The impact of positive agricultural income shocks on rural Chinese households*

Jessica Leight

February 7, 2018

* *Jessica Leight (corresponding author) is a professor at American University, Washington, D.C., U.S.; her email address is leight@american.edu.* *The author thanks Abhijit Banerjee, Loren Brandt, Esther Duflo, William Gentry, Doug Gollin, Rohini Pande, Albert Park, and Xiaobo Zhang for detailed feedback and seminar participants at the MIT political economy breakfast, Williams, Cornell, IFPRI, Beijing University, the Chinese Center for Agricultural Policy, and the Hong Kong University of Science and Technology for comments. A supplementary online appendix is available with this article at the World Bank Economic Review website.*

Abstract

In the post-collectivization period, rural Chinese households were required to sell part of their grain output to the state at a below-market price; however, increases in this quota price beginning in 1993 generated substantial positive income shocks. These income shocks also varied cross-sectionally in accordance with crop composition given that quotas were systematically larger for rice-producing households, generating a quasi-random source of variation in the size of the shock driven by climatic variation in suitability for rice cultivation. Households induced to experience relatively larger income shocks show evidence of decreased agricultural investment, increased investment in non-agricultural businesses, and increased migration as households gain increased income, consistent with the hypothesis that credit constraints may have constrained some households from entering non-agricultural production ex ante. In addition, there is evidence that these households were concentrated among households who had not previously diversified out of agriculture.

1 Introduction

A large literature in development economics has sought to analyze the effects of both positive and negative income shocks on poor households, and in particular to probe whether such shocks lead to any systematic shifts in the allocation of productive activities. In China, rural households in the post-1983 period were required to sell a fixed quota of grain to the state at a below-market price as part of the so-called Household Responsibility System implemented following decollectivization. However, the central government began to raise this price gradually in the mid-1990s (Huang and Rozelle, 2002). This shift was equivalent to a reduction in the size of the lump-sum tax imposed on rural households (Huang, 1998), or conversely, a positive income shock; moreover, given that grain quota prices remained high thereafter, the shock was plausibly viewed as an increase in both current and future expected income.

This paper estimates the impact of this policy-driven shock on investment in agricultural and non-agricultural production, seeking to identify whether a lump-sum increase in income leads to a reallocation of productive activities. Crucially, the size of this shock varied systematically as a function of the crop cultivated, given the stark difference in the treatment of different crops under the quota system (Huang, Rozelle, Ha and Li, 2002). Predominantly rice-growing areas systematically experienced larger income shocks, given that they were subject to larger mandatory quota quantities than predominantly wheat-growing areas. The identification strategy here exploits this cross-sectional variation in conjunction with shocks to the quota price over time. While cultivation of rice is an endogenous household decision, the specification of interest analyzes the impact of increasing quota income for areas that have a greater propensity to cultivate rice based on their climatic conditions, conditional on village and province-year fixed effects.

The results suggest that in an instrumental variables specification in which income is predicted by the interaction of the propensity to cultivate rice and quota price, an increase in quota income leads to a decrease in agricultural investment and an increase in non-agricultural investment and migration. Households experiencing increases in quota income also show evidence of a higher probability of accessing a non-zero amount of credit, and increased consumption of non-staple goods. While the results should be interpreted cautiously, the observed pattern is consistent with the hypothesis that credit constraints may have been binding ex ante for some households, limiting their diversification into non-agricultural productive activities. The positive income shock generated by the quota price increase relaxes this constraint.

The results are also robust to a series of specification checks. The effects of the quota shocks are observed primarily for households that have not diversified into non-agricultural production ex ante, i.e., those that are plausibly constrained. There is no evidence that these patterns reflect differential trends in areas with different climatic conditions, or differential shocks in policy or other output prices correlated with the quota price. There is also little evidence that variations in quota policy are endogenously driven by varying conditions in the local economy.

This paper contributes to several related literatures. First, a large literature has analyzed the effects of income shocks, both negative and positive, on developing country households' pattern of asset-holdings and livelihood diversification. A number of papers provide evidence that households facing negative shocks derived from adverse weather outcomes or natural disasters experience asset disinvestment as well as consumption declines, and these effects are particularly salient for poor households (Carter, Little, Mogues and Negatu, 2007; Carter and Lybbert, 2002; Janzen and Carter, 2013; Hoddinott, 2006). In analyzing positive shocks, by contrast, Barrett et al. (2001) evaluate the effects of exchange rate and food aid shocks on household diversification in Cote d'Ivoire and Kenya. Gertler, Martinez and Rubio-Codina (2012) and Sadoulet, de Janvry and Davis (2001) find that households benefiting from cash transfer programs in Mexico invest transfers in productive assets, resulting in a long-term increase in consumption; Gilligan, Hoddinott and Taffesse (2008) find similar results examining a government assistance program in Ethiopia. Blattman, Fiala and Martinez (2014) find that cash transfers to young adults for non-agricultural activities lead to increases in assets and income.[1]

Second, a number of papers have analyzed how households respond to evolution in the rate of return on agricultural investments. Foster and Rosenzweig (2004) estimate the impact of shocks to the returns to agriculture in India induced by the adoption of Green Revolution technology and find that industrial growth is fastest in areas where agricultural growth is lagging. Jedwab (2011) finds that positive price shocks to cocoa during cocoa booms in Ghana and the Ivory Coast lead to an increase in urbanization. Kaboski and Townsend (2011) evaluate the impact of a government-sponsored microcredit program as a major income shock in rural Thailand on consumption and investment. Most recently, Bustos et al. (2015) present evidence that the introduction of genetically

[1]While this empirical analysis does not directly address the question of poverty dynamics in rural households, and focuses instead on dynamics in sectoral allocation over time, there is also a large literature analyzing the presence or absence of poverty traps in developing country contexts. Kraay and McKenzie (2014) provide a useful overview. In particular, it may be useful to highlight that the evidence of steadily increasing income and substantial sectoral mobility for rural households in China renders this context distinct from many other developing countries, particularly in sub-Saharan Africa, where there is more evidence of income stagnation and/or structural immobility (Giesbert and Schindler, 2012; Kwak and Smith, 2013; Naschold, 2012, 2013; Quisumbing and Baulch, 2013).

engineered soybean seeds in Brazil leads to industrial growth.

Third, this paper contributes to the body of work focused on analyzing the grain quota system as it operated in China. Lu (1999) describes broad trends in the operation of the quota system over time, and Rozelle et al. (2000) describe reform over the decade of the 1990s, particularly the backlash following early attempts at liberalization. Huang, Rozelle and Wang (2006) evaluate the quota system as part of a broader analysis of capital flows in and out of the agricultural sector. Further discussion of existing work describing the grain quota system can be found in Section 2.2.

More recently, a broader literature in macroeconomics has documented large gaps in labor productivity between the agricultural and non-agricultural sectors in developing countries, and relatedly, a large gap in labor productivity comparing the agricultural sector in developing and developed countries (Gollin, Lagakos and Waugh, 2014; Lagakos and Waugh, 2013). While this paper does not directly analyze productivity gaps across sectors, the suggestive evidence of constraints on exit from agriculture is consistent with the presence of a large and relatively unproductive agricultural sector.

Relative to the existing literature, this paper provides new evidence about the effect of a quasi-exogenous income shock on productive diversification for rural households, and has the advantage of exploiting a plausibly permanent income shock. It is also one of the first papers to analyze the impact of variation in quota policy in China on investment in non-agricultural production. The paper proceeds as follows. Section 2 describes the data and institutional background, and Section 3 presents the identification strategy. Section 4 presents results and robustness checks, and Section 5 concludes.

2 Data and institutional background

This section briefly summarizes the data employed, provides an institutional overview of the grain quota system, and describes the information available in the dataset about the quota system.

2.1 Overview of data sources

The data employed here is the China Research Center for the Rural Economy panel dataset, collected in a sample of 206 villages in 13 provinces in China between 1986 and 2001, excluding 1992 and 1994.[2] Provinces observed include Shanxi, Jilin, Jiangsu, Zhejiang, Anhui, Henan, Hunan, Sichuan and Gansu. The surveys prior to 1993 were considerably briefer, and thus the primary analysis is restricted to the post-1993 panel; 1993 is identified as the baseline year.

In addition, the analysis employs two climatic or agronomic data sets. The first is monthly climatic data from a network of weather stations in China collected by the Carbon Dioxide Information Analysis Center (CDIAC). Using data on the latitude and longitude of the county centroid, climatic measurements for each county and year are constructed by interpolating using the inverse distance weighting method.[3] The second is an index of climatic suitability for rice cultivation generated by the FAO, incorporating precipitation and temperature as well as other soil and topographic features. More specifically, the analysis utilizes the FAO's index of suitability for high input-level irrigated rice, given that Chinese agriculture is characterized by a relatively high level of input use and irrigation, and estimate the mean index value within the county in which each sampled village lies. More details on how these indices are used in the analysis are provided in Section 3.

2.2 Grain policy in China

Prior to 1978, agricultural production in China was highly collectivized. The primary unit of production was the production team, a cluster of 20 to 30 households that jointly farmed agricultural land and sold the resulting output. After the death of Mao Ze-

[2] A randomly selected sample of households in each surveyed village forms the panel; the mean number of households in a village-year cell is 69.

