Adaptation of AI Use to Educational Inequality[†]

This Draft: March 7, 2026

Abstract

Educational artificial intelligence (EdAI) has the potential to mitigate urban-rural educational inequality, but this potential can be realized only if rural users increase their EdAI usage when educational inequality intensifies. We examine this user response through a natural experiment in a Chinese city where a sudden local-enrollment policy rendered rural students ineligible for top-tier urban senior high schools, thereby exogenously increasing urban-rural educational inequality. We focus on a river-based policy boundary, with an urban district on one bank and a rural county on the opposite bank, and use granular cellphone data aggregated to 1 km by 1 km grid cells within a corridor extending five kilometers from each riverbank. We estimate a spatial difference-in-differences model comparing changes in EdAI users across the two banks around the policy shock. We find that the number of active EdAI users on the newly disadvantaged bank declines immediately after the policy change but rebounds above its pre-policy baseline within one year, producing a J-curve pattern. We replicate this pattern along a comparable land-based urban-rural boundary elsewhere in the city, while two policy-neutral boundaries show no such pattern. Heterogeneity analyses indicate that perceived control, proxied by socioeconomic status (SES), shapes this response. Low-SES areas exhibit steeper short-run declines and sharper rebounds, whereas high-SES areas display relatively flat responses. Overall, the findings show that managers and policymakers should consider the full path of user responses when designing business strategies and interventions for marginalized users.

Keywords: Urban-Rural Educational Inequality, Educational Artificial Intelligence, Adaptation Behavior, Natural Experiment

[†]The authors declare that there was no financial support for this research and no conflicts of interest related to this work.

> "If machines produce everything we need, the outcome will depend on how things are distributed. Everyone can enjoy a life of luxurious leisure if the machine-produced wealth is shared, or most people can end up miserably poor if the machine-owners successfully lobby against wealth redistribution."
>
> — Stephen Hawking

1 Introduction

Educational inequality[1] refers to systematic disparities in the distribution of academic resources, including but not limited to school funding, qualified teachers, and physical facilities, to socially marginalized communities (Blanden et al., 2023; Ferreira and Gignoux, 2014). These disparities remain a central barrier to human capital development across the world. A large body of research shows that such disparities depress marginalized students' skill acquisition (Duflo et al., 2011; Zimmerman, 2003), reduce their future earnings (Chetty et al., 2014; Heckman et al., 2010; Hendel et al., 2005; Jackson and Holzman, 2020), and worsen health outcomes (Balaj et al., 2024; Clark and Royer, 2013; Davies et al., 2018). At the macro level, educational inequality constrains productivity growth (Hanushek and Woessmann, 2008; Heckman and Masterov, 2007) and contributes to the persistent poverty across generations (Ahluwalia, 1976; Neidhöfer et al., 2018).

Educational inequality also tends to follow a clear spatial pattern. One of its most prominent forms is the gap between urban areas and rural areas. Rural students often encounter weaker school infrastructure, fewer experienced teachers, and more constrained learning environments (Young, 2013; Zahl-Thanem and Rye, 2024). This urban-rural educational gap exists in both developed economies (Byun et al., 2012; van Maarseveen, 2021; Zahl-Thanem and Rye, 2024) and developing economies (Agrawal, 2014; Li et al., 1200; Lu, 2012; Zhang, 2017). In India, for example, 26.3% of the population remains illiterate. While the proportion of the population with below primary school level education has declined from 12.5% to 11.5% in 2016, this improvement is concentrated in urban areas, leaving urban-rural educational disparities largely unaddressed (Choudhury, 2025).

Educational artificial intelligence (EdAI) offers a promising avenue for mitigating such urban-rural educational inequality. By providing personalized guidance and feedback at scale, EdAI can replicate key functions of private tutoring for large numbers of students, thereby helping

[1] Throughout the paper, we use educational inequality, educational gap, and educational disparity interchangeably.

alleviate the scarcity of high-quality teachers in marginalized areas. Prior research demonstrates the potential of similar educational technologies to narrow educational gaps (Anderson and Kimball, 2019; Bianchi et al., 2022; Gao et al., 2021; Goodman et al., 2019; Keppler et al., 2022; Zheng et al., 2024). Motivated by this promise, many governments have recently expanded efforts to broaden rural students' access to EdAI (IndiaToday, 2025; UNESCO, 2024; Xinhua, 2025).

However, realizing this promise depends critically on how rural users actually engage with EdAI. Rising urban-rural educational inequality not only worsens educational outcomes but can also dampen users' beliefs about the returns to effort, weaken their sense of control over future outcomes, and reduce their willingness to use new technologies (Brosnan and de Waal, 2014; Cohn et al., 2014; van Wolkenten et al., 2007). If inequality discourages rural users from engaging with EdAI, then EdAI availability may not translate into meaningful use. In that case, EdAI usage may become concentrated among more advantaged groups, causing the benefits of EdAI to accrue disproportionately to those who are already ahead. Rather than narrowing urban-rural educational gaps as an equalizer, EdAI could become a polarizer that widens the very disparities they are expected to address. This type of equalizer-polarizer dilemma has been discussed in other technological settings (Agostinelli et al., 2022; Chen et al., 2015; Fu et al., 2025; Pswarayi-Riddihough and Klapper, 2025; Zhang et al., 2021). For example, Zhang et al. (2021) show that Airbnb's AI pricing algorithm narrows racial revenue gaps among hosts who adopt it, but because black hosts are significantly less likely to adopt the algorithm, the overall racial gap widens.

Taken together, understanding how rural users change their EdAI usage in the face of rising educational inequality is therefore essential. Yet existing research has largely examined how inequality shapes behaviors that are not technology-related (Bamberger et al., 2021; Bao et al., 2024; Browman et al., 2017; Mostagir and Siderius, 2023; Yoon and Kim, 2018; Zhong et al., 2025), leaving open the question of how urban-rural educational inequality affects rural users' EdAI usage. This gap motivates our central research question: *How do rural users change their EdAI usage in response to urban-rural educational inequality?*

In this study, we provide new empirical evidence on how increases in urban-rural educational inequality shape rural users' EdAI usage. Our analysis exploits a natural experiment created by an abrupt local-enrollment policy implemented in April 2023 in Fuzhou, a major Chinese city. The policy barred rural students from enrolling in senior high schools located in districts, where

the city's top-tier senior high schools are concentrated. We build our identification around a river-based urban-rural boundary that sharply divides the affected regions. To study how students respond to this shift, we construct a quarterly panel of EdAI usage using anonymized cellphone activity data from July 2022 to June 2024. We create a grid covering a 5-kilometer zone on each side of the river and aggregate EdAI usage into 1 km $\times$ 1 km cells. For each cell, we track the number of active users accessing EdAI in each quarter. Residents on the rural bank constitute the treatment group because they lost access to high-quality senior high schools after the policy, while residents on the urban bank serve as the control group because their access remained unchanged. We estimate a spatial difference-in-differences model comparing changes in EdAI usage between the rural and urban areas before and after the policy shock.

We find that increases in urban-rural educational inequality reshape EdAI usage in different ways in the short and long run. During the quarter of the policy change and the following two quarters, the number of active users of EdAI on the rural riverbank declined sharply. Starting in the third post-policy quarter, however, EdAI usage rebounded above its pre-policy baseline and remained persistently higher throughout the remainder of the study period. On average, the short-term effect corresponds to a 33.4% decline in active EdAI users, followed by a 186.1% increase in the longer term. Together, these dynamics form a pronounced J-curve pattern: an immediate decline in EdAI usage followed by a substantial recovery.

To assess whether the J-curve pattern truly reflects the impact of the rising urban-rural educational inequality, we conduct two falsification tests along river boundaries where no cross-bank policy differences existed. The first test examines an immediate upstream segment where both sides of the river are rural and were subject to the same enrollment restriction. The second examines a nearby tributary where both sides are urban and were entirely unaffected by the policy. In both rural-rural and urban-urban settings, the estimated effects are statistically indistinguishable from zero, which reinforces that the J-curve pattern identified in our main analysis arises from the policy shock rather than from inherent differences across river boundaries. We further evaluate the robustness of our findings using a comprehensive set of checks. These include replicating the analysis along a comparable land-based boundary elsewhere in the city, sub-sample tests that exploit the policy's exemption rule, alternative distance bands, synthetic difference-in-differences estimators, and adjustments for spatial correlation. All tests consistently reproduce the J-curve

pattern in EdAI usage. Finally, we rule out alternative explanations. We find no evidence of residents re-sorting across riverbanks and differential trends in mobile network infrastructure. Taken together, these results highlight a robust J-curve pattern in EdAI usage following a sudden rise in urban-rural educational inequality.

In exploring the underlying mechanism, we find that perceived control potentially shapes the user response. According to coping theory (Folkman, 1984, 2013), when confronted with the restricted educational opportunities created by the local enrollment policy, users in rural areas initially experience a diminished sense of perceived control, which prompts emotion-focused coping and a decline in their use of EdAI. Over time, as they reappraise their circumstances and recognize EdAI as a compensatory resource, they shift toward problem-focused coping, leading to a rebound in EdAI usage. To test this mechanism, we conduct sub-sample analyses by socioeconomic status (SES) using both cell-level housing prices and town-level industrial output. We find convergent evidence that low-SES areas, where perceived control is plausibly more constrained, exhibit a larger short-term drop and a sharper long-term rebound in EdAI usage. In contrast, high-SES areas display a relatively flat response. These patterns align with the view that perceived control plays a central role in shaping how rural users change their use of EdAI when educational inequality intensifies.

We contribute to two streams of research. First, we extend work on information technology and educational inequality by shifting the analytic lens from access to adaptation. Prior studies largely evaluate whether information technologies mitigate or exacerbate preexisting educational gaps across social groups (Anderson and Kimball, 2019; Bianchi et al., 2022; Figlio et al., 2013; Gao et al., 2021; Goodman et al., 2019; Keppler et al., 2022; Zheng et al., 2024). We move beyond this focus by showing how disadvantaged people change their engagement with EdAI when structural barriers tighten. Second, we contribute to the literature on human responses to inequality by documenting users' behavioral changes in technology usage. Existing research primarily emphasizes behavioral changes that are not related to technology (Bamberger et al., 2021; Bao et al., 2024; Browman et al., 2017; Mostagir and Siderius, 2023; Yoon and Kim, 2018; Zhong et al., 2025), such as consumption choices (Yoon and Kim, 2018; Zhong et al., 2025) and academic persistence (Browman et al., 2017). We show that widened educational inequality induces a distinct temporal J-curve in EdAI usage, thereby integrating technology-related behavior into research on how people

respond to rising inequality.

Below, we first introduce the research streams related to this paper in section 2. We then introduce the institutional background of the policy, identification strategy, data, and economic specification in section 3. Then, we present the main results, falsification tests, robustness checks, and exclusion of alternative explanations in section 4. We further present the analysis of the potential mechanism in section 5. Finally, we conclude with a discussion in section 6.

2 Literature Review

2.1 Information Technology and Educational Inequality

Our paper is closely related to the research stream linking information technology to educational inequality[2]. Prior research has examined two broad categories of technologies: general-purpose information technologies (e.g., broadband internet and personal computers) and education-related information technologies (e.g., online learning platforms and educational crowdfunding platforms).

Existing research offers mixed evidence regarding the effects of general-purpose information technologies on educational inequality (Belo et al., 2014; Chen et al., 2021; Fairlie and London, 2012; Fairlie and Robinson, 2013; Gao et al., 2021; Goolsbee and Guryan, 2006; Leuven et al., 2007). On the one hand, several studies suggest that access to such technologies can help reduce inequality across social groups. For instance, Fairlie and Robinson (2012) find that providing home computers slightly improves students' academic performance, with particularly strong benefits for those living farther from campus who previously faced greater access constraints. On the other hand, other research indicates that such information technologies may have little effect on, or even exacerbate, existing educational inequalities. For example, Belo et al. (2014) show that broadband expansion affects exam scores across all types of schools, regardless of preexisting performance levels. Leuven et al. (2007) further show that additional funding for computers and software can negatively affect girls' achievement, thereby widening gender gaps in education.

Compared with general-purpose technologies, education-related information technologies

[2]We consider urban-rural educational inequality as a specific type of educational inequality and discuss the literature of broad educational inequality in this section.

tend to yield more equalizing effects (Anderson and Kimball, 2019; Bianchi et al., 2022; Figlio et al., 2013; Gao et al., 2021; Goodman et al., 2019; Keppler et al., 2022; Zheng et al., 2024). For example, Zheng et al. (2024), leveraging the pandemic-induced shift to remote learning as a natural experiment, find that students in developing regions improved their learning outcomes more than peers in developed areas. In another context, Keppler et al. (2022) show that educational crowdfunding platforms help teachers enhance student learning outcomes, especially in low-income schools, thereby narrowing educational inequality.

While prior studies primarily treat information technology as an important factor influencing educational inequality, our study reverses this perspective by examining educational inequality as an antecedent of information technology usage. Specifically, we investigate how disadvantaged groups respond to heightened educational inequality through their engagement with EdAI. This reversal is crucial because it reveals information technology usage as an endogenous response mechanism, not just an external intervention, which has important implications for understanding when and why technologies succeed or fail in mitigating educational inequality. Therefore, our research contributes to this literature by providing causal evidence on how rural people change their information technology use in response to increased educational inequality, thereby offering a complementary perspective on technology usage under structural constraints.

2.2 Human Responses to Inequality

A growing body of research demonstrates that inequality functions as a powerful environmental stressor that triggers diverse responses. These responses generally fall into two broad categories: cognitive adjustments and behavioral adaptations.

First, inequality reshapes how individuals perceive their social standing, sense of control, and opportunity structures, thereby influencing beliefs about fairness, redistribution, and social mobility (Andretti et al., 2025; Browman et al., 2019; Chen et al., 2024; McCall et al., 2017; Payne et al., 2017; Sands, 2017). For example, Browman et al. (2019) show that exposure to inequality erodes young people's belief in socioeconomic opportunity, reducing motivation to pursue upward mobility. Moreover, McCall et al. (2017) find that exposure to rising economic inequality strengthens beliefs in structural rather than individual determinants of success. Also, inequality

exposure increases support for redistributive government policies.

Second, inequality also manifests in concrete behavioral changes, typically observed in non-technology domains (Bamberger et al., 2021; Bao et al., 2024; Browman et al., 2017; Mostagir and Siderius, 2023; Yoon and Kim, 2018; Zhong et al., 2025). For instance, Yoon and Kim (2018) show that low-SES people who perceive limited economic mobility tend to seek greater variety in consumption as a compensatory response to a diminished sense of personal control. Browman et al. (2017) find that perceptions about socioeconomic mobility directly influence low-SES students' academic persistence when confronted with difficulty.

While this literature has provided valuable insights into the cognitive and behavioral consequences of inequality, it offers a limited understanding of how people adapt their technology usage behaviors when confronted with widening educational inequality. Our study fills this gap by providing empirical evidence on how disadvantaged people change their use of EdAI following a policy shock that increases educational inequality, thereby extending the inequality-response literature into the technology domain.

2.3 External Shocks and Educational Technology Usage

We acknowledge two papers that are closely related to our study, both examining how external shocks reshape educational technology usage.

First, Bacher-Hicks et al. (2021) show that during the early months of the COVID-19 pandemic, online searches for learning resources surged, with much larger increases in higher-income areas with better internet access, suggesting that the pandemic amplified educational disparities. Second, Ko et al. (2023) find that students living in the epicenter of the COVID-19 outbreak initially reduced their use of an educational application but later recovered to match the learning trajectories of unaffected peers.

