Decentralization, Rural Development, and Growth: Evidence from India*

Veda Narasimhan
University of Zurich

November 13, 2021

CLICK HERE FOR LATEST VERSION

Abstract

Over the last few decades, many countries have undertaken a process of decentralization. The effects of comprehensive decentralization on village-level development are ex-ante ambiguous. In addition, there is little evidence on how this rural development influences aggregate outcomes. A 1991 federal reform in India resulted in the creation of local government units, with broad political, administrative, and financial powers. This paper uses a regression discontinuity design to study the effects of this reform on village development and district-level outcomes. It leverages population thresholds that the state of Uttar Pradesh used to determine which villages would receive their own village council. After the reform, treated villages have higher levels of amenities —education and roads—and durable consumption. Additionally, they experience patterns of structural change coupled with higher population growth, which is likely driven by treated villages becoming more attractive locations. At the aggregate level, districts with a higher share of treated villages have higher rural-urban migration, consistent with a reduction in educational and financial constraints at the individual-level, and higher GDP per capita. The results suggest that decentralization not only fosters economic development at the village level, but also has positive effects on urbanization and growth.

*E-mail: veda.narasimhan@econ.uzh.ch. I would like to thank my advisors, Lorenzo Casaburi, David Yanagizawa-Drott, and Joachim Voth. I would also like to thank Guilherme Lichand, Bruno Caprettini, Maria Alsina Pujols, Sili Zhang, Xiaoyue Shan, Andrea Hofer, Luca Bagnato, and all seminar participants at the University of Zurich for their helpful comments, suggestions, and feedback. I am also grateful to the following organizations for their grant support: the Swiss National Science Foundation through grant 187942, and the Forschungskredit of the University of Zurich, grant no. FK-19-017.

1 Introduction

Classic models of development highlight that population movements from rural areas to urban areas help to drive economic growth (e.g. Lewis et al., 1954; Harris and Todaro, 1970). Evidence from developed countries suggests that economies transition out of agriculture, converting rural populations away from subsistence agriculture towards manufacturing and services in urban areas. Low and middle-income countries today are in the early to middle stages of this transition. However, a significant portion of their development policies are targeted specifically towards rural development.

Starting in the 1980s, many countries used decentralization reforms to encourage rural development, devolving power to smaller political units (Gadenne and Singhal, 2014; Mookherjee, 2015). However, there is little evidence on how effective these programs were at driving village-level development, especially in rural areas. Improvements in the standards of living may make these rural areas more attractive locations for individuals to settle. This brings up a potential tradeoff, as the focus on rural development may detract from rural-urban migration, and, in turn, aggregate growth and development (Lagakos et al., 2018).

In this paper, I use a large reform decentralizing power to villages in India to causally identify the effect of decentralization on village-level development outcomes and to then measure its aggregate consequences. A regression discontinuity design shows that villages with a higher degree of decentralization have higher levels of public goods and durable consumption, and show signs of structural change away from agricultural activity. Additionally, more decentralized villages also experience higher population growth, as the village-level development improvements make these more attractive locations to settle. I then test whether these changes are driven by substitution away from rural-urban migration at the district-level. Surprisingly, decentralization intensity at the district level increases rural-urban migration: the changes in rural population appear to be driven by a reshuffling between rural areas, rather than substitution away from urban migration. In addition, decentralization intensity at the district-level has a positive impact on district GDP. The results suggest that decentralization not only fosters economic development at the village level, but also has positive effects on urbanization and growth.

My setting is a federal decentralization reform from the nineties in India, that devolved power to local political institutions and created a three-tier political system with the village council as its smallest elected unit. In 1991 the government passed the 73^{rd} Amendment to the Indian constitution. This Amendment established the legal and political existence of village-level councils (also known as Gram Panchayats, or GPs), and brought them into the democratic process whilst increasing their financial, administrative, and political independence. From introducing elections at the local government level, to devolving responsibility of multiple aspects of governance, such as taxation, education, and local economic development, this system resulted in a significant shift in the potential of these local government units.

For identification purposes, I focus on the state of Uttar Pradesh. Here, population thresholds were used to determine eligibility for a GP. Villages below the threshold ended up sharing a government unit with other villages, whilst villages above the threshold were more likely to receive their own local council. This setting allows me to measure the effect of different intensities of decentralization. Proximity to administrative centers can affect development outcomes, which means that control villages already have less access to the local institutions due to sharing the unit across multiple villages. In addition, shared councils have a greater number of preferences and villages to cover, which may reduce efficiency and local knowledge of people's needs.

I utilize a Regression Discontinuity Design (RDD) to causally identify village-level outcomes. More specifically, I use a fuzzy RDD, where the cutoff dummy is used as an instrument to predict the likelihood that a village receives its own council. An additional benefit of this empirical strategy is that I can then aggregate my RDD up to the district-level and use it to identify some of the causal impacts on the aggregate level, by generating a measure of the share of treated villages within the RDD population bandwidth. By using the RDD structure I can control for some of the endogeneity concerns when analyzing the district-level consequences of the rural development program. I test my population variable to ensure there are no discontinuous jumps at the thresholds, and that the cutoff is a strong predictor for treatment status. Additionally, I use optimal RDD bandwidths for all my results but check my outcomes across multiple bandwidths to ensure robustness of my results.

At the village-level, I focus on development outcomes associated with location choice - public goods, economic activity, and durable consumption. I find that relatively more decentralized villages have better provision of public goods, higher levels of durable consumption, and show evidence of movement away from agriculture. The higher population growth I observe at the village-level is consistent with improvements and changes in these development outcomes making treated villages more attractive locations to settle. Importantly, the public good effects that I find at the village-level, kick in prior to the population changes, which suggests that they are a driver rather than a consequence. Furthermore, exploiting the spatial distribution of villages, I also show that a higher share of treated villages around a village negatively impacts own village population growth, which is in line with the findings above.

The second half of my paper then looks at how the village-level effects impact aggregate outcomes. The changes in the rural population at the village-level could be driven by two potential channels. First, that individuals may be substituting away from rural-urban migration, which could have negative consequences on aggregate growth. Alternatively, it could be driven by a re-optimization of rural populations from control villages to treated ones. I test both these channels with district-level migration data using my aggregated RDD strategy, and find no evidence of substitution away from urban migration. This result, taken with my finding of increased educational provision as a public good and increased durable consumption at the village-level, is consistent with the idea that rural development may help to reduce migration frictions associated with urban migration. Combined with the findings of a positive effect on district-level GDP, these results suggest that rural development has indirect benefits to aggregate growth and migration that may not be captured when directly measuring development outcomes.

My paper bridges two strands of literature that have not been traditionally studied together. First, I contribute empirically to the research that has been conducted on the impact of decentralization (e.g. Faguet, 1999; Galasso and Ravallion, 2005; Zhang and Gong, 2005; Dahis and Szerman, 2020).[1] Additionally, work by Bardhan and Mookherjee (2006) provides an extensive review of advantages and disadvantages of decentralization, suggesting that there may be negative consequences policymakers do not take into account.[2] I contribute to this literature by showing empirical evidence at the village-level of the impacts of a comprehensive decentralization reform covering multiple aspects of decentralization.[3] My findings also provide micro-level support to the literature on democratization and growth (see e.g. Barro, 1996; Glaeser et al., 2004; Foster and Rosenzweig, 2004; Acemoglu et al., 2019).

The second main strand I contribute to is the literature on the dynamics of rural and urban development, including rural-urban migration. Seminal papers in development economics (Lewis et al., 1954; Harris and Todaro, 1970) highlight the importance of rural-urban migration and the transition from agriculture to manufacturing and services. Lagakos (2020) presents a recent overview of the determinants of rural-urban migration, but does not discuss political institutions, which is the focus of my paper. Additionally, a growing body of empirical work has delved into the barriers and impact of rural-urban migration (Bryan et al., 2014; Munshi and Rosenzweig, 2016; Bryan and Morten, 2019; Morten, 2019).[4] My results that village-level decentralization may increase rural-urban migration by improving education and relaxing financial constraints in rural areas are consistent with the findings of this literature.

This paper is structured as follows: section 2 gives a background of the reform, and discuss the conceptual framework, while section 3 describes the data sources used in this paper. Section 4 then lays out the empirical strategy in detail followed by section 5 which presents the village-level analysis on rural development. Finally, section 6 looks at the district-level impacts on migration and growth, and section 7 concludes.

[1] Decentralization has also been the subject of extensive theoretical work beginning with Oates et al. (1972), for more recent reviews on the literature, see Gadenne and Singhal (2014) and Mookherjee (2015).

[2] For a detailed discussion on corruption and accountability in decentralization, see Mookherjee et al. (2005).

[3] Several prominent papers have looked into the reservation quota aspect of the same reform I study: Besley et al. (2004), Chattopadhyay and Duflo (2004a), Chattopadhyay and Duflo (2004b), and Besley et al. (2007). These papers do not exploit the variation in decentralization intensity that I use in this paper.

[4] More generally, a large literature looks at the spatial misallocation of labor as a barrier to growth. See Gollin et al. (2014), Munshi and Rosenzweig (2016),Hsieh and Moretti (2019), and Donovan and Schoellman (2021)

2 Setting, Decentralization, and Population Movement

2.1 Background of the Reform

Throughout the nineties, India engaged in a series of extensive reforms, both political and economic, intended to bring about aggregate growth and reduce poverty. One of the major reforms was the devolution of power to village-level units as a means to better meet the needs of the large rural population that existed within the country. Many developing countries were engaging in decentralization programs at this time, with countries such as Indonesia and China conducting their own decentralization within this time period. The Indian decentralization differed from other programs in a few key dimensions however. First, the reform was at a much more granular level than the majority of other reforms happening around this period, and second, it was a comprehensive program that covered a large range of different responsibilities.

Rural village councils in India have existed since the pre-modern era. After independence, India introduced the first official legislation concerning the administration and existence of village councils - known as gram panchayats (GPs). Throughout the post-independence period there were discussions about the existence and purpose of these councils, but they were given relatively limited responsibilities, with generally no political oversight or elections. During the eighties, there was a growing movement in the central government to utilize these village councils more effectively and efficiently reach the large rural population that existed in the country.

Starting with a series of federal commissions, multiple recommendations were made to the Central government regarding the potential of these village councils. These recommendations culminated in the introduction of the 73^{rd} Amendment to the Indian Constitution. This reform was a comprehensive decentralization reform covering a multitude of different policies and purposes. There were however three key aspects of the reform that defined its purpose. First, the reform resulted in the formation of a three-tier system of governance at the state-level - district, sub-district, and village. The second key component was that there was a devolution of power and responsibility to the village-level across a set of different areas. This areas covered by the reform were covered by a formal list known as the 11^{th} schedule which is discussed in detail below. The final key component of the reform was the introduction of political elections at the village-level for seats on the villages councils. In addition to the introduction of elections, the federal government also introduced the reservation of seats within councils for women and individuals from backward castes.

By formally amending the constitution with this policy, the federal government ensured that all states would uniformly implement the key aspects of the reform. The federal government did however allow some flexibility in the decision on which exact areas would be decentralized and how the councils would be formed. These decisions were left up to the discretion of the individual states, and each state set up their own Acts detailing how the reform would be implemented in that state.

Target areas of the reform There are 29 different subjects covered under this the 11^{th} Schedule. The schedule was produced in order to give a list of potential areas of devolution to the state bodies. Since the federal government left the decision of which topics to decentralize to the discretion of the individual states, the schedule acted as a guide for which areas states could focus on. Table A.1 shows the list of potential activities that could be devolved to village councils. The topics range from economic development (agriculture, irrigation, technical training, cottage industries etc.) to social security (welfare programs, public distribution systems).

The GPs were also used as distribution sources. Federal programs often used GPs as the point of contact and administration in order to increase efficiency and targeting. This included programs such as the National Rural Employment Guarantee Scheme (NREGS) and public distribution systems (more specifically, the determination of individuals who qualified for below the poverty line benefits). The GPs became a focal point for investing in rural development through a multitude of different sources. As can be seen from Figure 1 shows the structure of the funding and the sources of direction and information flows.