[3] Each interpolation employs only data from stations within 150 kilometers of the county centroid. The average number of stations employed to construct climatic data for a county is three. While county-level weather data reported directly by counties to provincial weather bureaus is also available, using weather station data has the advantage of ensuring consistency of data quality and reporting methods across the sample.

dong, however, major changes in agricultural policy were introduced. The household was reinstated as the primary unit of production under a system known as the Household Responsibility System. Each household was provided with an allocation of land for its own use, while land title continued to be held by the village (Brandt et al., 2002).

In addition, households were mandated to deliver a fixed amount of quota grain to the state at a preset price. Excess production could be sold to the state at a higher, above-quota price, or at rural markets (Lin, 1992). Grain thus sold to the state was primarily funneled via official grain bureaus to urban consumers, who were entitled to purchase staple consumption goods at subsidized prices (Lu, 1999; Rozelle et al., 2000).

This mandatory grain procurement by the state, widely known as the grain quota system, remained a key dimension of agricultural policy through the early post-2000 period. The volume of grain quotas did decline after 1995 relative to earlier in the decade; subsequent surveys have found that state procurement accounts for 25–30% of rural households' grain output, though in this sample, quota sales constitute only around 10% of grain production on average (Wang et al., 2003; Sicular, 1995).[4]

More importantly, there was a substantial increase in quota prices beginning in 1993. The primary analysis in this paper will focus on the period 1993 to 2002; evidence in this sample suggests that at the beginning of the period, the quota price was 30% lower than the market price. In some years later in the period, however, the prices were nearly equal. This change is equivalent to a large reduction in the magnitude of the lump-sum tax imposed on farmers via the mandatory quota system (Huang, 1998). More recently, grain quotas have been widely eliminated. The timing of this change varied, but generally occurred post–2002, after the period of interest here (Huang, Wu and Rozelle, 2009).

[4]Sicular (1995) reports aggregate statistics on the percentage of farm output sold as quota sales between around 1985 and 1993; grain quotas fall rapidly in this period, accounting for the fact that the percentage of total output constituted by the quota in this sample, observed between 1993 and 2002, is lower. Wang et al. (2003) report data on the operation of the grain quota system assembled from secondary data sources (i.e., not household surveys) in 25 counties in Zhejiang, Jiangsu and Sichuan for the years 1980 to 1999. They report that the share of the quota in total output is less than 25% after 1995. Given that two of the three provinces included in the survey are relatively high-income, high-productivity provinces, the fact that quota sales constitute a relatively higher percentage of output may be unsurprising; there is in general a positive correlation between the level of agricultural production and the proportion of total output sold as quota sales.

It is also useful to note that the quota system treated different grain crops very differently, and rice was consistently the crop most heavily penalized. Procurement prices for rice were lower, and rose more slowly. In addition, the quota quantity constituted a higher proportion of total production for rice producers than for wheat producers (Huang, Rozelle, Ha and Li, 2002). The identification strategy in this paper will exploit the systematic cross-sectional variation in quota quantity that is correlated with cross-sectional variation in the propensity to cultivate rice, in conjunction with these fluctuations in the quota price over time, in order to identify variation in quota income.

In terms of the implementation of the quota system, Wang et al. (2003) provide a detailed analysis of the system's operation during the period examined here. Quota prices were set by the central government and were constant at the national or provincial level, while quota quantities in the form of grain deliverable per household were generally set by county leaders for all villages within their jurisdiction, and varied with village characteristics. Wealthier, more productive, and more agricultural villages were assigned higher quotas. Given that these analyses employ secondary data that report county- or village-level averages for quota quantity, there is no analysis of intra-village variation in quota quantity. Furthermore, quota quantities for a given village changed only incrementally once set by county leaders. By contrast, the central government had discretion over the grain quota price and would change it annually (Rozelle et al., 2000; Lin, 1991).

2.3 Descriptive statistics around quota income

This section provides an overview of the quota and market price over time, and describes the key sources of variation in the quota price.

Quota and market price over time The RCRE household data used in this analysis provides information on the price for both mandatory grain quota sales and market sales, as well as the volume of grain quota sales at the household level.[5] Figures 1a and 1b

[5] After they have fulfilled their grain quota, rural households also have the option to sell their excess production to the government at a higher, negotiated procurement price. Huang, Rozelle, Ha and Li

report the mean market and quota price of grain by year between 1986 and 2002, and additionally capture spatial variation in prices, showing in each year the 5th and 95th percentile of the within-year, cross-village distribution of prices; Figure 1a shows the cross-village variation in market prices, and Figure 1b shows the cross-village variation in quota prices.[6] Prior to 1991, both prices are low and stable. The market price begins to dramatically increase after 1991, and the quota price likewise increases, though more slowly. Cross-village variation in quota price is relatively limited, while the cross-village variation in market price is more substantial.

Figure 1c shows mean income from quota sales at the household level (reported in hundreds of yuan) by year, again reporting the 5th and 95th percentile; the 5th percentile is uniformly zero, corresponding to households that report no quota sales. In addition, it is useful to estimate the magnitude of the implicit quota tax, defined as the difference between the market and quota price multiplied by the quota quantity, as a percentage of total income. The evolution of the quota tax is summarized in Figure 1d; it peaks in 1995, at around 5% of income, and is lowest in 1998 and 1999, where it is nearly zero for all households.

In order to highlight the difference between the cross-household and cross-village distributions, Figures S1a and S1b in the on-line appendix (available with this article at the World Bank Economic Review website) reproduce these figures, but also add the corresponding 5th and 95th percentiles of the within-year distribution of village-level means of the quota income and the quota tax. (It is evident that the majority of the variation is cross-village, rather than cross-household and within-village. This is consistent with the evidence from the literature that the primary determinants of quota

(2002) show that these negotiated procurement prices are generally intermediate between the quota price and the wholesale market price. In this data, negotiated procurement prices are not reported, and no distinction is made between these two types of sales, as both represent the price that the rural producer would face for the marginal unit of production. The market price is thus used to denote the price for marginal grain sales.

[6]Price data is constructed at the household level for households that report market or quota sales, but a large number of observations are missing, corresponding to cases in which households do not report market sales in particular year. Accordingly, price data is collapsed to the village-year level, as every cell has sufficient data to estimate a market price and a quota price. Observations corresponding to the top and bottom 1% of observed prices across villages are trimmed to avoid undue influence of outliers.

policy vary at the village level.

Focusing on the quota tax, there is clearly considerable variation in the implicit tax rate posed by the quota, with some households reporting rates of up to 10%, or even higher in 1995. In addition, this tax rate is systematically higher for lower-income households; given that there is relatively little within-village variation in quota quantity, the quota tax is regressive. More specifically, the implicit tax posed by the quota system is always close to zero for households in the top quantile of income, but it averages 2% of income for households outside the top income quantile, and 5% of income for households outside the top income quantile prior to 1996.[7]

Finally, it is useful to highlight while the market price is fluctuating during this period, it is primarily shifts in the quota price that serve to close the gap between the two prices. In particular, the quota price is consistently below the market price in the early part of the period, but it converges upward to the market price after 1996. This convergence generates a positive income shock that is also plausibly expected to be permanent, given that quota prices subsequently remain high.

Defining an alternative quota price measure $\tilde{P}$ The graphical evidence presented here clearly suggests that some of the variation in the quota price is driven by variation in the market price; more specifically, the quota price can be described as a linear function of the market price, $\alpha + \beta P_m + \epsilon$. The parameters α and β reflect the average relationship between the market and quota price over time, and in general $P_q < P_m$.

Given that shifts in the market price can have many other complex effects, however, this analysis will not exploit variation in the quota price that reflects underlying fluctua-

[7]Previous papers in the literature have analyzed shocks that varied in size. Gilligan et al. (2008) in their analysis of an Ethiopian safety net program do not provide an estimate of the transfer's size relative to income, bur rather focus on measures of caloric intake. Blattman et al. (2014) analyze a much larger shock in northern Uganda, equal to approximately 100% of the participants' annual income; however, this paper also reports much larger effects than will be observed here. Barrett et al. (2001) analyzed a food-for-work program implemented in Kenya that provided wages equal to approximately 15% of annual income, and Gertler et al. (2012) analyze transfers provided by Progresa, for which there is substantial variation in the relative size of the payment depending on household structure. In general, the shock analyzed here is somewhat smaller in magnitude, but also differs in that it was plausibly permanent.

tions in the market price. Rather, the primary objective is to exploit variation in quota policy in which the government sets a quota price that is abnormally high (or low) relative to the market price — that is, increases (decreases) in the quota price **conditional** on the market price. A high quota price shock is characterized as a positive ϵ, resulting in the quota price exceeding the predicted quota price based on prevailing market prices. There will correspondingly be years in which the quota price is lower than the predicted quota price based on prevailing market prices, and ϵ is negative.

In order to generate estimates of these shocks, the mean quota price in each village v in province p in year t is regressed on the corresponding market price in the following equation, conditional on village fixed effects λ_{vp}.

$$P^q_{vpt} = \beta P^m_{vpt} + \lambda_{vp} + \epsilon_{vpt} \quad (1)$$

The residual from this equation is then constructed, denoted $\tilde{P}_{vpt}$. By construction, the residual is uncorrelated with the market price, and all subsequent analysis of the quota price will employ this variable $\tilde{P}_{vpt}$. Further specifications presented in Section 4.1 demonstrate that the primary results are robust to re-defining the quota price. The hypothesis that fluctuations in $\tilde{P}_{vpt}$ may be correlated with other observable covariates is explored further in the on-line appendix.[8]

Variation in quota quantity and price across time and space In the primary analysis, the quota price variable $\tilde{P}_{vpt}$ will be interacted with a cross-sectional climatic index and employed as an instrument for quota income, defined as the income received by households from their grain quota sales. The empirical strategy seeks to exploit variation across villages in quota quantity and variation over time in quota price given that the institutional background suggests these are the primary dimensions of variation for each variable. However, this hypothesis can also be confirmed in the data.

