ABSTRACT

Title of dissertation: ESSAYS ON ASPECTS OF EDUCATION AND ANTI-CORRUPTION POLICIES IN CHINA

Ming Fang, Doctor of Philosophy, 2023

Dissertation directed by: Professor Katharine G. Abraham
Professor Judith K. Hellerstein
University of Maryland Department of Economics

This dissertation comprises three chapters that examine the outcomes of government policies pertaining to education and anti-corruption measures in contemporary China. The study investigates the impacts of these policies on diverse domains such as the marriage market, innovation activities in higher education, and citizens' political attitude towards the government.

Chapter 1: The Effects of China's College Expansion on the Marriage Market (with Sai Luo)

Education policies can have crucial effects on the marriage market. In this chapter, we study the impacts of China's college expansion on the marriage market, with a special focus on its effects on the marriage outcomes of college-educated women and men. The empirical analysis is undergirded by a model featuring educational investment, marriage matching, and reductions in search frictions associated with the expansion. We estimate the effects of the expansion on marriage outcomes by exploiting geographic and birth-cohort variation in exposure to the expansion. Our analysis shows that, consistently with the predictions of the model, the expansion increased the marriage probability of college graduates. The expansion also increased the probability of college-college matches relative to the counterfactual of random matching and reduced the marriage age gap. Our findings highlight the important role of higher education institutions in shaping the marriage

market.

Chapter 2: (Mis)use of Power in the Ivory Tower: Evidence from Deans in Chinese Universities (With Yuyu Chen and Xuan Wang)

In a hierarchical academic system, power can distort the allocation of research resources and output ownership. We study the role of power in intellectual property acquisition. Using biographical information of deans in elite universities in China, we find that the deanship increases their patent applications by 14%. Further analysis suggests that the deanship effect is driven by misuse of power rather than ability or research resources. We provide causal evidence by showing that an anti-corruption campaign, which increases the cost of misusing power, contains the deanship effect. Finally, we find that misusing power distorts resource allocation.

Chapter 3: Anti-Corruption and Political Trust: Evidence from China (with Weizheng Lai)

How can anti-corruption campaigns influence political trust in government? We investigate this question through the lens of China's recent anti-corruption campaign, launched in 2013, which has unprecedentedly disclosed many corruption investigations to the public. By analyzing a large individual panel dataset, we show that on average, the campaign has significantly reduced political trust, particularly among groups less informed about corruption before the campaign. We document strong heterogeneity in trust changes, possibly driven by a pro- and anti-government cleavage, as captured by previous unpleasant experiences with the government, pro-government indoctrination, and Confucian norms. Our results fit in a model where polarization is rationalized by differences in priors about the government. We also rule out several alternative explanations for our findings.

ESSAYS ON ASPECTS OF EDUCATION AND ANTI-CORRUPTION POLICIES IN CHINA

by

Ming Fang

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2023

Advisory Committee:
Professor Katharine G. Abraham, Co-chair
Professor Judith K. Hellerstein, Co-chair
Professor Ethan Kaplan
Professor Jing Cai
Professor Susan Parker

Dedication

To my girlfriend, Huiyu, and my family.

Acknowledgments

It is impossible to thank everyone who helped me in my completion of this dissertation and my graduate research. I am deeply indebted to my advisors. Judy Hellerstein introduces labor economics to me in her fascinating class. Since then, she has provided invaluable guidance to my research. Katharine Abraham has also provided invaluable guidance and support for my research work. I cannot be more grateful to Judy and Katharine for all their support, encouragement, and the amount of time that they were willing to discuss with me about both the big pictures and details of my research. I am also grateful to Ethan Kaplan, whose feedback and discussions greatly improved all three chapters. I would like to thank Jing Cai and Susan Parker for serving on my dissertation committee.

I would like to thank other faculty members at the University of Maryland for their insightful comments and suggestions during discussions and seminars. Jessica Goldberg, Melissa Kearney, Guido Kuersteiner, Margaret Pearson, Nolan Pope, Sergio Urzua, and Chenyu Yang have greatly enhanced the quality of my research through their valuable contributions.

My peers at the University of Maryland have been a source of support and mentorship. Sai Luo, my co-author on the first chapter, has been a big brother and a role model throughout my Ph.D. studies. Hao Bo, my first-year peer advisor, provided me with helpful suggestions. Weizheng Lai's valuable comments improved the first two chapters of my dissertation, and he is also the coauthor on the third chapter. Runtian Tang, Zhenxun Liu, Eugene Oue, Mrin Chatterjee, Anu Sivaram, Nathalie Gonzalez, and John Soriano have been my friends and provided support during various stages. My research process was made much easier thanks to the invaluable assistance and support provided by Vickie Fletcher, Terry Davis, Amanda Statland, Mark Wilkerson, Heather Nalley, Brian McLoughlin, and other members of the Department of Economics staff.

I grateful acknowledge the financial support from the Chiang Ching-kuo Foundation for International Scholarly Exchange.

Lastly, I would like to thank my girlfriend, Huiyu, who has been my source of strength and encouragement throughout all the challenges I faced during my academic journey.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

Economic modernization is a complex process that requires a range of reforms and policies aimed at improving various aspects of institutions and society. Evaluating the impacts of these policies, particularly unintended consequences, is essential for a comprehensive understanding of the modernization process. In this dissertation, I study the effects of education policies and anti-corruption policies in China during the twenty-first century. Both types of policies are critical for developing countries. Education policies affect the accumulation and distribution of human capital, while anti-corruption policies affect the efficiency and accountability of the public sector.[1] This dissertation investigates various aspects of these policies to gain a better understanding of economic and societal modernization. These issues are particularly important for developing countries undergoing similar policy changes. Additionally, the economic implications derived from these chapters may have broader implications for other societies.

Education policies are unequivocally important in shaping economic and lifetime outcomes of individuals. While the labor market and health benefits of education have been extensively studied, education policies may also have an impact on marriage outcomes. Education not only affects attractive traits in the marriage market directly, but may also, at least in certain circumstances, education institutions may themselves work as places for marriage matching.[2] The first chapter studies the equilibrium effects of China's college expansion policy on the marriage market. China's

[1] It's worth noting that in countries where the government directly administers the higher education system, anti-corruption policies can have significant effects on universities, as discussed in Chapter 3.

[2] For example, Kirkebøen et al. (2021) document that in Norway colleges affect whom one marries directly because of attending the same institution.

college expansion, which started in 1999, greatly increased the availability of higher education to young women and young men. Although the policy was not intended to affect marriage outcomes, it may have had unintended consequences by altering the distribution of different types in the matching market and reducing search frictions in the marriage matching process.

To conceptualize these channels through which the college expansion may have affected the marriage market, we first develop a theoretical model with education investment and marriage matching. A key feature of the model is that colleges can act as a local marriage market, thereby reducing search frictions in the search and matching process. Motivated by the theoretical model, we analyze Chinese census data using a difference-in-differences approach to estimate the causal effects of the college expansion. Our empirical analysis leverages both geographical and cross-cohort variation in the extent of the expansion, with the former reflecting the institutional and historical characteristics of China's higher education system.

We find that the college expansion increased the marriage probabilities of college women and college men. Besides the extensive margin effects, we find that the expansion also led to changes in the patterns of marriage matching. Specifically, we document a reduction in the marriage age gap and an increase in the level of assortative mating by education, which reflects the underlying tendency of college graduates to marry each other. Our numerical analysis based on our theoretical marriage model suggests that this positive effect was driven by the fact that colleges can reduce search frictions in the marriage market, thus facilitating easier matches for both sexes. These results have important implications for education policy and its impact on marriage outcomes. First, our findings indicate that the college expansion policy, originally intended to increase access to higher education, also had the unanticipated benefit of increasing the marriage probability of college graduates in the marriage market equilibrium, even for women who have traditionally been at a disadvantage due to social norms.[3] Second, our finding of increased assortative mating raises

[3] The "left-over women" phenomenon in China refers to the fact that highly educated women are less likely to get married by a certain age (usually 27 years old).

concerns about potential effects on increasing income inequality resulting from an education policy designed to increase social welfare.

In addition to promoting human capital, higher education institutions also create new knowledge that has important positive externalities. However, this function can be hindered by institutional structures, especially in highly bureaucratic environments. In such environments, certain positions may wield excessive power and potentially distort the allocation of innovation output and incentives. In the second chapter, we investigate this issue in the context of Chinese universities, where deans of schools hold strong administrative power. Our event-study analysis reveals that deans obtained more patents after taking office. While alternative explanations cannot be ruled out, our extensive heterogeneity analysis provides suggestive evidence that this phenomenon is most likely due to the misuse of deans' power. This misuse could manifest in rent-seeking or exchange of favors given the power under the deans' control. For instance, we find that deanship had larger effects on patent applications for individuals with previous paper retraction experience, which could signal a lack of academic ethics.

In the second part of this chapter, we provide additional evidence on the potential misuse of power in universities and explore possible government policies that can address this issue. Specifically, we investigate the impact of China's anti-corruption policy on the deanship effects in patenting. By exploiting the recent anti-corruption campaign as a natural experiment, we find that the policy has reduced the deanship effects on patenting, further supporting the argument that deans' increased patent production is a result of misusing their administrative power rather than their own productivity. Our findings highlight an unintended positive consequence of the anti-corruption policy, which can potentially alleviate the distortion effects of power in the hierarchical structure of higher education institutions.

Beyond its impact on higher education institutions, the anti-corruption campaign in China may have had broader implications for citizens' perceptions of their government. Changes in political attitudes and trust could potentially affect the effectiveness of government policies in general. In the third chapter of this dissertation, we examine the effects of China's anti-corruption campaign on citizens' political trust in local government. We adopt a difference-in-differences framework and exploit regional and temporal variation in the intensity of the campaign. An important feature of our approach is the use of longitudinal data on beliefs, which allows us to study changes in reported attitudes within the same individual. By doing so, we can improve the reliability of self-reported attitude measures relative to previous studies, such as Wang and Dickson (2022).

We find that, on average, the anti-corruption campaign led to a reduction in political trust in local governments. However, our heterogeneity analysis reveals several interesting patterns. Specifically, individuals who had negative prior experiences with the government were more likely to lose trust in response to the campaign. Additionally, we found that higher education played a significant role in shaping the effects of the campaign on political attitudes. The campaign increased political trust among individuals who had received higher education. This effect was further moderated by Confucian norms, such that individuals in regions with strong Confucian traditions were more likely to respond positively to the anti-corruption campaign. Our findings suggest that Confucian norms may facilitate pro-government indoctrination via the education system, which may help explain the heterogeneous effects of the anti-corruption policy on political trust.

This dissertation examines the impact of government policies on various aspects, spanning education and anti-corruption. The first chapter uses both theoretical and empirical approaches to explore the equilibrium effects of a policy that expands access to higher education on the marriage market. The second chapter investigates the potential distortionary effects of administrative power on innovation activities in Chinese universities, and examines the impact of the anti-corruption campaign on the use of such power. The last chapter extends the analysis of the anti-corruption campaign to explore its effects on citizens' trust in local government, uncovering intriguing heterogeneous effects by education and prior experience with the government. Overall, the findings in this dissertation enhance our understanding of the unintended effects of government policies on critical outcomes in various domains.

Chapter 2: The Effects of China's College Expansion on the Marriage Market

2.1 Introduction

This chapter studies the impact on the marriage market of a radical reform of college education in China. In 1999, the Ministry of Education (MOE) sharply expanded access to higher education. Before the expansion, the rate of college attainment was minimal. The expansion led to a large and continuous increase in college attainment for young men and women. New enrollment in colleges increased by more than 400% in the first eight years after the onset of the expansion (Figure 2.1).[1] In addition to impacting the labor market and firms as shown in previous studies (Che and Zhang, 2018; Feng et al., 2018; Li et al., 2017), the reform could also have affected the structure of the marriage market in crucial ways. For example, massively expanding access to higher education altered the distribution of college- and non-college-educated women and men and may have reduced search frictions in the college marriage market.

Education not only provides a path to labor-market success but affects marriage outcomes. Returns to education in the marriage market are important influences on educational choices and lifetime welfare (Chiappori et al., 2009, 2018; Lafortune, 2013). Education policies, therefore, may substantially impact the marriage market. Given the nature of the marriage market as a two-sided matching market, such impacts hinge on not only individuals' own education but others' education. When evaluating the impacts of education policies on the marriage market, the general

[1] The dropout rate is extremely low at Chinese universities (Appendix 2.8.2). Therefore, the vast majority of the newly enrolled students became college graduates.

equilibrium effects must be taken into account. This is particularly important in a society whose overall education level has significantly increased, as in many countries in the past few decades (Schofer and Meyer, 2005; World Bank, 2017).

However, it is often difficult to evaluate the causal effects, either from a partial or a general equilibrium perspective, of education on marriage market outcomes because educational investment responds to returns to education in the marriage market. Previous studies have pointed out that marital returns are an important determinant of educational choices (Chiappori et al., 2009; Ge, 2011; Lafortune, 2013; Bruze, 2015; Attanasio and Kaufmann, 2017; Zhang, 2021). China's college expansion provides us with a great opportunity for studying the equilibrium impacts of education policies on the marriage market. First, by exploiting the exogenous timing and the geographic variation of the expansion, we can estimate the causal effects of expanding access to college education on the marriage market. Second, the unprecedented magnitude of the expansion enables us to more easily test its equilibrium effects on the marriage market.

To understand the potential effects of the college expansion on the marriage market and to discipline our empirical analysis, we first build a model of educational investment and marriage matching. In our two-period model, young women and men make choices regarding college education in the first period, and in the second period they match in the marriage market based on educational-attainment type and idiosyncratic preferences (Choo and Siow, 2006). Anticipation of market prospects in the second period affects educational choices in the first period (Chiappori et al., 2009, 2017).

Figure 2.1: National Trend of New Enrollment

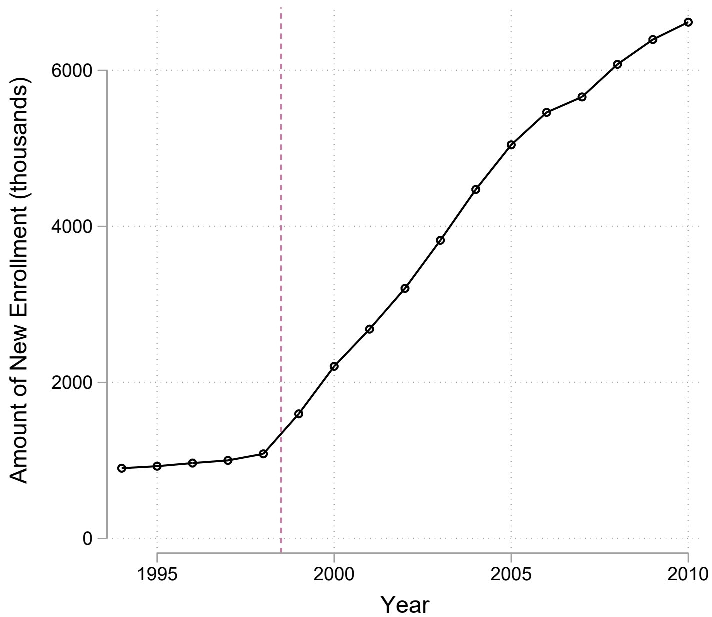

Notes: This figure reports national new enrollment in colleges for each year based on data from the National Bureau of Statistics of China. The college expansion started in 1999. Both four-year universities and three-year colleges are included.

Our model is further enriched by incorporating the role of colleges in reducing search frictions. Not only might having a college education be an attractive trait in the marriage market, but colleges themselves may serve as local marriage markets. A particular educational institution could directly affect who marries whom by providing a space for social interactions. This channel may operate via various social networks formed during college. Young women and men may meet in class, in student organizations, or via shared friends they met in college. These occasions for social interactions could significantly boost students' chances of getting to know each other and forming romantic relationships.[2] We refer to these potential opportunities as the "local college marriage

[2]For example, in Chinese universities, students are typically assigned to a series of classes based on their majors, and majors are determined upon admission to most universities. This arrangement creates chances for the formation of classmate relationships lasting years. Various student associations, meanwhile, are usually formed voluntarily based on common hobbies or aspirations, and they might also provide romantic opportunities based on sorting into the same one.

market" (LCMM).[3]

We model search frictions in the marriage market in a parsimonious way: everyone who enters the market has to pay a fixed cost, which represents search costs. The only exception is that some college students can randomly meet and potentially match. They can thus avoid the search costs yet still get married. The college expansion creates a thicker local marriage market. Therefore, it becomes easier for college students to form marriages without paying the search cost. This reduces search frictions for college graduates and raises their marriage probability. Based on certain assumptions about how the local college marriage market changed in response to the college expansion, we calibrate the model using microdata from before the college expansion. We then simulate the responses of the marriage market to the college expansion using the model. The simulation predicts that the expansion raised the probability of marriage for college women and men.

The college expansion impacted marital outcomes through two channels, according to the model. The first is an adjustment in matching outcomes that was induced by different marginal distributions of education types in the marriage market as also implied in classical matching models. This channel means the expansion potentially reduced college graduates' marriage probability. The second channel is a reduction in average search frictions in the local college marriage market, which is the key innovation of our enriched model. The second channel dominates the first in our model, leading to an overall positive effect of the expansion on college graduates' marriage prospects.

Motivated by the theoretical model, we estimate the causal effects of the college expansion on marriage-market outcomes using a difference-in-differences (DID) design. Exploiting the institutional fact that regions with more pre-existing college resources experienced a larger expansion, we use, as a proxy for college-expansion intensity, a measure that is proportional to historical college enrollment per capita. We compare birth cohorts affected and unaffected by the expansion

[3]Recent literature has documented the important role of colleges per se as local marriage markets (Kirkebøen et al., 2021). Indeed, we provide supportive evidence in our context using household survey data (Section 2.2.2).

in places with various levels of this proxy. Under the assumption that in the absence of the expansion, the dynamics of marriage outcomes in locations with different values of the expansion proxy followed parallel trends, our empirical strategy delivers the causal impacts of the expansion on marriage outcomes. We focus on the marriage outcomes of college women and men, but we also look at those of noncollege women and men, as the model also predicts changes in marriage outcomes for these groups because of spillover effects in equilibrium.

The DID estimates show that the college expansion led to a modest increase in marriage probability for both college women and men. When looking at cohorts that went to college five to eight years after the onset of the expansion, a one standard-deviation increase in our expansion-intensity proxy leads to a 2.7 percentage-point increase in marriage probability for college men and a 1.7 percentage-point increase for college-educated women. The results survive a battery of robustness tests. For noncollege groups, we find a relatively small and positive effect of the expansion on the probability of marriage for noncollege men and no effects on noncollege women.

By altering the marriage-market structure, the college expansion may have changed marriage-matching patterns in addition to affecting marriage probabilities. We first look at the effects of the expansion on assortative mating by education level; that is, we examine whether college women and men are now more likely to marry each other. To tease out mechanical effects of the enlarged college population, we construct an index for assortative mating: the difference between the actual probability of college-college matches and the probability of college-college marriages in the hypothetical situation of random matching. We find that the expansion indeed increased the level of assortative mating. Second, we show that the expansion reduced the marriage age gap for college graduates. This finding is also potentially consistent with the story of decreased search frictions in the college marriage market.

This chapter contributes to several strands of the literature. First, we contribute to theoretical models of marriage by parsimoniously embedding search frictions in classical matching models. We start with the Choo and Siow (2006) framework, in which marriage matching is transformed into a discrete-choice problem. Chiappori (2017) and Chiappori et al. (2018) add educational choices to this framework. We build into the model a college-specific matching technology. In our model, some people can meet their spouses in college without incurring the search costs that exist for other types of marriage. This reduction in search frictions plays a vital role in explaining what happens in the marriage market following the college expansion. This approach may potentially be used to characterize the marriage market in similar contexts.

Second, our study contributes to a burgeoning literature about how education reforms impact marriage markets. Hener and Wilson (2018) find that a compulsory-schooling reform in the UK reduced the marriage age gap for affected women. André and Dupraz (2019) study a school-construction program in Cameroon and show that a higher level of education leads to a higher likelihood of a polygamous marriage for both men and women. Salcher (2020) finds that girls in Zimbabwe who benefited from an education reform married younger and better educated husbands. Closest to our work is Ge and Huang (2020), which uses China's college expansion as an instrument to estimate the effects of one's own education on marriage and fertility. These studies focus on the partial equilibrium effects of upgrading one's own education, thus neglecting the intrinsic general equilibrium feature of the marriage market as a two-sided matching market. Our study stresses the importance of general equilibrium effects. If such effects matter, then we should interpret the partial equilibrium estimates of the effects of education policies on marriage outcomes with caution. One paper that adopts a similar general equilibrium perspective is Zha (2022), which investigates the effects of the school-construction program in Indonesia on the marriage age gap. This chapter differs from Zha (2022) in two respects. First, we study college education instead of primary education. Second, we show the role of college education in reducing search frictions in the marriage market, which is a novel channel.

Third, our study sheds light on a recent strand of the literature that uncovers the direct role of higher education institutions as marriage markets. Kirkebøen et al. (2021) find that attending a certain college in Norway raises the probability of marrying someone from the same institution. Artmann et al. (2021) document strong assortative mating by field of study in college in the Netherlands and provide causal evidence on the effects of access to specific marriage markets. Using Swedish data, Nybom et al. (2017) show that universities contribute to couples' skill sorting and their children's skill inequality. By exploiting the natural experiment of China's college expansion, we quantitatively show the significance of local college marriage markets when evaluating policies targeted at higher education. A novel finding in this chapter is the extensive-margin effect of the expansion on marriage probability. This chapter also confirms the importance of colleges as marriage markets by documenting the effects of the expansion on assortative mating by education and on the marriage age gap. Compared to previous studies in developed economies, our study also provides more relevant insights for less developed countries that are experiencing or will experience expansions in higher education.

Our findings also have implications for some critical issues in China. We find positive effects of the college expansion on the level of assortative mating, which suggests that the expansion could potentially increase inequality and intergenerational persistence of income and social status. Our results are also relevant to China's so-called "leftover women" phenomenon. "Leftover women" is a term used to describe educated women who marry later or are less likely to get married (Fincher, 2016; Magistad, 2013; To, 2015). The reason for the phenomenon is arguably that according to certain social norms, women should "marry up." We find positive effects of the expansion on the marriage probability of college women, suggesting that expanding access to higher education can remove some of the barriers for college women in the marriage market.

The remainder of the paper proceeds as follows. Section 2.2 develops the model and generates hypotheses. Section 2.3 describes the context and the data used for our empirical analysis. Section 2.4 introduces our empirical strategy. Section 2.5 reports our main results regarding the effects of the college expansion on the probability of marriage. Section 2.6 presents findings about the effects of the expansion on assortative-mating patterns. Section 2.7 concludes.

2.2 Model

In this section, we develop a marriage-matching model with educational investment and fixed search costs.

2.2.1 Overview of the Model

In the first of two periods, women and men draw their idiosyncratic costs of college education, after which they make the decision whether to attend college. The individual choices determine the distributions of different education types (college or noncollege). Following the education choices, individuals draw their idiosyncratic taste for a spouse, which depends on only the education type of their potential spouse and not the latter's identity. Women and men then match in the marriage market based on education type following a transferable-utility framework (Choo and Siow, 2006).[4]

The model is enriched by including the potential role of colleges in reducing search frictions in the marriage market. This feature of the model incorporates the idea that colleges serve as local marriage markets in which young women and men meet their potential future partners.[5] To enter the marriage market, everyone has to pay a fixed cost of searching for a partner. A subgroup of college women and men, nevertheless, randomly meet each other and form potential matches. If both sides of a potential match agree to marry, then they no longer need to search and therefore do not pay the search cost. We call the set of these meetings the LCMM.

[4] In our setting, transferable utility means that husbands and wives send or receive within-marriage utility transfers without loss.

[5] More generally, the social networks developed during college help one find a suitable potential match at one's own college, on a neighboring campus, or even after graduating from college while still in the network.

2.2.2 Model Details

The economy is populated by N_f women and N_m men. We use i as the index for individual woman i and j for man j. There are two periods: $t = 1$ for education choice, and $t = 2$ for marriage-market matching. In Figure 2.2, we graphically present the structure of the model. We discuss the model details in the rest of the subsection.

Figure 2.2: Model Structure

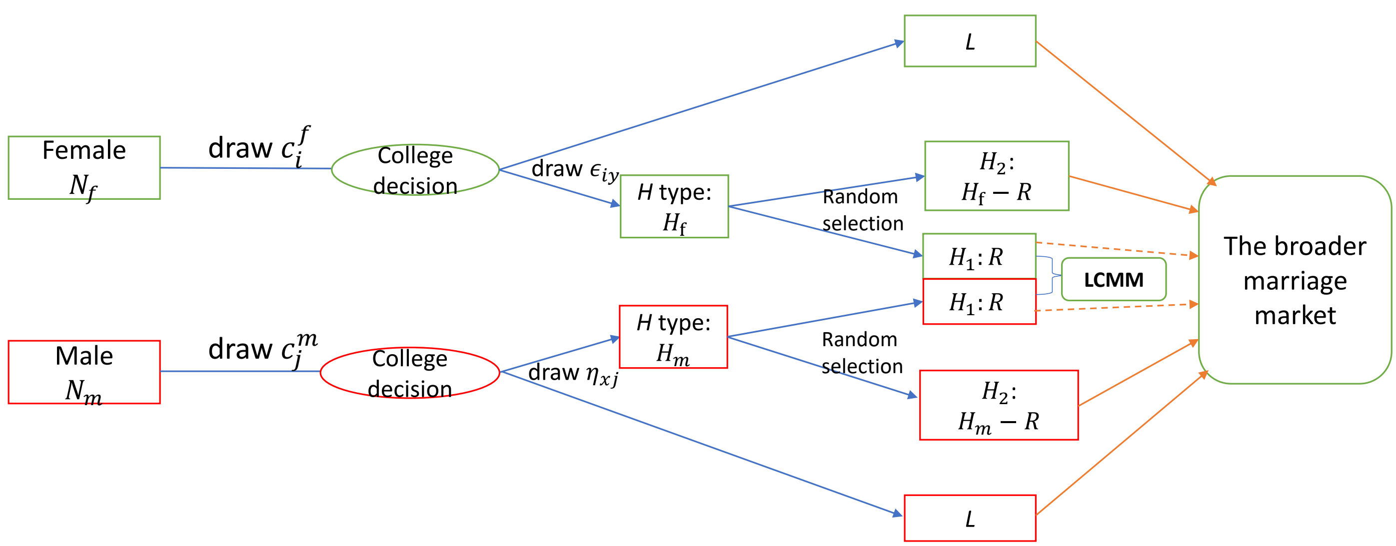

Notes: Individuals draw their costs of education before choosing whether to get a college education. Before entering the marriage market, they observe their idiosyncratic marital preferences. A certain fraction of college students randomly meet a potential partner in college (type H_1). If they and their potential partners agree to match (and marry) at this stage, they do not need to participate in the costly marriage-search process. If they choose not to match at this stage, they can choose either to stay single or enter the broader marriage market, in which they will have to bear the search cost.

Education Choice

There are two types of educational attainment: college (H) and noncollege (L). At the start of the first period, individuals first draw their idiosyncratic costs of college education and then decide whether to go to college. The cost of education of woman i (man j) is denoted as follows:

$$c_i^f = c_f + \theta_i^f$$
$$c_j^m = c_m + \theta_j^m$$

$c_f(c_m)$ is the average cost of education for the woman (man). θ_i^f (θ_j^m), the individual-specific-shock component of education cost, follows the distribution $G^f(\theta)$ ($G^m(\theta)$).

The group size of college women is denoted by H_f, and that of college men by H_m. We define the college ratios for women and men as $h_f = \frac{H_f}{N_f}$ and $h_m = \frac{H_m}{N_m}$, respectively.

The college expansion is modeled as a reduction in the average-cost parameters (c_f and c_m). We consider the situation in which c_f decreases more than c_m; that is, the education cost decreases faster for women than for men. This situation matches the empirical observation that the college ratio of women increased faster than that of men during the college expansion, as we discuss in more detail in Appendix 2.10.3.

Marriage Market in the Second Period

In the second period, conditional on all the individuals' educational choices, women and men meet and match in the marriage market. We adopt the framework developed by Choo and Siow (2006) to characterize the matching process. Individuals in the marriage market match based on their education type. For example, a college man could match with a college woman, match with a noncollege woman, or stay single. Before the matching process starts, each individual draws their idiosyncratic tastes for the education type of their potential partner. This individual taste shock depends only on the potential partner's education type, not on their identity. Moreover, we add to this framework the LCMM to characterize how colleges help reduce search frictions in the marriage market.

Evidence for the Local College Marriage Market

Our assumption about the existence of the LCMM is based on the observation that colleges play an important and direct role in marriage formation in China. We document evidence that supports this view using information from the China Family Panel Studies (CFPS). The CFPS is a nationally representative longitudinal survey. Launched in 2010, the survey covers extensive economic outcomes including family dynamics and relationships. Importantly, the respondents reported how they met with their spouses.

Table 2.1 reports the fraction of people who met their (first) spouse in school by education type. Comparing the top rows (for college graduates) to the bottom rows (for noncollege individuals), we see that a much higher fraction of college graduates met their spouses in school. On average, more than 20% of college graduates reported that they met their spouses in school. If we look at the statistics before and after the college expansion, 17% of pre-expansion college graduates and 26% of post-expansion college graduates met their spouses in school, and the difference is statistically significant. These statistics provide supportive evidence for the role of colleges as local marriage markets. They also suggest that the importance of LCMMs has increased since the college expansion.[6]

[6] One concern is that individuals in the post-expansion cohorts are younger and mechanically more likely to marry people they met in school. We show in Table 2.14 that the results are robust to confining the sample to those who married early. The second concern is that the before-versus-after comparison might be contaminated by other secular trends. In Appendix 2.9.1, we provide suggestive causal evidence on the effects of the college expansion by exploiting a DID strategy similar to our baseline econometric specification. Separately, a survey by Wang and Wang (2000) shows that 20.3% of college students in one city were in a romantic relationship in 2000. In a survey by Su et al. (2011), conducted 11 years after the previous survey, this ratio had increased to 46.9%. The numbers suggest that colleges may have played a bigger role in facilitating matching after the college expansion than before.

Table 2.1: Fraction of People Who Met Their Spouses in School

	Cohorts	Fraction	Observations
College	1975-80 (Pre-expansion)	0.17	313
	1981-88 (Post-expansion)	0.26	1,108
	Difference	0.09 ($p < 0.01$)	
Non-college	1975-80 (Pre-expansion)	0.03	3642
	1981-88 (Post-expansion)	0.05	5169
	Difference	0.02 ($p < 0.01$)	

Source: China Family Panel Studies 2010-2018. All results weighted using the CFPS survey weights.

The Local College Marriage Market

To model the way in which colleges reduce search frictions, we assume that college women and men can meet each other at college. Once they meet, they have the option to get married. We characterize this process using the following meeting function.

Definition 2.1. *The meeting function for the LCMM, $R(H_f, H_m)$, is the number of potential meetings between H_f college women and H_m college men.*

$$R = zH_f^a H_m^b$$

Among all the college women and men, R college men *randomly* meet with R college women. Channels for these potential matches, as discussed in the introduction, include various forms of social interactions that provide chances for students of different genders to meet each other. When they meet, they have the option to agree to get married. We assume that the probability of entering the LCMM is independent of student characteristics (including cost of education and marital preferences). Put another way, college students are randomly selected into the LCMM.[7] We denote the college students who are randomly selected into the LCMM (the R college women and men) as type H_1, and the rest of the college students as type H_2.

We assume that the matching function R follows a Cobb-Douglas form, in which z, a, b are constants.[8] We further assume that this matching function exhibits increasing returns to scale ($a + b > 1$). The latter assumption implies that college women and men are more likely to enter the LCMM after the college expansion. The intuition is that as more students enroll in college, the LCMM becomes a thicker market. This implies that with more people enrolled in college, it becomes more likely for a college student to randomly meet a potential match of the opposite gender via the social interactions in college and, more generally, via the social networks formed in college.

We embed search costs in the model in a parsimonious way by making the simplified assumption that individuals have to pay a fixed cost δ if they choose to enter the marriage market and form a match over staying single. The cost consists of various components, such as money and time spent searching. However, if college woman i and college man j meet in the LCMM and agree to match, they will get married without having to pay the search cost. If they decide not to form a match, then they may choose to either stay single or enter the broader marriage market (Figure 2.2), for which they will need pay the search cost.[9] We formally describe the payoffs and search cost in the marriage market in the following section.

[7] Many of the occasions for social interactions, such as class assignment, are random in Chinese universities.

[8] Similar forms of Cobb-Douglas matching functions are widely used in the literature to characterize the (potential) match between workers and vacancies.

[9] For example, if a college graduate draws a strong preference for a noncollege partner, then they will reject the potential college-college match even if they randomly enter the LCMM.

Payoffs in the Marriage Market

The payoff of marriage is determined by one's own type, the type of one's spouse, and one's unique taste for marriage. For woman i of type x who marries a man of type y, her individual marital payoff is as follows:

$$u_{ixy} = \alpha_{xy} + \tau_{xy} - \delta \mathbb{1}[y \neq 0] \max\{\mathbb{1}[x \neq H_1], \mathbb{1}[y \neq H_1]\} + \varepsilon_{iy}$$

The marital payoff to man j of type y is as follows:

$$v_{xyj} = \gamma_{xy} - \tau_{xy} - \delta \mathbb{1}[x \neq 0] \max\{\mathbb{1}[x \neq H_1], \mathbb{1}[y \neq H_1]\} + \eta_{yj}$$

Here, $x, y \in \{H_1, H_2, L, 0\}$. τ_{xy} is the utility transfer within the marriage, which represents how the joint marital surplus is allocated. The value of τ_{xy}, which could be positive or negative, is determined in equilibrium.[10]

Both sides of the match have to pay the fixed cost δ if they decide to enter the marriage market and search for a spouse rather than stay single. α_{xy} and γ_{xy} are the systematic marital payoffs. The only exception is when both sides of a match are in the LCMM *and* agree to marry each other. They can then avoid paying the cost. Their payoffs are thus determined as follows:

$$u_{iH_1H_1} = \alpha_{HH} + \tau_{H_1H_1} + \varepsilon_{iH}$$

$$v_{H_1H_1j} = \gamma_{HH} - \tau_{H_1H_1} + \eta_{Hj}$$

This is a special case of the payoff functions above. The H_1 and H_2 types generate the same

[10] We denote the case in which woman i stays single $y = 0$ and that in which the man j stays single $x = 0$. $\tau_{x0} = \tau_{0y} = 0$. In this transferable-utility framework, there is no loss of surplus in the transfer of τ_{xy}. The first subscript always refers to the wife's type, and the second always refers to the husband's type.

systematic marital returns except that there is no fixed search cost for the H_1H_1 match.[11] The H_1 type and the H_2 type are valued the same way by their potential spouses in terms of idiosyncratic marital preferences.[12]

2.2.3 Equilibrium

The Marriage Market

Given the preferences and the utility transfer (τ), each individual determines their preferred partner type. The aggregated demand and supply for a given type of match are equal to each other in equilibrium. Specifically, we assume that μ_f^{xy} women of type x choose type-y men and μ_m^{xy} men of type y choose type-x women. Then, in equilibrium,

$$\mu_f^{xy} = \mu_m^{xy} = \mu^{xy}.$$

The equilibrium conditions pin down the transfer τ_{xy}. The matching function μ_{xy} represents the number of type-xy matches in equilibrium. μ_{x0} is the number of type-x women that stay single, and μ_{0y} is the number of single type-y men.

The Educational Choice

Now we return to the educational investment choice. In the first period, after observing idiosyncratic costs of education, each individual decides whether to attend college, anticipating what will happen in the marriage market. Since they do not observe their idiosyncratic preferences for marriage types yet, they only take into account the expected payoff of a given education type (H or L).[13]

[11] That is, $\alpha_{xH_1} = \alpha_{xH_2} = \alpha_{xH}, \alpha_{H_1y} = \alpha_{H_2y} = \alpha_{Hy}, \gamma_{xH_1} = \gamma_{xH_2} = \gamma_{xH}$, and $\gamma_{H_1y} = \gamma_{H_2y} = \gamma_{Hy}$.

[12] That is, $\varepsilon_{iH_1} = \varepsilon_{iH_2} = \varepsilon_{iH}, \eta_{H_1j} = \eta_{H_2j} = \eta_{Hj}$.

[13] This also rules out the concern of multiple equilibria, which might occur if individuals observe their idiosyncratic marital preferences before the education decision.

For woman i who chooses education type H, the expected payoff is $U^H - c_i$. Her expected payoff of choosing education L is U^L. (U^H, U^L), which represent the expected marital payoffs to different education types (before paying the education costs) are determined in the following way. The expected payoff of choosing noncollege (L), U^L, is the maximal payoff determined by one's optimal choice of spouse:

$$U^L = E(u_{iLy} | y = arg \max_{y=0,H,L} u_{iLy})$$

If an individual chooses college education (H), the expected payoff follows the same structure but also depends on whether they randomly enter the LCMM:

$$U^H = \sum_{k=1,2} p^f(H_k) E(u_{iH_k y} | y = arg \max_{y=0,H_k,L} u_{iH_k y})$$

Here, $p^f(H_1) = \frac{R}{H_f}$ is the probability of being selected into the LCMM conditional on attending college.[14] $p^f(H_2)$ is the probability of not being selected into the LCMM conditional on attending college, and $p^f(H_2) = 1 - p^f(H_1)$.

Woman i therefore chooses college education (H) based on the following decision rule:

$$U^H - c_i - U^L \geq 0$$

[14] When making the education choice, individuals do not observe their marital preferences or whether they will enter the LCMM. Therefore, only the expected payoffs matter for their choices.

Man j faces a symmetric problem:

$$V^L = E(v_{xLj} | x = arg \max_{x=0,H,L} v_{xLj})$$
$$V^H = \sum_{k=1,2} p^m(H_k) E(v_{xH_kj} | x = arg \max_{x=0,H_k,L} v_{xH_kj})$$

Here, $p^m(H_1) = \frac{R}{H_m}$ is the probability of entering the LCMM and $p^m(H_2) = 1 - p^m(H_1)$. Man j follows a similar decision rule to woman i regarding college education:

$$V^H - c_j - V^L \geq 0$$

We describe the equilibrium of the model in detail in Appendix 2.10.1. The equilibrium is determined by individual educational choices and the marriage market equilibrium. Individual educational choices are determined by expected payoffs in the marriage market. Marriage-market returns are determined by the distribution of education types and by individual marriage choices. The equilibrium is characterized as a fixed point in which individual education choices are consistent with marital returns to education. The existence of the equilibrium is guaranteed by Brouwer's fixed-point theorem.

2.2.4 Comparative Statics

To illustrate the intuition about the impacts of the college expansion on the marriage market, we report marital outcomes as a function of the distributions of college women (h_f) and college men (h_m). For the full model, both the marital outcomes and education outcomes are functions of the exogenously shifted mean costs of education (c_f and c_m). In Appendix 2.10.4, we report the comparative statics of the full model, including how educational-attainment and marriage outcomes respond to the college expansion.

Calibrating Model Parameters

We start by calibrating the parameters in the model using pre-expansion marriage-market data. Specifically, the systematic-returns-to-marriage parameters are pinned down using marriage patterns on pre-expansion cohorts in the 2005 China mini-census data. We construct a data set for married couples aged between 27 and 40 years old in 2005, which provides us with a snapshot of the marriage market prior to the expansion. Crucial to our model are the parameters for payoffs of marriage. Following the Choo-Siow framework, we can show that for a given type of match xy that is *not* formed via the LCMM (that is, x or y is not type H_1), the joint surplus and the matching function is as follows:

$$\frac{(\mu^{xy})^2}{\mu^{x0}\mu^{0y}} = \exp(\alpha_{xy} + \gamma_{xy} - \alpha_{x0} - \gamma_{0y} - 2\delta) \tag{2.1}$$

The right-hand side is the exponential form of the joint systematic surplus of marriage over staying single. Intuitively, a higher surplus is associated with a larger measure of the corresponding type of marriage xy. For marriages formed via the LCMM ($x = H_1, y = H_1$), the search cost is avoided. The relationship is as follows:

$$\frac{(\mu^{H_1H_1})^2}{\mu^{H_10}\mu^{0H_1}} = \exp(\alpha_{HH} + \gamma_{HH} - \alpha_{H0} - \gamma_{0H}) \tag{2.2}$$

The LCMM, however, introduces additional complications in calibrating the parameter for the search costs (δ) because we do not distinguish between type H_1 and H_2 in the data. In order to calibrate δ, we use auxiliary information from the CFPS about the fraction of college-college marriages that are formed via meeting in school. Intuitively, conditional on the matching function R and the marginal distributions of education types, a higher fraction of college-college marriages that are formed by meeting in school should be associated with a higher fixed search cost. We draw from the CFPS a variable λ defined as the fraction of college-college marriages via the LCMM out of all college-college marriages:

$$\lambda \stackrel{\text{def}}{=} \frac{\mu^{H_1H_1}}{\mu^{HH}}$$

In addition, conditional on the LCMM meeting function R, we have the following definitions of variables:

$$J_f \stackrel{\text{def}}{=} \frac{R - \mu^{H_1H_1}}{H_f - R - \mu^{H_2H_2}}$$
$$J_m \stackrel{\text{def}}{=} \frac{R - \mu^{H_1H_1}}{H_m - R - \mu^{H_2H_2}}$$

It turns out that

$$\delta = \ln \frac{\lambda}{1-\lambda} - 0.5 \ln J_f J_m. \tag{2.3}$$

Equations 2.1, 2.2, and 2.3 enable us to identify the search-cost parameter together with payoffs to marriages conditional on the meeting function R. We describe the procedure in detail and prove Equation 2.3 in Appendix 2.10.2.

Meeting Function for Local College Marriage Market

The function $R = R(H_f, H_m)$ is important for both calibrating model parameters and simulating our model's comparative statics. Without direct information on the matching process, unfortunately, we cannot pin down its functional form. Based on the assumption of increasing returns to scale, we set the function as follows:

$$R = 0.5 H_f^{0.75} H_m^{0.75}$$

The choice of the meeting function is somewhat ad hoc: we choose a functional form that can generate the comparative statics that are largely consistent our main empirical findings. That is, under reasonable assumptions about the meeting function, the simulation based on the theoretical model can explain our empirical results.