Figure 1: Political Structure of the Panchayat System

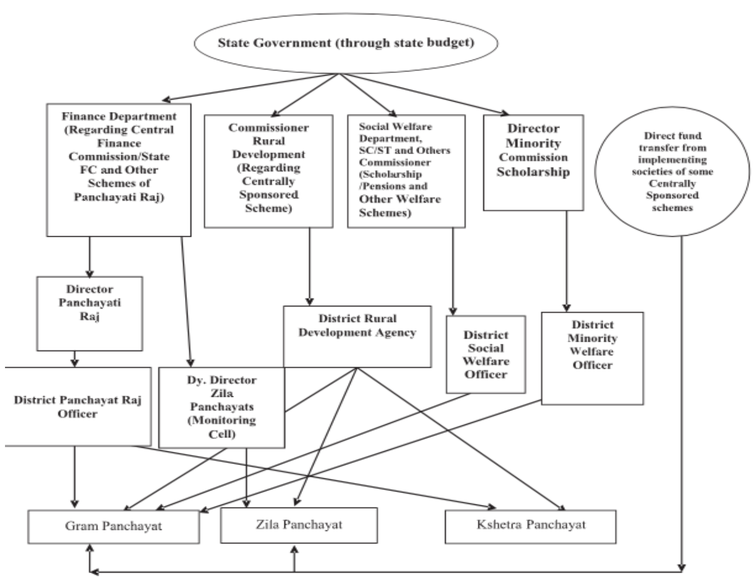

Source: NITI Aayog.

2.2 Implementation of the reform in Uttar Pradesh

The largest state in India by population, Uttar Pradesh (UP) often ranks near the bottom in terms of educations, health and economic outcomes. UP ranks 31^{st} out of 35 in terms of literacy, and as of the 2001 census, 79% of the population still lives in rural areas. It is also considered one of the less developed states in India, and the majority of the population are in the agrarian sector. Approximately 20.3% of the country's poor were found in Uttar Pradesh in 1993-1994, and the local governments were seen as potential means to help improve the standard of living in the state.

On the back of the federal constitutional amendment, by 1994 Uttar Pradesh had reviewed and amended their original Uttar Pradesh Panchayat Raj Act of 1947 to fall in line with the central government recommendation. They included the new list of devolved functions (see next section) as well as adjusting the village-council structure and many of the laws regarding the election and responsibilities of the panchayat. Uttar Pradesh made a concentrated effort to devolve significant functionalities to the village councils, giving them control over work related to education, drinking water, welfare and child development, medicine and health, agriculture, and village development.

Uttar Pradesh was one of the few states to highlight the use of population thresholds to determine which villages would be eligible for their own village council in the official state panchayat act, and additionally, effectively followed through with this policy, which is why I choose to focus on this state. In the official Panchayat Act, the state wished to implement village councils with populations of a thousand as far as it was feasible. I therefore use this population threshold to implement an RDD at the thousand person threshold. Given that the village council eligibility was not followed perfectly, I use the population cutoffs as an instrument for treatment, which I define as a village with their own village council.

(a) Proportion of treated villages

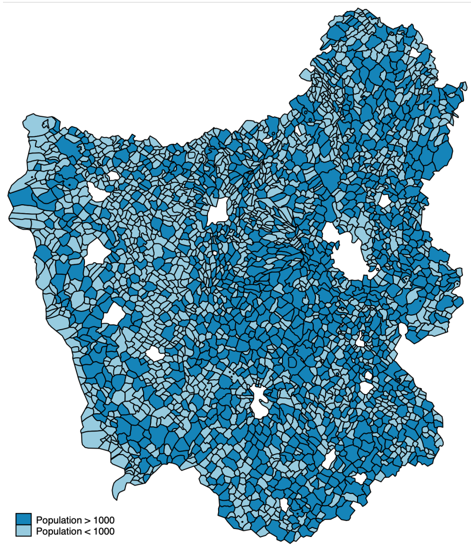

(b) District

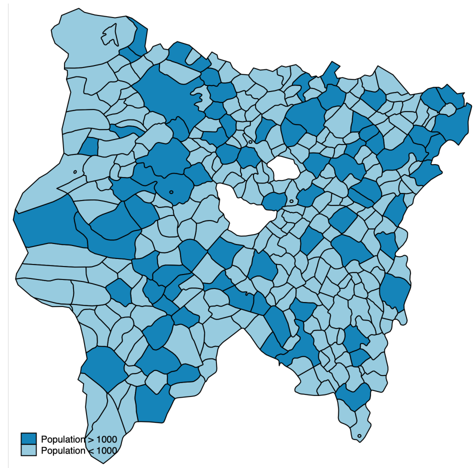

(c) Sub-district

Notes: This figure shows the distribution of villages above and below the 1000 threshold in (1) a single district (2) a single subdistrict. The boundaries are determined using the official census maps from 1991. For more information on the underlying spatial data, please see Meiyappan et al. (2018).

2.3 Conceptual Framework

Decentralization Decentralization can impact outcomes in a multitude of ways [5]. Work by Flèche (2021) has shown that centralization generates declines in welfare for individuals. Additionally, work by World Bank (2000) argues that decentralization allows local governments to take advantage of local knowledge. Alternatively, work by Bardhan and Mookherjee (2000) suggests the decentralization may increase changes of corruption and can cause local governments to lose out on economies of scale. Local units may not have the same degree of capabilities as higher echelons of government. Similarly for consumption, the reform intended to utilize local knowledge to increase the efficiency and reach of subsidy and social support programs. In addition, the reform targeted provision of infrastructure and encourage local economic activity, both of which could potentially improve consumption for individuals. Which effect dominates in this case becomes an empirical question that this paper answers.

Population Movement The exact effect of decentralization on local outcomes is an empirical question, but understanding village-level changes gives us important insights into how populations move. Models of location choice as seen in Bryan and Morten (2019), based on Eaton and Kortum (2002) and Allen and Arkolakis (2014) tell us that individuals value amenities and income opportunities. In order to capture some of these effects, I focus on de-

[5] For a detailed discussion. on this please see Gadenne and Singhal (2014)

velopmental outcomes associated with these indicators - public goods, durable consumption, and economic activity. Given improvements in these categories, these location choice models give us predictions on how we would expect individual's to move. Higher levels of development make these locations more attractive locations for individuals to move to, and given sufficient improvements, we may expect individuals to substitute away from urban migration if the benefits of remaining in rural areas is sufficiently high. In the second half of this paper I test this hypothesis looking at district-level data.

3 Data

I combine multiple datasets from a variety of sources in order to collect information on the universe of villages within Uttar Pradesh and the outcomes associated with them. The core datasets I use come from the Indian census information at the village-level, and I mix this with information on district-level outcomes, and spatial characteristics in order to build my foundational dataset.

Census This project uses three population census datasets - 1991,2001, and 2011. The 1991 census data provides population estimates, which are required for testing the empirical strategy this paper employs, as the 1991 census data is used for the baseline population cut-offs. I also use information from the 1991 census to measure baseline characteristics of these villages. In addition, I collect information on the The 2001 census data provides information on the employment status and village amenities at the village-level in 2001. I use this information to test the medium-term effects of the reform. Finally, the information from the 2011 census on village to gram panchayat allocations will be useful, as well as similar information as the 2001 census to look into long-term consequences. The core variables of interest from these datasets will be the education measures, healthcare variables, and the population data.

Housing Census I utilize housing census data from the 2011 Indian census. This dataset covers the entire universe of Indian villages, and collects information on housing quality, type and availability of housing amenities. This data unfortunately does not exist at the village-level for previous years, so I am not able to run detailed baseline tests. The dataset contains additional information on sanitary facilities, lighting sources and fuel sources.

SHRUG This paper uses the Asher et al. (2019) SHRUG database to connect all three rounds of the population census, as well as the three rounds of economic census. This data is merged with the above datasets to provide a consistent set of available outcomes to study.

By utilizing the set of consistent identifiers across all the censuses, I'm able to maximize the size of my dataset whilst ensuring consistency throughout.

Migration Census - 2011 The migration census from 2011 gives me a detailed breakdown of rural-rural versus rural-urban migration within-district for the full state of Uttar Pradesh. This migration data is an additional part of the 2011 Indian census, and is available at the district-level. I match the codes using the 2011 census codes, to include this information in my main data.

ICRISAT ICRISAT is the International Crops Research Institute for the Semi-Arid Tropics, and they collect detailed agricultural and socioeconomic information on Indian districts. From the ICRISAT dataset I get detailed information on district-level GDP, broken down by GDP per capita, and different sectors. This dataset also has extensive information on crop productivity for a multitude of different agricultural products.

SEDAC India Shapefiles For all the spatial analysis I conduct in this paper, I use the Columbia SEDAC database and their underlying shapefiles. This gives me the universe of all villages from 1991-2011, with all relevant census identifiers so I can link the files to the rest of my databases. More information on this database is available in Meiyappan et al. (2018).

4 Empirical Methodology

This paper uses a Regression Discontinuity Design (RDD) strategy to compare the causal effect of having your own village council to sharing one. I focus on the state of Uttar Pradesh, which set a population threshold in order to determine the distribution of villages across panchayats. Some panchayats consist of multiple villages, whilst others (e.g. larger villages) have their own.

I compare those villages just below the cut-off who are forced to share a local government unit, with those that are just above the cut-off and therefore get their own local government. I assume that there are no significant underlying differences between the villages just above and below the cut-off prior to the implementation of the reform, and test this with the set of outcome variables that have data available for my baseline (see Table 1).

Table 1: Baseline Test (1991) for key outcome variables

	(1) Primary School (1991)	(2) Middle School (1991)	(3) Good Quality Roads (1991)	(4) Bad Quality Roads (1991)	(5) Commercial Building Use (1990)
Own Village Council	-0.0904	0.0575	0.0880	-0.0648	1.505
	(0.0649)	(0.0360)	(0.0644)	(0.0632)	(2.199)
Bandwidth	228.9	457.8	460.1	410.9	423.3
Control Mean	0.71	0.07	0.39	0.66	7.8

Notes: This table shows the results of running the RDD on the baseline outcome to ensure that there are no discontinuous jumps at the baseline. This allows us to ensure that any effects we find can be attributed to the reform.

4.1 Identification - RDD

The key identification strategy of my natural experiment is that there was a population-based eligibility threshold for villages to have their own village council. Therefore I consider the following empirical model:

$$Y = \alpha_0 + \alpha_1 D + \beta_1 (X - c) + \beta_2 D(X - c) + \epsilon$$

The state intended to use a threshold of one thousand below which a village would be paired up with other nearby villages to share a council. However there were cases in my data where the cut-off rules were not followed exactly. The most likely reason for this is the geographic distribution of villages sometimes required smaller villages to be paired up with much larger villages. To account for this imperfect compliance, I adapt my RDD to a fuzzy RDD, running a two-stage RDD regression. Here my running variable acts as an instrument for treatment status, and the first-stage looks at the likelihood that a village above the cut-off is selected into being treated.

Therefore I use the following 2 estimable equations:

First-stage:
$D_i = \alpha_0 + \alpha_1 C(X_i \geq 0) + \delta W_i + \gamma_d + u_i$ if $X_i \in [\underline{X}, \overline{X}]$

Second-stage:
$Y_i = \beta_0 + \beta_1 (X_i - 0) + \delta D_i + \beta_2 (X_i - 0) D_i + \delta W_i + \gamma_d + \epsilon_i$ if $X_i \in [\underline{X}, \overline{X}]$

Where D_i is the status of being "treated" (which here is defined as a village with its own panchayat), Y_i is the outcome of interest, $C(X_i \geq 0)$ is a dummy that takes a value 1 when the village has a population greater than the cut-off, W_i are village-level controls, and γ_d are district-level fixed effects.

The effect I am interested in is the impact of treatment status on my outcomes of interest, β_1, which is identified under the condition that my running variable (in this case population) is smooth across the threshold (see next section). All discontinuous jumps I observe in my data can therefore be attributed to the treatment under these conditions.