If quota quantity is regressed on village fixed effects and year fixed effects, village

[8]These results can be found in Section S6.2.

fixed effects account for around 35% of the variation in quota quantity, while year fixed effects account for only 7% of variation. If household fixed effects are included instead of village fixed effects, the R-squared increases to 67%.[9] For the quota price, the pattern is inverted. If the quota price variable is regressed on year fixed effects, the R-squared is around 50%; if year fixed effects are replaced by province-year fixed effects, the R-squared again increases to 70%.

Accordingly, quota income is decomposed as the product of a quota price $\tilde{P}$ that is time-varying but does not systematically vary by village, and a quota quantity Q that varies across localities (and to a limited extent within them), but is approximately constant over time. Importantly, quota quantity not only varies across localities, but varies systematically, and will generally be higher for richer and more agriculturally productive villages; this comparative static will also be confirmed in the data.

3 Identification strategy

This section first identifies the key identifying variation in quota quantity that will be exploited, and then reports the first stage.

3.1 Endogeneity of quota quantity

As already noted, quota quantity is endogenous, correlated with many observable and unobservable local characteristics. My objective here is to identify a time-invariant, cross-sectional variable correlated with quota quantity and interact it with the price in order to generate an instrument for quota income. Given the widespread analysis in the literature of the systematic cross-sectional differences in quota policy implementation that correspond to variation in crop composition, variation in the propensity to cultivate rice will be the key source of variation this analysis will exploit.

[9]This is consistent with the hypothesis that there is some within-village variation in quota quantity, in addition to the cross-village variation discussed in Wang et al. (2003). Again, previous analyses of quota quantity have primarily employed secondary data reported at the village or county level, and thus have been unable to analyze within-village variation in quotas.

To demonstrate the robustness of this relationship and increase precision in the primary two-stage least squares results, the analysis will utilize two different variables capturing the propensity to cultivate rice based on agronomic and climatic conditions. The first measure is constructed directly from climatic data (referred to thereafter as the "climatic index"); the second measure is drawn from an index generated by the Food and Agriculture Organization (referred to as the "FAO index").

First, in China, as in other countries, the suitability of a region for rice cultivation is partially determined by temperature and precipitation. "Total temperature" is defined as the total accumulated temperature over a year for days with temperatures above 10 degrees Celsius; the agronomic literature suggests the total temperature must exceed 2000 degrees to cultivate rice in China (Shao et al., 2001). Similarly, "seasonal precipitation" is defined as mean precipitation observed between May and October, the key cultivation months for rice (Tang et al., 2010). For both variables, the mean total temperature or seasonal precipitation observed in each county over the period of interest is calculatde to generate a time-invariant climatic variable. The climatic index is constructed as the mean of both variables, and denote this index $Clim_{vp}$ for village v in province p.[10]

Second, the analysis utilizes a FAO-generated index of agronomic suitability for high-input irrigated rice. This index is time-invariant, and is calculated for each sampled locality as the mean value observed within the borders of the county in which the locality lies, denoted FAO_{vp}.

It is important to clarify the level at which both indices vary. The variables of interest ($Clim_{vp}$ and FAO_{vp}) are constructed using the geographic coordinates reported for the county in which each village lies, given that this is the lowest level at which linkable geographic indicators are reported for this dataset; village-level geocoding information is not publicly available. However, in this sample, cross-county heterogeneity is virtually equivalent to cross-village heterogeneity, as only one village is sampled per county.[11] The

[10]Precipitation is re-scaled to have the same mean as total temperature. The top and bottom 2% of climatic index observations are trimmed to minimize the effects of outliers.

[11]There are two pairs of villages located in the same county in the core sample.

underlying variation for the climatic index is drawn from geo-coded data reported by 730 weather stations across China. To take into account spatial correlation in climatic characteristics, the standard errors in all subsequent specifications will be estimated employing two-way clustering at the province and year level, to allow for arbitrary correlation in the climatic index across observations in the same province.[12]

The relationship of interest between climatic conditions and quota quantity can be estimated by regressing quota quantity on these climatic indices, conditional on province-year fixed effects; in addition, the direct correlation between quota quantity and rice area is estimated. Quota quantity for household i in village v, province p, and year t is denoted Q_{ivpt}; area cultivated in rice is denoted A_{ivpt}; and province-year fixed effects are denoted ν_{pt}. The equations of interest are thus written as follows:

$$Q_{ivpt} = \beta Clim_{vp} + \nu_{pt} + \epsilon_{ivpt} \quad (2)$$

$$Q_{ivpt} = \beta FAO_{vp} + \nu_{pt} + \epsilon_{ivpt} \quad (3)$$

$$Q_{ivpt} = \beta A_{ivpt} + \nu_{pt} + \epsilon_{ivpt} \quad (4)$$

Columns (1) through (3) of Table 2 report the results, suggesting that a one standard deviation increase in the propensity to cultivate rice increases the quota quantity by at least 40% for localities within the same province and year.[13] A one standard deviation increase in rice area increases the quota quantity by 50%. This is an effect of substantial magnitude, and it does not rely on the cross-provincial heterogeneity evident in China between southern rice-cultivating and northern wheat-cultivating provinces.

[12]More specifically, standard errors are estimated employing the ivreg2 command in Stata and the two-way clustering option.

[13]It should be noted that the sample for these and all subsequent regressions is limited to household-year observations that report both quota quantity and climatic data. Less than 1% of household-year observations do not report the quantity of grain sold for the quota. About 13% of household-year observations are missing climatic data. This reflects villages located in counties for which only outdated county codes are available, as the original county has since changed boundaries or dissolved; given that it is not possible to ascertain in which successor county the village of interest lies, these localities have been dropped. There is no correlation between the probability a given village is missing climatic data and the average quota quantity, quota price, or area reported cultivated in rice in that village.

3.2 First stage

Again, the objective of the identification strategy is to identify variation in quota quantity that is correlated only with variation within a province in the propensity to cultivate rice. This variation is then presumed to be uncorrelated with other economic or political variables that enter into the county leader's determination of quota quantity.

Accordingly, the two climatic indices are interacted with a time-varying measure of quota price to generate instruments for quota income. Columns 4 and 5 of Panel A of Table 2 report the two specifications of interest, regressing quota income on the interaction of climatic index and price conditional on village and province-year fixed effects. Standard errors are estimated employing two-way clustering at the province and year level. The estimating equation is as follows, where I_{ivpt} denotes quota income, and λ_{vp} and ν_{pt} are village and province-year fixed effects, respectively; a parallel specification is estimated employing the FAO index.

$$\begin{aligned} I_{ivpt} &= \beta Clim_{vp} \times \tilde{P}_{vpt} + \lambda_{vp} + \nu_{pt} + Clim_{vp} \times P^{sum}_{ivpt} + \eta^{clim}_{vp} \times P^{g}_{ivp,t-1} + \eta^{clim}_{vp} \times \gamma_t \\ &+ \eta^{ind}_{vp} \times \gamma_t + \epsilon_{ivpt} \end{aligned} \quad (5)$$

The price variable employed is uncorrelated with the market price of grain by construction, but controls are also included for the interaction of the climatic index and a summary variable of other agricultural prices P^{sum}_{ivpt}, and the interaction of dummy variables for quantiles of the overall climatic index η^{clim}_{vp} with the lagged market price of grain.[14] Additional control variables include climate quantile fixed effects interacted with year fixed effects and dummy variables corresponding to the quantiles of industrial employment in the village η^{ind}_{vp} interacted with year fixed effects to allow for varying trends over time in areas of differing levels of industrialization.

The results indicate that a one standard deviation increase in the interacted instru-

[14]The summary price variable constructed is the mean of the market prices of the most common agricultural products reported sold: rice, wheat, corn, soy, fruit, vegetables, and cotton.

ment leads to an increase in quota income of around 15%, and in both cases, the relationship is precisely estimated and of comparable magnitude. The magnitude is calculated by noting that a one standard deviation increase in the normalized instrument yields an increase of approximately .3 in quota income measured in hundreds of yuan, relative to a mean of 2.065.[15] To reiterate, the exclusion restriction for these specifications requires that an increase in the quota price has no differential impact across areas with varying propensity to cultivate rice, other than a varying lump-sum income shock. The objective is to capture the differential effect of an increase in the quota price on a locality where the quota quantity is higher by virtue of its climatic suitability for rice cultivation.

Several assumptions are embodied in this specification. First, a shift in the quota price is assumed to represent an income shock, rather than a price shock. Second, the implicit decomposition of quota income postulated is $Q \times \tilde{P}$, where Q is treated as time-fixed and endogenous and $\tilde{P}$ as time-varying and exogenous. Third, it is assumed that shifts in climate do not induce systematic changes in quota income other than those mediated by changes in rice area cultivated. Fourth, it is assumed that yearly fluctuations in climate are uncorrelated with shifts in the quota quantity. Evidence consistent with the first assumption is presented here, as these results are essential to defining the primary specification. Additional specification checks that suggest the secondary assumptions are also valid are presented in Section S2 of the on-line appendix.