Social Norms against Marrying Down among Women

An important feature of China's marriage market is the social norm of aversion to seeing women marry lower-status men ("marrying down"). It is therefore much less likely for highly educated women to marry less educated men than the reverse (Figure 2.14). This norm is formalized in our model by the marriage-payoff parameters: the marital surplus for college women who marry down is less than that for noncollege women who marry up (the inequality below). This is consistent with the marital-payoff parameters estimated using observed data.

$$\alpha_{HL} + \gamma_{HL} - \alpha_{H0} - \gamma_{0L} < \alpha_{LH} + \gamma_{LH} - \alpha_{L0} - \gamma_{0H}$$

2.2.5 Simulating the Probability of Marriage

Based on the estimated and calibrated parameters, we simulate marriage-market responses to the college expansion. The college ratio of women is set to be initially lower than that of men based on the data moments for pre-expansion marriages (0.08 and 0.09) used to calibrate the baseline model's parameters. We allow the college ratio of women to increase faster than that of men. Specifically, we set the ending values of the college ratios for women and men as 0.35 and 0.33, respectively. These values reflect the distribution of education types for post-expansion cohorts (1985–88) in high-expansion regions. High-expansion regions are defined based on the value (above the median) of our empirical proxy for the magnitude of the college expansion, which we define in Section 2.3.2. These data moments also reflect an important feature of the expansion: the college ratio of women increased faster than and overtook that of men (further supporting evidence is discussed in Appendix 2.10.3).

The data moments we use for the simulation mainly reflect the temporal variation in the national distribution of education types.[15] For our empirical analysis, however, a before-versus-after comparison does not serve as a reliable identification strategy because it is very likely confounded by other secular trends. Therefore, later in Section 2.4, we resort to a DID design to empirically estimate the effects of the expansion on the marriage market.

Figure 2.3 reports the simulated marriage probabilities of college graduates as a function of the college ratios of women and men based on our model. Overall, the simulated results predict an increase in the marriage probability of both college men and women, except for college women early on in the expansion. Our model can help us further disentangle different mechanisms.

Figure 2.3: College Expansion and Marriage Rates: Simulated Results

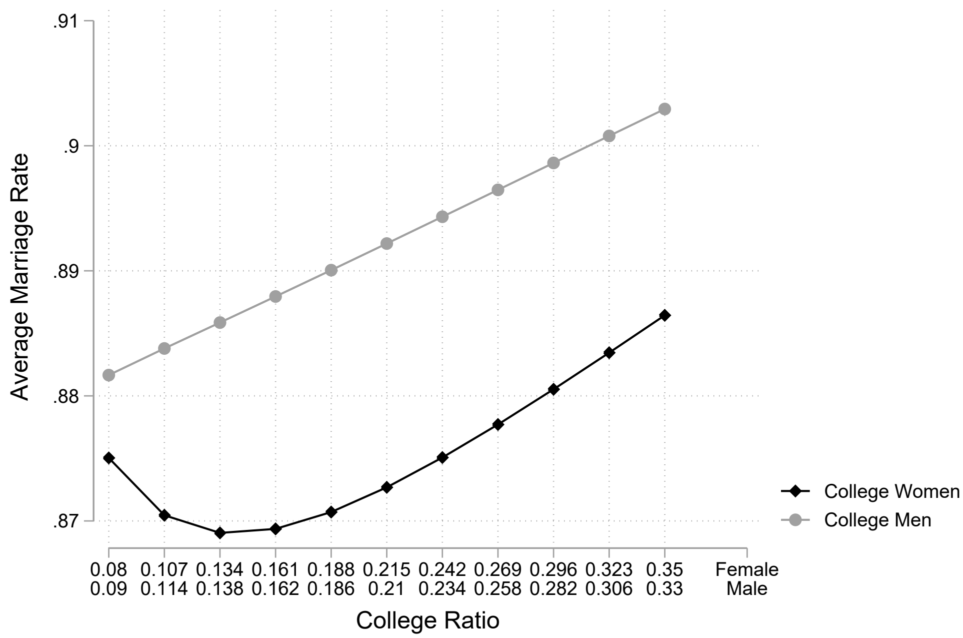

Notes: The x-axis displays the college ratios of women and men, which are allowed to evolve simultaneously.

[15]For the post-expansion cohorts, we use data in high-expansion regions. This partially incorporates cross-sectional variation.

Figure 2.4: College Expansion and Marriage Rates: Simulated Results without Local College Marriage Market

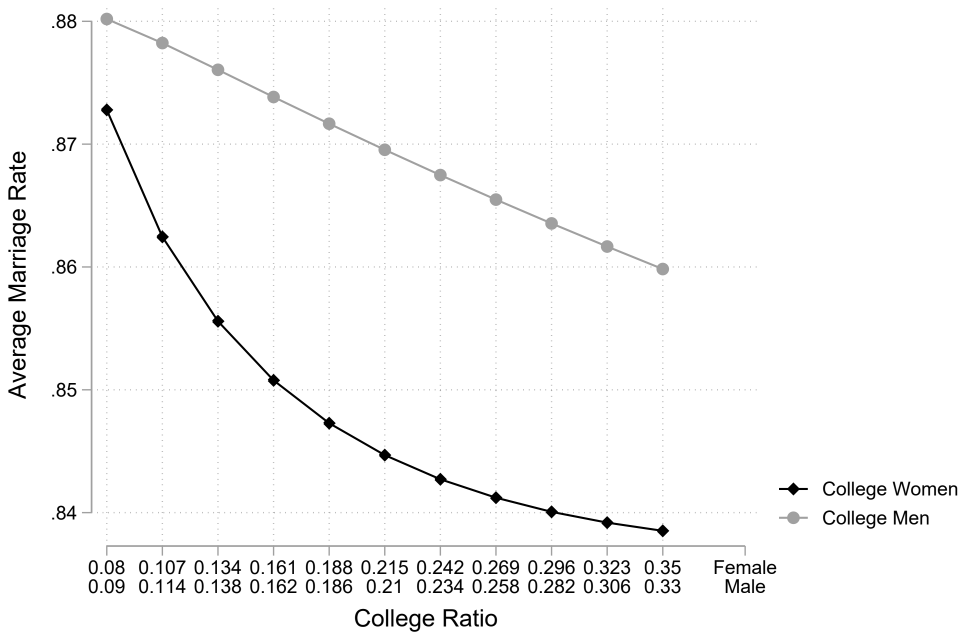

Notes: The x-axis displays the college ratios of women and men, which are allowed to evolve simultaneously.

The Local College Marriage Market Reduces Search Frictions. Through the LCMM channel, the college expansion boosts the marriage probability of college graduates by reducing search frictions. A thicker LCMM makes it easier for college students to find a preferable match within the pool of college-educated individuals. Ceteris paribus, this force pushes up the marriage probability for both college men and women. To evaluate the effects of this channel, we redo the simulation without the LCMM in the model. The results are reported in Figure 2.4. The marriage probabilities of both college women and men decrease following the college expansion. This is driven by the increasing relative supply of college types in the marriage market: because there are more college women and men, their bargaining power within marriage now decreases relative to that of noncollege types. The reduction in search frictions via the LCMM dominates the effects of changing distributions of education types in the marriage market, causing an overall positive effect as shown in Figure 2.3.

Change in Distributions of Education Types. As discussed above, a change in the distributions of education types leads to decreasing marriage probabilities for both college women and men (Becker, 1973; Choo and Siow, 2006). Importantly, the education of women increases faster than that of men. As the group size of college women grows faster than that of college men, the former become less scarce and hold less bargaining power in college-college marriages. As a result of this relative change in the group size in the college-specific marriage market, the marriage probability of college women tends to decrease by more than that of college men, as shown in Figure 2.4. Even in Figure 2.3, we also observe an initial decrease in the marriage probability of college women. The pattern implies that the effect of the LCMM dominates that of changing education distributions for college women only later, when the college expansion becomes intense enough.

How important is the gender difference in the rate of college expansion? To look into this, we simulate the baseline marriage model by allowing the college ratios of women and men to increase symmetrically. The results, as in Figure 2.5, show that the effect of the expansion is positive for college women and initially negative but later positive for college men. The reason is the social norms regarding marital preferences (Section 2.2.4). Following an increasing supply of college men, their bargaining power relative to noncollege men decreases. College women, however, rarely marry noncollege men and do not suffer from this devaluation. College men are now willing to transfer more to college women since college men are less attractive to noncollege women. As a result, the marriage rate of college women increases, while that of college men tends to decrease. The LCMM mechanism further contributes to the positive effect on college women and dominates the negative effect on college men later in the expansion.

Figure 2.5: College Expansion and Marriage Rates: Simulated Results under Gender-Symmetric Expansion

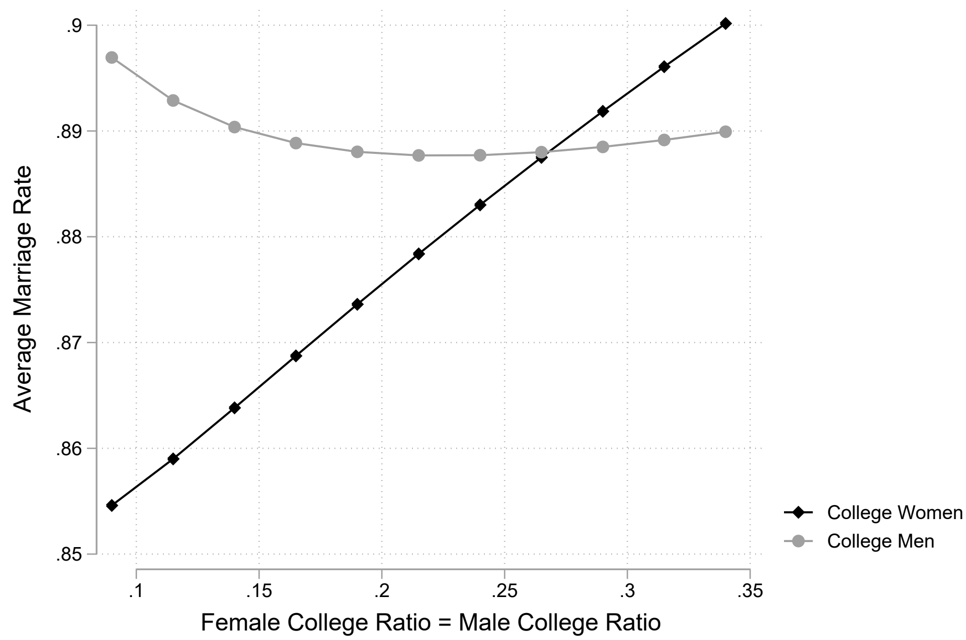

Notes: The x-axis displays the college ratios of women and men, which are allowed to evolve simultaneously.

Table 2.2: Model Predictions under Different Assumptions

		College Women	College Men
LCMM	Gender Asymmetric Expansion (Faster for Women)	$- \to +$	$+$
	Gender Symmetric Expansion	$+$	$- \to +$
No LCMM	Gender Asymmetric Expansion (Faster for Women)	$-$	$-$

Table 2.2 summarizes the model predictions about the effects of the college expansion on college graduates' marriage probabilities under different assumptions and hypothetical expansion scenarios. Both the LCMM and the change in relative distributions of education types matter for marriage outcomes according to our model and simulation. The enriched model with an LCMM predicts overall positive effects for both college women and men. As we show below, this is consistent with our empirical findings.

In Appendix 2.10.3, we discuss the simulation for noncollege groups. In Appendix 2.10.3, we report how the within-marriage transfers (τ) change in response to the college expansion in the three hypothetical situations discussed above, respectively. The changes in transfers are consistent with the discussed mechanisms. For example, the college expansion increases the transfer from college men to noncollege women and reduces the transfer from noncollege men to college women in all the three scenarios. When there is no search cost and the college enrollment increases faster for women than for me as we specified in the model, the transfer from college men to college women decreases. When the college ratio increases symmetrically between women and men, the transfer from college men to college women increases.

2.3 Background and Data

2.3.1 The Higher Education System and College Expansion in China

Most Chinese universities are public. The process of college admission is strictly controlled by the MOE, which also determines total college enrollment in the country (Feng, 1999). The amount of annual enrollment is set on the basis of the MOE's five-year plans. Each university closely adheres to its assigned quota when setting its admission plans each year. Among the different types of colleges, we focus on the regular college system ("regular colleges"), which consists of four-year and three-year colleges (the latter are analogous to two-year colleges in the United States). Though there is also a part-time postsecondary credential system, which mainly serves adults who are older than regular college students, only regular colleges experienced the expansion. In Appendix 2.8, we provide more details about the college system.

The expansion represented a sharp change in the MOE's plan. The government abruptly decided in 1999 to expand access to college in order to accommodate more youth at risk of unemployment in response to the 1997 Asian financial crisis (Wang, 2014). The policy continued even after the effects of the financial crisis subsided. Figure 2.1 displays the increase in new college enrollment nationwide before and after the college expansion. The expansion, which doubled the amount of new enrollment within three years after 1999 and even more afterward, has sharply changed the levels in education in relevant cohorts of young women and men.

The MOE implemented the college expansion mainly through scaling up enrollment in existing colleges. In Figure 2.6, we decompose the rise in total new enrollment into the increase in average enrollment per institution and the increase in the number of institutions. The overall increase is mainly driven by existing colleges. This pattern is also consistent with the assumption that the expansion has led to a thicker college marriage market (Sections 2.2.2 and 2.2.2).

Figure 2.6: Decomposing the Increase in New Enrollment

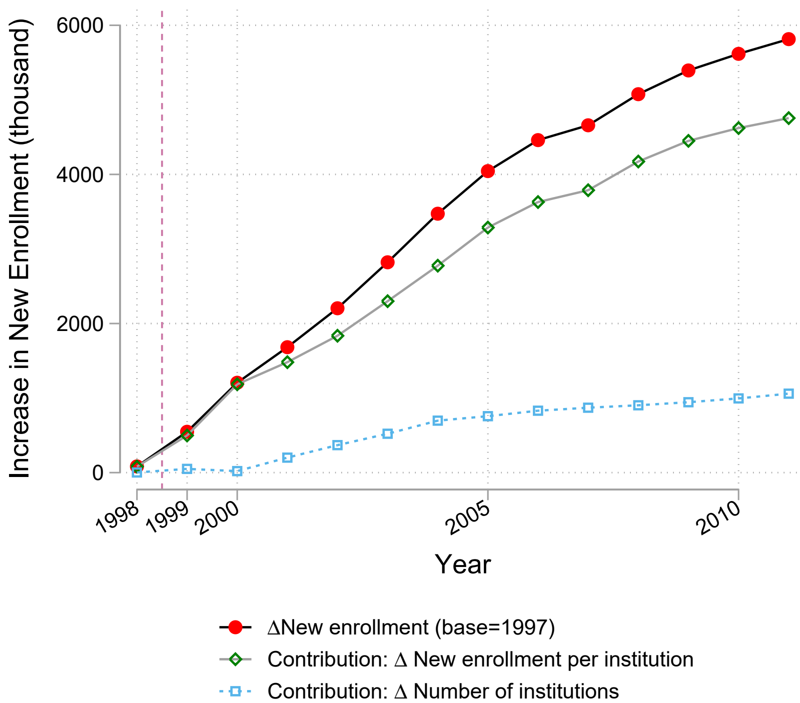

Notes: The data come from Chinese Education Yearbooks. We plot (1) the total increase in new enrollment relative to 1997; (2) the increase in new enrollment if the number of institutions did not change and the average new enrollment per institution increased as actually happened; (3) the increase in new enrollment if the average new enrollment did not change and the number of institutions increased as actually happened.

2.3.2 Measuring Exposure to the College Expansion

To identify the causal impact of the college expansion on the marriage market, we need to account for secular trends in socioeconomic and marriage-market conditions that might create a spurious relationship between being exposed to the expansion and marriage outcomes. We tackle this problem with a DID design that exploits geographic and birth-cohort variation in exposure to the expansion.

The typical age of college enrollment is 18. Because the college expansion started in 1999, we consider cohorts born in and after 1981 as post-expansion cohorts that were directly exposed to the expansion. We compare the marriage outcomes of these cohorts with those of the pre-expansion cohorts (those born before 1981). In our baseline analysis, the 1975–78 cohorts are used as the pre-expansion comparison cohorts. We discuss the choice of post- and pre-expansion cohorts in Section 2.4.

Critical to our empirical strategy is variation in the intensity of the college expansion across regions. The expansion was implemented mainly via scaling up enrollment in existing universities. As a result, regions with more historical higher education resources naturally benefited more from the expansion. This regional variation has been exploited in various previous studies on the expansion (Feng et al., 2018; Ge and Huang, 2020; Li et al., 2017; Ma, 2020). Motivated by this variation, we construct a proportional-expansion proxy based on the historical abundance of university resources.

In our baseline analysis, we look at local marriage markets at the prefecture level.[16] For each prefecture p, we construct the proportional index as follows:

$$ExpProxy_p = \frac{CollegeEnrollment_p^{1982}}{PopSize_p^{1982}} \tag{2.4}$$

[16]The prefecture is the subprovincial geo-unit in China. We constructed a consistent set of prefecture units to accommodate historical division changes. There are 340 prefecture units in our sample for analysis.

This proportional-expansion proxy measure is the ratio of college enrollment to population size as of 1982. It is constructed using microdata from the 1982 China census provided by IPUMS International. In Figure 2.7, we show the geographical distribution of the intensity of the expansion as measured by this proxy. Because there is substantial variation in historical college resources across prefectures, the same is true for the proxy. High-intensity regions (as measured by this proxy) do not show obvious geographic patterns, with the one exception that prefectures in the Northeast tend to have higher intensity.

This proxy is highly predictive of college enrollment in later years. Figure 2.8 plots the (log) enrollment ratio in 2005 against the (log) expansion proxy. Their correlation can be approximated using a straight fitted line with a slope of 0.67 and an R-squared of 0.44.[17] Therefore, the expansion proxy provides sufficient variation for identifying differing responses of local marriage markets to the expansion.

In Appendix 2.11, we provide several additional tests and arrive at two findings. First, the power of the expansion proxy to predict college enrollment in later years is robust across years. Second, the treatment proxy is not associated with economic growth or sex ratio, both of which can affect the marriage market in important ways (Burgess et al., 2003; Chu et al., 2018; Hankins and Hoekstra, 2011; Wei and Zhang, 2011; Ebenstein and Sharygin, 2009). The orthogonality of the treatment proxy in relation to these economic and marriage-market conditions also provides supportive evidence for the parallel-trends assumption of our DID design.[18] Previous studies on China's college expansion have also discussed the validity of exploiting variation in historical college endowments. Feng et al. (2018), Li et al. (2017), and Ma (2020) show that historical college endowment is highly predictive of the number of college graduates after the expansion, and Ma (2020) documents that college endowment in 1982 is not associated with changes in GDP

[17] If college enrollment in later years is perfectly proportional to the initial endowment measured using the 1982 data, the slope of the line should be 1.

[18] In our empirical analysis, we also provide robustness tests that directly control for local GDP per capita and sex ratio.

or population.

Figure 2.7: Geographical Distribution of the Proxy for College-Expansion Intensity

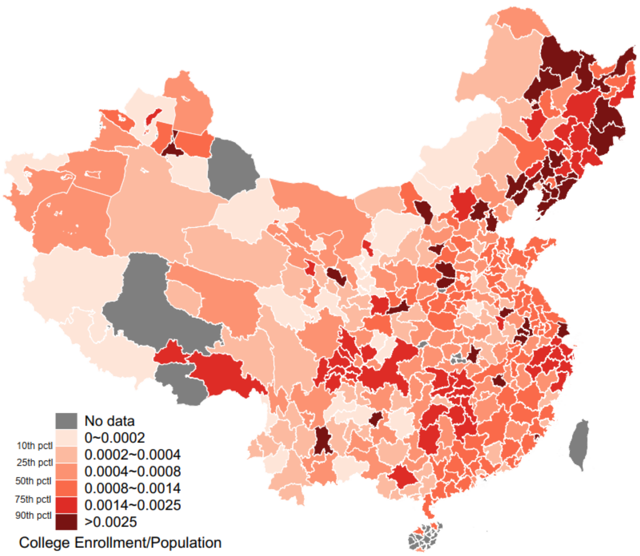

Notes: The proportional proxy for the magnitude of the college expansion is constructed using 1982 census microdata of China, obtained via IPUMS International.

Figure 2.8: Initial College-Expansion Proxy and College Enrollment in 2005

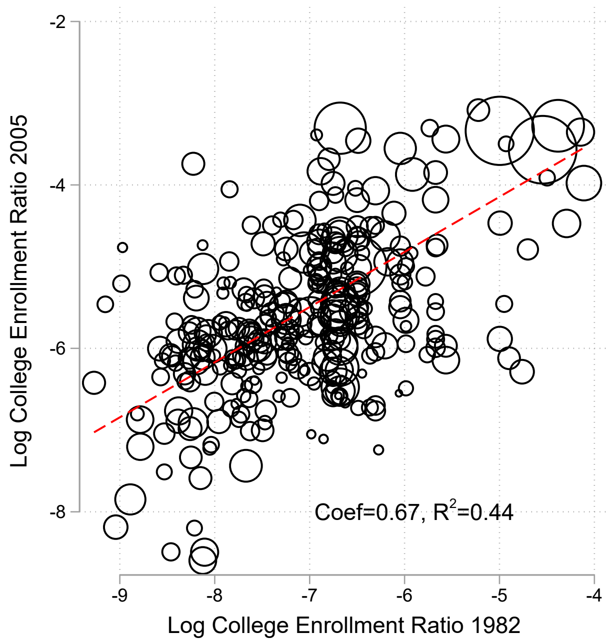

Notes: The 1982 ratio of college enrollment to population is constructed using 1982 Chinese-census microdata obtained via IPUMS International. The 2005 ratio of college enrollment to population comes from the 2005 Chinese mini-census. The plots are weighted using the population size of each prefecture.

2.3.3 Marriage Outcomes

We obtain information about marriage outcomes using the confidential 2010 Chinese census and the confidential 2015 Chinese mini-census. The 2015 mini-census sample that we use is a 0.15% random sample of the population. The 2010 census sample is a 0.35% random sample of the population.[19] For our post-expansion cohorts (born after 1981), we construct their marriage outcomes from the 2015 mini-census data. Corresponding marriage outcomes of the pre-expansion cohorts are constructed using the 2010 census data. The 2010 census data set contains information about individuals' marital history.[20] For example, when we choose the 1975–78 birth cohorts as the pre-expansion cohorts and the 1981–84 birth cohorts as the post-expansion cohorts, we impute

[19] Both data sets are accessed via the Shanghai University of Finance and Economics.

[20] Unfortunately, similar marital-history information does not exist in the 2015 mini-census data set.

the marriage outcomes for the pre-expansion cohorts based on their marital status as of 2009, when they were aged 31–34, the same ages as the post-expansion cohorts (born 1981–84) in 2015.

2.4 Empirical Strategy

2.4.1 Baseline Empirical Strategy

In our DID design, we compare cohorts exposed to the expansion (post-expansion cohorts) to cohorts not exposed to it (pre-expansion cohorts) in different local marriage markets. We consider cohorts born after 1981 as the post-expansion cohorts. We divide the post-expansion cohorts into two groups: (1) the 1981–84 birth cohorts ("early post-expansion cohorts"), and (2) the 1985–88 birth cohorts ("late post-expansion cohorts"). The former group was 31–34 years old as of 2015, while the latter was 27–30 years old. Therefore, these two groups capture the impacts of the expansion on different cohorts at different ages.

The choice of pre-expansion cohorts requires further discussion. For two reasons, cohorts born too close to 1981 do not constitute good comparison units. First, the age of college enrollment at 18 is a norm, not a binding constraint. Therefore, cohorts born slightly before 1981 were still partially exposed to the expansion because they enrolled in high school late or retook the college entrance examination. Second, and more importantly, as the marriage market is a two-sided matching market, cohorts that were not directly exposed to the expansion were very likely affected in general equilibrium, as some of their potential partners were exposed to the expansion. Therefore, the ideal pre-expansion cohorts are far enough from the onset of the expansion that they are not subject to such spillover effects. However, choosing control cohorts born too much earlier than the treatment may make the two groups less comparable. This is a concern since China experienced significant cultural, social, and economic shifts when the cohorts that we investigate grew up.

For our baseline analysis, we use the 1975–78 cohorts as the pre-expansion comparison group. These cohorts reach college age as close to the expansion shocks as seems sensible given the potential spillovers. These cohorts serve as the comparison group for both the 1981–84 and 1985–88 post-expansion cohorts. To make the marriage outcomes comparable between the post-expansion and the pre-expansion groups, we draw information from marital histories recorded in the 2010 census for the pre-expansion cohorts. The early post-expansion cohorts (born in 1981–84) were observed in 2015, when they were 31–34 years old. For the 1975–78 cohorts, we construct their marriage outcomes as of 2009, when they were similar in age to the post-expansion cohorts in 2015. Meanwhile, the late post-expansion cohorts (born in 1985–88) were 27–30 years old in 2015. Therefore, when we use the 1975–78 cohorts as their comparison group, we look at the marriage outcomes of the earlier cohorts as of 2005, when they were also comparable in age.

We estimate the causal impacts of the college expansion on marriage outcomes using the following specification:

$$y^k_{iapt} = \beta^k_0 + \beta^k_1 ExpProxy_p * Post_t + \beta^k_2 * Post_t + \lambda^k_p + \xi^k_a + \varepsilon^k_{iapt} \tag{2.5}$$

y^k_{iapt} is the outcome (for example, ever married) of individual i at age a in local marriage market (prefecture) p. The subscript t captures whether the observation belongs to a post-expansion or pre-expansion cohort. $Post_t$ is a dummy variable, with $Post_1 = 1$ indicating post-expansion cohorts. We also control for prefecture (λ^k_p) and age (ξ^k_a) fixed effects. $ExpProxy_p$ is the proxy for local exposure to the expansion (described in Section 2.3.2).

In the following analysis, we standardize the treatment proxy $ExpProxy_p$ so that one unit represents one standard deviation (SD) across all prefectures. Thus, the coefficient on the interaction term $ExpProxy_p * Post_t$ estimates the change in marital outcomes in a prefecture if we increase the treatment proxy by one SD. A one-SD increase in the treatment proxy is also approximately twice the size of the increase from the 25th percentile to the 75th percentile in the distribution across prefectures. As an example, a one-SD difference in the proxy is approximately the difference between Shanghai, one of the most developed and educated cities in China, and Wuhu, an inland prefecture that ranked 235th in population size and 82nd in GDP in 2015.[21]

2.4.2 Pre-expansion Cohorts and Graphical Evidence

We choose the 1975–78 cohorts as the pre-expansion cohorts in our baseline analysis in an attempt to resolve the tension between choosing pre-expansion cohorts that are too far from the expansion and those too close to it. However, it is helpful to check whether the secular trends in marriage patterns in the pre-expansion cohorts are parallel between regions with higher exposure to the expansion and regions with lower exposure. Largely parallel pre-trends can provide supportive evidence for the parallel-trends assumption. In addition, they may suggest that the tension regarding the choice of pre-expansion cohorts does not greatly affect our findings. In this section, we provide preliminary graphical evidence on the pre-trends.

We divide the local marriage markets (prefectures) into two groups: those with high expansion intensity and those with low intensity. The former group includes prefectures whose proxy for expansion is above the median, and the rest are included in the latter. The fraction of college graduates who were ever married at 27–30 years old is plotted in each year for both groups. The results are displayed in Figure 2.9 (we report the graphical evidence for parallel pre-trends of noncollege groups in Figure 2.26). Most college graduates in the early-1970s cohorts were already married at this age (nearly 90% of college women and about 75% of college men). The marriage probabilities declined overall in subsequent cohorts. For the 1985–88 cohorts (late post-expansion cohorts), the numbers are more than 70% for college women and more than 60% for college men. The marriage probabilities have consistently been higher in low-expansion regions than in high-expansion regions.

In Figure 2.9, left of the vertical line are the pre-expansion cohorts. Looking through these

[21] We do not weight using population size when standardizing the treatment proxy across prefectures. Weighting the proxy leads to very similar results: the treatment proxy is only rescaled by a constant factor, and the magnitudes are similar whether weighted or not. We execute all regressions at the individual observation level, effectively weighting the regressions by population size.

cohorts, it seems that the high- and low-expansion regions followed largely parallel trends before being affected by the expansion. We conduct formal tests regarding parallel pre-trends in Section 2.5.2, the results of which are consistent with the preliminary graphical evidence. After the expansion, the differences between the high- and low-expansion regions decreased. This decrease suggests the expansion increased marriage probabilities if we assume that the high- and low-expansion regions would have continued the parallel trends in the absence of the expansion.

Figure 2.9: Trend in Marriage Rates of College Graduates by College-Expansion Intensity

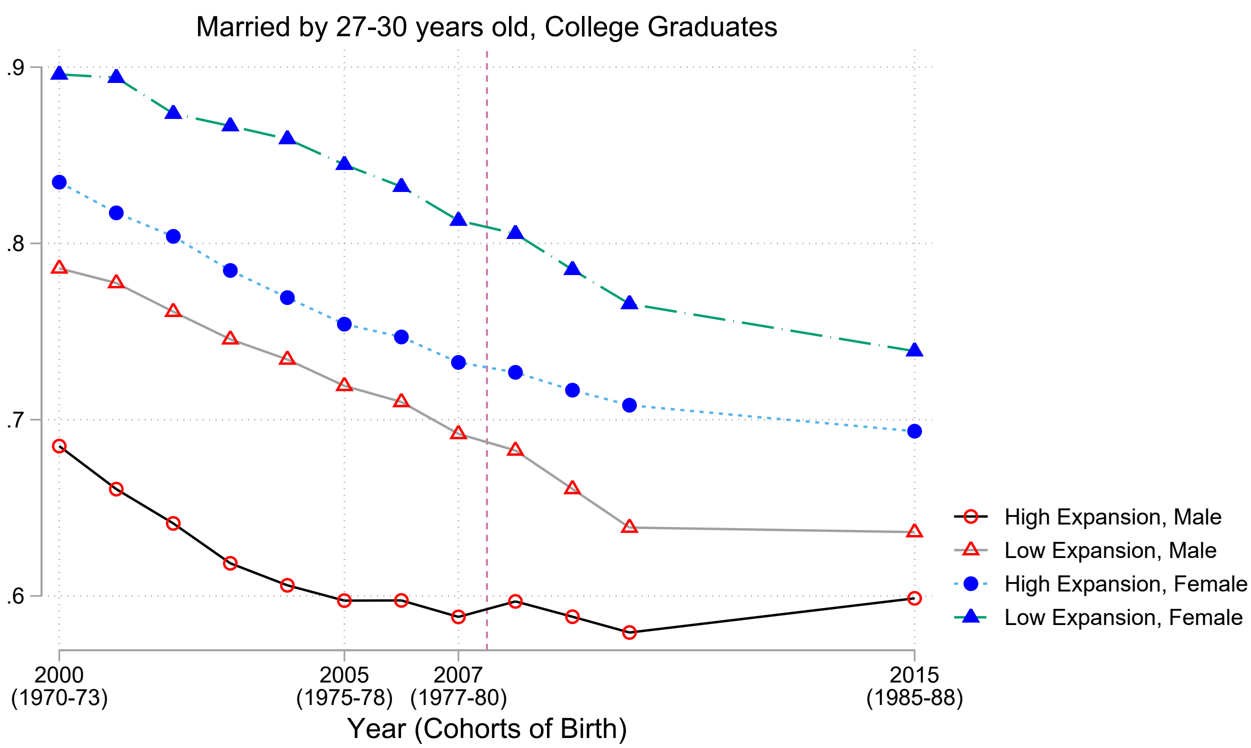

Notes: The sample is divided into two groups based on the value of the expansion treatment proxy below or above the median. Left of the vertical line are samples that consist of only pre-expansion cohorts. The marital-history information in the 2010 census allows us to impute the ever-married fractions before 2010. Using the 2015 mini-census, however, we can only know people's marital status as of 2015.

2.5 Impacts of the College Expansion on Marriage Probability

2.5.1 Baseline Results

In this section, we quantify the effects of the college expansion on marriage probability. We focus on college women and men. Our model predicts that their marriage rates increased in response to the expansion. We test this prediction using the DID strategy discussed in Section 2.4.1.

We report the baseline estimates using the DID specification (Equation 2.5) in Table 2.3. The first two columns present estimates (separately for college and noncollege) in which we examine the effects on the late post-expansion cohorts (1985–88) and measure their marriage rates at ages 27–30. Columns (3) and (4) present estimates for the early post-expansion cohorts (1981–84), whose marriage rates are measured at ages 31–34. Here we focus on the estimates for college graduates (Columns (1) and (3)). In Panel A, we present the estimates for men, and in Panel B, we present the estimates for women. The gender gap in the effects on the same cohorts is reported in Panel C.

The results suggest positive and modest impacts of the expansion on the probability of marriage for college graduates. Column (1) shows that when we look at the late post-expansion cohorts (27 to 30 years old in 2015), a one-SD increase in the expansion-treatment proxy raised the marriage probability by 2.7 percentage points for college men and 1.7 percentage points for college women.[22] If we compare a prefecture at the 75th percentile to one at the 25th percentile of the distribution of the expansion proxy, the expansion led to an increase of 1.44 percentage points in the marriage probability of college men from the late post-expansion cohorts and 0.66 percentage points for their college-women counterparts.

The magnitudes of these effects, though not very large, are still meaningful. For example, they are comparable to the changes in the mean marriage probabilities presented in Table 2.3.

[22] The magnitudes of these effects are approximately those of college graduates in Shanghai versus Wuhu, as discussed in Section 2.4.1.

For college men in the late post-expansion cohorts (born in 1985-88), the increase in marriage probability driven by a one-SD increase in the treatment proxy (2.7 percentage points) is 88% of the the overall decrease in marriage probability of college men aged 27–30 between the pre-expansion and late post-expansion cohorts (3.1 percentage points). Looking at college women from the late post-expansion cohorts, the effect of a one-SD increase in the treatment proxy (1.7 percentage points) is equivalent to 22% of the overall decline in marriage probability (7.8 percentage points).[23]

The effects are smaller for college graduates from the early post-expansion cohorts (31–34 years old in 2015). Column (3) shows that a one-SD increase in the treatment proxy leads to an increase of 1.24 percentage points in marriage probability for college men and 0.53 percentage points for college women. Moreover, the estimated effect on college women is statistically nonsignificant. The effect on college men is still meaningful, as it is equivalent to 50% of the decrease in marriage probability between the pre-expansion and early post-expansion cohorts (2.5 percentage points).

[23] Among all the groups in Table 1, college women aged 27 to 30 experienced the biggest drop in the ever-married rate.

Table 2.3: Effects of the College Expansion on the Probability of Marriage: Baseline Results

Dependent variable: Ever being married				
	(1)	(2)	(3)	(4)
Post-expansion cohorts	1985-88, 27-30 years old in 2015		1981-84, 31-34 years old in 2015	
Pre-expansion cohorts	1975-78, 27-30 years old in 2005		1975-78, 31-34 years old in 2009	
	College	Noncollege	College	Noncollege
A. Male				
ExpProxy∗Post	0.0273***	0.0059***	0.0124***	0.0008
	(0.0049)	(0.0020)	(0.0039)	(0.0014)
Observations	40196	187259	36486	181105
Marriage rate of				
Pre-expansion cohorts	0.644	0.793	0.883	0.882
Post-expansion cohorts	0.613	0.728	0.858	0.874
B. Female				
ExpProxy∗Post	0.0171***	0.0020	0.0053	-0.0009
	(0.0037)	(0.0034)	(0.0040)	(0.0012)
Observations	38477	182842	33894	176213
Marriage rate of				
Pre-expansion cohorts	0.787	0.916	0.919	0.961
Post-expansion cohorts	0.709	0.873	0.892	0.953
C. Female − Male				
	-0.0102*	-0.0039	-0.0072	-0.0017
	(0.0061)	(0.0039)	(0.00560)	(0.0018)

Notes: The marital outcome of pre-expansion cohorts is constructed using marriage history so that it is comparable to the post-expansion cohorts. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit represents one standard deviation across all prefectures. All regressions control for prefecture fixed effects and age fixed effects. Standard errors clustered at the prefecture level are in parentheses. There are 340 clusters. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The difference between the estimated effects for the early and late post-expansion cohorts may arise for two reasons. First, the magnitude of the expansion was much larger in later years than in the first few years (Figure 2.1). Therefore, the effects of the expansion are probably larger for the late post-expansion cohorts. Second, we observe the early post-expansion cohorts at older ages (31–34) than the late post-expansion cohorts (27–30). If the expansion reduced the average age of first marriage and did not indicate permanent changes in marriage probability, then the estimated effects should be smaller when we observe the post-expansion cohorts when they were older than when they were younger. This age effect may also (partially) explain the difference between the early and late post-expansion cohorts. However, because of a data limitation, we cannot distinguish between these two explanations: age and cohort are perfectly co-linear in the cross-sectional post-expansion data.[24]

We also observe gender differences in the estimated effects of the expansion, as shown in the last row of the table. Although most of the differences are imprecisely estimated, the point estimates suggest that the expansion had a larger effect on the probability of marriage for college men than women. In the table, we look at women and men from the same cohorts so they were exposed to the same degree of expansion. Because of the marriage age gap (women on average marry men older than themselves), however, women and men from the same cohorts are not in exactly the same marriage market. Some college women in the post-expansion cohorts may marry older men who were not exposed to the expansion or were exposed to a smaller expansion. The effects of the reduction in search frictions are arguably smaller for these women than men from the same cohorts because college enrollment of the women's potential spouses increased by a smaller magnitude. In Section 2.5.6, we further discuss this issue.

In Columns (2) and (4), we observe no effects of the expansion on the noncollege groups except for noncollege men from the late post-expansion cohorts. Our simulated comparative statics, how-

[24] With soon-available 2020 census microdata, we will be able to test these two explanations. For example, we will be able to observe the late post-expansion cohorts (1985–88) when they were 31–34 years old and test whether the effect on their marriage probability is still larger than the effect for the early post-expansion cohorts.

ever, predict that the marriage rate of noncollege men decreased in response to the expansion. The contrast between the theory and the empirical results implies that the model does not capture all key features of the noncollege marriage market. The expansion not only directly changed the distribution of college attainment but generated spillover effects on the distribution of below-college education types. In results not reported in this paper, for example, we find that the expansion increased the rate of high school graduation among those with less than a college education. To the extent that high schools might also serve as local marriage markets, the expansion likely also reduced search frictions in high school marriage markets. The mechanisms, however, are beyond the scope of this paper.

Overall, the results suggest an economically sizable impact of the expansion on college graduates' marriage probability. The positive estimated effects are in line with the simulation results of our model, in which the reduction in search frictions dominates the effects of only altering the relative distribution of education types. In Sections 2.5.2–2.5.6, we discuss a series of robustness checks for our main finding.

Education Penalty for Women in the Marriage Market

In recent years, the rising phenomenon of educated women marrying later and at a lower rate has attracted public attention and raised policy concerns (Fincher, 2016; Magistad, 2013; To, 2015). Professional women in their late 20s and 30s who are still unmarried have even been labeled with the derogatory term "leftover women" by both the general public and the government.[25] The rapid rise of college-educated women in combination with traditional gender norms (Section 2.2.4), arguably, is contributing to this trend. The mean values presented in Table 2.3 indeed show that the marriage probability of college women is much lower than that of noncollege women at the same age for all cohorts. Nevertheless, we show that expanding access to college education

[25] The national lexicon in China included the term "leftover women" when the Ministry of Education of the People's Republic of China made an official statement in 2007, defining it as any unmarried woman who is 27 years or older.

generates positive spillover effects on the marriage probability of college women. Our findings add to the field's understanding of the effects of college education on women's marriage outcomes and on gender inequality. Education policies that expand access to higher education, our results suggest, may actually raise the marriage prospects of highly educated women. Our stance is *not* that marriage is the better option for all educated women. Instead, our research indicates that the college expansion has broadened the choice set available to women who have received higher education.

Effects on Permanent Marriage Rates and Early Marriage

The earliest cohort that was directly exposed to the expansion (born in 1981) was only 34 years old when observed in the 2015 mini-census data. Therefore, we do not observe the marriage outcomes of the post-expansion cohorts at an old-enough age to determine whether the effects indicate permanent changes in marriage outcomes or temporary changes (for example, earlier marriages). Figure 2.10 reports the age profile of marriage probabilities by gender and college education. As suggested by the cross-sectional age profiles, most people eventually get married. If everyone in the post-expansion cohorts in our sample follows this pattern, then our results probably indicate effects of the expansion on earlier marriage rather than a permanent increase in marriage rates. Nevertheless, the age profile mostly stabilizes after the mid-30s. Therefore, the positive effects on the early post-expansion cohorts, who were 31–34 years old as of 2015, might at least partially reflect an increase in permanent marriage rates.

The profiles depicted in Figure 2.10 are based on cross-sectional data and do not tell us about the post-expansion cohorts when they grew older. After microdata from the 2020 Chinese census become accessible to us, we will be able to more confidently determine whether the expansion has affected the permanent marriage rates of the post-expansion cohorts.