Bandwidth Selection I follow the empirical methodology put forward by Calonico et al. (2014) and use an optimal bandwidth selection method. There are two main advantages to this methodology. First, this methodology allows me to develop more robust point estimations based on data-driven optimal bandwidth selections, and second, it generates more robust standard error predictions which I can compare across different types of estimation methods to ensure my results are robust. Throughout the paper I present the results based on the optimal bandwidth, however in Appendix C I present the results across different bandwidths to ensure that my findings are robust. In specifications where I include controls or fixed effects, I follow Calonico et al. (2014) and do not interact them with my treatment or cutoff dummy.

Concerns A potential concern here is whether this cutoff is also used for other programs. I want to make sure that the results that I find are driven purely by this reform. The main program that may potentially clash is the Pradhan Mantri Gram Sadak Yojana (PMGSY) which was a large-scale road-building program started in 2000. The impacts of this program have been studied extensively in Asher and Novosad (2020), and many states used a 1000 person cutoff to determine eligibility for the program. However, not only was the population threshold based on the 2001 census (which puts it 10 years after my baseline), but more importantly, Uttar Pradesh was one of the few states not to follow the rule (Asher and Novosad, 2020). Hence, these cut-offs do not coincide with those used by other government programs, which either take the 2001 census as a baseline, or use different cut-offs. This ensures that the effects I estimate are unlikely to be due to other programs.

Changes over time Since there has been some geographical and administrative shifts over the period that we are interested in, it is important to consider how this should be dealt with. For districts that are demarcated to the newly formed state of Uttarakhand in 2000, I leave them out of the analysis entirely. For other districts and villages, I use the SHRUG dataset identifiers to ensure consistency over the three time periods. [6]

Generating the treatment variable Unfortunately the cleanest, most reliable source of information on the GP-to-village comes from the 2011 census. One source of concern with using this dataset is that changes in treatment status over time may be correlated with actual outcomes. Therefore I collected information on the number of panchayats over time as listed

[6] In general I try to use baseline geographic characteristics when possible to ensure that any of the changes are not interacting with my outcomes.

Table 2: Census of village panchayats by year

Year	No. Of Village Panchayats
1995	58,620
2000	52,929
2005	52,001
2010	51,914

Source: Uttar Pradesh Ministry of the Panchayat.

by the official state election committee [7]. The biggest change in the number of GPs during my period of interest is between the data collection period in 1995-2000, where as the updated Amendment came into effect, nearly 4000 GPs are dropped. Starting from 2000 however, the number of GPs stays relatively consistent (Table 2). As I use a fuzzy RDD, and some of the imperfect compliance may be driven by some treatment changes over time, my instrument helps to deal with some of the endogeneity concerns.

Validity of the RDD This paper combines multiple databases to generate a complete dataset on village councils in Uttar Pradesh. The final dataset contains 96,000 villages. I start with the 1991 dataset as the basis (which contains a total of 123,847 villages) and merge in the 2001 and 2011 censuses. To ensure that there is no bias in which type of villages get merged over time, I test the distribution of villages above and below the cutoff. As I unfortunately have to take the information on the Gram Panchayats (and the treatment status) from a later census, I cannot test the distribution of treated versus untreated villages directly, I instead test the distribution based on the cutoffs.

Table 3: Breakdown of village councils by cutoff status

Panel A: Unrestricted Dataset		
	Full Dataset	Merged Dataset
Population < 1000	66.2	61.9
Population > 1000	33.8	38.1
Panel B: Restricted Dataset		
	Full Dataset	Merged Dataset
Population < 1000	59.2	58.8
Population > 1000	40.8	41.2

[7] `http://sec.up.nic.in/site/commission.aspx`

Table 3, Panel A shows the distribution of the entire dataset. There is a more signification drop in the population under 1000, however a large proportion of this is driven by low-population villages as can be seen from Table 3, Panel B where I restrict the analysis to the proportion of villages that would be used for my RDD. In the restricted datasets, there is no significant observable differences between the datasets.

Density check We look at baseline outcomes first to ensure that there are no significant differences between treated and untreated groups in the pre-treatment period. The first thing we do is run a RDD density check (Calonico et al., 2014) to ensure that our running variable (in this case the population scaled so that the cutoff is at zero) is smooth across the cutoff. As can be seen from Figure 4, the population is smooth across the 1000 cutoff.

First-stage In Figure 3 we plot the effective first-stage of our fuzzy RDD; as can see from the graph, the first-stage is fairly strong with a clear 20% jump at the threshold. One minor concern here is that the likelihood of being treated (receiving your own GP) is already trending positively before we hit the cutoff. Hence the second-stage will likely be a little noisier than if we had had perfect compliance.

Figure 3: First-Stage check

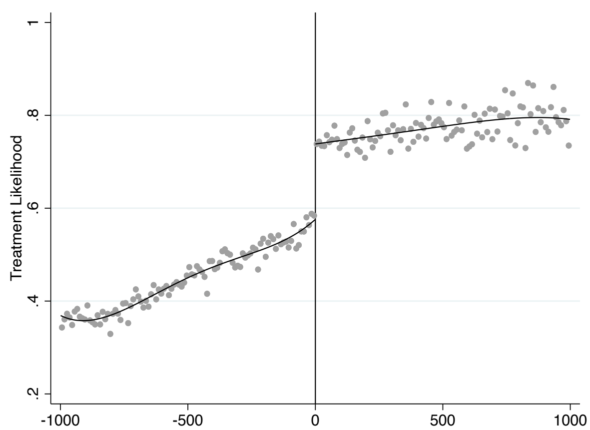

Notes: The figure shows the test for the jump at the cutoff of the treatment variable. The running variable has been restricted to below 2000 once again to be able to view the cutoff clearly. The y-axis is the treatment variable, and I use 100 bins on either side of the cutoff.

In order to ensure that the results are being driven by changes in the treatment variable and are not capturing underlying pre-trends, I run all my outcome variables at the baseline to ensure that they are balanced between treated and control villages. By checking that there are no significant differences between the treatment and control villages at the pre-reform year, I make sure that any significant effects I find can be attributed to the reform rather than anything else (see Table 1).

Figure 4: Manipulation of the cut-off

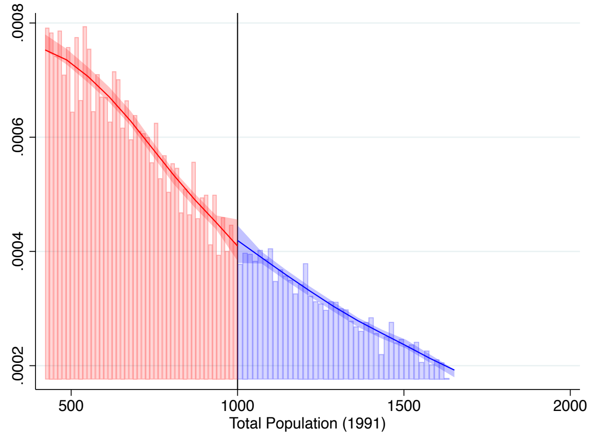

Notes: This figure is a visual representation of the density test of the running variable. The x-axis displays the population in 1991 (non-standardized), and for purpose of easy readability, it has been restricted to villages below 2000.

4.2 Empirical Strategy - District-level

In order to observe the effects of the reform on the district-level outcomes, I aggregate my RDD to the district-level. To do this I utilize the fuzzy RDD and generate a share of treated villages within a specific bandwidth, and then instrument for this using the proportion of the villages above the population cutoff within the same bandwidth.

Identification The core concern when looking at the district-level is that we may not be able to identify impacts of micro-level changes in a meaningful way. The major reason for this is that we often don't have exogenous variation at the district-level that allows us to perceive more than correlational patterns. However one unique aspect of my setting is that I have a method for generating some degree of exogenous variation, based on my RDD at the village-level.

I define "treatment" (independent variable) as the following:

$$TreatedShare_d = \frac{\sum \text{Number of Treated Villages}}{\sum \text{Total Number of Villages}} \text{ if } Pop_{vd} \in [\underline{X}, \overline{X}]$$

Where, v represents the village, and d the district.

To correct for potential endogeneity, and take advantage of the random variation induced by my RDD, I further instrument for the treatment share using a similarly aggregated version of the instrument used for the fuzzy RDD:

$$Instrument_d = \frac{\sum \text{Number of Villages above the Cutoff}}{\sum \text{Total Number of Villages}} \text{ if } Pop_{vd} \in [\underline{X}, \overline{X}]$$

First-stage I check the validity of my aggregate instrument to ensure that it is meaningful. The full results are available in Table B.1. The first stage is strong: the coefficient on the instrument is positive and significant, and the F-statistic is around 15. This instrument is highly correlated because I restrict the observations to ensure that the sample that is used for the RDD is what is underlying the variation in the aggregated independent variable, hence the strong first-stage. This empirical strategy allows me to take advantage of the exogenous variation induced by my RDD, but aggregate it to the district-level to generate meaningful estimates of the effect of the reform on aggregate outcomes.

5 Rural Development

The first of results I present are rural development outcomes. More specifically, I focus on outcomes associated with individual location choice. This means looking at outcomes related to public goods and amenities, generating measures of economic activity, and durable consumption. I choose outcomes that allow me to proxy for these categories, which gives me a stronger mechanism to link the development outcomes to the changes I observe in population movement.

5.1 Public Goods

There is a lot of evidence, both theoretical and empirical [8], in the literature on the importance of amenities and public goods in attracting individuals to a location. Furthermore, public goods are a key indicator of efficacy of local governments (Besley et al., 2004). To measure the impact of decentralization on public goods in my setting, I use two key outcomes - education and roads. Education is an important contributor to human capital, and focusing on this as a public good seems almost obvious. Better access to education has long been a key rural development policy [9], and education plays a key role in driving migration decisions which I discuss in more detail in the migration section of this paper. Roads are also considered a key public good, with extensive findings showing the importance of roads, especially in India (Asturias et al., 2019).

5.1.1 Education

Starting with education, Table 4 shows the impact of the reform on the provision of education at the primary and middle school level at the medium-term (2001) and long-term (2011). The independent variable is measured as a dummy variable indicating whether or not a school is

[8] e.g. Bryan and Morten (2019), Redding and Rossi-Hansberg (2017)

[9] This ranges from work done by both the federal government and international organizations in driving education rates up.

available in the village. I find that treated villages have a higher provision of middle schools in both 2001 and 2011, whilst in 2001, there appears to be a higher provision of primary school but this effect .

Table 4: Educational Provision

	(1) Primary School (2001)	(2) Middle School (2001)	(3) Primary School (2011)	(4) Middle School (2011)
Own Village Council	0.132**	0.0951**	0.0433	0.235***
	(0.0558)	(0.0483)	(0.0592)	(0.0648)
Observations	95,755	95,755	94,851	94,851
Effective Observations	10,849	17,564	12,384	17,310
Bandwidth	234.1	350	262.3	346.8
Control Mean	0.83	0.15	0.80	0.40

Notes: This table shows the coefficient on the RDD treatment variable on educational provision at the primary and middle school level for the years 2001 and 2011. Primary school covers children aged 6-11, and middle school is 11-14. The dependent variable is a dummy variable taking a value 0 if a school is unavailable in the village, and a value 1 if it is available. As always, the specification is run with district-level FE and the standard errors are clustered at the district-level. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

In 2001, treated villages had a 9 percentage point higher provision of middle schools, and a 13 percentage point higher provision of primary schools, both significant at the 5% level. By 2011 however, there was no significant difference in the provision of primary education between control and treated villages. Regarding the disappearance of a significant effect over time for primary schools, this is likely driven by the introduction of the Sarva Shiksha Abhiyan program introduced in November 2000, which is a federal education program focused on the provision of basic education throughout the country. The main focus of this program was the provision of elementary, or primary, education, which may explain the disappearance of effects post-2001, as the provision of primary education was no longer determined at the village-level. Unfortunately, data restrictions mean that it is not possible to directly test this. An alternative explanation is also that there was some degree of saturation in the provision of primary schooling, given that 80% of the control villages had primary schooling by 2001, which meant that there were no significant gains left by 2011, hence the lack of effect.