There are two primary channels through which the assumption that a shift in the quota price is an income effect rather than a price effect could be violated: substitution in and out of agriculture, and substitution between crops.[16] The first potential channel is relevant if shifts in the quota price induce households to substitute in and out of agriculture entirely. Households not cultivating grain are still required to provide grain they have purchased or an equivalent cash payment (Brandt, Rozelle and Turner, 2004); however, when they re-optimize their production decisions in the next year, the quota

[15] Unsurprisingly, the two instruments are highly correlated with each other (correlation coefficient of .88).

[16] The potential challenge of households for which quota production is equal to total production is discussed in the on-line appendix.

price will be the price of the marginal unit of grain production. Complete exit from grain cultivation is rare. Only 10% of households report even one year in which they do not cultivate grain, and on average these households still report cultivation in about half the years surveyed. For analytical clarity, however, all households that do not report grain cultivation in every year have been dropped from the analysis. Importantly, the probability that a household reports complete exit from grain cultivation is uncorrelated with the propensity to cultivate rice, and thus, dropping these households does not create differential patterns of selection into the sample in treatment and control areas.[17]

A second channel through which an income effect could be a price effect is if households switch crops in response to changes in the quota price and begin selling a larger quantity of rice as their mandated quota (rather than the smaller mandated quantity of wheat) when the price increases. While this may not be optimal if the quota price remains below the market price, as it does on average, it could be locally optimal if the quota price is close to the market price in some village-years. In this case, if households have some discretion over the quantity they sell, then a change in the quota price can no longer be plausibly interpreted as a change in a lump-sum tax.

In order to test this hypothesis, villages are classified as heterogeneous or homogeneous in the primary grain crops of interest (rice or wheat) using a simple rule: any village in which the total amount of both rice and wheat cultivated over all observations exceeds zero are denoted as heterogeneous cultivators.[18] The remaining villages (constituting approximately 60% of all observations) are classified as homogeneous cultivators. Columns (6) and (7) of Table 2 show the results of estimating the following regression to test whether quota quantity varies year-on-year with changes in quota price. This equation is estimated for both homogeneous and heterogeneous villages, denoted "Hom"

[17]This strategy may pose a challenge for external validity, as the resulting estimates cannot be extrapolated to households for which exiting agriculture entirely is a meaningful counterfactual. However, given the extremely small number of households that show this pattern and the fact that a much more common empirical regularity – as will be elaborated further below – is households that simultaneously pursue agricultural production, non-agricultural household production and/or employment outside the household – this does not seem to be a major concern.

[18]Observations reporting cultivation of less than .01 hectare rounded down to zero.

and "Het," respectively.

$$Q_{ivpt} = \beta \tilde{P}_{vpt} + \lambda_{vp} + \epsilon_{ivpt} \tag{6}$$

The results show that the relationship in homogeneous villages is close to zero and insignificant, while the coefficient in heterogeneous villages is positive and significant. This suggests either that households are crop-switching or there are other inconsistencies in quota implementation in heterogeneous areas. For example, some households in heterogeneous production areas may be opting to produce a crop other than rice to avoid the high rice quota. If this choice is costly, when the quota price increases, households may opt to return to rice production.

Accordingly, the primary sample will be restricted to villages that are homogeneous in grain production of rice or wheat; in these villages, there is no evidence that quota quantity is endogenously determined by the price. The specifications of interest reporting the first stage in the restricted sample of homogeneous villages are presented in Columns (8) and (9) of Table 2, and suggest a one standard deviation increase in the interacted instrument leads to an increase in quota income of around 10%. The exclusion restriction for this analysis requires that fluctuations in quota price are not correlated with any other shocks that vary systematically across areas with varying propensity to cultivate rice based on their climatic conditions. In other words, the only channel through which an increase in the quota price differentially affects areas more and less likely to cultivate rice is via the differential impact of a quota price increase on quota income.

4 Results

First, it is useful to briefly characterize the primary sample using the summary statistics provided in Table 1. The average household consists of four individuals, cultivating an area of around 1.4 hectares primarily in grain.[19] 90% of household-year observations report ownership of at least one productive asset for use in agriculture (e.g., animals, tools

[19]In China, mu is the traditional unit employed for land area; 1.4 hectares is equal to 21 mu.

or machinery). 26% of household-year observations report ownership of non-agricultural capital, and about 40% report that at least one household member is engaged in wage labor outside the household.

The analysis will focus on the effect of increased quota income on investment in and income derived from agriculture, investment in and income derived from non-agricultural household businesses, outside employment, migration, borrowing, and consumption. For two outcomes of interest – investment in agriculture, and investment in non-agricultural businesses – the results are reported employing a summary variable constructed using principal-component analysis, standardized to have mean zero and standard deviation one, as well as disaggregated results for a number of separate variables. Additional details about variable construction, particularly for agricultural income and consumption, are reported in Section S1 of the on-line appendix.

For agricultural investment, a summary variable is reported constructed employing six variables: area sown, labor invested in days, value of fertilizer employed, value of seeds employed, investment in animals, and investment in tools. Fertilizer, seeds and agricultural investment are reported as expenditure in yuan, and deflated using a summary price index for agricultural inputs published by the China National Bureau of Statistics. Principal-component analysis is conducted on these six variables, and the first component is employed as the summary index.

Agricultural income is calculated valuing all agricultural production (including livestock, fish, and forestry products) at the market price observed in each village-year cell; the market price of each crop or product is calculated as a sales-weighted average of unit prices reported by households. This income measure is then deflated using a weighted Laspeyres price index constructed from the same prices, employing 1993 as the base year. (Note this calculation abstracts from any variation introduced by quota policy itself, as all production is valued at the market price rather than the quota price.)

To analyze non-agricultural production, results are again reported for variables capturing non-agricultural investment and income. For non-agricultural investment, a summary outcome measure is constructed using four variables: a dummy variable equal to one if the household reports any new cash investment in a non-agricultural business, a dummy variable equal to one if the household reports any labor invested in non-agricultural machinery, and the amount of labor and cash investment reported in non-agricultural businesses. It is important to note that both agricultural and non-agricultural investment as constructed here are flow measures: they capture new investments in the year of interest. Again, the summary index is constructed via principal component analysis, retaining the first component as the index of interest. Non-agricultural income is the sum of income from non-agricultural household businesses and wage labor.

In order to deflate non-agricultural investment and income to constant prices, an index of ex-factory prices for industrial products published by the National Bureau of Statistics is employed. Unfortunately, the RCRE panel itself does not report disaggregated sales for any non-agricultural product, and accordingly it is not possible to construct a price index for non-agricultural inputs or outputs. The use of an index of factory prices may not be ideal in the analysis of household-level businesses that may produce very different products; this point will be addressed further in the discussion of the results below.

In addition, total income is constructed equal to the sum of agricultural and non-agricultural income. Additional analysis reports a dummy variable for outside labor equal to one if the household reports any days worked outside the household and associated wage income; and a dummy variable for migration equal to one if the household reports days worked outside the township. For borrowing, a dummy variable equal to one if the household reports any access to credit (formal or informal) is utilized, as well as a continuous variable equal to the amount borrowed.

Finally, results are reported for consumption of grain staples, non-grain consumption, and total consumption. Grain consumption is the sum of expenditure on grain and the imputed value of own-grain consumption, while non-staple consumption is the sum of expenditure on all consumption items excluding staple grains (both food and non-food), and the imputed value of consumption of own-farm non-grain products. Consumption expenditure is deflated employing a province- and year-specific consumer price index generated by Brandt and Holz (2006).

4.1 Two-stage least squares

The specification of interest can be written as follows, where Y_{ivpt} denotes economic outcomes of interest and the primary independent variable is $I_{ivp,t-1}$, lagged income from grain quota sales. Standard errors are estimated employing two-way clustering at the province and year level.

$$\begin{aligned} Y_{ivpt} &= \beta I_{ivp,t-1} + \lambda_{vp} + \nu_{pt} + Clim_{vp} \times P^{sum}_{vpt} + \eta^{clim}_{vp} \times P^{g}_{ivp,t-2} + \eta^{clim}_{vp} \times \gamma_t \\ &+ \eta^{ind}_{vp} \times \gamma_t + \epsilon_{ivpt} \end{aligned} \tag{7}$$

Again, the specification includes village and province-year fixed effects, the interaction of industrial employment quantile dummies and year fixed effects, and the interaction of climatic index quantile dummies and year fixed effects.[20] Additional controls include the interaction of the climatic index and the summary agricultural price variable, measured contemporaneously, and the interaction of quantiles of the climatic index and the lagged grain market price; as the market price of grain is lagged relative to the quota price, this is a two-year lag relative to the outcomes of interest.

The results from estimating (7) are reported in Table 3, employing two instruments for lagged quota income (the interaction of the climatic index and the quota price and the interaction of the FAO index and the quota price); summary variables are reported in Panels A and B, followed by disaggregated variables in Panels C and D, and all of the dependent variables other than the dummy variables are standardized to have mean zero and standard deviation one. (The ordinary least squares results are discussed in

[20] The fixed effects included are defined with respect to quantiles of the climate index constructed employing temperature and rainfall; subsequent robustness checks will demonstrate that the primary results are robust to the addition of fixed effects defined using quantiles of the FAO index. The quantiles of the climatic index are identified within the subsample of homogeneous rice- or wheat-cultivating villages included in the analysis.

