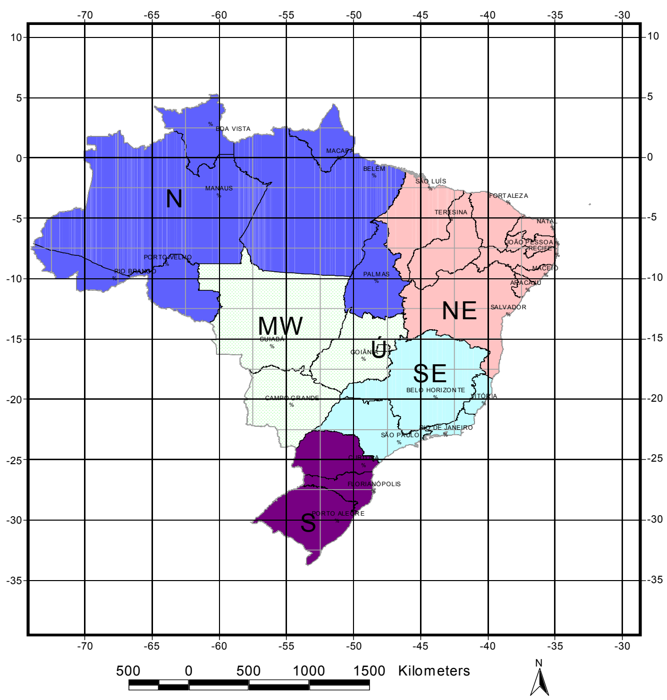

Figure 1: Brazil

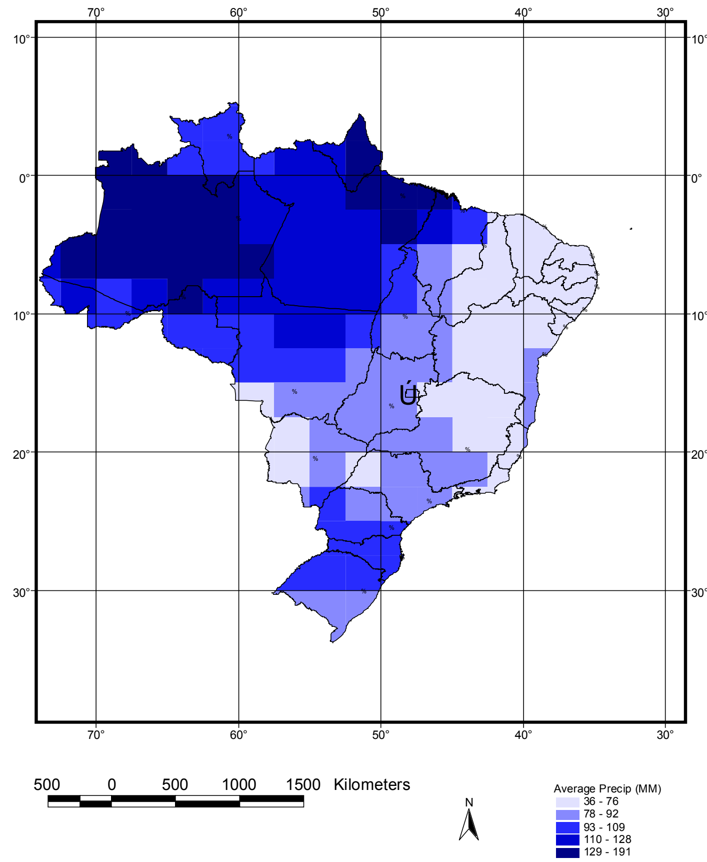

Figure 2: Average Annual Precipitation (Jan 1979-Sep 2004)