In contrast, middle school effects continue to persist through 2011. More decentralized villages were 23 percentage points more likely to have a middle school at the 1% significance level, which was a significant jump from 2001. This suggests that the effect of education provision at the middle school level was compounded over time, and more importantly, that left up to the local councils, this effect persisted through time. One concern may be that population differences are driving these effects, however in Appendix H I show that controlling for the ratio of schools-to-population has no effect on the magnitude or the significance of the effects I find.

Combined, the effects suggest that individuals in the treated villages have better access to education, especially at a higher level. In addition, the effects begin in 2001, and persist over time for higher levels of education. Though we can't say much about the quality of education provided at these villages, just the availability of this public good, may already make these locations more attractive.

5.1.2 Roads

The second public good I look into is the impact of having an own village council on road provision. Road quality is divided into two types - good quality versus bad quality. A good approach is a road that has some kind of gravel/tar structure, and is easily accessible by car. Bad roads are usually simple mud roads, or roads that do not have defined structure or building. For 2011, unfortunately only information on blacktop roads was available, so there is no "bad" comparison. These variables measure the availability of these roads, but I can't do a more detailed breakdown of the quality of these roads apart from blacktop versus mud due to data limitations. Therefore roads that are blacktop but badly maintained would still show up as "good" quality roads.

Table 5 shows that there is evidence to suggest that treated villages initially did have higher quality provision of roads. Treated villages having a 16 percentage point higher likelihood of having a good road, significant at the 1% level. In addition, villages with their own GP were 12 percentage points less likely to have a bad road. However, similar to the primary education program, there was a federal road program introduced post-2001. This program has been studied extensively previously by Asher and Novosad (2020), which may explain the lack of observable effect by 2011. Uttar Pradesh was one of the few states to not follow specific population rules in implementing the road building road program, hence it is unlikely that my analysis would be able to capture any significant results after the introduction of the program with my RDD.

Table 5: Roads

	(1) Good Road Quality (2001)	(2) Bad Road Quality (2001)	(3) Good Road Quality (2011)
Own Village Council	0.162***	-0.122**	-0.0567
	(0.0577)	(0.0538)	(0.0486)
Observations	87,883	87,883	87,706
Effective Observations	23,643	17,951	20,952
Bandwidth	459.3	367.5	417.6
Control Mean	0.58	0.68	0.69

Notes: This table shows the coefficient on the RDD treatment variable on the provision of roads. The dependent variable is measured as a dummy variable with 0 if that type of road is not available, and 1 if it is available. The census measures two types of roads, *pucca* (which is good quality) and *kuccha* (which is bad quality). A good quality road is usually one that is permanent, usually made up or tar or gravel. Bad quality roads are ones that are just mud. As always, the specification is run with district-level FE and the standard errors are clustered at the district-level. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Road-building has significant impacts on a multitude of factors. Work by Donaldson (2020) and Asturias et al. (2018) are just some of the papers that highlight how important transport infrastructure is for reducing spatial misallocations, improving productivity, and welfare improvements. Improving access to markets for these villages could have significant positive effects, especially on economic activity, and the evidence suggests that in 2001, there were significant benefits to being in a decentralized village which may have made treated villages more attractive locations. These effects mostly disappear by 2011, however any effects on population could be driven by the initial differences in 2001.

Overall, there is some evidence to suggest that public goods such as educational provision and roads were directly affected by having an own GP, however in some cases federal programs superseded the local governments in managing these amenities. The federal governments seem to have been more effective in some of these cases in reducing the village-level differences between control and treated villages, which is in many ways a good thing, however it is unclear whether the aggregate level effects are overall positive, and I will not be able to compare directly effects between federal control with village-level control. The effects are clearly positive in 2001, and to test whether it had an impact on how attractive these locations were, I look at population change between 2001 and 2011 in the migration section.

5.2 Economic Activity

Given the improvements we see in roads and education at the village-level, it would stand to reason that these changes may have effects on the economic activity in these areas. Higher levels of education drive up human capital, and better access through roads improve market access and reduce mobility frictions. To look into this in more detail, I look at changes in economic activity in these villages.

I look at the number of commercial enterprises available in villages, as well as the share of firms engaged in agricultural activities. Table 6 shows that treated villages show a growth in the number of commercial enterprises in 2011 relative to control villages, although the magnitude are relatively small. Treated villages in 2013 have 5 more commercial establishments. This is measured by the proportion of buildings within a village that is utilized for commercial business purposes. Furthermore, treated villages show a 4.7 percentage point decline in the share of agricultural enterprises compared to control villages. This evidence is suggestive of changes in the type of economic activity in treated villages, however the magnitudes are relatively small, and the changes marginal.

Table 6: Economic Activity

	(1) No. of Commercial Properties (1990)	(2) No. of Commercial Properties (2013)	(3) Agricultural Activity (%) (1990)	(4) Agricultural Activity (%) (2013)
Own Village Council	1.505	7.040**	0.510	-4.797***
	(2.199)	(2.960)	(1.852)	(1.479)
Observations	88,100	88,100	88,100	88,100
Effective Observations	21,364	20,808	22,869	18,940
Bandwidth	423.3	414.3	447.4	383.9
Control Mean	7.81	10.98	4.40	3.70

Notes: This table shows the coefficient on the RDD treatment variable on a set of economic characteristics. Column (1) measured the effect on the use of commercial properties in the baseline year 1990. This is measured in terms of number of properties. Commercial properties are defined as buildings used for commercial purposes such as retail. Column (2) measures the same variable in 2013. Column (3) measures the effect on agricultural activity as a share of total activity in 1990, and (4) measures the same variable in 2013. Agricultural activity is defined as firms or businesses engaged in agricultural practices as defined by the economic census. As always, the specification is run with district-level FE and the standard errors are clustered at the district-level. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

The combination of improved road access and higher education may contribute to this finding that agricultural activity seems to be declining. These villages are undergoing some degree of structural change at a micro-level, and the evidence seems to imply that decentralization is driving improvements in public goods that help to change the type of economic activity that is taking place in these villages. The treated villages start acting like rural centers for economic activity, but further research is required to understand the degree to which this happens. It's additionally possible that the functionalities of the village councils directly (e.g. encouraging local development) may be a contributing factor to increasing economic activity, unfortunately there is no way to directly test this channel. In the next section however I look at the effects of the decentralization on individual consumption choices, which may give some indicators of how the changes are impacting individual incomes.

5.3 Durable Consumption

So far I have focused purely on village-level outcomes that are driven by direct impacts from decentralization. It is also important however to measure standard of living in these villages, to gain a better understanding of how the reform affects individuals more directly. Although I don't have detailed individual-level data, I can use information on durable (long-term) consumption measures from the village-level. The main source of this data is the housing census data from 2011 [10], which gives detailed breakdowns of the type of material quality utilized in house building in villages. I additionally collect information on sanitation availability, banking access, and lighting sources as additional information on standard of living for these individuals.

Figure 7 shows that treated villages are more likely to use high-quality material such as stone and brick rather than low- quality material such as grass and corrugated metal. Villages with their own GP are -7.6 percentage points less likely to use grass/thatch roofing material, but 9.2 percentage points more likely to use brick. Grass makes up 25% of households in control villages, whilst brick (the most popular material) makes up 32%. Additionally there is a 3.3 percentage point higher chance that a treated village uses stone building material, and a -1.6 percentage point decrease in the chance of using corrugated metal roofs, although the overall proportion of control villages using metal is only 3%. In Appendix D I present the full breakdown by roof and wall material covering all potential types of material. A somewhat surprising result was that concrete is 4.1 percentage points less in treated villages - one potential explanation for this is that brick and concrete act as substitutes. Finally, there seems to be no significant difference in the use of tiles in roofing between control and treated villages.

[10] The independent variables in these regressions are all measured as the share of homes within the village that fall within that specific category. Therefore no one household should appear in multiple categories.

Table 7: Durable Consumption (2011)

	Roofing Material					Wall Material		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Brick	Tiles	Stone	Grass	Concrete	Brick	Mud	Grass
Own Village Council	9.222***	1.645	3.288**	-7.560***	-4.135***	8.278***	-4.291**	-2.738**
	(2.516)	(2.096)	(1.539)	(2.239)	(1.602)	(2.134)	(1.954)	(1.197)
Observations	94,856	94,856	94,856	94,856	94,856	94,856	94,856	94,856
Effective Observations	18,796	17,505	20,556	16,206	19,887	19,552	19,122	13,880
Bandwidth	370	349.5	397.6	328.5	386.2	381.5	374.2	289.5
Control Mean	32.06	12.45	11.27	25.53	14.03	61.38	26.66	6.34

Notes: This table shows the coefficient on the RDD treatment variable for two categories of housing material - roof and wall. The data is broken down by the proportion of households that use that type of material in the village. For the full results please see Table D.2 and Table D.1 in Appendix D. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

The material used in treated villages are not only more expensive, but also tend to be long-lasting. In addition to being suggestive evidence that individuals in treated villages are capable of investing in higher-quality material, this evidence also suggests that individuals in treated villages are making more long-term housing investments, which is suggestive of a long-term desire to remain in these areas. These results are supported by the evidence from the wall material quality. Figure 7 shows that treated villages are more likely to utilize brick materials over the more traditional and low-cost option of grass and/or mud. Treated villages are 8.3 percentage points more likely to use brick for their wall material. In addition, they are 2.7 percentage points less likely to use grass, and 4.3 percentage points less likely to use mud. In general however, grass is a relatively rarer material for walls as only 6.3% of control villages use this material. Mud is a more common material in control villages, with 26.7% of households in control villages using it, and brick is the most commonly used material at 61.4%.

As a final check on housing quality, I also look into the type of housing that is prevalent in these villages. The results can be found in Table D.5 in the Appendix, and shows the impact of the reform on the type of housing available. I find that treated villages have a significantly higher proportion of permanent housing, which is defined by the use of more long-lasting materials. This results serve as a check with the previous results of this section, as housing status is defined by the type of material used in roofing and walls. I find that treated villages have a 7.8 percentage point higher change of having permanent houses, where control villages already have 58% of households living in permanent houses. In line with this, villages with their own GP are 6.6 percentage points less likely to have temporary housing.

Given that there are significant differences in individual investment in housing, I also check other alternative measures of consumption, in the form of source material used for lighting. Lighting sources are traditionally based on kerosene or other similar fuels, but large electrification programs have slowly been rolled out across this period. However although a village may be connected to the grid, it isn't always the case that the houses within the village are actively connected to this grid. Looking at the choice of lighting (Table D.4), I find that treated villages are 7.9 percentage points more likely to be utilizing electricity rather than kerosene as the main source of light usage. Although kerosene is subsidized, it is still considered less efficient, more unhealthy, and less preferred compared to electricity. Electricity in general still only makes up less than a quarter of households main light source in treated villages. Kerosene is often additionally used even in households that are connected if the connection is poor and the usage of back-up generators is high. The other lighting sources are negligent in terms of the proportion of households using them, and in terms of the differences between control and treated villages. Kerosene remains the most used type of lighting, with 75% of all households in control villages using it as their main light source. This finding suggests that houses in treated villages not only have better quality housing, but they also seem to have greater capabilities to access infrastructure improvements.

This also holds true when I look at the impacts on the provision of sanitary facilities *within* homes. Table 8 shows that across the board, treated villages not only have a higher proportion of households with latrine facilities, but that they have better access in their homes for multiple types of sanitary facilities. Column (1) is measured in the number of households, whilst the other categories are measured as before with the share of households reporting which of the potential options applies to them. These results likely come from a mixture of access to better sanitary service provision at the village-level and higher capability to afford to construct such facilities within the household. The results suggest that treated villages have on average 18 households more with latrine facilities. Given that the control average is 17 households, this is a pretty significant difference between treated and untreated villages. In addition, the most commonly found sanitary facility appears to be a septic tank, and treated villages have a 5.5 percentage point increase on a baseline of 9% of households in control villages. Like the number of households with a latrine, this is a fairly large magnitude.

