Banking Beyond the Branch: How does increased bank presence affect the gender literacy gap in India?

EC331 Quantitative Thesis
Word Count: 9734

Abstract

This paper investigates the causal impact of increased banking access on gender literacy gaps in India, exploiting the exogenous variation induced by the 2005 Reserve Bank of India (RBI) bank expansion policy. Employing a fuzzy regression discontinuity design on 2011 census data across Indian districts, the analysis reveals no statistically significant average effect of increased bank presence on the gender literacy gap. However, subgroup analyses demonstrate heterogeneity: greater banking access significantly reduces the gender literacy gap in districts with more balanced child sex ratios, indicating that increased financial access for girls' education is more effective in contexts with more gender-equal norms. The evidence in areas of varying female labour force participation and agricultural dependency is less conclusive, offering limited support for the role model effect or increased female agency in the household as significant pathways of causation. Overall, the findings suggest that the impact of financial access on reducing gender literacy gaps is contingent upon a supportive cultural foundation characterised by more gender-equal norms, implying that financial inclusion alone is insufficient to overcome deeply entrenched gender biases.

1 Introduction

Financial inclusion is a well established driver of economic growth, development and increased economic opportunity in developing countries such as India [Burgess and Pande, 2005, Bruhn and Love, 2009]. However, the areas which benefit from these initiatives often lack progressive gender views and equality [Jayachandran, 2015]. Despite the economic advantages, they often fail to offer meaningful support or advancement for women, as men typically serve as breadwinners and control household income and finances. Many women are financially dependent on the men of their household, and prefer storing wealth in tangible assets like gold, and hence remain unengaged with formal banking institutions [Garg and Goyal, 2018].

I intentionally choose to focus on the gender literacy gap as a proxy for gender progress. Whilst immediate benefits may be limited for current financially dependent women, transformative potential does lie in their children's ability to break the cycle of dependency [Duflo, 2012]. Entrenched gender barriers and inequalities require longer-term, inter-generational change, and education serves as a critical marker of sustainable progress. This allows us to truly understand whether financial inclusion initiatives can help catalyse lasting social change through a new generation of educated and financially independent women.

Several key channels demonstrate how banking services can positively impact women's education. First, credit access and savings facilities provide consumption smoothing and financial planning capabilities, helping families manage the costs of girls' education, an area often de-prioritised when resources are scarce. Women's exposure to banking services could improve financial autonomy and increase their decision-making power within households. Visibility of banks creates role model effects as children observe women in professional financial roles. Finally, bank branches create formal employment opportunities that demonstrate tangible returns to female education. I explore these pathways more thoroughly in the mechanisms section.

To identify causal effects, I exploit the 2005 RBI bank expansion policy, which grants banks licenses in more favourable locations conditional on opening branches in 'underbanked' districts. Using a fuzzy regression discontinuity design, I examine the effect of being eligible for the bank expansion program in 2005 on gender literacy gaps measured in the 2011 census across Indian districts. This approach allows me to overcome potential endogeneity concerns by leveraging the quasi-random variation around the population-to-branch ratio cut-off that determined underbanked status.

I find that there is no direct effect of increased bank presence on the gender literacy gap in the full sample. However, when examining different subgroups of districts that vary in key characteristics like female labour force composition, agricultural dependency and child sex ratios, the results reveal heterogeneity. Districts with more balanced child sex ratios show statistically significant reductions in the gender literacy gap when exposed to increased banking access, reducing absolute gender literacy gaps by 0.639 percentage points and relative gaps compared to 2001 by 4.24 percent. Districts with skewed ratios show no improvement. Districts with high female labour force participation show directionally consistent decreases in the gender literacy gap compared to districts with low female participation, though these differences don't reach significance. The relationship with agricultural dependency follows a similar pattern but with less consistency, as non-agricultural districts show slightly better outcomes, but overall null.

By identifying the specific characteristics that enable women to benefit from financial inclusion initiatives, my research highlights how cultural context and existing progressive attitudes serve as critical prerequisites for effective policy implementation. The significant heterogeneity in banking access effects across districts demonstrates that financial inclusion initiatives alone cannot overcome entrenched gender biases. Rather, these initiatives require fertile social ground, embodied in more balanced child sex ratios, to effectively reduce educational gender disparities. This finding challenges the conventional assumption that economic interventions naturally lead to social progress through pathways such as increased female economic agency and role model effects as I predicted. It suggests instead that cultural transformation may need to precede or accompany financial inclusion efforts, whose benefits to female literacy are likely seen through improved financial access and labour market incentives. Without this cultural foundation, financial inclusion may fail to benefit women or potentially reinforce existing disparities by disproportionately benefiting males in such communities with strong son preference.

2 Literature Review

Existing research on the impact of bank branch expansion policies in India focuses heavily on the effects on poverty and welfare. The discourse starts with Burgess and Pande [2005], establishing the foundational finding that branch expansion into previously unbanked rural locations significantly reduced rural poverty in India, through increased deposit mobilisation and credit disbursement. They estimated that opening a bank branch in an additional rural location per 100,000 persons lowers aggregate poverty by 4.10 percentage points, and that that agricultural wages, an important income source for the rural poor, increased with branch expansion. The approach exploited the "1:4 branch licensing policy," as a source of exogenous variation, introduced in 1977 and in place until 1990, as a source of exogenous variation. This policy mandated banks open four branches in eligible unbanked locations for every one branch opened in a banked location. key identifying assumption is that this policy caused more rapid rural branch expansion in financially less developed states (which had more unbanked locations) between 1977 and 1990, with the reverse being true before 1977 and after 1990.

However shortly after this contribution, Kochar [2005] and Panagariya [2006] offer some critiques of Burgess and Pande [2005] methodology. Panagariya argues that that the actual branch expansion during 1977-90 exceeded the prescribed ratio, suggesting it was likely driven by the government and RBI's active efforts within wider anti-poverty programs like the Integrated Rural Development Program (IRDP) and the Bank Licensing Program (BLP), rather than the ratio rule. A lot of Panagariya's critiques build upon and agree with Kochar, however Kochar offers a new finding that the expansion of the bank network in rural India during the 1980s significantly increased consumption inequality. Both Kochar and Panagariya emphasise that the branch-expansion program was designed to go hand-in-hand with the IRDP program, making it "almost impossible to separate the effect of the expansion of the banking infrastructure on outcomes such as poverty from that of the Government's anti-poverty programs"[Kochar, 2005]. This highlights the need to understand the wider policy context of such programs to ensure the effects are truly from the policy of interest.

Outside of India, Bruhn and Love [2009] explore explores the role of access and to finance and banks on poverty alleviation in Mexico, specifically providing evidence on the labour market channels this effect occurs. Exploiting the simultaneous opening of 80 banks in Mexico in 2002, they find, in support of Burgess and Pande [2005] results, a 7 percent increase in income levels over two years. Subsequently, Dupas et al. [2018] performed a randomised control trial across Uganda, Malawi and Chile (all different stages of development) to test the impact on expanding access to basic banks account on savings and welfare. Their main finding was that the expanded access was unlikely to noticeably improve welfare on average. Despite offering accounts for free with fee waivers, take-up and active usage were relatively low, particularly in Chile. The constraints varied by context in Malawi and Uganda, insufficient income primarily limited savings capacity, while in Chile, being unbanked often reflected a rational choice given the country's robust social safety nets and consequently reduced need for individual saving. The self-reported barriers also included needing to withdraw shortly, not accumulating enough for a trip, the illiquidity of bank accounts, and distance to the bank branch. The fact that distance to the bank branch was a self-reported driver of the low take up strongly suggests that bank branch presence is an important part of usage, and branch expansion policies likely improves bank usage, offering reassurance that bank presence is a worthwhile indicator of financial inclusion in my research question.

Only two other papers use a similar methodological approach to me in my knowledge, are Young [2017] and Cramer [2021], who both exploit the 2005 RBI Bank Branch expansion policy and employing a regression discontinuity design, however look at different outcomes to me. Young focuses on the increase in private sector banks, showing that each additional private bank branch increased local GDP by approximately 0.33 percent and treated districts experiencing about 1.65 percent GDP growth based on average branch increases. Cramer examines health outcomes, finding that households in districts treated by the policy experienced significantly improved health indicators six years after implementation.

From my understanding, there is a gap in the literature of how bank branch expansion programs impact women differentially to men, largely looking at broad indicators of poverty and welfare. In the Indian context, it is important to understand whether such financial inclusion policies to improve welfare, actually improves the welfare of women as well, considering how Indian society is entrenched with patriarchal values, a lack of female financial independence and household decision making power, low female literacy. [Jayachandran, 2015, Garg and Goyal, 2018]

More recent research has addressed this gap. For instance, Karlan et al. [2014] exploring microfinance interventions have provided some insights into the potential for financial access to empower women, though these often focus on credit rather than the broader impact of traditional banking infrastructure. Furthermore, research on the impact of women's self-help groups (SHGs) in India suggests that collective financial participation can enhance women's agency and decision-making abilities [Armendáriz and Morduch, 2010]. However, the specific pathways through which large-scale bank branch expansion policies influence women's educational attainment and literacy remain largely unexamined. This study aims to contribute to this area by focusing on the gender literacy gap as a key outcome, providing insights into the differential impact of financial inclusion on human capital development for women in India.

3 Conceptual Framework

3.1 Channels and Mechanisms

I primarily examine four channels of causation through which I predict a decrease in the gender literacy gap, or in other words, how increased bank presence would specifically impact and improve female literacy differentially to male literacy. I then talk about how this informs my subgroup selection.

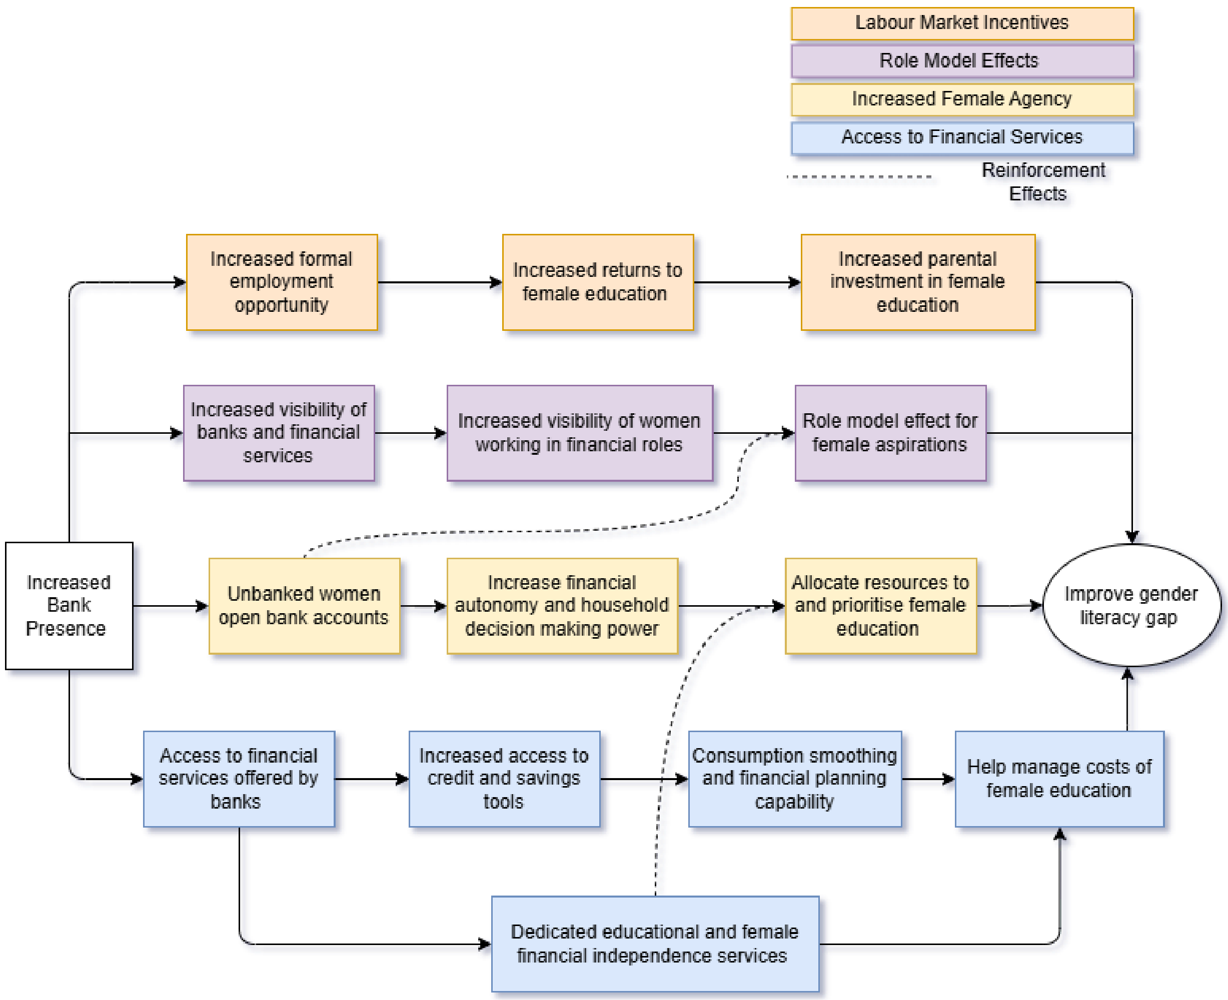

Figure 1: A flowchart demonstrating four main pathways identified to predict how increased bank presence could improve female literacy differentially to male literacy, and reduce the gender literacy gap in the context of India

The '**Access to Financial Services Pathway**' examines how increased bank branch presence enhances access to financial services. Banks primarily function as savings institutions that facilitate interest accumulation and secure wealth storage, while also providing access to credit through business and personal loans. These services enable consumption smoothing and enhance household financial planning capabilities. [Friedman, 2018]

This improved financial access effectively expands household budget constraints. When resources are limited, families often face difficult choices in educational investment. For instance, if a household has only enough resources to educate one child, cultural norms and economic considerations typically favour sons over daughters [Duflo, 2012]. However, with access to financial services, families can overcome these temporary constraints through savings, loans, or more efficient resource allocation, enabling investment in both sons' and daughters' education simultaneously [Dupas and

Robinson, 2013].