While these studies demonstrate how health-related disruptions influence educational technology engagement, our work differs in the nature of the shock. The COVID-19 pandemic represents a universal and temporary shock that disrupted schooling for all students. This shock imposed short-term constraints on access to schools, which were later recovered. In contrast, the policy intervention in our study (i.e., local-enrollment policy) constitutes a persistent change that

exogenously heightened urban-rural educational inequality. This setting allows us to isolate behavioral responses to a durable increase in structural disadvantage, rather than responses driven by temporary disruptions or access constraints.

3 Empirical Design

3.1 Local-Enrollment Policy

We exploit a policy change that occurred in Fuzhou, a city in China. Fuzhou is a prefecture-level city in southeastern China and the capital of Fujian Province. As one of the major cities in China, Fuzhou had a total population of 7.336 million in 2023. According to official statistics, its total gross domestic product (GDP) reached about 1.293 trillion yuan in 2023, with the tertiary (service) sector accounting for the largest share of output.[3] Fuzhou includes six districts and seven counties or county-level cities (hereafter collectively referred to as counties). Although districts and counties are formally classified at the same county-level administrative tier in China, they differ sharply in their degree of urbanization, economic structures, and access to educational resources. In prefecture-level cities such as Fuzhou, districts typically constitute the urban core, characterized by service-oriented economic structures, and a high concentration of public investments—especially in education. Counties, by contrast, encompass predominantly rural or peri-urban areas, with lower urbanization rates, more manufacturing- and agriculture-oriented economies, and substantially fewer high-quality educational resources. Throughout this paper, we use the district-county boundary in Fuzhou as an institutionalized and policy-relevant proxy for the urban-rural divide. In Appendix A, we provide institutional background and quantitative evidence documenting the differences between districts and counties in Fuzhou.

We focus on the local-enrollment policy in Fuzhou as the central institutional setting of our analysis. Before April 2023, students from the counties were allowed to enroll in senior high schools located in the districts. In April 2023, a sudden local-enrollment policy fundamentally altered this arrangement by restricting senior high school admissions along the district-county (urban-rural) boundary. Specifically, cross-district enrollment within districts and cross-county enrollment within counties remained permitted, but cross-boundary enrollment between districts

[3]Source: https://www.fuzhou.gov.cn/book/fztjnj/2024tjnj/indexch.htm

and counties was prohibited. A screenshot of the official policy document is presented in Appendix B. Under the new policy, only students whose academic records were registered in junior high schools located within the districts were eligible to attend senior high schools in the districts. Students from the counties were limited to senior high schools within the counties. This policy change generated a sudden educational disparity along the district-county (urban-rural) boundary because most of the city's top-tier or high-quality senior high schools are concentrated in the districts. We will introduce the institutional details of high-quality senior high schools in the next subsection.

This local-enrollment policy provides a valuable natural-experimental setting for identifying the causal effects of heightened educational inequality for two main reasons. First, the policy was implemented without prior public announcement or a transition period, which thus minimizes anticipatory behavioral responses. Specifically, residents in the affected riverbank area had no opportunity to adjust their educational or residential choices before the policy shock. Second, the policy's strict implementation largely rules out self-selection across riverbanks. To obtain eligibility for admission to top-tier senior high schools in the districts, students must have academic records registered in junior high schools located within the districts before the senior high school entrance examination. Because this criterion is based on school enrollment rather than household registration, relocating after the policy announcement would not help families bypass the restriction. In practice, there are two possible stages at which county-to-district admission might occur, but each faces strong institutional barriers. (a) During the primary school or junior high school stage, transferring to a district school is administratively cumbersome, requires multiple layers of approval, and is seldom granted. (b) During the transition from primary to junior high school, admissions typically follow a feeder-school arrangement, under which most pupils advance automatically to affiliated junior high schools within the same county, making county-to-district admission rare. Taken together, these constraints render relocation both inefficient and costly, making large-scale self-selection highly unlikely.

3.2 Policy-Induced Educational Inequality

The policy-induced restriction intensified pre-existing disparities in access to high-quality education (i.e., top-tier senior high schools) along the district-county (urban-rural) boundary. High quality in this context reflects multiple dimensions, including stronger resource inputs, higher teaching quality, and notably higher university admission rates.

One observable indicator of school quality is the designation of "Model High Schools" by the Fujian Provincial Department of Education. Among the top thirty schools recognized between 2018 and 2021,[4] seven are located in Fuzhou's districts, while none are located in the counties.[5] Model High School status brings substantial public investment. According to the provincial policy document, local education authorities are required to coordinate across government departments to increase fiscal allocations, upgrade facilities, and provide preferential support in admissions, teacher recruitment, and teacher promotion.

A similar pattern arises under another provincial quality designation, the "Provincial First-Class Standard High School." High schools can apply for this status, and the Fujian Provincial Department of Education evaluates applicants across nine dimensions, including school management, teaching quality, and campus facilities. Schools that receive the designation must prepare a three- to five-year development plan, and local governments commit to supporting its implementation through preferential treatment in funding, staffing, and equipment allocation.[6] As of 2024, eighteen First-Class Standard High Schools were located in Fuzhou's districts, compared with twelve in the counties.[7] Using resident population figures from the 2024 Fuzhou Statistical Yearbook,[8] this translates into approximately 0.043 First-Class Standard High Schools per ten thousand residents in districts, versus 0.031 in counties.[9]

These resource advantages generally translate into higher academic performance, most notably in the rates at which students gain admission to better universities through the National

[4]Source: https://jyt.fujian.gov.cn/xxgk/zywj/201901/t20190104_4734352.htm

[5]One exception is Fuzhou No. 1 High School, which is physically located in a county but is administered under the urban-district admissions system and does not enroll county students.

[6]Source: https://jyt.fujian.gov.cn/xxgk/zdjc/202301/t20230131_6101460.htm

[7]Source: httpsjyj.fuzhou.gov.cn/zfxxgkzl/gkml/gzdt/202409/P020240913650105413303.pdf

[8]Source: https://www.fuzhou.gov.cn/book/fztjnj/2024tjnj/indexch.htm. We acknowledge that the 2024 Fuzhou Statistical Yearbook reports resident population data for 2023, which are the most recent data available.

[9]The disparity is even more pronounced in our main analysis county, Minhou, which has only one First-Class Standard High School. This corresponds to roughly 0.010 schools per ten thousand residents.

College Entrance Examination (NCEE). We manually collect data on the first-tier university admission rates of NCEE in 2023 and 2024 from the top-5 senior high schools in districts and counties. We find that senior high schools in districts maintained consistently high admission rates, averaging 93.6% in 2023 and 94.6% in 2024. Senior high schools in counties exhibited lower performance, with an average of 73.5% in 2023 and 77.8% in 2024.[10] This translates to a 20.1 percentage point gap in 2023 and a 16.8 percentage point gap in 2024. Since the NCEE remains the most decisive and competitive channel for access to higher education in China (Li et al., 2012; Liu, 2013; Zhang et al., 2005), placement in a top-tier senior high school substantially enhances students' long-term educational opportunities and social mobility. Taken together, a local-enrollment policy that restricts rural students' access to such schools therefore exacerbates urban-rural educational inequality.

3.3 Identification Strategy

Our main analysis focuses on a segment of the district-county boundary between Cangshan, a district, and Minhou, a neighboring county, indicated by the red dashed line in Figure 1. The two areas are separated by the Min River, which flows from the northwest to the southeast. The northeastern bank of the river belongs to Cangshan District, while the southwestern bank belongs to Minhou County. The river-based boundary in our main analysis extends for 32.20 km. Following the implementation of the local-enrollment policy in April 2023, students residing on the southwestern riverbank lost eligibility to enroll in senior high schools located in the districts, where most of the city's top-tier senior high schools are concentrated. In contrast, students living on the northeastern riverbank retained full access to these high-quality senior high schools. We exploit this policy-induced contrast using a spatial difference-in-differences framework that compares changes over time in EdAI usage between areas on opposite sides of the river. Users in the rural bank constitute the treatment group because they became educationally disadvantaged after the policy, whereas users in the urban bank constitute the control group, whose access to high-quality schools remained unchanged. By tracking both groups before and after the policy shock, we estimate how an exogenous rise in urban-rural educational inequality influences rural users' EdAI usage.

[10] In our main analysis county, Minhou, senior high schools exhibited much lower performance, with an average of 33.0% in 2023 and 38.2% in 2024.

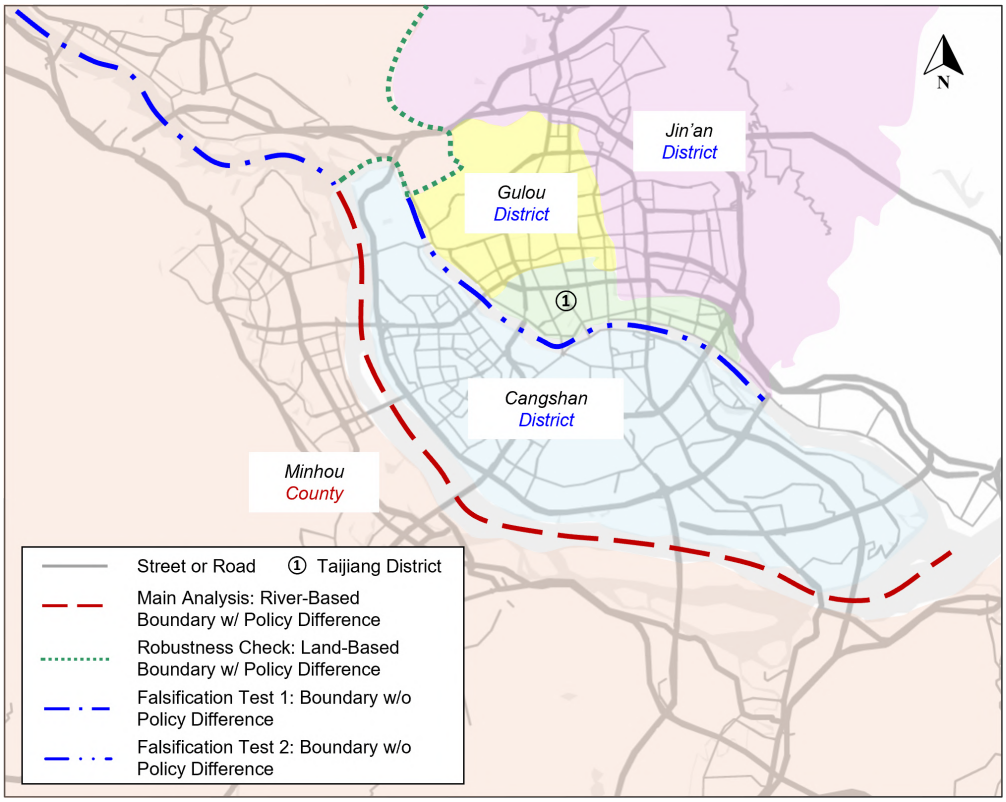


Figure 1: Research Design

Notes: The map displays the spatial layout of Fuzhou City and illustrates the research design. Each shaded area represents a district or a county. The map displays the four boundaries used in the main analysis, falsification tests, and robustness checks. We show the corresponding satellite image in Appendix C.

3.4 Data

3.4.1 EdAI Usage Data

We obtain anonymized cellphone usage records from China's largest telecommunications provider. As of 2023, the provider serves approximately 991 million subscribers nationwide. These data capture user activity through connections to communication towers, which allows us to measure the spatial distribution of users who use EdAI at a fine-grained geographic scale. We collect data for a corridor spanning nine kilometers on each side of the Min River. In the main analysis, we focus on observations located very close to the river-based boundary, specifically within five kilometers on each side. Additional distance bands are included later as part of the robustness checks. We also collect EdAI activity data for the other boundaries used in falsification tests and robustness checks, as introduced in later sections. The area is divided into 1 km $\times$ 1 km grid cells,

which serve as the unit of analysis.[11] We process all analytical grid-cell data based on the World Geodetic System 1984 (WGS84)[12] and Python geospatial libraries. To measure EdAI usage within each grid cell, we construct a comprehensive list of EdAI applications, as detailed in Appendix D. For each grid cell, we aggregate the number of resident users who accessed any of these EdAI applications quarterly from July 2022 to June 2024, which yields a balanced panel of EdAI usage covering eight quarters.[13] The final sample for the main analysis contains 192 grid cells. Among these, 97 cells lie on the treated (rural) county riverbank and 95 cells lie on the control (urban) district riverbank.[14]

3.5 Resident Data

We source resident population data from China's Seventh National Population Census and use estimates produced by the Institute of Geographic Sciences and Natural Resources Research at the Chinese Academy of Sciences.[15] The dataset provides gridded resident estimates at a 1 km × 1 km resolution. Because the census grid cells and the analytical grid cells (i.e., grid constructed for the EdAI usage data) are not perfectly aligned, we use area-weighted interpolation to harmonize them, assigning resident counts to each analytical grid cell according to the proportional overlap between the two systems. For each analytical grid cell g, the interpolated residents count is calculated as

$$P_g = \sum_i w_{ig} \times P_i, \tag{1}$$

where P_i is the number of residents in census grid cell i, and w_{ig} is the share of census grid cell i overlapping with analytical grid cell g with $0 \leq w_{ig} \leq 1$ and $\sum_g w_{ig} = 1$ for each analytical grid cell g. For example, if 60% of a census grid cell overlaps with a given analytical grid cell, we assign 60% of that census cell's residents to the analytical grid cell. This interpolation method preserves

[11] We do not aggregate to grid cells smaller than 1 km × 1 km because the telecommunications provider indicates that measurement noise increases substantially at finer spatial resolutions.

[12] We acknowledge that there is a known misalignment between WGS84 and China's GCJ-02 coordinate system. All spatial data originally recorded in GCJ-02 were converted to WGS84 using standard coordinate transformation procedures to ensure spatial consistency across datasets.

[13] To identify resident users, we classify a user as residing in a given grid cell if their mobile device was located in that cell for at least fifteen days between 7:00 p.m. and 7:00 a.m. during the month preceding the policy implementation.

[14] To limit potential bias from university populations, we exclude all grid cells that contain any university. Since university-driven activity does not reflect the population directly affected by the policy, including these areas would distort the measurement of EdAI usage and weaken the comparability between rural and urban users.

[15] Source: https://www.resdc.cn/doi/doi.aspx?doiid=32

the total residents count and ensures spatial consistency across two grid systems. The resident population distribution across the study area in our main analysis is shown in Appendix E.

3.5.1 Housing Price Data

We obtain proprietary housing price data from a major second-hand real estate platform in China. The dataset covers 774 commercial housing developments located within a five-kilometer distance band along the river boundary used in our main analysis. Each record contains precise geographic coordinates and transaction information.[16] We aggregate these development-level data to match the spatial resolution of our main dataset for the same period. Specifically, we compute quarterly averages of transaction prices at the 1 km $\times$ 1 km grid-cell level from July 2022 to June 2024. [17]

3.5.2 Other Data

We further compile data from several public sources to capture other grid-cell-level characteristics. First, nighttime light intensity is obtained from the National Aeronautics and Space Administration (NASA).[18] Second, information on communication towers is collected from OpenCelliD.[19] Third, environmental characteristics are sourced from NASA's Earthdata platform.[20] Fourth, climatic and surface features are obtained from the National Earth System Science Data Center.[21]

3.6 Variables

3.6.1 Dependent Variable

To account for variation in local resident density, we normalize the number of active EdAI users in each 1 km $\times$ 1 km grid cell by the resident population and multiply the ratio by 10,000 (10K). This

[16] Among the 192 grid cells in the main sample, 111 contain valid housing price data. The absence of housing price data does not imply a lack of residents; rather, it indicates the absence of commercial housing transactions in those areas. In China, commercial housing refers to market-oriented residential properties built and sold by private or state-owned developers, typically located in urbanized or newly developed zones. Areas without such transactions often consist of non-commercial housing, such as village residences, collective housing, or public institutional compounds, where property is not traded on the open market.