Figure 2.10: Age Profiles of Ever-Married Rate

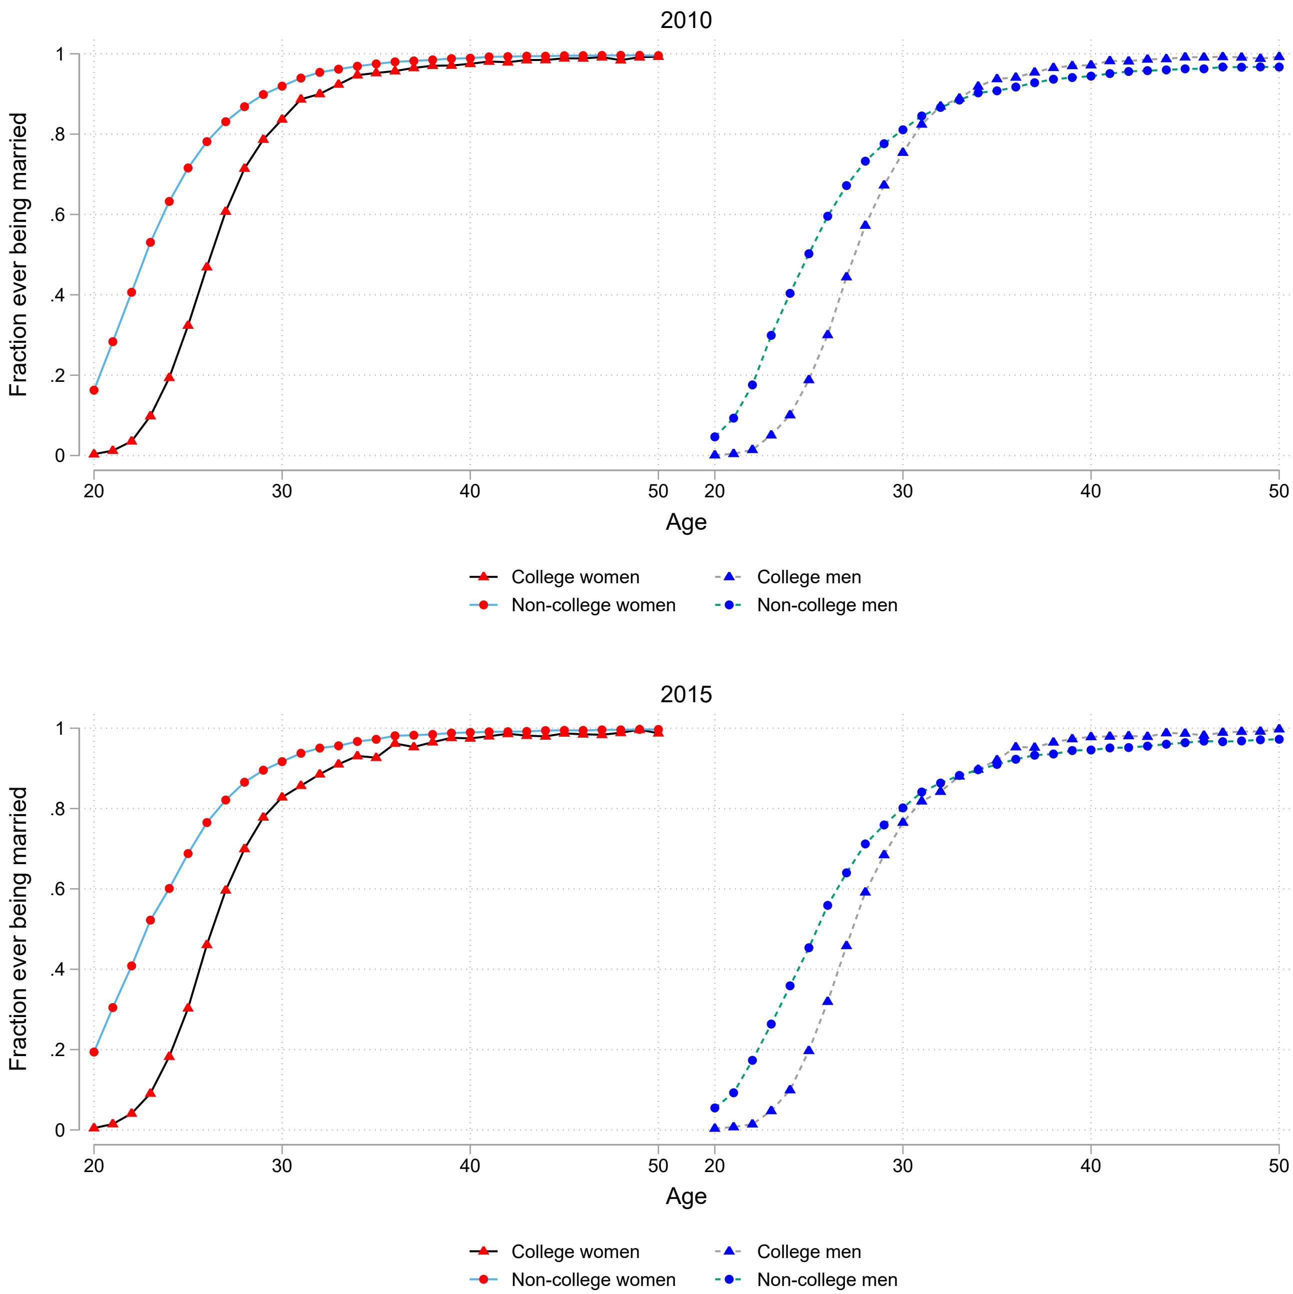

Notes: Data come from random samples of the 2010 census and 2015 mini-census of China. Each plot shows the fraction of people who were ever married by age in the cross-sectional data.

2.5.2 Falsification Test and Sensitivity to Pre-expansion Cohorts

We chose the 1975–78 birth cohorts as our control cohorts for the baseline analysis. The choice, as discussed in Section 2.4.1, reflects the tension between avoiding potential spillovers in the marriage market and ensuring that the pre-expansion and post-expansion cohorts are comparable. In this section, we investigate whether our findings are sensitive to the choice of pre-expansion cohorts; this exercise also provides suggestive evidence on whether the pre-trends are parallel. To test the sensitivity of estimated results at a more granular level and probe potential spillover effects in the marriage market, we further divide the post-expansion cohorts into four groups: 1987–88 (27–28 years old in 2015), 1985–86 (29–30 years old in 2015), 1983–84 (31–32 years old in 2015), and 1981–82 (33–34 years old in 2015). This also enables us to better track the different treatment effects across post-expansion cohorts.

Let us take the 1987–88 post-expansion cohorts as an example. They were observed at the ages of 27 and 28 in the 2015 census. For the sensitivity test, we use different pre-expansion cohorts and consider their marital status at the same ages (for example, the 1977–78 cohorts in 2005). We then apply our DID model (Equation 2.5) to these pre-expansion and post-expansion cohorts. Varying the choice of the pre-expansion cohorts, we obtain a series of estimates for the effects of the expansion on the marriage probability of the 1987–88 cohorts in 2015.

The results of the sensitivity test are displayed in Figure 2.11. We report results for college women and men. Each dot represents the estimated treatment effect for one DID regression. We plot the coefficients and 95% confidence intervals by holding the post-expansion cohorts fixed and varying the pre-expansion cohorts. Two patterns emerge. First, the coefficients are mostly stable. Therefore, our baseline findings are not sensitive to the choice of pre-expansion cohorts.

Second, when estimating the model with pre-expansion cohorts that reach college age just before the college expansion (born in 1979 or 1980), we observe much smaller estimated effects for college men compared to estimated effects using the pre-expansion cohorts as in our baseline analysis (shown in bold on the x-axis in Figure 2.11). We do not observe a similar pattern for college women. This is consistent with the possibility of spillover effects in the marriage market. Men in China usually marry younger women (the average husband-wife age gap is about two years). Therefore, college men born during 1979–80 were partially treated since their potential spouses were directly exposed to the expansion, while college women born during 1979–80 were much less likely to be treated. This again justifies our choice of the pre-expansion cohorts in our baseline analysis. We further discuss this issue below in a falsification test based on the pre-expansion data.

Figure 2.11: Sensitivity Test for Pre-expansion Cohorts

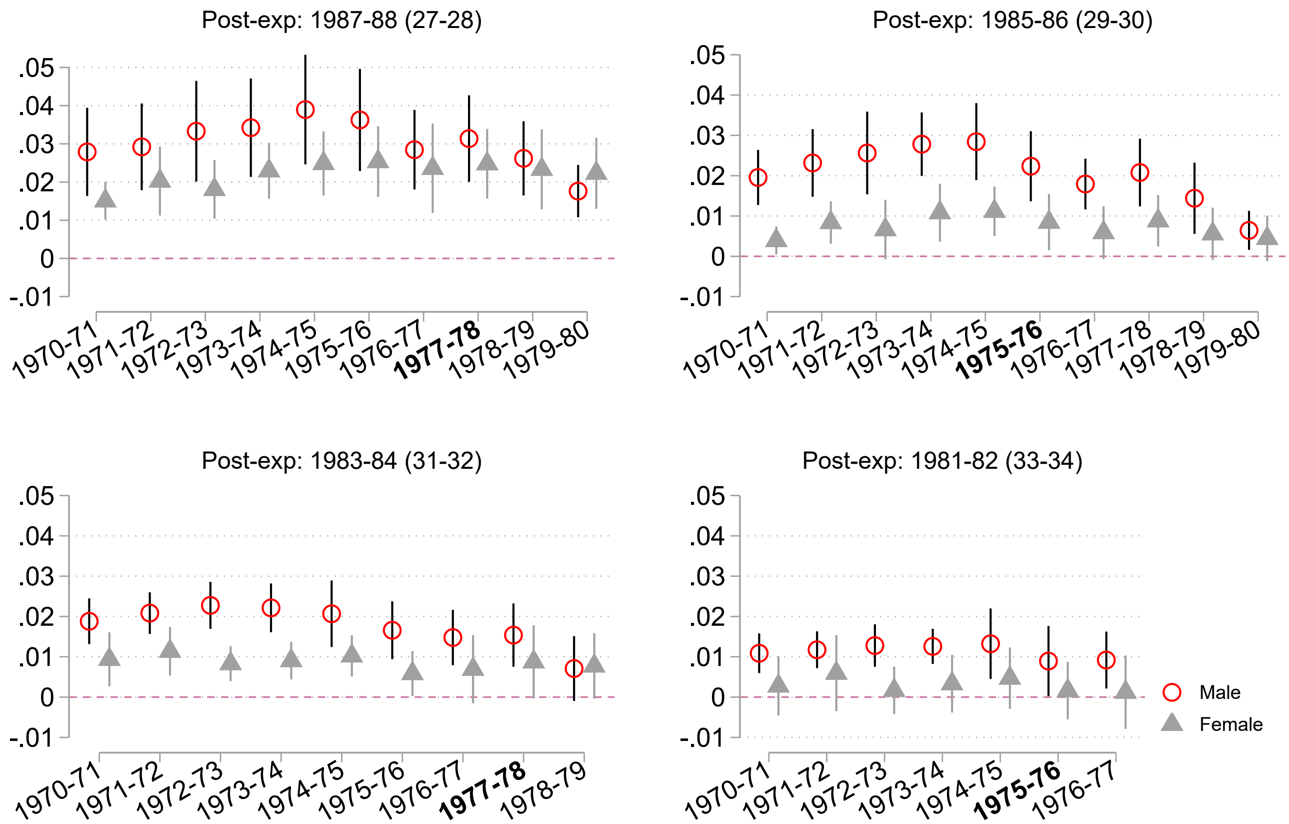

Notes: Each panel plots the difference-in-differences coefficients estimated by combining the post-expansion cohorts in its title and different pre-expansion cohorts on its x-axis. The specification is the same as Equation 2.5. Point estimates for college women and men and 95% confidence intervals, estimated by clustering at the prefecture level, are plotted. Pre-expansion cohorts in **bold** on the x-axis are the cohorts used for our baseline results.

We also conduct a falsification test considering only the pre-expansion cohorts. In the test, we assume that the 1975–76 cohorts were not exposed to the expansion while other cohorts were, and therefore we take the 1975–76 cohorts as the benchmark pre-expansion cohorts. The placebo treatment effects are then estimated by using different placebo "post-expansion" cohorts and the 1975–76 pre-expansion cohorts.[26] If our observed effects on marital outcomes are indeed caused

[26]We choose 1976 as the cutoff because it is the midpoint of our baseline pre-expansion cohorts (1975–78). We can therefore formally conduct the falsification test using the baseline pre-expansion cohorts.

by the expansion, then the treatment should have no effects in these placebo tests. These tests can help us formally detect whether the parallel-trends assumption is violated in the pre-expansion cohorts.

Figure 2.12: Falsification Test

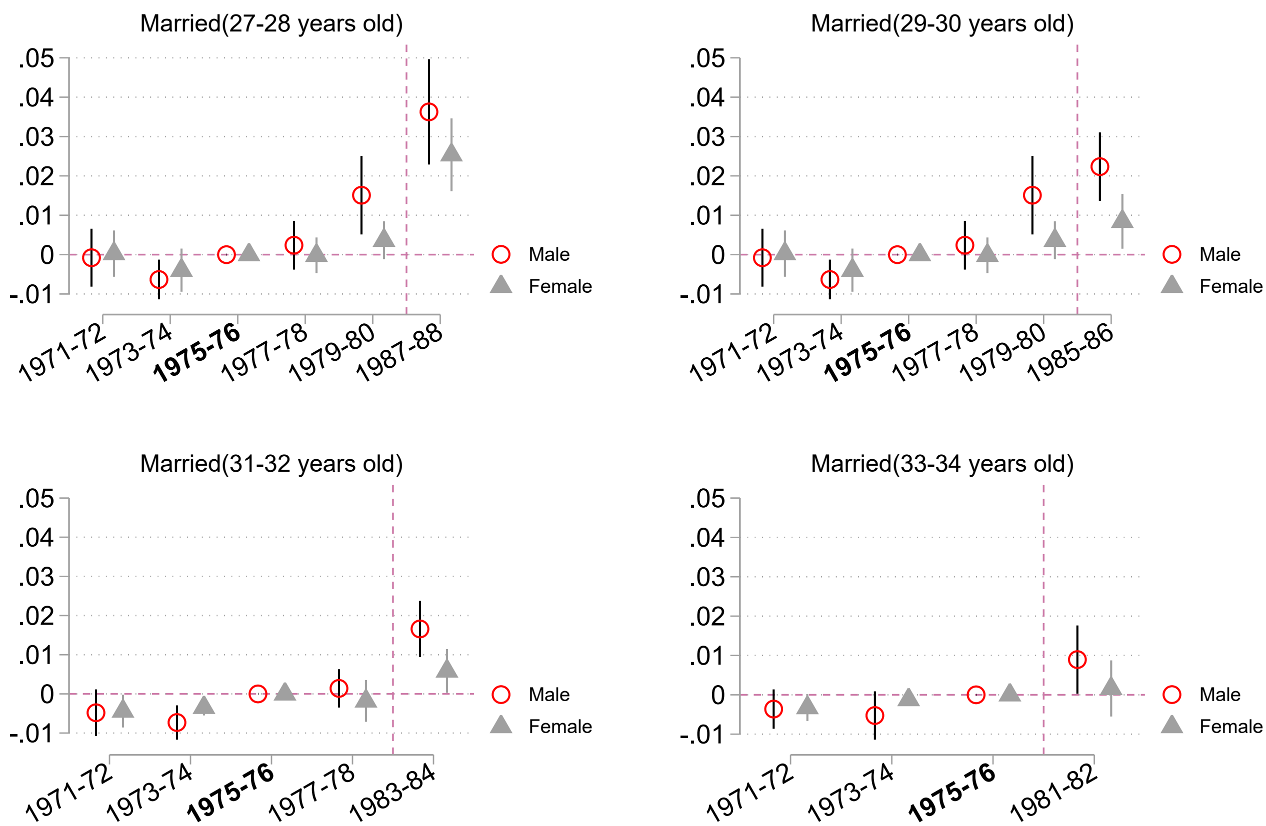

Notes: Each panel plots the difference-in-differences coefficients. The pre-expansion cohorts are fixed at 1975–76 (in bold). The x-axis displays different falsification "post-expansion" cohorts. The specification is the same as Equation 2.5. Point estimates for college women and men and 95% confidence intervals, estimated by clustering at the prefecture level, are plotted. Right of the vertical line are estimated coefficients using the actual post-expansion cohorts (cohorts born later than 1980). Because we can only obtain marital-history information from the 2010 census, we cannot use the 1979–80 cohorts when the outcome is ever being married at the ages of 31–32 *or* the 1977–80 cohorts when the outcome is ever being married at the ages of 33–34.

The results are reported in Figure 2.12. We fix the pre-expansion cohorts at 1975–76 (in bold) and estimate the DID specification using falsification "post-expansion" cohorts after or before them. The estimates are displayed in a similar way as in Figure 2.11). Consistently with the suggestive patterns from the previous sensitivity test, we largely find no effects of the placebo expansion. One exception is the 1979–80 cohorts. The placebo expansion seems to have positive effects on college men from these cohorts. Considering the marriage age gap, the potential spouses of these men were directly impacted by the (actual) college expansion. This can raise the marriage prospects of these men because of (1) a reduction in search frictions and (2) a relative increase in the number of college women in the marriage market.[27] We only observe positive spillover effects on the marriage probability of college men from these pre-expansion cohorts. This pattern further supports the view that our empirical results are driven by the college expansion.

2.5.3 Continuous-Treatment DID

Recent studies of DID designs have pointed out potential issues with continuous treatment variables. First, the traditional OLS estimator of the two-way fixed-effects model may load negative weights on treatment effects of some units with a multivalued or continuous treatment variable. If there are heterogeneous treatment effects, the negative weights lead to bias in the estimated causal effects (de Chaisemartin et al., 2022a). To address this concern, we adopt a heterogeneity-robust DID estimator (de Chaisemartin et al., 2022a, 2019). The results are reported in Table 2.4.

The alternative estimator delivers largely qualitatively robust findings. For college men and women from the late post-expansion cohorts, the new point estimates are larger than our baseline estimates. For college men from the early post-expansion cohorts (1981–84), however, the new estimate is still positive, but it is smaller and no longer statistically significant.

The second problem is selection bias associated with different treatment intensities (Callaway and Sant'Anna, 2021). If we interpret the DID estimator as the weighted average of treatment

[27] The college expansion did not directly increase the college ratio of men in the 1979–80 cohorts. However, it raised the college ratio of women who are younger yet potentially in the same marriage market. This tended to raise the marriage prospects of college men by reducing the sex ratio of college-educated people.

Table 2.4: Fuzzy-DID Estimator (de Chaisemartin et al., 2022)

Dependent variable: Ever being married				
	(1)	(2)	(3)	(4)
Post-expansion cohorts	1985-88, 27-30 years old in 2015		1981-84, 31-34 years old in 2015	
Pre-expansion cohorts	1975-78, 27-30 years old in 2005		1975-78, 31-34 years old in 2009	
	College	Non-college	College	Non-college
A. Male				
ExpProxy$*$Post	0.0532***	0.0066	0.0054	-0.0011
	(0.0133)	(0.0109)	(0.0090)	(0.0067)
	[0.0273]	[0.0059]	[0.0124]	[0.0008]
Observations	40196	187259	36486	181105
B. Female				
ExpProxy$*$Post	0.0373***	0.0010	0.0080	-0.0022
	(0.0133)	(0.0080)	(0.0050)	(0.0038)
	[0.0171]	[0.0020]	[0.0053]	[-0.0009]
Observations	38477	182842	33894	176213

Notes: This table reports the estimates using the fuzzy-DID estimator for all college graduates by gender. The marital outcome of pre-expansion cohorts is constructed using marriage history so that it is comparable to the post-expansion cohorts. *ExpProxy* is the proxy for college expansion, which is standardized such that one unit represents one standard deviation across all prefectures. All regressions control for prefecture fixed effects and age fixed effects. The bottom quintile is used as the benchmark "untreated" group for the estimator. Baseline estimates are in brackets. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

effects when moving between contiguous treatment values, there might be a selection bias unless all the treated units would have shared the same path of outcomes if they had received treatment of the same value. In our case, there could be a positive selection bias if the treatment effect in a high-treatment prefecture would have been higher than that in a low-treatment prefecture if the former had also received the low treatment. Because we can only observe one realized treatment value for each prefecture, there is no easy fix for this potential problem. Such a selection bias is likely to appear, for example, if the treated units can select their treatment values and if units with larger treatment effects tend to select a higher treatment value. Nevertheless, our treatment proxy is based on higher education endowment well before the college expansion. Therefore, prefectures in our sample do not actually select based on treatment effects. This setup alleviates our concern about the unique selection bias associated with continuous treatment variables.

2.5.4 Measurement Errors in College Types

Parallel to the regular college system in China is a postsecondary credential system (also referred to as adult higher education) that mainly serves adults older than typical college ages. The admission bar to the latter system is very low, and students are not required to regularly study on-site (Kai-Ming et al., 1999; Wang, 2011). The college expansion, as mentioned, has been concentrated on the regular college system. Unfortunately, we cannot distinguish between regular college degrees and other postsecondary credentials in the census data. To address the issue of potential measurement errors in college types, we use CFPS data to correct for the potential errors. The CFPS data set, though much smaller compared to the census, contains detailed information about respondents' college-degree types. We predict whether an individual possesses a regular college degree using basic demographic characteristics in the CFPS data with a logit model.[28] The estimated model is then applied to the census data, and we classify individuals whose predicted

[28]The predictors include interactions between (a) birth cohort and (b) the province of residence, gender, ethnicity, and urban-rural residency status.

probability of having a regular college degree is above 0.5 as regular college graduates. Details about the prediction model are provided in Appendix 2.12.

The estimated effects on the predicted regular college graduates are comparable to our baseline estimates in Columns (1) and (3) of Table 2.3. The estimates are quantitatively similar to the baseline estimates in Table 2.3 and also exhibit similar levels of statistical significance. The results show that our baseline findings are robust to considering the potential difference between regular college graduates and other postsecondary-credential holders.

Table 2.5: Results for Predicted Regular College Graduates

Dependent variable: Ever being married				
	(1)	(2)	(3)	(4)
Post-expansion cohorts	1985-88, 27-30 years old in 2015		1981-84, 31-34 years old in 2015	
Pre-expansion cohorts	1975-78, 27-30 years old in 2005		1975-78, 31-34 years old in 2009	
	Male	Female	Male	Female
ExpProxy∗Post	0.0260***	0.0174***	0.0130***	0.00658
	(0.00538)	(0.00369)	(0.00366)	(0.00494)
	[0.0273]	[0.0171]	[0.0124]	[0.0053]
Observations	29039	21703	26757	18628
Marriage rate of				
Pre-expansion cohorts	0.640	0.747	0.881	0.909
Post-expansion cohorts	0.629	0.720	0.859	0.890

Notes: All models are estimated using the sample of predicted regular college graduates. Baseline estimates are in brackets. The marital outcome of pre-expansion cohorts is constructed using marriage history so that it is comparable to the post-expansion cohorts. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit represents one standard deviation across all prefectures. All regressions control for prefecture and age fixed effects. Standard errors clustered at the prefecture level are in parentheses. There are 340 clusters. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Baseline estimates are in brackets.

2.5.5 Local Marriage Markets and Migration

We define each prefecture (the subprovincial geo-unit) as a local marriage market. Whether this definition captures actual marriage markets is crucial for interpreting our empirical results. One potential concern is that the marriage market is larger than a single prefecture, leading to mis-specification in the empirical model. We discuss this issue from two aspects.

First, there is good reason to believe that the marriage markets are largely local. The typical prefecture in China is a very large geo-unit. The average population size of a prefecture was four million in 2015.[29] Given the large size, it is probably reasonable to believe that most marriages are formed within prefectures. In a recent work, Chen et al. (2022) document that commuting-based metropolitan areas in China usually do not cross prefectural boundaries. The authors use this information to delineate China's local labor markets, which also likely reflects how localized China's marriage markets are.

Second, we relax the assumption that local marriage markets are confined to prefectures. Instead, we consider the province as the unit of analysis.[30] The treatment proxy is defined at the province level. We re-estimate the baseline DID specification (Equation 2.5) with the province-level treatment variable and province fixed effects. The results are reported in Table 2.6. As shown in Columns (1) and (3), we still find that the expansion boosted the marriage probabilities of college women and men. For example, Column (1) shows that increasing the treatment proxy by one SD across provinces leads to a 4 percentage-point increase in the marriage probability of college men and a 2.4 percentage-point increase for college women in the late post-expansion cohorts. The results are qualitatively consistent with our baseline results.[31] Therefore, our conclusions are robust to taking a province as a local marriage market.

[29]The average population size of a US state, for comparison, was 6.4 million in 2015.

[30]There are 31 provincial-level units in mainland China.

[31]The results are also quantitatively comparable when defining the marriage market at the prefecture versus province level. The SD of the prefecture-level treatment proxy is 1.57 times that of the province-level proxy. If we rescale the estimates in Table 2.6 using this factor, for example, the treatment effect on college men in Column (1) is 0.025, while that for college men in Column (1) of Table 2.3 is 0.027.

Table 2.6: College Expansion and Marriage: Province Level Marriage Markets

Dependent variable: Ever being married				
	(1)	(2)	(3)	(4)
Post-expansion cohorts	1985-88, 27-30 years old in 2015		1981-84, 31-34 years old in 2015	
Pre-expansion cohorts	1975-78, 27-30 years old in 2005		1975-78, 31-34 years old in 2009	
	College	Non-college	College	Non-college
A. Male				
ExpProxy∗Post	0.0398***	0.0050	0.0185***	-0.0020
	(0.0047)	(0.0046)	(0.0063)	(0.0030)
Observations	40198	187259	36489	181105
Marriage rate of				
Pre-expansion cohorts	0.644	0.793	0.883	0.882
Post-expansion cohorts	0.613	0.728	0.858	0.874
B. Female				
ExpProxy∗Post	0.0239***	0.0038	0.0104*	0.0005
	(0.0049)	(0.0085)	(0.0056)	(0.0012)
Observations	38478	182842	33897	176213
Marriage rate of				
Pre-expansion cohorts	0.787	0.916	0.919	0.961
Post-expansion cohorts	0.709	0.873	0.892	0.953

Notes: The marital outcome of control cohorts is constructed using marriage history so that it is comparable to the treatment cohorts. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit is one standard deviation of the treatment proxy across provinces. All regressions control for province fixed effects and age fixed effects. Standard errors clustered at the province level are in parentheses. There are 31 clusters. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A related concern is migration between local marriage markets. For two reasons, this turns out to be not as consequential as it might appear. First, while college students in one prefecture do not solely come from that prefecture, this fact does not change the interpretation of our results. Our cross-sectional variation comes from the fact that college enrollment increased much more in places with high preexisting higher education resources, regardless of where the additional enrollment came from.

Second, after they graduate, college students do not necessarily stay in the prefecture where they go to college. This might create potential measurement errors if our treatment proxy is not an accurate measure of the stock of college graduates in the local marriage market. But the mobility of college graduates after graduation seems low. For example, based on the 2015 mini-census data, the probability of staying in the same prefecture as five years ago is 90% for college women and 88% for college men aged 27–35 years old.

If the migration decision after graduation is uncorrelated with either the treatment or local marriage-market conditions, that would at most attenuate our estimates by introducing additional measurement errors. The concern would be that college graduates who have a stronger preference for marriage are more likely to move to cities that have more college graduates and can provide more abundant opportunities for marriage. But, given the low rate of migration after college graduation, this potential confounding channel is unlikely to drive our main results. The robustness of the results using provincial-level local marriage markets also alleviates this concern because cross-province migration is less common than cross-prefecture migration.[32]

[32]The probability of staying in the same province as five years ago is 95% for college women and 93% for college men aged 27–35 years old. Separately, our data set provides information on (1) whether the current prefecture of residence is the same as five years ago and (2) the province of residence five years ago. Based on this, we conduct two additional tests: (1) we re-estimate the baseline model with only individuals that resided in the same prefecture as five years ago; (2) we re-estimate the model that takes provinces as local marriage markets but take the province of residence five years ago as the current province of residence. The results are quantitatively very similar to (and qualitatively consistent with) our original results (Table 2.3 for the first test and Table 2.6 for the second test). These tests further suggest that selection into migration is not likely to bias our results.

2.5.6 Marriage Age Gap

We investigate women and men from the same cohorts in the baseline results so that the analyzed cohorts experienced the same intensity of treatment and were observed at the same ages. However, women on average marry men older than themselves. The average husband-wife age gap in China is about two years. Women and men from the same cohorts, therefore, might not be in exactly the same marriage market. The dynamics are further complicated by the fact that the intensity of the expansion increased over time. When we look at men and women from the same cohorts, the women are actually in a marriage market with older men, and the effects of the expansion on college attainment are smaller for their potential spouses.

Table 2.7: Compare Treatment Effects on Different Genders Allowing for Two-year Age Gap

Dependent variable: Ever being married				
	(1)	(2)	(3)	(4)
	College	Non-college	College	Non-college
A. Male				
Post-expansion cohorts	1985–88, 27-30 years old in 2015		1981–84, 31-34 years old in 2015	
Pre-expansion cohorts	1975–78, 27-30 years old in 2005		1975–78, 31-34 years old in 2009	
ExpProxy∗Post	0.0273***	0.0059***	0.0124***	0.0008
	(0.0049)	(0.0020)	(0.0039)	(0.0014)
Observations	40196	187259	36486	181105
Marriage rate of				
Pre-expansion cohorts	0.644	0.793	0.883	0.882
Post-expansion cohorts	0.613	0.728	0.858	0.874
B. Female				
Post-expansion cohorts	1987–90, 25-28 years old in 2015		1983–86, 29–32 years old in 2015	
Pre-expansion cohorts	1977–80, 25-28 years old in 2005		1977–80, 29–32 years old in 2009	
ExpProxy∗Post	0.0285***	0.0106**	0.0066**	-0.0011
	(0.00596)	(0.00489)	(0.0033)	(0.0026)
Observations	43630	178921	38302	169558
Marriage rate of				
Pre-expansion cohorts	0.580	0.841	0.864	0.935
Post-expansion cohorts	0.504	0.788	0.825	0.924
C. Female − Male				
	0.00125	0.00467	-0.00589	-0.00186
	(0.00774)	(0.00527)	(0.00511)	(0.00291)

Notes: The marital outcome of control cohorts is constructed using marriage history so that it is comparable to the treatment cohorts. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit represents one standard deviation across all prefectures. All regressions control for prefecture and age fixed effects. Standard errors clustered at the prefecture level are in parentheses. There are 340 clusters. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

To test the effects of the expansion for women and men in the same marriage market, we re-estimate the econometric model by comparing men from the same cohorts as in the baseline model but also comparing women two years younger than the corresponding men. The results are reported in Table 2.7.[33] Column (1) shows that when we compare college women that are two years younger than college men in the late post-expansion cohorts, the expansion has quantitatively similar effects on their marriage probabilities. If we look at college women who are two years younger than college men from the early post-expansion cohorts, as in Column (3), the impact of the college expansion is statistically significant. The gender difference in treatment effects, however, is still sizable (and nonsignificant). Taken together, these findings suggest that at least part of the gender differences observed in our baseline findings (Table 2.3) are explained by the positive average husband-wife age gap.

2.5.7 Potential for Selection and Composition Change

We focus on the equilibrium effects of the expansion on college graduates. One concern, however, is that the expansion may have shifted the composition of college students. Such an endogenous composition change could bias our results if students that enrolled in college as a result of the expansion are very different from other college students in terms of their marriage decisions. For example, if those who went to college because of the expansion were more likely to get married than other college students with or without the expansion, then we may observe a spurious positive impact of the expansion on the marriage probability of an average college graduate.[34]

[33] These results need to be interpreted with caution because age is another important trait to match on in the marriage market. Therefore, the marriage age gap is endogenous to change in marriage-market conditions. In Section 2.6.4, we show that the college expansion indeed reduced the marriage age gap. The effects, nevertheless, are small relative to the preexisting marriage age gap. Therefore, we take the two-year gap as approximately given in this exercise.

[34] It is ex ante unclear in which direction the potential bias will be. For example, our results may reflect a negative bias if the students that went to college because of the expansion came from relatively low-socioeconomic-status families and are less likely to marry. On the other hand, the bias could be positive if these students had less human capital and therefore lower opportunity costs in the labor market from marrying early.

In our model, the only trait relevant for marriage matching is education type. Marital preferences do not systematically differ for individuals with a given education type. The model, therefore, does not accommodate this potential selection story. From the model perspective, we focus on the expansion's average equilibrium effects that shifted the overall education distribution. Therefore, as long as the marginal students who enrolled as a result of the expansion were not systematically different from other college graduates in terms of marital preferences or other traits in the marriage market, our empirical results will still be largely consistent with the effects that we predict in the theoretical model.

Still, it is helpful to empirically test whether the endogenous composition change threatens our findings. One critical feature of China's college system is that a single score from the college entrance examination (gaokao) is the sole determinant of college admission.[35] Only students whose scores were above a certain threshold had the opportunity to be admitted each year. The college expansion increased the quota for college admission and therefore drew more students from the relatively lower score distribution. If our main findings are driven by a composition change, then we expect that students from the lower score distribution are also more likely to get married.

We use a unique data set, the China Household Income Project (CHIP), to conduct a test on the gaokao score and marriage probability. The CHIP is a nationally representative household survey that covers income and expenditure information of Chinese households.[36] The 2013 wave of CHIP contains information on respondents' gaokao score. Using the sample of college graduates, we estimate the correlation between individuals' score and marriage probability with the following specification:

$$Married_{iars} = \rho Score_{iqrs} + \xi_a + \phi_r + \varphi_s + \varepsilon_{iars} \tag{2.6}$$

[35] The gaokao includes multiple subjects. It takes place only once a year in each summer.

[36] The surveys were organized in multiple years by the China Institute for Income Distribution at Beijing Normal University and conducted by the National Bureau of Statistics of China.

The dependent variable is a dummy for being married for individual i of age a. r indicates the region (province) where i took the gaokao exam, and s stands for the year of the exam. *Score* is i's total exam score.[37] We also control for age, region, and gaokao-year fixed effects when applicable. The sample includes cohorts born between 1975 and 1986, which are roughly the same as the cohorts in our main analysis. These cohorts were at least 27 years old in 2013, so their marriage outcomes are largely comparable to those in our baseline analysis.

The results are shown in Table 2.8. Columns (1) and (2) show that there is no significant association between the gaokao score and marriage probability. The estimates are very imprecise as a result of the small sample size, but the magnitudes are small. For example, Column (1) suggests that increasing the gaokao score by 100 is only associated with a 0.8 percentage-point higher marriage probability for college men and 0.3 percentage-point higher probability for college women.[38]

We further look into cohorts that were and were not exposed to the expansion. If our results are driven by selection bias, then we should see a more negative association between the score and marriage probability for the post-expansion cohorts. As shown in Columns (3) and (4), we actually observe a negative association for the pre-expansion cohorts and a positive association for the post-expansion cohorts. Based on the point estimates, it seems that the composition-change bias, if anything, goes in favor of our main findings and suggests that our baseline estimates provides a lower-bound of the expansion's positive effect on marriage probabilities. None of the estimates are statistically significant, so they should be interpreted with caution. Nevertheless, these patterns suggest that the composition change is not likely to be driving our results.

[37] The score ranges from 0 to 750, and we divide the score by 100 in the regressions.

[38] One hundred points represent a giant increase in the test score and can catapult a student to a much higher tier of universities.

Table 2.8: Correlation between College-Entrance-Exam Score and Marriage Probability

Dependent variable: Ever being married				
	(1)	(2)	(3)	(4)
A. Male				
Cohorts	1975-86		Pre-expansion	Post-expansion
$\frac{Score}{100}$	0.0082	0.0032	-0.0175	0.0155
	(0.0143)	(0.0151)	(0.0150)	(0.0257)
Observations	494	468	219	249
Age FE	Yes	Yes	Yes	Yes
Exam Year FE		Yes	Yes	Yes
Province FE		Yes	Yes	Yes
B. Female				
Cohorts	1975-86		Pre-expansion	Post-expansion
$\frac{Score}{100}$	0.0027	-0.0019	-0.0105	0.0098
	(0.0189)	(0.0165)	(0.0185)	(0.0252)
Observations	490	467	193	274
Age FE	Yes	Yes	Yes	Yes
Exam Year FE		Yes	Yes	Yes
Province FE		Yes	Yes	Yes

Notes: The sample consists of college graduates in the 2013 wave of China Household Income Project. The maximum of the original exam score is 750. Robust standard errors are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

2.5.8 Other Robustness Tests

In this section, we report several additional tests that further confirm the robustness of our findings. Our results are robust to looking at four-year or three-year institutions, using an alternative proxy for college-expansion treatment, and controlling for prefecture-level covariates.

Four-Year versus Three-Year Institutions

There are two types of college in China: four-year universities (benke) and three-year colleges (zhuanke). The four-year universities have higher quality and higher admission standards, and both types of college degree bring very high wage returns compared to lower education levels (Zhong, 2011). Both four-year and three-year institutions experienced the expansion. In our baseline analysis, we focus on the whole higher education system and therefore do not distinguish between these two levels of higher education. It might be informative to explore whether the effects of the expansion differ between them. Table 2.19 reports the effects of the expansion on four-year and three-year institutions' graduates, respectively. While the expansion increased the marriage probabilities of both groups, the effects seem to be larger for four-year university graduates.

Using 1990 Data to Construct the Expansion Proxy

In the baseline analysis, we use the historical endowment of higher education resources as of 1982 as a proxy for treatment intensity. Because variation in this measure was mostly predetermined during the pre-expansion era, this proxy is unlikely to be confounded by the growth of higher education that is associated with subsequent economic development. This benefit, however, comes with a potential cost: the proxy might be less strongly correlated with the actual intensity of the expansion than a proxy based on variation closer to the expansion period. As a robustness check, we construct an analogous treatment proxy using data from the 1990 Chinese census and redo the analysis. The results, reported in Table 2.20, are qualitatively similar to our baseline findings.

Additional Covariates

Our identification assumption requires that the marriage outcomes follow parallel trends in the absence of the expansion. The parallel-trends assumption could be violated if our treatment proxy is associated with other key factors that influence the marriage market and these factors cause different secular trends. We have shown, in Appendix 2.11, that our treatment variable is not associated with GDP-per-capita growth or local sex ratio. We further control for baseline prefecture characteristics interacted with a time dummy. In the first exercise, we control for GDP per capita and the sex ratio of each prefecture in 2000.[39] The results are reported in Table 2.21. In the second exercise, we include province-by-time fixed effects, which control for regional differences in economic growth and marriage-market conditions at the province level and exploit only within-province variation to identify the effects of the expansion on local marriage markets. The results are reported in Table 2.22. In both tests, our main findings are robust, both qualitatively and quantitatively.

2.6 How the Expansion Affected Matching Patterns

2.6.1 Effects on Marriage-Matching Patterns

In this section, we further assess the implications of the college expansion for marital sorting. It is perhaps unsurprising that there are more college-college marriages as a result of the increasing supply of college graduates in the marriage market. It is unclear, however, whether this reflects a change in underlying marital sorting—that is, whether educated women and men are more likely to marry each other conditional on the marginal distribution of education types. Research has shown that increasing positive assortative mating by education could amplify household income inequality (Greenwood et al., 2014; Eika et al., 2019). It is therefore important to understand how

[39] We calculate the sex ratio of each prefecture in the age range of 10–30 years old using the 2000 China census data from IPUMS. Data on GDP per capita come from the City Statistics Year Book of China.

the underlying sorting patterns in China have changed following the expansion.

We first directly estimate the impacts of the expansion on matching patterns by focusing on the proportions of different types of match. We focus on the simple dichotomy between college and noncollege graduates and estimate the DID model in Equation 2.7 with married couples:

$$z_{ijpt} = \beta_3 + \beta_4 ExpProxy_p * Post_t + \beta_5 * Post_t + \lambda_p + \varepsilon_{ijpt} \tag{2.7}$$

The dependent variable is a series of dummies for different match types between wife i and husband j in local marriage market p at time t. These dummies include whether the marriage consists of (1) a college wife and a college husband, (2) a college wife and a noncollege husband, (3) a noncollege wife and a college husband, and (4) a noncollege wife and a noncollege husband. The parameter of interest is the coefficient on the interaction between the expansion proxy and the dummy for post-expansion cohorts ($Post_t$).[40] We also include local-marriage-market fixed effects λ_p.[41]

The results, reported in Table 2.9, confirm that the college expansion has led to more college-college marriages.[42] For example, a one-SD increase in the treatment proxy caused a 2.1 percentage-point increase in the probability of college-college marriages for the early post-expansion cohorts. Comparing the national averages of different marriage types in the pre-expansion and post-expansion cohorts also leads to similar conclusions. These patterns, however, could partially be mechanical results of more college women and men in the marriage market. In order to distinguish the underlying changes in the matching structure from the mechanical effects, we have to adjust

[40] We adjust the age of couples in the pre-expansion cohorts to make them comparable to the post-expansion cohorts. For example, when we choose the 1981–84 cohorts (who were 31–34 years old when observed in 2015) as the post-expansion cohorts, couples in the pre-expansion cohorts (1975–78) only include those who got married by 2009 (when they were at the same age). The estimates are similar and do not qualitatively alter our conclusion if we do not make such age adjustment.

[41] Controlling for age fixed effects of the husband and wife does not alter the estimates by much and does not affect our qualitative conclusion.

[42] By construction, the DID coefficients of the four columns in the same row add up to zero. In Panel A, the coefficients in Column (1) (Column (2)) and Column (4) (Column (3)) appear exactly opposite due to rounding errors.

for the changes in the marginal distributions. This is the goal of the next subsection.

Table 2.9: Effects of the College Expansion on Matching Patterns

	(1)	(2)	(3)	(4)
Dependent variable (Wife-Husband)	C-C	NC-C	C-NC	NC-NC
A. Early Post-expansion Cohorts				
Post-expansion cohorts	1981–84, 31–34 years old in 2015			
Pre-expansion cohorts	1975–78, 31–34 years old in 2009			
ExpProxy$*$Post	0.0210***	0.00165	-0.00165**	-0.0210***
	(0.00537)	(0.00101)	(0.000826)	(0.00609)
Observations	224384	224384	224384	224384
Dep. var. mean of				
Pre-expansion cohorts	0.097	0.043	0.023	0.836
Post-expansion cohorts	0.185	0.048	0.036	0.731
B. Late Post-expansion Cohorts				
Post-expansion cohorts	1985–88, 27–30 years old in 2015			
Pre-expansion cohorts	1975–78, 27–30 years old in 2005			
ExpProxy$*$Post	0.0204***	0.00232***	-0.00262***	-0.0201***
	(0.00421)	(0.000826)	(0.000671)	(0.00441)
Observations	202855	202855	202855	202855
Dep. var. mean of				
Pre-expansion cohorts	0.083	0.040	0.020	0.856
Post-expansion cohorts	0.172	0.050	0.039	0.738

Notes: Each observation is a married couple. Dependent variables are dummies for corresponding matching type. C: college. NC: non-college. For example, C-NC refers to a dummy for the wife having a college degree while the husband does not. The marital outcome of pre-expansion cohorts is constructed using marriage history so that it is comparable to the post-expansion cohorts. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit is equivalent to one standard deviation across all prefectures. All regressions control for prefecture fixed effects. Standard errors clustered at the prefecture level are in parentheses. There are 340 clusters. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

2.6.2 Assortative Mating

In order to address the mechanical effects of a larger college-educated population, we construct a new measure for assortative mating: the difference between the actual probability of college-college matching and the probability of college-college matching under the counterfactual of random matching. This measure, which we refer to as the absolute-difference measure, accounts for the fact that there would be more college-college marriages following the college expansion even under random matching. Formally, for the realized matching outcomes in a marriage market (as characterized by a contingency table, Table 2.10), the index is defined as follows:

Table 2.10: Contingency Table

		Education of Husband	
		C	NC
Education	C	k_1	k_2
of Wife	NC	k_3	k_4

Notes: The number in each cell k_j is the number of marriages of the corresponding type. $K = k_1 + k_2 + k_3 + k_4$.

$$\begin{aligned} AbsDiff &= \frac{k_1}{K} - \frac{k_1+k_2}{K} * \frac{k_1+k_3}{K} \\ &= \frac{k_1k_4 - k_2k_3}{K^2} \end{aligned}$$

In the hypothetical situation in which marriage matching by education is totally random, this index will always be zero. An increase in this index implies that college-college marriages are now more likely relative to the benchmark of random matching. An increase also implies a potential increase in income inequality because highly educated individuals are more likely to sort into matches with other highly educated individuals.