To formalise this choice, I consider a household utility maximisation problem where parents must allocate limited resources between male and female children's education. Denote:

- E_m = Investment in male child's education
- E_f = Investment in female child's education
- R = Total household resources available for education
- α = Cultural preference parameter (where $\alpha > 1$ indicates son preference)
- $u(\cdot)$ = Utility function (increasing and concave)

Without financial services, the household solves:

$$\max_{E_m, E_f} \alpha \cdot u(E_m) + u(E_f) \tag{1}$$

$$\text{subject to } E_m + E_f \leq R \tag{2}$$

With son preference ($\alpha > 1$) and binding resource constraint, corner solutions often emerge where $E_f = 0$ if R is sufficiently small.

With financial services access:

1. Resource constraint relaxes to R' where $R' > R$ due to:
    - Consumption smoothing allowing reallocation of emergency funds
    - Access to education loans with parameter L
    - Increased savings returns with parameter i

The new constraint becomes:

$$E_m + E_f \leq R' = R + L + i \cdot S \tag{3}$$

where S represents household savings, increasing the likelihood of a solution where $E_f > 0$.

Banks can also provide specialised educational loans, scholarships, and schemes for female education, which is common in India [Karlan et al., 2014]. They also facilitate microcredit and help groups for women, which creates a reinforcement into the female agency pathway. [Prina, 2015]

The **'Increased Female Agency'** pathway examines how increased bank presence is likely to increase the number of women who own bank accounts, especially previously unbanked women. The mere presence of the bank increases curiosity and financial literacy in women, but also provides a close, walkable, and accessible opportunity for women to store wealth [Allen et al., 2016]. Many Indian women traditionally store wealth in jewellery [Garg and Goyal, 2018], so having banks nearby would likely encourage formal savings. Once women start engaging with formal banking and have their own wealth, they gain more financial autonomy in the household [Ashraf et al., 2010] and greater decision-making power [Anderson et al., 2009]. Women who prioritise financial independence and have a say in household resource allocation would likely want to educate their daughters, so they prioritise female education. I expect this effect to be stronger in areas with high female labour participation, where women have their own income to maintain their own accounts.

The increased visibility of banks and financial services also drives the **'Role Model Effect'** [Porter and Serra, 2020]. When children, especially girls, walk past banks or accompany family members to banking facilities, they see women working in professional roles and may be inspired to pursue similar paths, leading them to recognise the value of education and remain in school. This effect may be more pronounced in agriculturally dependent districts where few women work in knowledge or service sector jobs.

The final pathway, **'Labour Market Incentives',** suggests that increased bank branches create formal employment opportunities, especially in more rural districts. This demonstrates to the community that there are tangible returns to female education. I theorise that effect is more likely present in areas with less anti-female gender norms and values. In communities where educating women is not prioritised because there are no visible nearby employment options beyond agricultural or manual labour, women are more likely to stay at home. However, if the option to work is a well-paid, dignified job with a good salary, parents are more likely to recognise the increased returns to female education, boosting investment in girls' schooling. If limited job opportunities are the barrier rather than cultural opposition to women working outside the home, then this pathway would significantly improve outcomes. Otherwise, this pathway might simply improve education for all, not females in particular.

3.2 Subgroup Selection

To help test these mechanisms, I perform a subgroup analysis, dividing my sample into subgroups in order to understand which pathways are likely at play.

- Female labour force participation
- Agricultural Dependency
- Child Sex Ratio

3.2.1 Female Labour Force Participation

$$\text{FLFPR} = \frac{\text{Working Females}}{\text{Total Female Population}} \tag{4}$$

Female labour force participation (FLFP) allows me to investigate how banking access interacts with women's pre-existing economic agency. Districts with high FLFP typically demonstrate greater female economic autonomy, potentially enabling women to redirect newly accessible financial resources toward girls' education when banking services expand [Anderson et al., 2009]. These districts likely have complementary infrastructure and services that simultaneously support female work and educational access [Duflo, 2012]. In contrast, districts with low female labour market participation may face persistent cultural and infrastructural barriers that constrain the effectiveness of financial access alone [Jayachandran, 2015]. High FLFP districts also demonstrate clear economic incentives for investing in girls' education by making the returns visible and tangible. When women actively participate in the economy, households observe concrete examples of how education translates into economic opportunities. This visibility effect is reinforced by intergenerational role model dynamics, where working mothers inspire higher educational aspirations in girls [Porter and Serra, 2020]. These effects create fertile ground for banking access to translate into educational investments, while such motivational factors may be largely absent in communities where women's economic roles remain limited to unpaid household work.

Pathway Prediction: If I find significant effects of banking access on reducing gender literacy gaps in high FLFP districts but not in low FLFP districts, this would provide evidence supporting both the Female Agency pathway (women redirecting resources to daughters' education) and the Role Model pathway (visibility of returns to female education).

3.2.2 Agricultural Dependency

$$\text{Agri Dependency} = \frac{\text{Agricultural Workers (Cultivators + Laborers)}}{\text{Total Workers}} \tag{5}$$

Agricultural dependency helps understanding of how banking impact varies with the underlying economic structure of districts, as this shapes incentives for female education [Foster and Addy, 2010, Jayachandran, 2015]. Non-agricultural economies offer more visible employment opportunities requiring literacy, whereas highly agricultural districts tend to see women performing unpaid household labour [Schultz, 2002]. Consequently, the perceived value of female literacy varies between these settings, with lower incentives in agriculturally dominant districts. Non-agricultural districts typically exhibit greater modernisation, extending to attitudes about female education and roles [Goldin, 1995]. Agricultural districts face seasonal labour demand affecting schooling patterns, with girls often pulled from school during peak seasons [Jacoby, 2000]. This contrasts with non-agricultural districts experiencing more stable labour demand patterns. From an access perspective, agricultural districts often have poorer educational infrastructure and greater distance to schools, potentially reducing policy effectiveness [Muralidharan and Prakash, 2017].

Pathway Prediction: If banking access significantly reduces gender literacy gaps in less agricultural districts but not in highly agricultural areas, this would support the Labour Market Incentives pathway (banking creates visible returns to female education through formal employment) and the Role Model Effect (women in non-agricultural professional roles serving as inspirational examples).

3.2.3 Child Sex Ratio

$$\text{Child Sex Ratio} = \frac{\text{Girls Aged 0-6}}{\text{Boys Aged 0-6}} \tag{6}$$

I use the Child Sex Ratio from the 2001 Population Census as an indicator for cultural gender norms in India. A skewed ratio (fewer girls than boys) strongly suggests daughter aversion, resulting from sex-selective abortions, female infanticide, or differential care [Jha et al., 2011, Sen, 1990]. This ratio often persists regardless of economic development, as even wealthy districts may maintain highly skewed ratios, making it a good measure of gender attitudes less influenced by economic opportunities [Klasen and Wink, 2003]. Areas with more balanced ratios typically show greater investment in girls' education, reflecting that cultural norms affecting child sex ratios likely influence schooling decisions too [Bharadwaj et al., 2020]. Districts with more balanced child sex ratios also tend to have fewer barriers to female labor market participation, as the same cultural factors that lead families to value daughters equally are often associated with more progressive attitudes toward women's economic roles [Jayachandran, 2015, Duflo, 2012].

Pathway Prediction: If banking access significantly reduces gender literacy gaps in districts with more balanced child sex ratios but not in districts with skewed ratios, this would provide evidence for both the Financial Access pathway and the Labour Market Incentives pathway. In such cases, banking expansion likely: (1) helps households overcome financial constraints to educate daughters they already value, and (2) creating employment opportunities that families are culturally willing to let their daughters pursue, increasing the perceived returns to girls' education.

4 Data and Policy

4.1 Policy and Institutional Background

The policy I exploit to simulate exogenous variation in bank presence is the 2005 policy by India's central bank, the Reserve Bank of India (RBI), used to promote financial inclusion. The policy used an incentive to encourage banks to increase the number of branches in under-banked districts. The incentive was that banks had the opportunity to increase their chance of obtaining licenses to set up branches in locations of their preference, as long as they increased their bank presence in the districts classified as 'underbanked'. The RBI released a list of districts that they classified as 'underbanked' in 2006, which were districts with a population-to-branch ratio above the national average ratio. [Reserve Bank of India, 2006]

$$\frac{\text{Population}_{District}}{\#\text{ Bank Branches}_{District}} > \frac{\text{Population}_{National}}{\#\text{ Bank Branches}_{National}} \tag{7}$$

The banking sector in India doesn't not permit free entry of banking firms or branches. New bank licences are granted infrequently [Young, 2017] and banks must also acquire licences prior to opening all new branches, and receive permission to close or shift branches in most markets. Prior to policy, banks applied for each of these changes on a case-by-case basis through the regional office of the RBI. Unlike the 4:1 policy studied in Burgess and Pande [2005], which required branches to be opened strictly in unbanked markets, banks could choose among any markets within under-banked districts to satisfy the rule. The randomisation of districts around the national average cut-off provides the variation in treatment exploited in my analysis.

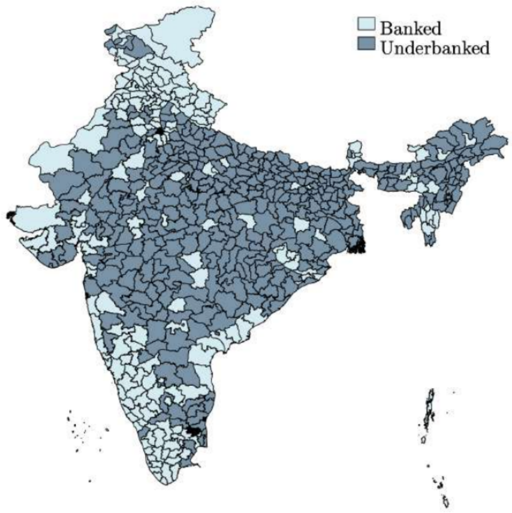

Figure 2: Map of India demonstrating the RBI's underbanked and banked district assignment [Cramer, 2021]

This incentive can be assumed to work most effectively during periods of high demand for bank branches in wealthy areas, which was true for India experiencing a period of rapid growth between 2003 and 2008. [World Bank, 2025]

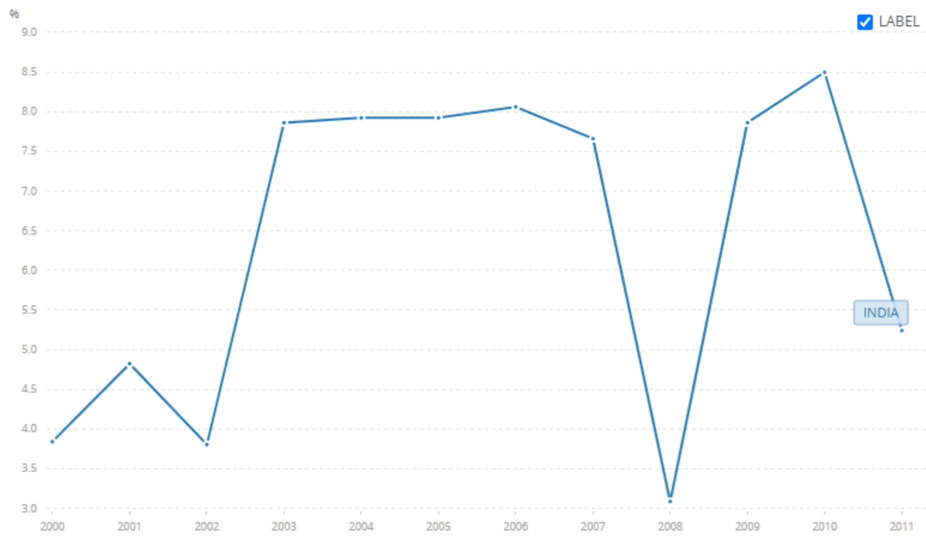

Figure 3: India's GDP growth between 2000 and 2011 [World Bank, 2025]

This then begs the question, did being on the 'underbanked' list have any tangible impact on the actual number of bank branches that opened in that district over the following years? I lean on the existing literature to justify this claim, and Cramer [2021] provides evidence that banks did, in fact, react to this policy. Comparing 2004 (pre-policy) and 2010 (post-policy), they find that treatment districts have on average 21 percent more branch licenses and 19 percent more branches than control districts. Reassuringly, the responses of the banks match up with the timing of the policy, which the reaction in licenses being immediate but the reaction in branches being delayed by roughly a year. This justifies my usage of being on the list as an actual treatment effect, as opposed to an intention-to-treat effect.