[17] The resulting dataset forms an unbalanced panel consisting of 877 observations over eight quarters across these 111 grid cells. The total number of observations is smaller than the theoretical maximum of 1,136 (= 142 $\times$ 8) because not every grid cell has housing transactions in each quarter.

[18] Source: https://ladsweb.modaps.eosdis.nasa.gov/

[19] Source: https://opencellid.org/

[20] Source: https://www.earthdata.nasa.gov/topics/land-surface/

[21] Source: http://wdcrre.data.ac.cn/

construction yields our main dependent variable, the number of unique active EdAI users per 10K residents in each grid cell (*EdAI*).

3.6.2 Control Variables

We include a set of control variables measured at the 1 km × 1 km grid-cell level. First, we control for nightlight intensity, derived from the VIIRS/NPP satellite, which serves as a proxy for local economic activity (*NightLight*). Second, we include the number of 4G communication towers within each cell to capture local digital infrastructure and network accessibility (*Towers*).[22] This variable reflects the availability and quality of mobile internet services, which directly affect users' ability to access EdAI. Third, we control for environmental characteristics, including the Normalized Difference Vegetation Index (*NDVI*), which measures surface greenness and can indicate urbanization levels. Fourth, we control for climatic and surface characteristics, including average precipitation (*Precipitation*), average air temperature (*Temperature*), surface albedo (*Albedo*),[23] and potential evapotranspiration (*PET*).[24] We show the descriptive statistics of the above-mentioned variables in Table 1.

3.7 Econometric Specification

3.7.1 Event-Study Spatial Difference-in-Differences Model

We first employ an event-study spatial difference-in-differences (SDID) model to estimate the dynamic effects of rising urban-rural educational inequality on rural users' EdAI usage. The model is specified as follows:

$$EdAI_{it} = \alpha + \sum_{k=1}^{8} \gamma_k Inequality_i \times \mathbf{1}(t = k) + \boldsymbol{\theta} \boldsymbol{X}_{it} + \mu_i + \tau_t + \varepsilon_{it} \quad (2)$$

[22] The numbers of 2G, 3G, and 5G communication towers show relatively little variation during the study period and are therefore omitted from the regression model.

[23] Surface albedo represents the fraction of sunlight reflected from the ground.

[24] All spatial datasets are matched to the 1 km × 1 km analytical grid using area-weighted interpolation to ensure spatial consistency.

Table 1: Descriptive Statistics of Panel Data in Main Analysis

		Obs.	Mean	S.D.	Min.	Max.
Dependent Variables						
EdAI	(Number of active users of EdAI in the grid cell/Number of residents in the grid cell) × 10,000	1,536	15.549	45.207	0.000	903.846
Independent Variables						
Inequality	= 1 for treated bank that lost eligibility to attend top-tier senior high schools; 0 otherwise	1,536	0.505	0.500	0.000	1.000
Control Variables						
NightLight	Nightlight intensity (nW/cm2/sr)	1,536	25.124	13.599	2.067	94.667
Towers	Number of 4G communication towers	1,536	4.643	6.949	0.000	34.000
NDVI	Normalized Difference Vegetation Index	1,536	0.388	0.129	-0.097	0.820
Precipitation	Average precipitation amount (mm)	1,536	112.120	73.195	36.470	252.099
Temperature	Average air temperature (°C)	1,536	21.524	5.782	13.050	30.082
Albedo	Fraction of sunlight reflected from the surface	1,536	0.060	0.008	0.029	0.081
PET	Potential Evapotranspiration (mm)	1,536	1021.155	32.444	942.250	1076.167

Notes: All variables are measured at the cell-quarter level, except for *Albedo* and *PET*, which are only available at the cell-year level. The cell-year data for *Albedo* and *PET* were merged into the cell-quarter panel based on the corresponding year of each quarter. The panel does not include any cross-year quarters.

where i indexes the grid cell and t indexes the quarter. $EdAI_{it}$ denotes the number of active EdAI users per 10k residents in cell i during quarter t. $Inequality_i$ is an indicator equal to 1 for cells located on the southwestern riverbank that lost eligibility to attend top-tier senior high schools after the policy and zero otherwise. $\mathbf{1}(t = k)$ is equal to 1 if calendar quarter t is in k-th quarter in our observation window ($t \in \{1, 2, ..., 8\}$). Our EdAI usage data span the period from July 2022 to June 2024, which corresponds to eight quarters. The fourth quarter in our dataset (i.e., the period from April to June 2023) is the policy shock quarter, as the policy was implemented in April 2023. We designate quarter $k = 1$ as the baseline and omit it from the estimation. Interaction terms $Inequality_i \times \mathbf{1}(t = k)$ capture time-specific treatment effects, with γ_k denoting the estimated coefficients. The vector $\boldsymbol{X}_{it}$ includes cell-quarter-varying control variables described in the previous subsection, and $\boldsymbol{\theta}$ is the corresponding coefficient vector. μ_i and τ_t represent cell and quarter fixed effects. ε_{it} is the error term. Standard errors are clustered at the cell level to account for serial correlation within units over time.

3.7.2 Non-Dynamic Spatial Difference-in-Differences Model

In addition to the event-study specification, we follow prior research (Datta et al., 2018; Norris and Xiong, 2023) in estimating average treatment effects over grouped post-policy periods using a non-dynamic SDID model. The specification is given by:

$$\begin{aligned} EdAI_{it} = \alpha &+ \gamma^S Inequality_i \times \mathbf{1}(t \in 4,5,6) \\ &+ \gamma^L Inequality_i \times \mathbf{1}(t \in 7,8) \\ &+ \boldsymbol{\theta} \boldsymbol{X}_{it} + \mu_i + \tau_t + \varepsilon_{it} \end{aligned} \tag{3}$$

This specification groups the post-policy periods into two intervals to capture average changes in EdAI usage following the rise in urban-rural educational inequality. The short-term period $\mathbf{1}(t \in 4,5,6)$ consists of quarters 4 through 6, which include the first quarter of policy implementation and the two subsequent quarters. The long-term period $\mathbf{1}(t \in 7,8)$ consists of quarters 7 and 8, which correspond to the 3rd and 4th quarters after implementation.[25] The coefficients γ^S and γ^L therefore measure the average short-term and long-term impacts of the policy-induced urban-rural educational inequality on EdAI usage, respectively. As in the event-study specification, the model includes cell-quarter-varying controls $\boldsymbol{X}_{it}$, the full set of cell fixed effects μ_i, and quarter fixed effects τ_t, with standard errors clustered at the grid-cell level.

4 Results

4.1 Model-Free Evidence

4.1.1 Cell-Level Model-Free Evidence

We begin by presenting the spatial distribution of EdAI usage at the grid cell level across four representative quarters in Figure 2. These grid-based choropleth maps depict the quarter immediately preceding the policy, the quarter of policy implementation, the first quarter following implementation, and the fourth quarter after implementation. Each map reports the number of EdAI users per 10K residents within each 1 km × 1 km grid cell, which allows us to visualize how

[25] We acknowledge that the grouping of post-policy periods into short-term and long-term intervals is chosen ex post, informed by the J-curve pattern observed in the event study estimates.

usage patterns evolve over space and time on the treated and control riverbanks.

We find that, in the pre-policy period, EdAI usage on the control riverbank is generally lower than on the treated riverbank. During the policy implementation quarter and the first post-policy quarter, EdAI usage declines in grid cells located on the treated riverbank, as reflected by the emergence of darker blue cells, while the spatial pattern on the control riverbank remains largely similar to that observed in the pre-policy period. By the fourth post-policy quarter, the treated riverbank exhibits a pronounced increase in EdAI usage, as indicated by the expansion of yellow-colored grid cells, reaching a substantially higher level relative to the pre-policy baseline. This evolution in spatial patterns provides preliminary visual evidence of a J-curve response, with an initial decline in EdAI usage on the treated riverbank followed by a strong rebound.

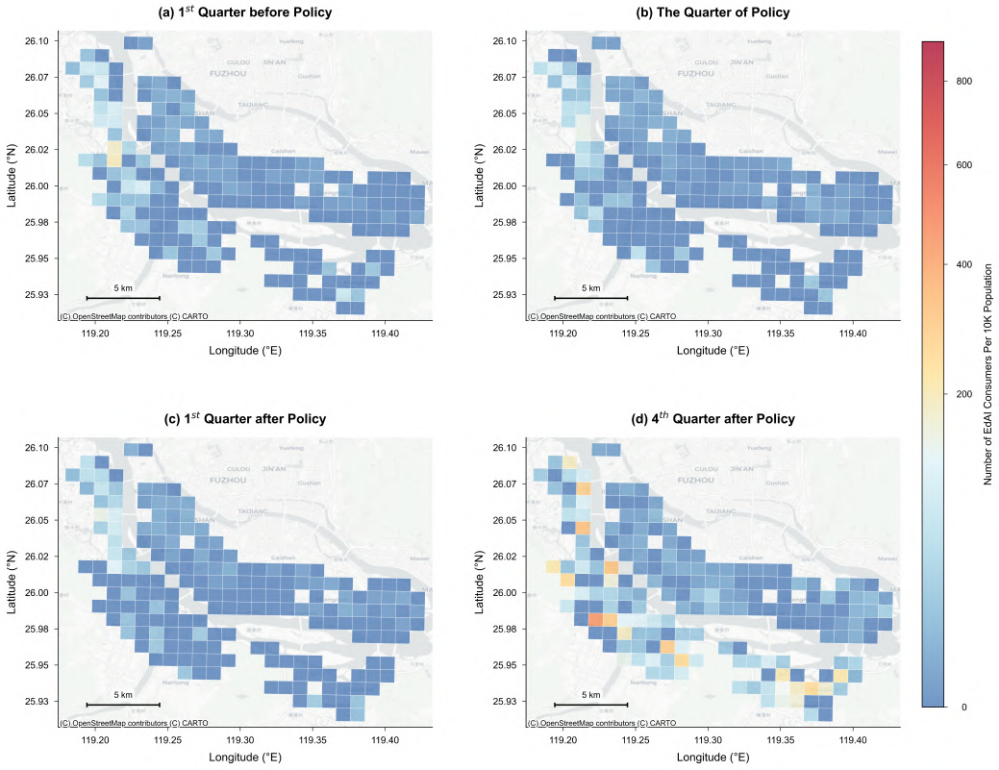


Figure 2: Cell-Level Model-Free Evidence

Notes: The figures illustrate the spatial distribution of EdAI usage within the 5 km corridor surrounding the river that defines the policy boundary in the main analysis.

4.1.2 Group-Level Model-Free Evidence

We next present a group-level model-free comparison of the number of active EdAI users per 10K residents in grid cells located on the treated and control riverbanks along the Min River, as shown in Figure 3. To emphasize relative temporal patterns, we follow prior research (Chen and Zhong, 2023) and demean the outcome by subtracting each cell's average value in the pre-policy quarters.

We find that, before the policy change, EdAI usage in treated and control cells followed a closely aligned trajectory. After the implementation of the local-enrollment policy, the trajectories diverged sharply. During the policy quarter and the following two quarters, treated cells experienced a pronounced decline in EdAI usage relative to control cells. Beginning in the third post-policy quarter, use in treated cells exceeded pre-policy levels and remained higher through the end of the sample. This group-level pattern provides additional model-free evidence of a J-curve pattern, which reflects an initial decline followed by a sustained increase in EdAI usage among users residing in the county that lost access to top-tier senior high schools.

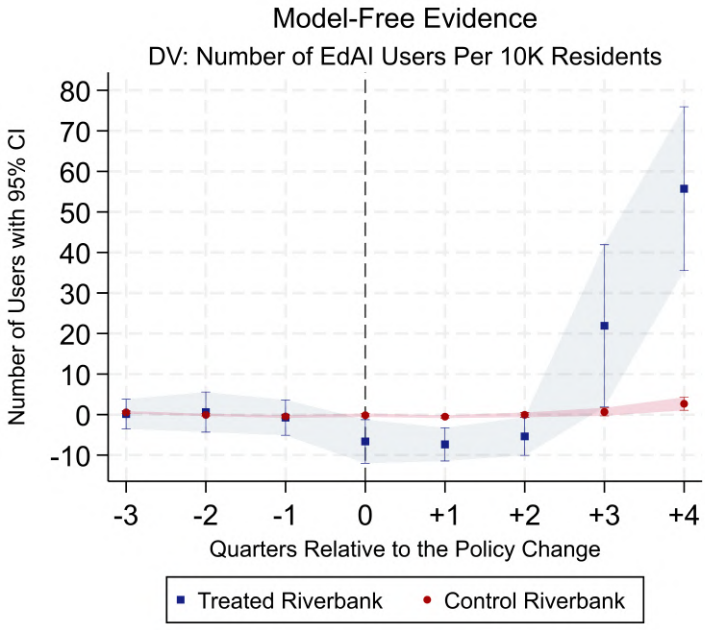


Figure 3: Group-Level Model-Free Evidence

Notes: The figures plot the demeaned number of EdAI users per 10k residents for cells in the treated riverbank and the control riverbank. The dashed vertical line marks the quarter in which the policy change occurred. The capped spikes represent the 95% confidence intervals.

4.2 Main Results

4.2.1 Estimation Results of Event-Study SDID Model

We then estimate Model 2, and the results are shown in Table 2 and Figure 4. Column (1) of Table 2 shows that EdAI usage in treated cells declined by about 5.166 users per 10K residents ($p > 0.05$) in the policy quarter and by 10.636 users per 10K residents ($p < 0.01$) in the first post-policy quarter. The negative effect diminished in the second post-policy quarter. EdAI usage then reversed direction and increased by 21.654 users per 10K residents ($p < 0.05$) in the third post-policy quarter and by 54.252 users per 10K residents ($p < 0.001$) in the fourth quarter, surpassing pre-policy levels. This pattern forms a clear J-curve pattern, with an initial short-term decline in EdAI usage among rural users followed by a longer-run increase.

4.2.2 Estimation Results of Non-Dynamic SDID Model

We show the estimation results of Model 3 in Column (2) of Table 2. We find that, EdAI usage in treated cells on average declined by 6.806 users per 10k residents ($p < 0.05$) in the short term (i.e., from the quarter of policy to 2^{nd} quarter after the policy) and by 37.899 ($p < 0.001$) in the long term (i.e., 3^{rd} to 4^{th} quarter after the policy). Given that the pre-policy mean number of EdAI users in treated cells was 20.368 per 10k residents, the short-term effect corresponds to a 33.4% (= 6.806 / 20.368) decline, followed by a 186.1% (= 37.899 / 20.368) rebound in the longer term.

4.3 Parallel Trends Assumption

To support a causal interpretation of the SDID estimates, we verify the parallel trends assumption, that in the absence of the policy, outcome trajectories between treatment and control groups would have evolved similarly. Estimates in Table 2 show no significant pre-policy differences, confirming this assumption. Moreover, we conduct an F-test to formally examine whether the pre-treatment coefficients are jointly equal to zero. The test yields an F-statistic of 0.020 with p-value of 0.976, indicating that we cannot reject the null hypothesis of joint insignificance for the pre-treatment period indicators. This result provides strong statistical evidence that the treatment and control groups exhibited parallel trends before the policy intervention, thereby validating the key identifying assumption of our spatial SDID framework.