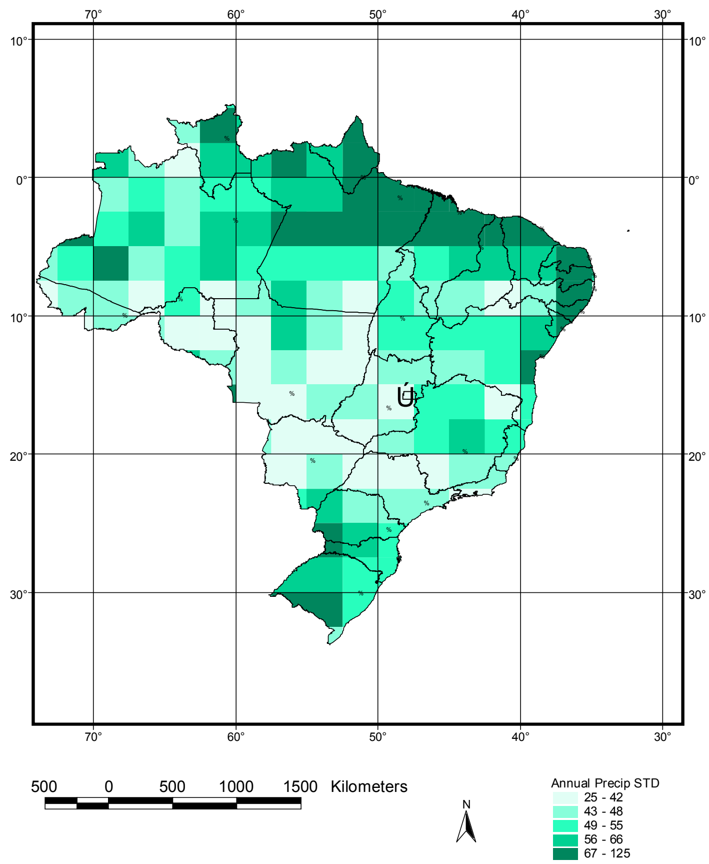

Figure 3: Variation in Annual Precipitation (Jan 1979-Sep 2004)