I examine several policies that occur around the same time and have particular relevance to my outcome of interest, the gender literacy gap, to contextualise my analysis and also understand what could confound my relationship of interest. The Right to Education Act in 2009 mandated free and compulsory education for all children aged 6–14 . While this likely improved the gender literacy gap, it is important to note that this policy was effective from April 2010 [Maheshwari, 2021], which is the same time data collected for the 2011 Population Census started. This suggests that, due to anticipated lag effects of such policies, it wouldn't confound the gender literacy gap I calculate using the census data. The Mid-Day Meal Scheme in 2004 provided free cooked meals in government primary schools, increasing school attendance and nutrition. There is no evidence suggesting that this affected female educational attendance differentially to male. [Deodhar et al., 2010] The Kasturba Gandhi Balika Vidyalaya scheme launched in 2004 established residential schools for girls from disadvantaged caste groups (SC/ST) in educationally backward areas to improve girls' upper primary school participation. [NITI Aayog, 2016] Since literacy is primarily developed in lower primary school, and because I control for SC/ST populations, KGBV's effects are likely attenuated in my context. The Sarva Shiksha Abhiyan in 2001 was a flagship program to achieve universal elementary education, explicitly targeting gender and social gaps. [Department of School Education and Literacy, 2019] The focus on bridging the gender gap could inflate measured effects of the RBI policy on female literacy, further motivating my regression discontinuity design to isolate causal effects around the RBI bank policy cut-off.

4.2 Data and Preparation

The three datasets I use are:

- Indian Population Census of 2001
- Indian Population Census of 2011
- RBI Bank Branches Dataset

I accessed the census data [of India, 2011] from the SHRUG Platform, an open access repository currently comprising dozens of datasets covering India. I also accessed the RBI bank branch data from the SHRUG, which provides information for 154,505 bank branches at the branch level, contributed by Garg and Gupta [2020].

Since the population-to-branch ratios are not publicly available, I chose to reconstruct these ratios from existing data. I use the population data from the 2001 Population Census, and bank branch data from the RBI Bank Branch Dataset. Since the list of underbanked districts was released in July of 2006, I restrict the bank branch dataset till July 2006, to only account for branches that existed prior to the list being released. [1]

One of the key issues I faced is that the numerical district coding differs between the 2001 and 2011 censuses, without any key to match district codes. To overcome this I chose to match by district names. I aggregate the branches by district, merging these counts with 2011 census district identifiers to attach district names. I then processes 2001 population census data, standardising district codes and preserving required variables for controls, subgroups and pre-existing literacy gaps. I integrate the bank branch data with both 2001 and 2011 census figures by matching on lowercase district names. Throughout the process, I implement checks such as: duplicate detection, consistent variables, and careful preservation districts which likely changed names or spelling. Despite best efforts, a loss of the number of districts in my dataset from 575 to 521 (9.4 percent reduction) is likely due to fuzzy matching of district names, which may not account for all possible boundary changes and spelling variations.

For each district, I then calculate the population-to-branch ratios, and the national average. I construct the gender literacy gap from the 2011 census for my outcome, and from the 2001 census for the control:

$$\text{Literacy Gap (\%)} = \left(\frac{\text{Literate Males}}{\text{Total Male Population}} \times 100\right) - \left(\frac{\text{Literate Females}}{\text{Total Female Population}} \times 100\right) \tag{8}$$

As a sense-check of the data quality, I examined the districts and states with the highest and lowest gender literacy gaps. The results align with established findings in the literature. Rajasthan

[1] It is important to note that I have simply deduced the datasets used to calculate this ratio, but I have no way of knowing exactly from what sources the RBI chose to calculate these ratios, and also whether that was the pure input factor into deciding whether the district made the underbanked list.

emerged as the state with the largest gap, consistent with historical patterns of gender inequality [Drèze and Sen, 2013], and Kerala with the lowest gap as expected [Jeffrey, 2016].

In order to validate my reconstruction of the official underbanked district classification, I compare my deisgnation to the list published by the RBI in 2006 [Reserve Bank of India, 2006]. Using my calculated national population-to-branch ratio as my threshold, my method correctly identifies 79.5 percent of districts in their official designation. Whilst this is a moderately high level of agreement, 20.5 percent of districts not matching correctly with the official list casts doubt on the accuracy of my threshold. This could be the case due to the RBI's opacity on exactly what data sources were used to construct the ratios. This difference is likely in the bank branch data, as opposed the the population census data. Secondly, the RBI could have used discretion, such as deciding to include a district despite not reaching the threshold ratio, or vice versa. The substantial agreement isn't a threat to identification, but provides confidence in the validity of fuzzy regression discontinuity approach. I share my first stage results in Section 5 to validate my calculated assignment as an instrument for official assignment.

5 Empirical Methodology

5.1 Fuzzy Regression Discontinuity Design

The challenge in estimating the causal effect of banking access on gender literacy gaps comes from two sources of endogeneity: (1) omitted variable bias from unobserved district characteristics affecting both underbanked status and education, and (2) confounding from concurrent improvements in economic development and social progressiveness (as detailed in the Section 4.1).

To address these concerns, I employ a fuzzy regression discontinuity design that exploits the RBI's population-to-branch ratio as a source of quasi-random variation. This policy rule creates a natural experiment because districts' position relative to the cutoff is as good as random around the threshold Lee and Lemieux [2010]. The design is fuzzy rather than sharp due to imperfect compliance: whilst crossing my calculated threshold increases the probability of being on the official RBI list, not all districts perfectly comply.

I use:

- Actual Treatment: Official RBI list 'underbanked' status
- Instrument: My threshold-based 'underbanked' assignment

I use the official RBI list status as treatment, since Cramer [2021] established its increase in 19 percent more branches than control districts. When we incorrectly apply sharp RD methods to a setting with imperfect compliance, our estimates will be biased toward zero, since we're comparing groups that don't perfectly align with treatment status [Lee and Lemieux, 2010]. The fuzzy RD estimator is essentially an instrumental variables estimator, estimating the local average treatment effect for compliers [Imbens and Lemieux, 2008]. In this context, compliers are the districts which match both my designation and the official designation.

5.2 Identifying Assumptions

5.2.1 First Stage

The existence of a strong first stage validates that the instrument is relevant, a crucial IV assumption that the instrument increases the probability of treatment. [Angrist and Pischke, 2009].

Table 1: First Stage: Effect of calculated assignment on official assignment

	Bandwidth (Thousands of People per Branch)				
	(1)	(2)	(3)	(4)	(5)
	7	10	12	15	20
Eligible	0.355***	0.467***	0.519***	0.543***	0.568***
	(0.065)	(0.054)	(0.049)	(0.045)	(0.042)
Constant	0.463***	0.333***	0.301***	0.277***	0.260***
	(0.048)	(0.037)	(0.034)	(0.032)	(0.030)
Observations	185	267	305	347	384
F-statistic	29.48	74.65	113.43	146.32	184.72

Notes: The table reports first-stage estimates of the effect of crossing the eligibility threshold on actual treatment status. Robust standard errors in parentheses. *** denotes significance at the 1% level.

Crossing the eligibility threshold increases treatment probability by 35.5-56.8 percentage points across bandwidths. All specifications produce F-statistics exceeding the critical threshold of 10 suggested by Stock and Yogo [2005], indicating a strong first stage that avoids weak instrument concerns. The F-statistics range from 29.48 at a bandwidth of 7 to 184.72 at a bandwidth of 20. While the coefficient magnitudes increase with bandwidth, this pattern is consistent with heterogeneous compliance rates at different distances from the threshold. [Imbens and Lemieux, 2008]

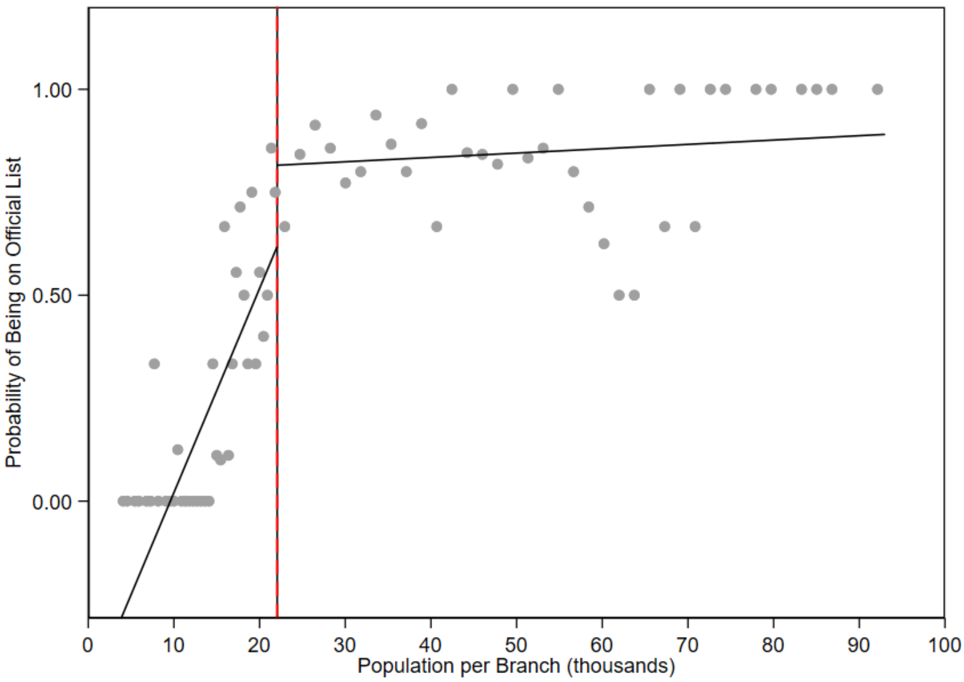

Figure 4: Probability of being on the official list (treatment) against population per branch (in thousands). The vertical red dashed line indicates the eligibility threshold of 22.05025 thousand population per branch.

Figure 4 shows a clear discontinuity at the threshold, supporting the validity of the instrument. Since the districts are grouped into bins, a data point at 0.5 for example, shows that 50 percent of districts in that bin matched the designation. Figure 4 demonstrates large clusters of bins at 0 and 1, demonstrating a concentration of bins in which all districts match the official designation.

5.2.2 Manipulation of the Running Variable

Districts could theoretically manipulate their population-to-branch ratios through selective reporting of population figures or branch counts, especially if they anticipated benefits from being classified as "underbanked." However, manipulation appears implausible for two key reasons. First, the precise threshold was determined after data collection and not announced in advance. Second, the calculation required combining 2001 Census data with banking statistics collected prior to the policy announcement in 2005-2006, making retroactive manipulation practically impossible.

To formally test for manipulation, I implement the Cattaneo et al. [2020b] density test, which examines whether there is an unusual discontinuity in the density of observations around the threshold. The robust test statistic (T = -0.2521) with a p-value of 0.8009 indicates no statistically significant discontinuity in the density at the cut-off. Additionally, binomial tests across various window lengths show no significant discontinuities, with p-values ranging from 0.2112 to 1.0000.

Visual inspection of Figure 5 shows that while the density of observations exhibit some natural variation throughout the distribution, there is no distinctive discontinuity or unusual pattern at the threshold value. This visual evidence, combined with the formal test results, strongly suggests that districts did not (and could not) systematically manipulate their position, reinforcing the validity of our identification strategy.

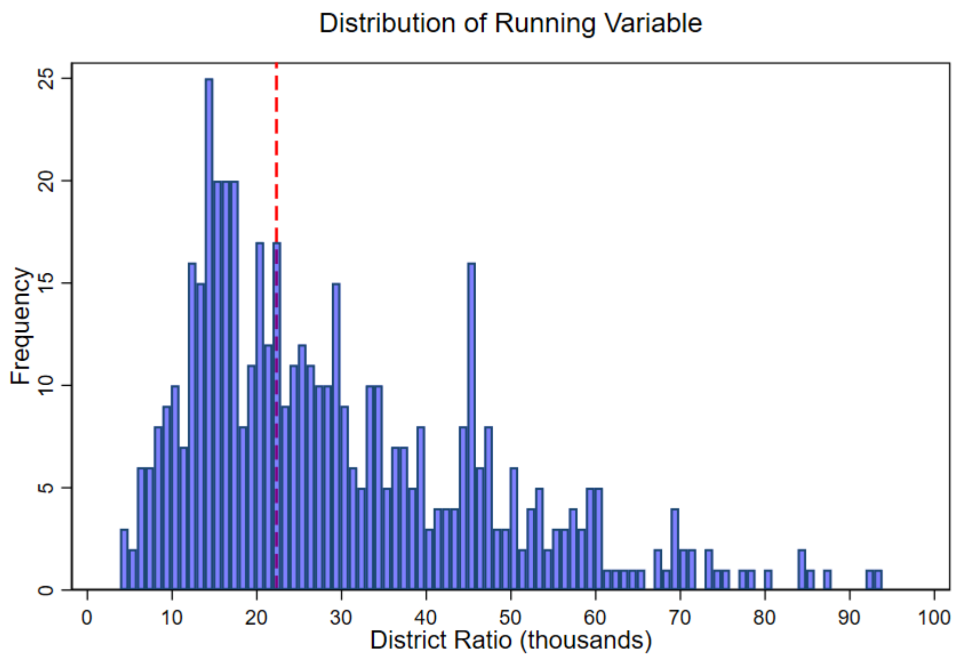

Figure 5: Distribution of the running variable (district ratio) with cutoff at 22.05. The histogram shows no visible discontinuity at the cutoff threshold, suggesting no systematic manipulation of the running variable. Bin width = 1 unit.