Section S3 of the on-line appendix.) For each specification in Table 3, the J statistic corresponding to the Hansen test of overidentifying restrictions is reported. This test uniformly fails to reject for the primary outcome variables.[21]

In Columns (1) and (2) of Panel A, we observe that the estimated coefficients for agricultural investment and income are negative and significant: in other words, the marginal complier household, induced to experience a larger increase in quota income because of the climatic suitability of its land for rice, actually invests less in agriculture. A 100-yuan increase in quota income in the prior year leads to a decline of around .2 standard deviations in agricultural inputs, and .1 standard deviations in agricultural income. Disaggregated results for agricultural inputs reported in Panel C show a decline in sown area of 15% that is not statistically significant, a 10% decline in labor, and a large (more than 50%) decline in the value of seeds employed. There is no evidence of significant shifts in the value of tools or animals.

By contrast, Column (3) of Panel A of Table 3 reports a positive effect of a quota income shock on non-agricultural investment of around .2 standard deviations. The disaggregated results in Panel D show increases in labor and cash investment in non-agricultural businesses that are noisily estimated, and a significant increase in the probability of investing capital or labor in non-agricultural production. This suggests that households are more likely to be establishing new non-agricultural businesses than expanding existing businesses. There is no significant increase in non-agricultural income (including income from non-agricultural businesses and wage labor), as reported in Column (4). If the same specification is estimated for income from non-agricultural businesses, however, there is a significant increase in income from this source of around .1 standard deviations. Given the differential effects for agricultural and non-agricultural income, no significant effect is observed for total income, as reported in Column (5).[22]

[21]The Kleibergen-Paap F statistic for the first stage is 5.4. The test of overidentifying restrictions fails for one of the disaggregated variables: animal investment, as reported in Column (5) of Panel C.

[22]It is also useful to briefly return to the deflation of non-agricultural investment and income employing an index of ex farm prices. Given that the significant effects here are primarily observed for the dummy variables for investment of cash or labor in a non-agricultural business, rather than the level of investment, the deflation of the level measures may not be of first-order importance. The reported results are also

In Column (6), we observe that there is no significant shift in the probability of engaging in outside labor, though the coefficient of interest is positive. (The disaggregated results in Panel D, Column (5) report the coefficient for the number of days worked outside the household as a continuous variable, and again the estimated coefficient is insignificant.) There is, however, a significant increase in the probability of migration observed in Column (7), suggesting that the positive income shock may also allow households to fund migration costs.[23]

Finally, Panel B of the same table reports the results for credit access and consumption. In Columns (1) and (2), we observe a significant increase in the probability of accessing credit of 10 percentage points on a base probability of around 20%, as well as a significant increase in the amount of credit in levels, estimated at .3 standard deviations. Unfortunately, the purpose of the loan accessed is not reported; however, if the reported magnitude of loans accessed is compared to the reported magnitude of investments in non-agricultural businesses, the median loan accessed is in fact around 30% higher than the median level of non-agricultural investment, suggesting that the increase in credit access is large enough to fund the observed expansion in non-agricultural assets. In Columns (3) through (5), there is evidence of a decrease in staple consumption that is insignificant, and an increase in all other reported consumption as well as total consumption.[24] Non-staple and total consumption increase by .1 standard deviations.

In the on-line appendix, additional results are reported for alternate measures of agricultural investment and income, as well as simple measures of human capital investment. The evidence suggests that the pattern for agricultural investment is quite consistent across different variables employed. There is no evidence of any significant effect on

consistent if the level measures are deflated using a consumer price index reported by Brandt and Holz (2006) — an approach that may be unsatisfactory given that this index also incorporates prices of agricultural products — or if nominal values are employed.

[23]The mean value of the outside labor dummy is .40, and the mean value of the migration dummy is .20.

[24]Previous research in Jensen and Miller (2008) presented evidence that rice and wheat are Giffen goods for households in poverty in urban areas in China, suggesting a negative income effect. These results do not indicate a significant negative income effect, but it should be noted that the pattern may be very different for rural vis-a-vis urban households.

household human capital investment.[25]

Taken together, these results suggest that rural households in China that experience income shocks show a clear pattern of behavior. They disinvest in agriculture – i.e., the income effect for agriculture is negative – and increase investment in non-agricultural household businesses. They also consume more non-staple goods, and access new sources of credit. This evidence is broadly consistent with the hypothesis that households are credit constrained at baseline, and the positive income shock relaxes this constraint, allowing households to enter new sectors in which the return to investments may be higher. Further evidence consistent with this pattern will be explored in the next section.

A number of alternate specifications are also explored. First, the climatic index is re-constructed using only geographic information within 50 kilometers of the county centroid, rather than 150 kilometres. Second, the specification is expanded to include additional controls for village-level migration as well as other village characteristics.[26] Third, an alternate quota price measure is constructed that is the linear difference between the quota and market prices and interact this difference with the climatic indices to construct the instruments. These specifications are reported in Table S4 of the on-line appendix, and are generally consistent; for concision, only the first set of outcome variables encompassing agricultural and non-agricultural investment and income are reported.

Further secondary checks include restricting the sample to households in which quota sales are unambiguously less than total grain sales; dropping any village-years in which there is evidence of quota phase-out; adding interactions between climatic index quantile fixed effects and the leads of the market price, rather than the lags; and adding the interactions between quantile dummy variables defined using the FAO index of propensity for rice cultivation and year fixed effects. Again, the primary results reported in Table S5 of the on-line appendix are consistent.

---

[25]These results are discussed in Section S4 of the on-line appendix.

[26]The village-level control variables include lagged controls for migration levels as a percentage of population as well as a number of additional village characteristics: the number of residents, the number of village enterprises, the size of the village labor force, value of productive assets, the quantity of arable land, area sown in grain and cash crops, and total production of grain and cotton.

It is also useful to return to the primary sample restriction included in the main results, in which the sample is limited to the sample of homogeneous rice-cultivating or wheat-cultivating villages in which there is no evidence of quota manipulation. Given the evidence that households are switching in and out of crops in heterogeneous cultivation villages, it is plausible to believe that the quota price may be the price of the marginal unit of production in these villages. In that case, the effect of a higher quota price is not merely a positive income shock, but also a price shock for agriculture. This should generate upward bias on measures of investment in agriculture, and downward bias on variables capturing non-agricultural investment. In fact, this is exactly the pattern observed in the sample of heterogeneous villages.[27]

4.2 Channels

The postulated channel for the observed effect is primarily the relaxation of ex ante credit constraints. If non-agricultural investments are at least partly lumpy and households are credit-constrained, they may be unable to diversify into new sectors ex ante; the positive shock generated by the quota price increase then relaxes this constraint, facilitating diversification. In this case, the effects should be more salient for households that are more plausibly constrained.

In order to further explore the channels through which the quota shock affects household economic behavior, it is useful to evaluate evidence of heterogeneous effects along two dimensions. First, households that had diversified into non-agricultural production prior to 1993 — households are prima facie unconstrained — are compared to households that were purely agricultural at baseline. A diversified household is defined as one that reports investment in non-agricultural assets in at least 25% of the observed years prior to 1993; at baseline, only 10% of households are already diversified employing this criteria. Unsurprisingly, households that are already diversified are characterized by significantly higher levels of total income at baseline, as well as greater baseline access to credit.

[27] Tabulations are reported in Panel E of Table S5 in the on-line appendix.

A second analysis evaluates evidence of a differential response for households in the top income quantile, as calculated using average income observed prior to 1993. These households are more likely to be already diversified. In addition, evidence has already been presented that the implicit tax posed by the quota system, as a fraction of income, is consistently low for these high-income households.

For ease of interpretation, the reduced form specification is utilized, defining a dummy variable $Nagri_{ivp}$ equal to one if the household reports any ownership of non-agricultural assets in 1993 or in any preceding year, and zero otherwise. The following specification is then estimated, interacting this dummy variable with the instrument of interest.[28]

$$\begin{aligned} Y_{ivpt} = {} & \beta_1 Clim_{vp} \times \tilde{P}_{vp,t-1} + \beta_2 Clim_{vp} \times \tilde{P}_{vp,t-1} \times Nagri_{ivp} + \beta_3 Nagri_{ivp} \\ & + \lambda_{vp} + \nu_{pt} + Clim_{vp} \times P^{sum}_{vpt} + \eta^{clim}_{vp} \times P^{g}_{ivp,t-2} + \eta^{clim}_{vp} \times \gamma_t + \eta^{ind}_{vp} \times \gamma_t + \epsilon_{ivpt} \end{aligned} \tag{8}$$

If the diversified households are less likely to be constrained, β_2 should be of opposite sign to β_1 for measures of investment in agriculture and non-agricultural production. The results of estimating equation (8) employing a dummy variable for already diversified households are reported in Panel A of Table 4. Though the results are somewhat noisy, the postulated pattern is observed for agricultural investment, total income and migration; the interaction terms for non-agricultural investment and income are of the expected sign, but insignificant. The final row of the panel reports the linear combination $\beta_1 + \beta_2$, and we can observe that there is almost no evidence of significant effects of the quota price shock on patterns of diversification for households that are already diversified ex ante.

In Panel B of Table 4, parallel results are reported employing a dummy variable for households in the highest quantile of income at baseline; the quantile is calculated employing the average income observed over the years in which a household appears in the panel prior to 1993.[29] Here, the interaction terms for agricultural variables are

[28] To increase precision, an additional 2% is trimmed from the top and bottom of the distribution of the climatic index times price instrument.