Table 2: Main Results

	(1) *EdAI*	(2) *EdAI*
3rd Quarter before Policy	-baseline-	-baseline-
2nd Quarter before Policy	0.698	0.777
	(4.251)	(4.182)
1st Quarter before Policy	-0.236	-0.192
	(3.443)	(3.441)
the Quarter of Policy	-5.166	
	(3.276)	
1st Quarter after Policy	-10.636**	
	(3.657)	
2nd Quarter after Policy	-5.173	
	(3.358)	
3rd Quarter after Policy	21.654*	
	(10.298)	
4th Quarter after Policy	54.252***	
	(10.005)	
the Quarter of Policy to 2nd Quarter after Policy		-6.806*
		(2.892)
3rd to 4th Quarter after Policy		37.899***
		(7.912)
Controls	-included-	-included-
Cell FEs	-included-	-included-
Time FEs	-included-	-included-
Observations	1,536	1,536
Adjusted R^2	0.325	0.317

Notes: $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Robust standard errors clustered at the cell level are shown in parentheses.

4.4 Falsification Tests

4.4.1 Two Policy-Neutral River Boundaries

We further evaluate the validity of our main results using two falsification tests. The first test focuses on an upstream segment of the Min River where both riverbanks were subject to the same local enrollment policy restriction. As indicated by the blue single dot dashed line in Figure 1, although a river boundary exists in this area, neither bank experienced differential treatment because both belong to a county and were therefore equally affected by the enrollment restriction. We construct 1 km $\times$ 1 km grid cells along this upstream segment following the same procedure used in the main analysis. For each grid cell, we aggregate the number of resident users who

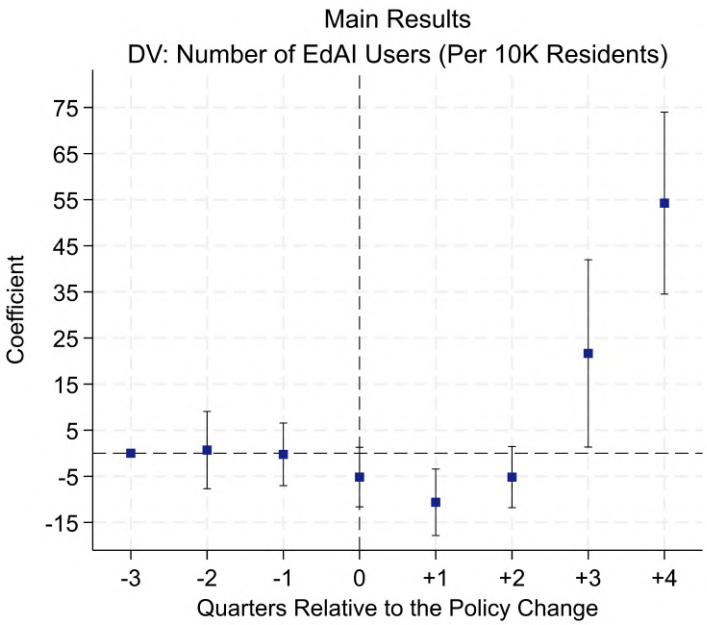


Figure 4: Main Results

Notes: The figure presents the event-study SDID coefficients estimated from Model 2 with 95% confidence intervals. The vertical dashed line indicates the policy quarter.

accessed any of the EdAI applications on a quarterly basis from July 2022 to June 2024 and normalize this measure by the resident population.[26] We then re-estimate Model 2 using grid cells located within a 5 km distance band along this upstream segment as an alternative sample. If the increase in urban-rural educational inequality indeed drives our main results, we should observe no significant differences in EdAI usage between the two sides of the river along this upstream section.

The second falsification test examines a nearby tributary where both banks were entirely unaffected by the local-enrollment policy. As shown in the blue double-dot-dashed line in Figure 1, the two riverbanks are both districts and thus students on both banks retained equal access to top-tier senior high schools. We re-estimate Model 2 using cells within a 5 km distance band along this tributary as another alternative sample. Similarly, because no cross-bank disparity in school eligibility existed in this setting, we should observe no significant differences in EdAI usage between the two banks if our main results are driven by the widened urban-rural educational

[26]We also construct all control variables for these grid cells following the same procedures described in the main analysis. The same procedures are applied to the grid cells used in the second falsification test and in the robustness check based on an alternative policy boundary.

inequality. We show the estimation results of both falsification tests in Figure 5. We find no significant treatment effects ($p > 0.05$) in either case, which supports the reliability of our main results.

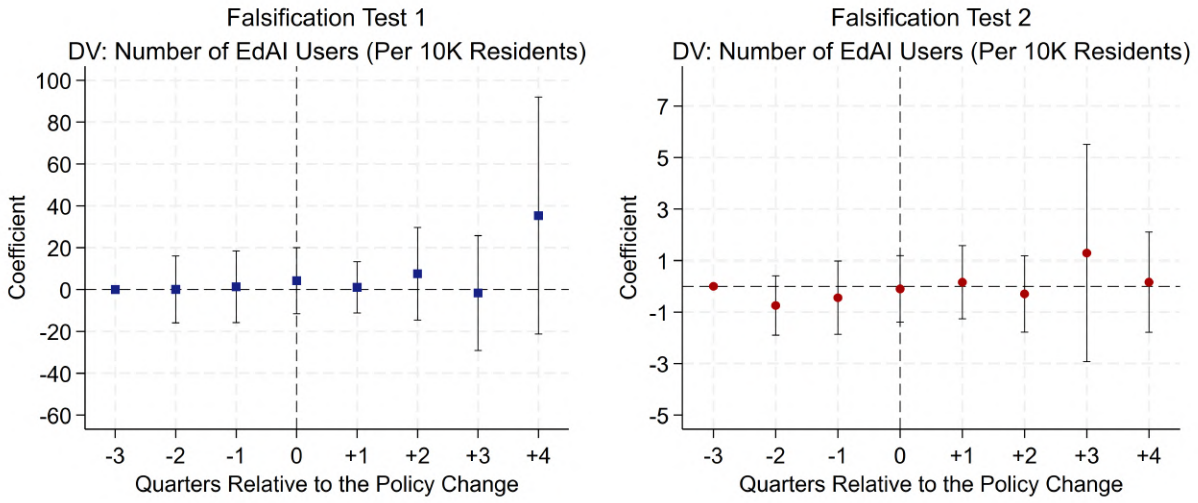


Figure 5: Falsification Tests

Notes: The figure presents the event-study SDID coefficients estimated from Model 2 with 95% confidence intervals. The vertical dashed line indicates the policy quarter.

4.4.2 Sub-sample Analysis Based on Policy Exemptions

We further exploit an exemption rule embedded in the policy to test the validity of our main results. The local enrollment policy in our setting includes an exemption clause for a narrowly defined group, as detailed in Appendix B. Specifically, students who attended Shangjie Experimental School and had at least one parent employed at a university were exempt from the local enrollment restrictions. To assess the reliability of our findings, we conduct a sub-sample analysis based on spatial proximity to Shangjie Experimental School (as indicated in Figure 6). The underlying assumption is that users residing closer to the school are more likely to belong to households eligible for the exemption. If the observed change in EdAI usage among treated cells is driven by the policy shock, we should observe minimal effects among users living nearer to Shangjie Experimental School.

We divide treated grid cells into two subgroups, namely cells located within 5 km of the

school and cells located more than 5 km away. We then re-estimate Model 2 for each sub-sample, combined with the full set of control cells. The results are presented in Figure 7. The estimated coefficients are statistically insignificant for grid cells within a 5 km radius ($p > 0.05$), which supports the validity of our main findings and suggests that the observed urban-rural usage response reflects the policy shock.

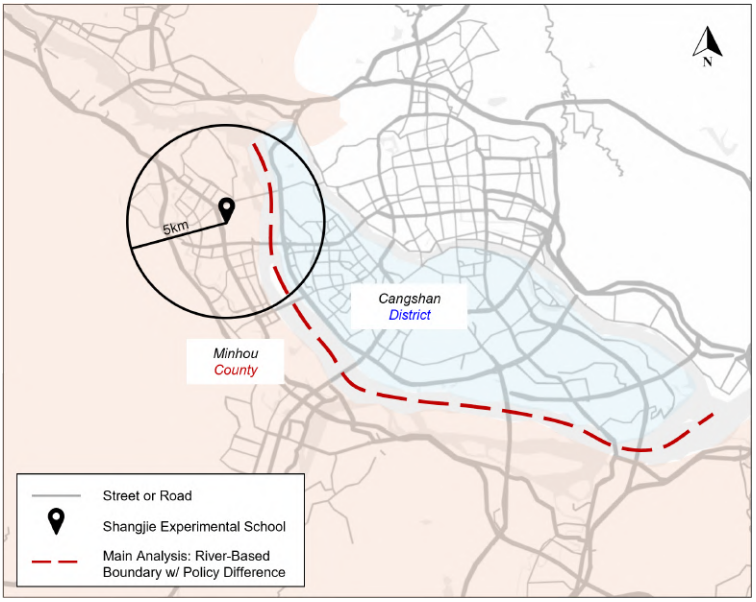


Figure 6: Location of Shangjie Experimental School

Notes: The figure shows the location of Shangjie Experimental School and the area within a 5 km radius.

4.5 Robustness Checks

4.5.1 Alternative Policy Boundary

We assess the robustness of our main findings by exploiting an alternative geographic boundary within the same city. As shown by the green dotted line in Figure 1, this land-based boundary separates Minhou County on the west from three districts on the east: Jin'an, Gulou, and Cangshan. Similar to the river-based boundary in our main analysis, after the implementation of the local-enrollment policy, students residing in the west-side county were ineligible to attend top-tier senior high schools located across this line. We re-estimate Model 2 using this alternative boundary to redefine treated and control cells and restrict the sample to a 5 km distance band along the boundary. If our main findings are indeed driven by the rising urban-rural educational inequality

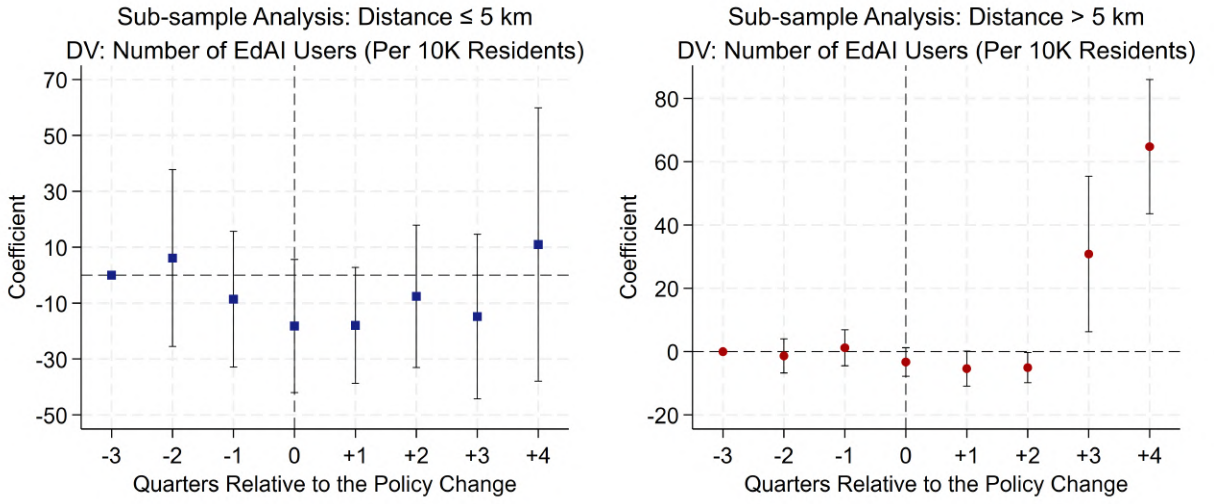


Figure 7: Exemption-Based Sub-sample Analysis

Notes: The figure presents the event-study SDID coefficients estimated from Model 2 with 95% confidence intervals. The vertical dashed line indicates the policy quarter.

after the implementation of the local-enrollment policy, we should observe a comparable J-curve pattern for rural users in the west-side county. Figure 8 reports the corresponding event-study estimates, which reproduce the same J-curve pattern documented in our main analysis, with an initial decline followed by a sustained rise in EdAI usage in the treated area. This finding reinforces the robustness of our results and suggests that the observed usage response is not specific to the river-based boundary used in the main analysis.

4.6 Other Robustness Checks

We further verify the robustness of our findings through a series of robustness checks.

First, we re-estimate Model 2 using alternative distance bands along the Min River. Following prior research (Norris and Xiong, 2023), we vary the bandwidth from three to nine kilometers on each side of the boundary and re-estimate the Model 2 for each spatial window. Across all distance bands, the post-policy dynamics consistently display the same J-curve pattern, indicating that our results are not sensitive to the choice of spatial bandwidth. The details of the results are shown in Appendix F.

Second, we implement a synthetic difference-in-differences estimator that constructs a

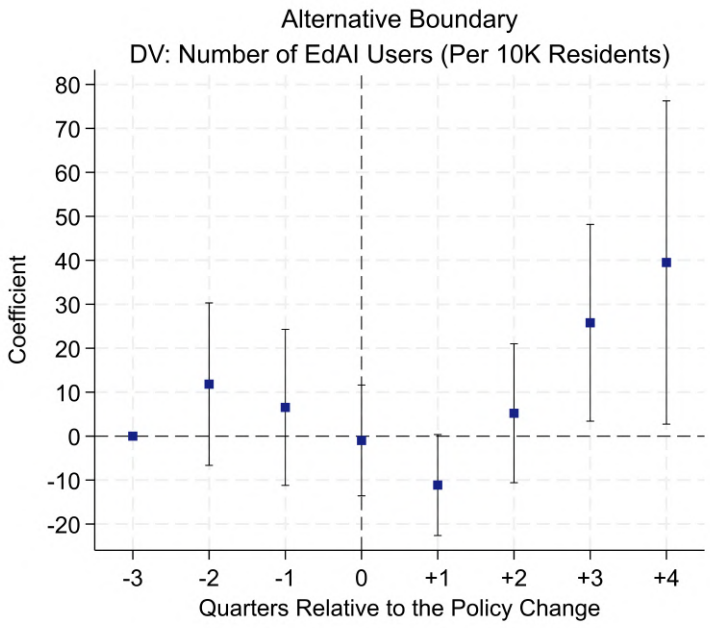


Figure 8: Robustness Check Based on an Alternative Boundary

Notes: The figure presents the event-study SDID coefficients estimated from Model 2 with 95% confidence intervals. The vertical dashed line indicates the policy quarter.

weighted combination of control grid cells to more closely match the pre-policy trajectory of treated cells (Arkhangelsky et al., 2021; Berman and Israeli, 2022; Farronato et al., 2024; Guyt et al., 2025; Pachali and Datta, 2025). As shown in Appendix G, the synthetic difference-in-differences estimates reproduce the same J-curve pattern observed in the main analysis. The magnitude of the long-run effect is comparable to that obtained from the main specification.

Third, we address potential concerns about spatial dependence in the error structure by re-estimating the model using standard errors that explicitly allow for spatial correlation across neighboring grid cells (Butts, 2023; Colella et al., 2019). As shown in Appendix H, the resulting event-study SDID estimates closely mirror those in the main analysis, suggesting that our conclusions are not sensitive to spatial correlation.

4.7 Ruling Out Alternative Explanations

4.7.1 Residents Re-sorting across Riverbanks

A potential alternative explanation for the observed effects is a change in the underlying pool of residents on either side of the river. The policy may have prompted selective movement or re-sorting of households that sought better educational opportunities for their children. Under this scenario, the short-term decline in EdAI usage in treated cells could reflect the relocation of families who are active EdAI users to the control bank rather than the user response we identify. As discussed in subsection 3.1, institutional constraints make relocation both costly and ineffective as a strategy for bypassing the policy, rendering large-scale self-selection through residential re-sorting highly improbable. We further consider this explanation unlikely for three reasons.