5.2.3 Smoothness of Pre-Treatment Covariates

Secondly, I test the assumption of covariate balance at the cutoff, as if the RD design effectively mimics a local randomised experiment, pre-treatment characteristics should exhibit smooth distributions across the threshold with no discontinuities. [Imbens and Lemieux, 2008]

Table 2 presents balance tests for key pre-treatment covariates measured in 2001, before the policy implementation. For each covariate, I estimate the discontinuity at the threshold using a local linear regression with MSE-optimal bandwidths. The results are largely reassuring—most covariates show no statistically significant discontinuities at the 5 percent level. The literacy gap, SC/ST population shares, sex ratios, female labor force participation, and child sex ratio all exhibit smooth distributions across the cutoff. Total population shows a marginally significant difference (p-value = 0.064) and agricultural dependency shows potential marginal significance in the robust specification (p-value = 0.091). However, given the number of covariates tested, finding two marginally significant differences at the 10 percent level is likely chance. The absence of discontinuities in these pre-treatment covariates provides evidence that districts just above and just below the threshold are comparable along observables.

Table 2: Balance Test for Pre-Treatment Covariates at RD Cutoff

Pre-treatment Covariate (2001)	Coefficient	Standard Error	p-value
Literacy gap (%)	-0.891	1.187	0.453
Total population	-700,000	380,000	0.064*
SC population share	-110,000	99,374	0.261
ST population share	-44,898	58,254	0.441
Sex ratio	-0.015	0.011	0.194
Female Labour Force Participation	-0.042	0.029	0.154
Agricultural dependency	-0.048	0.035	0.168†
Child sex ratio	-0.008	0.011	0.430

Notes: Bandwidth selection method is MSE-optimal. * indicates significance at 10% level for conventional estimates. † The robust p-value for agricultural dependency is 0.091, marginal significance at the 10% level.

5.2.4 Placebo Tests

I conduct placebo tests to further validate the regression discontinuity design [Imbens and Lemieux, 2008]. These tests examine whether "false" thresholds—values where no policy discontinuity exists—show significant treatment effects. Table 3 presents results from fuzzy RD estimates at the true cutoff (22.05) and six placebo cutoffs ranging from 18 to 25 thousand population per branch.

The results strongly support the validity of my design. All placebo cut-offs yield statistically insignificant treatment effects with p-values much greater than the accepted 5 percent significance levels [Lee and Lemieux, 2010]. Both conventional and robust p-values across all placebo thresholds exceed 0.13, with most above 0.45, indicating no discontinuities at these arbitrary thresholds.

Notably, the estimate at the true cutoff (22.05) shows a positive coefficient of 3.532, though it remains statistically insignificant (robust p-value = 0.657). The lack of statistical significance at both the true and placebo cut-offs suggests modest effects given the small sample size of 521, but doesn't threat validity.

Table 3: Fuzzy RD Treatment Effect at True Cutoff and Placebo Tests

Cutoff	Coefficient	p-value		95% Robust CI
		Conventional	Robust	
18	0.146	0.922	0.986	(-3.311, 3.371)
20	5.461	0.353	0.450	(-7.851, 17.715)
21	-0.729	0.669	0.835	(-4.103, 3.313)
22.05	**3.532**	**0.629**	**0.657**	**(-12.045, 19.100)**
23	3.512	0.133	0.169	(-1.513, 8.609)
24	0.694	0.931	0.901	(-16.222, 18.431)
25	3.340	0.469	0.553	(-6.881, 12.859)

Notes: All specifications control for baseline literacy gap, population characteristics, and state fixed effects. The treatment effect is estimated using actual treatment status as the endogenous variable. Confidence intervals are based on robust bias-corrected inference.

5.3 Empirical Specification

The baseline specification is:

$$D_i = \gamma + \delta T_i + g(X_i - c) + \nu_i \quad \text{(First Stage)} \tag{9}$$

$$Y_i = \alpha + \tau D_i + f(X_i - c) + \mathbf{Z_i'}\beta + \epsilon_i \quad \text{(Second Stage)} \tag{10}$$

where:

- Y_i is the gender literacy gap in district i;
- T_i is an indicator for eligibility equal to 1 if the district's population per branch ratio exceeds the calculated cutoff;
- D_i is the actual treatment status (whether the district was on the official list);
- X_i is the running variable (population per branch in thousands);
- $f(\cdot)$ and $g(\cdot)$ are flexible polynomial functions of the running variable, allowing for different slopes on either side of the cutoff;

- $\mathbf{Z_i}$ is a vector of district-level covariates including total population, scheduled caste and scheduled tribe (SC/ST) population shares, sex ratio, and state fixed effects;
- ϵ_i and ν_i are error terms.

I examine two different outcomes:

1. **Level Effects:** $Y_i = \text{GenderGap}_{i,2011}$, the female-male literacy gap in 2011 (measured in percentage points). For this specification, I include the baseline gender gap ($\text{GenderGap}_{i,2001}$) as an additional control variable in $\mathbf{Z_i}$. The coefficient τ represents the causal effect of banking expansion on the post-treatment literacy gap, holding constant pre-treatment disparities.
2. **Percentage Change:** The percentage change in the gender literacy gap between 2001 and 2011, defined as:

$$Y_i = \left(\frac{\text{GenderGap}i, 2011 - \text{GenderGap}i, 2001}{\text{GenderGap}_{i,2001}} \right) \times 100 \tag{11}$$

   For this specification, the baseline gap is implicitly accounted for in the outcome variable and is therefore excluded from the controls. Here, τ represents the causal effect on the relative improvement in gender literacy disparities, which is informative for understanding how banking expansion affects the rate of convergence in gender literacy rates.

The parameter τ in both specifications represents the local average treatment effect (LATE) for compliers at the threshold [Angrist and Imbens, 1995], but with different interpretations. In the level specification, it captures the effect on absolute literacy gap differences, while in the percentage change specification, it measures the effect on relative convergence rates.

I estimate the model using local polynomial regression with orders 1, 2, and 3. The optimal bandwidth is selected using the CER-optimal procedure. [Calonico et al., 2014]. Standard errors are clustered at the state level using nearest neighbor variance estimation with 3 neighbors, and all specifications include state fixed effects. These choices involve trade-offs, I provide detailed justification for these specifications in Section 5.5.

5.3.1 Subgroup Analysis

For each heterogeneity dimension, I define binary subgroups based on district characteristics as follows:

- **Female Labor Force Participation (FLFP):** Districts above the median FLFP rate are classified as "high FLFP."
- **Agricultural Dependency:** Districts above the 75th percentile are classified as "high agricultural dependency," focusing on the most agriculture-dominated economies.
- **Child Sex Ratio:** Ratio of female to male children aged 0-6. Districts above the median are classified as "high gender equality," indicating relatively lower gender discrimination.

I calculate the distributional statistics across the full sample of districts before assigning districts to subgroups, to ensure that the subgroup definitions are based on pre-treatment characteristics and are exogenous to the treatment assignment.

I employ three estimation methods for each subgroup in order to try and fully understand the findings across different modelling choices:

Separate Subgroup RD Estimation For each dimension, I divide the sample into binary subgroups (high/low) and estimate separate RD models:

$$Y_i = \alpha_g + \tau_g D_i + f_g(X_i - c) + \gamma_g' Z_i + \epsilon_{i,g} \tag{12}$$

where $g \in \{0, 1\}$ indicates the subgroup, D_i is the treatment indicator, and Z_i includes controls for the other 2 dimensions along with baseline covariates.

The coefficients τ_0 and τ_1 represent the causal effects of banking access for districts with low and high values of the characteristic, respectively. This approach allows for completely different data-generating processes across subgroups. Each subgroup can have its own optimal bandwidth, functional form, and parameter values, accommodating potentially different underlying relationships. The specification is also more flexible in allowing all parameters, not just the treatment effect, to vary by subgroup. However as a limitation, the binary split imposes an arbitrary threshold that simplifies a continuous characteristic, potentially discarding valuable variation. With smaller sample sizes in each subgroup, precision may be reduced. Additionally, if the subgroups have different optimal bandwidths, comparing treatment effects becomes more challenging as they represent effects at different levels of locality around the threshold. [Cattaneo et al., 2024, Lee and Lemieux, 2010, Angrist and Pischke, 2009, Imbens and Kalyanaraman, 2012]

Binary Interaction Model This model incorporates heterogeneity directly into a single regression framework:

$$\begin{aligned} Y_i = \alpha + \beta D_i + \delta G_i + \theta(D_i \times G_i) + \gamma(X_i - c) + \lambda[D_i \times (X_i - c)] \\ + \phi[G_i \times (X_i - c)] + \psi[D_i \times G_i \times (X_i - c)] + \gamma' Z_i + \epsilon_i \end{aligned} \tag{13}$$

where G_i is a binary indicator equal to 1 if the district has a 'high' value of the heterogeneity dimension. The coefficient β represents the treatment effect for districts with low values of the characteristic ($G_i = 0$), while $\beta + \theta$ gives the effect for districts with high values ($G_i = 1$). The interaction coefficient θ directly measures the differential effect between subgroups. The coefficient δ captures level differences between high and low characteristic districts. This provides a formal statistical test of heterogeneity in effects between high and low through the significance of θ. It maintains the full sample for estimation, potentially improving precision over the separate subgroup approach. Using a single model makes it easier to control for multiple dimensions of heterogeneity simultaneously and uses a consistent bandwidth across all observations, making treatment effect comparisons more intuitive. The binary split still splits continuous characteristics, losing potentially valuable variation. The model imposes parameter restrictions by assuming that only the intercept and slope of the treatment effect can vary across subgroups, while other aspects of the functional form remain constant. This may be too restrictive if the subgroups truly have different data-generating processes. [Brambor et al., 2006, Athey and Imbens, 2016, Carneiro and Lee, 2009, Imbens and Kalyanaraman, 2012]

Continuous Interaction Model This approach leverages the full variation in the subgroup:

$$Y_i = \alpha + \beta D_i + \delta G_i + \theta(D_i \times G_i) + \gamma(X_i - c) + \lambda[D_i \times (X_i - c)] \\ + \phi[G_i \times (X_i - c)] + \psi[D_i \times G_i \times (X_i - c)] + \gamma' Z_i + \epsilon_i \tag{14}$$

where G_i is now the continuous measure of the heterogeneity dimension, centred at its mean.

The coefficient β represents the treatment effect for a district with the average value of the characteristic. The interaction coefficient θ indicates how the treatment effect changes with a one-unit increase in the subgroup characteristic. The marginal effect of treatment for a district with characteristic value G_i is $\beta + \theta G_i$. This utilises the full variation in the subgroup dimension rather than imposing an binary cut-off, capturing more nuanced patterns of effect heterogeneity. Centering the continuous variable at its mean makes β the effect at the average value of the characteristic. Some limitation may be that the model imposes a linear relationship between the heterogeneity dimension and the treatment effect, which may not capture non-linear patterns of effect moderation. As with the binary interaction model, it restricts certain aspects of the functional form to be constant across different values too. [Heckman et al., 1997, Imbens and Lemieux, 2008, Brambor et al., 2006]

Implementation Details For all three approaches, I estimate models for both level effects (gender literacy gap in 2011, controlling for the 2001 gap) and percentage change outcomes (percentage change in the gap between 2001 and 2011). In the interaction models, I test sensitivity to bandwidth selection by estimating models with multiple fixed bandwidths (5, 7, and 10). When examining heterogeneity along one dimension, I control for the other two dimensions to ensure that the estimated heterogeneity is not confounded by correlated district characteristics. For example, when analysing heterogeneity by female labour force participation, I include agricultural dependency and child sex ratio as control variables to isolate the specific moderating effect of FLFP. I include state fixed effects, across all specifications and employ state-level clustered standard errors with nearest neighbour adjustment. For the separate subgroup estimations, I use the same CER-optimal bandwidth selection method as in the baseline specification, further expanded on in Section 5.5.

5.4 Covariate Choice

For my baseline district-level covariates, I included the literacy gap, total population, scheduled caste/tribe population shares, and sex ratio from 2001. The baseline literacy gap is crucial as it accounts for pre-existing gender disparities in education before the policy was implemented, allowing me to isolate the changes attributable to increased bank presence. Total population controls help address the mechanical relationship with the population-to-branch ratio that determines treatment status. This is particularly important given that my running variable involves population, to ensure I'm not picking up size effects.

The inclusion of scheduled caste/tribe population shares captures socioeconomic factors affecting educational access. These marginalised groups historically face greater educational barriers, especially for females, and their presence varies systematically across Indian districts [Borooah and Iyer, 2005]. These shares can influence both banking access and literacy outcomes independently of the policy.

Sex ratios directly affect the calculation of gender-based statistics. In districts with highly skewed sex ratios, the literacy gap may appear different due to compositional effects rather than true differences in educational outcomes . In subgroup models, I also control for Child Sex Ratio, which demonstrates collinearity with the sex ratio, so Stata removes it. This multicollinearity is expected given that both measures reflect underlying gender composition and preferences, though they capture different demographic cohorts. The adult sex ratio primarily affects the current calculation of literacy gaps, while the child sex ratio might better indicate contemporary gender attitudes that influence educational investment decisions [Jayachandran, 2015].