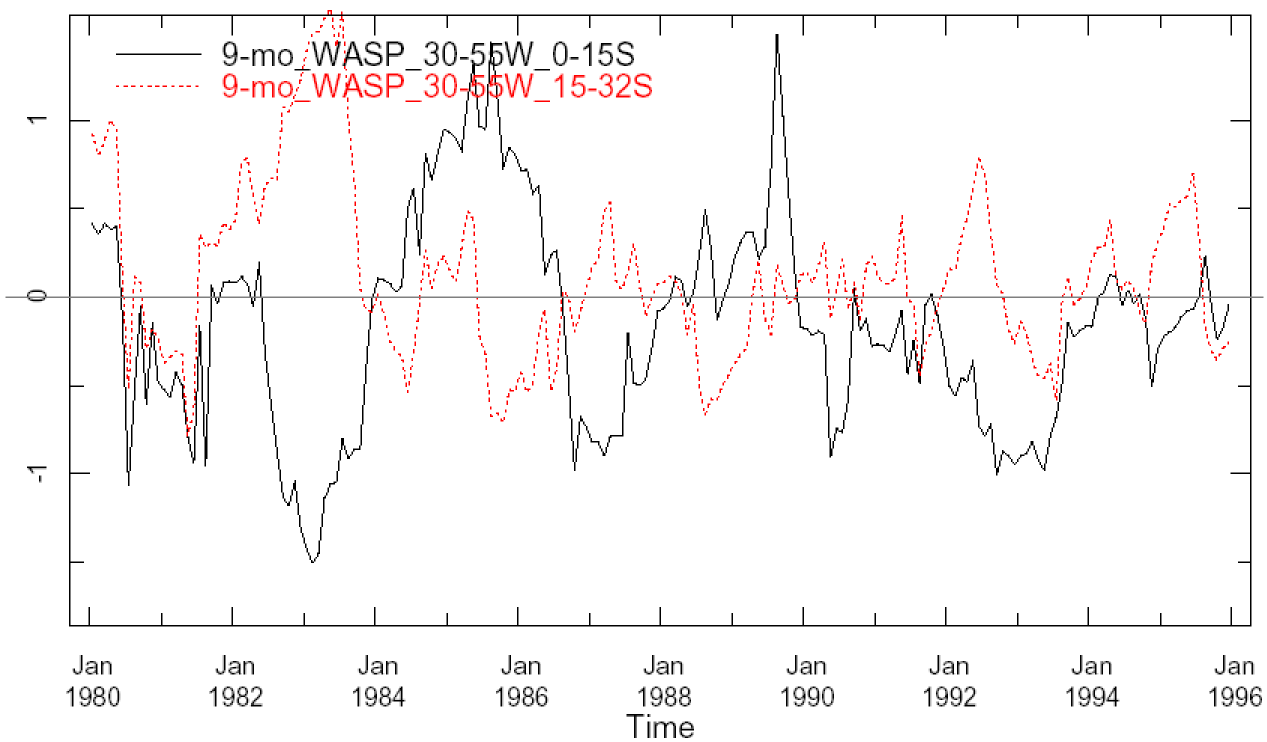

Figure 4: Precipitation anomalies in the North and South of Brazil

Table 8: Wage Regression Used to Impute Household Income in Place of Origin

Variable	Parameter	Std.Error (Std. Dev.)
Wage Equation		
Intercept	2.7783	(0.0396)
Male	0.5347	(0.0810)
Black	-0.1484	(0.0762)
Education	0.1201	(0.0151)
Age	0.0768	(0.0136)
Age-squared	-0.0008	(0.0002)
State Fixed Effect	0.1792	(0.2764)
Selection Equation		
Intercept	1.7102	0.0259
Age		
18 to 25	0.8722	0.0143
26 to 40	0.9800	0.0166
41 to 55	0.7875	0.0196
greater than 55	0.3462	0.0247
Male	-0.1772	0.0124
Relationship to Head of Household		
Spouse	-1.2118	0.0178
Child	-0.9824	0.0179
Relative	-0.6536	0.0237
Other	-0.2367	0.0454
Household Size	-0.0242	0.0021
Receipt of Auxiliary Income		
Pension	-0.3024	0.0200
Remittance	-0.5766	0.0617
Dividends	0.3592	0.0387
Region		
North	-0.1825	0.0220
Northeast	-0.5399	0.0126
South	-0.3823	0.0144
Midwest	-0.2883	0.0181
Lambda	-0.3621	0.0075
Sigma	0.7416	0.0020
Rho	-0.4883	0.0095
LR test: Rho=0	1825.67	
LLF	-178837	
Observations	125,190	

Table 9: Household Income Regression with Fixed Effects

Variable	No FE	Origin FE	Origin, Destination FE	Origin, Destination, Origin×Destination FE
Intercept	0.3831***	0.5508***	0.5541***	0.5260***
	(0.1231)	(0.1586)	(0.1585)	(0.1620)
Age	-0.0060***	-0.0059***	-0.0065***	-0.0065***
	(0.0011)	(0.0011)	(0.0011)	(0.0011)
Precipitation mean	-0.0110	-0.0693**	-0.0658**	-0.0558
	(0.0127)	(0.0331)	(0.0336)	(0.0347)
Precipitation variance	0.1798***	0.2859***	0.2410***	0.2442***
	(0.0717)	(0.0928)	(0.0929)	(0.0938)
Precipitation shock	-0.0280	-0.0279	-0.0249	-0.0278
	(0.0214)	(0.0215)	(0.0216)	(0.0215)
Precipitation shock squared	-0.0615***	-0.0641***	-0.0668***	-0.0695***
	(0.0222)	(0.0222)	(0.0221)	(0.0223)
Log of imputed wage	0.9927***	0.9909***	0.9972***	0.9942***
	(0.0179)	(0.0183)	(0.0184)	(0.0184)
R squared	0.58	0.58	0.58	0.59
F test: All coefficients	546.02***	328.93***	241.76***	120.74***
F test: Shock variables	3.96**	4.28***	4.59***	4.92***
F test: Origin Fixed effects		1.06	1.24	2.22*
F test: Destination Fixed effects			6.17***	0.68
F test: Origin×Destination Fixed effects				2.24***

1 Heteroskedasticity-robust standard errors are in parentheses. ***, **, and * indicate the parameters are significant at the 1, 5, and 10 percent critical levels.

Table 10: Additional Specifications of the Household Income Regression

Variable			
Intercept	0.3780**	0.7956***	0.5705***
	(0.1642)	(0.1817)	(0.1722)
Number of household members			
aged 11 to 17		0.0062	
		(0.0163)	
aged 18 to 25		0.0318*	
		(0.0167)	
aged 26 to 40		-0.0129	
		(0.0217)	
aged 41 to 64		-0.0584**	
		(0.0244)	
aged greater than 64		-0.1929***	
		(0.0447)	
Education			
5 to 8 years		-0.0177	
		(0.0339)	
9 to 12 years		-0.0407	
		(0.0340)	
greater than 12 years		0.3546***	
		(0.0604)	
Age			-0.0066***
			(0.0012)
Distance			-0.0076
			(0.0091)
Distance squared			0.0002
			-0.0003
Precipitation mean	-0.0628*	-0.0480	-0.0553
	(0.0344)	(0.0340)	(0.0349)
Precipitation variance	0.2504***	0.2590***	0.2548***
	(0.0948)	(0.0933)	(0.1005)
Precipitation shock	-0.0019	-0.0029	-0.0015
	(0.0269)	(0.0261)	(0.0265)
Precipitation shock squared	-0.0647***	-0.0638***	-0.0670***
	(0.0247)	(0.0238)	(0.0245)
Log of imputed wage	0.9820***	0.9099***	0.9964***
	(0.0184)	(0.0275)	(0.0185)
R squared	0.58	0.60	0.59
F test: All coefficients	95.70***	87.48***	90.30***
F test: Shock variables	3.67**	3.83**	4.03**
F test: Distance variables			0.35

1 Heteroskedasticity-robust standard errors are in parentheses. ***, **, and * indicate the parameters are significant at the 1, 5, and 10 percent critical levels.

2 The distance variable is divided by 100.

3 Regional destination fixed effects, regional origin fixed effects, fixed effects interacting regional destination and regional origin dummy variables, and time fixed effects are included in all wage model specifications.