Table 8: Sanitary Facilities

	(1) No. of HH with latrine	(2) Piped Latrine	(3) Septic Tank	(4) Slab pit	(5) No slab pit
Own Village Council	18.41***	2.633***	5.536***	6.244***	1.337***
	(2.756)	(0.927)	(1.489)	(1.612)	(0.357)
Observations	94,856	94,856	94,856	94,856	94,856
Effective Observations	15,274	14,107	19,887	15,517	23,720
Bandwidth	312.4	293.8	386.2	316.7	445.7
Control Mean	16.73	1.688	8.974	3.122	0.689

Notes: This table shows the coefficient on the RDD treatment variable on the availability of certain types of sanitary facilities. For households with sanitary facilities, the facilities are broken down into the following categories: Piped sewer system, Septic tank, Other system, Pit latrine with slab, Pit latrine without slab, Night soil disposed into open drain, Night soil removed by human, Night soil serviced by animal. The full breakdown by treatment and control villages is available in Appendix D. As always, the specification is run with district-level FE and the standard errors are clustered at the district-level. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

As with the findings on lighting, these results seem to point towards individuals in treated villages not only having a better ability to afford higher quality housing, but also able to access the infrastructure required to make those improvements happen. Not only do villages with their own council show higher levels of connection to the electric network, but they also have a greater availability of sanitation facilities across the board, and at a higher quality level. These results together seem to suggest that individuals within these villages are better off, but what remains unclear is whether this is driven by improvements in development that directly impact incomes, or whether it is driven by changes in the population composition of these villages. In the section on rural population growth I discuss these two channels at greater length. One potential channel that is in line with the findings however is that the improved development outcomes at the village-level for public goods combined with the changing economic activity, contribute to improvements in income at the individual-level.

These results are also reflected when I look at banking access (see Appendix D, Table D.3). Banking access is measured as the total number of households per village who are availing banking. Treated villages have a significantly higher number of households using banking at the 5% significance level, however the magnitude is relatively small. Villages with their own GP have 5 more households engaging in banking relative to the baseline of 78% of households in control villages that are engaged in banking. Taken with the previous results, this also suggests more engagement in the formal economic sector.

In general, villages that have been treated exhibit a higher proportion of houses with more permanent and high-quality materials. The higher cost of such materials as well as the choice of individuals to invest in longer-lasting material is suggestive evidence that individuals in treated villages are more able, and more willing, to invest in long/term, high-quality housing. This evidence also suggests that these individuals have higher durable consumption patterns given the type of investments they are choosing to make in their houses.

Overall, the findings in this section suggest that not only do we find significant benefits to being more decentralized in terms of public goods, but that these improvements generated changes in economic activity and individual utility. Taken together with the findings that individuals had higher levels of durable consumption across multiple measures, the treated villages seem to be doing better on a multitude of measures. This suggests that they were in general more attractive locations for individuals to settle in, and we do see evidence that there were people investing into long-term housing in the area. These improvements would suggest that there could be reallocation of individuals across space, and the next section tests this hypothesis more thoroughly.

5.4 Rural Population Growth

Models of location choice such as Bryan and Morten (2019) would tell us that given the improvements in development outcomes that I observe at the village and individual-level these rural locations may become more attractive for individuals to settle in or move to. In order to test whether this is the case, and whether the results from the first half of this paper have an impact on individual choice, I look first at rural population change at the village-level using my RDD. Developments in public goods, plus evidence suggesting that economic activity and standard of living in these treated villages are going up, may suggest that these villages are now more attractive for people.

A key indicator of preference for a location comes from population movement. Since there is a lack of detailed village-level data on migration, this paper uses population movement to proxy for migration to try and capture some of the underlying population trends. Using the population totals at the village-level from the 1991, 2001, and 2011 census data, I calculate the population change between 1991 and 2011, as well as between 1991-2001, and 2001-2011. Given the reform in Uttar Pradesh was officially started in 1995, and that population changes are measuring long-term locational preferences, I expected to see the majority of the effect, if any, kicking in towards the latter half of the twenty year period. In addition, changes to public good provision seem to take effect by 2001, so any changes in population would be driven by these observable differences.

Populations changes can be thought of as revealed preferences. Table 9 shows the RDD results of population change over the period this paper covers. As we would expect, we see the effects on population change kicking in after the changes to public goods have already been observed. From column 2 we can see that there is a significantly higher population growth in villages that are treated. These villages have 9.4 percentage points higher growth compared to control villages, and this explains the majority of the population change over the twenty year period [11].

Table 9: Population Change

Population Change	(1) 1991-2011	(2) 2001-2011	(3) 1991-2001
Own Village Council	10.30	9.454***	-0.0196
	(6.353)	(2.695)	(3.983)
Observations	95,755	95,755	95,755
Effective Observations	9273	12,210	10,034
Control Mean	50	21	26
Bandwidth	204.2	258.7	219

Notes: This table shows the coefficient on the RDD treatment variable on population change in percentage. Column (1) shows the change over the entire period of interest, 1991-2011, column (2) shows the change between 2001-2011, and column (3) shows the population change between 1991-2001. I wincorize all the population data to ensure that there are no errors affecting the results. To wincorize I take the bottom and top 1% value and replace any observations outside those with the 1% value. All the regressions are run with district-level fixed effects, and district-level clustered standard errors. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

The delayed onset of population change is likely driven by individuals observing the impacts of decentralization before making the decision to move. The provision of amenities in the earlier period may encourage individuals to move to villages where they perceive that people have better access to public goods, or better ability to engage with the political system.

One concern however is that the change in population may be driven my changes in fertility between treated versus untreated village. In order to ensure that the population changes are generated by population movements rather than differential fertility rates, I also check for changes in population aged 0-6 both in levels and in changes, and find no significant differences between the treated and untreated villages (Table F.1 presents the results of the changes). The standard errors on the coefficient estimates are relatively large, however the point estimates are still less than the changes we observe at for the total population, suggesting that the change can't be driven purely by fertility changes.

Given the estimates on changes in fertility rates however, it is possible that in the future we may observe changes in the fertility rate as the effect of the population change between 2001-2011 starts to kick in. The results from the gender breakdown (Appendix H) also suggest that there is an similar effect on population change for both genders, which suggests that this is not just single adult males moving for work, but more likely to be families or groups

[11] The population change is reflected also in the individual gender breakdowns (Table E.1 & Table E.2)), though an interesting result is that the effect is larger for women versus men, although both are significant.

of individuals relocating. We may therefore observe changes in the fertility rates in future censuses.

5.4.1 Spillover Effects

One concern given the population results may be that the development effects from the rural development section may be driven by population changes rather than direct impacts of decentralization. Results on public goods are already in effect by 2001, which pre-dates the population growth, hence it's unlikely to have been driven by the population changes; if anything it is the likely driver of the population change. On the other hand, the durable consumption results could be driven by wealthier individuals moving into treated villages from control villages. This would suggest that we overestimate the effects on durable consumption, as there are spillovers on control villages. To test for this, I run a robustness check which controls for potential spillovers.

Empirical Strategy I use the spatial distribution of villages to disentangle some of the spillover effects. I do this by calculating the share of villages in the surrounding areas for different distance ranges. I restrict the villages used in the calculation to within a specific population bandwidth to take advantage of the RDD structure and control for some of the endogeneity concerns associated with population distribution (the main results presented are for population within 300 below and above the cutoff). I calculate the share of villages in the surrounding areas that are above the cutoff for the three radii: 10km, 20km, and 50km. I use just the cutoff results and focus on the reduced form results. I then include these shares as controls and re-run the RDD.

Durable Consumption Spillovers I present the results in Appendix G. The results show that controlling for share of treated villages in the surrounding areas has little to no-effect on the coefficient on durable consumption across all measures. There is no significant difference in the magnitude of the effect, and the effect remains significant. This suggests that the effects are not being driven just through population movements, and the limited changes in the magnitude on the treatment coefficient suggests that it's unlikely that the effects are driven purely through the changes in population.

Summary To sum up the findings at the village-level, I find first, that there are significant developments at the village-level in terms of provision of public goods. This is in line with findings from other papers (e.g. Besley et al., 2004; Chattopadhyay and Duflo, 2004a) who find that preferences are better met after decentralization, and that public good provision goes up. My results take these findings a step further however, as I also have empirical evidence of more indirect changes at the village-level, such as structural change and higher standards of living as measured through durable consumption. Overall a greater intensity of decentralization has a positive effect on village-level outcomes, and treated villages become more attractive locations for people to settle. Looking at the results through the eyes of location choice models, we can then measure revealed preferences by observing changing patterns of population growth at the village-level. I find that villages with a greater degree of decentralization have higher rates of population growth, especially towards the latter half of the period of focus on, specifically after observing development improvements at the village-level. This is an important finding, as it provides micro-level evidence on the effect of development on population movement, [12] and tells us that micro-level changes can trigger spatial changes in population distribution. To understand more deeply the impact that these spatial changes have, I next observe district-level effects.

6 Migration and Growth

The reform that I focus on provides me with a unique setting to be able to not just observe empirically the effects at the local-level, but to use my empirical strategy to identify some of the aggregate effects on district-level outcomes. Given that we observe rural population changes at the village-level, there are two potential channels which could be driving these effects. First, the population growth could come from people substituting away from urban migration. This is the main concern with the effects that we find, as we have evidence that higher urbanization rates are positively correlated with growth. The second potential channel is that the rural population changes are driven by a reshuffling of the rural population, as treated villages become rural centers, attracting people from untreated villages. To test both these channels, I use migration data at the district-level that allows me to break down within-district migration by rural-rural and rural-urban. Unfortunately data limitations and the use of district-level data, means that I only have 42 observations available for analysis, so this limitation should be kept in mind when interpreting the results that follow.

6.1 District Migration

As i discussed in the empirical strategy section of my paper, one of my main contributions in this paper is that I can use an aggregated version of my RDD to generate empirical estimates at the district-level. Using my aggregated independent variable - share of treated villages - and instrumenting for it with the share of villages above the cutoff, allows me to use the exogenous variation inherent in my RDD.

[12] Work by Bryan and Morten (2019) for example find similar effects but at a much more regional level.

I run the following regression at the district-level:

$$\frac{\text{Migrants}_d}{\text{Total migration population}}_s = \alpha_0 + \alpha_1 Treat_d + \alpha_2 Pop_d^{91} + \epsilon$$

I calculate the proportion of migrants as a share of the rural population in order to account for differences in population distribution across districts, and run the regression at the district level. I also run a robustness check for my results by changing the output variable to the share of migrants out of the total migration population, and include these results in the following tables.

Rural-Urban Migration Table 10 shows that there is a positive effect of treatment share at the district-level on rural-urban migration. A one standard deviation increase in the share of treated villages within a district leads to a 1.7% increase in the share of total migration. The effect is also positive and significant when broken down by gender.

Table 10: Total Migration: Within District (Rural-Urban)

	Share of Total Population			Share of Total Migration		
	(1)	(2)	(3)	(4)	(5)	(6)
	Total	Male	Female	Total	Male	Female
Share of Treated Villages	0.155***	0.0803***	0.0748**	0.207***	0.108***	0.0997**
	(0.0581)	(0.0296)	(0.0293)	(0.0661)	(0.0274)	(0.0411)
Mean	2.75	1.04	1.71	0.21	0.28	0.14
Observations	54	54	54	54	54	54

Notes: This table shows the district-level regressions on within-district rural-urban migration. In all the regressions, treatment status is instrumented by the cutoff dummy. I restrict the calculation of the treatment share to three hundred below and above the bandwidth. The dependent variable is migration share, which is measured as a share of the total number of migrants out of the rural population. Standard errors in parentheses, *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Though the initial concern was that improvements in development may encourage substitution away from rural-urban migration, I find quite the opposite. The increase in rural-urban migration is quite surprising, but could potentially be explained by two potential channels. First, the improvements in availability of education starting from 2001, suggests that individuals in treated villages may have higher levels of education. This is in line with findings from other studies such as Schultz (1960) and others who find that education reduces barriers to migration. Greater availability of educational institutions at the primary and secondary level in 2001 in treated villages (and carried onto 2011) may provide a higher proportion of individuals with access to basic education. This may improve their access to jobs in urban areas and increase their expected wage.

Additionally, another potential channel is that financial migration constraints may be mitigated, given that we observe higher levels of durable consumption in treated villages. Higher levels of durable consumption may suggest that households in treated villages face lower financial constraints and risks. Given that migration often has high associated financial risks and costs (e.g. Munshi and Rosenzweig, 2016), some of these risks may be mitigated from treated villages. This drives up the likelihood and the ability of individuals from these villages to migrate to urban areas.