I intended to include educational controls at a primary school level in 2001, but ran into some data availability issues. I could only find state-level educational indicators (number of primary schools, gender parity index in education, gross enrolment ratios), rather than the district-level data. Using state-level controls (a higher level than the district level where the RD is implemented) would be problematic for my design, especially since the sample size is already limited (n=521). [Angrist and Pischke, 2009]I wouldn't be able to distinguish the treatment effect from the effect of state-level controls as they're constant for all districts in a state. It may reduce the identifying variation I can exploit, especially if the bandwidth includes districts only from one or two states. So instead, I opted to use state fixed effects which control for all time-invariant state characteristics while keeping statistical power intact. [Imbens and Rubin, 2015, Cattaneo et al., 2020a, Gelman and Hill, 2007]

5.5 Specification Choices

I employ state-level clustering with nearest neighbor variance estimation (`vce(nncluster stateid 3)`) to account for the hierarchical structure of my data, where districts are nested within states. This acknowledges that districts within the same state share unobserved characteristics related to educational policies, institutional factors, and implementation of banking programs [Angrist and Pischke, 2009]. The nearest neighbour approach is particularly well-suited to my RD design as it addresses the challenges of boundary estimation while accounting for spatial correlation within state boundaries [Calonico et al., 2014]. Given that I have approximately 35 states in my sample, this approach is more robust compared to conventional clustering methods, which can lead to over-rejection of the null hypothesis with few clusters[Cameron and Miller, 2015]. This choice follows established practice in development economics literature examining sub-national policy effects in India, where state-level clustering is standard. [Muralidharan and Prakash, 2017]

I employ coverage error rate optimal bandwidths with the `certwo` option, which calculates separate optimal bandwidths below and above the cutoff [Cattaneo et al., 2020a]. This choice is critical given my dataset of approximately 500 districts with substantial heterogeneity across Indian states. Since my primary research objective is to determine the statistical significance of banking access effects on gender literacy gaps rather than precisely quantifying effect magnitudes, the CER-optimal approach offers better properties compared to MSE-optimal alternatives. As Cattaneo et al. [2020a] demonstrate, MSE-optimal bandwidths tend to be excessively large for confidence interval construction when working with medium-sized samples clustered at higher administrative levels. The `certwo` option is particularly appropriate given the asymmetric density of districts around the banking branch ratio threshold and the distinct relationships observed in my subgroup analyses. For robustness, I also present results using MSE-optimal bandwidths (`msetwo`), which consistently yield similar directional effects but with wider confidence intervals – supporting my preference for CER-optimal bandwidths in this specific empirical context.

For my agricultural dependency subgroup analysis, I classify districts in the top quartile (75th percentile and above) rather than using a median split, due to sample size considerations. For robust heterogeneity analysis, Lee and Lemieux [2010] recommend at least 30 effective observations in each subgroup , while Cattaneo et al. [2020a] suggest a minimum of 20-30 effective observations per side for minimal credibility. In preliminary analyses using a median split, the non-agricultural districts had as few as 5-6 effective observations on one side of the cut-off, well below these minimums. Using the top quartile approach creates more balanced comparison groups with sufficient observations to support reliable inference, particularly given that Gelman and Imbens [2019] suggest that for interaction models, you need roughly 2-4 times the sample size of main effects models.

6 Results and Discussion

6.1 Baseline Effect

Table 4: Regression Discontinuity Estimates of the Effect on Gender Literacy Gap

	Level of Literacy Gap			*Percent Change in Literacy Gap*		
	(1)	(2)	(3)	(4)	(5)	(6)
	$p = 1$	$p = 2$	$p = 3$	$p = 1$	$p = 2$	$p = 3$
RD estimate	0.124	0.044	0.532	0.888	3.175	2.707
	(0.277)	(0.329)	(0.450)	(1.554)	(2.260)	(2.612)
Observations	521	521	521	521	521	521

Notes: This table reports regression discontinuity estimates of the treatment effect on the female-male literacy gap. Columns (1)–(3) present level effects; columns (4)–(6) report percentage changes. All specifications control for state fixed effects and district-level covariates. Robust standard errors are shown in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The first-order polynomial results show that districts just above the banking access threshold have a gender literacy gap that is 0.124 percentage points wider than similar districts just below the threshold, though this difference is not statistically significant. Similarly, in relative terms, districts with greater banking access experienced a 0.888% larger increase in the gender gap, but this too is statistically indistinguishable from zero. These findings contradict my prediction that improved banking access would significantly affect gender literacy gaps, suggesting that bank presence alone does not meaningfully impact educational gender disparities across all districts.

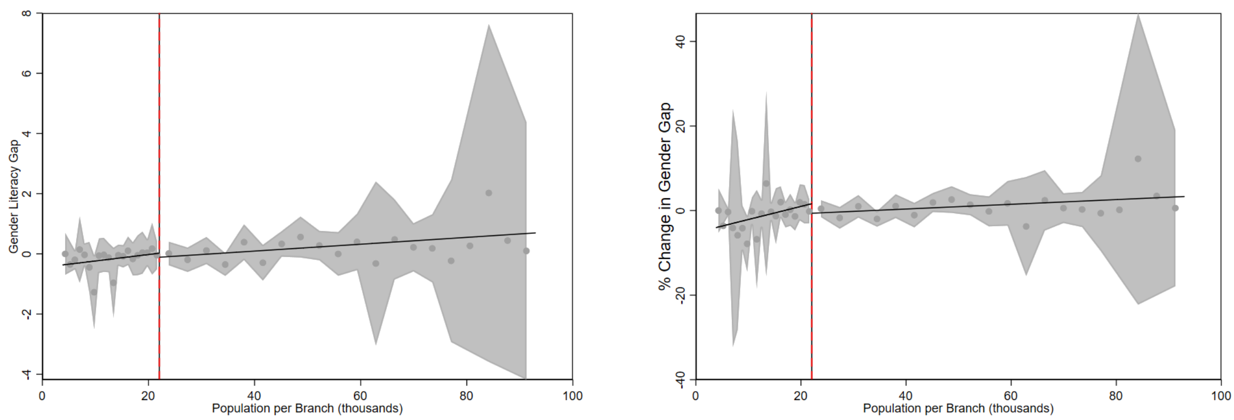

(a) Level effects for Full Sample

(b) Percentage Change for Full Sample

Figure 6: RD Plots for Full Sample

Figure 6 shows the regression discontinuity plots of the effect on the residualised gender literacy gap after controlling for covariates, not the raw values, which improves the precision of the visualisation. The confidence intervals widen considerably as we move away from the cut-off, particularly on the right, indicating less precision in estimates at higher values of the running variable due to fewer observations and greater variability in those regions. The overlapping confidence intervals from the left and right sides at the cut-off point confirm that the discontinuity is not statistically significant.

Despite the absence of a significant overall effect, I still examine whether banking access might have differential impacts across subgroups. This approach allows me to explore the specific channels outlined in my theoretical framework empirically, but also determine whether certain district characteristics create environments where banking presence can more effectively reduce gender disparities in education.

6.2 Subgroup 1: Female Labour Force Participation

Table 5: Banking Access and Gender Literacy Gap: Heterogeneous Effects by Female LFPR

	Level of Literacy Gap			% Change in Literacy Gap		
	(1)	(2)	(3)	(4)	(5)	(6)
	$p = 1$	$p = 2$	$p = 3$	$p = 1$	$p = 2$	$p = 3$
Panel A: High Female LFPR						
RD estimate	-0.304	-0.115	0.162	-1.182	0.018	2.315
	(0.247)	(0.203)	(0.394)	(0.986)	(1.265)	(2.217)
Observations	111	160	161	121	162	157
Bandwidth	5.93	6.48	6.96	5.19	5.67	5.73
Panel B: Low Female LFPR						
RD estimate	0.544	0.503	0.273	3.315	3.895	4.166
	(0.347)	(0.427)	(0.558)	(2.237)	(2.998)	(3.789)
Observations	85	128	152	79	115	143
Bandwidth	3.69	5.86	6.80	3.55	4.48	6.21

Notes: All specifications use CER-optimal bandwidth selection (CERTWO) and include state-level clustered standard errors (nearest-neighbor approach). The running variable is district population per bank branch (in thousands), with a cutoff at 22.05.

Robust standard errors clustered at the state level are in parentheses.

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

In Table 5, districts with high female labour force participation show a gender literacy gap that is 0.304 percentage points narrower with increased banking access, though not statistically significant. These districts experienced a 1.182 percent reduction in the relative size of the gap. In contrast, districts with low female labour force participation show an increase in the gender literacy gap by 0.544 percentage points and a 3.315 percent increase in the relative gap, though these effects are also not individually significant.

The approximately 0.85 percentage point difference in absolute effects (from +0.544 to -0.304) and 4.5 percentage point difference in relative effects between high and low FLFP districts highlight how the same policy could have different outcomes depending on women's existing economic engagement. While these effects do not reach thresholds of statistical significance, their consistent direction and substantial magnitude across multiple specifications are promising and economically meaningful. However whilst my patterns align with the Female Agency and Role Model pathways outlined in my theoretical framework, the statistical insignificance of my results disproves them.

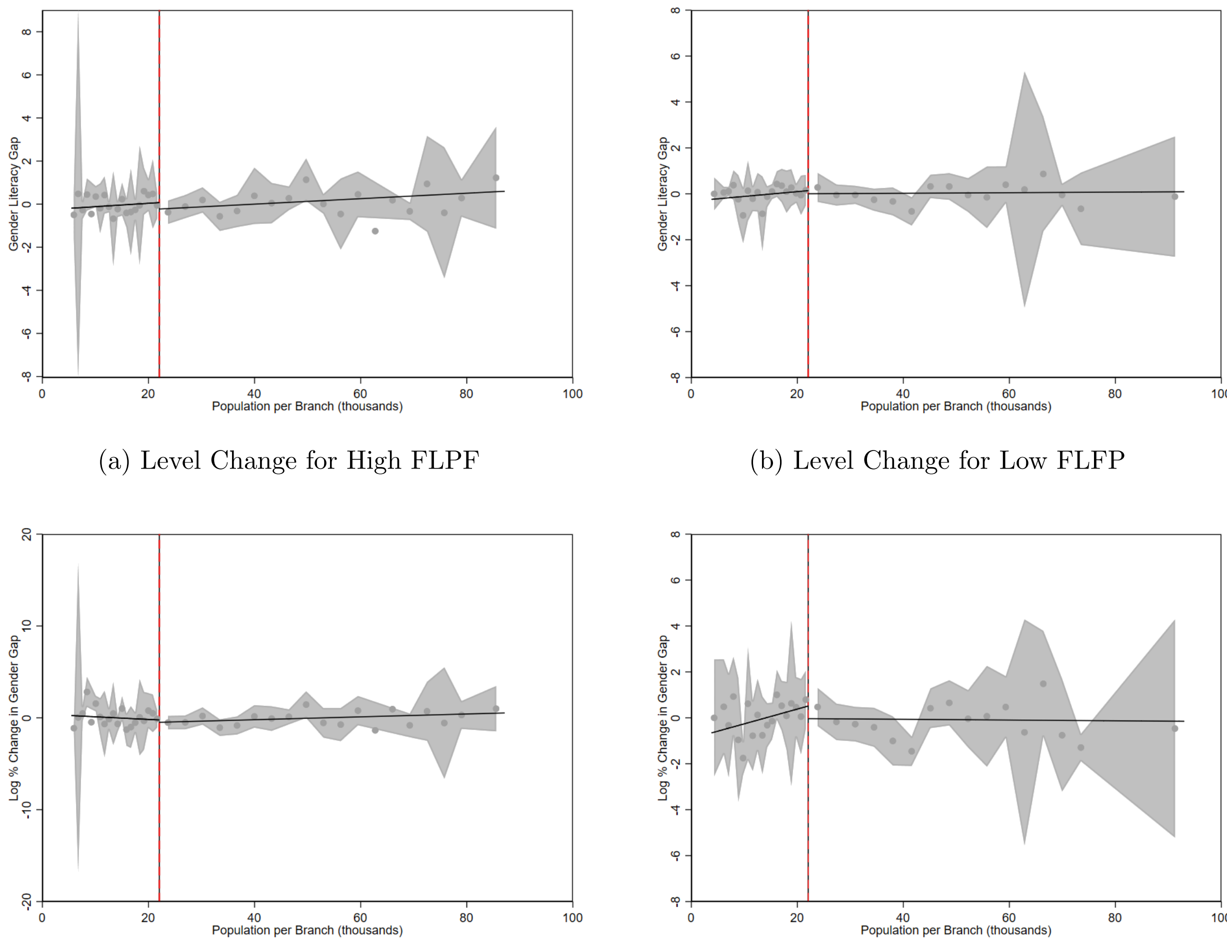

(a) Level Change for High FLPF

(b) Level Change for Low FLFP

(c) Percentage Change for High FLFP

(d) Percentage Change for Low FLPF

Figure 7: RD Plots for the Binary Subgroup Split of Female Labour Force Participation (FLFP)

The plots in Figure 7 visually confirm this heterogeneity. For high FLFP districts (panels a and c), the fitted lines show a visible downward jump at the threshold, indicating a reduction in the gender gap, while low FLFP districts (panels b and d) show a slight upward jump. The confidence intervals are narrower for high FLFP districts, suggesting more precise estimation, but still support statistical insignificance. For percentage change visualisations, I applied a logarithmic transformation that preserves the sign of the original values while compressing the scale, allowing for clearer visualisation of the discontinuity without discarding extreme values.