To estimate the causal effects of the expansion on assortative mating, we again employ a DID design. We divide the national sample into high- and low-expansion regions based on the value of the expansion proxy. For the pre-expansion and post-expansion cohorts, we estimate the matching indexes in the high- and low-expansion regions, respectively. The estimated assortative-mating indexes are then used for constructing the DID estimate. We also adjust the pre-expansion cohorts so that their age range when their marriage outcomes were observed is comparable to the age range of the post-expansion cohorts, as in Section 2.6.1.[43]

Table 2.11 reports estimates for the effects of the expansion on assortative mating by college education. The expansion increased the level of assortative mating. If we look at the late post-expansion cohorts, the expansion increased the probability of college-college marriages relative to random matching by 4.1 percentage points in the high-expansion regions compared to the low-expansion regions. The magnitude of this effect seems large: it is half of the pre-expansion assortative-mating index in the high-expansion regions and more than 80% of the pre-expansion index in the low-expansion regions. We find quantitatively similar effects for the early post-expansion cohorts. The results imply that the expansion potentially increased income inequality across households by increasing assortative mating by college-educated individuals.

Compared to graduates of three-year colleges, graduates of four-year universities on average have more human capital and higher earnings. In Table 2.23, we estimate how the expansion changed the level of assortative mating by four-year university degrees versus lower education levels. We again find that the expansion had a positive and non-negligible impact on assortative mating.

[43] If we do not make the age adjustment, the results are similar and do not qualitatively change our conclusion.

Table 2.11: Effects of the College Expansion on Assortative Mating: Absolute-Difference Index

		Late Post-expansion Cohorts: 1985-88		Early Post-expansion Cohorts: 1981-84	
	Region by Expansion Intensity	High	Low	High	Low
Cohorts	Pre-expansion: 1975-78	0.082	0.050	0.095	0.054
	Post-expansion	0.149	0.075	0.158	0.075
	Diff-in-diff	**0.041*****		**0.042*****	
		(0.002)		(0.002)	

Notes: The index is calculated using the sample of married couples in the random samples of 2010 census and 2015 mini-census. The sample includes married couples with either side falling in the specified cohort range. Ages of couples in pre-expansion cohorts are adjusted to be comparable to those in the post-expansion cohorts. The national sample is divided into high vs. low regions based on whether the treatment proxy is above or below median. Standard errors in parentheses are estimated by boostrapping from the original sample 1,000 times. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

2.6.3 Comparison with Previous Measures of Assortative Matching

Many methods have been developed to measure assortative matching. The goal of most previous studies is to compare the underlying tendencies for assortative mating (by education) of two marriage markets (characterized by two contingency tables) that have different marginal distributions of education types. Put another way, the indexes are designed to fully adjust for the difference in marginal distributions. That goal slightly differs from ours because we are curious about the impact of the expansion, which itself is a change in marginal distributions of education, on matching patterns and its implications for inequality. Put another way, we are investigating the effects of changing marginal distributions of education types rather than trying to fully adjust for it. Therefore, the existing measures do not perfectly fit our purpose.

To illustrate the point, we compare our index with several measures that have been widely used and examined: log odds ratio (Siow, 2015; Chiappori et al., 2020; Ciscato and Weber, 2020), minimum distance (Liu and Lu, 2006; Fernández and Rogerson, 2001), and rank correlation. Given the matching outcomes characterized by Table 2.10, the definitions of these indexes are provided in Table 2.12.

Table 2.12: Compare Different Matching Indexes

	Definition	Range	Scale invariance	Symmetry	Monotonicity	Perfect PAM
Log Odds Ratio	$\ln \frac{k_1 k_4}{k_2 k_3}$	$[0, \infty)$	Y	Y	Y	Y
Rank Correlation	$\frac{k_1 k_4 - k_2 k_3}{\sqrt{(k_1+k_2)(k_3+k_4)(k_1+k_3)(k_2+k_4)}}$	$[0, 1]$	Y	Y	Y	Y
Minimum Distance	$\frac{k_1 k_4 - k_2 k_3}{(k_1+min\{k_2,k_3\})*(k_4+min\{k_2,k_3\})}$	$[0, 1]$	Y	Y	Y	Y
Absolute Difference	$\frac{k_1 k_4 - k_2 k_3}{K^2}$	$[0, 0.25]$	Y	Y	Y	N

Notes: $K = k_1 + k_2 + k_3 + k_4$.

Chiappori et al. (2021) assess indexes of assortative mating used in the literature and propose that a satisfactory matching index should have the following properties. First, it should be invariant to the scale of the population. Second, it should be symmetric between the categories (in our case, college versus noncollege). Third, the monotonicity condition requires that it should increase with more people in diagonal cells (k_1 and k_4) when the marginal distributions are held constant. Fourth, the perfect-PAM (positive assortative mating) condition requires that a contingency table under perfectly positive assortative mating ($k_2 = k_3 = 0$, no off-diagonal matches) should exhibit the maximal value of the index. We compare the various indexes' properties in Table 2.12. In Appendix 2.14, we discuss the properties of these indexes.

Our index (absolute difference) fails only the perfect-PAM condition. We argue that this failure is not a fatal threat to our index. To see why, consider the hypothetical situation of perfect assortative mating. A measure that satisfies the perfect-PAM condition should always achieve the maximum value no matter how the marginal distribution of education changes. However, if access to higher education was initially very low but then expanded to a much higher level, such an expansion would still contribute to inequality, which is the outcome of the joint forces of a high level of assortative mating and the higher education expansion. Therefore, we choose the absolute-difference index as the measure for the level of assortative mating.

Using the alternative indexes, we also explore whether the expansion affected the underlying assortative-mating patterns. The results are reported in Table 2.24. For none of the three alternative indexes are there any significant estimates for differential trends in the high-expansion versus low-expansion regions after the expansion. The results suggest that the expansion did not significantly change the underlying assortative-matching tendency, although it led to more college-college marriages relative to the benchmark of random matching.

2.6.4 Marriage Age Gap

In addition to education, age is another important trait for marriage matching. One often-used measure for matching patterns regarding age is the age gap between the husband and the wife. Our model does not incorporate matching on age and therefore does not provide theoretical guidance on how to think of the effects of the college expansion on matching by age. Still, investigating the expansion's impact on matching patterns by age helps us to understand how the expansion changed the marriage market. It may also provide insights about the channels for our main finding about the positive effects of the expansion on the marriage probability of college graduates: if the positive effects are driven by the LCMM, then we should probably expect that the expansion shrank the marriage age gap of college graduates because they increasingly met people from similar backgrounds at college.

In this section, we estimate the effects of the expansion on the marriage age gap using the DID specification in Equation 2.5. Table 2.13 reports the results for college women and men. Our estimates suggest that the expansion had a small but statistically significant negative impact on the marriage age gap. For example, Column (1) shows that a one-SD increase in the treatment proxy led to a 0.03-year (0.36 month) drop in the marriage age gap for college women in the late post-expansion cohorts. The results support the story that the expansion reduced search frictions in the LCMM. The expansion also had a negative effect on the marriage age gap of the early post-expansion cohorts: in Column (3), a one-SD increase in the treatment proxy decreased the age gap for college women by 0.07 years (0.84 months).[44]

In Table 2.25, we report the estimated effects of the expansion on the marriage age gap of noncollege groups. We again find a negative effect, implying that the expansion had spillover effects on the matching patterns of noncollege individuals beyond the (mostly) null effects on their marriage probabilities. We discuss the results and potential explanations in Appendix 2.13.4.

Table 2.13: Effects of the College Expansion on the Marriage Age Gap of College Graduates

	(1)	(2)	(3)	(4)
Dependent variable: Age gap (Husband - Wife)				
Post-expansion cohorts	1985-88, 27-30 years old in 2015		1981-84, 31-34 years old in 2015	
Pre-expansion cohorts	1975-78, 27-30 years old in 2005		1975-78, 31-34 years old in 2009	
	Female	Male	Female	Male
ExpProxy$*$Post	-0.0284**	-0.0309*	-0.0732***	-0.0866***
	(0.0142)	(0.0171)	(0.0172)	(0.0262)
Observations	20724	18683	22987	24347
Average age gap of				
Pre-expansion cohorts	2.104	1.331	1.995	1.839
Post-expansion cohorts	1.737	0.760	1.733	1.369

Notes: The marital outcome of pre-expansion cohorts is constructed using marriage history so that it is comparable to the post-expansion cohorts. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit represents one standard deviation across all prefectures. All regressions control for prefecture fixed effects and age fixed effects. Standard errors clustered at the prefecture level are in parentheses. There are 340 clusters. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

[44] The effect on age gap seems to be larger for the early post-expansion cohorts than for the late post-expansion cohorts. Nevertheless, the effect sizes are very small, so the difference between the late and early post-expansion cohorts may not be economically meaningful.

2.7 Conclusion

China's college expansion that began in 1999 has provided access to higher education for millions of young women and men. This chapter has studied the impacts of that radical education reform on the marriage market. We combined a theoretical model with empirical analysis to uncover equilibrium effects of the expansion and understand the underlying mechanisms.

We first developed a marriage-matching model with educational investment and search frictions in the marriage market. One key assumption is that expanding access to higher education reduces average search frictions in the college marriage market. The main forces at work in the model are changes the relative distributions of different education types in the marriage market and the reduction in search costs in the LCMM due to the expansion. Our model suggests the expansion has had important general equilibrium effects on the marriage market. When the LCMM channel dominates the effects of changing the relative distributions of education types, the expansion raises the marriage probabilities of both college women and men.

Exploiting regional and cross-cohort variation in exposure to the expansion, we empirically estimated causal impacts of the expansion on the marriage market. We showed that the expansion indeed meaningfully increased the marriage probabilities of college women and men. We also found effects of the expansion on marriage-matching patterns. The expansion increased the level of assortative mating by education, as measured by the difference between the probability of actual college-college marriages and the hypothetical probability of such marriages under random matching. It also reduced the marriage age gap. Our empirical findings have important implications for critical issues in contemporary China such as the so-called "leftover women" phenomenon and income inequality.

This chapter adds to our understanding of the equilibrium effects of education institutions and education reforms on marriage outcomes. The findings of this chapter are potentially important for policy makers when considering the impacts of education policies on critical lifetime outcomes, and are especially relevant for countries that are expanding or will expand access to higher education.

2.8 Appendix A: More Details about The Higher Education System in China and the College Expansion

2.8.1 College Types

There are different types of colleges in China. We focus on regular colleges in this paper. Regular colleges admit students through the national college entrance examination (a.k.a "Gaokao"), and the exam score is the sole determinant of which colleges the students are admitted into. Students are required to be fulltime and on-site. There are two types of regular college education. One is four-year universities (*Benke*), the other is junior colleges (*Zhuanke*), which usually requires 3 years to accomplish. Four-year universities are better funded by the central and the local government, have more teachers, and are considered as having higher quality and prestige. Four-year universities can be regarded as Tier 1 and junior colleges can be regarded as Tier 2. Students admitted into four-year universities have much higher scores. Most students take the Exam right after high school, around the age of 18.

There is also a post-secondary credential system (sometimes referred to as *adult higher education*) in China. The major difference of this system is that students do NOT need to take the same college entrance exam as the regular college students and the admission bar is close to open-enrollment. Students are not required to study regularly on-site. Instead, they may attend classes on a part-time basis (e.g. at night or over weekend), and they may also study remotely (e.g. online). The degrees obtained are different from those of regular colleges, and the requirements for graduation are also much lower compared to those of regular colleges. Unsurprisingly, the degrees are not rewarded as much as regular college degrees in the labor market (and possibly also in the marriage market).

One related issue in the China census data (and most household surveys) is that when respondents report their educational attainment, they only report whether they have a college degree, and whether it is *Benke* or *Zhuanke*, but do not distinguish between regular and special college degrees. The only census year that this information is collected in 2000. CFPS data collects this information. We rely on CFPS data to show the relative importance of different colleges in China's higher education system, and over the period of college expansion.

2.8.2 Dropout in Chinese Universities

Unlike in some other contexts (e.g. the US), the college dropout rate is extremely low in China. For example, Marioulas (2017) states that "China has one of the lowest college dropout rates in the world, with sources from the ministry of education, who state that less than 1% of students fail to complete their degrees." Using confidential data from a random sample of 646 universities, Wu et al. (2016) documented that the average graduation rate was 96.91% and only 2.63% among these universities had a graduation rate that was below 90%. Moreover, the graduation rate did not systematically differ by college quality. For example, the graduation rate in Project 985 universities (the top 39 universities in China) was 95.51%, only slightly lower than the average value.

To check if the pattern of dropout changed after the college expansion, we plot in Figure 2.13 the ratio of graduates to all exits from higher education institutions in each year using data from the Chinese Education Yearbooks.[45] Ideally, we would like to calculate the rate of graduation for students that newly enrolled in a given year. However, the Yearbooks do not distinguish between four-year universities and three-year colleges when reporting the total number of graduates each year. Nevertheless, Figure 2.13 shows that most college students successfully graduated and that pattern did not change following the college expansion.

[45] Other potential reasons for exiting include completion without a degree, being suspended, being expelled, quitting, and death.

Figure 2.13: The Ratio of Graduates to All Exits from Higher Education Institutions

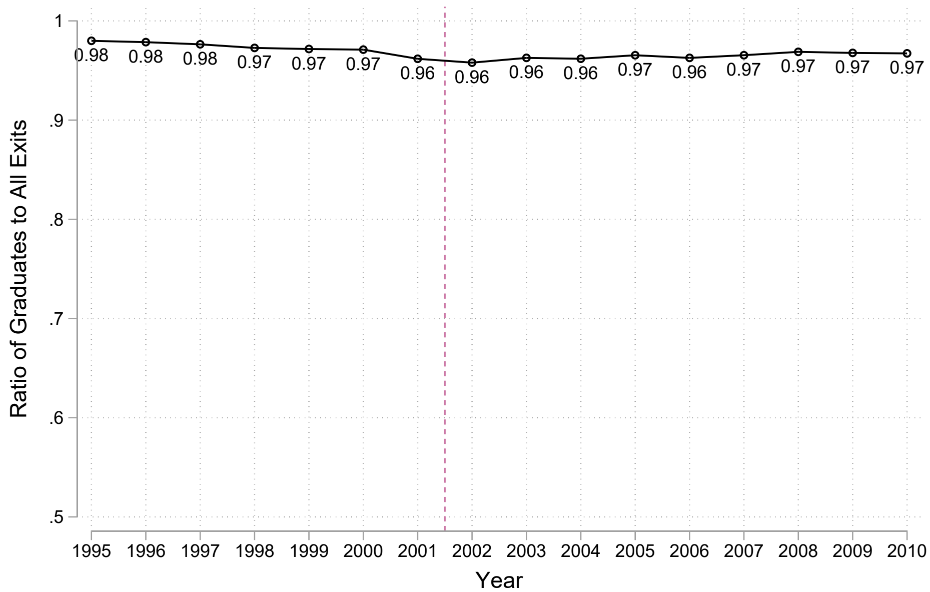

Source: The Chinese Education Yearbooks. This figure reports the ratio of graduates to all exits from higher education institutions in each year. Both four-year universities and three-year colleges are included. The college expansion started in 1999. Three-year-college students who were newly enrolled in 1999 normally graduated in 2002 (as indicated by the vertical line), and four-year-university students in 2003.

2.9 Appendix B: College Local Marriage Markets: Suggestive Evidence from the CFPS

We investigate the role of colleges as local marriage markets using the CFPS data. In the CFPS, married respondents were asked one question about how they met their (first and current) spouse. Among other options, they were asked whether they met their spouse (both current and first spouse) "at school" by themselves. The specific question is as follows:

How did you get to know your current/first spouse? [Select only one response]

1. Knew each other at school by ourselves

2. Knew each other at workplace by ourselves

3. Knew each other at place of residence by ourselves

4. Met each other at other places by ourselves

5. Introduced by friends/relatives

6. Through marriage agency

7. Arranged by parents

8. Through the Internet

77. Other [Please specify]

We take the answer of "meeting spouse at school" as a proxy for how school/college serves as a local marriage market.[46] Using the answers to this information, we define a dummy variable for married individuals meeting their spouse on campus, which equals to 1 if the answer to the question is "met her/him at school by myself". For each cohort, we can then look at what how many marriages arose from acquaintance in college.

In Table 2.1, we have shown the fractions of people who met spouses in school by college education and before versus after the expansion. One potential concern is that the observed couples in the post-expansion cohorts were younger than the pre-expansion cohorts. It could be that mechanically younger couples are more likely those who meet on campus. To adjust for this potential bias, we further restrict the sample to those who got married by the age of 27 (i.e. those who married early) in both the pre- and the post-expansion cohorts. The results are reported in Table 2.14. After adjusting for early marriages, we still observe similar patterns: college graduates are much more likely to meet their spouses in school and the fraction increased after the college expansion.

[46]One advantage of this proxy is that we directly observe how individuals met each other, while previous literature use only their common experience to gauge the role of common school/workplace in marriage matching.

Table 2.14: Fraction of People Who Met Their Spouses in School: Early Marriages

	Cohorts	Fraction	Observations
College	1975-80 (Pre-expansion)	0.25	183
	1981-88 (Post-expansion)	0.32	643
	Difference	0.07 ($p < 0.1$)	
Non-college	1975-80 (Pre-expansion)	0.04	2449
	1981-88 (Post-expansion)	0.06	4038
	Difference	0.02 ($p < 0.01$)	

Source: China Family Panel Studies 2010-2018. All results weighted using the CFPS survey weights. The sample includes only individuals who were married by 27.

2.9.1 Suggestive Causal Evidence

We show that the probability of meeting one's spouse in college has increased after the college expansion. However, this before-after comparison could be contaminated by other secular trends. We provide some suggestive causal evidence in this section by exploiting a DID design analogous to our baseline econometric model. We know individuals' provinces of residence. Therefore, we estimate the following DID model:

$$meetschool_{ipb} = \rho_1 ExpProxy_p * Post_b + \rho_2 * Male_i + \lambda_p + \xi_b + \varepsilon_{ipb} \tag{2.8}$$

The dependent variable is a dummy variable for meeting one's spouse in school for individual i of cohort of birth b in province (local marriage market) p. $ExpProxy_p$ is the same as the expansion proxy used in our baseline model but measured at the province level (Section 2.5.5). $Post_b$ is a dummy for the post-expansion cohorts (1981–88).[47] We further control for gender, cohort fixed

[47] For the DID model, we include pre-expansion cohorts in 1975–78 and post expansion cohorts in 1981–88 that are

effects, and province fixed effects. The results are reported in Table 2.15. Column (1) shows the results for college graduates, which suggest the college expansion indeed increased the probability of meeting one's spouse in school. We do not find effects of the college expansion on non-college individuals, as reported in Column (2).

Table 2.15: Effects of the College Expansion on Meeting Spouses in School

	(1)	(2)
	College	Non-college
ExpProxy$*$Post	0.0281**	-0.0176
	(0.0119)	(0.0151)
Bootstrapped p-value	[0.154]	[0.375]
N	1690	7658

Dependent variable is a dummy for meeting one's spouse in school. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit represents one standard deviation across all provinces. Standard errors clustered at the province level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Boostrapped p-values are reported in brackets. Results are weighted using the CFPS survey weights.

2.10 Appendix C: Appendix to the Model and Simulation

2.10.1 Characterization of the Full Model Equilibrium

Educational Choices

In this subsection, we characterize the full model including educational choices and the marriage matching process. The full equilibrium of the model is determined by the fixed point of educational choices and marriage market matching outcomes. To capture this, let's first reconsider the first step of the model: educational choices. Recall that the costs of college education are

consistent with our baseline econometric model.

assumed to follow:

$$c_i^f = c_f + \theta_i^f$$
$$c_j^m = c_m + \theta_j^m.$$

Consider the choice problem of female i (we use i for individual woman and j for individual man). If she chooses not to go to college, then the expected marital payoff conditional on education type L is U^L, which is determined as below (E is the Euler's constant ≈ 0.577):[48]

$$\begin{aligned} U^L =& Emax\{u_{iLH}, u_{iLL}, u_{iL0}\} \\ =& E(u_{iLy}|y = arg \max_{y=0,H,L} u_{iLy}) \\ =& E + \ln(\sum_{y=0,H,L} \exp(\alpha_{Ly} - \delta + \tau_{Ly})). \end{aligned} \tag{2.9}$$

The equation is derived based on the type-I extreme value distribution of the idiosyncratic marital preferences ε_{iy} (Choo and Siow, 2006).

Similarly, denote her expected marital payoff (without paying the education cost) conditional on choosing education type H is U^H. U^H follows a similar structure but also depends on whether she randomly enters the LCMM or not.

$$\begin{aligned} U^H =& \sum_{k=1,2} p^f(H_k) E(u_{iH_ky}|y = arg \max_{y=0,H_k,L} u_{iH_ky}) \\ =& p(H_1)[E + \ln(\exp(\alpha_{HH} + \tau_{HH}^1) + \sum_{y=0,L} \exp(\alpha_{Hy} - \delta + \tau_{Hy}))] \\ &+ (1 - p(H_1))[E + \ln(\exp(\alpha_{HH} - \delta + \tau_{HH}^2) + \sum_{y=0,L} \exp(\alpha_{Hy} - \delta + \tau_{Hy}))]. \end{aligned} \tag{2.10}$$

$p^f(H_1) = \frac{R(H_f, H_m)}{H_f}$ is the probability of entering the LCMM.

[48]In the choice of an L type, she does not distinguish between H_1 and H_2 because these two are identical in her preferences.

Woman i compares the expected payoff of college education ($U^H - c_i^f$) and that of not choosing college education (U^L). Conditional on the expected payoffs U^H and U^L, the distribution of c_i^f determines the distribution of education. The college ratio of women is determined by

$$h_f = G^f(U^H - U^L - c_f). \tag{2.11}$$

Similarly, for the education choice problem of men, V^L is the expected martial payoff conditional on being the noncollege type, and V^H is the expected martial payoff conditional on college education (without paying the education cost), we have:

$$\begin{aligned} V^L =& E(v_{xLj} | x = arg \max_{x=0,H,L} v_{xLj}) \\ =& E + \ln(\sum_{x=0,H,L} \exp(\gamma_{xL} - \delta - \tau_{xL})). \end{aligned} \tag{2.12}$$

$$\begin{aligned} V^H &= \sum_{k=1,2} p^m(H_k) E(v_{xH_kj} | x = arg \max_{x=0,H_k,L} v_{xH_kj}) \\ &= p^m(H_1)[E + \ln(\exp(\gamma_{HH} - \tau_{HH}^1) + \sum_{x=0,L} \exp(\gamma_{xH} - \delta - \tau_{xH}))] \\ &+ (1 - p^m(H_1))[E + \ln(\exp(\gamma_{HH} - \delta - \tau_{HH}^2) + \sum_{x=0,L} \exp(\gamma_{xH} - \delta - \tau_{xH}))]. \end{aligned} \tag{2.13}$$

$$h_m = G^m(V^H - V^L - c_m). \tag{2.14}$$

Marriage Market Choices

Conditional on educational choices and educational distribution, men and women make their choices of desired partners. The choices are made given the systematic marital returns, marital transfer (utility price) τ, and individual taste shocks (which follow type-I extreme value distribution). This can be transformed into a discrete choice problem, where each individual chooses their preferred spouse education type (including staying single). Assume the measure of women (men) belonging to type x (y) is μ_f^x (μ_m^y), and among them, μ_f^{xy} (μ_m^{xy}) choose the type y (x) spouse. $y = 0$ ($x = 0$) denotes staying single. This discrete choice problem leads to the following conditions for individual choices (McFadden et al., 1973).

For women type H_1 ($\mu_f^{H_1} = \mu_f^{H_1H_1} + \mu_f^{H_1L} + \mu_f^{H_10}$):

$$\frac{\mu_f^{H_1H_1}}{\mu_f^{H_1}} = \frac{\exp(\alpha_{HH} + \tau_{HH}^1)}{\exp(\alpha_{HH} + \tau_{HH}^1) + \exp(\alpha_{HL} + \tau_{HL} - \delta) + \exp(\alpha_{H0})} \tag{2.15}$$

$$\frac{\mu_f^{H_1L}}{\mu_f^{H_1}} = \frac{\exp(\alpha_{HL} + \tau_{HL} - \delta)}{\exp(\alpha_{HH} + \tau_{HH}^1) + \exp(\alpha_{HL} + \tau_{HL} - \delta) + \exp(\alpha_{H0})} \tag{2.16}$$

$$\frac{\mu_f^{H_10}}{\mu_f^{H_1}} = \frac{\exp(\alpha_{H0})}{\exp(\alpha_{HH} + \tau_{HH}^1) + \exp(\alpha_{HL} + \tau_{HL} - \delta) + \exp(\alpha_{H0})} \tag{2.17}$$

For women type H_2 ($\mu_f^{H_2} = \mu_f^{H_2H_2} + \mu_f^{H_2L} + \mu_f^{H_20}$):

$$\frac{\mu_f^{H_2H_2}}{\mu_f^{H_2}} = \frac{\exp(\alpha_{HH} + \tau_{HH}^2 - \delta)}{\exp(\alpha_{HH} + \tau_{HH}^2 - \delta) + \exp(\alpha_{HL} + \tau_{HL} - \delta) + \exp(\alpha_{H0})} \tag{2.18}$$

$$\frac{\mu_f^{H_2L}}{\mu_f^{H_2}} = \frac{\exp(\alpha_{HL} + \tau_{HL} - \delta)}{\exp(\alpha_{HH} + \tau_{HH}^2 - \delta) + \exp(\alpha_{HL} + \tau_{HL} - \delta) + \exp(\alpha_{H0})} \tag{2.19}$$

$$\frac{\mu_f^{H_20}}{\mu_f^{H_2}} = \frac{\exp(\alpha_{H0})}{\exp(\alpha_{HH} + \tau_{HH}^2 - \delta) + \exp(\alpha_{HL} + \tau_{HL} - \delta) + \exp(\alpha_{H0})} \tag{2.20}$$

For women type L ($\mu_f^L = \mu_f^{LH} + \mu_f^{LL} + \mu_f^{L0}$):

$$\frac{\mu_f^{LH}}{\mu_f^L} = \frac{\exp(\alpha_{LH} + \tau_{LH} - \delta)}{\exp(\alpha_{LH} + \tau_{LH} - \delta) + \exp(\alpha_{LL} + \tau_{LL} - \delta) + \exp(\alpha_{L0})} \tag{2.21}$$

$$\frac{\mu_f^{LL}}{\mu_f^L} = \frac{\exp(\alpha_{LL} + \tau_{LL} - \delta)}{\exp(\alpha_{LH} + \tau_{LH} - \delta) + \exp(\alpha_{LL} + \tau_{LL} - \delta) + \exp(\alpha_{L0})} \tag{2.22}$$

$$\frac{\mu_f^{L0}}{\mu_f^L} = \frac{\exp(\alpha_{L0})}{\exp(\alpha_{LH} + \tau_{LH} - \delta) + \exp(\alpha_{LL} + \tau_{LL} - \delta) + \exp(\alpha_{L0})} \tag{2.23}$$

Similarly, we have 9 equations governing the choices of men.

For men type H_1 ($\mu_m^{H_1} = \mu_m^{H_1H_1} + \mu_m^{LH_1} + \mu_m^{0H_1}$):

$$\frac{\mu_m^{H_1H_1}}{\mu_m^{H_1}} = \frac{\exp(\gamma_{HH} - \tau_{HH}^1)}{\exp(\gamma_{HH} - \tau_{HH}^1) + \exp(\gamma_{LH} - \tau_{LH} - \delta) + \exp(\gamma_{0H})} \tag{2.24}$$

$$\frac{\mu_m^{LH_1}}{\mu_m^{H_1}} = \frac{\exp(\gamma_{LH} - \tau_{LH} - \delta)}{\exp(\gamma_{HH} - \tau_{HH}^1) + \exp(\gamma_{LH} - \tau_{LH} - \delta) + \exp(\gamma_{0H})} \tag{2.25}$$

$$\frac{\mu_f^{0H_1}}{\mu_m^{H_1}} = \frac{\exp(\gamma_{0H})}{\exp(\gamma_{HH} - \tau_{HH}^1) + \exp(\gamma_{LH} - \tau_{LH} - \delta) + \exp(\gamma_{0H})} \tag{2.26}$$

For men type H_2 ($\mu_m^{H_2} = \mu_m^{H_2H_2} + \mu_m^{LH_2} + \mu_m^{0H_2}$):

$$\frac{\mu_m^{H_2H_2}}{\mu_m^{H_2}} = \frac{\exp(\gamma_{HH} - \delta - \tau_{HH}^2)}{\exp(\gamma_{HH} - \delta - \tau_{HH}^2) + \exp(\gamma_{LH} - \tau_{LH} - \delta) + \exp(\gamma_{0H})} \tag{2.27}$$

$$\frac{\mu_m^{LH_2}}{\mu_m^{H_2}} = \frac{\exp(\gamma_{LH} - \tau_{LH} - \delta)}{\exp(\gamma_{HH} - \delta - \tau_{HH}^2) + \exp(\gamma_{LH} - \tau_{LH} - \delta) + \exp(\gamma_{0H})} \tag{2.28}$$

$$\frac{\mu_m^{0H_2}}{\mu_m^{H_2}} = \frac{\exp(\gamma_{0H})}{\exp(\gamma_{HH} - \delta - \tau_{HH}^2) + \exp(\gamma_{LH} - \tau_{LH} - \delta) + \exp(\gamma_{0H})} \tag{2.29}$$

For men type L ($\mu_m^L = \mu_m^{HL} + \mu_m^{LL} + \mu_m^{0L}$):

$$\frac{\mu_m^{HL}}{\mu_m^L} = \frac{\exp(\gamma_{HL} - \tau_{HL} - \delta)}{\exp(\gamma_{HL} - \tau_{HL} - \delta) + \exp(\gamma_{LL} - \tau_{LL} - \delta) + \exp(\gamma_{0L})} \tag{2.30}$$

$$\frac{\mu_m^{LL}}{\mu_m^L} = \frac{\exp(\gamma_{LL} - \tau_{LL} - \delta)}{\exp(\gamma_{HL} - \tau_{HL} - \delta) + \exp(\gamma_{LL} - \tau_{LL} - \delta) + \exp(\gamma_{0L})} \tag{2.31}$$

$$\frac{\mu_m^{L0}}{\mu_m^L} = \frac{\exp(\gamma_{0L})}{\exp(\gamma_{HL} - \tau_{HL} - \delta) + \exp(\gamma_{LL} - \tau_{LL} - \delta) + \exp(\gamma_{0L})} \tag{2.32}$$

Marriage Market Equilibrium

The marriage market equilibrium must satisfy the following conditions for consistent choices between women and men (equaling "demand" and "supply").

$$\mu_f^{H_1H_1} = \mu_m^{H_1H_1} \tag{2.33}$$

$$\mu_f^{H_2H_2} = \mu_m^{H_2H_2} \tag{2.34}$$

$$\mu_f^{H_1L} + \mu_f^{H_2L} = \mu_m^{HL} \tag{2.35}$$

$$\mu_f^{LH} = \mu_m^{LH_1} + \mu_m^{LH_2} \tag{2.36}$$

$$\mu_f^{LL} = \mu_m^{LL} \tag{2.37}$$

In addition, the quantities must satisfy the following accounting identities.

$$\mu_f^{H_1H_1} + \mu_f^{H_1L} + \mu_f^{H_10} = R \text{ ...} H_1 \text{ female} \tag{2.38}$$

$$\mu_f^{H_2H_2} + \mu_f^{H_2L} + \mu_f^{H_20} = H_f - R \text{ ...} H_2 \text{ female} \tag{2.39}$$

$$\mu_m^{H_1H_1} + \mu_m^{LH_1} + \mu_m^{0H_1} = R \text{ ...} H_1 \text{ male} \tag{2.40}$$

$$\mu_m^{H_2H_2} + \mu_m^{LH_2} + \mu_m^{0H_2} = H_m - R \text{ ...} H_1 \text{ male} \tag{2.41}$$

$$\mu_f^{LL} + \mu_f^{LH} + \mu_f^{L0} = N_f - H_f \text{ ...} L \text{ female} \tag{2.42}$$

$$\mu_m^{LL} + \mu_m^{HL} + \mu_m^{0L} = N_m - H_m \text{ ...} L \text{ male} \tag{2.43}$$

Equilibrium Conditions

Given the exogenous variables (education costs parameters and population size), the final set of endogenous variables are:

$$U^H, U^L, H_f, \mu_f^{H_1H_1}, \mu_f^{H_1L}, \mu_f^{H_10}, \mu_f^{H_2H_2}, \mu_f^{H_2L}, \mu_f^{H_20}, \mu_f^{LH}, \mu_f^{LL}, \mu_f^{L0}$$
$$V^H, V^L, H_m, \mu_m^{H_1H_1}, \mu_m^{H_1L}, \mu_m^{H_10}, \mu_m^{H_2H_2}, \mu_m^{H_2L}, \mu_m^{H_20}, \mu_m^{LH}, \mu_m^{LL}, \mu_m^{L0}$$
$$R, \tau_{HH}^1, \tau_{HH}^2, \tau_{HL}, \tau_{LH}, \tau_{LL}$$

We have 30 endogenous variables. To pin down the system, we have:

(1) Expression of U, V as a function of utility parameters and educational distributions (Equation 2.9, 2.10, 2.12 & 2.13): 4 equations.

(2) Discrete marriage choices (Equation 2.15–Equation 2.32): 18 equations. However, 6 out of these 18 equations are redundant because we do not explicitly list μ_f^x and μ_m^y as the final endogenous variables. To see this, let's substitute the expression $\mu_f^{H_1} = \mu_f^{H_1H_1} + \mu_f^{H_1L} + \mu_f^{H_10}$ into Equation 2.15–Equation 2.17, then one of the three equations (e.g. Equation 2.17) is redundant.

Therefore, we effectively have 12 equations based on the discrete choice conditions.

(3) Marriage market clearing conditions and account identities (Equation 2.33–Equation 2.43: 11 equations.

(4) Definition of the LCMM meeting function

$$R = R(H_f, H_m). \tag{2.44}$$

(6) Individual educational choices that determines the distribution of education types

$$H_m = N_m * h_m = N_m * G^m(V^H - V^L - c_m) \tag{2.45}$$

$$H_f = N_f * h_f = N_f * G^f(U^H - U^L - c_f) \tag{2.46}$$

We have 30 equations for 3 endogenous variables that totally pin down the equilibrium educational choices and marriage matching functions.

2.10.2 Estimate/Calibrate the Marital Payoffs and Search Cost Parameter

This subsection describes how we estimate or calibrate important payoff and cost parameters in the marriage market. Suppose we observe the following statistics:

(1) The distribution of different types of marriages (but cannot distinguish between H_1 and H_2).

(2) Within the type HH marriage, the fraction of H_1H_1 marriages (estimated using auxiliary information from the CFPS).

From the individual choice functions, we can derive that:

$$2\ln\frac{\mu^{H_1H_1}}{\sqrt{\mu^{H_10}\mu^{0H_1}}} = \alpha_{HH} + \gamma_{HH} - \alpha_{H0} - \gamma_{0H} \tag{2.47}$$

$$2\ln\frac{\mu^{H_2H_2}}{\sqrt{\mu^{H_20}\mu^{0H_2}}} = \alpha_{HH} + \gamma_{HH} - \alpha_{H0} - \gamma_{0H} - 2\delta \tag{2.48}$$

$$\frac{1}{2}\ln\frac{\mu^{H_1L}\mu^{H_2L}(\mu^{H_1L}+\mu^{H_2L})^2}{\mu^{H_10}\mu^{H_20}(\mu^{0L})^2} = \alpha_{HL} + \gamma_{HL} - \alpha_{H0} - \gamma_{0L} - 2\delta \tag{2.49}$$

$$\frac{1}{2}\ln\frac{\mu^{LH_1}\mu^{LH_2}(\mu^{LH_1}+\mu^{LH_2})^2}{\mu^{0H_1}\mu^{0H_2}(\mu^{L0})^2} = \alpha_{LH} + \gamma_{LH} - \alpha_{L0} - \gamma_{0H} - 2\delta \tag{2.50}$$

$$2\ln\frac{\mu^{LL}}{\sqrt{\mu^{L0}\mu^{0L}}} = \alpha_{LL} + \gamma_{LL} - \alpha_{L0} - \gamma_{0L} - 2\delta \tag{2.51}$$

These are equivalent to Equation 2.1 & 2.2 and can be derived using the individual choice functions following Choo and Siow (2006). For the simulation, we need to at least identify the model parameters at the right-hand-side of these equations. However, these values are not readily available as we do not distinguish between H_1 and H_2.

Identify Parameters Associated with Noncollege Types

It turns out this does not matter for the martial surplus parameters involving noncollege types. Combining Equation 2.16, 2.17, 2.19, 2.20, we can easily show that:

$$\frac{\mu^{H_1L}}{\mu^{H_10}} = \frac{\mu^{H_2L}}{\mu^{H_20}}. \tag{2.52}$$

As a result:

$$\frac{\mu^{H_1L}}{\mu^{H_10}} = \frac{\mu^{H_2L}}{\mu^{H_20}} = \frac{\mu^{H_1L}+\mu^{H_2L}}{\mu^{H_10}+\mu^{H_20}} = \frac{\mu^{HL}}{\mu^{H0}}. \tag{2.53}$$

Plug Equation 2.53 back to Equation 2.49:

$$2\ln\frac{\mu^{HL}}{\sqrt{\mu^{H0}\mu^{0L}}} = \alpha_{HL}+\gamma_{HL}-\alpha_{H0}-\gamma_{0L}-2\delta. \tag{2.54}$$

The intuition is that noncollege types do not care about the distinct between H_1 and H_2. This is because the search costs do not affect the trade off between H and L or between L and L: H_1 and H_2 are equally valuable to an L type individual. Analogously, we have:

$$2\ln\frac{\mu^{LH}}{\sqrt{\mu^{L0}\mu^{0H}}} = \alpha_{LH}+\gamma_{LH}-\alpha_{L0}-\gamma_{0H}-2\delta. \tag{2.55}$$

Therefore, we can identify the marital surplus (net of the search cost) for HL, LH, and LL types of marriages.

Calibrate Search Cost and College-College Marital Surplus

Because we cannot distinguish between H_1 and H_2, additional information is required for knowing the surplus to HH marriages with and without paying the search cost.

Assume we can observe $\lambda = \frac{\mu^{H_1H_1}}{\mu^{HH}}$ (estimated using the CFPS data). We have the following condition for the marriage market equilibrium which is analogous to Equation 2.52 :

$$\frac{\mu^{LH_1}}{\mu^{0H_1}} = \frac{\mu^{LH_2}}{\mu^{0H_2}} \tag{2.56}$$

Recall that we defined the following quantities:

$$
\begin{aligned}
\lambda &\stackrel{\text{def}}{=} \frac{\mu^{H_1H_1}}{\mu^{HH}} \\
J_f &\stackrel{\text{def}}{=} \frac{R-\mu^{H_1H_1}}{H_f-R-\mu^{H_2H_2}} \\
J_m &\stackrel{\text{def}}{=} \frac{R-\mu^{H_1H_1}}{H_m-R-\mu^{H_2H_2}}.
\end{aligned}
$$

Rearrange Equation 2.52 , we get:

$$
\frac{\mu^{H_1L}}{\mu^{H_2L}} = \frac{\mu^{H_10}}{\mu^{H_20}} = \frac{\mu^{H_1L}+\mu^{H_10}}{\mu^{H_2L}+\mu^{H_20}} = \frac{R-\mu^{H_1H_1}}{H_f-R-\mu^{H_2H_2}} = J_f
$$

Similarly, we can get

$$
\frac{\mu^{LH_1}}{\mu^{LH_2}} = \frac{\mu^{0H_1}}{\mu^{0H_2}} = J_m
$$

Note that $\mu^{H_1H_1} = \lambda\mu^{HH}, \mu^{H_2H_2} = (1-\lambda)\mu^{HH}$

These conditions imply:

$$
\begin{aligned}
\mu^{H_10} &= \frac{J_f}{1+J_f}\mu^{H0},\ \mu^{H_20} = \frac{1}{1+J_f}\mu^{H0} \\
\mu^{0H_1} &= \frac{J_m}{1+J_m}\mu^{0H},\ \mu^{0H_2} = \frac{1}{1+J_m}\mu^{0H}
\end{aligned}
$$

The point of these equations above is that we can transform terms that we do not observe (the LFS) to things we can observe under certain functional form assumptions (RHS).

Plug these equations back to equations 2.47 and 2.48, we can get

$$\alpha_{HH}+\gamma_{HH}-\alpha_{H0}-\gamma_{0H}-2\delta = 2\ln\frac{\mu^{HH}}{\sqrt{\mu^{H0}\mu^{0H}}}+2\ln[(1-\lambda)\sqrt{(1+J_f)(1+J_m)}]$$
$$2\delta = 2\ln\frac{\lambda}{1-\lambda}-\ln J_f J_m$$

To pin down the value of these model parameters, we need to known the functional form of R. As noted in the paper, the form of R is not identified. Instead, we choose a Cobb-Douglas function that satisfies increasing returns to scale. We can then identify these values based on observed data moments.