Taken together these findings suggest that improvements in rural development may actually contribute to reducing migration frictions, which could have a positive effect on aggregate growth. Given that the reform measures the intensity of decentralization rather than a direct comparison of decentralization versus none, the magnitude of the effects seem quite plausible. Given the relatively small sample size these results are based on, I can confidently rule out a negative effect on rural-urban migration, however I can't say much about how large the positive gains may be.

Rural-rural Migration The finding that there is no negative effect on urbanization rates suggests that the change is coming instead from rural reshuffling. To test this, I look at within-district rural-rural migration. Table 11 shows that both the total and the gender-breakdown show a negative effect of treatment share at the district-level on share of rural migration. This fits with the hypothesis that the population growth at the village-level is driven by a reshuffling of the rural population from control villages to treated villages. This fits with the literature on location choice models, as population growth in treated villages suggests that improvements in public goods and changes in economic activity as well as the evidence that individuals in these villages experience a higher level of durable consumption, all suggest that treated villages become more attractive locations for individuals to move to. The developments at the village-level seems to result in a greater concentration of population in treated areas.

Table 11: Total Migration: Within District (Rural-Rural)

	Share of Total Population			Share of Total Migration		
	(1)	(2)	(3)	(4)	(5)	(6)
	Total	Male	Female	Total	Male	Female
Share of Treated Villages	-0.494***	-0.0247**	-0.469***	-0.958***	-0.0521**	-0.906***
	(0.113)	(0.0101)	(0.106)	(0.253)	(0.0206)	(0.240)
Mean	17.54	1.32	16.23	42.49	3.05	39.45
Observations	54	54	54	54	54	54

Notes: This table shows the district-level regressions on within-district rural-rural migration. In all the regressions, treatment status is instrumented by the cutoff dummy. I restrict the calculation of the treatment share to three hundred below and above the bandwidth. The dependent variable is migration share, which is measured as a share of the total number of migrants out of the rural population. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

To test whether this rural reshuffling hypothesis holds at the village-level, I regress village-level population change between 2001 and 2011 on the share of villages in the surrounding area that has a population above the cutoff. (See Appendix G for more details).

Figure 5: Spillover: Effect of treated villages in the surrounding area on population change (2001-2011)

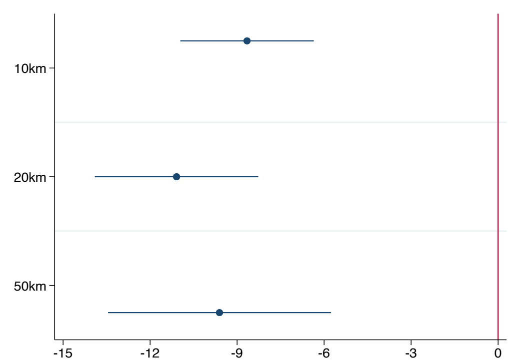

Notes: This figure shows the coefficient on the independent variable measuring share of villages in the surrounding area whose population is above the cutoff. The outcome is the population change between 2001-2011.

Figure 5 shows the coefficients on the share of villages above the cutoff for three distance ranges - 10km, 20km, 50km. The interpretation of this coefficient is that as the share of villages above the cutoff in the surrounding areas grows, the population change in the villages itself is lower. This aligns with the idea that individuals are moving to treated villages from control villages. As the number of surrounding villages that are "treated" (since this is a reduced form variable) increases, the villages grow more slowly. This result suggests that the effects may have an inverted-U shape, as at some point once all the villages are treated, we would not expect any population change in any villages. In Appendix G I also show that the RDD population change results are robust to the inclusion of the shares as a control, and that the interaction term between the village treatment status and share of villages above the cutoff in the surrounding areas is negative as we would expect, as the effect of having treated villages in the surrounding area is lower for villages that are themselves treated.

Summary Taken with the findings on rural-urban migration, there are two key takeaways from this section. First is that village-level developments seem to have no negative consequences on urban migration rates, as they may help to reduce migration frictions. Second, the population changes we observe at the village-level are likely driven by reshuffling of the rural population, as I find evidence of a negative correlation between the share of treated villages in a district and the rural migration share. In addition, there is also evidence at the village-level, that having a lot of treated villages in the surrounding area may reduce the rate at which a village grows, however overall treated villages still grow more, even when controlling for the share of treated villages in the area. Given these changes in the spatial distribution of population, especially the increase in rural-urban migration, the final outcome we need to look at is the impact on district-level GDP.

6.2 District GDP

The combination of the improvements in development combined with the changes in the spatial distribution of population, are likely to drive changes in aggregate growth. We know from work by Gadenne and Singhal (2014) that rural-urban migration contributes to economic growth, and the developments at the village-level, especially the shift away from agriculture, may itself increase aggregate growth. To test this, I proxy for growth using district-level GDP per capita data, as well as sectoral data, all taken from the ICRISAT dataset. I run the same regression as the migration outcomes, replacing the migration share with GDP per capita and GDP in the primary, secondary, and tertiary sector.

Figure 6: Decentralization Intensity and GDP at the District-level

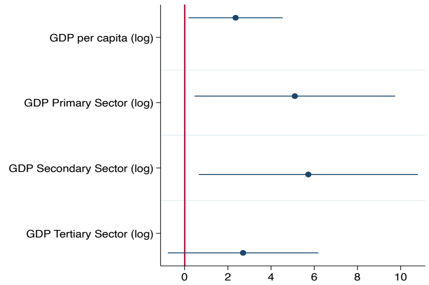

Notes: This figure shows the district-level regressions on four key outcomes - GDP per capita, Primary Sector GDP, Secondary Sector GDP, Tertiary sector GDP, all measured in current prices. The regressions instrument for treatment status with the cutoff dummy. This figure restricts the cutoff aggregations to using only villages with a population that is within three hundred of the cutoff.

The results are presented in Figure 6. I find that despite a lack of precision in the estimates, the results are fairly sizable and positive, with the GDP per capita increasing by at most 29%, and at least 0.3% for a one standard deviation increase in treated share of villages within a district. Based on these results, we can rule out a negative impact on GDP, and this effect holds across all three sectors as well. [13] Unfortunately I am unable to disentangle the exact channel driving this increase in GDP, however given the fairly large standard errors and the small sample size, it is more meaningful to treat this result as a non-negative rather than a firm positive outcome.

Summary The main concern in the research question was that a sufficiently high level of rural development may negatively impact urban migration, which in turn would have negative consequences for aggregate growth. In contrast, I find that not only does rural development have positive effects on rural-urban migration, but it also seems to have a positive effect on GDP per capita, or at the very least a non-negative effect. This is an important finding, because I am able to link micro-level changes to aggregate changes in an empirical framework.

7 Conclusion

In this paper I provide micro-level evidence on the impact of a decentralization reform on village-level development using an RDD. I then empirically analyze the impact this reform had on district-level outcomes, using an aggregated version of my RDD strategy. I find that treated villages, those with a higher degree of decentralization, have a higher provision of public goods before federal control, show signs of structural change away from agricultural activity, and have higher levels of durable consumption. These improvements at the village-level also lead to an increase in population growth at the village level for treated villages,

[13] I also breakdown the results by year and category in Appendix I

which is consistent with the the theoretical models of location choice which suggest that these treated villages become more attractive locations for people to settle.

In the second half of my paper, I use my RDD to generate a measure of the share of treated villages at the district-level, to analyze aggregate outcomes. I find that there are surprisingly no negative consequences on rural-urban migration, or on GDP. Instead my findings are consistent with the result that decentralization-led development helps to relax some of the educational and financial frictions associated with rural-urban migration. Additionally, the population changes are driven by a reallocation of the population within rural areas, and the combined effect of development at the village-level and increased rural-urban migration is consistent with the findings on increased GDP per capita at the district-level.

The results of this paper suggest that there are indirect consequences to decentralization and development, that are not necessarily captured if we focus on just the village-level outcomes. Rural development, rural-urban migration and growth are closely linked, and this paper shows that micro-level changes need to be studied alongside aggregate outcomes in order to meaningfully measure the impact of policy such as the decentralization reform.

Given the results in this paper, there are still some open questions that need to be answered with further research. First, the exact channels through which rural development impacts aggregate outcomes, for example, what are the drivers of the district-level GDP effects, and can we disentangle the contribution of rural-urban migraition versus village-level development? In addition, understanding the flows of migration, most importantly what the compositional effects are, if left to future work.

References

Acemoglu, D., S. Naidu, P. Restrepo, and J. A. Robinson (2019). Democracy does cause growth. *Journal of political economy 127*(1), 47–100.

Allen, T. and C. Arkolakis (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics 129*(3), 1085–1140.

Asher, S., T. Lunt, R. Matsuura, and P. Novosad (2019). The Socioeconomic High-resolution Rural-Urban Geographic Dataset on India (SHRUG). Working paper.

Asher, S. and P. Novosad (2020). Rural roads and local economic development. *American economic review 110*(3), 797–823.

Asturias, J., M. García-Santana, and R. Ramos (2019). Competition and the welfare gains from transportation infrastructure: Evidence from the golden quadrilateral of india. *Journal of the European Economic Association 17*(6), 1881–1940.

Bank, W. (2000). Overview of rural decentralization in india.

Bardhan, P. K. and D. Mookherjee (2000). Capture and governance at local and national levels. *American economic review 90*(2), 135–139.

Bardhan, P. K. and D. Mookherjee (2006). *Decentralization and local governance in developing countries: A comparative perspective*. Mit Press.

Barro, R. J. (1996). Democracy and growth. *Journal of economic growth 1*(1), 1–27.

Besley, T., R. Pande, L. Rahman, and V. Rao (2004). The politics of public good provision: Evidence from indian local governments. *Journal of the European Economic Association 2*(2-3), 416–426.

Besley, T., R. Pande, and V. Rao (2007). Political economy of panchayats in south india. *Economic and Political Weekly*, 661–666.

Bryan, G., S. Chowdhury, and A. M. Mobarak (2014). Underinvestment in a profitable technology: The case of seasonal migration in bangladesh. *Econometrica 82*(5), 1671–1748.

Bryan, G. and M. Morten (2019). The aggregate productivity effects of internal migration: Evidence from indonesia. *Journal of Political Economy 127*(5), 2229–2268.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust data-driven inference in the regression-discontinuity design. *The Stata Journal 14*(4), 909–946.

Chattopadhyay, R. and E. Duflo (2004a). Impact of reservation in panchayati raj: Evidence from a nationwide randomised experiment. *Economic and political Weekly*, 979–986.

Chattopadhyay, R. and E. Duflo (2004b). Women as policy makers: Evidence from a randomized policy experiment in india. *Econometrica 72*(5), 1409–1443.

Dahis, R. and C. Szerman (2020). Development via administrative redistricting: Evidence from brazil. *Available at SSRN 3125757*.

Donovan, K. and T. Schoellman (2021). The role of labour market frictions in structural transformation. Technical report, Working paper, Yale University.

Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica 70*(5), 1741–1779.

Faguet, J.-P. (1999). *Does decentralization increase responsiveness to local needs? Evidence from Bolivia*. The World Bank.

Flèche, S. (2021). The welfare consequences of centralization: Evidence from a quasi-natural experiment in switzerland. *Review of Economics and Statistics 103*(4), 621–635.

Foster, A. D. and M. R. Rosenzweig (2004). Democratization and the distribution of local public goods in a poor rural economy.

Gadenne, L. and M. Singhal (2014). Decentralization in developing economies. *Annu. Rev. Econ. 6*(1), 581–604.

Galasso, E. and M. Ravallion (2005). Decentralized targeting of an antipoverty program. *Journal of Public economics 89*(4), 705–727.

Glaeser, E. L., R. La Porta, F. Lopez-de Silanes, and A. Shleifer (2004). Do institutions cause growth? *Journal of economic Growth 9*(3), 271–303.

Gollin, D., D. Lagakos, and M. E. Waugh (2014). The agricultural productivity gap. *The Quarterly Journal of Economics 129*(2), 939–993.