[29] The specification also includes the dummy variable and a linear control for average income at baseline.

counterintuitive in sign, but as expected there is no evidence of an increase in non-agricultural investment, income or migration for higher-income households.

Taken together, these results suggest that the observed heterogeneity in the primary effects is generally consistent with the hypothesis that the main channel for the observed effects is a relaxation of credit constraints. The on-line appendix also briefly explores the potential relationship between the positive income shock and risk tolerance, as well as the hypothesis that the shock stimulates exit from agriculture by households that are differentially productive in non-agricultural production. There is no evidence that either of these channels are relevant.[30]

4.3 Placebo tests

The fundamental identifying assumption of the main analysis requires that there is no unobserved variable correlated with fluctuations in the quota price that also has a disparate impact across areas with different climatic conditions. A useful test of this assumption is to evaluate whether trends in major economic outcomes are parallel between those areas in a period without major changes in the quota price.

Between 1986 and 1993, the quota price, again defined as the unexplained residual in a regression of market on quota price, showed no major fluctuations; this is evident in Figure 1b. Accordingly, it is feasible to evaluate whether parallel trends are observed across areas with different climatic conditions in this period, using a more limited set of outcomes reported in these earlier surveys.[31]

[30] These results are reported in Section S5 of the on-line appendix.

[31] In this analysis, agricultural investment is constructed using principal component analysis on a restricted set of variables: sown area, agricultural labor, and the value of investment in animals and tools. Agricultural income is the value of grain and pork production. Non-agricultural investment is constructed using principal component analysis on a restricted set of variables: a dummy variable equal to one if the household reports any new cash investment in non-agricultural assets, a dummy variable equal to one if the household reports any labor invested in non-agricultural businesses, and the amount of labor reported invested in non-agricultural businesses. Non-agricultural income is income from non-agricultural household businesses and wage income. Outside labor is not reported in this period; the other variables of interest are defined identically to the primary analysis. Given that only agricultural income from a restricted set of sources (grain and pork production) is reported, total income is not calculated. In addition, given that non-grain consumption is not reported in this period, consumption results are omitted.

Two sources of evidence are presented on this point. First, the primary specification, equation (7), is re-estimated using the earlier data. Referring to the years in which the dependent variables are observed, the primary analysis examines the years 1995 to 2002, and the placebo analysis examines the years 1986 to 1993; again, quota income is measured in the year prior to the dependent variables of interest. The results can be found in Panel C of Table 4. It is evident that there is no significant first stage in the pre-period, and the second-stage results are likewise small in magnitude and generally statistically insignificant; in many cases, the coefficients are opposite in sign relative to the coefficients for the main results reported in Panel C of Table 3.

Even in the absence of any significant correlation with the instrument, however, there may be diverging trends over time in areas that are and are not favorable for rice cultivation. In order to test this hypothesis, trends over time between 1986 and 1991 in the economic outcomes of interest are evaluated, comparing villages where the average propensity to cultivate rice is above or below the median.[32] Given that data from 1992 and 1994 are missing, 1991 is employed as the endpoint in order to examine a continuous period. The graphs are shown in Figure 2; the graphs capture trends over time in the variables of interest standardized to have mean zero and standard deviation one. There is no evidence of diverging trends that would be a systematic source of bias.

In addition, further robustness checks are reported in Section S6 of the on-line appendix, demonstrating that the quota price is not correlated with any other cross-sectional policy variation that could be a source of bias. In addition, tests are conducted to explore potential endogenous determination of the quota quantity and the quota price, and the results suggest that lagged values of covariates do not predict cross-sectional variation in either quota quantity or price.

[32]The median is defined using the average of the climatic index and the FAO index employed in the main analysis.

5 Conclusion

This paper analyzes the evolution of an unusual institution in rural China, the grain quota system, in order to estimate the impact of a gradual positive income shock on household economic behavior. This system effectively imposed a lump-sum tax on rural households that declined in magnitude over time as the quota price increased, and also varied in magnitude for counties producing different crops. The identification strategy exploits cross-sectional variation in the climatic suitability of different areas for rice cultivation in conjunction with time variation in the quota price to generate a source of quasi-exogenous variation in quota income.

The results indicate that the effect of a positive income shock on investment in agriculture is negative, while there are positive effects on investment in non-agricultural household businesses, migration, borrowing, and consumption of non-staple goods. These effects are concentrated among households that are not diversified into non-agricultural production at baseline, as well as relatively low-income households.

The empirical evidence is consistent with the hypothesis that households may be credit-constrained ex ante, and face required initial investments in non-agricultural production that are lumpy. Accordingly, an increase in agricultural income that is particularly large for low-income households allows households to invest in non-agricultural activities for the first time. The implications of this shift for the long-run welfare of Chinese rural households remains an interesting topic for future exploration.

References

Barrett, Christopher, Mesfin Bezuneh, and Abdillahi Aboud, "Income diversification, poverty traps and policy shocks in Cote d'Ivoire and Kenya," *Food Policy*, 2001, *26*, 367–384.

Blattman, Christopher, Nathan Fiala, and Sebastian Martinez, "Generating skilled self-employment in developing countries: Experimental evidence from Uganda," *Quarterly Journal of Economics*, 2014, (2), 697–752.

Brandt, Loren and Carsten Holz, "Spatial price differences in China: Estimates and implications," *Economic Development and Cultural Change*, 2006, *55* (1), 43–86.

_ , Jikun Huang, Guo Li, and Scott Rozelle, "Land rights in rural China: Facts, fictions and issues," *The China Journal*, 2002, pp. 67–97.

_ , Scott Rozelle, and Matthew Turner, "Local government behavior and property rights formation in rural China," *Journal of Institutional and Theoretical Economics*, 2004, *160* (4), 627–662.

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli, "Agricultural productivity and structural transformation: Evidence from Brazil," 2015.

Carter, Michael R. and Travis J. Lybbert, "Consumption versus asset smoothing: testing the implications of poverty trap theory in Burkina Faso," *Journal of Development Economics*, 2002, *99*, 255–264.

_ , Peter D. Little, Tewodaj Mogues, and Workneh Negatu, "Poverty traps and natural disasters in Ethiopia and Honduras," *World Development*, 2007, *35*, 835–856.

Foster, Andrew and Mark Rosenzweig, "Agricultural productivity growth, rural economic diversity, and economic reforms: India, 1970–2000," *Economic Development and Cultural Change*, 2004, *52* (3), 509–542.

Gertler, Paul, Sebastian Martinez, and Marta Rubio-Codina, "Investing cash transfers to raise long-term living standards," *American Economic Journal: Applied Economics*, 2012, *4* (1), 164–192.

Giesbert, Lena and Kati Schindler, "Assets, shocks, and poverty traps in Rural Mozambique," *World Development*, 2012, *40*, 1594–1609.

Gilligan, Daniel, John Hoddinott, and Seyoum Taffesse, "The impact of Ethiopia's productive safety net programme and its linkages," *Journal of Development Studies*, 2008, *45* (10), 1684–1706.

Gollin, Douglas, David Lagakos, and Michael Waugh, "The agricultural productivity gap," *Quarterly Journal of Economics*, 2014, *129* (2), 939993.

Hoddinott, John, "Shocks and their consequences across and within households in rural Zimbabwe," *The Journal of Development Studies*, 2006, *42*, 301–321.

Huang, Jikuan, Scott Rozelle, and Honglin Wang, "Fostering or stripping rural China: Modernizing agriculture and rural to urban capital flows," *The Developing Economies*, 2006, *44* (1), 1–26.

Huang, Jikun and Scott Rozelle, "The nature of distortions to agricultural incentives in China and implications of WTO accession," 2002.

_ , _ , Ruifa Ha, and Ninghui Li, "China's rice economy and policy: Supply, demand and trade in the 21st century," in M. Sombilla, M. Hossain, and B. Hardy, eds., *Developments in the Asian Rice Economy*, Manila: International Rice Research Institute, 2002.

_ , Yunhua Wu, and Scott Rozelle, "Moving off the farm and intensifying agricultural production in Shandong: A case study of rural labor market linkages in China," *Agricultural Economics*, 2009, pp. 203–18.

Huang, Yipin, *Agricultural reform in China: Getting institutions right*, Cambridge University Press, 1998.

Janzen, Sarah A. and Michael R. Carter, "After the drought: The impact of microinsurance on consumption smoothing and asset orotection," 2013. NBER Working Paper No. 19702.

Jedwab, Remi, "Cash crop windfalls and development: On cocoa and cities in Ivory Coast and Ghana," 2011.

Jensen, Robert and Nolan Miller, "Giffen behavior and subsistence consumption," *American Economic Review*, Sep. 2008, *98* (1), 1553–1577.

Kaboski, Joseph and Robert Townsend, "A structural evaluation of a large-scale quasi-experimental microfinance initiative," *Econometrica*, 2011, *79* (5), 1357–1406.

Kraay, Aart and David McKenzie, "Do poverty traps exist? Assessing the evidence," *The Journal of Economic Perspectives*, 2014, *28*, World Development.

Kwak, Sungil and Stephen C. Smith, "Regional agricultural endowments and shifts of poverty trap equilibria: Evidence from Ethiopian panel data," *The Journal of Development Studies*, 2013, *49* (7), 955–975.