Estimated/Calibrated Model Parameters

The value of λ is estimated based on the CFPS data: among the college-college marriages (pre-expansion cohorts), the fraction of couples that reported they met in school. We have $\lambda = 0.3$. Based on the value of λ and the conditions described above, we can estimate the parameters governing marital surplus and the search cost. The parameter values are reported in Table 2.16. Based on the the marriage model, we cannot separately identify α_{x0} and γ_{0y}. We normalize them as zero.[49]

Table 2.16: Model Parameter Values

Marriage Types	Parameter	Estimated Value
HH	$\alpha_{HH}+\gamma_{HH}-\alpha_{H0}-\gamma_{0H}$	6.04
HL	$\alpha_{HL}+\gamma_{HL}-\alpha_{H0}-\gamma_{0L}$	2.32
LH	$\alpha_{LH}+\gamma_{LH}-\alpha_{L0}-\gamma_{0H}$	4.29
LL	$\alpha_{LL}+\gamma_{LL}-\alpha_{L0}-\gamma_{0L}$	8.34
	δ	1.88

[49]This does not affect our results as only the marital surplus matters in the model for individual decisions.

Using these parameter values, we can then simulate (1) the marriage matching functions conditional on the distribution of education types and (2) the change in education choices and marriage matching functions following the change in the average costs of college education.

2.10.3 Additional Results for the Marriage Model

Distribution of Relative Education Within Married Couples

Figure 2.14: Relative Education of Wife and Husband

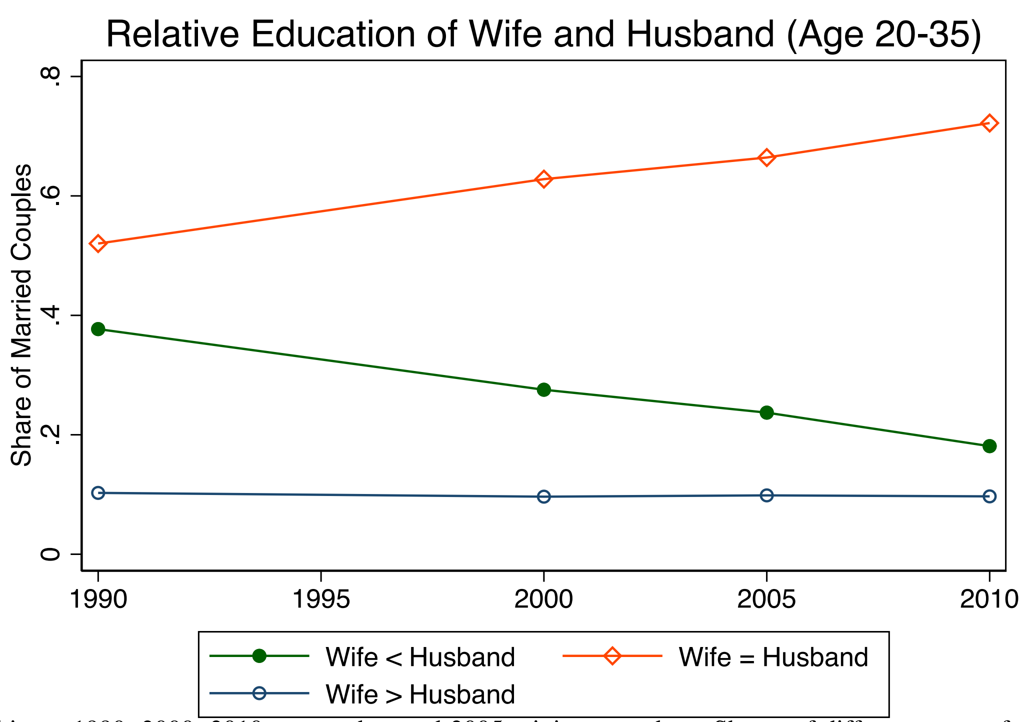

Source: Chinese 1990, 2000, 2010 census data and 2005 mini-census data. Shares of different types of couples by relative education are plotted.

Female College Ratios Overtook Male College Ratios in the College Expansion

Figure 2.15 reports the college ratio of women and men by cohort of birth. Post-expansion cohorts are those who were born after 1980. We divide the national sample into two groups based on the value of the college expansion proxy (Section 2.3.2). Three important patterns emerge from the figure. First, college ratios increased by more in high expansion regions relative to low expansion regions after the college expansion. This provides validation for our expansion proxy. Second, before the college expansion, the college ratio of men is higher than that of women. This gender difference has been reversed after the expansion. Third, the college ratio of women overtook that of men earlier and by a larger extent in high expansion regions. This lends supportive evidence to the effects of the college expansion on the reverse of the gender education gap.

Figure 2.15: College Ratio by Gender and College Expansion Intensity

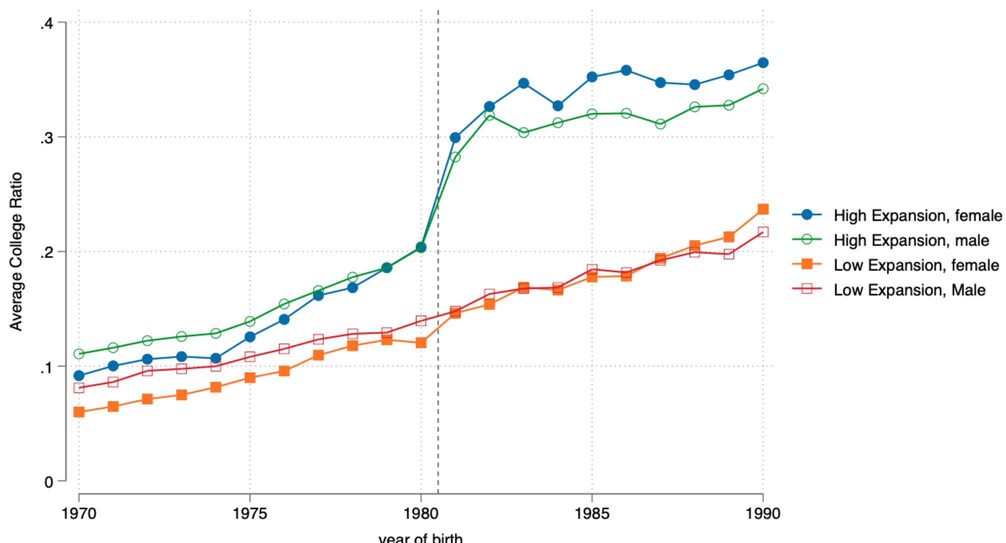

Source: Chinese 2010 census data and 2015 mini-census data. High Expansion refers to regions with the expansion intensity proxy above median. Low Expansion refers to regions with the expansion proxy below median.

Simulation Results for Noncollege Groups

Figure 2.16 presents the simulated comparative statics for the noncollege groups. The spillover effects in the marriage market also matter for the noncollege individuals. Although their search frictions are not affected in the model, their marriage outcomes are affected by the college expansion through changes in marginal distributions in the marriage market. The change in the relative distribution is such that both noncollege women and noncollege men become more scarce.

However, to the extent that noncollege men are less likely than noncollege women to marry up, noncollege men will face stronger competition from the increasing supply of college men. As a result, the model predicts that the marriage prospect of noncollege men will decrease. At the same time, the increasing relative demand for noncollege women from both college men and noncollege men drives up their marriage prospects.

Figure 2.16: College Expansion and Marriage Rates: Simulated Results

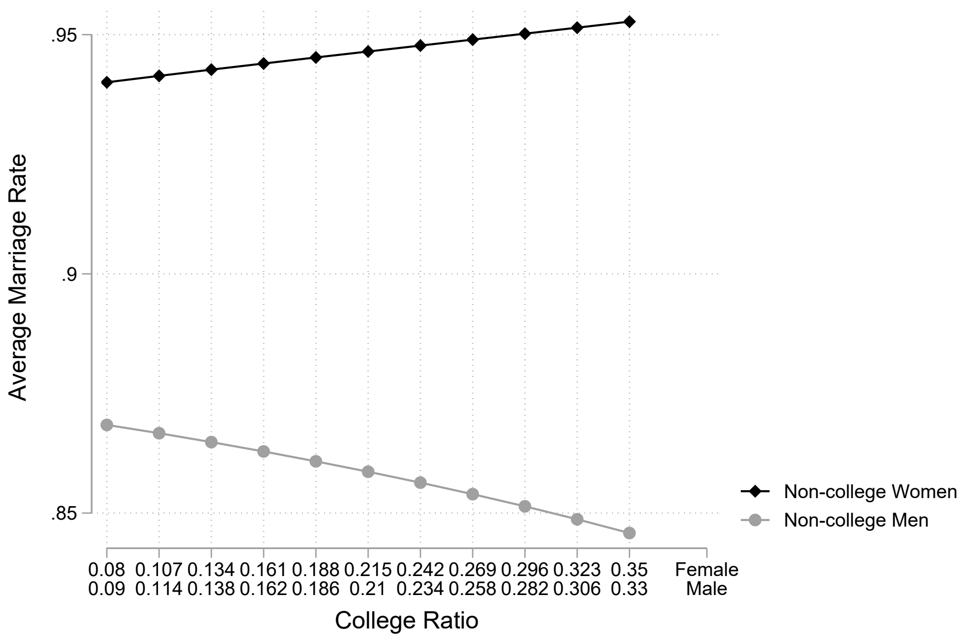

Notes: The x-axis displays the college ratios of women and men, which are allowed to evolve simultaneously.

Change in Within-marriage Transfers

We report the change in simulated within-marriage transfers following the college expansion in 2.17. Panel B shows that, consistent with the intuition, college men pay noncollege women a higher transfer and college women receive from noncollege men a lower transfer following the college expansion. The case for within-college type marriages is more complicated. As the college expansion reduces the relative "marriageability" of college men by more than that of the college women (because college men are more likely to marry down), college men are willing to pay a higher transfer to college women. On the other hand, as the college ratio increases faster for college women than for college men, creating a relative "over-supply" of college women, they lose some bargaining power and might receive a lower transfer from college men. To see this more clearly, in Figure 2.19 we plot the within marriage transfers under the hypothetical situation of gender symmetric expansion (Figure 2.5). In this case, we can observe a clear increase in the within-marriage transfer from college-educated husband to college-educated wife. In Figure 2.18, we plot the change in transfers if there was no search frictions in the model. In this case, the within-college-marriage transfer decreases following the college expansion. This implies that the increasing supply of college women relative to that of college men plays a more important role in this no-friction scenario.

Figure 2.17: College Expansion and Within-marriage Transfer: Simulated Results

A. Transfers in within-college type marriages

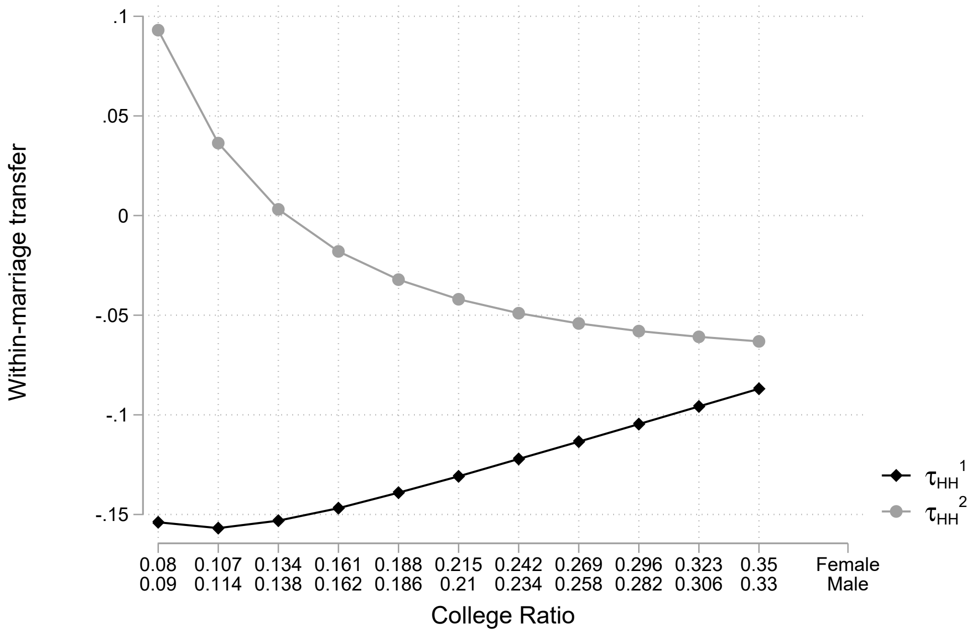

B. Transfers in cross-education type marriages

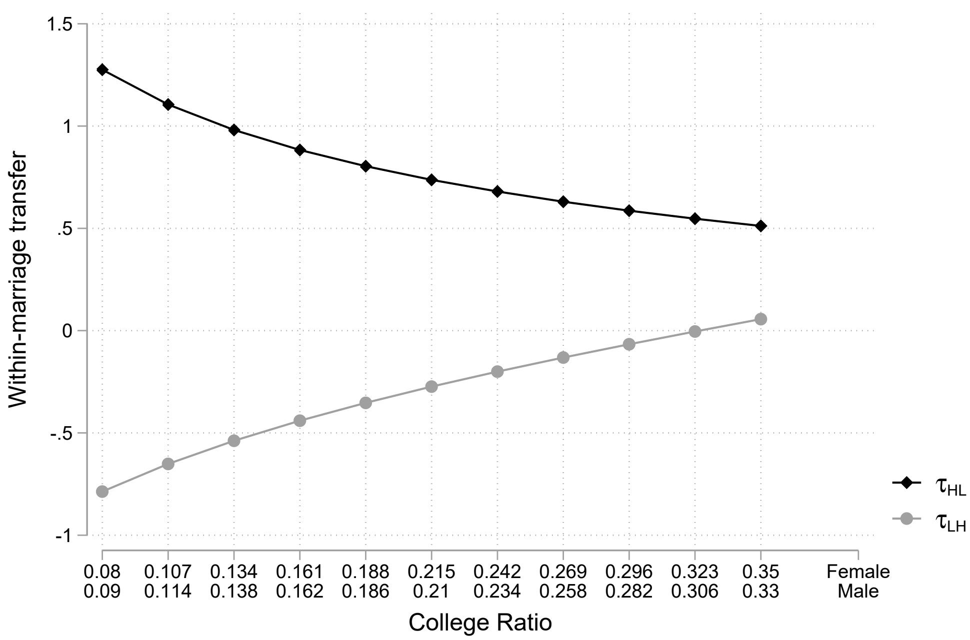

Notes: The x-axis displays the college ratios of women and men, which are allowed to evolve simultaneously.

Figure 2.18: College Expansion and Within-marriage Transfer: Simulated Results Without Local College Marriage Market

A. Transfers in within-college type marriages

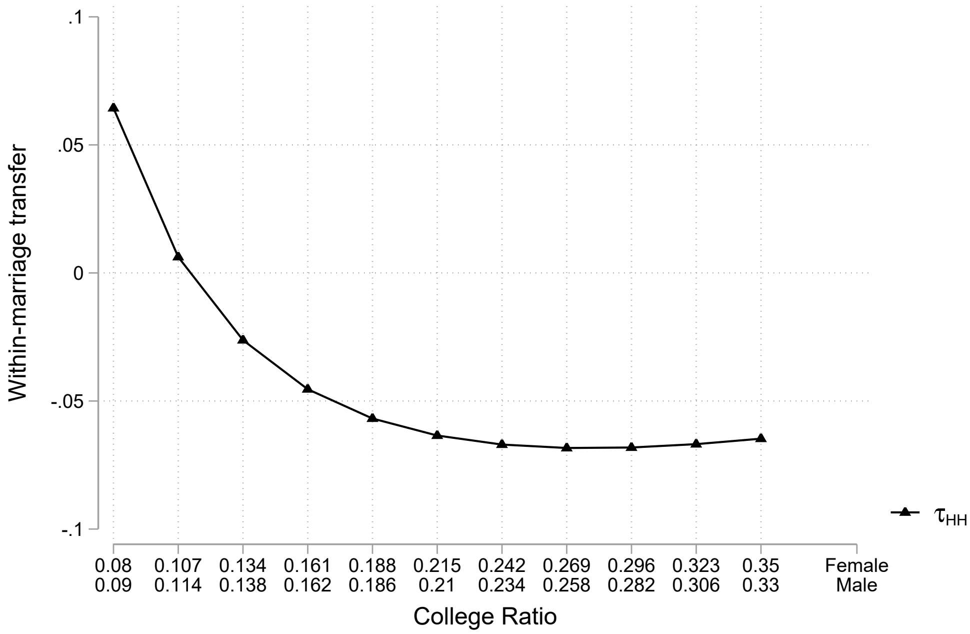

B. Transfers in cross-education type marriages

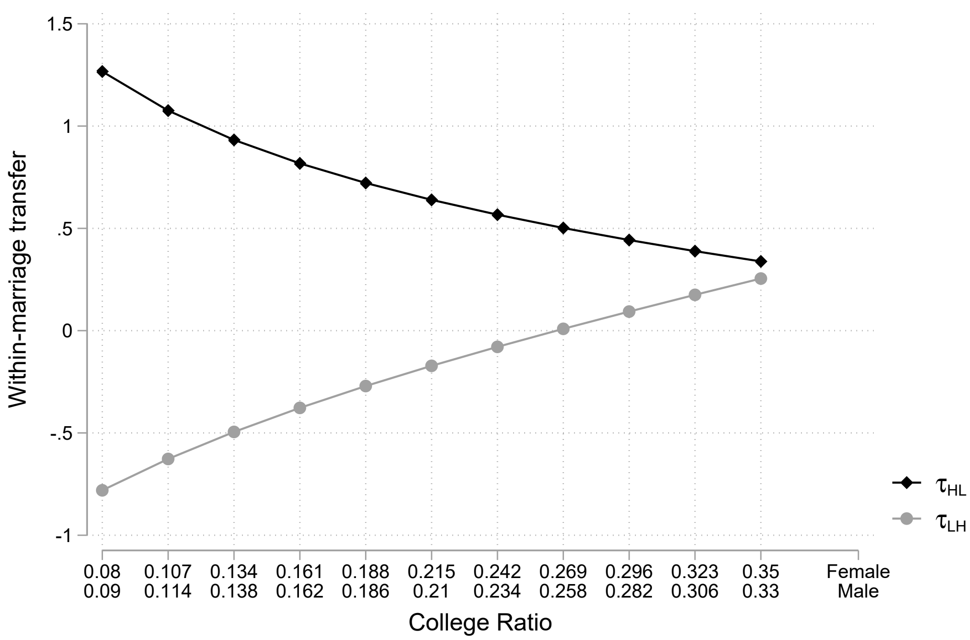

Notes: The x-axis displays the college ratios of women and men, which are allowed to evolve simultaneously.

Figure 2.19: College Expansion and Within-marriage Transfer: Simulated Results Under Gender-symmetric Expansion

A. Transfers in within-college type marriages

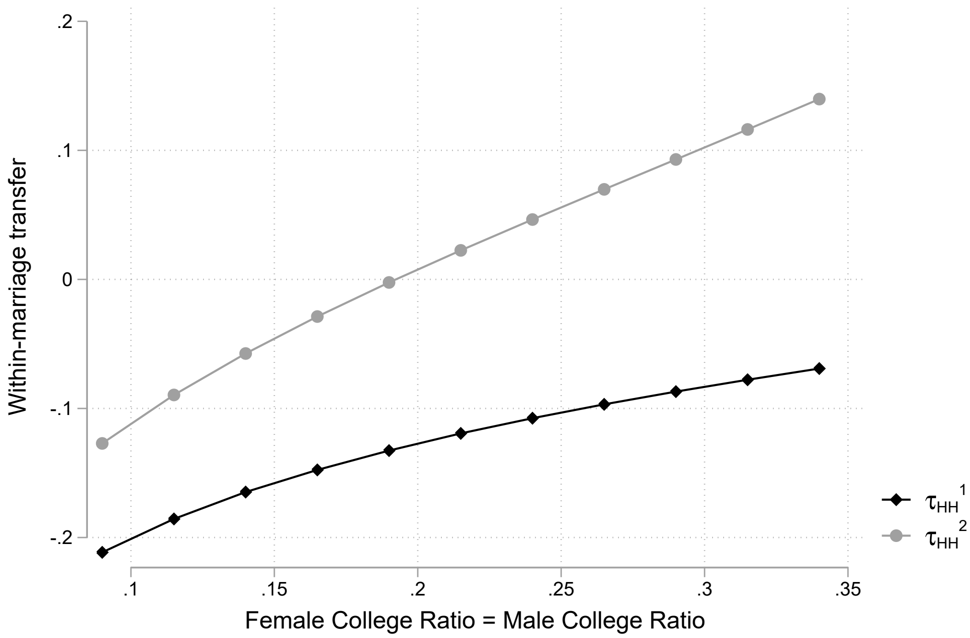

B. Transfers in cross-education type marriages

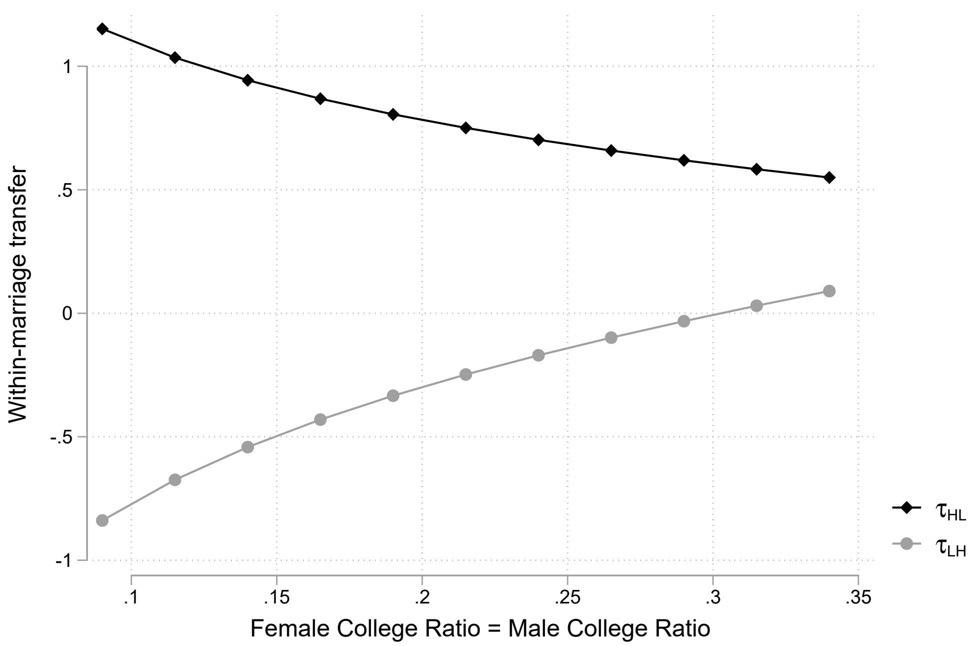

Notes: The x-axis displays the college ratios of women and men, which are allowed to evolve simultaneously.

2.10.4 Simulate the Full Model

This subsection simulates the full model by changing the education cost parameters. To fully describe individual choices, we need to know the values of preferences parameters α_{xy} and γ_{xy}. We have identified the joint surplus from the empirical matching functions. However, we cannot separately identify these specific preference parameters. Without loss of generality, we make the following normalization:

(1) $\alpha_{x0} = \gamma_{0y} = 0$;[50] (2) $\alpha_{xy} = \gamma_{xy}$.[51]

We can then pin down the expected marital returns to education U^H, U^L, V^H, V^L (Equation 2.9, 2.10, 2.12 & 2.13). The change in average education cost parameters (c_f, c_m) are set to be consistent with the education distributions in our baseline simulation exercise (Section 2.2.5). Without loss of generality, we assume that the distributions of idiosyncratic the education cost for both men and women follow a standard normal distribution. In Figure 2.20, we report the change in average educational attainment by gender following the change in the cost parameters. Following the college expansion, the college attainment increases faster for women than men. Figure 2.21 further display the marriage rates of college women and college men as a function of the reduction in the average education cost parameters. Consistent with our simulated results in the paper, the college expansion has an overall positive effects on the marriage probabilities of both college women and college men.

[50] With this normalization, α_{xy} and γ_{xy} represent the net systematic surplus from marriage.

[51] Because we cannot separately identify $alpha$, γ, and τ, this assumption simply states that by default the wife and the husband share the joint surplus from marriage. The within-marriage transfer τ determines the actual division of the joint surplus.

Figure 2.20: Reduction in Education Cost and College Attainment: Simulated Results

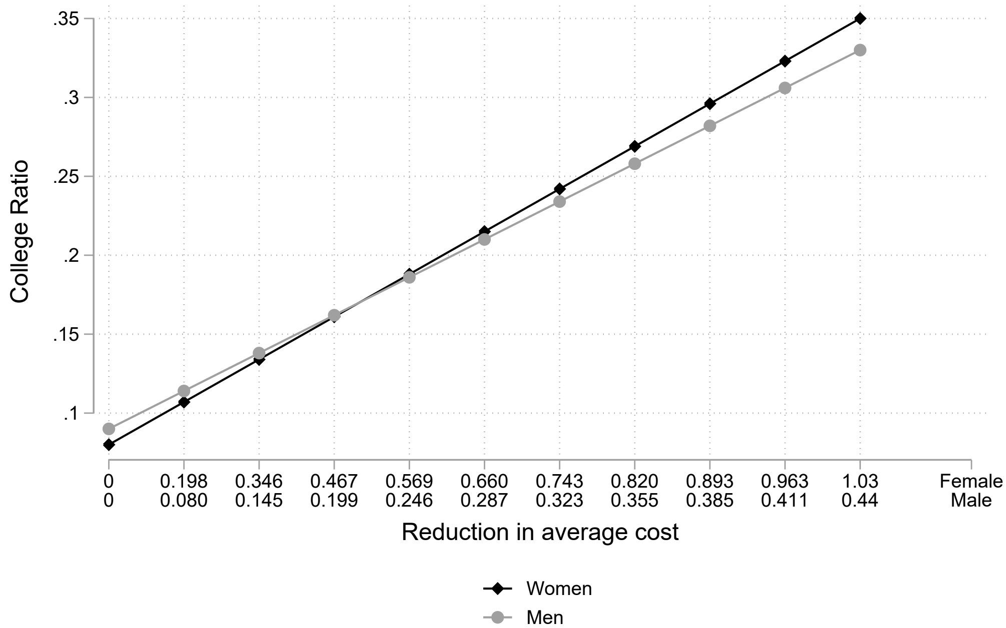

Notes: The x-axis displays the reduction in the mean college education cost, which are allowed to evolve simultaneously for both women and men.

Figure 2.21: Reduction in Education Cost and Marriage Rate: Simulated Results

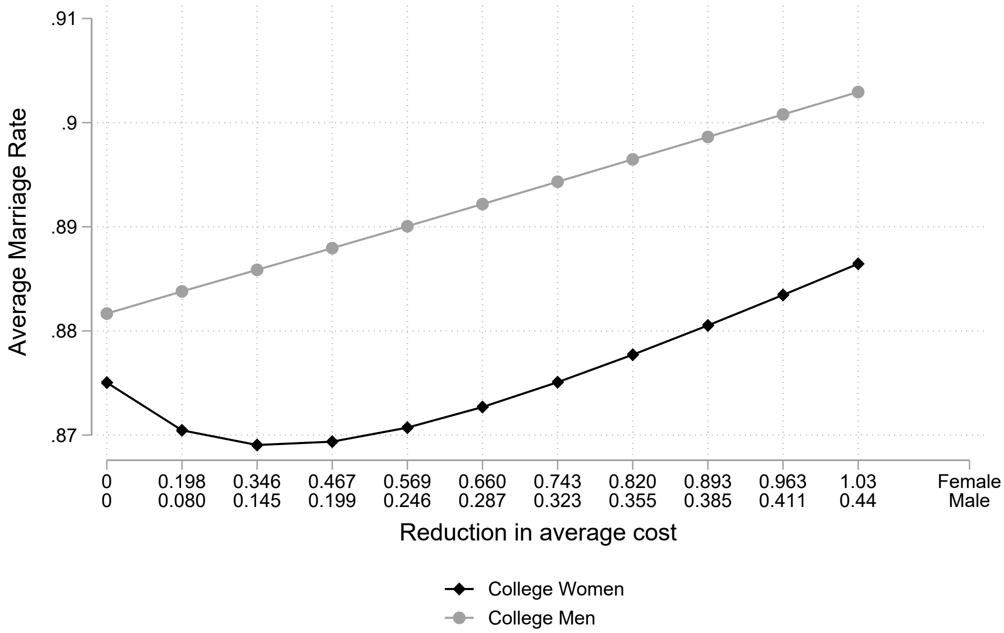

Notes: The x-axis displays the reduction in the mean college education cost, which are allowed to evolve simultaneously for both women and men.

2.11 Appendix D: Additional Tests for the College Expansion Proxy

In this section, we run several additional tests that shows: (1) the college expansion proxy is relevant for increase in college enrollment following the college expansion, and (2) the college expansion proxy is exogenous to economic growth and local sex ratio.

To examine whether our proxy predicts the local intensity of the college expansion, we run the following dynamic DID model:

$$(\frac{CollegeEnroll}{PopSize})_{pt} = \sum_{\tau} \alpha_{\tau} ExpProxy_p * (t = \tau) + \theta_p + \mu_t + \varepsilon_{pt} \tag{2.57}$$

The dependent variable is the fraction of college enrollment over population in prefecture (city) p. θ_p and μ_t stand for prefecture and year fixed effects, respectively. The data is at each prefecture-year level, which is collected from City Year Books. In Figure 2.22, we plot the dynamic coefficients on the expansion proxy. The results are consistent with the notion that the college expansion drastically increased college degrees in regions with a larger proxy. One unit of the estimated coefficient implies that for one more college student per capita in 1982, the college enrollment per capita in a prefecture will increase by one between 1998 and the given year in the x-axis. The triangle dots in Figure 2.22 are the hypothetical dynamic coefficients under the assumption that the college expansion has been perfectly proportional to the initial college enrollment in 1982.[52] The actual dynamic coefficients and the predicted coefficients closely trace each other, with the former being larger after the college expansion. This pattern seems to suggest that the regional inequality of higher education has been further amplified relative to the benchmark measured using the 1982 data.

[52] That is, we assume that the college enrollment in each prefecture is the product of the national enrollment in a given year and the fraction of the prefecture's enrollment to national enrollment in 1982.

Figure 2.22: Dynamic Effects of the College Expansion on College Enrollment

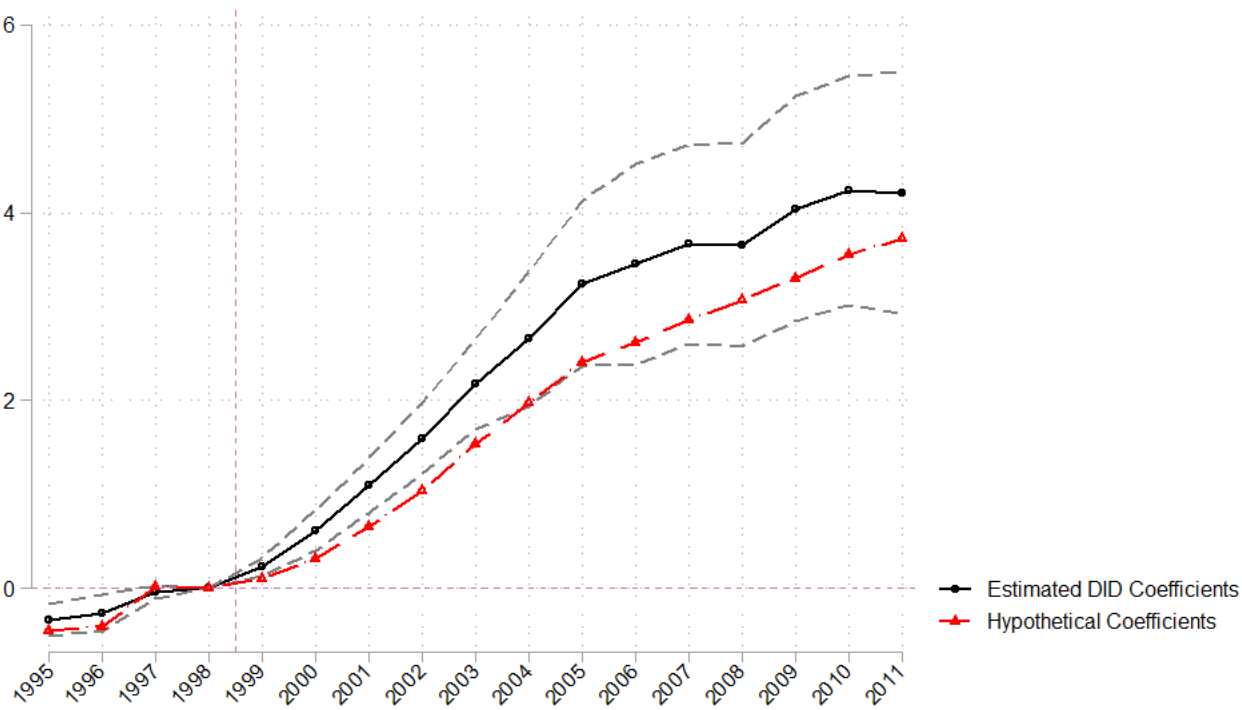

Notes: This figure plots the dynamic coefficients (α_τ) from Equation2.57 and associated 95% confidence intervals. The red triangle dots are dynamic coefficients assuming that the college expansion has been perfectly proportional to the initial college enrollment in 1982.

In order for our DID strategy to plausibly identify the causal impacts of the expansion, the expansion proxy must not be associated with other omitted factors that affect the marriage market. One concern is that the expansion proxy might be correlated with changes in regional economic performance. That potential correlation would bias our results if income directly affects marriage outcomes (Burgess et al., 2003; Chu et al., 2018; Hankins and Hoekstra, 2011). In Figure 2.23, however, we show that the expansion proxy is not systematically associated with the growth of GDP per capita in subsequent years. The growing confidence interval is also consistent with the notion that regional economic inequality has been enlarged in China. Nevertheless, the increasing regional inequality is not associated with our proxy for the college expansion.

Figure 2.23: Dynamic Effects of the College Expansion on GDP per capita

Notes: The dependent variable is the log GDP per capita of each prefecture obtained from City Statistical Yearbooks. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit is equal to one SD of the treatment proxy across prefectures. This figure plots the dynamic coefficients and the 95% confidence intervals.

Another factor is the sex ratio, which is potentially important given that sex imbalance has become increasingly severe and has affected China's marriage market in various ways (Wei and Zhang, 2011; Ebenstein and Sharygin, 2009). We test the correlation between the expansion proxy and local sex ratio for the cohorts of birth between 1970–1990 in Figure 2.24. Based on the univariate regression between sex ratio and the expansion proxy, we find no evidence that the proxy is significantly associated with sex ratio.

Figure 2.24: Correlation between College Expansion and Sex Ratio

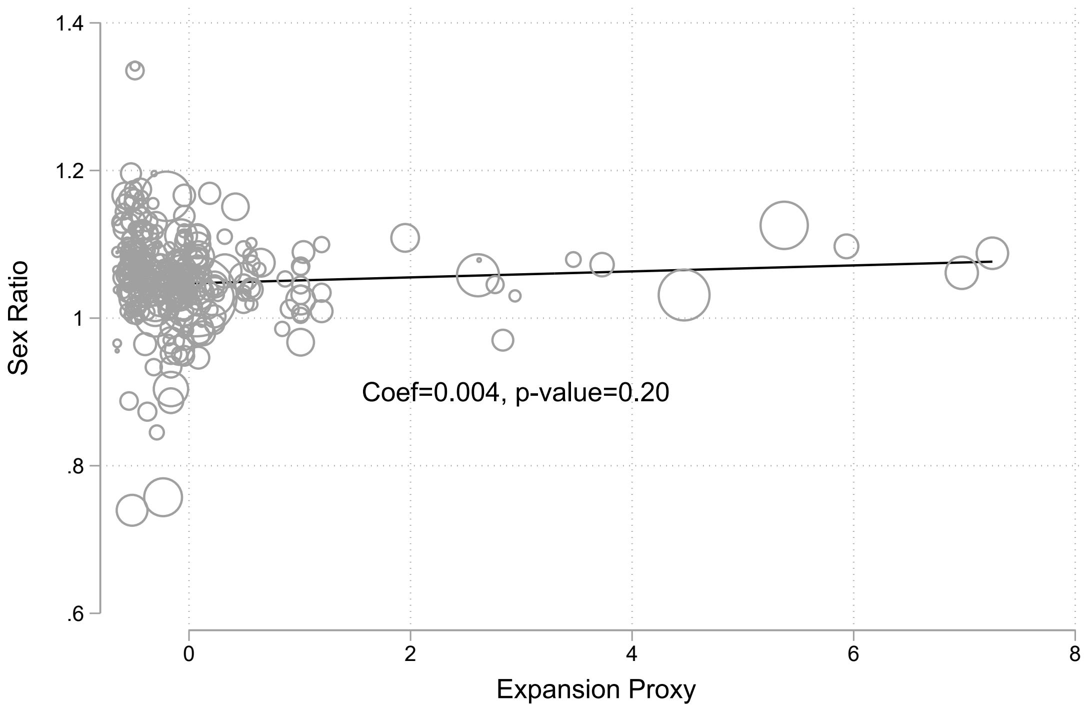

Notes: The dependent variable is sex ratio of cohorts of birth between 1970 and 1990. Data is from the microdata of 2000 China census obtained via IPUMBS International. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit is equal to one SD of the treatment proxy across prefectures. Each dot represents a prefecture weighted using population size. The solid line is fitted line of the univariate linear regression. P-values are computed based on robust standard errors.

2.12 Appendix E: Predicting Regular College Degrees

2.12.1 Regular college vs. other post-secondary credentials

There is one specific form of measurement errors in the census data used for this research: we do not distinguish between regular college students and other post-secondary credentials that serve adults older than regular college ages in the census data. As the college expansion was all about the regular college system, introducing this form of measurement errors in our analysis might confound our estimates. Intuitively, this would tend to attenuate our estimates if the measurement errors in classifying the relevant group are not correlated with the treatment of the policy. To further rule out the confounding effects, we use a prediction exercise to address this concern. We predict the status of regular college graduates using information available in another smaller dataset where we can distinguish different types of college graduates.

We resort to the China Family Panel Studies (CFPS). The CFPS is a representative household survey of China. Among rich information about individuals and their families, the CFPS contains detailed information on educational attainment including whether the college degree comes from a regular college or the post-secondary credential system. We use the education and demographic information to predict regular college degree status, and then apply that prediction model to census data to obtain a sample of individuals who are (more likely) regular college system graduates. The rationale of this exercise is that it could reduce the measurement errors in college types by picking a group of individuals who are much more likely regular college graduates.

2.12.2 The Prediction Model

To apply the prediction model to census data, we need information shared by the CFPS and the census data. The common variables include: year of birth, province of residence, gender, ethnics, rural-urban residency registration (*Hukou*) status,

in the CFPS data, the sample size is very small for people with a college degree in particular. Therefore, for a given cohort-of-birth-by-province cell, there are very likely very few or even no observations. To deal with that, we group the years of birth into cohort groups of birth in the following way:

In the first step, we estimate the following Logit model:

$$\begin{aligned} Prob[Regular_i = 1] = F(&\beta_0 + \beta_1 Male_i + \beta_2 Urban_i + \beta_3 Han_i + \gamma_b + \eta_{Prov} + \gamma_b * \eta_{Prov} \\ &+ (Male_i * \gamma_b)\Pi_1 + (Han_i * \gamma_b)\Pi_2 + (Urban_i * \gamma_b)\Pi_3) \end{aligned}$$

$Male_i$, $Urban_i$, and Han_i are dummy variables for being male, having urban residency, being of the Han ethnicity. γ_b represents dummies for the cohort of birth, η_{Prov} is dummies for the province of residence. We also include the interactions between the first three demographic variables and cohort dummies as well as the interactions between cohort dummies and province dummies. The model is estimated using the CFPS data, where we can observe whether a respondent holds a regular college degree or just a post-secondary credential if they reports having a post-secondary degree. The estimated model coefficients are then applied to respondents who have a post-secondary degree in the census data. If the predicted probability is above 50%,

2.12.3 Internal Accuracy in the CFPS Data

We first evaluate the accuracy of the prediction model when it is applied to the original CFPS data (the dataset used to "train" the prediction model). We use two measures that are typically used to assess machine learning models: the *precision* rate and the *recall* rate. The precision rate captures, among those who are predicted to have a regular degree, what fraction is truly regular college graduates. The recall rate is the fraction of regular college graduates who are predicted as so. Table 2.17 reports these numbers: the precision rate is 78.8% and the recall rate is 92.5%. For our empirical analysis, this implies that when we look at the predicted regular college graduates, the vast majority of them are indeed regular college graduates.

We further assess the performance of this prediction model using a K-fold cross validation. Specifically, we divide the sample into five random sub-samples. Each time, we use four of them as the training dataset for the prediction model, and evaluate the accuracy of the prediction using the remaining sub-sample as the test dataset. As shown in Table 2.17, the average precision rate is 76.3%, and the average recall rate is 88.2%. Although we only include in the predict model a limited set of basic demographic information, it seems to predict the regular college degree status pretty accurately.

Table 2.17: Internal Accuracy in the CFPS Data

	Full Sample	5-Fold Cross-validation
Precision Rate	77.8%	76.5%
Recall Rate	93.7%	89.2%

Note: The precision rate is the ratio of actual regular college graduates in the predicted sample. The recall rate is the ratio of predicted regular college graduates in the sample of actual regular college graduates.

2.12.4 External Accuracy in the Census Data

We cannot directly test the accuracy of this prediction model in the census data. To provide some suggestive evidence on how accurately the prediction fits the actual data, we first compare the average regular college ratio of the predicted sample to the actual college ratio that we can gauge using public data. For each cohort of birth, we impute the ratio of college enrollment using public information on the number of births and the number of college enrollment, assuming that everyone goes to college at the age of 18. We then compare this ratio to the predicted regular college ratio by the cohort of birth. The results are reported in Figure 2.25.