First, resident movement cannot explain the pronounced long-term rebound in EdAI usage within the treated cells. If the initial decline were driven by the out-migration of households with high EdAI usage, the treated riverbank would experience a persistent reduction in the number of EdAI users, and use levels would not subsequently rise above pre-policy baselines. The recovery and eventual increase in EdAI usage follow a J-curve pattern and are inconsistent with an explanation based solely on urban-rural re-sorting of residents.

Second, if the policy had triggered substantial movement of residents across the riverbanks, we would expect to observe corresponding adjustments in the local housing market (Bayer et al., 2007; Black, 1999). In particular, out-migration from the treated bank would reduce housing demand in that area, which would place downward pressure on local property prices. We re-estimate Model 2 using quarterly housing prices as the dependent variable. The results, presented in Figure I1 of Appendix I, show no significant differences in both listing and transaction housing prices between the treated and control riverbanks before or after the policy change. This evidence supports the conclusion that cross-river relocation is unlikely to explain the observed J-curve pattern in EdAI usage after the policy shock.

Third, to address the possibility that the results are driven by cross-boundary re-sorting rather than by the policy shock itself, we conduct a donut-type estimation that removes the grid cells located closest to the policy boundary. This procedure is a common approach for mitigating concerns that users may cross the boundary (Dube et al., 2010; Ehrlich and Seidel, 2018; Jia et al.,

2021; Kline and Moretti, 2014; Norris and Xiong, 2023). As shown in Figure I2 of Appendix I, we exclude the inner-most grid cells within one kilometer and two kilometers of the boundary from the main analysis sample and re-estimate Model 2. The resulting estimates continue to follow the same J-curve pattern as in the baseline results. This consistency suggests that the main findings are not driven by immediate cross-boundary re-sorting.[27]

4.7.2 Changes in Mobile Network Infrastructure

A potential alternative explanation is that the observed effects are driven by changes in mobile network infrastructure across the two riverbanks, particularly the expansion or contraction of communication towers. For example, one might argue that the short-term decrease in EdAI usage reflects a temporary reduction in network capacity if communication towers were removed or upgraded on the treated riverbank, and that the subsequent long-term increase in EdAI usage could be the result of a later expansion of tower coverage or improvements in network quality that enhanced access to mobile applications. Under this interpretation, the observed pattern in EdAI usage would arise from fluctuations in mobile connectivity rather than from policy-induced urban-rural educational inequality. We consider this explanation unlikely for two reasons.

First, the number of 4G towers in each grid cell is already included as a control variable in all baseline specifications, which accounts for differences in local network capacity.

Second, the data provide no evidence that changes in tower deployment could generate the observed pattern in EdAI usage. The number of 4G towers remains essentially flat on both riverbanks throughout the study period, and the only noticeable variation is a small increase in the control riverbank during the second quarter after the policy change. If variation in network infrastructure were driving the results, this increase should coincide with a rise in EdAI usage among users in the control group beginning in that quarter. However, we find that EdAI usage on the control riverbank remains stable over time, with no upward shift corresponding to the

[27] We acknowledge an important limitation in interpreting these donut-type estimates. Existing applications of this approach typically operate over larger geographic units, such as provincial or state borders, where removing observations near the boundary creates a meaningful buffer and sharply reduces the scope for cross-boundary re-sorting (Dube et al., 2010; Ehrlich and Seidel, 2018; Jia et al., 2021; Kline and Moretti, 2014; Norris and Xiong, 2023). In contrast, our study examines a much narrower urban-rural riverbank setting. Excluding only one or two kilometers on either side of the boundary can only mitigate the bias from spillover behavior that occurs across the riverbanks. For this reason, the donut estimates should be viewed as supportive robustness checks rather than as definitive evidence for excluding the alternative explanation.

timing of the tower increase. These observations indicate that changes in communication tower deployment are unlikely to explain the J-curve pattern in EdAI usage.

5 Underlying Mechanism

We propose that the mechanism underlying the observed J-curve pattern operates through perceived control as users respond to increasing educational inequality. We argue that, when faced with the restricted educational opportunities created by the local enrollment policy, users in rural areas likely perceive the policy shock as a threat to their future opportunity of growth (Folkman, 1984). Crucially, because they possess limited control over their educational circumstances, they are more likely to adopt emotion-focused coping strategies, such as avoidance, rather than problem-focused approaches that involve active attempts to address the stressor in the very beginning of policy shock (Carver et al., 1989). This emotion-focused coping manifests as an initial decline in EdAI usage. Over time, as they reappraise their circumstances and recognize EdAI as a compensatory resource that can restore some degree of control, they shift toward problem-focused coping, leading to a rebound in use (Ko et al., 2023).

We argue that, if this mechanism holds, socioeconomic status (SES) should moderate rural users' behavioral responses to educational inequality. To be specific, users from high-SES backgrounds generally possess greater perceived control over their circumstances, which shapes how they cope with constraint and uncertainty (Caplan and Schooler, 2007). In high-SES areas, access to richer educational resources and stronger support networks (Holahan and Moos, 1987) reinforces a sense of control over educational outcomes, dampening the magnitude of behavioral change (Carver et al., 1989; Terry and Hynes, 1998). Consequently, we expect users in high-SES areas to display relatively flat responses, while users in low-SES areas display a deeper short-run decline and a sharper rebound in EdAI usage.

We first operationalize SES using housing prices as a resource-based indicator. Housing prices capture accumulated household wealth and local economic security, which buffer against shocks and facilitate access to educational resources (Krieger et al., 1997). They also reflect neighborhood-level socioeconomic conditions (Matthews and Gallo, 2011; Krieger et al., 1997). To test the mechanism, we calculate the average cell-level transaction housing price before the pol-

icy shock and divide the sample of treated grid cells into high-SES and low-SES subgroups based on the median price among treated cells. We then re-estimate Model 2 for each subgroup, pairing each with the full set of control cells. If perceived control indeed shapes behavioral adjustment, the J-curve pattern should be more pronounced among users in low-SES areas and relatively flat among users in high-SES areas. The estimation results, shown in Figure 9, confirm this prediction: the J-curve pattern is stronger among users in areas with housing prices below the median, suggesting that users in low-SES areas exhibit more pronounced compensatory adaptation to the rise of urban-rural educational inequality.

In Appendix J, we further operationalize SES using town-level industrial output of enterprises above a designated size per capita and again find similar results, which strengthens the robustness of this mechanism.

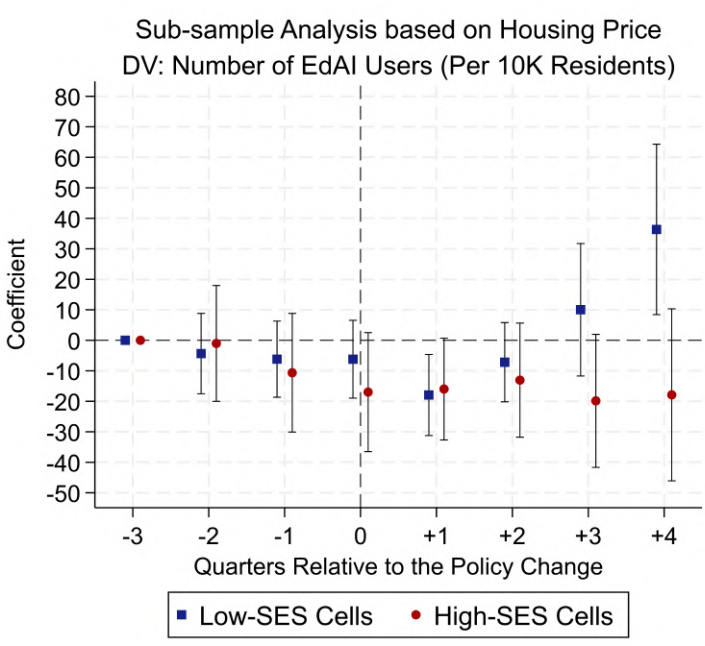


Figure 9: Sub-sample Analysis by SES Proxied by Housing Price

Notes: The figure presents the event-study SDID coefficients estimated from Model 2 with 95% confidence intervals. The vertical dashed line indicates the policy quarter.

6 Discussion

6.1 Summary of Findings

As urban-rural educational inequality continues to widen (Byun et al., 2012; van Maarseveen, 2021; Zahl-Thanem and Rye, 2024; Agrawal, 2014; Li et al., 1200; Lu, 2012; Zhang, 2017), societies face growing challenges in ensuring that rural users have equitable access to high-quality educational resources. These disparities have fueled interest in EdAI as a potential tool for closing urban-rural educational gaps. Yet this equalizing promise depends critically on how rural users engage with such technologies. If the use of EdAI becomes concentrated among urban users, EdAI could inadvertently amplify rather than reduce educational inequality. Exploiting a natural experiment in a major Chinese city, this study provides causal evidence on how a sudden increase in urban-rural educational inequality affects EdAI usage among rural users. We find that EdAI usage in rural areas fell sharply immediately after the policy shock but subsequently recovered and rose above its pre-policy baseline within one year, generating a J-curve pattern.

6.2 Theoretical Contributions

Our findings contribute to two lines of research.

First, we extend the literature on information technology and educational inequality by shifting the analytical focus from access to adaptation. Existing work has largely asked whether access to information technologies mitigates or exacerbates inequality across social groups (Anderson and Kimball, 2019; Bianchi et al., 2022; Figlio et al., 2013; Gao et al., 2021; Goodman et al., 2019; Keppler et al., 2022; Zheng et al., 2024). In contrast, we examine educational inequality as a driving force that shapes technology use. By treating educational technology engagement as an endogenous behavioral response to educational inequality, we highlight how rural users actively change their interaction with EdAI when structural barriers intensify.

Second, we contribute to research on human responses to inequality by identifying a technology-related behavioral change. Prior studies have primarily examined behaviors that are not related to technology (Bamberger et al., 2021; Bao et al., 2024; Browman et al., 2017; Mostagir and Siderius, 2023; Yoon and Kim, 2018; Zhong et al., 2025). Our findings reveal that widened urban-rural educational inequality also changes people's technology-related behavior, inducing a

temporal J-curve pattern in EdAI usage. In doing so, our study integrates technology use into the broader framework of human responses to inequality.

6.3 Practical Implications

Our findings carry important implications for both managers and policymakers.

For managers of EdAI products or companies, the results highlight the importance of recognizing temporal dynamics in user responses to large shocks in educational inequality. The short-term decline in EdAI usage following the policy change does not reflect permanent disengagement but instead appears to be part of an adjustment process that culminates in a strong rebound. Accordingly, business strategies should be designed to maintain user engagement through the initial downturn, anticipate delayed adoption or re-engagement, and position EdAI products to meet emerging demand as users adapt to new educational constraints. Such strategies may include communicating with users to reduce uncertainty, highlighting features that help compensate for resource disadvantages, and timing promotional efforts to align with the recovery phase. Moreover, our mechanism analyses provide insights that can inform the long-term targeting of EdAI products toward users from low-SES backgrounds, for whom the rebound in usage is particularly pronounced.

Second, for policymakers, the observed J-curve pattern suggests that when educational inequality rises, disadvantaged users may initially withdraw from EdAI before later re-engaging and eventually exceeding their baseline usage levels. Policymakers should therefore anticipate the temporal lag inherent in this adaptation process and design interventions that support users through the transition. From this perspective, policy efforts should emphasize resilience building by creating conditions that help disadvantaged users move through the initial decline and accelerate recovery. This may include providing targeted support during the early stage of the J-curve, such as subsidized access to digital devices and localized AI literacy programs that help disadvantaged users sustain motivation until their usage rebounds. Moreover, our mechanism analyses indicate that interventions targeted toward low-SES communities are likely to be particularly effective, as these users exhibit the strongest initial decline and subsequent recovery in EdAI usage.

6.4 Limitations

There are several limitations to this work.

First, our analysis relies on cellphone-based data on EdAI usage, which, while granular and behaviorally rich, may not capture all forms of educational technology engagement. Students might also use offline AI systems that are not reflected in our dataset, potentially leading to an underestimation of overall adaptation. Future research could combine multiple data sources, such as offline system logs and survey evidence, to validate and extend these behavioral insights.

Second, our study focuses on a single city context in China. Cultural, institutional, and technological factors shape how educational inequality interacts with EdAI use, so generalizing the J-curve pattern to other areas or countries requires caution.

Third, while coping theory provides a useful psychological lens for interpreting behavioral adaptation, our analysis is observational and does not directly measure psychological states in coping processes. Future work can integrate survey or experimental measures of coping processes to enrich our understanding of the mechanisms underlying the observed J-curve pattern.

References

Agostinelli, F., Doepke, M., Sorrenti, G., and Zilibotti, F. (2022). When the Great Equalizer Shuts down: Schools, Peers, and Parents in Pandemic Times. *Journal of Public Economics*, 206:104574.

Agrawal, T. (2014). Educational Inequality in Rural and Urban India. *International Journal of Educational Development*, 34:11–19.

Ahluwalia, M. S. (1976). Inequality, Poverty and Development. *Journal of Development Economics*, 3(4):307–342.

Anderson, S. W. and Kimball, A. (2019). Evidence for the Feedback Role of Performance Measurement Systems. *Management Science*, 65(9):4385–4406.

Andretti, B., Vieites, Y., Elmor, L., and Andrade, E. B. (2025). How Socioeconomic Status Shapes Food Preferences and Perceptions. *Journal of Marketing*, 89(6):33–56.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic Difference-in-Differences. *American Economic Review*, 111(12):4088–4118.

Bacher-Hicks, A., Goodman, J., and Mulhern, C. (2021). Inequality in Household Adaptation to Schooling Shocks: Covid-Induced Online Learning Engagement in Real Time. *Journal of Public Economics*, 193:104345.

Balaj, M., Henson, C. A., Aronsson, A., Aravkin, A., Beck, K., Degail, C., Donadello, L., Eikemo, K., Friedman, J., Giouleka, A., Gradeci, I., Hay, S. I., Jensen, M. R., Mclaughlin, S. A., Mullany, E. C., O'connell, E. M., Sripada, K., Stonkute, D., Sorensen, R. J. D., Solhaug, S., Vonen, H. D., Westby, C., Zheng, P., Mohammad, T., Eikemo, T. A., and Gakidou, E. (2024). Effects of Education on Adult Mortality: A Global Systematic Review and Meta-Analysis. *Lancet Public Health*, 9(3):e155–e165.

Bamberger, B., Homburg, C., and Wielgos, D. M. (2021). Wage Inequality: Its Impact on Customer Satisfaction and Firm Performance. *Journal of Marketing*, 85(6):24–43.

Bao, W., Ni, J., and Singh, S. (2024). Educational Inequality and Reservation Policy in Developing Markets. *Management Science*, 70(5):3162–3181.

Bayer, P., Ferreira, F., and McMillan, R. (2007). A Unified Framework for Measuring Preferences for Schools and Neighborhoods. *Journal of Political Economy*, 115(4):588–638.

Belo, R., Ferreira, P., and Telang, R. (2014). Broadband in School: Impact on Student Performance. *Management Science*, 60(2):265–282.

Berman, R. and Israeli, A. (2022). The Value of Descriptive Analytics: Evidence from Online Retailers. *Marketing Science*, 41(6):1074–1096.

Bianchi, N., Lu, Y., and Song, H. (2022). The Effect of Computer-Assisted Learning on Students' Long-Term Development. *Journal of Development Economics*, 158:102919.

Black, S. E. (1999). Do Better Schools Matter? Parental Valuation of Elementary Education. *Quarterly Journal of Economics*, 114(2):577–599.