Lagakos, David and Michael Waugh, "Selection, agriculture, and cross-country productivity differences," *American Economic Review*, 2013, *103* (2), 948–80.

Lin, Justin Yifu, "The household responsibility system reform and the adoption of hybrid rice in China," *Journal of Development Economics*, 1991, *36*, 353–372.

_ , "Rural reforms and agricultural growth in China," *American Economic Review*, 1992, *82* (1), 34–51.

Lu, Feng, "Three grain surpluses: Evolution of China's grain price and marketing policies," 1999.

Naschold, Felix, "The poor stay poor: Household asset poverty traps in rural semi-arid India," *World Development*, 2012, *40*, World Development.

_ , "Welfare dynamics in Pakistan and Ethiopia Does the estimation method matter?," *The Journal of Development Studies*, 2013, *49* (7), 936–954.

Quisumbing, Agnes R. and Bob Baulch, "Assets and poverty traps in rural Bangladesh," *The Journal of Development Studies*, 2013, *49* (7), 898–916.

Rozelle, Scott, Albert Park, Jikun Huang, and Hehui Jin, "Bureaucrat to entrepreneur: The changing role of the state in China's grain economy," *Economic Development and Cultural Change*, 2000, *48* (2), 227–252.

Sadoulet, Elizabeth, Alain de Janvry, and Benjamin Davis, "Cash transfer programs with income multipliers: PROCAMPO in Mexico," *World Development*, 2001, *29* (6), 1043–1046.

Shao, Yun, Xiangtao Fan, Hao Liu, Jianhua Xiao, S. Ross, B. Brisco, R. Brown, and G. Staples, "Rice monitoring and production estimation using multitemporal RADARSAT," *Remote Sensing of Environment*, 2001, *76*, 310–25.

Sicular, Terry, "Redefining state, plan and market: China's reforms in agricultural commerce," *The China Quarterly*, 1995, *144*, 1020–1046.

Tang, Shengxiang, Li Ding, and Alain Bonjean, "Rice production and genetic improvmeent in China," in Zhonghu He and Alain Bonjean, eds., *Cereals in China*, D.F., Mexico: CIMMYT, 2010, pp. 15–34.

Wang, Dewen, Jikun Huang, Daniel A. Sumner, and Scott D. Rozelle, "Quotas and grain production in China," in Scott Rozelle and Daniel A. Sumner, eds., *Agricultural trade and policy in China*, Burlington, VT: Ashgate, 2003, pp. 111–137.

6 Figures and Tables

Figure 1: Descriptive statistics: Grain quota

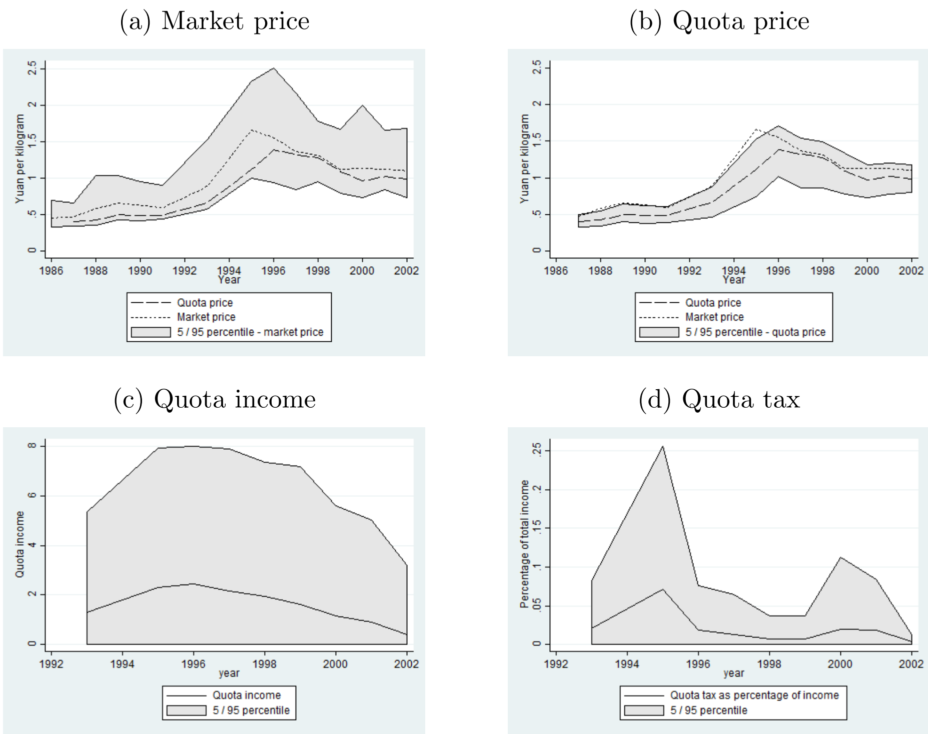

Notes: Figures 1a and 1b show the market price and the quota price of grain, measured in yuan per kilogram, between 1986 and 2002; Figure 1a also shows the 5th and 95th percentile of the within-year, cross-village distribution of the market price, and Figure 1b shows the same percentiles for the within-year, cross-village distribution of the quota price. Figure 1c shows mean household-level quota income by year in hundreds of yuan, as well as the 5th and 95th percentile across households by year. Figure 1d shows the mean implicit tax posed by the quota system by year, calculated as the difference between the market and the quota price multiplied by the quota quantity, as a percentage of total income; again, the 5th and 95th percentiles across households by year are also included.

Source: Data from the RCRE household panel analyzed by the author.

Figure 2: Pre-trends prior to quota price increase

(a) Agri. investment

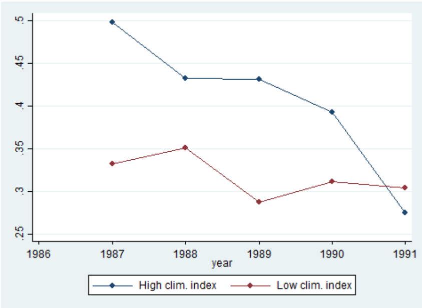

(b) Agri. income

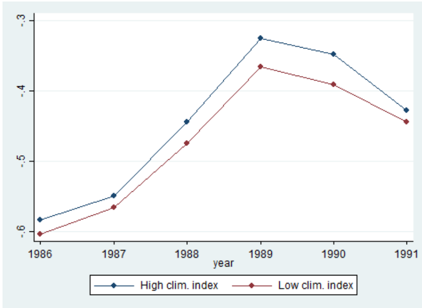

(c) Non-agricultural inv.

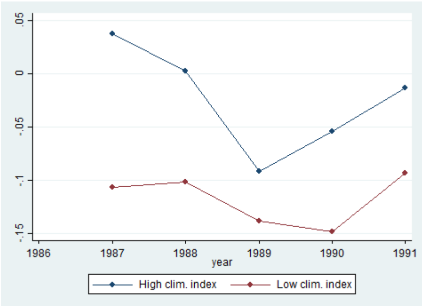

(d) Non-agricultural inc.

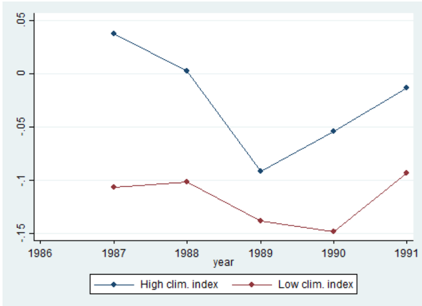

(e) Migration

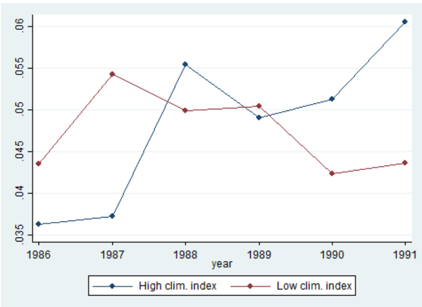

(f) Borrowing

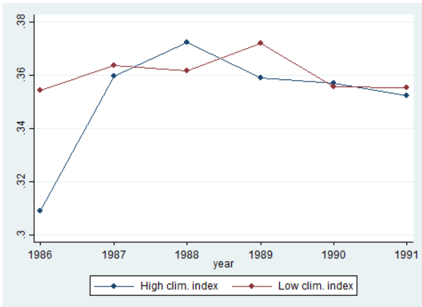

Notes: These figures show trends over time in the primary economic outcomes of interest for villages above and below the median of the climatic index capturing suitability for rice cultivation. In this analysis, agricultural investment is constructed using principal component analysis on a restricted set of variables: sown area, agricultural labor, and the value of investment in animals and tools. Agricultural income is the value of grain and pork production. Non-agricultural investment is constructed using principal component analysis on a restricted set of variables: a dummy variable equal to one if the household reports any new cash investment in non-agricultural assets, a dummy variable equal to one if the household reports any labor invested in non-agricultural businesses, and the amount of labor reported invested in non-agricultural businesses. Non-agricultural income is income from non-agricultural household businesses and wage income. The migration dummy is a dummy variable for any labor worked as a migrant, and borrowing is a dummy variable equal to one if the household reports any access to credit. All variables are standardized to have mean zero and standard deviation one.

Source: Data from the RCRE household panel analyzed by the author.