Figure 2.25: Compare Predicted and Actual College Enrollment Rates

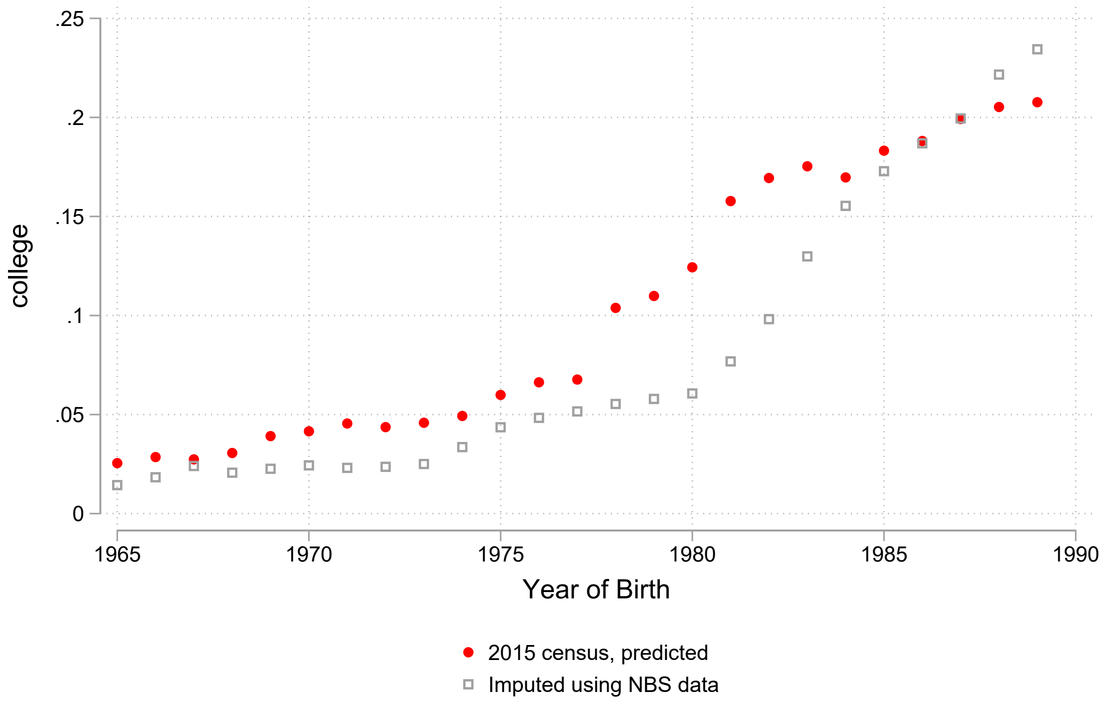

Note: *2015 census, predicted* is the regular college enrollment ratio based on the predicted results. *Imputed using NBS data* is the regular college enrollment ratio calculated using yearly new enrollment data and birth cohort size published by the National Bureau of Statistics of China.

In the second exercise, we evaluate the performance of this model in the 2000 census micro-data which contains information on whether a college degree is regular or just a post-secondary credential. Applying the prediction model parameters estimated with the CFPS data to the 2000 census, we are able to compare the predicted regular degree to the actual regular degree. The results are reported in Table 2.18. We can still obtain a meaningfully large precision rate.

Table 2.18: External Accuracy in 2000 Census

	2000 Census College Subsample
Precision Rate	73.9%
Recall Rate	69.2%

2.13 Appendix F: Additional Results

In this section, we report additional empirical results that supplement main results presented in the paper.

2.13.1 Parallel Pre-trends for Non-college Groups

Figure 2.26: Trend in Marriage Rates of College Graduates by College Expansion Intensity

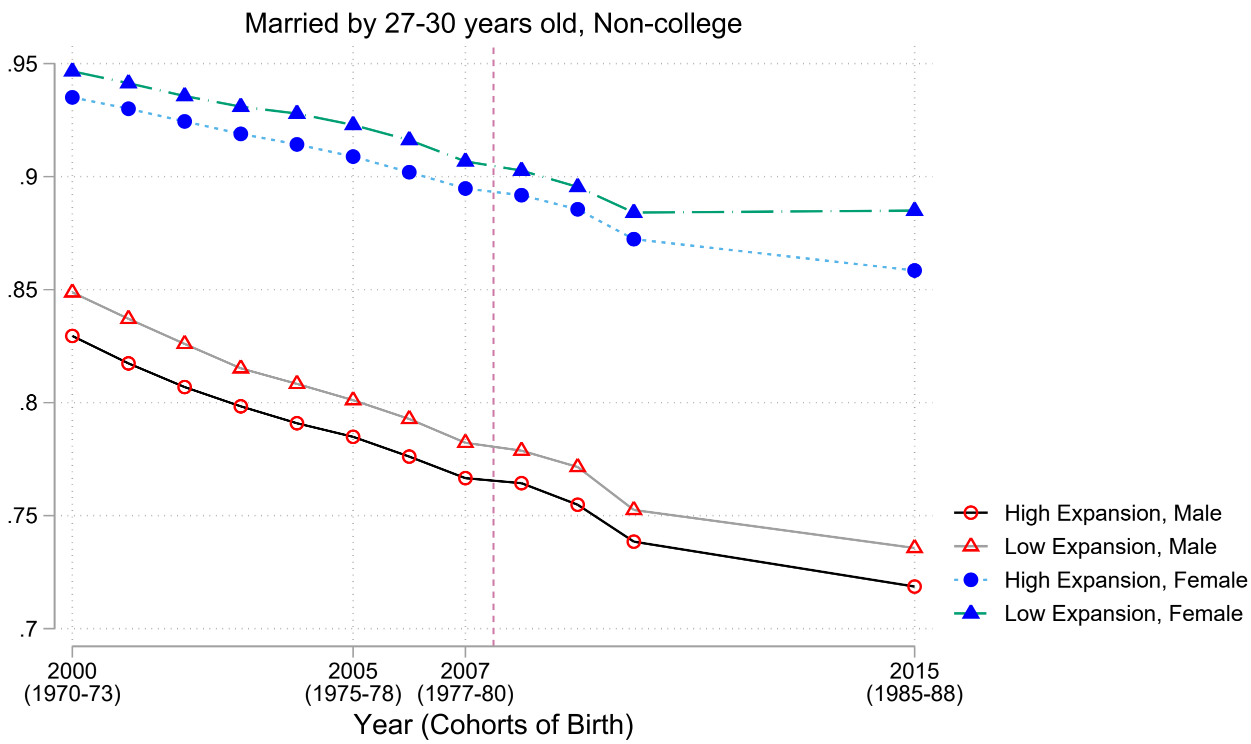

Notes: The sample is divided into two groups based on the value of the expansion treatment proxy below or above the median. The marital history information in the 2010 census allows us to impute the ever-married fractions before 2010. Using the 2015 mini-census, however, we can only know people's marital status as of 2015.

2.13.2 Additional Robustness Checks

Table 2.19: Heterogeneity by College Types: Four-year Universities vs. Three-year Colleges

Dependent variable: Ever being married				
Post-expansion cohorts	1985-88, 27-30 years old in 2015		1981-84, 31-34 years old in 2015	
Pre-expansion cohorts	1975-78, 27-30 years old in 2005		1975-78, 31-34 years old in 2009	
	Four-year	Three-year	Four-year	Three-year
A. Male	(1)	(2)	(3)	(4)
ExpProxy$*$Post	0.0277***	0.0235***	0.0161***	0.00708**
	(0.00469)	(0.00521)	(0.00465)	(0.00286)
Observations	18153	22043	17377	19109
Marriage rate of				
Pre-expansion cohorts	0.587	0.688	0.874	0.890
Post-expansion cohorts	0.575	0.647	0.857	0.860
B. Female				
ExpProxy$*$Post	0.0198***	0.0146***	0.00701	0.00217
	(0.00396)	(0.00401)	(0.00524)	(0.00274)
Observations	17375	21102	15911	17983
Marriage rate of				
Pre-expansion cohorts	0.745	0.816	0.906	0.928
Post-expansion cohorts	0.667	0.749	0.887	0.899

Notes: The marital outcome of pre-expansion cohorts is constructed using marriage history so that it is comparable to the post-expansion cohorts. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit represents one standard deviation across all prefectures. All regressions control for prefecture fixed effects and age fixed effects. Standard errors clustered at the prefecture level are in parentheses. There are 340 clusters. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.20: Effects of College Expansion on Marriage Probability: Treatment Proxy using 1990 Census

Dependent variable: Ever being married.				
Post-expansion cohorts	1985-88, 27-30 years old in 2015		1981-84, 31-34 years old in 2015	
Pre-expansion cohorts	1975-78, 27-30 years old in 2005		1975-78, 31-34 years old in 2009	
	College	Non-college	College	Non-college
A. Male	(1)	(2)	(3)	(4)
ExpProxy∗Post	0.0220***	-0.0009	0.0132***	-0.0006
	(0.0080)	(0.0028)	(0.0045)	(0.0017)
	[0.0273]	[0.0059]	[0.0124]	[0.0008]
Observations	40196	187259	36486	181105
Marriage rate of				
Pre-expansion cohorts	0.644	0.793	0.883	0.882
Post-expansion cohorts	0.613	0.728	0.858	0.874
B. Female				
ExpProxy∗Post	0.0139**	-0.0008	0.0062	-0.0007
	(0.0056)	(0.0034)	(0.0041)	(0.0012)
	[0.0171]	[0.0020]	[0.0053]	[-0.0009]
Observations	38477	182842	33894	176213
Marriage rate of				
Pre-expansion cohorts	0.787	0.916	0.919	0.961
Post-expansion cohorts	0.709	0.873	0.892	0.953

The marital outcome of pre-expansion cohorts is constructed using marriage history so that it is comparable to the post-expansion cohorts. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit represents one standard deviation across all prefectures. All regressions control for prefecture fixed effects and age fixed effects. Standard errors clustered at the prefecture level are in parentheses. There are 340 clusters. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Baseline estimates are in brackets.

Table 2.21: Results Controlling for Baseline City Characteristics

Dependent variable: Ever being married.				
Post-expansion cohorts	1985-88, 27-30 years old in 2015		1981-84, 31-34 years old in 2015	
Pre-expansion cohorts	1975-78, 27-30 years old in 2005		1975-78, 31-34 years old in 2009	
	College	Non-college	College	Non-college
A. Male	(1)	(2)	(3)	(4)
ExpProxy∗Post	0.0299***	0.00806***	0.0122***	0.00108
	(0.00552)	(0.00224)	(0.00357)	(0.00146)
	[0.0273]	[0.0059]	[0.0124]	[0.0008]
Observations	36988	160626	33678	154910
Marriage rate of				
Pre-expansion cohorts	0.641	0.793	0.882	0.882
Post-expansion cohorts	0.615	0.728	0.858	0.874
B. Female				
ExpProxy∗Post	0.0176***	0.00239	0.00414	-0.000987
	(0.00382)	(0.00314)	(0.00342)	(0.00124)
	[0.0171]	[0.0020]	[0.0053]	[-0.0009]
Observations	35561	157529	31519	151635
Marriage rate of				
Pre-expansion cohorts	0.784	0.916	0.916	0.961
Post-expansion cohorts	0.709	0.873	0.892	0.953

Notes: The marital outcome of control cohorts is constructed using marriage history so that it is comparable to the treatment cohorts. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit represents one standard deviation across all prefectures. City level control variables include log GDP per capita and sex ratio, both measured in 2000 and interacted with *Post* dummy. All regressions control for prefecture and age fixed effects. Standard errors clustered at the prefecture level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Baseline estimates are in brackets.

Table 2.22: Results Controlling Province by Year Fixed Effects

Dependent variable: Ever being married				
Post-expansion cohorts	1985-88, 27-30 years old in 2015		1981-84, 31-34 years old in 2015	
Pre-expansion cohorts	1975-78, 27-30 years old in 2005		1975-78, 31-34 years old in 2009	
	College	Non-college	College	Non-college
A. Male	(1)	(2)	(3)	(4)
ExpProxy$*$Post	0.0239***	0.00717***	0.0069***	0.0033**
	(0.00448)	(0.00253)	(0.00222)	(0.00153)
	[0.0273]	[0.0059]	[0.0124]	[0.0008]
Observations	40196	187259	36486	181105
Marriage rate of				
Pre-expansion cohorts	0.644	0.793	0.883	0.882
Post-expansion cohorts	0.613	0.728	0.858	0.874
B. Female				
ExpProxy$*$Post	0.0216***	0.00149	-0.00064	-0.00188
	(0.00305)	(0.00145)	(0.00269)	(0.00191)
	[0.0171]	[0.0020]	[0.0053]	[-0.0009]
Observations	38477	182842	33894	176213
Marriage rate of				
Pre-expansion cohorts	0.787	0.916	0.919	0.961
Post-expansion cohorts	0.709	0.873	0.892	0.953

The marital outcome of control cohorts is constructed using marriage history so that it is comparable to the treatment cohorts. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit represents one standard deviation across all prefectures. All regressions control for province fixed effects interacted with the *Post* dummy. All regressions control for prefecture and age fixed effects. Standard errors clustered at the prefecture level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

2.13.3 Assortative Mating

Table 2.23: Effects of the College Expansion on Assortative Mating: Four-year Universities vs. Below

		Post-expansion: 1985-88		Post-expansion: 1981-84	
	Region by Expansion Intensity	High	Low	High	Low
Cohorts	Pre-expansion: 1975-78	0.032	0.014	0.040	0.016
	Post-expansion	0.075	0.024	0.094	0.029
	Diff-in-diff	**0.033******		**0.040******	
		(0.002)		(0.002)	

Notes: The index is calculated using the sample of married couples in the random samples of 2010 census and 2015 mini-census. The sample include married couples with either side falling in the specified cohort range. Ages of couples in pre-expansion cohorts are adjusted to be comparable to those in the post-expansion cohorts. The national sample is divided into high vs. low regions based on the treatment proxy being above or below median. Standard errors in parentheses are estimated by boostrapping from the original sample 1000 times. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.24: Effects of the College Expansion on Assortative Mating: Various Indexes

		Log Odds Ratio	Minimum Distance	Correlation	Absolute Difference
A. Post-expansion cohorts: 1985–1988, Pre-expansion cohorts: 1975–78					
1975–78	High	4.347	0.774	0.707	0.082
	Low	4.623	0.769	0.679	0.050
1985–88	High	4.030	0.757	0.741	0.149
	Low	4.264	0.744	0.693	0.075
Diff-in-diff		**0.042**	**0.007**	**0.020***	**0.041*****
		(0.086)	(0.013)	(0.011)	(0.002)
B. Post-expansion cohorts: 1981–1984, Pre-expansion cohorts: 1975–78					
	High	4.308	0.780	0.719	0.095
	Low	4.524	0.764	0.678	0.054
1981–84	High	4.202	0.784	0.762	0.158
1975–78	Low	4.488	0.768	0.715	0.075
Diff-in-diff		**-0.071**	**0.001**	**0.007**	**0.042*****
		(0.083)	(0.012)	(0.010)	(0.002)

Notes: See Table 2.12 for the definition of different indexes. The indexes are calculated using the sample of married couples in the random samples of 2010 census and 2015 mini-census. The sample include married couples with either side falling in the specified cohort range. Ages of couples in pre-expansion cohorts are adjusted to be comparable to those in the post-expansion cohorts. The national sample is divided into high vs. low regions based on the treatment proxy being above or below median. Standard errors in parentheses are estimated by boostrapping from the original sample 1000 times. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

2.13.4 Marriage Age Gap for Non-college Groups

In Table 2.25, we show that the college expansion also reduced the marriage age gap of non-college women and non-college men. The magnitude of these effects is small. For example, Column (1) shows that a SD higher treatment proxy caused 0.06 years drop in the marriage age gap of non-college women. Nevertheless, this is unlike the mostly null effects on their marriage probabilities. Our model, which does not incorporate matching on age (and the interaction between education and age in matching), is unable to provide a theoretical guidance. We discuss some possible intuitions here.

Assume non-college women have preferences for both younger age (physical desirability) and higher income (or human capital). Before the college expansion, an older age is possibly a proxy for a higher income among the large pool of non-college men. Therefore, non-college women who have a stronger preferences for spouses' income may accept a larger age gap. Following the college expansion, those women with a stronger preferences for income are more likely to match with increasingly available college men. This is consistent with Column (2) of Table 2.9 that shows a small but positive increase in non-college women marrying college men. On the other hand, before the expansion, one useful strategy for non-college men that have a preference for younger wives might be to accumulate more human capital and wait longer before they marry. This tends to create a larger marriage age gap. However, following the college expansion, with non-college women that prefer a higher income increasingly marry college men, the marginal benefit of waiting longer for human capital accumulation depreciates for non-college men. As a result, they choose to marry earlier and accept a smaller age gap. This story might explain our findings. However, it is just a conjecture that is not formally test either theoretically or empirically. It can be a fruitful future line of research to dig into the interaction between (college) education and age in the marriage market.

Table 2.25: Effects of the College Expansion on Marriage Age Gap of Non-college Groups

	(1)	(2)	(3)	(4)
Dependent variable: Age gap (Husband - Wife)				
Post-expansion cohorts	1985-88, 27-30 years old in 2015		1981-84, 31-34 years old in 2015	
Pre-expansion cohorts	1975-78, 27-30 years old in 2005		1975-78, 31-34 years old in 2009	
	Female	Male	Female	Male
ExpProxy$*$Post	-0.0586**	-0.0344**	-0.0504*	-0.0373*
	(0.0232)	(0.0135)	(0.0271)	(0.0221)
Observations	123683	105990	128557	117749
Average Age Gap				
Pre-expansion cohorts	2.143	1.296	2.121	1.566
Post-expansion cohorts	2.294	1.082	2.264	1.626

Notes: The marital outcome of pre-expansion cohorts is constructed using marriage history so that it is comparable to the post-expansion cohorts. *ExpProxy* is the proxy for college expansion, which is standardized so that one unit represents one standard deviation across all prefectures. All regressions control for prefecture fixed effects and age fixed effects. Standard errors clustered at the prefecture level are in parentheses. There are 340 clusters. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

2.14 Appendix G: Additional Discussions on Assortative Mating Index

We start by summarzing whether these indices satisfy the properties proposed in (Chiappori et al., 2021).

Properties	Absolute Difference	Log Odds Ratio	Min Distance	Correlation
Scale invariance	Y	Y	Y	Y
Symmetry	Y	Y	Y	Y
Monotonicity	Y	Y	Y	Y
Perfect PAM	N	Y	Y	Y

The first two properties are intuitive and we can easily observe that the two indices satisfy them.[53]

Monotonicity means that, when we hold the identical marginal distributions of college women ($\frac{k_1}{k_1+k_2}$) and college men $\frac{k_1}{k_1+k_3}$, then adding more people in each diagonal cell ($k_1$or k_4) will always increase assortativeness. The two indices both satisfy this property.

2.14.1 Perfect PAM

A more complicated issue arises with the Perfect PAM. This property basically states that a contingency table T_{PAM} under perfect assortativeness (meaning that $k_2 = 0$ or $k_3 = 0$) should achieve the maximum value of the index and no other table under imperfect assortativeness ($k_2 * k_3 > 0$) should lead to an index value than T_{PAM}.

They proposed two versions of PAM: the strong version means the stated property hold as long as one of $\{k_2, k_3\}$ is 0, while the weak version says the property holds when $k_2 = k_3 = 0$. The weak version is therefore a necessary condition for the strong version.

The log odds ratio satisfies both strong and weak PAM condition because the ratio is $+\infty$ for a contingency table with perfect assortativeness. The absolute difference measure, however, does not. Let's consider the case when $k_2 = k_3 = 0$, we have

$$\begin{aligned} AbsDiff &= \frac{k_1 k_4 - k_2 k_3}{K^2} \\ &= \frac{k_1(K - k_1)}{K^2} \end{aligned}$$

which achieves a maximal value of 0.25 when $k_1 = 0.5K$ (the educated and the uneducated are balanced) but goes to 0 when k_1 approaches 0 or 1. Therefore, we can easily construct a table whose index value is greater than the index of a perfect assortativeness table, as shown in Table

[53] Symmetry means that if we "re-label" college as non-college and non-college as college, the resulted index should not change.

2.26.

Table 2.26: Explore the Perfect PAM property of the Absolute Difference Index

Wife-Husband Type	T1: perfect assortativeness	T2
C-C(k_1)	0.02	0.2
C-NC(k_2)	0	0.05
NC-C(k_3)	0	0.05
NC-NC(k_4)	0.98	0.7
Absolute Diff.	0.0196	0.1375

T1 and T2 represent two matching contingency tables. T1 exhibits perfect assortativeness, while the matching in T2 is not perfect but the college ratio is much higher than that in T1.

Chapter 3: (Mis)use of Power in the Ivory Tower: Evidence from Deans in Chinese Universities

3.1 Introduction

Universities produce fundamental knowledge that can spur long term economic growth (Cantoni and Yuchtman, 2014; Hausman, 2012). A key factor in the rise of high-tech sectors is breakthroughs originating at first-rate universities, such as innovations in information technology, biotechnology and biomedicine, and aerospace. The incentives that motivate university researchers, however, differ from those that are available to inventors in the private sector. The literature increasingly focuses on academic researchers' incentives to generate innovations (Hvide and Jones, 2018), but protection of researchers' intellectual property has received little attention. Academia is not a utopia that stands apart from the influence of power (Jia et al., 2019) . Can university leaders (mis)use power for their own benefits at the cost of other researchers' intellectual property? We assess this issue by providing new evidence that university leaders benefit from of the power the wield and discuss how they benefit.

This chapter shows that administrative power enhances the intellectual property rights (as measured by patent applications) that school deans enjoy at top Chinese universities. The deans are powerful administrative leaders and they control personnel and critical research resources. We construct a new dataset consisting of detailed biographical information about deans. This new dataset enables us to investigate the link between power and innovation by comparing deans' patent-application records before and after taking office. We use an event-study design to estimate the dynamics of patenting activities before and after dean appointments.

We start by documenting a strong "deanship effect" on patenting: Patent applications increase immediately following a professor's promotion to a deanship. We see no trends in patenting prior to the professor's ascension to dean. This positive deanship-patenting association (hereafter "deanship effects") could be driven by any of several channels, which have different welfare implications. Plausible candidate channels include: (1) positive selection into a deanship based on ability: a researcher is appointed as a dean because of an increase in research productivity; (2) a deanship opens access to previously unavailable research resources; (3) misuse of political power: a dean could request authorship or coauthorship of a patent or be invited to join a patent as a favor in light of her political power. We perform additional heterogeneous-effects analyses to identify the source of the deanship effects. Moreover, by exploiting the recent anti-corruption campaign in China as a natural experiment, we provide causal evidence that misuse of political power is an important channel for explaining the deanship effects.

In this chapter, we use the term "misuse of power" to refer to the potential issue of deans being listed on more patent applications than they would if they did not hold the political power of the deanship. We consider this effect of power as "misuse" for two reasons. First, as a general principle, researchers should not claim intellectual property ownership of work to which they have not contributed effort or resources. Second, deans in China's higher education system are very powerful due to their control over resources and personnel. If they obtain extra patents *because of* their power, these benefits are very likely to result from rent-seeking behaviors that potentially distort the incentives and innovation activities of local researchers. Moreover, such behaviors are inconsistent with the duties of deans in the academic institution. However, due to data limitations, we cannot observe whether the "misuse of power" has distortionary effects on innovation. Therefore, the focus of this chapter is on whether the power of deans directly results in more patents.

We present five main findings from our exploration of the sources of deanship effects. First, the increase in patent applications filed by deans occurs in research fields in which deans were not working prior to their appointments. If we consider patent applications in fields in which deans had filed more than five years before taking office, we observe fewer patents following deanship appointments.[1] Divergence between total patents and patents in fields with which deans were familiar prior to their appointments suggests that the increase in patenting cannot be attributed to deans' own research efforts. Second, the deanship effects we find are on average larger for deans that experience paper retractions. We treat paper retractions as a correlate of weak academic ethics, which implies that deanship effects are less likely driven by positive selection or more access to research resource.

Third, the deanship effects we observe are smaller in higher-quality universities and in those whose locations include more neighboring universities. Professors working in these universities presumably enjoy more outside labor-market options as a result of either greater vertical mobility (for high-quality universities) or greater horizontal mobility (with more neighboring universities). The outside options would reduce deans' bargaining power over other researchers. These results are consistent with the misuse-of-power channel. Fourth, patent applications fall after a dean loses power, but do not change if the dean is further promoted. These results are again consistent with the role of the power of administrative heads. These two findings cannot, however, rule out resources as another channel. It could be that resources attached to deanships are more important in lower-quality universities. The effects of outside option and losing a deanship can also be explained by the resource channel of deanship. In our fifth exercise, we provide direct evidence of the (non-)association between deanships and resources by showing that a deanship is not followed by more funding from the National Natural Science Foundation of China.

The various pieces of evidence mentioned above suggest that the misuse of political power rather than superior ability or easier access to resources better explains our findings. We next provide evidence from another angle: if deanship effects reflect the misuse of power, then political

[1] Every patent falls under a particular international patent classification (IPC). We define a patent's field using the first two digits of its IPC code.

efforts to purge a university's bureaucracy would mitigate the deanship effects. By exploiting the recent anti-corruption campaign in China as a natural experiment, we show that the anti-corruption efforts have indeed reduced the deanship effects. Moreover, the anti-corruption campaign has not changed the association between deanship and patenting in a dean's previous field of expertise. These findings cannot be readily explained by channels other than the misuse of political power.

Finally, we discuss the spillover effects of power on resource allocation. We document that, after deans take office, their previous co-inventors receive more funding from the National Natural Science Foundation of China (NSFC). This result provides supportive evidence that deans exert considerable influence over the allocation of resources that are crucial to local researchers. By exerting their power, deans could seek benefits for themselves via favors or political influence. Others who are connected to deans could also benefit from a dean's power. These findings imply that a dean's power could distort resource allocation and thereby also affect innovation activities in academia.

This chapter contributes to the literature on scientific production and innovation by documenting the role of political power in academic institutions. Previous research has demonstrated the importance of material incentives (Borjas and Doran, 2015b; Hvide and Jones, 2018; Ytsma, 2019) as well as cooperation and spillover effects (Azoulay et al., 2010; Borjas and Doran, 2015a; Borjas et al., 2018; Jaravel et al., 2018). This chapter shows that political power per se plays a role in driving innovation output, while further analyses suggest that this reflects the misuse of power. One important study that first provides a similar political-economy perspective on knowledge production is Jia et al. (2019). They document that promotion to a deanship is associated with an increase in journal publications for economists at elite Chinese universities. Our study further shows that the power of a deanship matters also for patenting. Patents, arguably, play a more substantial and direct role in knowledge production that drives economic growth. Moreover, given the economic benefits associated with patents, it is much more likely that the deanship effects on patents imply more than merely having names added to patents.

Broadly speaking, our study is related to the literature on misallocation of resources (Hsieh and Klenow, 2009; Hsieh et al., 2019), with a focus on knowledge production in academia. The allocation of talent has attracted attention since Murphy et al. (1991) because of its role in economic growth. Without an efficient market in which to allocate resources, connections or political power may easily distort the allocation of resources for research. Scientists in China have long criticized the part that connections with powerful people play in allocating research funding (Shi and Rao, 2010). Fisman et al. (2018) document that connections and favoritism distort the allocation of scientific talent in China. We show that power generates patents and funding allocation through connections. Because the productivity of the deans themselves is not likely to explain this pattern, it is more likely that political power has distorted innovation activities in our context.

Recently, the issue of academic bullying has received increasing attention in the scientific community. The hierarchical structure of academic institutions enables bullying behaviors by powerful individuals (Else, 2018). Bullying in academia, including violations of intellectual property (Gewin, 2021; Mahmoudi, 2019), negatively impacts researchers' welfare and incentives and impedes scientific advances. As a result of data limitations, we don't know that benefits to deans are not from the dean offering carrots rather than sticks or even deans offering researchers nothing but researchers asking deans to put their name on the patent in hope of future benefits. Nevertheless, our findings point to how power potentially distorts knowledge production in academia, echoing the growing concern with academic bullying.

Our study also contributes to the literature that examines power and internal politics in organizations. Power has been identified as a cause of internal politics in firms (Prendergast, 1993; Prendergast and Topel, 1996; Tirole, 1986) and government corruption (Banerjee et al., 2013). Theoretical models have considered various ways in which internal politics affects organizations, including information collection (Prendergast, 1993) and even collusion (Tirole, 1986). Empirical evidence for these issues has been scarce. This chapter provides suggestive evidence that power in a hierarchy could impact the key function of academic organizations perhaps through voluntary or involuntary collusion.[2]

This chapter is organized as follows. In Section 3.2 we describe the institutional setting and the data. In Section 3.3 we introduce our empirical strategy. In Section 3.4 we present results pertaining to deanship effects and various heterogeneity analyses that shed light on the channels through which the power works. In Section 3.5 we present results pertaining to the impacts of the anti-corruption campaign on the deanship effects. Additional evidence on the spillover effects of deanships on local researchers is presented in Section 3.6. Section 3.7 concludes.

3.2 Data and Institutional Backgrounds

3.2.1 Deans in China's Universities

Deans of schools in Chinese universities work as administrative leaders as well as public officials in the country's bureaucratic system. A school's dean is designated by the university administration and promoted within the school. A dean exercises considerable control over researchers and resources via several channels including: (1) selection to and promotion of researchers within the school, and (2) allocating eligibility for subsidies, grants, and awards provided by the government, which normally consist of major resources to support research. Deans have the final say on personnel issues, including selection, evaluation, and promotion of researchers. These personnel decisions in turn greatly affect wages and benefits for the faculty. Some deans might make appointments on the basis of favoritism rather than merit (Yang, 2005). In addition to exercising direct personnel control, the dean can also exert influence on other resources that are available to researchers in the school. For example, both national-level grants and subsidies from the central and local governments require recommendation by the dean.[3] The dean therefore plays a key role

[2]Our study also yields general implications for broader contexts. The issue of managerial supervisors (bosses) taking credit for their subordinates' work has received considerable attention in the popular press, while little theoretical or empirical research has addressed this problem (Smith, 2013; Shellenbarger, 2019).

[3]The major national funding organizations, the National Natural Science Foundation of China and the National Social Science Fund of China accept applications only via a school. The school evaluates and allocates opportunities

in determining researchers' eligibility for crucial resources.

3.2.2 Importance of Innovation in Chinese Universities

Universities are among the major contributors to China's rapid growth in innovation. In Figure 3.1 we show that university patent applications grew from about 5,000 in 2000 to 369,000 in 2017. In 2017, university patent applications accounted for around 11% of total patents in China and 16% of the more valuable invention patents. All of these innovative activities are accomplished at 1,243 universities. We focus on the deans of schools that specialize in research fields that represent the science, technology, engineering, and math (STEM) domain at all Project 211 universities in China. Project 211 was initiated by the Ministry of Education of China in 1995. The intention was to raise the research standards of high-level universities.[4] As shown in Figure 3.1, Project 211 universities carry out more than one-third of China's university innovation activities, especially for invention patents. Another reason we focus on these universities is that we need to collect biographical information on deans, which is missing for most deans at non-Project-211 universities.

3.2.3 The Anti-corruption Campaign

A large scale anti-corruption campaign was launched by the Chinese Communist Party in 2013. The campaign started by sending inspection teams to lower level governments. Many officials at various hierarchical levels of government ("tigers" and "flies") were investigated. Although speculation about the ultimate goal of this anti-corruption campaign has circulated, it is perceived as the largest and most intensive in the PRC's history. Many individuals and institutions, including but not limited to the government sector, have been affected. The campaign has significantly reduced rent-seeking behaviors by local officials and politically-connected firms (Chen and Kung, 2019a).

for researchers within its purview.

[4]By 2006, 116 universities were "211" universities, which received 70% of scientific research funding from the government. Over 10 billion yuan has been invested in the "211 Project" (see http://en.people.cn/90001/6381319.html).

State-owned and private firms have also been impacted by the anti-corruption campaign in various ways (Ding et al., 2020; Kong and Qin, 2019; Kong et al., 2020; Xu and Yano, 2017a). University officials, as part of the bureaucratic system, might also be held accountable for corruption. Many anti-corruption investigations are actually targeted at bureaucrats and professors in universities.[5] It seems likely that the anti-corruption campaign would affect deans' behaviors, especially if they are guilty of misconduct.

3.2.4 Data and Sample Construction

Deans and Patents

We combined multiple data sources for our analysis. We hand-collected biographical information on incumbent (as of July 2019) and previous deans of STEM schools in all Project 211 universities from their online resumes and biographies. Key information about a dean includes the date of taking office, educational experience, whether the dean has retired and the date of retirement.

We collected data for patent applications from the universe of patent applications obtained from the National Intellectual Property Administration (NIPA) before 2019. The patent records include patent name, patent number, patent type, application date, the names of inventors, the name of the institution owning the patent, the International Patent Classification (IPC) code, and other patent characteristics including citations and the number of claims. There are three types of patents: invention, utility, and design. Invention patents are more likely to indicate substantial innovation activities. In our analysis, we focus on the number of applications for both total patents and invention patents. Because it can take up to 18 months before the NIPA publishes records of a patent application, we drop patent records after 2017 to avoid issues with missing data. We consider the number of (1) total patent applications in each year and (2) invention patent applications in

[5] https://www.chronicle.com/article/government-targets-chinese-universities-in-anti-corruption-drive/

each year. The latter account for the majority of all patent applications. Because there are many zeros in these outcome variables (about 57%), we use two sets of dependent variables for total and invention patents: (1) the inverse hyperbolic sine transformation of the number of applications $IHS(y) = ln(y + \sqrt{1+y^2})$ (Bellemare and Wichman, 2020); and (2) a dummy variable for having filed non-zero applications during a given year, which captures the extensive margin.

We focus on a sample period that runs from 2006 through 2016, mainly because of data availability: it is hard to find detailed information for deans in earlier years. We also wanted to avoid confounding effects: the State Council of China introduced the National Medium- and Long-Term Program for Science and Technology Development, an ambitious plan to bolster innovation through for example, subsidizing patent applications and providing more research funding to universities, in 2006. Therefore, our final sample starts in 2006. We chose 2016 as the end-point because one of our measures of anti-corruption efforts, the Tencent corruption investigations database, extends only through 2016.[6] For our main analysis, we consider deans who had not retired by 2016. In our further analysis regarding the effects of leaving a deanship, we look at deans who were already incumbent deans as of 2006. To test whether there were spillover effects of a dean's power, we also construct a sample of local researchers as well as their co-inventorship with deans. In Appendix 3.8.1, we describe in detail the construction of our datasets. Descriptive statistics for these deans are reported in Table 3.1.

Other Information

We use the distribution of funding from the National Natural Science Foundation of China (NSFC) as our main measure of research-related resource allocation. The NSFC is affiliated with China's State Council for the purpose of managing the National Natural Science Fund. The NSFC

[6]Other sample period options, such as 2005–2017, 2007–2015, etc, lead to very similar results.

is one of the major sources of funding for scientific research in China.[7] We scraped detailed information pertaining to all the projects funded by the NSFC from its official website.[8] The projects and funds are then matched to deans and other local researchers.

We use paper-retraction records to measure academic professionalism. The information is provide by the Retraction Watch Database. The database collects retractions of scientific articles published across all fields.[9] As academic misconduct has become an increasing concern for the scientific community, this databases represents a very useful source for studying relevant issues (Oransky, 2020; Zuckerman, 2020). We collect information on retracted papers from Retraction Watch, which was matched to deans in our sample. Among the deans in our sample, 8% had at least one paper retracted.

Anti-corruption Efforts

To further explore how the anti-corruption campaign affected deanship effects, we construct several measures of anti-corruption efforts. We collect data on anti-corruption investigations released by the Central Commission for Discipline Inspection (CCDI) of China.[10] The CCDI, as the internal control and supervision institution of the Chinese Communist Party (CCP), is responsible for combating corruption in the party. It has led the anti-corruption campaign as most government officials are also members of the CCP. The CCDI has been publishing records of investigations into government officials since the onset of the anti-corruption campaign. These data enable us to capture anti-corruption investigations in every prefecture. We compute the cumulative number of investigations into local officials at the prefectural (sub-provincial) level since the onset of the

[7] From 2006 to 2016, 167.6 billion *yuan* was provided by the NSFC to support researchers in academic institutions (https://www.nsfc.gov.cn/publish/portal0/tab505/).

[8] https://isisn.nsfc.gov.cn/egrantindex/funcindex/prjsearch-list

[9] See http://retractiondatabase.org/RetractionSearch.aspx for more details.

[10] See http://www.ccdi.gov.cn/scdc/ for the released investigation records. Because records of investigations conducted in earlier years have been removed from the website, we use Internet Archive to collect all the released investigations since the start of the anti-corruption campaign. http://web.archive.org/web/*/http://www.ccdi.gov.cn/

anti-corruption campaign as our measure of local anti-corruption efforts.

We also use two alternative measures of anti-corruption to test the robustness of our results. The first is a dummy variable for "central inspection", which takes the value of 1 if a central inspection team had previously been sent to the province where the university is located Chen and Kung (2019a). Since 2013, 11 waves of inspection teams have been sent to randomly selected provinces. Each inspection lasted for approximately two months. The inspection teams talked to local government officials, received complaints from residents, and then provided feedback to the central government. The second measure is based on anti-corruption investigations published by Tencent, a large Chinese technology firm. Tencent launched this comprehensive online database in collaboration with the Chinese government to inform propaganda for the anti-corruption campaign. The original database is currently not accessible, so we use the data scraped by Wang and Dickson (2022) in August 2016.[11] See Wang and Dickson (2022) for more details about the dataset.

Figure 3.2 displays the cumulative national cases from the Tencent database and the CCDI database. The difference between these two databases lies in their coverage. The CCDI publishes records of investigations only into government officials who have been appointed and administered by the Central Organization Department or the Provincial Organization Department, while the Tencent database covers investigations into all levels of government officials.[12] The CCDI database provides data that pertain more directly to our sample, as all the Project 211 universities fall under the control of either a provincial government or the central government. The correlation between the measures based on the CCDI releases and those based on the Tencent database is 0.84.[13]

[11] See https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/9QZRAD for the dataset.

[12] In local Chinese governments, the personnel in certain key positions are under the control of the Provincial Organization Department or even the Central Organization Department. The remaining positions are under the control of the Organization Department in each local government (prefecture or county).

[13] For the Tencent database, we include only investigations into officials of a certain hierarchical level (*Chuji*) because they are more directly relevant. Including all levels of government officials delivers similar conclusions. The CCDI database by default contains only government officials above the *Chuji* level.

Table 3.1: Summary Statistics

Panel A. Patents of Deans					
		All Patents		Invention Patents	
	Obs.	Mean	SD	Mean	SD
Baseline Analysis Sample					
Patent Applications	9878	2.941	7.675	2.507	6.453
Patent Applications > 0	9878	0.425	0.494	0.412	0.492
Leaving Deanship Sample					
Patent Applications	2431	3.896	7.891	3.325	6.568
Patent Applications > 0	2431	0.558	0.497	0.539	0.499

Panel B. Individual Characteristics			
	Obs.	Mean	SD
Baseline Analysis Sample			
Year of Taking Office	898	2013.2	3.96
Year of Graduation	898	1996.1	7.70
Having Paper Retraction Experience	898	0.080	0.27
Leaving Deanship Sample			
Year of Retirement from Deanship	221	2014.4	2.07
Year of Graduation	221	1993.3	5.43
Losing Power after Retirement	221	0.63	0.48

Panel C. University Characteristics			
	Obs.	Mean	SD
Project 985 University*	97	0.38	0.49
# of Universities in the Same City	97	23.24	12.35

Panel D. Measures for Local Anticorruption Efforts (2013-2016)**					
		CCDI Data		Tencent Data	
	Obs.	Mean	SD	Mean	SD
Cumulative Investigations into Local Officials	156	10.15	11.78	15.01	22.55
Cumulative Investigations into University Officials	156	0.67	1.36	1.31	2.27
IHS(Cumulative Investigations into Local Officials)***	156	2.20	1.49	2.48	1.49
IHS(Cumulative Investigations into University Officials)	156	0.42	0.70	0.71	0.88

Notes: * Project 985 universities are the high-quality universities identified by the Ministry of Education of China. ** By construction, measures of anti-corruption efforts before 2013 are zero. Investigations into local officials reflect the number of reported anti-corruption cases in each year in the city where the university is located, as found in the corresponding dataset (Tecent or CCDI). Only government officials who were above certain hierarchical level (*Chuji*) are counted. Investigations into university officials reflect to reported cases in which university officials were investigated. *** Inverse hyperbolic transformation: $IHS(y) = ln(y + \sqrt{1+y^2})$.

Figure 3.1: Time Trends of Patent Applications

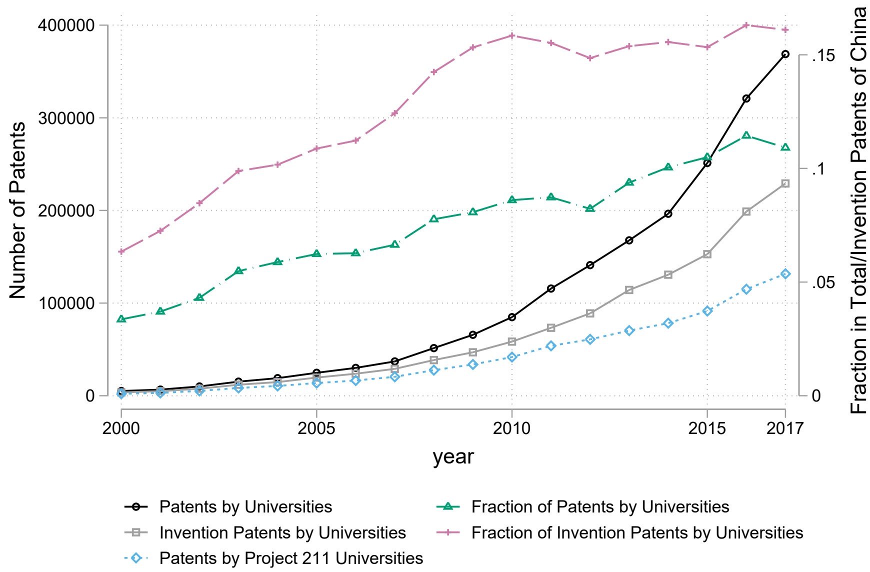

Notes: Calculated from China Patent Database.