Table 6: Interaction Model: Effects of Banking Access on Gender Literacy Gap

	Binary Interaction			Continuous Interaction		
Specification	(1)	(2)	(3)	(4)	(5)	(6)
A. Level Effects						
Treatment Effect	−0.012	0.166	−0.089	−0.122	−0.189	−0.366
	(0.736)	(0.568)	(0.448)	(0.424)	(0.334)	(0.276)
Treatment × High LFPR	−0.070	−0.726	−0.475			
	(0.876)	(0.680)	(0.598)			
Treatment × LFPR (cont.)				0.521	−1.511	−1.627
				(4.040)	(2.798)	(2.414)
B. Percentage Change						
Treatment Effect	0.020	1.029	−1.291	0.330	−0.737	−0.305
	(4.070)	(3.209)	(2.663)	(2.609)	(2.008)	(2.332)
Treatment × High LFPR	1.316	−4.051	2.603			
	(5.269)	(4.010)	(5.096)			
Treatment × LFPR (cont.)				11.943	−10.124	0.427
				(23.452)	(16.764)	(18.137)
Observations	127	185	267	127	185	267
Bandwidth	5	7	10	5	7	10
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Table presents RD estimates across bandwidth specifications. Columns (1)-(3) show binary interaction models (treatment × high LFPR dummy); columns (4)-(6) show continuous interaction models (treatment × LFPR). All specifications control for running variable (linear), baseline characteristics, and state fixed effects. CER-optimal bandwidths range from 5.8 to 8.3. Robust standard errors in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In the first-order polynomial specification (Column 1), districts with low FLFP show a negligible treatment effect of -0.012 on literacy gap levels, while the interaction term (-0.070) suggests high FLFP districts experience a slightly larger reduction. For percentage changes, the main effect and interaction follow similar patterns. However, the extremely large standard errors relative to point estimates indicate very low precision across all estimates.

As bandwidths widen (Columns 2-3), coefficients remain inconsistent and statistically insignificant. The continuous interaction models (Columns 4-6) similarly show unstable results, with FLFP interaction coefficients ranging from -1.627 to 0.521 for level effects and -10.124 to 11.943 for percentage changes, all with standard errors frequently exceeding coefficient magnitudes. These statistically insignificant findings contradict my prediction that banking access would more effectively reduce gender literacy gaps in high FLFP districts. The weak results could stem from insufficient statistical power (127-267 observations), high outcome variance, model specification constraints. While separate subgroup analyses showed promising directional patterns, these more rigid interaction models fail to confirm statistically significant heterogeneity in how banking access affects gender literacy gaps across districts with varying levels of female economic participation.

6.3 Subgroup 2: Agricultural Dependency

Table 7: Banking Access and Gender Literacy Gap: Heterogeneous Effects by Agricultural Dependency

	Level of Literacy Gap			% Change in Literacy Gap		
	(1)	(2)	(3)	(4)	(5)	(6)
	$p = 1$	$p = 2$	$p = 3$	$p = 1$	$p = 2$	$p = 3$
Panel A: High Agricultural Dependency						
RD estimate	0.017	0.068	0.187	3.515	3.107	2.966
	(0.340)	(0.594)	(0.827)	(3.228)	(4.601)	(5.995)
Observations	39	44	84	51	45	83
Bandwidth	7.037	6.837	7.932	10.238	6.154	8.993
Panel B: Low Agricultural Dependency						
RD estimate	-0.011	0.042	0.076	-0.212	0.874	0.202
	(0.321)	(0.414)	(0.445)	(1.758)	(2.461)	(2.514)
Observations	161	195	239	175	195	222
Bandwidth	4.191	5.152	6.157	4.632	5.193	5.148

Notes: All specifications use CER-optimal bandwidth selection (CERTWO) and include state-level clustered standard errors (nearest-neighbor approach). The running variable is district population per bank branch (in thousands), with a cutoff at 22.05.

Robust standard errors clustered at the state level are in parentheses.

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

Table 7 tests my prediction that banking access would significantly reduce gender literacy gaps in less agricultural districts but not in highly agricultural areas, which would support the Labour Market Incentives and Role Model Effect pathways. Less agricultural districts show a very small reduction in the gender literacy gap by -0.011 percentage points and -0.212 percent in relative terms, while highly agricultural districts show a slight increase of 0.017 percentage points and 3.515 percent.

While the direction of these coefficients aligns weakly with my prediction—showing marginally better outcomes in less agricultural districts—the magnitude of difference (0.028 percentage points) is trivial, and both effects are statistically indistinguishable from zero due to large standard errors. Higher-order polynomials show inconsistent patterns, with some specifications contradicting the hypothesised direction. The negligible difference between district types suggests that agricultural dependency may not substantially moderate how banking access influences gender literacy gaps, not showing support for the predicted pathways.

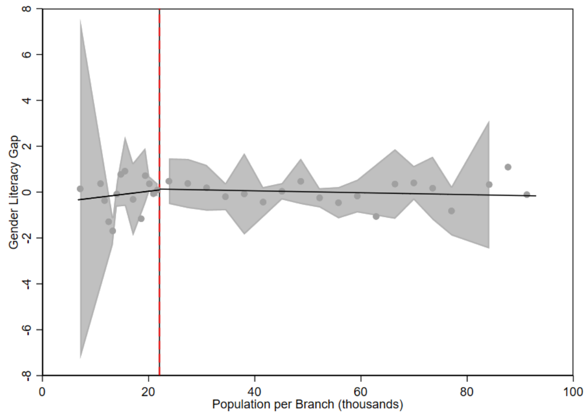

(a) Level Effects for High Agricultural Dependency

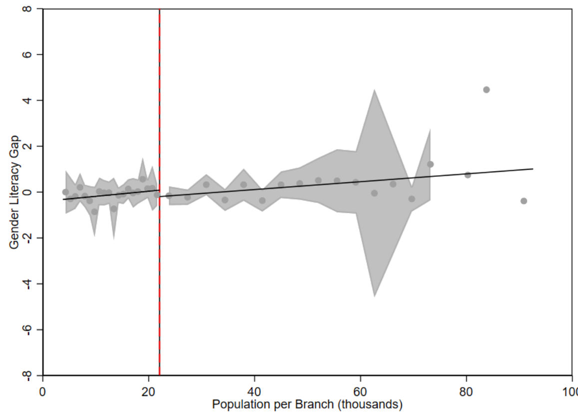

(b) Level Effects for Low Agricultural Dependency

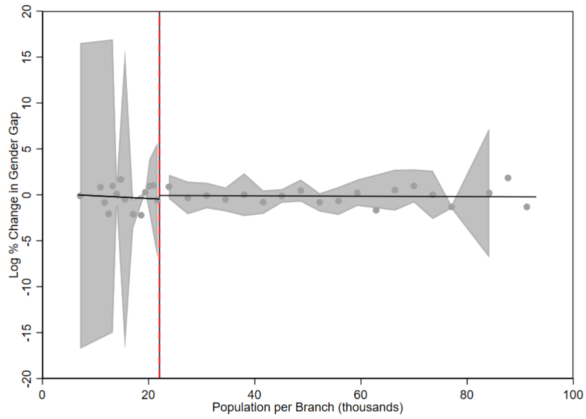

(c) Percentage Change for High Agriculutral Dependency

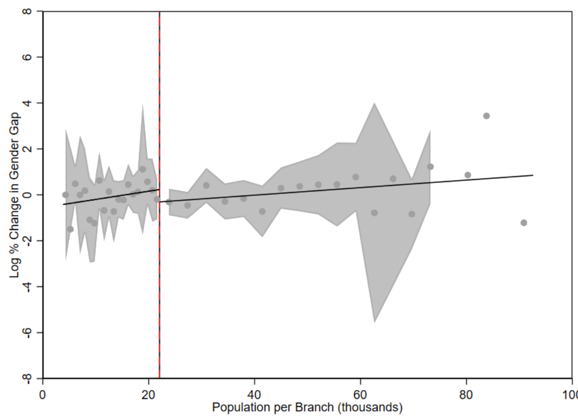

(d) Percentage Change ofr Low Agricultural Dependency

Figure 8: RD Plots by Agricultural Dependency

The plots in Figure 8 visualise these differences, and illustrate the limited precision of these estimates. Districts with low agricultural dependency (panels b and d) show slight reductions in the gender gap at the threshold, while highly agricultural districts (panels a and c) show essentially no change or slight increases. The overlapping confidence intervals for both groups indicate the lack of statistically significant effects.

Table 8: Interaction Model: Effects of Banking Access by Agricultural Dependency

	Binary Interaction			Continuous Interaction		
Specification	(1)	(2)	(3)	(4)	(5)	(6)
A. Level Effects						
Treatment Effect	−0.131	−0.501	−0.716	−0.055	−0.155	−0.400
	(0.873)	(0.596)	(0.487)	(0.363)	(0.305)	(0.277)
Treatment × Top 25% Agri.	−0.137	0.479	0.407			
	(0.984)	(0.705)	(0.593)			
Treatment × Agri. Dep. (cont.)				6.168	4.727	3.786
				(4.970)	(3.719)	(2.918)
B. Percentage Change						
Treatment Effect	−0.676	−2.792	−4.202	0.635	−0.713	0.643
	(4.257)	(3.134)	(2.596)	(2.479)	(1.928)	(2.837)
Treatment × Top 25% Agri.	0.743	3.007	6.339			
	(5.351)	(3.996)	(4.872)			
Treatment × Agri. Dep. (cont.)				35.227	24.558	47.435*
				(27.892)	(21.820)	(25.127)
Observations	127	185	267	127	185	267
Bandwidth	5	7	10	5	7	10
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Table presents RD estimates across bandwidth specifications. Columns (1)-(3) show binary interaction models (treatment × top 25% agricultural dependency); columns (4)-(6) show continuous interaction models (treatment × agricultural dependency). All specifications control for running variable (linear), baseline characteristics, and state fixed effects. Robust standard errors in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Agricultural Dependency Interaction Models

Table 8 examines how agricultural dependency moderates banking access effects through interaction models. In the binary model, the treatment effect (-0.131) for less agricultural districts and interaction term (-0.137) suggest banking access reduces the gender literacy gap more in highly agricultural districts (combined effect -0.268), contradicting the direction found in Table 7's separate subgroup analysis. This inconsistency persists across bandwidths, with interaction terms changing direction from negative to positive as bandwidth increases. The continuous interaction models show positive coefficients (6.168, 4.727, 3.786) indicating that higher agricultural dependency is associated with less beneficial effects on gender gaps, more aligned with our theoretical predictions. For percentage changes, this relationship reaches marginal significance at the widest bandwidth.

Despite these directional patterns, it's important to note that all coefficients remain statistically insignificant at the 5 percent level with large standard errors relative to point estimates, consistent with Table 7's findings. This persistent statistical insignificance across different modelling approaches reinforces our conclusion that agricultural dependency has limited to no moderating influence on how banking access affects gender literacy gaps.

6.4 Subgroup 3: Child Sex Ratio

Table 9: Banking Access and Gender Literacy Gap: Heterogeneous Effects by Child Sex Ratio

	Level of Literacy Gap			% Change in Literacy Gap		
	$p = 1$	$p = 2$	$p = 3$	$p = 1$	$p = 2$	$p = 3$
Panel A: Balanced Child Sex Ratio (CSR)						
RD estimate	-0.639**	-0.544*	-0.375	-4.249**	-3.918*	-2.010
	(0.254)	(0.289)	(0.388)	(1.762)	(2.193)	(2.885)
Observations	92	115	122	132	168	173
Bandwidth	4.38	4.82	4.42	5.91	6.29	6.77
Panel B: Skewed Child Sex Ratio (CSR)						
RD estimate	0.100	0.327	0.622	2.337	3.428	4.718
	(0.399)	(0.512)	(0.598)	(2.442)	(3.023)	(3.673)
Observations	123	179	197	119	182	204
Bandwidth	6.48	9.00	8.89	5.79	9.33	9.36

Notes: All specifications use CER-optimal bandwidth selection (CERTWO) and include state-level clustered standard errors (nearest-neighbor approach). The running variable is district population per bank branch (in thousands), with a cutoff at 22.05.

Robust standard errors clustered at the state level are in parentheses.

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

The child sex ratio analysis reveals evidence for the importance of cultural context in mediating banking effects. In Table 9, districts with more balanced child sex ratios (Panel A) show banking access significantly reduces the gender literacy gap by 0.639 percentage points and the relative gap by 4.249 percent in first-order polynomial specifications. These represent the largest and most statistically significant effects observed across all analyses for the linear polynomial at the 5 percent level. In contrast, districts with more skewed child sex ratios (Panel B) show a positive effect of 0.100 percentage points on the absolute gap and a 2.337 percent increase in the relative gap, neither statistically significant. However, these significant effects are not fully consistent across polynomial orders. The magnitude and significance diminish in higher-order polynomials, with the effect falling to -0.544* in second-order specifications and becoming statistically insignificant (-0.375) in third-order models. This sensitivity to polynomial specification suggests some degree of model dependence in these findings.

Despite this caveat, the pattern of results suggests that banking access likely helps reduce gender educational disparities primarily in districts where cultural norms already value girls more equally. The finding demonstrates how deeply cultural attitudes toward gender may influence the effectiveness of economic policies, supporting my theoretical prediction that the child sex ratio captures fundamental cultural factors that mediate both the Financial Access and Labour Market Incentives pathways.