Harris, J. R. and M. P. Todaro (1970). Migration, unemployment and development: a two-sector analysis. *The American economic review*, 126–142.

Hsieh, C.-T. and E. Moretti (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics 11*(2), 1–39.

Lagakos, D. (2020). Urban-rural gaps in the developing world: Does internal migration offer opportunities? *Journal of Economic perspectives 34*(3), 174–92.

Lagakos, D., A. M. Mobarak, and M. E. Waugh (2018). The welfare effects of encouraging rural-urban migration. Technical report, National Bureau of Economic Research.

Lewis, W. A. et al. (1954). Economic development with unlimited supplies of labour.

Meiyappan, P., P. S. Roy, A. Soliman, T. Li, P. Mondal, S. Wang, and A. K. Jain (2018). India village-level geospatial socio-economic data set: 1991, 2001 [data set].

Mookherjee, D. (2015). Political decentralization. *economics 7*(1), 231–249.

Mookherjee, D., P. Bardhan, et al. (2005). Decentralization, corruption and government accountability: An overview. *International Handbook on the Economics of Corruption*.

Morten, M. (2019). Temporary migration and endogenous risk sharing in village india. *Journal of Political Economy 127*(1), 1–46.

Munshi, K. and M. Rosenzweig (2016). Networks and misallocation: Insurance, migration, and the rural-urban wage gap. *American Economic Review 106*(1), 46–98.

Oates, W. E. et al. (1972). Fiscal federalism. *Books*.

Redding, S. J. and E. Rossi-Hansberg (2017). Quantitative spatial economics. *Annual Review of Economics 9*, 21–58.

Schultz, T. W. (1960). Capital formation by education. *Journal of political economy 68*(6), 571–583.

Zhang, Y. and L. Gong (2005). The fenshuizhi reform, fiscal decentralization, and economic growth in china. *CHINA ECONOMIC QUARTERLY-BEIJING- 5*(1), 75.

A Areas of Decentralization

Table A.1: The 29 subjects under the 11^{th} Schedule

11^{th} **Schedule**		
• Agriculture, including agricultural extension • Land improvement, implementation of land reforms, land consolidation and soil conservation • Minor irrigation, water management and watershed development • Animal husbandry, dairying and poultry • Fisheries • Social forestry and farm forestry • Minor forest produce • Small scale industries, including food processing industries • Khadi, village and cottage industries • Rural housing	• Drinking water • Fuel and fodder • Roads, culverts, bridges, ferries, waterways and other means of communication • Rural electrification, including distribution of electricity • Non-conventional energy sources • Poverty alleviation program • Education, including primary and secondary schools • Technical training and vocational education • Adult and non-formal education • Libraries	• Cultural activities • Market and fairs • Health and sanitation, including hospitals, primary health centers and dispensaries • Family welfare • Women and child development • Social, welfare, including welfare of the handicapped and mentally retarded • Welfare of the weaker sections, and in particular, of the Scheduled Castes and the Scheduled Tribes • Public Distribution system • Maintenance of community asset

B First-Stage

This section presents the first-stage results for the district-level instrument. I instrument for the independent variable defined as the share of villages within a district who have their own village council, restricting the subset of villages to those whose population is within 300 of the cutoff. The instrument is also defined within a population bandwidth, and calculates the share of villages above the cutoff within a district with the same population restrictions.

Table B.1: First-Stage: District Instrument

	(1)
	Treatment Share
Cutoff Share	1.223***
	(0.316)
Observations	42
R-squared	0.316
F-Stat	14.54

Notes: This table displays the first-stage for the aggregated-version of the RDD instrument. I regress the share of treated villages within a district on the share of villages above the cutoff within the district, restricting the villages used in my calculation to 300 above and below the population cutoff. The instrument and the independent variable are calculated on a restricted bandwidth to increase the strength of the instrument.

As can be seen from Table B.1, there is a positive and significant first-stage. The F-state is almost 15, which is strong given the small sample size of only 42 observations.

C Robustness Check

This section presents the results of the fuzzy RDD allowing for different bandwidths.

	(1)	(2)	(3)	(4)
Population Change (1991-2011)	BW: 100	BW: 200	BW: 300	BW: 400
Own Village Council	13.07	4.619	7.070	4.018
	(12.68)	(6.425)	(5.233)	(4.517)
Observations	7,928	16,167	24,707	33,660
Control Mean	49.45	50.16	50.38	50.71
FE	Y	Y	Y	Y
Clusters	Y	Y	Y	Y

Population Change (2001-2011)	(1) BW: 100	(2) BW: 200	(3) BW: 300	(4) BW: 400
Own Village Council	14.45**	6.141*	8.324***	9.350***
	(6.758)	(3.134)	(2.761)	(2.443)
Observations	7,928	16,167	24,707	33,660
Control Mean	20.47	20.80	20.80	21.02
FE	Y	Y	Y	Y
Clusters	Y	Y	Y	Y

Population Change (1991-2001)	(1) BW: 100	(2) BW: 200	(3) BW: 300	(4) BW: 400
Own Village Council	0.784	-1.761	-0.569	-3.814
	(7.845)	(3.854)	(3.305)	(2.622)
Observations	7,928	16,167	24,707	33,660
Control Mean	25.40	25.79	26.04	26.22
FE	Y	Y	Y	Y
Clusters	Y	Y	Y	Y

The results presented here, do not significantly differ from the optimal bandwidth results presented in the text.

D Housing Census

In this section I present the full set of results for the housing material data, including all available materials. Housing material is defined as material used in buildings within the village for either a roof or a wall. For roofing material (as seen in Table D.1), there are six potential types of material used, with brick and grass being the most common. Grass is a traditional material, and comes in the form of thatch.

Table D.1: Roofing Material

	(1) Grass	(2) Tiles	(3) Brick	(4) Stone	(5) Metal	(6) Concrete
Own Village Council	-7.560***	1.645	9.222***	3.288**	-1.609***	-4.135***
	(2.239)	(2.096)	(2.516)	(1.539)	(0.588)	(1.602)
Observations	94,856	94,856	94,856	94,856	94,856	94,856
Bandwidth	328.5	349.5	370	397.6	304.8	386.2
FE	Y	Y	Y	Y	Y	Y
Clusters	Y	Y	Y	Y	Y	Y
Control Mean	25.53	12.45	32.06	11.27	3.149	14.03

Standard errors in parentheses

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Notes: This table shows the effect of having a GP on the proportion of houses using a specific type of roofing material. The data is broken down by proportion, and included here are some of the main outcomes. As with all specifications, I control for district FE, and cluster my standard errors at the district-level. For a full breakdown of the distribution of the different materials in treated and control villages, see Appendix D

The results for the wall material are presented in Table D.2. There are three potential types of material, grass, mud, and brick, with grass and mud being the traditionally-used materials.

Table D.2: Wall Material

	(1)	(2)	(3)
	Grass	Mud	Brick
Own Village Council	-2.738**	-4.291**	8.278***
	(1.197)	(1.954)	(2.134)
Observations	94,856	94,856	94,856
Bandwidth	289.5	374.2	381.5
FE	Y	Y	Y
Clusters	Y	Y	Y
Control Mean	6.335	26.66	61.38

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows the effect of having a GP on the proportion of houses using specific types of wall material. Like the roofing material, the data is broken down by proportion, and the full breakdown by treatment and control villages is available in Appendix D. As always, the specification is run with district-level FE and the standard errors are clustered at the district-level.

Banking access, and more important banking usage, gives us a good indicator of individual involvement in the formal financial sector. It is also a good indicator of people's ability to access banking services. Table D.3 shows the effect of having an own village council on banking access.

Table D.3: Banking Access

	(1)
	HH availing banking facilities
Own Village Council	5.492**
	(2.683)
Observations	94,856
Bandwidth	305.7
FE	Y
Clusters	Y
Control Mean	78.40

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows the impact of having a GP on the number of households availing banking facilities. The dependent variable is measured in terms of the number of households. As always, the specification is run with district-level FE and the standard errors are clustered at the district-level.

An alternative proxy for measuring infrastructure access, is lighting, as to utilize electricity, houses need to be connected to the electric grid. Table D.4, shows that treated villages in general use more electricity, although the most-used source remains kerosene.

Table D.4: Lighting Sources

	(1)	(2)	(3)	(4)	(5)	(6)
	Electricity	Kerosene	Solar	Other Oil	Other	None
Own Village Council	7.934***	-6.445**	-0.252	-0.592*	-0.00476	-0.439*
	(2.630)	(2.575)	(0.227)	(0.311)	(0.125)	(0.259)
Observations	94,856	94,856	94,856	94,856	94,856	94,856
Bandwidth	317.7	324.7	362.6	326.3	412.6	337.3
FE	Y	Y	Y	Y	Y	Y
Clusters	Y	Y	Y	Y	Y	Y
Control Mean	23.84	74.97	0.579	0.295	0.193	0.123

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows the effect of having a GP on the proportion of houses using a certain type of lighting. The full breakdown by treatment and control villages is available in Appendix D. As always, the specification is run with district-level FE and the standard errors are clustered at the district-level.

Finally, Table D.5 shows the results on the type of housing structures found in the village. In line with the findings on the housing material, there are more permanent houses in treated villages. This is to be expected however, as housing status is measured based on the type of materials used for the roof and wall.

Table D.5: Housing Status

	(1)	(2)	(3)
	Permanent	Semi-permanent	Temporary
Own Village Council	7.769***	-0.859	-6.612***
	(2.020)	(1.879)	(2.017)
Observations	94,856	94,856	94,856
Bandwidth	368.6	370.9	336
FE	Y	Y	Y
Clusters	Y	Y	Y
Control Mean	58.14	19.29	19.82

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows the prevalence of different types of housing structures. The households within the village are broken down by: Permanent, Semi-Permanent, Temporary, Serviceable, Non-Serviceable, and Unclassifiable. The category a house is assigned to is dependent on the type of material that is utilized in the roof and wall [14]. As always, the specification is run with district-level FE and the standard errors are clustered at the district-level.

E Population Change by Gender

This section presents the results on the rural population change at the village-level, broken down by gender. Table E.1, shows that population change was significant higher in treated villages between 2001-2011, which matches the total effects we saw in the main text. The magnitude of the coefficient is also relatively similar.

Table E.1: Population Change - Male

Population Change (male)	(1) 1991-2011	(2) 2001-2011	(3) 1991-2001
Robust	7.488	9.728***	-2.014
	(7.305)	(3.665)	(4.556)
Observations	95,754	95,754	95,753
Effective Observations	9380	11,263	10,351
Control Mean	47.15	20.19	24.13
FE	Y	Y	Y
Clusters	Y	Y	Y
Bandwidth	206.6	243	225.8

Notes: This table shows the impact of own village council status on population change in percentage for males. Column (1) shows the change over the entire period of interest, 1991-2011, column (2) shows the change between 2001-2011, and column (3) shows the population change between 1991-2001. I wincorize all the population data to ensure that there are no errors affecting the results. To wincorize I take the bottom and top 1% value and replace any observations outside those with the 1% value. All the regressions are run with district-level fixed effects, and district-level clustered standard errors. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Females, somewhat surprisingly, show a much higher population change in treated villages relative to both males and the total, with the population growth over the entire period being significantly higher than males and the total.

Table E.2: Population Change - Female

Population Change (female)	(1) 1991-2011	(2) 2001-2011	(3) 1991-2001
Robust	17.45**	12.32***	4.355
	(7.356)	(4.602)	(4.958)
Observations	94,736	94,505	94,838
Effective Obsservations	9009	8236	8188
Control Mean	52.42	21.37	26.68
FE	Y	Y	Y
Clusters	Y	Y	Y
Bandwidth	199.9	184.9	183.2

Notes: This table shows the impact of own village council status on population change in percentage for females. Column (1) shows the change over the entire period of interest, 1991-2011, column (2) shows the change between 2001-2011, and column (3) shows the population change between 1991-2001. I wincorize all the population data to ensure that there are no errors affecting the results. To wincorize I take the bottom and top 1% value and replace any observations outside those with the 1% value. All the regressions are run with district-level fixed effects, and district-level clustered standard errors. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Taken together, these results suggest that the total effects we observe are not driven purely by movement of a single gender, although females grow at a much higher rate. These findings also imply that the population movements we observe are not just driven my young males migrating for economic reasons.