Figure 3.2: National Anti-corruption Investigations Over Time

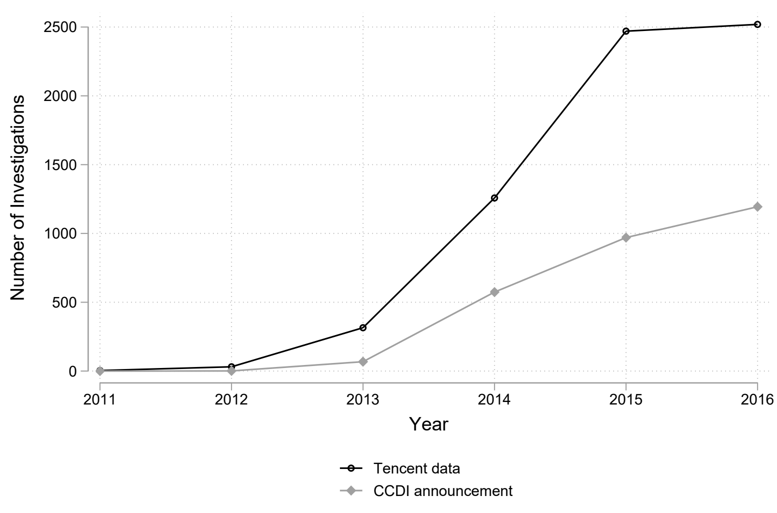

Notes: Reports of anti-corruption investigations from Tencent are obtained via https://news.qq.com/zt2016/fanfu_ccdi/index.htm. Announcements of anti-corruption investigations from the Central Commission for Discipline Inspection (CCDI) are obtained via http://www.ccdi.gov.cn/scdc/. Only cases involving public sector officials who had achieved the hierarchical level of *Chuji* and above are included.

3.3 Estimation Strategy

3.3.1 Deanship Effects

In the baseline model, we estimate the effects of deanships on patent applications using the following event-study specification while controlling for both individual and time fixed effects.

$$y_{it} = \sum_{k=-5}^{-2} \delta_k Dean_{it}^k + \sum_{k=0}^{5} \delta_k Dean_{it}^k + \phi_1 exp_{it} + \phi_2 exp_{it}^2 + \mu_i + \eta_t + \varepsilon_{it} \tag{3.1}$$

where k denotes k years relative to the year of taking office. $Dean_{it}^k$ is a dummy variable indicating that it has been k years since individual i took office as a dean. We stack more than 5 years before becoming dean as k = -5, and more than 5 years after becoming dean as k = 5. μ_i is individual fixed effects, and η_t stands for year fixed effects. We use as dependent variables the inverse hyperbolic sine transformation of patent applications by each individual in each year, as well as a dummy for having nonzero patent applications. exp_{it} is years of potential experience, which equals to t minus the year of graduation. Standard errors are clustered at the individual level. Not controlling for potential experience or clustering at the university level has minimal effects on our estimates and does not qualitatively change our findings.

OLS estimates for the two-way fixed effects model might be biased if deanship effects are heterogeneous for deans that took office in different years (Abraham and Sun, 2021; Callaway and Sant'Anna, 2021). To address this concern, we use the interaction-weighted estimator proposed in Abraham and Sun (2021) for our event study estimates.

We also examine the association between deanships and patent quality. In those models, we run the regression at the patent level where the dependent variables are measures of patent quality, such as patent citations. The independent variables are the same as in Equation 3.1. We still cluster standard errors at the individual dean level.

3.3.2 Anti-corruption and Deanship Effects

In the second part of our analysis, we explore whether the positive deanship effects we observe have been dampened by the recent anti-corruption campaign. We implement a triple-differences (DDD) design to test this hypothesis, exploiting variations in local anti-corruption efforts.

We use the number of anti-corruption investigation cases as a measure of prefecture-level anti-corruption efforts. Specifically, we measure local anti-corruption efforts using the total number investigations of government officials in each prefecture c in year t.[14] In our main specification, we use the cumulative number of cases since 2013 to capture the cumulative impacts of the anti-corruption campaign.

Denote $CuInvestigation_{ct}$ as the (inverse hyperbolic sine of the) cumulative number of investigations in prefecture c in year t. We estimate the following DDD specification:

$$\begin{aligned} y_{\text{it}} = {} & \tilde{\beta}_1 * Dean_{it} + \tilde{\beta}_2 * Dean_{it} * CuInvestigation_{ct} + \phi_1 exp_{it} + \phi_2 exp_{it}^2 \\ & + \mu_i + \eta_t + \mu_i * CuInvestigation_{ct} + \eta_t * CuInvestigation_{ct} + \varepsilon_{it} \end{aligned} \tag{3.2}$$

We allow the impacts of the anti-corruption efforts to vary with individual invariant characteristics and time effects, which controls for unobserved factors in a very flexible way. The coefficient of interest is the coefficient on the interaction term between the deanship dummy and the variable for local anti-corruption efforts, which captures the extent to which anti-corruption efforts reduced the deanship effects on patenting.

The data for anti-corruption investigations come from reports by the China Central Commission for Discipline Inspection. We also carry out robustness checks using anti-corruption data collected from the alternatives sources discussed in Section 3.2.4.

To check whether the triple-differences exercise indeed captures the impacts of the anti-corruption

[14] We count only cases involving officials who achieve at least the least *Chuji* level. The results are similar when we include government officials from all levels.

campaign on universities, we also restrict our measure for anti-corruption efforts to include only investigations into university officials. Counting these university-related investigations, we construct a similar variable for the cumulative number of investigations in the same city since the start of the campaign.

3.4 Empirical Results: Deanship Effects

In this section, we discuss our empirical results regarding the deanship effects on patenting. We present several heterogeneity analyses that shed light on the mechanisms that might drive the positive association between deanship and patenting.

3.4.1 Baseline Results: the Impact of Deanship on Patent Applications

We first present the baseline findings regarding the deanship effects. One threat to our empirical estimation is the potential bias in two-way fixed effects models in the presence of cohort-of-treatment heterogeneous treatment effects. As Abraham and Sun (2021) have pointed out, when there is substantial heterogeneity in treatment effects across cohorts defined by the time of taking up the treatment (in our case, the year of promotion to a deanship), the estimated coefficients of the event-study design could be contaminated by negative weights on some cohort-specific treatment effects or picking up treatment effects from other relative time periods. This might be a concern for us if the deans promoted in the early years differ from those promoted in more recent years. The deanship effects could also differ across cohorts if deans are more powerful early on than in later years due to the anti-corruption campaign, or more powerful in later years as a result of control over more resources. The potential correlation between the types of individuals promoted and the timing of promotion raises the concern of heterogeneous treatment effects. To address this concern, we employ the Abraham and Sun interaction-weighted estimator in our event study estimators.

We report the baseline results using Equation (1) and Equation (2) in Table 3.2 and Figure 3.3, respectively. We find a significant jump in patent applications by deans immediately after they take office at both the intensive and the extensive margins. Based on results reported in column (1) of Table 3.2, the total number of patent applications increased by 7.5% in the year of taking office, and 13% in the year after taking office. The 7.5% rise in the year of taking office reflects deanship for only a partial year. Therefore, the jump in patent applications upon promotion should have been even sharper for a full year under a deanship. We find similar patterns at the extensive margin. The probability that patent applications are filed increased by 5% in the year of taking office, and rises to and stabilizes at around 9% above the level of filings before taking office. On average, a deanship is associated with a 14% increase in patent applications, and an 8% higher probability of filing any patent applications. In column (3) and column (4) of Table 3.2, we also report the results for invention patents only, which occupy the majority of total patent applications and are more directly relevant to innovation. The deanship effects on invention patents are very similar to those indicated by results for total patents.

Table 3.2: Deanships and Patent Applications: Event Study Results

	(1)	(2)	(3)	(4)
	IHS(Patents)	Patents>0	IHS(Invention Patents)	Invention Patents>0
5+ years before	-0.028	-0.026	-0.018	-0.009
	(0.079)	(0.030)	(0.076)	(0.031)
4 years before	-0.047	-0.009	-0.036	0.001
	(0.051)	(0.022)	(0.049)	(0.022)
3 years before	-0.024	-0.008	-0.014	0.004
	(0.041)	(0.019)	(0.039)	(0.019)
2 years before	-0.036	-0.008	-0.035	-0.002
	(0.033)	(0.016)	(0.031)	(0.016)
Year of promotion	0.075^{**}	0.051^{***}	0.047	0.039^{**}
	(0.037)	(0.018)	(0.036)	(0.018)
1 year after	0.134^{***}	0.094^{***}	0.122^{**}	0.092^{***}
	(0.050)	(0.023)	(0.048)	(0.023)
2 years after	0.171^{***}	0.081^{***}	0.148^{***}	0.077^{***}
	(0.060)	(0.026)	(0.056)	(0.026)
3 years after	0.192^{***}	0.089^{***}	0.171^{**}	0.080^{***}
	(0.072)	(0.030)	(0.068)	(0.030)
4 years after	0.225^{***}	0.117^{***}	0.215^{***}	0.114^{***}
	(0.084)	(0.035)	(0.079)	(0.035)
5+ years after	0.223^{***}	0.067^{***}	0.202^{***}	0.064^{***}
	(0.064)	(0.024)	(0.062)	(0.024)
Average deanship effects				
	0.142^{***}	0.080^{***}	0.116^{***}	0.070^{***}
	(0.038)	(0.016)	(0.037)	(0.016)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dep. Var. Mean	$2.941^{\dagger}$	0.425	$2.507^{\dagger}$	0.412

Notes: All regressions control for a quadratic polynomial of potential experience. Event study coefficients are estimated using the interaction-weighted estimator proposed by Abraham and Sun (2021). Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, (4) a dummy for having positive invention patent applications.Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

Figure 3.3: Event-study Estimates for Dynamic Deanship Effects

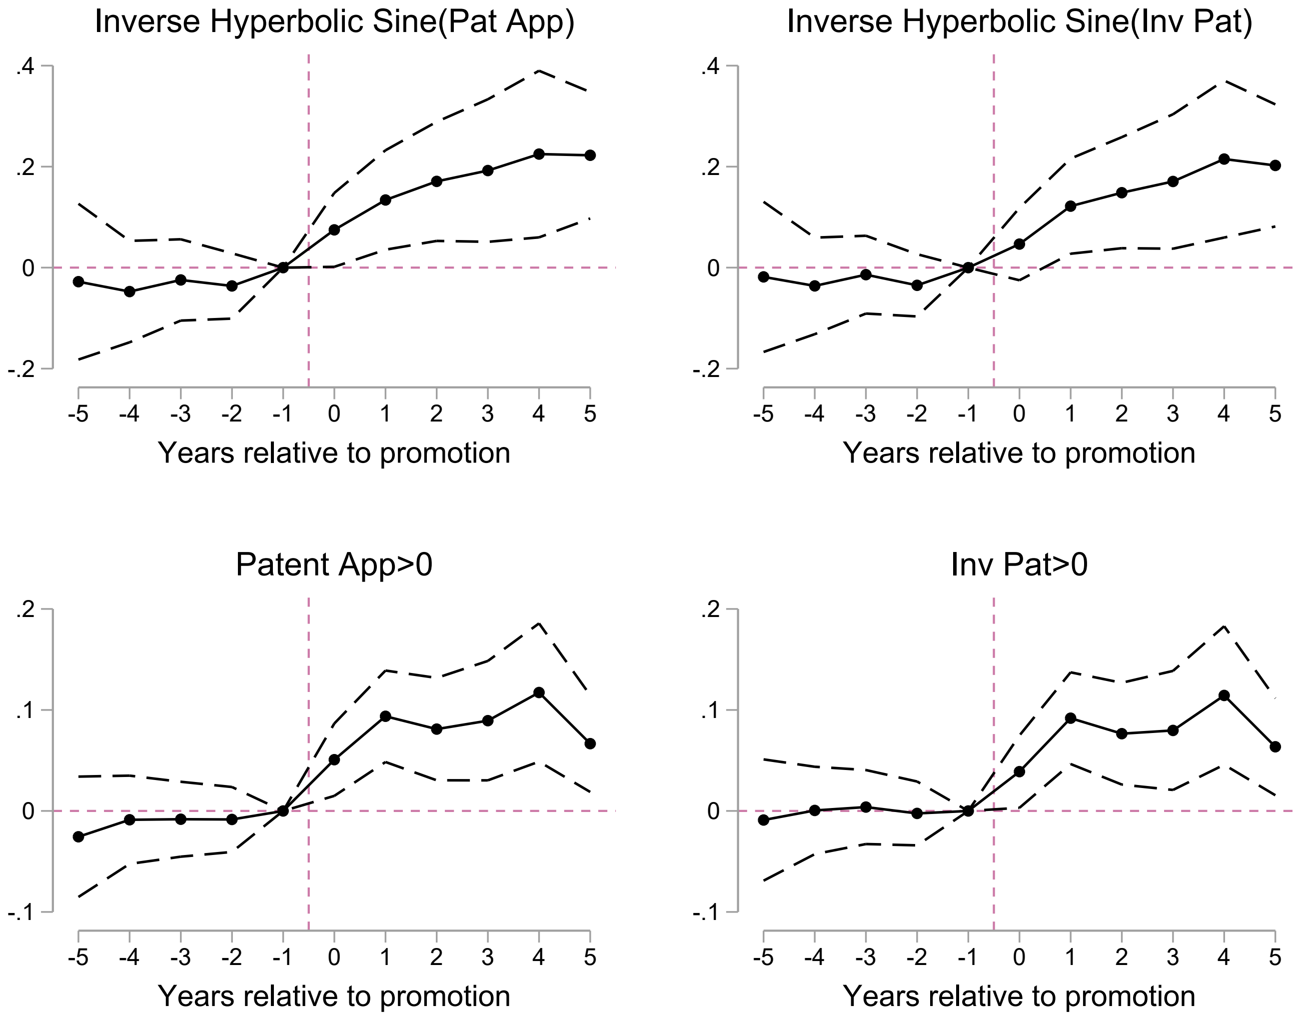

Notes: Event study estimates using the Abraham and Sun (2021) estimator and 95 percent confidence intervals are plotted. Regressions control for individual fixed effects, year fixed effects, and the quadratic polynomial of potential experience. The dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) the inverse hyperbolic sine of invention patent applications, (3) a dummy for having positive patent applications, and (4) a dummy for having positive invention patent applications. Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5.

In appendix Figure 3.12 and appendix Table 3.16 we report our baseline estimates for the dynamic deanship effects using the traditional OLS two-way fixed-effects estimator. The general patterns and point estimates are very similar to our baseline estimates using the Abraham and Sun interaction-weighted estimator. These results eliminate most of our concerns regarding bias that might be generated by cohort-heterogeneous deanship effects.

Patent Quality

If deans file for and are granted more patents, are they receiving better patents? We also compare the quality of patents before and after a professor becomes a dean. Measures of patenting quality that we consider include citation counts one, two and three years after a patent application as well as the total number of claims in the application. The estimated associations between patenting quality measures and deanship status are reported in Figure 3.4 and Table 3.3. We find no significant relationship between deanship and patent quality measures, suggesting that the positive association between deanships and patenting is not explained by the ability or resource channel. More productive researchers or researchers with access to more resources should produce higher-quality patents. Note that here we cannot rule out the ability or resource channel entirely because quality did not drop.

We also investigate the number of inventors as an informative feature of a patent. Ceteris paribus, a patent with fewer inventors implies higher average expected benefits to each inventor. In Figure 3.13 and Table 3.17, we report correlations between deanship and the number of listed inventors of patents. There is no systematic change in the number of inventors of patents associated with an individual after promotion to the deanship. In Figure 3.14, we further divide the patents into two categories: the number of inventors above or below median (5). The “effects” of deanships on patent applications do not seem to differ by the number of inventors. If anything, deanships have a larger effects on patents with fewer inventors. This pattern suggests that deans are not just putting their names on patents, which would presumably be easier for patents with more co-inventors. This argument, of course, is only suggestive because it hinges on the assumption that each inventor can benefit more from a patent with fewer co-inventors.

Table 3.3: Deanships and Patent Quality: Event Study Results

	(1)	(2)	(3)	(4)
	1 Year Citations	2 Year Citations	3 Year Citations	Number of Claims
5+ years before	0.009	0.082	0.077	0.254
	(0.033)	(0.066)	(0.103)	(0.270)
4 years before	0.014	0.039	0.017	0.196
	(0.030)	(0.055)	(0.086)	(0.208)
3 years before	0.085**	0.106*	0.145*	0.017
	(0.037)	(0.060)	(0.084)	(0.153)
2 years before	-0.003	0.012	0.006	0.142
	(0.023)	(0.036)	(0.049)	(0.127)
Year of promotion	-0.017	-0.013	-0.001	0.125
	(0.022)	(0.035)	(0.054)	(0.130)
1 year after	0.057**	0.079	0.161**	0.128
	(0.025)	(0.051)	(0.075)	(0.169)
2 years after	0.003	-0.008	0.055	0.090
	(0.027)	(0.057)	(0.094)	(0.232)
3 years after	0.012	-0.021	0.045	0.025
	(0.037)	(0.075)	(0.116)	(0.251)
4 years after	0.012	0.055	0.191	-0.123
	(0.042)	(0.086)	(0.129)	(0.274)
5+ years after	0.007	-0.033	0.018	0.070
	(0.029)	(0.055)	(0.088)	(0.239)
Observations	29052	29052	29052	29052
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dep. Var. Mean	0.220	0.575	1.011	5.927

Notes: All regressions control for quadratic polynomial of potential experience. Dependent variables are cumulative citations one, two, and three years after application, and the number of exclusive claims. Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$.

Figure 3.4: Deanships and Patent Quality

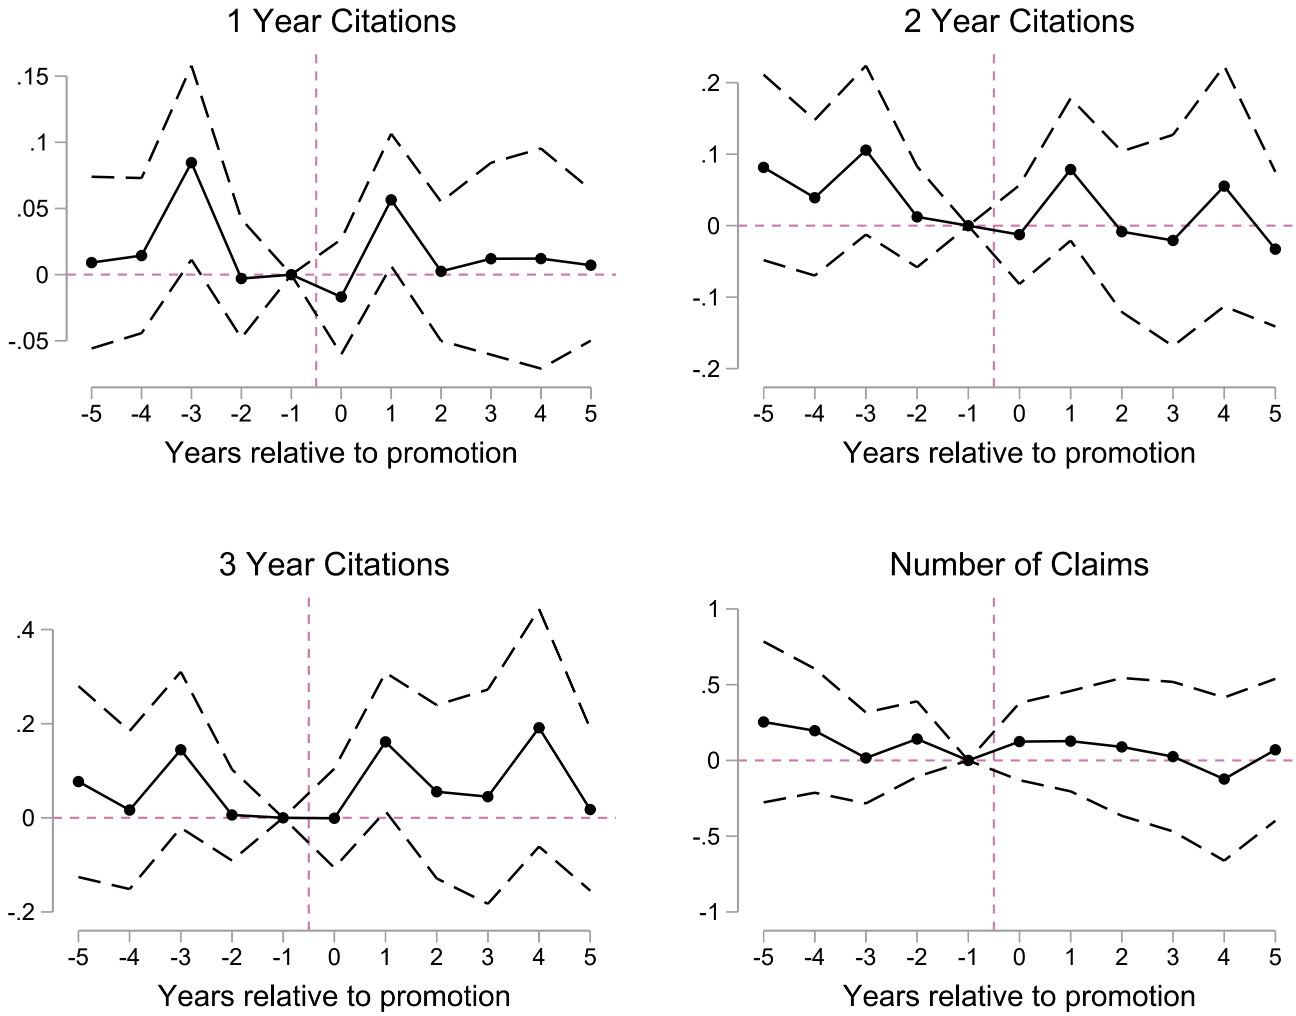

Notes: Event study estimates using the Abraham and Sun (2021) estimator and 95 percent confidence intervals are plotted. Regressions control for individual fixed effects, year fixed effects, and the quadratic polynomial of potential experience. Dependent variables are cumulative citations one, two, and three years after application, and the number of exclusive claims. Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5.

3.4.2 Granted Patents

We use patent applications as the main outcome in our analysis two reasons. Firstly, this outcome captures all the patenting activities by the dean, including patent applications that were not granted but had positive expected value at the time of application.[15] Secondly, it takes time for a patent to be granted, particularly for invention patents. According to a consultancy with the NIPA, it takes an average of about three years before an invention patent application is granted. Focusing on patent applications ensures that we can observe all the patenting behaviors by the deans in our sample period.[16] However, it is still reasonable to be concerned with whether patent applications can represent the actual benefits received by deans, as applications do not necessarily result in granted patents.

To address this concern, we also implement analysis using information on granted patents. First, we find that the ratio of successfully granted patents is high for deans in our sample. For patent applications by deans in our sample during the sample period, the average probability of being successfully granted a patent is 72% for all types of patents and 67% for invention patents. Figure 3.5 shows that the successful grant ratio of invention patents is much higher for deans than for the average inventor.[17] The grant ratio for deans is also relatively stable during the sample period. These results alleviate concerns that we are missing important information regarding granted patents by focusing on patent applications.

[15]For example, if a patent application encompasses cutting-edge technology and faces competition with other researchers, it may be scooped and fail to secure a granted patent. However, the patent application is still ex ante valuable.

[16]We observe the universe of patent applications and granted patents reported by the NIPA before 2019.

[17]The pattern seems to be reversed for all types of patents. The reason is that the fraction of invention patents in all applications is much higher for deans (86%) than for the average inventor (40%).

Figure 3.5: The Ratio of Granted Patents to Applications

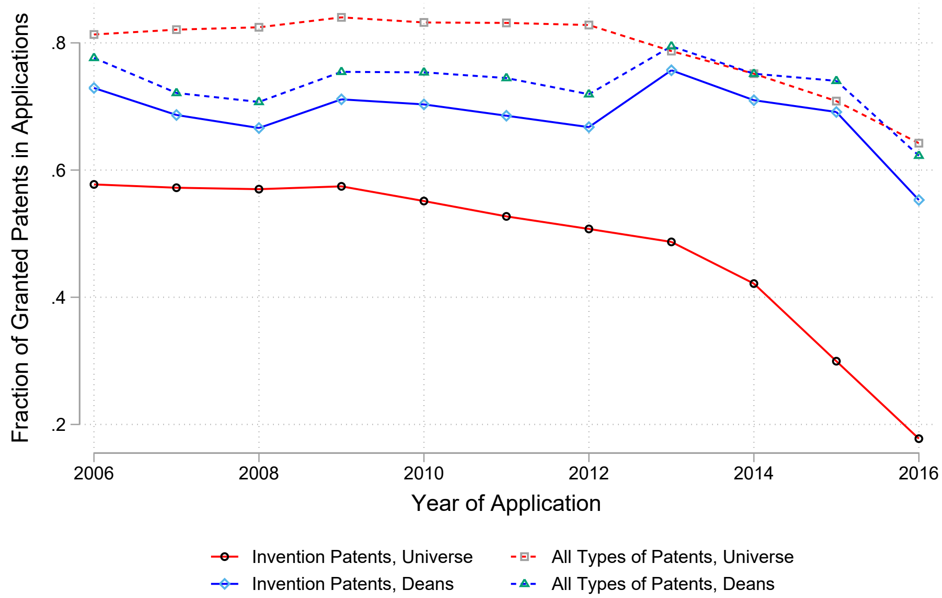

Notes: This figure plots the fraction of granted patents in patent applications by the year of application. The ratio is plotted by the type of applicants (deans vs. the universe of applicants) and the type of patents (invention patents vs. all the patent types). One mechanical reason for the decline in later years is that the grant process for invention patents usually takes several years.

Second, we also examine the association between deanship and granted patents instead of applications. Specifically, we re-estimate Equation 3.1 using granted patents as outcome variables. The event-study estimates are reported in Figure 3.6, which shows that deanships are indeed associated with more granted patents and a higher probability of possessing any granted patents. This pattern indicates that deanships lead to an increase in valuable granted patents, not just low-value patent applications that are less likely to be granted.

Figure 3.6: Event-study Estimates for Dynamic Deanship Effects on Granted Patents

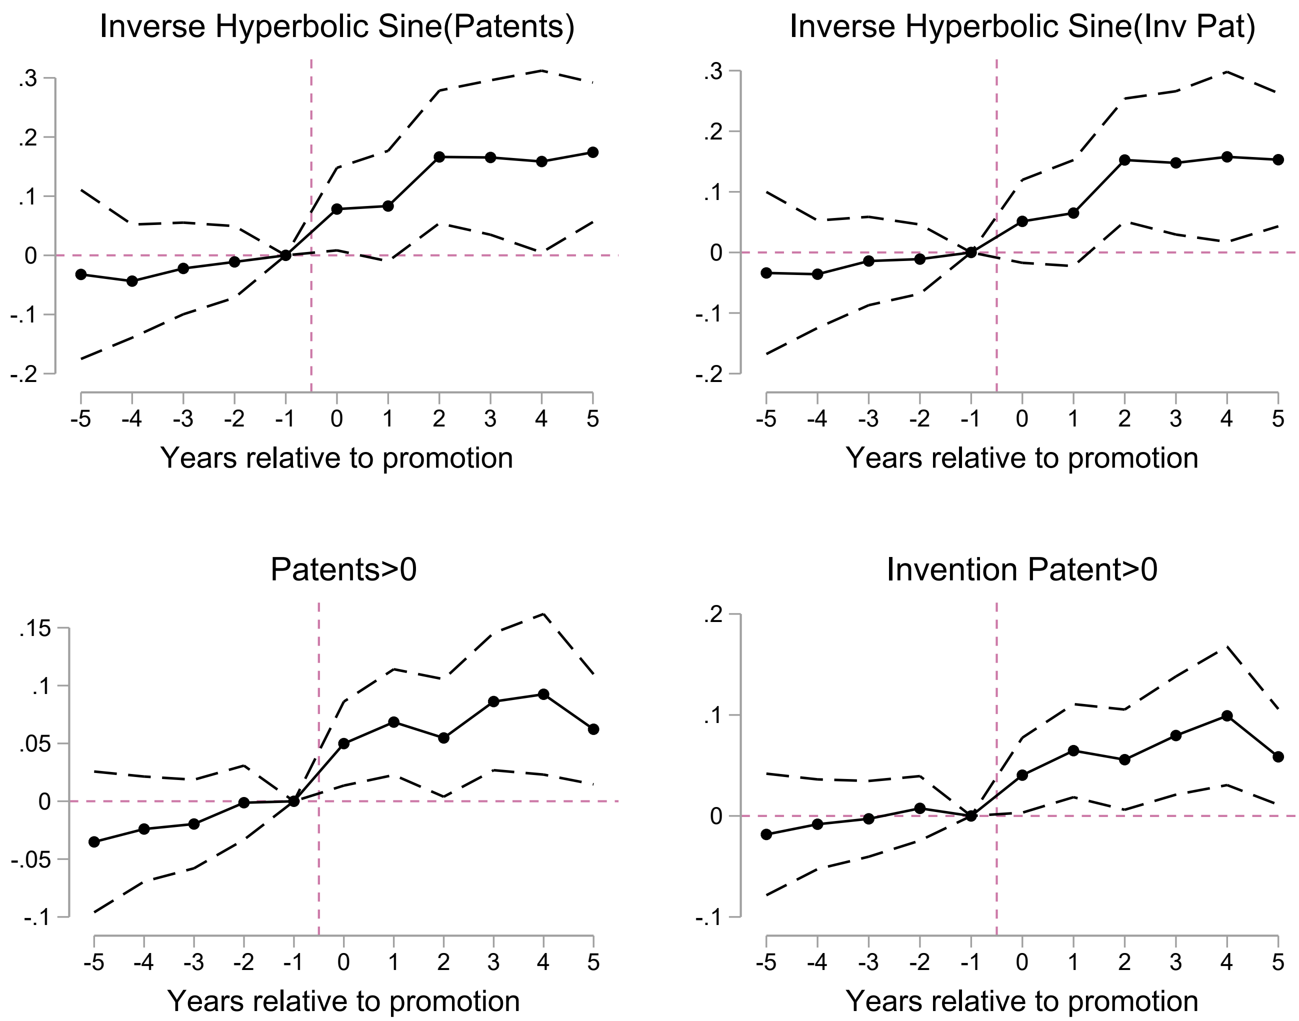

Notes: Event study estimates using the Abraham and Sun (2021) estimator and 95 percent confidence intervals are plotted. Regressions control for individual fixed effects, year fixed effects, and the quadratic polynomial of potential experience. The dependent variables are (1) the inverse hyperbolic sine of total granted patents, (2) the inverse hyperbolic sine of granted invention patents, (3) a dummy for having any granted patents, and (4) a dummy for having any granted invention patents. Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5.

Last, we also investigate the impact of deanships on the likelihood of patent applications being granted by re-estimating our baseline model using the ratio of granted patents to all patent applications as the outcome variable. Figure 3.7 reports the event study coefficients. The left panel uses our full sample of analysis, while the right panel uses only patent applications before 2013 to minimize the problem of missing patents that were not yet granted. Figure 3.7 illustrates that deanships did not have a significant impact on the likelihood of being granted patents. Specifically, as deans obtain more patent applications, there is no noticeable decrease in their probability of being granted patents conditional on applying for a patent.

Figure 3.7: Deanships and the Probability of Patents Being Granted

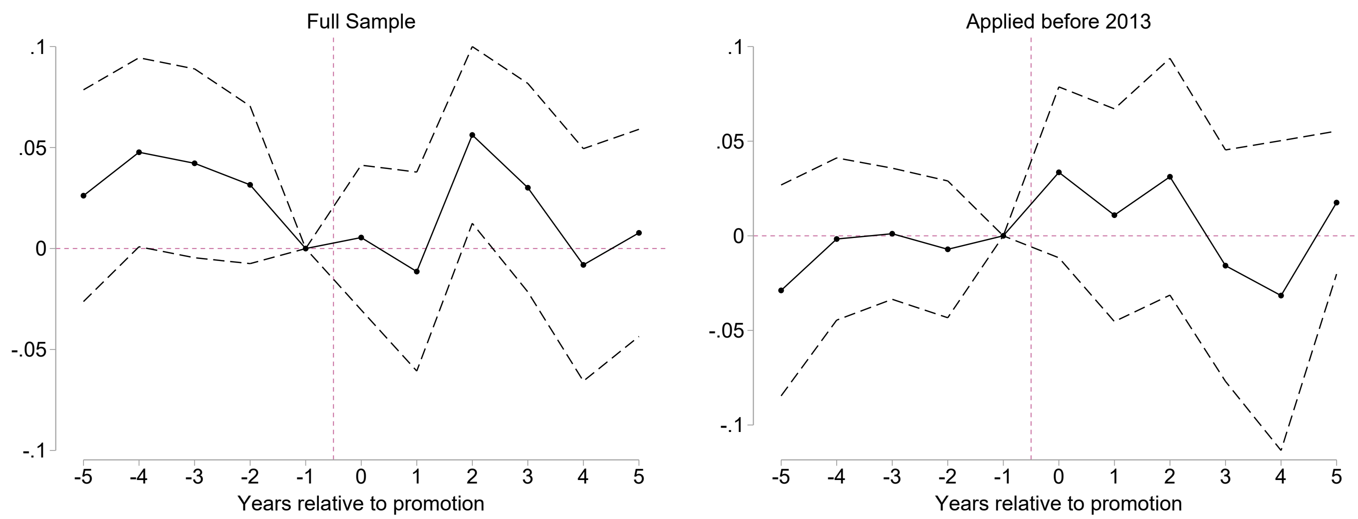

Notes: Event study estimates using the Abraham and Sun (2021) estimator and 95 percent confidence intervals are plotted. Regressions control for individual fixed effects, year fixed effects, and the quadratic polynomial of potential experience. Dependent variable is the ratio of granted patents to all applications. In the left panel, the regression is estimated using the full sample. In the right panel, the regression is estimated using patent applications before 2013. Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5.

3.4.3 Mechanisms That Might Drive Deanship Effects

As discussed in the introduction, there are three candidate explanations of the positive association between deanships and patent applications: (1) a professor's ability improved and promotion to dean occurred at the same time; (2) appointment to the position of dean brings more resources for research and innovation; or (3) with political power and control, deans can win more patents via exchanges of favors or political influence.

We control for individual fixed effects in all of our regressions. Therefore, time-invariant individual innate ability should not confound our estimates. The positive correlations between deanship and patents could also be generated by increased research productivity. The large jump in patent applications upon promotion suggests, however, that the deanship effects are less likely to be driven by increased productivity. Changes in individual academic productivity should be gradual unless there is an interaction effect between individual ability and the position of dean, which seems not very likely.

In this section, we further examine the channels that might drive the deanship effects. Our results rule out the ability and resource explanations. First, we show that patent applications in one's previous field of expertise actually decrease significantly following promotion to a deanship. Second, we look at heterogeneity in individuals' characteristics. We find that the deanship effects tend to be larger for individuals who have ever experienced a paper retraction, which is a sign of poor academic professionalism. Third, exploiting heterogeneous characteristics of universities, we find that the deanship effects are much smaller at universities with weaker monoposony power, which is proxied using high-quality status of the university and more neighboring universities.

Fourth, we find that patent applications decrease greatly after leaving a deanship and loss of political power, but this figure does not change when a dean experiences further promotion up the bureaucratic hierarchy. Combining these pieces of evidence, it seems that the positive deanship effects on patents are driven by the (mis-)use of political power, rather than improved ability or greater access to resources to support a dean's own research. Finally, using data from the Natural Science Foundation of China, we show that a deanship does not bring more research funding.

Shifts in Fields of Expertise

If deans attained additional patent applications as a result of higher productivity or greater access to resources for their own research work, we should expect them to have more patents in their own fields of expertise. We therefore look at the shifts in patents from a dean's previous fields of knowledge following promotion. We define the field of a patent based on the first two digits of the International Patent Classification (IPC) code. A field is defined as a "previous" field of expertise for a dean if the dean had filed patent applications in this field more than five years before taking office. We count only patents in these old fields of expertise and re-estimate the deanship effects on patents.

The results are displayed in Figure 3.8 and Table 3.4. Surprisingly, we see a drop rather than an increase in the quantity of patent applications in deans' previous fields of expertise following promotion. It seems that the increase in patents of deans occurs mainly in fields of knowledge with which they were previously unfamiliar. If deans are indeed getting more patents as a reflection of enhanced ability or resources, they should enjoy such an advantage in their own fields of expertise to at least the same or a greater extent than in other fields with which they were previously unfamiliar. Our empirical findings, however, are inconsistent with this prediction. Specifically, the findings rule out the possibility of the interaction effects between individual ability and the dean's position.

Table 3.4: Deanships and Patents in Old Fields of Study

	(1) IHS(Patents)	(2) Patents>0	(3) IHS(Invention Patents)	(4) Invention Patents>0
5+years before	-0.038	0.015	-0.048	0.014
	(0.062)	(0.024)	(0.061)	(0.024)
4 years before	-0.023	0.015	-0.022	0.017
	(0.038)	(0.016)	(0.037)	(0.016)
3 years before	-0.019	-0.002	-0.021	0.001
	(0.031)	(0.014)	(0.030)	(0.014)
2 years before	-0.006	0.006	-0.001	0.011
	(0.024)	(0.011)	(0.024)	(0.011)
Year of promotion	-0.053**	-0.015	-0.053**	-0.014
	(0.022)	(0.009)	(0.022)	(0.010)
1 year after	-0.070***	-0.008	-0.063**	-0.001
	(0.027)	(0.012)	(0.026)	(0.012)
2 years after	-0.030	-0.009	-0.026	-0.005
	(0.033)	(0.015)	(0.032)	(0.015)
3 years after	-0.046	-0.017	-0.042	-0.013
	(0.040)	(0.016)	(0.039)	(0.017)
4 years after	-0.028	-0.014	-0.022	-0.009
	(0.046)	(0.019)	(0.046)	(0.019)
5+years after	-0.098***	-0.029**	-0.096***	-0.027*
	(0.038)	(0.015)	(0.037)	(0.015)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dep. Var. Mean	1.548†	0.238	1.347†	0.233
$(\frac{\text{Old fields}}{\text{Total}})^{\#}_{before}$	0.750		0.759	
$(\frac{\text{Old fields}}{\text{Total}})^{\#\#}_{after}$	0.230		0.237	

Notes: "Old" fields are defined as fields in which the individual had patent applications more than 5 years before becoming dean. Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, (4) a dummy for having positive invention patent applications. Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications. # The ratio of patents in old fields to all patents before promotion; ## the ratio of patents in old fields to all patents after promotion.

Figure 3.8: Deanship and Cognitive Mobility: Patents in Old Fields of Expertise

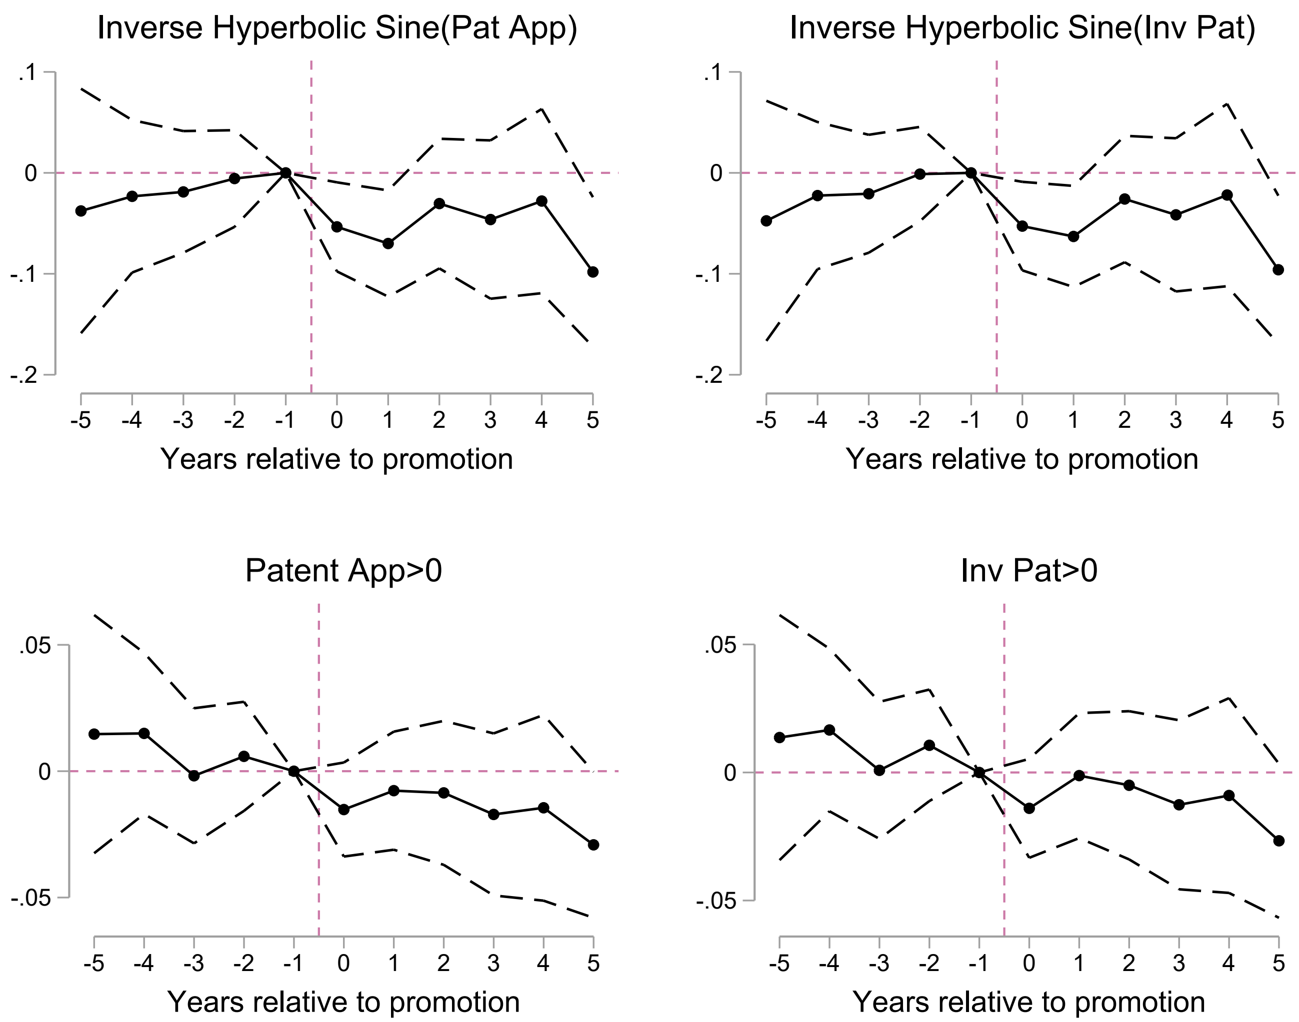

Notes: Event study estimates using the Abraham and Sun (2021) estimator and 95 percent confidence intervals are plotted. Regressions control for individual fixed effects, year fixed effects, and quadratic polynomial of potential experience. ”Old” fields are defined as fields in which the individual had patent applications more than 5 years before becoming dean. The field is determined by the first two digits of the International Patent Classification Code. We only include patent applications in these ”old” fields. Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5.