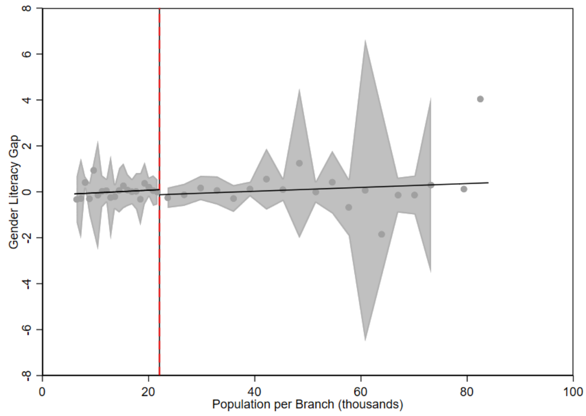

(a) Level Effects for Balanced CSR

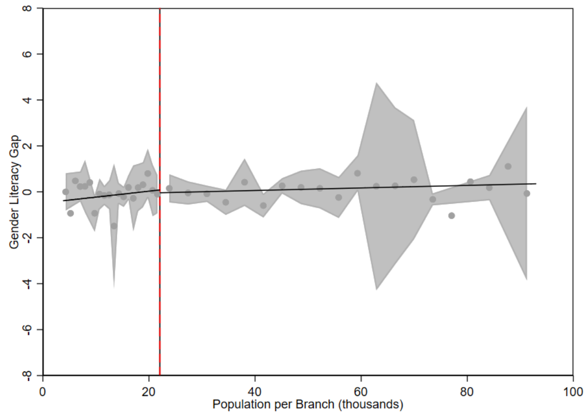

(b) Level Effects for Skewewd CSR

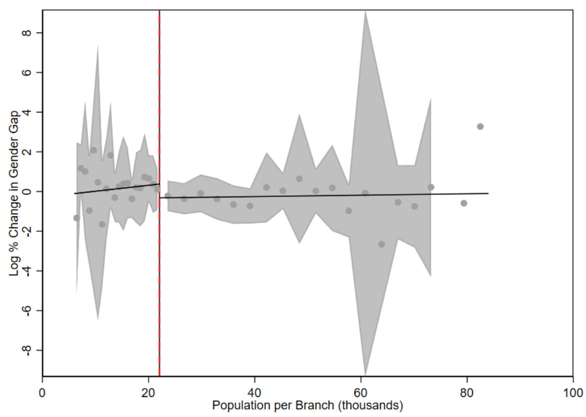

(c) Percentage Change for Balanced CSR

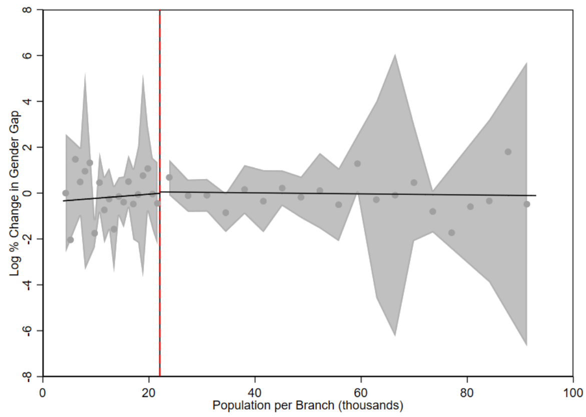

(d) Percentage Change for Skewed CSR

Figure 9: RD Plots depending on Child Sex Ratio (CSR)

The RD plots in Figure 9 visually confirm this pattern. For districts with balanced child sex ratios (panels a and c), there is a clear negative jump at the cut-off, while districts with skewed ratios (panels b and d) show slight positive jumps or essentially no change. The narrower confidence intervals for balanced CSR districts, particularly near the cut-off, indicate more precise estimation.

Table 10: Interaction Moel: Effects of Banking Access by Child Sex Ratio

	Binary Interaction			Continuous Interaction		
Specification	(1)	(2)	(3)	(4)	(5)	(6)
	$BW = 5$	$BW = 7$	$BW = 10$	$BW = 5$	$BW = 7$	$BW = 10$
A. Level Effects						
Treatment Effect	0.242	−0.124	−0.183	−0.096	−0.142	−0.400
	(0.709)	(0.531)	(0.440)	(0.423)	(0.302)	(0.270)
Treatment × High CSR	−0.567	−0.005	−0.373			
	(0.846)	(0.617)	(0.543)			
Treatment × CSR (cont.)				−3.548	1.072	−2.054
				(9.663)	(5.537)	(5.145)
B. Percentage Change						
Treatment Effect	2.667	0.783	−0.129	0.420	−0.565	−0.455
	(3.693)	(2.982)	(2.575)	(2.566)	(1.823)	(2.351)
Treatment × High CSR	−3.598	−2.199	−0.248			
	(4.612)	(3.713)	(4.325)			
Treatment × CSR (cont.)				−23.982	−3.638	33.207
				(52.402)	(31.439)	(43.767)
Observations	127	185	267	127	185	267
Bandwidth	5	7	10	5	7	10
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Table presents RD estimates across bandwidth specifications. Columns (1)-(3) show binary interaction models (treatment × high child sex ratio dummy — above median); columns (4)-(6) show continuous interaction models (treatment × continuous child sex ratio). All specifications control for running variable (linear), baseline characteristics, and state fixed effects. CSR = Child Sex Ratio (females per male, ages 0–6). Robust standard errors in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10 presents interaction models examining how child sex ratio moderates banking effects on gender literacy gaps. In the binary model with bandwidth 5 (column 1), districts with skewed ratios show a treatment effect of 0.242, suggesting banking access widens the gender gap, while the interaction term (-0.567) indicates balanced-ratio districts experience a reduction (-0.325 combined effect). For percentage changes, similar patterns emerge with skewed-ratio districts showing increased gaps (2.667 percent) and interaction terms (-3.598 percent) indicating better outcomes in balanced-ratio districts.

Importantly, unlike the separate subgroup analysis in Table 9 where balanced-ratio districts showed statistically significant improvements, none of these interaction coefficients reach statistical significance. This inconsistency in statistical significance between modelling approaches is notable and suggests limited robustness. Additionally, the continuous interaction model shows inconsistent directions, with coefficients ranging from -3.548 to 1.072 across bandwidths, further undermining confidence in these findings.

Despite the insignificance of these interaction models, the directional consistency with Table 9's findings provides some complementary evidence for the moderating role of cultural gender norms. The separate regression approach in Table 9 could still be reliable for several reasons: (1) it allows for potential threshold effects, where a minimum level of gender equality is necessary before banking access becomes beneficial; (2) it permits different bandwidths optimised for each subgroup; and (3) it allows the full data-generating process to differ between subgroups rather than constraining all parameters except the interaction term. While the interaction models provide additional context, the stronger evidence from separate regressions suggests that cultural contexts with more balanced gender norms create environments where banking access can effectively reduce educational gender disparities.

6.5 Discussion

This quantitative economics research question I explored was how banking access affects gender literacy gaps in India by exploiting the 2005 RBI bank branch expansion policy. While my theoretical framework proposed multiple causal pathways: Female Agency, Role Model Effects, Financial Access and Labour Market Incentives, the empirical results offer a more nuanced picture of how such financial inclusion initiatives truly impact educational gender disparities, and challenged many of my predicted mechanisms.

The baseline results reveal no statistically significant effect of increased bank presence on gender literacy gaps across all districts, challenging the fundamental assumption that financial inclusion initiatives inherently benefit women's educational outcomes differentially to male. This null result is important in itself, suggesting that banking access alone, without complementary interventions or favourable contextual factors, may be insufficient to catalyse meaningful reductions in gender educational disparities.

The subgroup analyses provide more revealing insights. The evidence most strongly supports the pathway where cultural context serves as a critical moderator, with child sex ratio emerging as the only characteristic producing statistically significant heterogeneous effects. In districts with more balanced child sex ratios—indicating more gender-equal cultural norms—banking access significantly reduces gender literacy gaps by 0.639 percentage points and relative gaps by 4.249 percent . This finding suggests that banking access can effectively improve female educational outcomes perhaps through improved financial access or labour market incentives, but primarily in environments where cultural attitudes already value girls' education and well-being.

The evidence for the Female Agency pathway, measured through female labour force participation, is directionally consistent but statistically insignificant. The separate subgroup analysis and interaction models revealed extremely imprecise estimates with large standard errors. Similarly, agricultural dependency showed only weak and inconsistent relationships with banking effectiveness. These patterns suggest that economic structure and women's labour market participation influence banking effects less directly than initially theorised, or that the relationship is more complex than captured in my models.

These empirical findings necessitate a reconsideration of my theoretical framework in several important ways:

The strong moderating effect of child sex ratio suggests that cultural norms may be more fundamental than economic factors in determining whether banking access benefits women. The hypothesised Financial Access pathway—where banking services expand household budget constraints and reduce resource competition between male and female children—likely operates differently depending on underlying cultural preferences. In contexts where son preference remains strong, expanded resources may continue to be allocated disproportionately toward boys, regardless of financial service availability.

The lack of evidence for the Role Model Effect pathway suggests that visibility of women in formal economic roles may be insufficient to change educational aspirations and investments. Enhanced role model effects may require more direct interventions that explicitly connect financial services to female education, rather than relying on implicit connections through observed employment patterns.

The relatively stronger findings for child sex ratio compared to economic factors indicate that cultural change may need to precede economic interventions for the latter to effectively reduce gender disparities. This challenges the sometimes implicit assumption in development economics that economic interventions can effectively drive social changes even in culturally resistant contexts.

The consistently large standard errors across specifications point to substantial unexplained heterogeneity in gender literacy gaps, suggesting that banking access represents just one factor in a complex system of influences. The formal modelling of banking impacts likely oversimplifies the relationship by failing to account for complementary factors, interaction effects, and non-linear relationships that shape educational outcomes.

These findings have important implications for financial inclusion policies, if policymakers want to improve female educational outcomes. Financial inclusion initiatives should be designed with careful attention to local cultural contexts. In regions with strong son preference or gender-biased norms, complementary interventions addressing these cultural factors may be necessary prerequisites for banking access to benefit women. Given resource constraints, policymakers might achieve greater impact by initially targeting financial inclusion efforts in regions with more gender-equal cultural norms, where the soil is 'fertile' for such interventions to benefit women. The limited effects of banking access alone suggest the need for integrated policy approaches that combine financial services with explicit gender-focused components, such as financial literacy programs specifically targeting women or educational subsidies for girls. Since the Female Agency pathway showed limited evidence despite theoretical promise, policies might need to more directly target women's financial decision-making power within households, rather than assuming bank presence automatically enhances this agency.

6.5.1 Limitations and Future Research Directions

By conducting district-level analysis, I inevitably obscure within-district heterogeneity in both banking access and educational outcomes. Banking branches cluster unevenly within districts, and gender literacy gaps likely vary substantially based on urban-rural divides and other local factors that my aggregate measures cannot capture. The six-year window between the 2005 policy implementation and my 2011 outcome measurement likely provides an insufficient time frame to fully capture banking effects on educational outcomes, especially since such changes often require more time to materialise through intergenerational mechanisms. This relatively short period may explain some of the statistically insignificant findings in my analysis.

While my subgroup approach offers suggestive evidence for certain theoretical pathways, I cannot directly test the proposed mechanisms. I infer rather than observe household decision-making processes and educational investment patterns. This indirect mechanism testing limits my ability to conclusively establish the causal channels through which banking access might affect gender literacy gaps. My relatively small sample of 521 districts, further reduced in subgroup analyses, substantially limits statistical power, especially for interaction models. This constraint may explain why many economically meaningful point estimates fail to reach statistical significance in my study. The sample size limitation is particularly acute for highly agricultural districts and in specifications requiring narrower bandwidths.

By focusing primarily on basic literacy as my outcome measure, I examine only a foundational threshold of educational achievement. This approach may miss banking effects on higher educational attainment or quality that could reveal more nuanced impacts on gender disparities. Similarly, my focus on bank branch presence rather than actual banking service usage captures only geographic access without measuring effective financial inclusion. The limited support for several of my theorised pathways suggests the need to reconsider my conceptual framework and explore alternative mechanisms not captured in my current analysis. Future research could address these limitations through more granular data, longer time frames, direct mechanism testing, larger samples, broader educational outcomes, and alternative theoretical approaches to better understand how financial inclusion affects gender educational disparities.

7 Conclusion

After exploring the relationship between banking access and gender educational disparities in India, the findings suggest that financial inclusion initiatives likely don't help reduce gender literacy gaps, but perhaps do in contexts with relatively gender-equal cultural norms. The significant heterogeneity in effects across different district characteristics highlights that banking access can't tackle gender disparities alone, but rather as a contextually-dependent intervention whose effectiveness varies with the social environment.

These results challenge overly simplistic narratives about the transformative potential of financial inclusion for women's outcomes and call for more sophisticated approaches that account for cultural, economic, and institutional factors. Future research and policy should focus on understanding the complex interactions between financial access and local context, designing interventions that explicitly address gender barriers, and evaluating impacts over time frames appropriate for capturing meaningful social change.