Blanden, J., Doepke, M., and Stuhler, J. (2023). Educational Inequality. In *Handbook of the Economics of Education*, volume 6, pages 405–497. Elsevier.

Brosnan, S. F. and de Waal, F. B. M. (2014). Evolution of Responses to (un)fairness. *Science*, 346(6207):1251776.

Browman, A. S., Destin, M., Carswell, K. L., and Svoboda, R. C. (2017). Perceptions of Socioeconomic Mobility Influence Academic Persistence Among Low Socioeconomic Status Students. *Journal of Experimental Social Psychology*, 72:45–52.

Browman, A. S., Destin, M., Kearney, M. S., and Levine, P. B. (2019). How Economic Inequality Shapes Mobility Expectations and Behaviour in Disadvantaged Youth. *Nature Human Behaviour*, 3(3):214–220.

Butts, K. (2023). Difference-in-Differences Estimation with Spatial Spillovers.

Byun, S.-y., Meece, J. L., and Irvin, M. J. (2012). Rural-Nonrural Disparities in Postsecondary Educational Attainment Revisited. *American Educational Research Journal*, 49(3):412–437.

Caplan, L. J. and Schooler, C. (2007). Socioeconomic Status and Financial Coping Strategies: The Mediating Role of Perceived Control. *Social Psychology Quarterly*, 70(1):43–58.

Carver, C. S., Scheier, M. F., and Weintraub, J. K. (1989). Assessing Coping Strategies: A Theoretically Based Approach. *Journal of Personality and Social Psychology*, 56(2):267–283.

Chen, N. and Zhong, Z. Z. (2023). History and Country-of-Origin Effects. *Marketing Science*.

Chen, Q., Wang, Y., and Zhang, Y. (2024). Developing Strengths or Remedying Weaknesses? How Perceived Social Mobility Affects Parents' Purchase Preferences for Children's Educational Products. *Journal of Marketing*, 88(5):46–62.

Chen, Y., Mittal, V., and Sridhar, S. H. (2021). Investigating the Academic Performance and Disciplinary Consequences of School District Internet Access Spending. *Journal of Marketing Research*, 58(1):141–162.

Chen, Z., Alcorn, B., Christensen, G., Eriksson, N., Koller, D., and Emanuel, E. J. (2015). Who's Benefiting from MOOCs, and Why. *Harvard Business Review*.

Chetty, R., Friedman, J. N., and Rockoff, J. E. (2014). Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood Per 10k. *American Economic Review*, 104(9):2633–2679.

Choudhury, S. (2025). Why Educational Inequality Remains High in India.

Clark, D. and Royer, H. (2013). The Effect of Education on Adult Mortality and Health: Evidence from Britain. *American Economic Review*, 103(6):2087–2120.

Cohn, A., Fehr, E., Herrmann, B., and Schneider, F. (2014). Social Comparison and Effort Provision: Evidence from a Field Experiment. *Journal of the European Economic Association*, 12(4):877–898.

Colella, F., Lalive, R., Sakalli, S. O., and Thoenig, M. (2019). Inference with Arbitrary Clustering. Technical report, IZA - Institute of Labor Economics.

Datta, H., Knox, G., and Bronnenberg, B. J. (2018). Changing Their Tune: How Consumers' Adoption of Online Streaming Affects Music Consumption and Discovery. *Marketing Science*, 37(1):5–21.

Davies, N. M., Dickson, M., Davey Smith, G., van den Berg, G. J., and Windmeijer, F. (2018). The Causal Effects of Education on Health Outcomes in the UK Biobank. *Nature Human Behaviour*, 2(2):117–125.

Dube, A., Lester, T. W., and Reich, M. (2010). Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties. *Review of Economic Studies*, 92(4):945–964.

Duflo, E., Dupas, P., and Kremer, M. (2011). Peer Effects, Teacher Incentives, and the Impact of Tracking: Evidence from a Randomized Evaluation in Kenya. *American Economic Review*, 101(5):1739–1774.

Ehrlich, M. and Seidel, T. (2018). The Persistent Effects of Place-Based Policy: Evidence from the West-German Zonenrandgebiet. *American Economic Journal: Economic Policy*, 10(4):344–374.

Fairlie, R. W. and London, R. A. (2012). The Effects of Home Computers on Educational Outcomes: Evidence from a Field Experiment with Community College Students. *Economic Journal*, 122(561):727–753.

Fairlie, R. W. and Robinson, J. (2013). Experimental Evidence on the Effects of Home Computers on Academic Achievement among Schoolchildren. *American Economic Journal: Applied Economics*, 5(3):211–240.

Farronato, C., Fong, J., and Fradkin, A. (2024). Dog Eat Dog: Balancing Network Effects and Differentiation in a Digital Platform Merger. *Management Science*, 70(1):464–483.

Ferreira, F. H. G. and Gignoux, J. (2014). The Measurement of Educational Inequality: Achievement and Opportunity. *World Bank Economic Review*, 28(2):210–246.

Figlio, D., Rush, M., and Yin, L. (2013). Is It Live or Is It Internet? Experimental Estimates of the Effects of Online Instruction on Student Learning. *Journal of Labor Economics*, 31(4):763–784.

Folkman, S. (1984). Personal Control and Stress and Coping Processes: A Theoretical Analysis. *Journal of Personality and Social Psychology*, 46(4):839–852.

Folkman, S. (2013). Stress: Appraisal and Coping. In *Encyclopedia of Behavioral Medicine*, pages 1913–1915. Springer, New York, NY.

Fu, R., Huang, Y., Mehta, N., Singh, P. V., and Srinivasan, K. (2025). Unequal Impact of Zestimate on the Housing Market. *Marketing Science*, 44(6):1407–1427.

Gao, Q., Lin, M., and Wu, D. J. (2021). Education Crowdfunding and Student Performance: An Empirical Study. *Information Systems Research*, 32(1):53–71.

Glaeser, E. (2013). A Review of Enrico Moretti's The New Geography of Jobs. *Journal of Economic Literature*, 51(3):825–837.

Glaeser, E. L., Kallal, H. D., Scheinkman, J. A., and Shleifer, A. (1992). Growth in Cities. *Journal of Political Economy*, 100(6):1126–1152.

Goodman, J., Melkers, J., and Pallais, A. (2019). Can Online Delivery Increase Access to Education? *Journal of Labor Economics*, 37(1):1–34.

Goolsbee, A. and Guryan, J. (2006). The Impact of Internet Subsidies in Public Schools. *Review of Economic Studies*, 88(2):336–347.

Guyt, J. Y., van Lin, A., and Keller, K. O. (2025). Banning Unsolicited Store Flyers: Does Helping the Environment Hurt Retailing? *Marketing Science*, 44(5):1104–1124.

Hanushek, E. A. and Woessmann, L. (2008). The Role of Cognitive Skills in Economic Development. *Journal of Economic Literature*, 46(3):607–668.

Heckman, J. J. and Masterov, D. V. (2007). The Productivity Argument for Investing in Young

Children. *Applied Economic Perspectives and Policy*, 29(3):446–493.
Heckman, J. J., Moon, S. H., Pinto, R., Savelyev, P. A., and Yavitz, A. (2010). The Rate of Return to the HighScope Perry Preschool Program. *Journal of Public Economics*, 94(1):114–128.
Hendel, I., Shapiro, J., and Willen, P. (2005). Educational Opportunity and Income Inequality. *Journal of Public Economics*, 89(5):841–870.
Holahan, C. J. and Moos, R. H. (1987). Personal and Contextual Determinants of Coping Strategies. *Journal of Personality and Social Psychology*, 52(5):946–955.
IndiaToday (2025). India's First AI-Powered Anganwadi Brings Digital Learning to Rural Children. https://www.indiatoday.in/education-today/news/story/indias-first-ai-powered-anganwadi-brings-digital-learning-to-rural-children-2765872-2025-08-04.
Jackson, M. and Holzman, B. (2020). A Century of Educational Inequality in the United States. *Proceedings of the National Academy of Sciences*, 117(32):19108–19115.
Jia, J., Liang, X., and Ma, G. (2021). Political Hierarchy and Regional Economic Development: Evidence from a Spatial Discontinuity in China. *Journal of Public Economics*, 194:104352.
Keppler, S. M., Li, J., and Wu, D. A. (2022). Crowdfunding the Front Lines: An Empirical Study of Teacher-Driven School Improvement. *Management Science*, 68(12):8809–8828.
Kline, P. and Moretti, E. (2014). Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority. *Quarterly Journal of Economics*, 129(1):275–331.
Ko, G. Y., Shin, D., Auh, S., Lee, Y., and Han, S. P. (2023). Learning Outside the Classroom During a Pandemic: Evidence from an Artificial Intelligence-Based Education App. *Management Science*, 69(6):3616–3649.
Krieger, N., Williams, D. R., and Moss, N. E. (1997). Measuring Social Class in US Public Health Research: Concepts, Methodologies, and Guidelines. *Annual Review of Public Health*, 18:341–378.
Leuven, E., Lindahl, M., Oosterbeek, H., and Webbink, D. (2007). The Effect of Extra Funding for Disadvantaged Pupils on Achievement. *Review of Economic Studies*, 89(4):721–736.
Li, H., Loyalka, P., Wu, B., Xie, J., and Rozelle, S. (Tue, 05/22/2018 - 12:00). Unequal Access to College in China: How Far Have Poor, Rural Students Been Left Behind? *China Quarterly*.
Li, H., Meng, L., Shi, X., and Wu, B. (2012). Does Attending Elite Colleges Pay in China? *Journal of Comparative Economics*, 40(1):78–88.
Liu, Y. (2013). Meritocracy and the Gaokao: A Survey Study of Higher Education Selection and Socio-Economic Participation in East China. *British Journal of Sociology of Education*, 34(5-6):868–887.
Lu, Y. (2012). Education of Children Left Behind in Rural China. *Journal of Marriage and Family*, 74(2):328–341.
Matthews, K. A. and Gallo, L. C. (2011). Psychological Perspectives on Pathways Linking Socioeconomic Status and Physical Health. *Annual Review of Psychology*, 62:501–530.
McCall, L., Burk, D., Laperrière, M., and Richeson, J. A. (2017). Exposure to Rising Inequality Shapes Americans' Opportunity Beliefs and Policy Support. *Proceedings of the National Academy of Sciences*, 114(36):9593–9598.
Moretti, E. (2012). *The New Geography of Jobs*. Houghton Mifflin Harcourt.
Mostagir, M. and Siderius, J. (2023). Social Inequality and the Spread of Misinformation. *Management Science*, 69(2):968–995.
Neidhöfer, G., Serrano, J., and Gasparini, L. (2018). Educational Inequality and Intergenerational Mobility in Latin America: A New Database. *Journal of Development Economics*, 134:329–349.
Norris, J. J. and Xiong, H. (2023). Ride-Sharing and the Geography of Consumption Industries. *Economic Journal*, 133(654):2449–2482.

Pachali, M. J. and Datta, H. (2025). What Drives Demand for Playlists on Spotify? *Marketing Science*, 44(1):54–64.

Payne, B. K., Brown-Iannuzzi, J. L., and Hannay, J. W. (2017). Economic Inequality Increases Risk Taking. *Proceedings of the National Academy of Sciences*, 114(18):4643–4648.

Pswarayi-Riddihough, I. and Klapper, L. (2025). Mobile Phone Ownership Is Widespread. Why Is Digital Inclusion Still Lagging? https://blogs.worldbank.org/en/voices/Mobile-phone-ownership-is-widespread-Why-is-digital-inclusion-still-lagging.

Sands, M. L. (2017). Exposure to Inequality Affects Support for Redistribution. *Proceedings of the National Academy of Sciences*, 114(4):663–668.

Terry, D. J. and Hynes, G. J. (1998). Adjustment to a Low-Control Situation: Reexamining the Role of Coping Responses. *Journal of Personality and Social Psychology*, 74(4):1078–1092.

UNESCO (2024). Can AI Close the Learning Gap in Thailand's Schools? https://www.unesco.org/en/articles/can-ai-close-learning-gap-thailands-schools.

van Maarseveen, R. (2021). The Urban-Rural Education Gap: Do Cities Indeed Make Us Smarter? *Journal of Economic Geography*, 21(5):683–714.

van Wolkenten, M., Brosnan, S. F., and de Waal, F. B. M. (2007). Inequity Responses of Monkeys Modified by Effort. *Proceedings of the National Academy of Sciences*, 104(47):18854–18859.

Xinhua (2025). AI Opens up New Opportunities as It Closes Urban-Rural Education Gap. https://www.chinadaily.com.cn/a/202504/03/WS67edcb0ca3104d9fd381d5f5.html.

Yoon, S. and Kim, H. C. (2018). Feeling Economically Stuck: The Effect of Perceived Economic Mobility and Socioeconomic Status on Variety Seeking. *Journal of Consumer Research*, 44(5):1141–1156.

Young, A. (2013). Inequality, the Urban-Rural Gap, and Migration. *Quarterly Journal of Economics*, 128(4):1727–1785.

Zahl-Thanem, A. and Rye, J. F. (2024). Spatial Inequality in Higher Education: A Growing Urban–Rural Educational Gap? *European Sociological Review*, 40(6):1067–1081.

Zhang, H. (2017). Opportunity or New Poverty Trap: Rural-Urban Education Disparity and Internal Migration in China. *China Economic Review*, 44:112–124.

Zhang, J., Zhao, Y., Park, A., and Song, X. (2005). Economic Returns to Schooling in Urban China, 1988 to 2001. *Journal of Comparative Economics*, 33(4):730–752.

Zhang, S., Mehta, N., Singh, P. V., and Srinivasan, K. (2021). Frontiers: Can an Artificial Intelligence Algorithm Mitigate Racial Economic Inequality? An Analysis in the Context of Airbnb. *Marketing Science*, 40(5):813–820.

Zheng, Z., Jiang, S., and Cui, R. (2024). Remote Learning and Educational Inequality. *Production and Operations Management*, 33(12):2421–2439.

Zhong, H., Zhong, Z. Z., and Chen, N. (2025). Gender Inequality And Household Purchase Decisions: The Case of Automobiles in China.

Zimmerman, D. J. (2003). Peer Effects in Academic Outcomes: Evidence from a Natural Experiment. *Review of Economic Studies*, 85(1):9–23.

Online Appendix to

"Adaptation of AI Use to Educational Inequality"

A China's Administrative Systems and Details of Fuzhou City

China's administrative system follows a five-tiered structure that begins with the provincial level, which includes provinces, autonomous regions, municipalities, and special administrative regions. Below this level are prefecture-level units such as cities and prefectures. These are followed by county-level units that include districts and counties, then township-level units that include towns, townships, and subdistricts, and finally village-level units.

Fuzhou administers several types of county-level units, including both districts and counties. Fuzhou has six districts, namely Gulou, Taijiang, Cangshan, Jin'an, Mawei, and Changle, and seven counties, namely Fuqing, Minhou, Lianjiang, Luoyuan, Minqing, Pingtan, and Yongtai (as shown in Figure A1). Although districts and counties are classified at the same county-level administrative tier in China, they often differ systematically in (1) urbanization, (2) economic structure, and (3) educational resources.

We manually collect data from the 2024 Fuzhou Statistical Yearbook[1] to characterize and compare districts and counties. For urbanization, the average urbanization rate across Fuzhou's districts is 91.6%, compared with 56.4% for counties. Focusing on the boundary areas used in our main analysis, Cangshan District is fully urbanized, with an urbanization rate of 100%, whereas Minhou County has an urbanization rate of only 61.2%.