F Fertility Levels

As a check to make sure that the population changes are not driven by changes in fertility, I also present below the results on population of children aged between 0-6, and show both the changes and the levels. Table F.1 shows that there is no significant difference between treated and control villages in the growth of children aged 0-6, whilst Table F.2 shows that the same holds true for the levels regression.

Table F.1: Fertility

	(1)	(2)	(3)
Population Change (Age 0-6)	1991-2011	2001-2011	1991-2001
Own Village Council	6.656	5.496	3.855
	(6.916)	(4.823)	(6.059)
Effective Obs	8345	9036	8345
Bandwidth	186.8	200.9	186.9
FE	Y	Y	Y
Clusters	Y	Y	Y

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows the impact of having an own GP on changes in fertility rates over the same periods as the previous table. As before, I include district FE and district-level clustered standard errors. I also run the same specification on the levels, and find no significant effects.

Table F.2: Fertility

Population Aged 0-6 (levels)	(1) 1991	(2) 2001	(3) 2011
Robust	1.539	-5.464	13.28
	(4.573)	(10.07)	(12.46)
Observations	95,754	94,960	94,857
Effective Observations	12,579	16,038	13,101
Bandwidth	265.3	325.2	275.9
Control Mean	193.9	225.7	220.6
FE	Y	Y	Y
Clusters	Y	Y	Y

Notes: This table shows the impact of having an own GP on fertility rates in levels over different periods. As before, I include district FE and district-level clustered standard errors.

G Spillovers

Presented below are the robustness tests to ensure that the effects on durable consumption are not being driven by population spillovers. I test this by controlling for the share of villages in the surrounding area that have a population above the cutoff, and reproduce the main RDD results. Column (1) in each table reproduces the RDD without the control.

Table G.1: Roof Material - Brick

	Roofing Material - Brick			
	(1)	(2)	(3)	(4)
	No Controls	10km Buffer	20km Buffer	50km Buffer
Own Village Council	9.222***	7.700**	8.732***	10.15***
	(2.516)	(3.104)	(2.970)	(3.017)
Bandwidth	370	308.5	317.6	314.6
Control Mean	32.06	31.99	31.99	31.99

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows the evolution of the RDD coefficient for the brick roof material including the share of villages above the cutoff in the surrounding areas. I restrict the calculation of the share to villages that fall within a specific population bandwidth - 300 - and for 3 distances, 10km, 20km, and 50km.

Table G.2: Roof Material - Grass

	Roofing Material - Grass			
	(1)	(2)	(3)	(4)
	No Controls	10km Buffer	20km Buffer	50km Buffer
Own Village Council	-7.560***	-9.431***	-9.137***	-9.214***
	(2.239)	(2.453)	(2.402)	(2.426)
Bandwidth	328.5	362.2	362	345.3
Control Mean	25.53	25.61	25.61	25.61

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows the evolution of the RDD coefficient for the grass roof material including the share of villages above the cutoff in the surrounding areas. I restrict the calculation of the share to villages that fall within a specific population bandwidth - 300 - and for 3 distances, 10km, 20km, and 50km.

Table G.3: Wall Material - Brick

	Wall Material - Brick			
	(1)	(2)	(3)	(4)
	No Controls	10km Buffer	20km Buffer	50km Buffer
Own Village Council	8.278***	8.394***	8.767***	9.394***
	(2.134)	(2.609)	(2.538)	(2.481)
Bandwidth	381.5	332.5	334.4	355.8
Control Mean	61.38	61.53	61.53	61.53

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows the evolution of the RDD coefficient for the brick wall material including the share of villages above the cutoff in the surrounding areas. I restrict the calculation of the share to villages that fall within a specific population bandwidth - 300 - and for 3 distances, 10km, 20km, and 50km.

Table G.4: Wall Material - Grass

	Wall Material - Grass			
	(1)	(2)	(3)	(4)
	No Controls	10km Buffer	20km Buffer	50km Buffer
Own Village Council	-2.738**	-3.897***	-3.962***	-4.234***
	(1.197)	(1.205)	(1.186)	(1.202)
Bandwidth	289.5	281.2	276.5	264.9
Control Mean	6.335	6.360	6.360	6.360

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows the evolution of the RDD coefficient for the grass wall material including the share of villages above the cutoff in the surrounding areas. I restrict the calculation of the share to villages that fall within a specific population bandwidth - 300 - and for 3 distances, 10km, 20km, and 50km.

H Population Spillovers

Table H.2 shows the effects of controlling for the share of villages above the cutoff in the surrounding areas. As can be seen from the coefficients, the magnitude doesn't change significantly with the inclusion of these controls, nor are the coefficients significantly different from each other.

Table H.1: Population Change - Controlling for surrounding villages

	(1)	(2)	(3)
Buffer Zones	10km	20km	50km
Own Village Council	8.889*	8.070*	8.135*
	(4.902)	(4.560)	(4.479)
Observations	95,317	95,317	95,317
Bandwidth	258.4	285.4	286.6
Control Mean	25.43	25.43	25.43

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

I also look at the interaction between the surrounding share of villages and the treatment status, and I find that the coefficient on the interaction is negative as we would expect. I run this as a fuzzy RDD, controlling for the bandwidth and instrumenting for treatment status with the cutoff dummy.

Table H.2: Population Change - Controlling for surrounding villages (Interacted)

	(1)	(2)	(3)
Population Change (2011-2001)	10km	20km	50km
Surrounding share of villages	-4.889***	-9.212***	-12.33***
	(1.419)	(1.820)	(2.575)
Cutoff Dummy	5.205***	3.490***	1.508
	(0.988)	(1.142)	(1.486)
Interaction	-8.361***	-3.928	1.492
	(2.344)	(2.872)	(3.929)
Population (1991)	-0.00149**	-0.00138*	-0.00147**
	(0.000727)	(0.000728)	(0.000727)
Observations	42,355	42,355	42,355

Robust standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

I GDP

This section shows the full set of GDP results, and presents some of the robustness checks. Figure I.1 reproduces the GDP per capita results, controlling for total consumption in 1991 at the district-level.

The following figures show the results of the four key GDP outcomes (GDP per capita and Primary, Secondary and Tertiary Sector GDP (in current prices)) for all available years.

Figure I.1: GDP per Capita - Controlling for Total Consumption (1991)

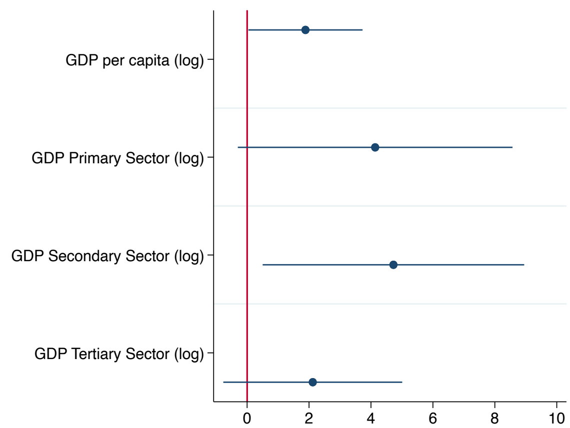

Notes: This figure shows the district-level regressions on GDP per capita, controlling for total consumption at the district-level in 1991. The regressions instrument for treatment status with the cutoff dummy. This figure restricts the cutoff aggregations to using only villages with a population that is within three hundred of the cutoff.

Figure I.2: GDP per Capita (all years)

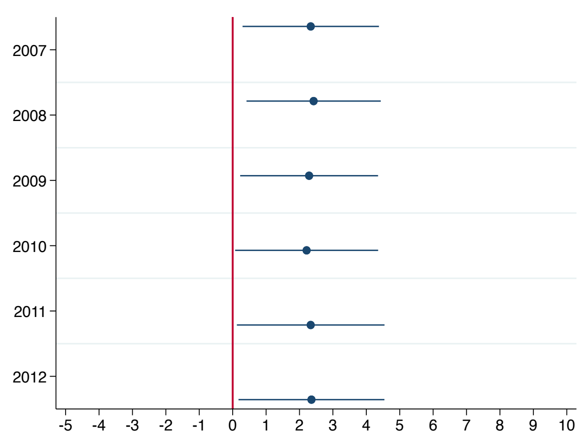

Notes: This figure shows the district-level regressions on GDP per capita, controlling for total consumption at the district-level in 1991. The regressions instrument for treatment status with the cutoff dummy. This figure restricts the cutoff aggregations to using only villages with a population that is within three hundred of the cutoff.

Figure I.3: Primary Sector GDP (all years)

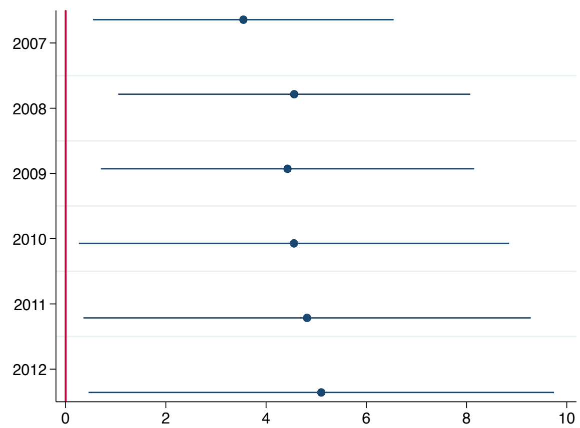

Notes: This figure shows the district-level regressions on Primary Sector GDP, controlling for total consumption at the district-level in 1991. The regressions instrument for treatment status with the cutoff dummy. This figure restricts the cutoff aggregations to using only villages with a population that is within three hundred of the cutoff.

Figure I.4: Secondary Sector GDP (all years)

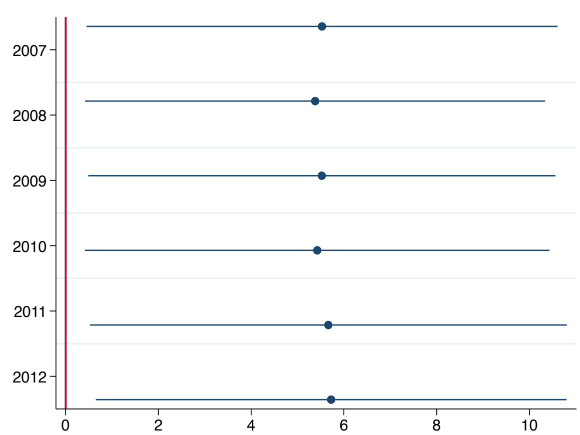

Notes: This figure shows the district-level regressions on Secondary Sector GDP, controlling for total consumption at the district-level in 1991. The regressions instrument for treatment status with the cutoff dummy. This figure restricts the cutoff aggregations to using only villages with a population that is within three hundred of the cutoff.

Figure I.5: Tertiary Sector GDP (all years)

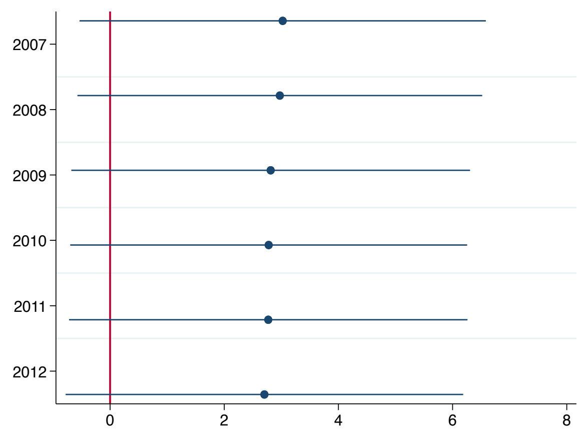

Notes: This figure shows the district-level regressions on Tertiary Sector GDP, controlling for total consumption at the district-level in 1991. The regressions instrument for treatment status with the cutoff dummy. This figure restricts the cutoff aggregations to using only villages with a population that is within three hundred of the cutoff.