8 Robustness Checks

Table 11: Regression Discontinuity Estimates of the Effect on Gender Literacy Gap - MSE-Optimal Bandwidth

	Level of Literacy Gap			*Percent Change in Literacy Gap*		
	(1)	(2)	(3)	(4)	(5)	(6)
	$p = 1$	$p = 2$	$p = 3$	$p = 1$	$p = 2$	$p = 3$
RD estimate	0.124	0.044	0.532	0.888	3.175	2.707
	(0.277)	(0.329)	(0.450)	(1.554)	(2.260)	(2.612)
Observations	521	521	521	521	521	521

Notes: This table reports regression discontinuity estimates of the treatment effect on the female-male literacy gap. Columns (1)–(3) present level effects; columns (4)–(6) report percentage changes. All specifications control for state fixed effects and district-level covariates. Robust standard errors are shown in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 12: Banking Access and Gender Literacy Gap: Heterogeneous Effects by Female LFPR - MSE-Optimal Bandwidth

	High LFPR			Low LFPR		
	$p = 1$	$p = 2$	$p = 3$	$p = 1$	$p = 2$	$p = 3$
	Level of Literacy Gap					
RD estimate	-0.274	-0.371*	0.055	0.487	0.318	0.193
	(0.206)	(0.196)	(0.296)	(0.324)	(0.396)	(0.547)
Observations	125	187	190	103	164	184
Bandwidth	7.101	7.966	8.509	4.440	7.250	8.354
	% Change in Literacy Gap					
RD estimate	-1.690*	-0.845	1.153	3.045	4.167	3.316
	(0.907)	(1.280)	(1.818)	(2.167)	(2.650)	(3.236)
Observations	140	186	177	93	143	176
Bandwidth	6.216	6.974	6.998	4.273	5.540	7.633

Notes: All specifications use MSE-optimal bandwidth selection (msetwo) with state-level clustering. Conventional RD estimates with robust p-values reported. Standard errors clustered at state level using nearest neighbor approach.

Cutoff value: 22.05025 (district population ratio in thousands)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 13: Banking Access and Gender Literacy Gap: Heterogeneous Effects by Agricultural Dependency - MSE-Optimal Bandwidth

	High Agri Dependency			Low Agri Dependency		
	$p = 1$	$p = 2$	$p = 3$	$p = 1$	$p = 2$	$p = 3$
	Level of Literacy Gap					
RD estimate	0.009	0.120	0.700	-0.260	-0.364	-0.465
	(0.277)	(0.360)	(0.470)	(0.568)	(0.775)	(0.837)
Observations	163	201	259	44	57	68
Bandwidth	6.821	7.260	6.521	3.069	4.792	5.148
	% Change in Literacy Gap					
RD estimate	0.045	2.942	4.337	0.324	1.008	0.430
	(1.494)	(2.260)	(2.877)	(2.756)	(3.521)	(4.230)
Observations	211	248	300	49	63	65
Bandwidth	7.834	6.926	8.358	4.259	5.209	5.045

Notes: All specifications use MSE-optimal bandwidth selection (msetwo) with state-level clustering. Conventional RD estimates with robust p-values reported. Standard errors clustered at state level using nearest neighbor approach.

Cutoff value: 22.05025 (district population ratio in thousands)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 14: Banking Access and Gender Literacy Gap: Heterogeneous Effects by Child Sex Ratio - MSE-Optimal Bandwidth

	High Child Sex Ratio (CSR)			Low Child Sex Ratio (CSR)		
	$p = 1$	$p = 2$	$p = 3$	$p = 1$	$p = 2$	$p = 3$
			Level of Literacy Gap			
RD estimate	-0.570**	-0.817***	-0.447	0.100	0.327	0.622
	(0.252)	(0.290)	(0.356)	(0.399)	(0.512)	(0.598)
Observations	125	147	164	123	179	197
Bandwidth	5.788	6.619	6.016	6.484	9.002	8.889
			% Change in Literacy Gap			
RD estimate	-4.249**	-3.918*	-2.010	2.337	3.428	4.718
	(1.762)	(2.193)	(2.885)	(2.442)	(3.023)	(3.673)
Observations	132	168	173	119	182	204
Bandwidth	5.909	6.289	6.770	5.789	9.331	9.364

Notes: All specifications use MSE-optimal bandwidth selection (msetwo) with state-level clustering. Conventional RD estimates with robust p-values reported. Standard errors clustered at state level using nearest neighbor approach.

Cutoff value: 22.05025 (district population ratio in thousands)

CSR = Child Sex Ratio (females per 1000 males, ages 0-6)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

References

Franklin Allen, Asli Demirguc-Kunt, Leora Klapper, and Maria Soledad Martinez Peria. The foundations of financial inclusion: Understanding ownership and use of formal accounts. *Journal of Financial Intermediation*, 27:1–30, 2016.

Siwan Anderson, Jean-Marie Baland, and Karl Ove Moene. Measuring the economic impact of microfinance: Moving beyond loan default and profit. *Journal of Development Effectiveness*, 1 (2):1–27, 2009.

Joshua Angrist and Guido Imbens. Identification and estimation of local average treatment effects, 1995.

Joshua D Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics*. Princeton University Press, 2009.

Claude (Anthropic). Claude ai llm. https://www.anthropic.com/. Utilised Claude AI tool to aid with debugging code in Stata, understanding concepts from academic papers, and offering suggestion for limitations or rewording of my ideas.

Beatriz Armendáriz and Jonathan Morduch. *The economics of microfinance*. MIT press, 2010.

Sam Asher, Tobias Lunt, Ryu Matsuura, and Paul Novosad. Development research at high geographic resolution: an analysis of night-lights, firms, and poverty in india using the shrug open data platform. *The World Bank Economic Review*, 35(4), 2021.

Nava Ashraf, Dean Karlan, and Wesley Yin. Female empowerment: Impact of a commitment savings product in the philippines. *World Development*, 38(3):333–344, 2010.

Susan Athey and Guido Imbens. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360, 2016.

Prashant Bharadwaj, Leah K Lakdawala, and Nicholas Li. The perverse consequences of well-intentioned regulation: Evidence from india's child labor ban. *Journal of the European Economic Association*, 18(3):1158–1195, 2020.

Vani K Borooah and Sriya Iyer. Vidya, veda, and varna: The influence of religion and caste on education in rural india. *The Journal of Development Studies*, 41(8):1369–1404, 2005.

Thomas Brambor, William Roberts Clark, and Matt Golder. Understanding interaction models: Improving empirical analyses. *Political analysis*, 14(1):63–82, 2006.

Miriam Bruhn and Inessa Love. The economic impact of banking the unbanked: evidence from mexico. *World bank policy research working paper*, (4981), 2009.

Robin Burgess and Rohini Pande. Do rural banks matter? evidence from the indian social banking experiment. *American Economic Review*, 95(3):780–795, 2005.

Sebastian Calonico, Matias D Cattaneo, and Rocio Titiunik. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326, 2014.

A Colin Cameron and Douglas L Miller. A practitioner's guide to cluster-robust inference. *Journal of human resources*, 50(2):317–372, 2015.

Pedro Carneiro and Sokbae Lee. Estimating distributions of potential outcomes using local instrumental variables with an application to changes in college enrollment and wage inequality. *Journal of Econometrics*, 149(2):191–208, 2009.

Matias D Cattaneo, Nicol'as Idrobo, and Roc'ıo Titiunik. A practical introduction to regression discontinuity designs: Foundations. *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2020a.

Matias D Cattaneo, Michael Jansson, and Xinwei Ma. Simple local polynomial density estimators. *Journal of the American Statistical Association*, 115(531):1449–1455, 2020b.

Matias D Cattaneo, Nicolás Idrobo, and Rocío Titiunik. *A practical introduction to regression discontinuity designs: Extensions*. Cambridge University Press, 2024.

Kim Fe Cramer. *Bank presence and health*. National Stock Exchange of India Limited, 2021.

Satish Y Deodhar, Sweta Mahandiratta, KV Ramani, Dileep Mavalankar, Sandip Ghosh, Vincent Braganza, et al. An evaluation of mid day meal scheme. *Journal of Indian School of Political Economy*, 22(1-4):33–49, 2010.

Department of School Education and Literacy. Sarva shiksa abhiyan, manual for planning and appraisal. Technical Report Manual/2019/PA, Ministry of Education, Government of India, New Delhi, India, 2019. URL https://dsel.education.gov.in/sites/default/files/2019-05/Manual_Planning_and_Apprisal.pdf. Technical manual for education planning processes.

Jean Drèze and Amartya Sen. *An uncertain glory: India and its contradictions*. Princeton University Press, 2013.

Esther Duflo. Women empowerment and economic development. *Journal of Economic literature*, 50(4):1051–1079, 2012.

Pascaline Dupas and Jonathan Robinson. Savings constraints and microenterprise development: Evidence from a field experiment in kenya. *American Economic Journal: Applied Economics*, 5 (1):163–192, 2013.

Pascaline Dupas, Dean Karlan, Jonathan Robinson, and Diego Ubfal. Banking the unbanked? evidence from three countries. *American Economic Journal: Applied Economics*, 10(2):257–297, 2018.

Joana Foster and Nii Antiaye Addy. Barriers to education for street and working children in india. *UKFIET Oxford Conference on Education and Development*, pages 1–20, 2010.

Milton Friedman. Theory of the consumption function. 2018.

Sandhya Garg and Samarth Gupta. Financial access of unbanked villages in india from 1951 to 2019: A spatial approach. *IEG Working Paper No. 403*, 2020.

Sugandha Garg and Pankaj Goyal. Gold monetization in india: An empirical study of household gold purchase behaviour. *Vikalpa: The Journal for Decision Makers*, 43(1):61–71, 2018.

Andrew Gelman and Jennifer Hill. *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press, 2007.

Andrew Gelman and Guido Imbens. Why high-order polynomials should not be used in regression discontinuity designs. *Journal of Business & Economic Statistics*, 37(3):447–456, 2019.

Claudia Goldin. The u-shaped female labor force function in economic development and economic history. *Investment in Women's Human Capital and Economic Development*, pages 61–90, 1995.

James J Heckman, Jeffrey Smith, and Nancy Clements. Making the most out of programme evaluations and social experiments: Accounting for heterogeneity in programme impacts. *The Review of Economic Studies*, 64(4):487–535, 1997.

Guido Imbens and Karthik Kalyanaraman. Optimal bandwidth choice for the regression discontinuity estimator. *The Review of economic studies*, 79(3):933–959, 2012.

Guido W Imbens and Thomas Lemieux. Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, 142(2):615–635, 2008.

Guido W Imbens and Donald B Rubin. *Causal inference for statistics, social, and biomedical sciences: An introduction*. Cambridge University Press, 2015.

Hanan G Jacoby. Is there an intrahousehold 'flypaper effect'? evidence from a school feeding programme. *The Economic Journal*, 110(462):163–183, 2000.

Seema Jayachandran. The roots of gender inequality in developing countries. *Annual Review of Economics*, 7(1):63–88, 2015.

Robin Jeffrey. *Politics, women and well-being: How Kerala became'a model'*. Springer, 2016.

Prabhat Jha, Maya A Kesler, Rajesh Kumar, Faujdar Ram, Usha Ram, Lukasz Aleksandrowicz, Diego G Bassani, Shailaja Chandra, and Jayant K Banthia. Trends in selective abortions of girls in india: analysis of nationally representative birth histories from 1990 to 2005 and census data from 1991 to 2011. *The Lancet*, 377(9781):1921–1928, 2011.

Dean Karlan, Aishwarya Lakshmi Ratan, and Jonathan Zinman. Savings by and for the poor: A research review and agenda. *Review of Income and Wealth*, 60(1):36–78, 2014.

Stephan Klasen and Claudia Wink. Missing women: Revisiting the debate. *Feminist Economics*, 9(2-3):263–299, 2003.

Anjini Kochar. Social banking and poverty: a micro-empirical analysis of the indian experience. *Stanford Center for International Development, Stanford University. Unpublished*, 2005.

David S Lee and Thomas Lemieux. Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2):281–355, 2010.

Sara Maheshwari. The right of children to free and compulsory education act, 2009. *Indian JL & Legal Rsch.*, 2:1, 2021.

Karthik Muralidharan and Nishith Prakash. Cycling to school: Increasing secondary school enrollment for girls in india. *American Economic Journal: Applied Economics*, 9(3):321–350, 2017.

NITI Aayog. Policy document on [title of the document]. Policy Document SNC1608P000302, NITI Aayog, Government of India, New Delhi, India, August 2016. URL https://www.

nitiforstates.gov.in/public-assets/Policy/policy_files/SNC1608P000302.pdf. Accessed from NITI Aayog's policy documents repository.

Census of India. Census of india, 2011. Accessed via SHRUG https://www.devdatalab.org/shrug_download/, v2.1.pakora.

Arvind Panagariya. Bank branch expansion and poverty reduction: A comment. *Economy, Indian Political*, pages 1–12, 2006.

Catherine Porter and Danila Serra. Gender differences in the choice of major: The importance of female role models. *American Economic Journal: Applied Economics*, 12(3):226–254, 2020.

Silvia Prina. Banking the poor via savings accounts: Evidence from a field experiment. *Journal of Development Economics*, 115:16–31, 2015.

Reserve Bank of India. Master circular on guarantees and co-acceptances. Technical Report 1280, Reserve Bank of India, Mumbai, India, July 2006. Updated on July 1, 2013.

T Paul Schultz. Why governments should invest more to educate girls. *World Development*, 30(2): 207–225, 2002.

Amartya Sen. More than 100 million women are missing. *The New York Review of Books*, 37(20): 61–66, 1990.

James Stock and Motohiro Yogo. Asymptotic distributions of instrumental variables statistics with many instruments. *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*, 6:109–120, 2005.

World Bank. Gdp growth (annual %) - india. World Bank Data, 2025. URL https://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG?locations=IN&start=2001&end=2010. Accessed: 2025-05-08.

Nathaniel Young. Banking and growth: Evidence from a regression discontinuity analysis. 2017.