Table 1: Summary statistics

Variable	Mean	Standard deviation	Obs.
Household size	4.093	1.31	15657
Sown area	1.403	.92	16070
Grain sown area	1.233	.816	16115
Productive assets	.909	.287	16203
Non agri. assets	.263	.44	16203
Outside labor	.388	.487	16203

Notes: Summary statistics are reported for variables of interest in the primary analysis sample, including only households in homogeneous cultivation villages. Productive assets is a dummy equal to one if the household reports any agricultural assets (tools, animals or machinery). Non-agricultural assets is a dummy equal to one if the household reports any non-agricultural assets. Outside labor is a dummy equal to one if the household reports any days worked outside the household, and associated wage income.

Source: Data from the RCRE household panel analyzed by the author.

Table 2: First stage

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
		Quota quan.		Quota income		Quota quan.		Quota income	
Clim. index	142.936 (45.623)***								
FAO index		70.833 (25.530)***							
Rice area			91.357 (17.026)***						
Quota price						-6.376 (4.774)	17.917 (7.791)**		
Clim. index x price				.314 (.071)***				.190 (.049)***	
FAO index x price					.321 (.092)***				.213 (.085)**
Mean dep. var.	195.22	195.22	195.22	2.065	2.065	204.813	224.527	2.225	2.225
Obs.	23558	23558	23511	27779	29009	16174	11539	16203	16203
Fixed effects		Prov.-year		Village + prov.-year		Village		Village + prov.-year	
Sample selected		Full		Full		Hom.	Het.	Hom.	

Notes: The dependent variable is quota quantity in Columns (1) to (3) and Columns (6) to (7), and quota income in Columns (4) and (5) and Columns (8) and (9). The independent variables are the climatic indices for propensity to cultivate rice derived from weather data and from FAO data; the area cultivated in rice; the quota price; and both climatic indices interacted with the quota price. Fixed effects and clustering are as reported in the table. “Hom.” denotes villages that are homogeneous in rice and wheat cultivation, while “Het.” denotes villages that are heterogeneous. Asterisks indicate significance at the ten, five, and one percent level, respectively.

Source: Data from the RCRE household panel analyzed by the author.

Table 3: Two-stage least squares results

Panel A: Agricultural and non-agricultural investment and income

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Agri. input	Agri inc.	Non-agri. inv.	Non-agri. inc.	Tot. inc.	Outside labor	Migration
Quota income (lagged)	-.205 (.074)***	-.123 (.071)*	.197 (.112)*	.066 (.058)	-.062 (.043)	.053 (.057)	.175 (.089)**
Hansen J-statistic	.239	.056	.017	.935	.083	.007	.087
Obs.	16203	16203	16203	16203	16203	16203	16203

Panel B: Credit and consumption

	Credit dummy	Amount borrowed	Grain cons.	Other cons.	Total cons.
Quota income (lagged)	.099 (.049)**	.317 (.188)*	-.067 (.084)	.136 (.050)***	.119 (.044)***
Hansen J-statistic	.175	.002	.979	.263	.204
Obs.	16203	16203	16201	16177	16201

Panel C: Agriculture investment

	Sown area	Labor	Fertilizer value	Seeds value	Animal inv.	Tools inv.
Quota income (lagged)	-.114 (.170)	-9.714 (2.034)***	-27.029 (52.041)	-25.946 (6.866)***	25.717 (96.532)	-16.160 (14.318)
Mean dep. var.	1.59	188.42	570.88	87.630	-41.43	7.39
Hansen J statistic	.247	.322	.609	1.319	3.125	.801
Obs.	16202	16177	16076	16203	15511	15941

Panel D: Non-agricultural investment

	Non-agri. labor	Labor dummy	Inv.	Inv. dummy	Days worked outside hh
Quota income (lagged)	17.067 (12.612)	.080 (.045)*	2.412 (2.479)	.098 (.044)**	-9.032 (12.514)
Mean dep. var.	75.79	.41	2.98	.06	133.60
Hansen J-statistic	.001	.755	.374	.448	1.034
Obs.	16203	16203	16201	16177	16104

Notes: All specifications include village and province-year fixed effects, the interaction of the climatic index and a summary measure of market prices, climatic index quantile fixed effects interacted with the two year lagged market price, climatic quantile index fixed effects interacted with year fixed effects, and industrial quantile fixed effects interacted with year fixed effects. Standard errors are estimated employing two-way clustering at the province and year level. The independent variable is quota income, lagged, instrumented by the lagged interactions of two climatic indices of the propensity to cultivate rice (derived from weather data and FAO data), interacted with the quota price. Asterisks indicate significance at the ten, five, and one percent level, respectively.

In Panel A, the dependent variable in Column (1) is a summary variable of agricultural investment, calculating using principal component analysis on six variables (area sown, agricultural labor, value of fertilizer employed, value of seeds employed, investment in animals, and investment in tools). The dependent variable in Column (2) is agricultural income. The dependent variable in Column (3) is a summary variable of non-agricultural investment, calculated using principal component analysis on four variables (a dummy variable for any new cash investment in non-agricultural machinery, a dummy variable for any labor invested in a non-agricultural business, and the amount of labor and cash investment reported in non-agricultural businesses). The dependent variable in Column (4) is income from non-agricultural household businesses and wage labor, and the dependent variable in Column (5) is total income. The dependent variable in Column (6) is a dummy variable for any outside labor; the dependent variable in Column (7) is a dummy variable for any labor worked as a migrant. In Panel B, the dependent variable in Column (1) is a dummy variable if the household reports any access to credit, and the dependent variable in Column (2) is the amount of credit. The dependent variables in Columns (3) through (5) are consumption in yuan on non-staple and staple items, and total consumption. All non-dummy dependent variables are standardized to have mean zero and standard deviation one.

The dependent variables in Panel C include sown area in hectares; days reported worked in agriculture; value of fertilizer and seeds reported used, in yuan; and new investment in tools and animals, reported in yuan. Expenditure reported in yuan is deflated using deflators for agricultural inputs provided by the National Bureau of Statistics. The dependent variables in Panel D include days worked in non-agricultural businesses; a dummy variable for working in a non-agricultural business; cash investment in a non-agricultural business; a dummy for positive cash investment; and days worked outside the household. Cash investment is deflated using indices for factory output compiled by the NBS.

Source: Data from the RCRE household panel analyzed by the author.

Table 4: Channels and robustness checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Agri. input	Agri inc.	Non-agri. inv.	Non-agri. inc.	Tot. inc.	Outside labor	Migration
Panel A: Heterogeneous effects for previously unconstrained households							
Clim. index x price (lagged)	-.008 (.002)***	-.005 (.005)	.004 (.003)	.003 (.001)*	-.002 (.002)	.003 (.002)	.006 (.004)*
Asset ownership int.	.004 (.002)***	-.002 (.001)	-.004 (.005)	-.001 (.001)	-.002 (.0006)***	.002 (.001)	-.002 (.0006)***
$\beta_1 + \beta_2$	-.004 (.003)	-.007 (.005)	.000 (.005)	.002 (.001)	-.004 (.002)*	.005 (.003)	.004 (.004)
Obs.	15118	15118	15118	15118	15118	14672	10791
Panel B: Heterogeneous effects for high-income households							
Clim. index x price (lagged)	-.007 (.001)***	-.004 (.006)	.004 (.003)	.003 (.001)*	-.001 (.003)	.003 (.002)	.006 (.003)*
Top quantile int.	-.004 (.001)***	-.004 (.003)	-.003 (.001)**	-.001 (.0004)**	-.003 (.000)	-.0004 (.001)	-.001 (.0008)*
$\beta_1 + \beta_2$	-.010 (.002)***	-.008 (.006)	.001 (.003)	.002 (.001)	-.004 (.002)**	.002 (.003)	.004 (.003)
Obs.	15118	15118	15118	15118	15118	14611	13290
Panel C: First stage and two-stage least squares, 1987–1993							
	Quota inc.	Agri. input	Agri inc.	Non agri. inv.	Non agri. inc.	Migration	Borrowing
Climatic index x price	.145 (.127)						
Quota income (lagged)		-.085 (.134)	-.125 (.078)	-.073 (.106)	-.024 (.063)	.009 (.027)	.036 (.109)
Obs.	16749	16106	16749	16748	16749	16749	16749

Notes: All specifications include village and province-year fixed effects, the interaction of the climatic index and a summary measure of market prices, climatic index quantile fixed effects interacted with the two year lagged market price, climatic quantile index fixed effects interacted with year fixed effects, and industrial quantile fixed effects interacted with year fixed effects. Standard errors are estimated employing two-way clustering at the province and year level. Standard errors are estimated employing two-way clustering at the province and year level. Asterisks indicate significance at the ten, five, and one percent level, respectively. The dependent variables are defined in the notes to Table 3.

In Panels A and B, the independent variables include the climatic index-price interaction (employing the weather-derived climatic index), and the triple interaction including a dummy variable for households already owning non-agricultural assets (in Panel A), or for a household identified as in the top income quantile pre-1993 (in Panel B). In Panels A and B, there are some cells defined by village, province-year and the dummy variable on the right hand side in which there is no variation in outside labor or migration; these cells are dropped, accounting for the lower number of observations. Additional controls include the non-agricultural assets dummy entering linearly in Panel A; and the high income dummy entering linearly in Panel B.

In Panel C, the first column shows the first stage, where the dependent variable is quota income; Columns (2) through (7) report the two-stage least squares specification for the subset of outcomes reported in the pre-1993 data.

Source: Data from the RCRE household panel analyzed by the author.