Economic activity also differs markedly between districts and counties. We aggregate gross domestic product (GDP) across all districts and all counties within Fuzhou. The total GDP of the districts amounts to 776.076 billion yuan, whereas the combined GDP of the counties is approximately 516.772 billion yuan.[2] In the districts, output from the tertiary sector reaches 541.058 billion yuan, accounting for 69.72% of total GDP, and clearly dominates the local economy. The secondary sector contributes 224.280 billion yuan, or 28.90%, while the primary sector accounts for only 10.740 billion yuan, or 1.38%. This industrial structure suggests a labor market concentrated in service activities such as finance and technology. In contrast, the counties exhibit a markedly different economic structure. Output from the tertiary sector totals 212.118 billion yuan, representing 41.05% of GDP. The secondary sector plays a more prominent role, contributing 243.235 billion yuan, or 47.07%, while the primary sector produces 61.418 billion yuan, or 11.89% of total output. Compared to the districts, counties rely more heavily on manufacturing, construction, and agriculture.

In subsection 3.2 of the manuscript, we provide a detailed discussion of educational disparities between districts and counties, focusing on both resource inputs and high-education admission rates.

[1] Source: https://www.fuzhou.gov.cn/book/fztjnj/2024tjnj/indexch.htm

[2] Due to rounding, the sum of sectoral outputs may not exactly equal total GDP. Minor discrepancies reflect rounding errors rather than substantive measurement differences.

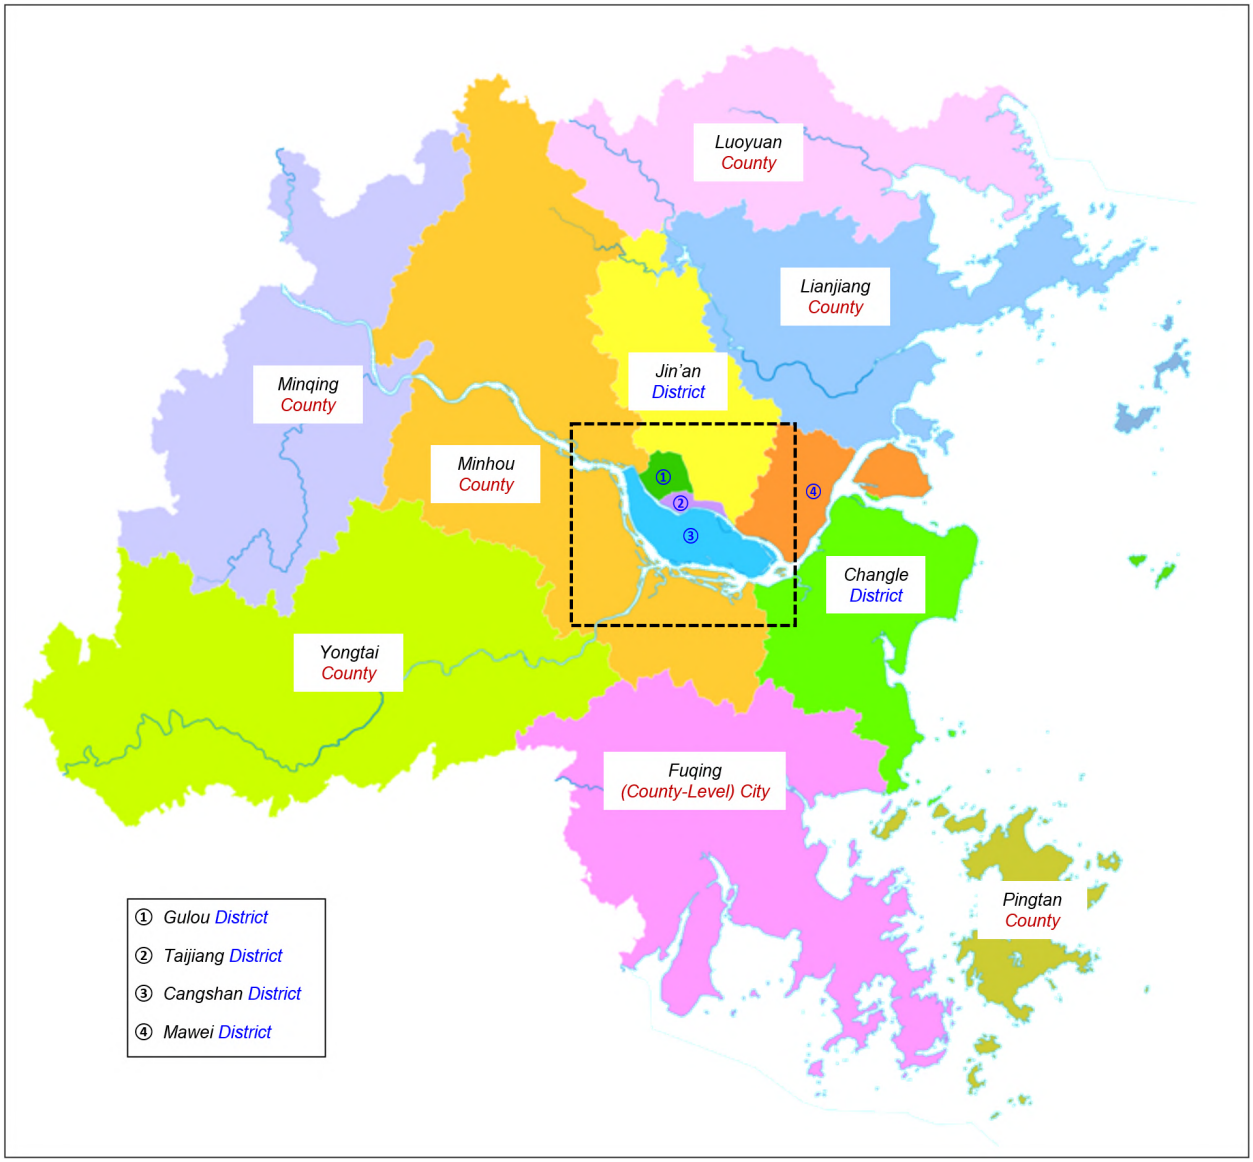
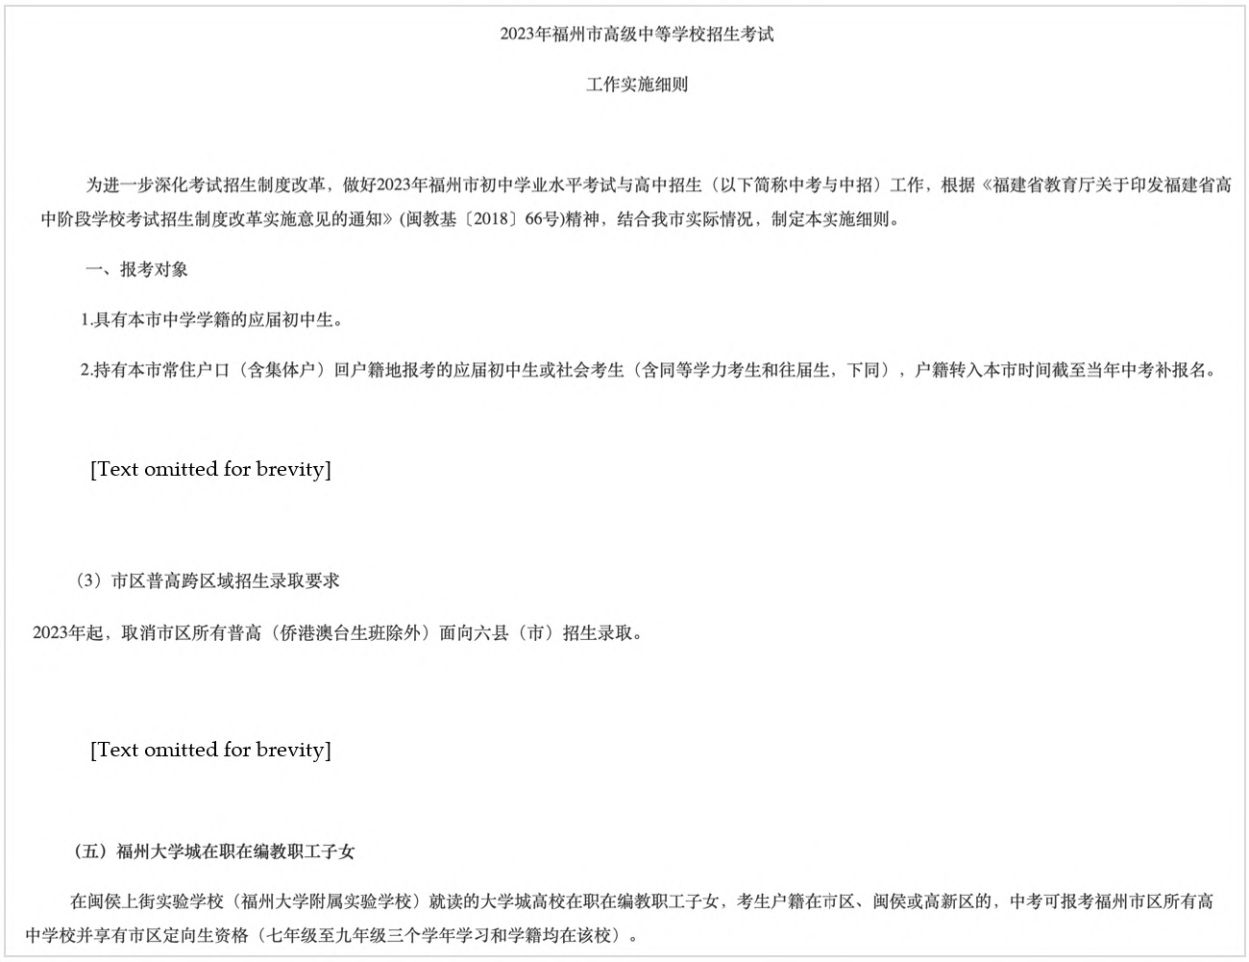


Figure A1: Administrative Divisions of Fuzhou City

Notes: The figure illustrates the administrative divisions of the core area of Fuzhou City. The black dotted box highlights the focal area of the empirical analysis. Islands located far from the core area are excluded from the map.

B Details of the Local-Enrollment Policy



2023年福州市高级中等学校招生考试

工作实施细则

为进一步深化考试招生制度改革，做好2023年福州市初中学业水平考试与高中招生（以下简称中考与中招）工作，根据《福建省教育厅关于印发福建省高中阶段学校考试招生制度改革实施意见的通知》(闽教基〔2018〕66号)精神，结合我市实际情况，制定本实施细则。

一、报考对象

1.具有本市中学学籍的应届初中生。

2.持有本市常住户口（含集体户）回户籍地报考的应届初中生或社会考生（含同等学力考生和往届生，下同），户籍转入本市时间截至当年中考补报名。

[Text omitted for brevity]

（3）市区普高跨区域招生录取要求

2023年起，取消市区所有普高（侨港澳台生班除外）面向六县（市）招生录取。

[Text omitted for brevity]

（五）福州大学城在职在编教职工子女

在闽侯上街实验学校（福州大学附属实验学校）就读的大学城高校在职在编教职工子女，考生户籍在市区、闽侯或高新区的，中考可报考福州市区所有高中学校并享有市区定向生资格（七年级至九年级三个学年学习和学籍均在该校）。

Figure B1: Official Notice on Fuzhou's Local-Enrollment Policy

Notes: The figure shows the official policy notice issued by the Fuzhou Municipal Education Bureau on April 7, 2023, which established the local-enrollment rule governing admission to senior high schools.

Implementation Guidelines for the 2023 Fuzhou Senior High School Entrance Examination

To further deepen the reform of the examination and enrollment system, and to ensure the proper organization of the 2023 Fuzhou Junior High School Academic Proficiency Examination and Senior High School Enrollment (hereinafter referred to as the "Entrance Examination and Admission"), these implementation guidelines are formulated in accordance with the *Notice of the Fujian Provincial Department of Education on the Implementation Opinions for the Reform of the Senior High School Examination and Enrollment System* (Document No. Minjiaoji [2018] 66), and in light of the actual conditions of Fuzhou.

Eligibility

1. Current junior high school graduates with registered student status in Fuzhou.
2. Current junior high school graduates or social candidates (including those with equivalent academic qualifications and previous graduates) holding permanent household registration (including collective household registration) in Fuzhou who return to their place of registration for the exam. The transfer of household registration into Fuzhou must be completed by the supplementary registration deadline for the entrance examination of the current year.

[Text omitted for brevity]

Admission Requirements for Enrollment in Urban General High Schools

Starting from 2023, all urban general senior high schools shall cease admitting students from the counties under Fuzhou's jurisdiction, except for classes designated for overseas Chinese, Hong Kong, Macao, and Taiwan students.

[Text omitted for brevity]

Children of In-service and On-staff Faculty Members in Fuzhou University Town

Children of in-service and on-staff faculty members from universities located in Fuzhou University Town who are studying at Shangjie Experimental School of Minhou County (Fuzhou University Affiliated Experimental School), and whose household registration is in the urban districts of Fuzhou, Minhou County, or the High-Tech Zone, may apply for admission to all senior high schools in Fuzhou's urban districts in the entrance examination. They are also eligible for the quota reserved for designated urban students, provided that they have studied and maintained student registration at the school for all three academic years from Grade 7 to Grade 9.

Figure B2: English Translation of the Local-Enrollment Policy

Notes: The figure shows the English translation of the local enrollment policy issued by the Fuzhou Municipal Education Bureau on April 7, 2023.

C Satellite Image



Figure C1: Satellite Image of Area in Research Design

Notes: The figure shows a satellite image of Fuzhou City obtained from Google Maps. The lines correspond to the lines in our research design.

D List of EdAI Applications

We aggregate the list of EdAI applications using the following procedures. We first searched for AI education applications on an app traffic tracking platform.[3] From the search results, we identified the top fifty applications as of April 2023. We then excluded applications that were not relevant to our context, including those focused on civil service exam preparation, minority language education, music education, dictionaries, and note-taking. This screening yielded a final list of twenty-three applications. We present an example screenshot from one EdAI application, Palfish Learning, in Figure D1.

Table D1: List of EdAI Applications

AI Foreign Teacher	Palfish Learning	Clever Elephant AI Education
Hi Echo	AI English Education Platform	Keli English
TalkAI	Dora AI Classroom	Hatch AI Speaking
New Concept English (AI Edition)	Speak AI	Speak Talk
Dolphin AI Learning	Xizhi AI Tutor	Exam Champion AI Classroom
Doubao Learning	AI Dictation	SpeakGuru
iFlytek AI Learning	Jinti AI	Zebra AI
TalkAI	Moss English	
Observations: 23		

[3]Link: https://app.diandian.com/search/



Figure D1: Screenshot of Palfish Learning

Notes: The figure shows the screenshot of Palfish Learning, translated and edited using Nano Banana.

E Resident Distribution

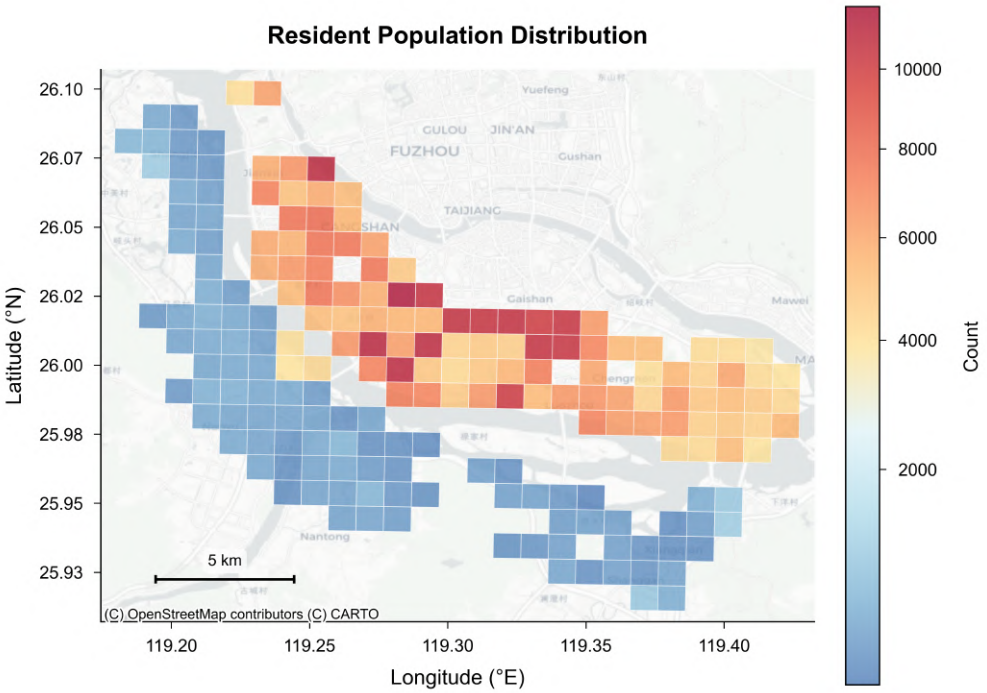


Figure E1: Resident Distribution

Notes: The figure shows the spatial distribution of resident density within the 5 km corridor surrounding the river that defines the policy boundary. This corridor contains the main residential zones along the river, while areas outside this range are sparsely populated. The analysis grid divides the region into 1 km × 1 km cells, which function as the unit of analysis. To reduce potential bias from university populations, all grid cells that contain any university are excluded.

F Robustness Check Based on Alternative Distance Bands

The main analysis restricts the sample to grid cells located within five kilometers on each side of the river-based policy boundary. Following prior research (Norris and Xiong, 2023), we test the robustness of our results by using alternative distance bands ranging from three to nine kilometers, widening the spatial window in two-kilometer increments. For each distance band, we re-estimate the event-study specification in Model 2. The resulting estimates are shown in Figure F1. Across all spatial windows, the pre-policy trends remain parallel, and the post-policy dynamics exhibit the same J-curve pattern documented in the main analysis. These results confirm that our main findings are not sensitive to the specific five-kilometer window used in the main analysis.

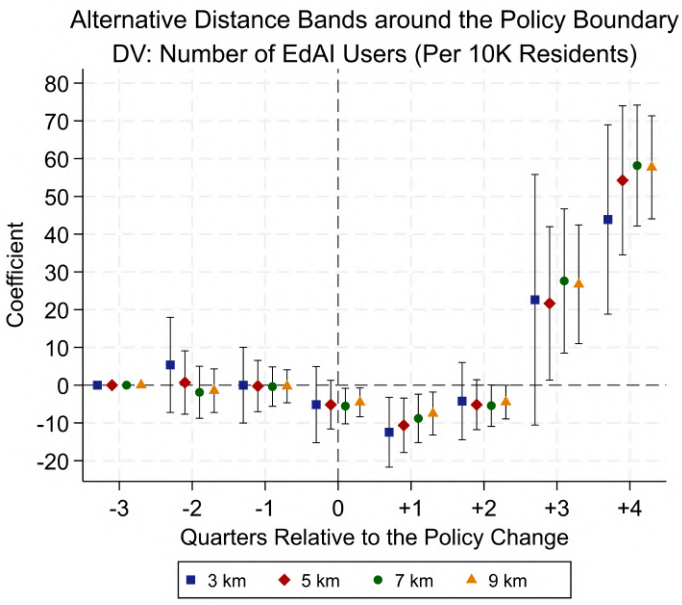


Figure F1: Robustness Check Based on Alternative Distance Bands

Notes: The figure presents the event-study SDID coefficients estimated from Model 2 with 95% confidence intervals. The vertical dashed line indicates the policy quarter.

G Robustness Check Based on Synthetic DID Estimates

As an additional robustness test, we implement a synthetic difference-in-differences estimator (Arkhangelsky et al., 2021; Berman and Israeli, 2022; Farronato et al., 2024; Guyt et al., 2025; Pachali and Datta, 2025). This method constructs a data-driven weighted combination of control units that more closely replicates the pre-policy trajectory of the treated units than a simple average of control cells. Figure G1 reports the resulting event-study SDID coefficients based on Model 2. The synthetic DID estimates exhibit parallel pre-policy trends, confirming that the method achieves a close match between treated cells and their synthetic counterparts. After the policy is implemented, the estimates reveal the same qualitative pattern observed in the baseline analysis: an initial decline in EdAI usage followed by a sharp and persistent increase beginning in the third post-policy quarter. The magnitude of the long-term increase is comparable to that found in the main SDID model, reinforcing the evidence of a J-curve pattern.

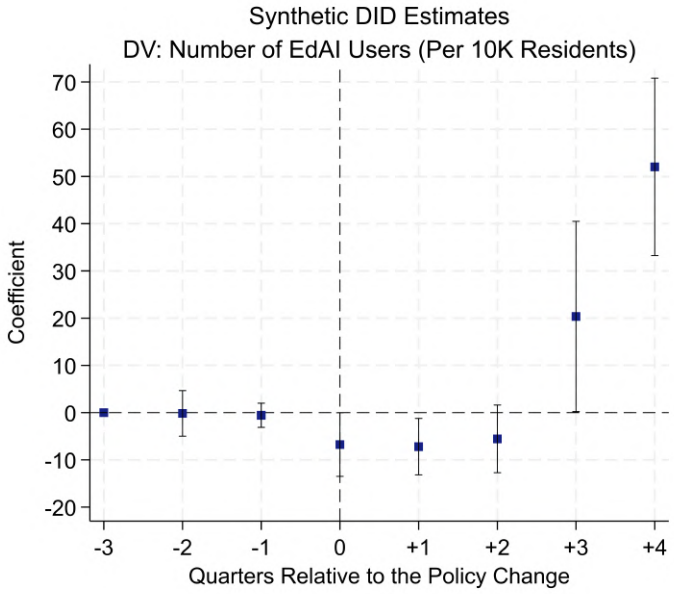


Figure G1: Robustness Check Based on Synthetic DID Estimates

Notes: The figure presents the event-study SDID coefficients estimated from Model 2 with 95% confidence intervals, when using the synthetic DID approach. The vertical dashed line indicates the policy quarter.

H Robustness Check Based on Spatial-Correlation Standard Errors

A potential concern in our setting is that grid cells located near one another may experience correlated shocks, which can generate spatial dependence in the error structure. Standard error clustered at the grid-cell level accounts for serial correlation over time but does not capture correlation across neighboring spatial units. To assess the robustness of our inference to this possibility, we re-estimate Model 2 using standard errors that explicitly allow for spatial correlation across cells (Butts, 2023; Colella et al., 2019). We apply the *acreg* command in Stata, which adjusts the variance-covariance matrix based on the geographical proximity of observations. Figure H1 displays the event-study estimates using these spatial-correlation standard errors. The resulting coefficient pattern closely matches the main findings. These results indicate that our conclusions are not driven by unmodeled spatial dependence and remain robust when allowing for spatially correlated errors.

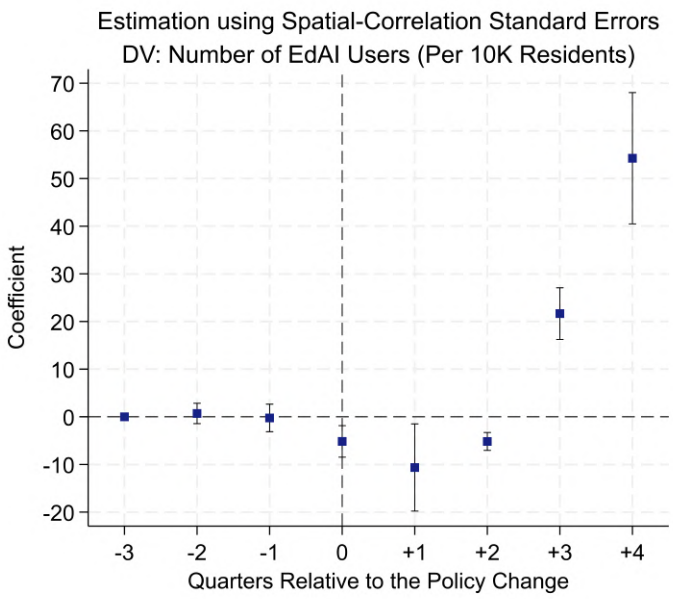


Figure H1: Robustness Check Based on Spatial-Correlation Standard Errors

Notes: The figure presents the event-study SDID coefficients estimated from Model 2 with 95% confidence intervals, when using standard errors that allow for spatial correlation across cells. The vertical dashed line indicates the policy quarter.

I Alternative Explanation: Residents Re-sorting across Riverbanks

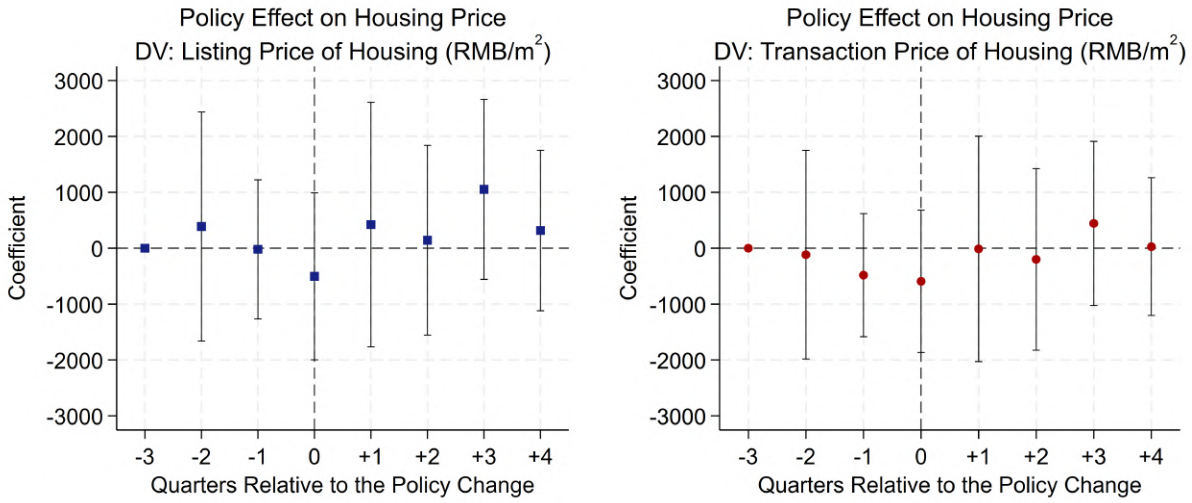


Figure I1: Event-Study SDID Estimates for Housing Prices

Notes: The figure presents the event-study SDID coefficients estimated from Model 2 with 95% confidence intervals. The vertical dashed line indicates the policy quarter. The dependent variable is quarterly averages of listing and transaction prices in each cell.

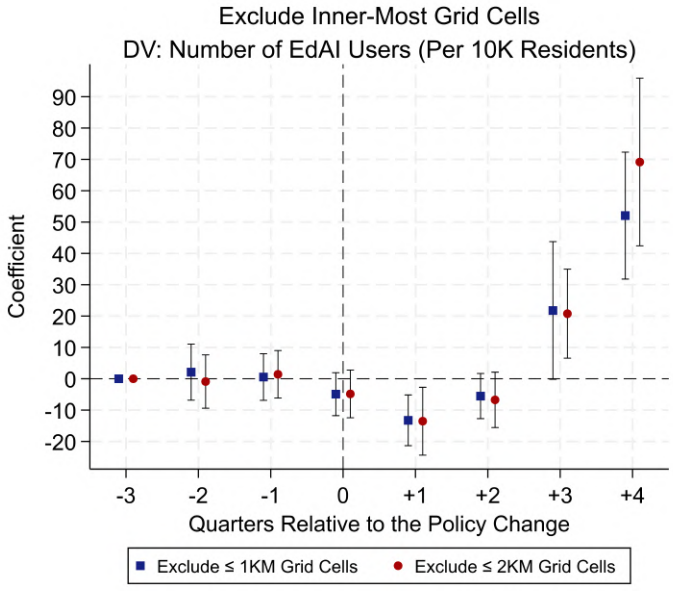


Figure I2: Donut Estimates by Excluding Inner-Most Grid Cells

Notes: The figure presents the event-study SDID coefficients estimated from Model 2 with 95% confidence intervals. The vertical dashed line indicates the policy quarter.

J SES Proxied by Industrial Output Per Capita

We further operationalize SES using the per capita industrial output of enterprises above a designated size (hereafter referred to as industrial output). Whereas housing prices primarily reflect accumulated household wealth, industrial output serves as a proxy for the broader regional economic environment (Glaeser et al., 1992; Glaeser, 2013; Moretti, 2012). Regions with higher industrial output tend to possess stronger local economic foundations that support public goods, including educational infrastructure, and provide more stable employment opportunities. These structural advantages enhance users' perceived control when facing external shocks, particularly shocks related to urban-rural disparities in educational opportunities.

We collect data on town-level per capita industrial output of enterprises above a designated size from the 2023 Fuzhou Statistical Yearbook. Each grid cell is assigned to its corresponding town, and the treated cells are divided into high- and low-SES subgroups based on the median value of industrial output per capita. We then re-estimate Model 2 separately for these two subgroups while pairing each with the full set of control cells.

The results, shown in Figure J1, are noisier than those from the housing price analysis, as reflected in wider confidence intervals. A plausible explanation is that town-level industrial output, as an aggregate measure, is less able to capture fine-grained variation in local socioeconomic status than cell-level housing prices. Nevertheless, even with the greater noise introduced by town-level aggregation, the J-curve is still more evident in low-SES towns, while the effects in high-SES towns remain consistently indistinguishable from zero. These findings support the interpretation that the local enrollment policy generates a behavioral response to urban-rural educational inequality through the mechanism of perceived control.

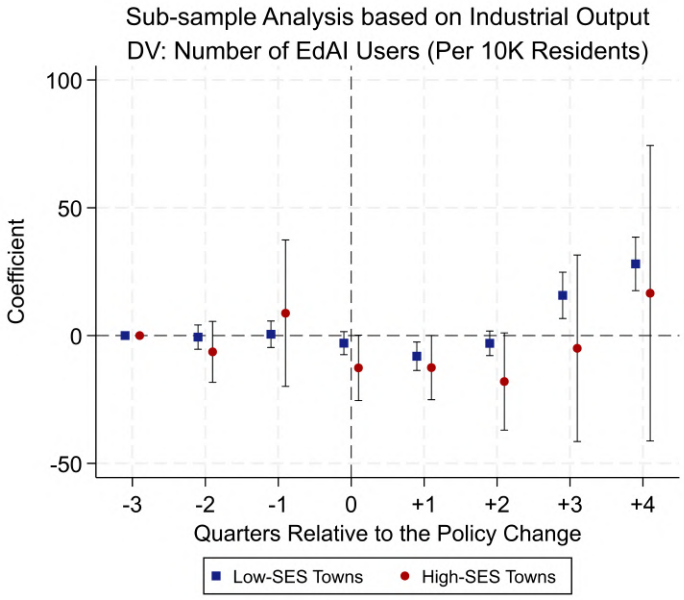


Figure J1: Sub-sample Analysis by SES Proxied by Industrial Output Per Capita

Notes: The figure presents the event-study SDID coefficients estimated from Model 2 with 95% confidence intervals. The vertical dashed line indicates the policy quarter.