Academic Ethics

In this subsection we study the relationship between dean effects and academic ethics, which we measure using paper retractions. The retraction of a published papers is a sign of academic misconduct. We collect information on paper retractions from Retraction Watch (https://retractionwatch.com) and match retraction records to the deans in our sample using their biographical information. Details on the dataset are described in Section 3.2.4. The results are reported in Table 3.5. We show that researchers who experienced paper retractions also field a higher number of patent applications following promotion. The pattern suggests that the positive deanship effects on patents are less likely to be driven by ability and are more likely correlated with some sort of misconduct.

It may come as a surprise that, according to Retraction Watch, 8% of deans in our sample have a history of paper retractions. This fraction is sizable since retractions are usually attributable to academic misconduct (Biagioli et al., 2019; Fang et al., 2012). Although this figure is not directly associated with patent-related misconduct, it implies that academic misconduct is a nontrivial issue in our context.

Table 3.5: Deanship Effects and Academic Ethics

	(1) IHS(Patents)	(2) Patents>0	(3) IHS(Invention Patents)	(4) Invention Patents>0
Dean	0.122***	0.077***	0.095**	0.066***
	(0.040)	(0.016)	(0.038)	(0.016)
Dean*Retraction experience	0.224*	0.024	0.227**	0.046
	(0.121)	(0.048)	(0.115)	(0.045)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dep. Var. Mean	2.941†	0.425	2.507†	0.412

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, (4) a dummy for having positive invention patent applications. *Dean* is dummy for the dean has taking office. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications. 8% deans have records of paper retraction.

Local Researchers' Outside Options in the Academic Market

We also investigate how deanship effects vary with researchers' outside options in the labor market. We expect that, when local researchers enjoy greater bargaining power, deans' misuse of power will be limited. If deans are using their political influences to obtain intellectual property rights from other researchers, then stiffer competition from prospective employers will give local researchers more bargaining power, mitigating the deanship effects. We find that the deanship effects tend to be smaller when researchers have more bargaining power.

We first measure opportunities for horizontal mobility using the number of local universities within a given city. The rationale is that neighboring universities could compete for researchers, whose moving costs would be relatively low because their families do not have to move within the city. We also measure opportunities for vertical mobility using university quality, which we measure using status as a Project 985 university. Project 985 universities are identified by the Ministry of Education as top-tier universities. Status as a Project 985 university is a strong predictor of university quality, which reflects prestige, government support, superior financial and academic resources, and higher-achieving researchers and students.

The results are reported in Table 3.6. As can be seen in panel A, we find that deanship effects are much smaller and in fact nonsignificant in Project 985 universities compared with non-Project 985 universities. Panel B indicates that deanship effects indeed drop with the number of neighboring universities located in the same city.

Table 3.6: Deanship Effects and Local Researchers' Labor Market

A. Number of neighboring universities and deanship effects

	(1)	(2)	(3)	(4)
	IHS(Patents)	Patents>0	IHS(Invention Patents)	Invention Patents>0
Dean	0.300***	0.154***	0.221***	0.130***
	(0.081)	(0.032)	(0.074)	(0.030)
Dean*No. of Neighbors	-0.007**	-0.003***	-0.005	-0.003**
	(0.003)	(0.001)	(0.003)	(0.001)

B. Project 985 vs. non-985 universities

	(1)	(2)	(3)	(4)
	IHS(Patents)	Patents>0	IHS(Invention Patents)	Invention Patents>0
Dean	0.235***	0.131***	0.186***	0.115***
	(0.054)	(0.022)	(0.052)	(0.021)
Dean*Project 985	-0.181***	-0.100***	-0.136**	-0.087***
	(0.069)	(0.027)	(0.066)	(0.026)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dep. Var. Mean	2.941†	0.425	2.507†	0.412

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, (4) a dummy for having positive invention patent applications. *Dean* is dummy for the dean has taking office. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

These results are consistent with our interpretation that deanship effects are weakened as the outside options for other local researchers multiply, implying that misconduct drives these effects. The effects could be driven by either voluntary collusion or involuntary "collusion" between deans and local researchers, both potentially impeding the incentives available to local researchers. The fact that deanship effects vary with outside conditions provides further evidence against the possible interaction effects between ability and the position of the dean.

An alternative explanation for these patterns is that research resources matter more in low-quality universities or in less fully-developed regions, where resources are relatively scarce. If that is true, then our results may reflect enhanced access to resources attached to the position of a dean. Although we cannot reject this possibility entirely, it is ex ante unclear whether the deanship effects on resources should be smaller in lower-quality universities. It is also very likely that there are more abundant resources available at the high-quality universities that would be at the dean's disposal.

The Effects of Losing Deanship and Further Promotion

The effects of leaving a deanship, either by losing power completely or being further promoted, also shed light on the role of power in obtaining patents. We consider individuals whose deanships ended during the sample period.[18] We classify these deans into two groups. The first group includes deans who lost power but still kept their academic positions. These individuals do not hold administrative positions after leaving office but keep their academic positions. The second group consists of individuals who were further promoted in the administrative hierarchy. For example, some were promoted to principal or vice principal of their universities, while others became heads of other university administrations. We do not include in our analysis deans who fully retired after leaving deanships, because it is natural to expect that their research project will drop significantly.

[18]This sample is different from our main sample of analysis. With this sample, we include only individuals whose deanships ended during the sample period to identify the effects of losing deanship.

We analyze the association between deanships and patent applications for these two groups of individuals respectively. The event study estimates for patent applications before and after leaving deanships are provided in Table 3.7 and 3.8 as well as Figure 3.9 and 3.10 . We find large decreases in patent applications for individuals who actually lose administrative power after leaving their deanships. On the other hand, we do not find significant changes in patent applications for the group of deans who were further promoted. Considering that those who were further promoted assume responsibility for additional administrative duties and are much less likely to remain focused on research, the fact that their patenting activities do not drop off also suggests that political power plays an important role. Taken together, these results imply that the positive associations between deanships and patent applications are more likely driven by political power.

Table 3.7: Effects of Losing Power on Patenting

	(1)	(2)	(3)	(4)
	IHS(Patents)	Patents>0	IHS(Invention Patents)	Invention Patents>0
5+years before	0.036	0.046	0.060	0.062
	(0.200)	(0.072)	(0.191)	(0.074)
4 years before	-0.017	-0.009	-0.025	-0.012
	(0.141)	(0.060)	(0.140)	(0.060)
3 years before	0.130	0.062	0.127	0.092*
	(0.128)	(0.053)	(0.125)	(0.054)
2 years before	0.078	0.044	0.091	0.055
	(0.104)	(0.046)	(0.102)	(0.045)
Year of leaving deanship	-0.084	-0.019	-0.078	-0.043
	(0.115)	(0.054)	(0.111)	(0.053)
1 year after	-0.030	-0.027	-0.023	-0.039
	(0.132)	(0.060)	(0.126)	(0.061)
2 years after	-0.419**	-0.132*	-0.402**	-0.163**
	(0.172)	(0.080)	(0.169)	(0.081)
3 years after	-0.382*	-0.096	-0.366*	-0.133
	(0.222)	(0.095)	(0.218)	(0.103)
4 years after	-0.705***	-0.209*	-0.671**	-0.238*
	(0.271)	(0.111)	(0.260)	(0.123)
5+years after	-0.543	-0.304*	-0.386	-0.302**
	(0.368)	(0.160)	(0.327)	(0.130)
Observations	1529	1529	1529	1529
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dep. Var. Mean	3.577†	0.535	3.038†	0.516

Notes: All regressions control for a quadratic polynomial of potential experience. Event study coefficients are estimated using the interaction-weighted estimator proposed by Abraham and Sun (2021). Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, and (4) a dummy for having positive invention patent applications.Years more than 5 years before leaving deanship are stacked to -5. Years more than 5 years after leaving deanship are stacked to 5. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

Table 3.8: Effects of Further Promotion on Patenting

	(1)	(2)	(3)	(4)
	IHS(Patents)	Patents>0	IHS(Invention Patents)	Invention Patents>0
5+years before	-0.067	0.033	0.010	0.042
	(0.339)	(0.134)	(0.339)	(0.139)
4 years before	0.008	0.005	0.084	0.026
	(0.254)	(0.102)	(0.253)	(0.104)
3 years before	0.042	-0.015	0.096	0.002
	(0.217)	(0.092)	(0.210)	(0.095)
2 years before	-0.013	-0.005	0.010	0.024
	(0.154)	(0.069)	(0.150)	(0.070)
Year of leaving deanship	-0.076	-0.017	-0.010	0.009
	(0.160)	(0.065)	(0.154)	(0.066)
1 year after	-0.039	0.114	-0.092	0.053
	(0.251)	(0.097)	(0.244)	(0.097)
2 years after	0.055	0.040	0.017	0.079
	(0.343)	(0.126)	(0.339)	(0.130)
3 years after	0.137	-0.073	0.123	-0.053
	(0.459)	(0.161)	(0.463)	(0.174)
4 years after	0.147	-0.003	0.050	-0.025
	(0.509)	(0.180)	(0.499)	(0.192)
5+years after	-0.189	-0.128	-0.128	-0.066
	(0.644)	(0.225)	(0.640)	(0.219)
Observations	902	902	902	902
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dep. Var. Mean	4.416†	0.594	3.791†	0.576

Notes: All regressions control for a quadratic polynomial of potential experience. Event study coefficients are estimated using the interaction-weighted estimator proposed by Abraham and Sun (2021). Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, and (4) a dummy for having positive invention patent applications.Years more than 5 years before leaving deanship are stacked to -5. Years more than 5 years after leaving deanship are stacked to 5. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

Figure 3.9: Retirement and Patenting: Losing Power

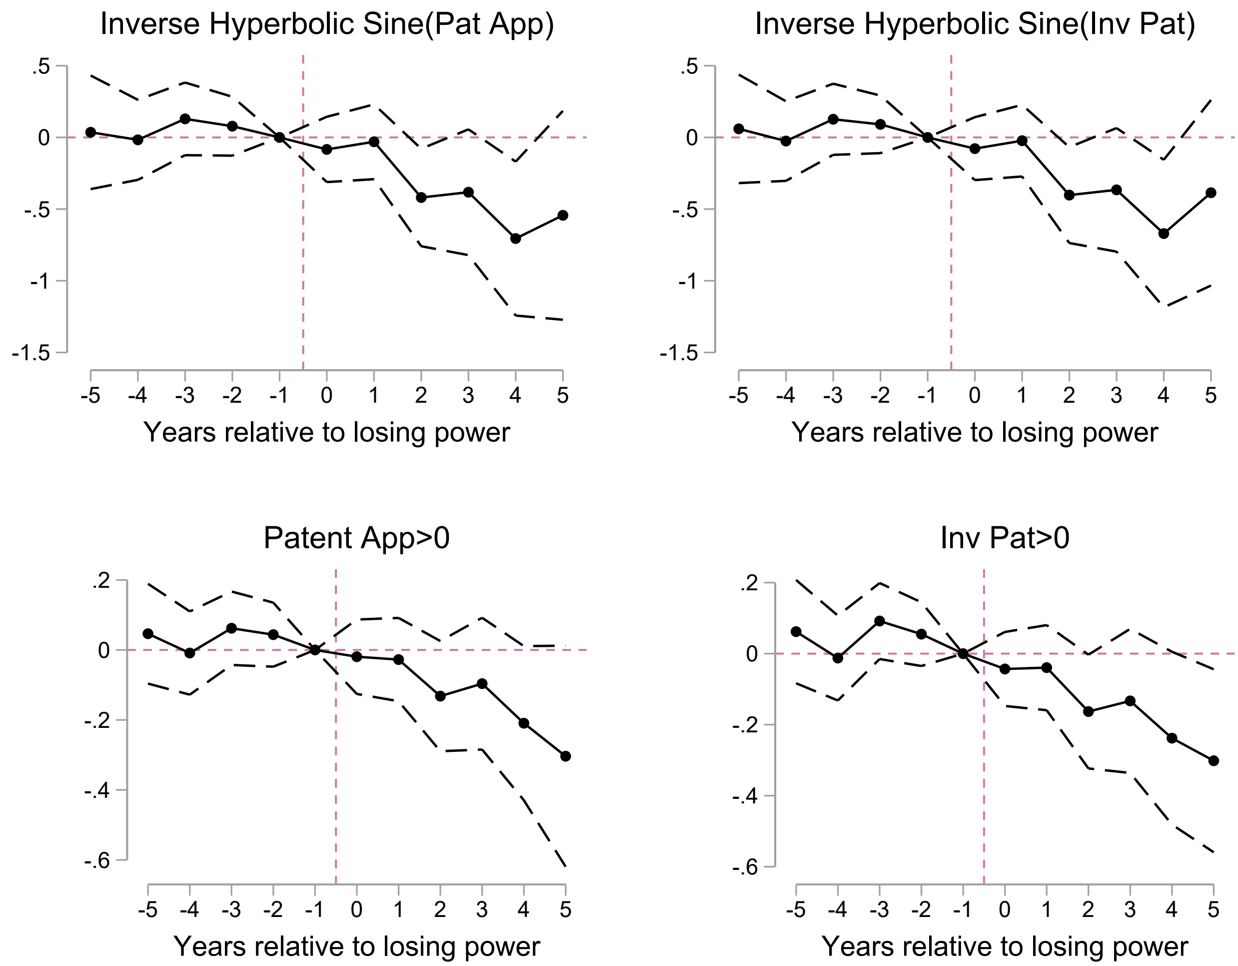

Notes: Event study estimates using the Abraham and Sun (2021) estimator and 95 percent confidence intervals are plotted. Regressions control for individual fixed effects, year fixed effects, and quadratic polynomial of potential experience. Years more than 5 years before leaving deanship are stacked to -5. Years more than 5 years after leaving deanship are stacked to 5.

Figure 3.10: Further Promotion and Patenting

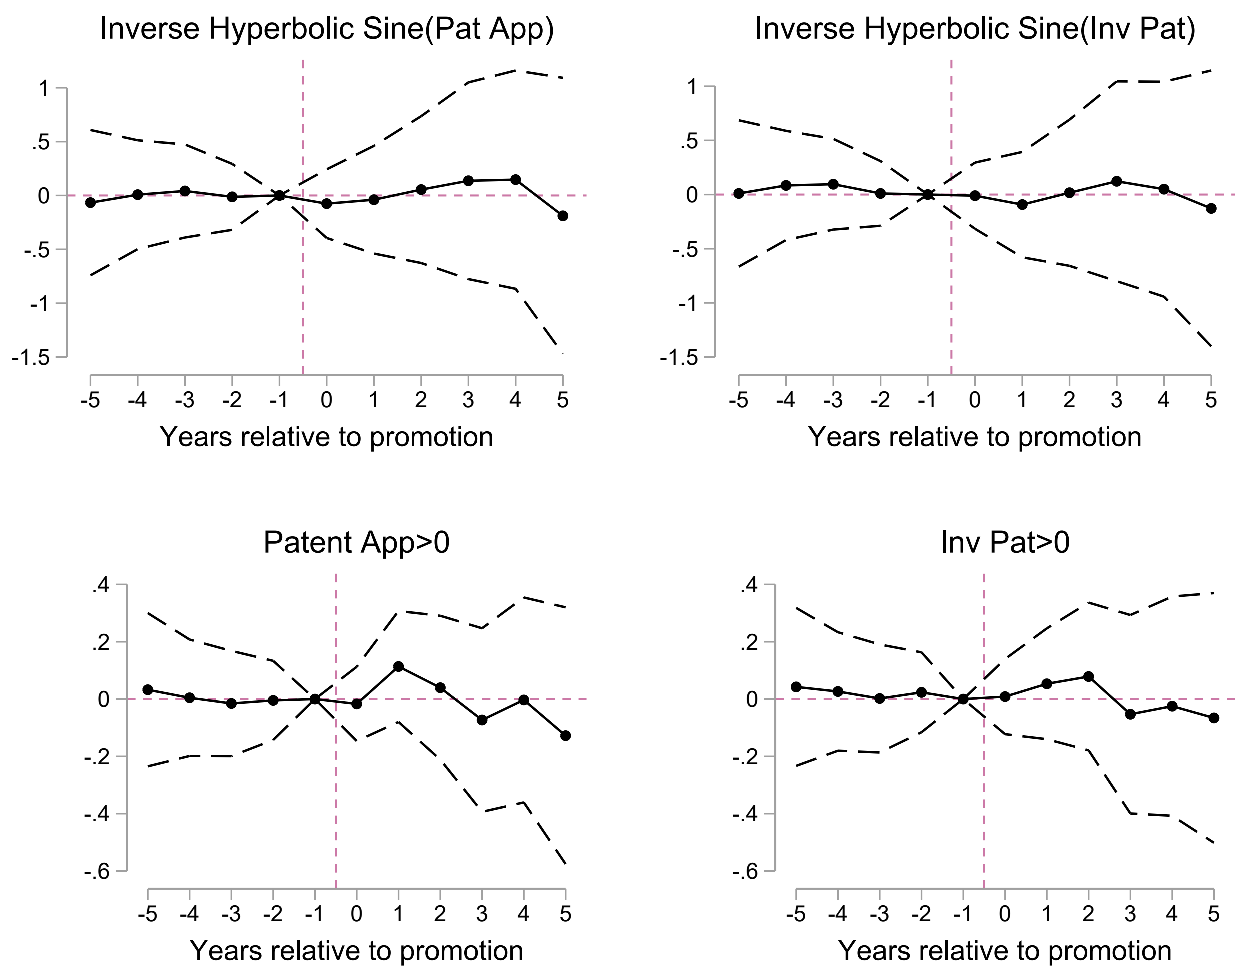

*Notes:*Event study estimates using the Abraham and Sun (2021) estimator and 95 percent confidence intervals are plotted. Regressions control for individual fixed effects, year fixed effects, and quadratic polynomial of potential experience. Years more than 5 years before leaving deanship are stacked to -5. Years more than 5 years after leaving deanship are stacked to 5.

However, it is also possible there is a systematic productivity gap between deans who have been promoted to higher positions and those who have not. To probe into this alternative explanation, we estimate the correlation between losing power and observable characteristics of individuals that left deanship using the following model:

$$y_{iut} = \theta_1 LosingPower_{iut} + \theta_1 exp_{it} + \theta_2 exp_{it}^2 + \delta_u + \nu_t + \varepsilon_{iut} \tag{3.3}$$

The dependent variables for individual i from university u who left in deanship in year t are their potential experience at the time of leaving deanship as well as measures for patent applications before leaving deanship. $LosingPower_{iut}$ is an indicator for individual i not being further promoted in the university hierarchy. We also control for the quadratic polynomial of potential experience, university fixed effects, and year-of-leaving fixed effects. The results are reported in Table 3.9. Column (1) shows that deans who were not promoted were older than those who were promoted, but the gap is not large (1.2 years) and statistically non-significant. The following columns show that, whether conditional on potential experience or not, there is no strong association between losing power and productivity of deans. Therefore, based on the observed characteristics, it seems that the heterogeneous effects of leaving deanship on patenting are not driven by the productivity gap between the two groups of deans.

The correlation between the position of deanship and patenting still does not fully reject the resource channel. It could be that higher administrative positions are associated with additional resources to support research, which is consistent with both the positive effects on patenting of obtaining deanships and the negative effects of losing deanships. However, our analysis suggests that the positive association between deanships and patent applications is driven by the administrative position rather than deans' own productivity. In the next part, we use direct information on research funding to investigate the potential role of resources.

Table 3.9: Observable Characteristics of Losing-power vs. Further Promoted Deans before Leaving Deanship

	(1)	(2)	(3)	(4)	(5)
	Exp	IHS(Patents)	Patents>0	IHS(Invention Patents)	Invention Patents>0
A. Unconditional on Potential Experience					
Losing Power	1.224	0.00498	0.0202	-0.0103	0.0163
	(0.961)	(0.220)	(0.0614)	(0.209)	(0.0626)
B. Conditional on Potential Experience					
Losing Power		0.00915	0.00816	-0.00734	0.00334
		(0.220)	(0.0639)	(0.210)	(0.0658)
Exp		0.0460	0.0412*	0.0542	0.0446*
		(0.0665)	(0.0199)	(0.0636)	(0.0208)
Exp^2		-0.00136	-0.000865*	-0.00156	-0.000938*
		(0.00121)	(0.000376)	(0.00114)	(0.000379)
Observations	217	217	217	217	217
Year FE	Yes	Yes	Yes	Yes	Yes
University FE	Yes	Yes	Yes	Yes	Yes
Dep. Var. Mean	20.10	1.557	0.682	1.439	0.664

Notes: Robust standard errors in parentheses. Dependent variables are the potential experience of the dean in the year of retirement and measures for average patent applications during the two years before retirement and the year of retirement. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

Research Resources

In this section we present a direct test of the research-resource channel. The resource channel could be ruled out if we were to find no increase in all resources (funds, equipment, etc.) granted to a dean before to after the promotion. Unfortunately, we have no access to such comprehensive information. We do, however, have information on funding from the National Natural Science Foundation of China (NSFC), which is a major source of government support for natural science research.[19]

We collect administrative records for all applications to the NSFC and match them with our dean-level dataset. We estimate the deanship effects on funding acquisition again using the two-way fixed effects event-study design. The results are reported in Figure 3.11 and Table 3.10. There is no increase in NSFC grants following promotion. This provides direct evidence to rule out the resource channel. If the dean is obtaining additional patents as a result of enhanced access to resources, then this channel would be partially reflected in obtaining more funding from NSFC, which represents a major source of scientific funding from the public sector. It is worth noting that we do not know whether the NSFC is more prone to the misuse of power or less. It is also unclear whether the NSFC funding would be correlated with ability. Therefore, we take the results regarding the NSFC funding as a test on the resource channel without making any claims regarding the other two channels. The limitation of this analysis is that we cannot observe grants from the private sector. This is not a major concern for us because the private sector should have stronger incentives to care about ability than the position of a dean. The influence of a dean's position per se should be weaker outside the academic system.

[19] In 2017, 64.54% of SCI papers published by Chinese researchers noted funding from the NSFC (He et al., 2018).

Table 3.10: Deanships and NSFC Grants: Event Study Results

	(1)	(2)
	NSFC>0	IHS(NSFC Funding)
5+ years before	0.023	0.091
	(0.020)	(0.097)
4 years before	-0.009	-0.052
	(0.019)	(0.089)
3 years before	-0.005	-0.031
	(0.018)	(0.084)
2 years before	0.001	0.003
	(0.018)	(0.085)
Year of promotion	0.010	0.053
	(0.020)	(0.094)
1 year after	0.019	0.083
	(0.022)	(0.102)
2 years after	-0.004	-0.022
	(0.023)	(0.108)
3 years after	0.010	0.062
	(0.024)	(0.109)
4 years after	0.007	0.057
	(0.026)	(0.121)
5+ years after	0.018	0.111
	(0.016)	(0.071)
Observations	9878	9878
Individual FE	Yes	Yes
Year FE	Yes	Yes
Dep. Var. Mean	0.220	83.866†

Notes: All regressions control for quadratic polynomial of potential experience. Event study coefficients are estimated using the interaction-weighted estimator proposed by Abraham and Sun (2021). Dependent variables are (1) a dummy for obtaining any NSFC awards, (2) the inverse hyperbolic sine of NSFC awards .Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average amount of NSFC awards (in thousand RMB yuan).

Figure 3.11: Dynamic Deanship Effects on NSFC Grants

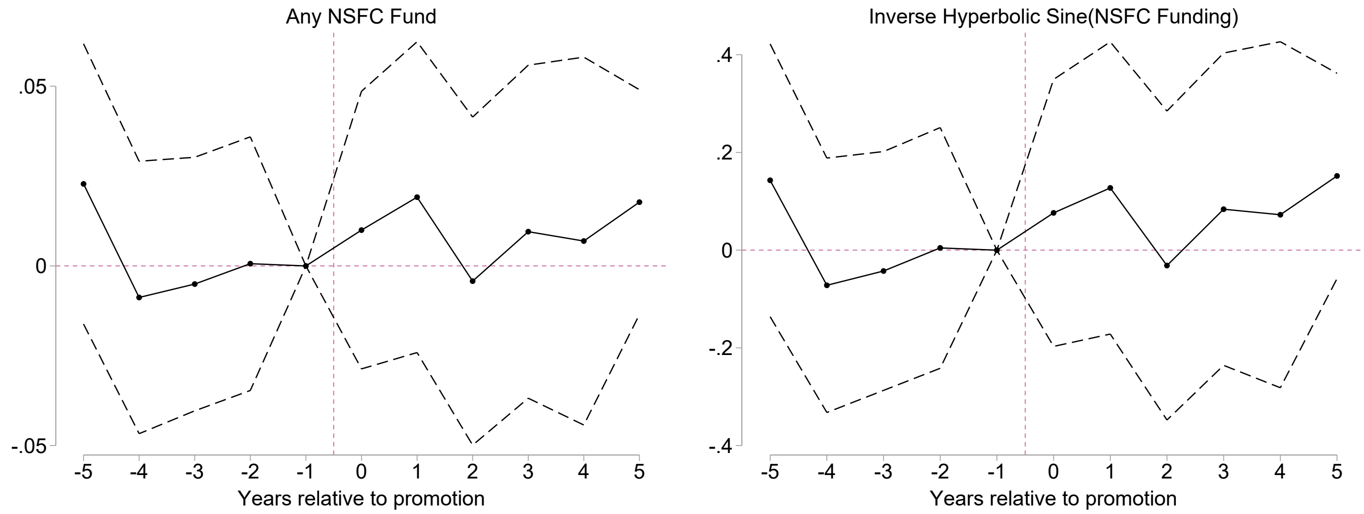

Notes: The dependent variables are (1) a dummy for having any NSFC funding and (2) the inverse hyperbolic sine transformation of total NSFC funding in the given year. We estimate the changes in NSFC funding before and after an individual is promoted to a deanship. Event study estimates using the Abraham and Sun (2021) estimator and 95 percent confidence intervals are plotted. Regressions control for individual fixed effects, year fixed effects, and quadratic polynomial of potential experience. Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5.

Do NSFC funds really matter for innovation? We show the correlation between gaining NSFC funding and patenting in Table 3.11. We regress patent applications on two measures of NSFC funds: a dummy for having any NSFC funding in a given year and the inverse hyperbolic transformation of the amount of NSFC funding awarded. We see that NSFC funding is highly predictive of patent outputs, but only for researchers who have not yet been appointed as deans. For example, before taking office, having any NSFC awards is correlated with filing approximately 10% more patent applications. The positive correlation between NSFC funding and patent applications disappears following appointment to a deanship. These results both confirm the importance of support from the NSFC for innovation and provide further evidence that the deanship effects on patenting are not driven by greater availability of resources for deans.

Table 3.11: Correlation between NSFC Funding and Patenting

	(1)	(2)	(3)	(4)
	IHS(Patents)	IHS(Patents)	IHS(Invention Patents)	IHS(Invention Patents)
Any NSFC	0.101***		0.086***	
	(0.031)		(0.030)	
Any NSFC∗Dean	-0.161***		-0.148***	
	(0.050)		(0.048)	
IHS(NSFC Funding)		0.016***		0.014***
		(0.005)		(0.005)
IHS(NSFC Funding)∗Dean		-0.026***		-0.023***
		(0.007)		(0.007)
Dean	0.164***	0.166***	0.136***	0.137***
	(0.039)	(0.039)	(0.038)	(0.038)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dep. Var. Mean	2.941†	0.425	2.507†	0.412

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are(1) the inverse hyperbolic sine of total patent applications, (2) the inverse hyperbolic sine of invention patent applications. *Dean* is dummy for the dean has taking office. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

3.5 Anti-corruption and the Deanship Effects

In the previous section, we provide evidence that misuse of power is the major source of deanship effects. If deanship effects are indeed a result of misuse of political power, then political reforms aimed at disciplining the behavior of officials should dampen such effects. We consider the recent anti-corruption campaign in China. The results obtained after exploiting the anti-corruption campaign show that local anti-corruption efforts have greatly weakened deanship effects on patents. This is unlikely to be consistent with the explanations based on deans' ability or (proper) use of resources for their own research. The fact that the political campaign matters suggests that the positive deanship effects on patents are more likely a result of political power associated with the bureaucratic position.

3.5.1 Anti-corruption Efforts and Deanship Effects

The local intensity of the anti-corruption campaign is measured using the cumulative number of anti-corruption investigations conducted in each city since the onset of the anti-corruption campaign (Section 3.2.4). Exploiting regional variations in the anti-corruption efforts, we estimate a DDD specification as Equation 3.2. For our baseline analysis, we use information on investigations released by the Central Committee of Discipline Inspection (CCDI), which is more directly relevant in our case as both officials covered in the CCDI database and the elite universities in our sample are under the control of either a provincial or the central government in the administrative hierarchy. We also use alternative measures for anti-corruption to confirm the robustness of our results.

The estimates of the DDD model are reported in Table 3.12. For panel A of Table 3.12, we consider investigations into all kinds of government officials. We find a significant reduction in the deanship effects as anti-corruption efforts intensify. The coefficients reported in column (1) of panel A suggest that doubling local anti-corruption efforts would reduce the deanship effect by 4.9%, which is large in magnitude relative to the 14% average deanship effects on patent applications (Table 3.2). In panel B, we restrict our attention only to anti-corruption investigations into officials affiliated with the higher education system. Presumably, these cases should be much more directly relevant to the deans. We do find larger impacts of these university-related anti-corruption efforts on deanship effects. These results further support our hypothesis that the anti-corruption campaign have weakened the deanship effects on patenting, which can be plausibly explained only by events in which the misuse of power is disciplined.

Table 3.12: Local Anti-corruption Efforts and Deanship Effects

A. All Local Investigations

	(1)	(2)	(3)	(4)
	IHS(Patents)	Patent App>0	IHS(Invention Patents)	Inv Pat>0
Dean	0.175***	0.086***	0.151***	0.078***
	(0.050)	(0.022)	(0.048)	(0.022)
Dean*Local Investigations	-0.049**	-0.019*	-0.049**	-0.021**
	(0.023)	(0.011)	(0.022)	(0.011)

B. College-related Local Investigations

	(1)	(2)	(3)	(4)
	IHS(Patents)	Patent App>0	IHS(Invention Patents)	Inv Pat>0
Dean	0.145***	0.077***	0.119***	0.070***
	(0.042)	(0.018)	(0.040)	(0.018)
Dean*Local Investigations	-0.155**	-0.062*	-0.166**	-0.084**
	(0.068)	(0.036)	(0.068)	(0.037)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Individual FE*Local Investigations	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Year FE*Local Investigations	Yes	Yes	Yes	Yes
Dep. Var. Mean	2.941†	0.425	2.507†	0.412

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, (4) a dummy for having positive invention patent applications. *Local Investigations* is the (inverse hyperbolic sine of) cumulative number of anti-corruption investigations for officials with level *Chuji* or above in the city since the start of 2013. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

In addition to the CCDI database, we also use alternative proxies for the anti-corruption campaign to check the robustness of our findings. The first alternative is the Tencent database (Section 3.2.4), which is the most comprehensive dataset covering anti-corruption investigations. The results using the Tencent database are reported in Table 3.18, confirming our findings for anti-corruption and deanship effects. The second alternative proxy we use is a dummy for having ever had central inspection teams sent to the province where the university is located. This measure captures whether the anti-corruption campaign was ever enforced at the province level, which is also used in previous literature such as Chen and Kung (2019a). The results are reported in Table 3.20 Although imprecisely estimated, the model exploiting the inspection teams confirms our previous findings. For these results, we exploit both temporal and regional variations in the campaign. This strategy provides more convincing identification of the impacts of anti-corruption. The anti-corruption campaign, however, was implemented nationwide. Therefore, we also estimate a version of the DDD design where we exploit the national level variation of the campaign by defining a dummy for anti-corruption campaign which equals one for the years after 2013. That model captures the overall variation before and after the onset of the anti-corruption campaign, at the cost of potentially suffering from other contemporaneous factors. The results, which are reported in Table 3.22, again confirms our findings regarding the impact of the anti-corruption campaign on deanship effects.

3.5.2 Placebo Test for Patents in Previous Fields

As shown above, deans did not obtain more patents in their previous fields of expertise following promotion. If the deanship effects on patents are driven by political power, we should observe that the anti-corruption campaign does not affect deans' innovative activities in their previous fields of expertise. We look at only deans' patent applications in their previous research areas and test the effects of the anti-corruption campaign. In Table 3.13 we report the empirical results, again using both local investigations and university-related local investigations as measures of anti-corruption efforts. The coefficient on the interaction term Dean*Investigation is neither economically nor statistically significant. This finding implies that the anti-corruption campaign shock limited the misuse of power but did not affect deans' real productivity that was not otherwise associated with deanships. In addition, we investigate the robustness of our placebo test by employing alternative measures for the local intensity of anti-corruption efforts. The results are reported in Table 3.19, Table 3.21, and Table 3.23, all of which are consistent with our findings using the CCDI measure for anti-corruption.

Table 3.13: Local Anti-corruption Efforts and Patenting in Deans' Previous Fields of Expertise

A. All Local Investigations

	(1)	(2)	(3)	(4)
	IHS(Patents)	Patent App>0	IHS(Invention Patents)	Inv Pat>0
Dean	-0.099***	-0.033***	-0.099***	-0.032***
	(0.026)	(0.012)	(0.026)	(0.012)
Dean*Local Investigations	0.008	0.007	0.008	0.005
	(0.014)	(0.006)	(0.014)	(0.006)

B. College-related Local Investigations

	(1)	(2)	(3)	(4)
	IHS(Patents)	Patent App>0	IHS(Invention Patents)	Inv Pat>0
Dean	-0.089***	-0.029***	-0.088***	-0.027**
	(0.025)	(0.010)	(0.025)	(0.010)
Dean*Local Investigations	-0.015	-0.001	-0.029	-0.022
	(0.041)	(0.020)	(0.041)	(0.022)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Individual FE*Local Investigations	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Year FE*Local Investigations	Yes	Yes	Yes	Yes
Dep. Var. Mean	1.548†	0.238	1.347†	0.233

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, (4) a dummy for having positive invention patent applications. *Local Investigations* is the (inverse hyperbolic sine of) cumulative number of anti-corruption investigations for officials with level *Chuji* or above in the city since the start of 2013. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

3.6 Spillover Effects on Other Local Researchers

If deans are obtaining additional patents as a result of their power, would that affect the allocation of key research resources within their institutions? We care about the potential production distortion and the associated welfare implications. Misuse of political power could reduce productivity in many ways, such as by discouraging young researcher's output by reducing their expected returns, or misallocating research resources to less productive researchers who are willing to enter into exchanges of interests with deans.

To shed light on this issue, we investigate whether there are spillover effects of deanships on local researchers that are connected to a dean. Given data limitations, though, we cannot test all types of resource allocation. Luckily, NSFC application records enable us to investigate some of the spillover effects of a dean's power on resource allocation. Specifically, we consider local researchers who had formed co-inventorship connections with a dean before the dean's promotion. Note that the sample we use here includes only researchers who were connected to the dean before the promotion. We find that these connected local researchers obtained more NSFC funding after the dean took office. We estimate the following DID model:

$$NSFC_{jit} = \beta_0 + \alpha * Dean\_Promoted_{it} + \lambda_j + \theta_t + \varepsilon_{ijt} \tag{3.4}$$

The dependent variable is either a dummy for having any NSFC awards or the inverse hyperbolic sine of the total amount of NSFC awards to local researcher j in year t. i is the dean with whom j had co-invention experience before i took office. $Dean\_Promoted$ is a dummy for the connected dean's having been promoted to the deanship, the coefficient of which is the parameter of interest. λ_j and θ_t are individual and year fixed effects, respectively.[20] The results are reported in Table 3.14. After a dean takes office, their previously connected researchers are 1.2% more likely

[20] We remove local researchers connected to multiple deans, so dean fixed effects are absorbed by λ_j.

to receive NSFC funding and the average amount of funding increases by approximately 7.6%. These effects are large relative to the average level of funding awarded to all researchers. Taken together with previous evidence of the misuse of deans' power, this additional evidence of spillover effects of deanships implies that the power that deans wield could distort resource allocation for research.

Table 3.14: Deanship Effects on NSFC Funding of Co-inventors before Promotion

	(1)	(2)
	NSFC>0	IHS(NSFC Funding)
Coinventor Promoted to Dean	0.012**	0.076*
	(0.006)	(0.042)
Observations	27159	27159
Individual FE	Yes	Yes
Year FE	Yes	Yes
Dep. Var. Mean	0.103	47.571†

Notes: Dependent variables are (1) a dummy for obtaining any NSFC awards, and (2) the inverse hyperbolic sine of NSFC awards. Standard errors in parentheses are clustered at the dean's level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average amount of NSFC awards (in thousand RMB yuan).

One alternative explanation is that the dean must care about the entire output of the schools they lead and so they help other faculty in their research efforts voluntarily. If that is the case, we should not observe the deanship effects in the first place and we should not observe any anti-corruption effects. Note here that we treat the exchange of power for other researchers' grants and obtaining more patents through this channel as a misuse of power. These pieces of evidence together alleviate our concern regarding the alternative explanation.

3.7 Conclusions

In this chapter, we show that the power that deans wield exerts large impacts on innovation activities in Chinese universities. We first document a large and immediate positive association between patenting and deanships. Various pieces of evidence are provided which suggest that the deanship effects are driven by political power rather than deans' ability or enhanced access to research resources.

Our argument is supported by a series of empirical tests. The sharp increase in patenting activities following promotion suggests that institutional factors rather than individual abilities are the main reason, unless promotion coincides with dramatic increase in research productivity. Additional evidence further supports the political-power interpretation. First, promotion to a deanship is associated with a decrease rather than an increase in patenting in the dean's previous fields of innovation, suggesting a deviation from one's own areas of expertise. Second, the deanship effects we find vary with personal traits: paper-retraction experiences are associated with larger deanship effects on patenting. This association, although indirect, further suggests that the deanship effects on patenting are more likely driven by misconduct than by ability or resources.

Third, the deanship effects we observe vary with university characteristics. Competition from neighboring universities weakens the deanship effects, which suggests that competition disciplines deans' behaviors. We also document that deanship effects are much smaller in higher-quality universities, possibly because it is easier for local researchers to move down the "university ladder". These two findings, however, cannot fully rule out the role of resources. Fourth, we show that patent applications drop dramatically after deanships are lost, but do not change following further promotion along the bureaucratic hierarchy. This suggests that the deanship effects we observe are likely not driven by individual ability, although this finding does not rule out the resources channel entirely. Fifth, we provide direct evidence, based on funding from the National Natural Science Foundation of China, that deanships do not generate additional research resources.

Furthermore, we exploit a natural experiment made possible by the recent anti-corruption campaign and show that (1) local anti-corruption efforts reduced deanship effects on patenting, (2) local anti-corruption efforts targeted at the higher education system reduced the deanship effects to a greater extent than general anti-corruption efforts did, and (3) the anti-corruption campaign had no effect on deans' patents in their own fields of expertise, which did not increase with deanships in the first place. Although these findings do not constitute direct evidence, the associated empirical patterns imply that the deanship effects are more likely to be driven by misconduct related to political power rather than by ability or the proper allocation of resources.

We summarize all the results in Table 3.15 below and seek to identify the interpretation that is most consistent with all of them. The three columns correspond to the enhanced ability, enhanced access to resources, and the misuse of power channels, respectively. All of these findings taken together allow us to conclude that misuse of power best explains the deanship effects.

Table 3.15: Channels for the Deanship Effects on Patenting

	Ability	Resource	Power Misuse
Research Field	×	×	✓
Paper Retraction	×	×	✓
Outside Options	×	✓	✓
Losing Power	×	✓	✓
NSFC Funding		×	
Anti Corruption	×	×	✓

Note: ✓if empirical findings support the channel. ×if empirical findings cannot be explained by the channel.

For our final exercise, we explore the spillover effects of deanships on the allocation of research resources. We show that promotion to a deanship improves previous co-inventors' access to funding provided by the Natural Science Foundation of China. These results may help to explain why deans can add their name to other researchers' patents, suggesting how the misuse of power affects the allocation of resources within an academic institution.

Unpacking the role of political power in innovation activities helps us better understand the institutional factors related to the process of innovation. Our findings suggest that it is crucial to devise policies that can limit the misuse of power and reduce misallocation of resources in academia. As a result of data limitations, we could not evaluate the impact of misuse of power on researchers' productivity. A comprehensive evaluation of the effects of political power on resource allocation and productivity in academia would provide a meaningful and rich area for future work.

3.8 Appendix

3.8.1 Data Appendix

Sample of Deans

We first match the deans to patent records based on the names of patent owners and institutions. We then exclude from the sample deans who graduated before 2006 or those whose date of graduation is missing.[21] Using information about date of graduation, we can construct the potential experience of the dean and control for the life cycle of productivity. We also exclude deans who retired before 2016. Some retired deans got promoted to higher positions while others lost power. We therefore focus on deans who were still incumbent as of 2016 to identify the deanship effects. These criteria leave us with a balanced panel of 898 deans from 97 universities located in 39 prefectures. Finally, we match this dataset to measures of anti-corruption efforts based on geographical location (provinces for inspection teams, prefectures for number of investigations)

[21] We use primarily the year of receiving doctoral degree as the year of graduation. However, in certain situations, some individuals started their career with a master's degree, and got a doctoral degree on the job many years after the master's degree. The reason is that the higher education system was not as developed in the 1970s or 80s as nowadays. Therefore, these individuals could work as a professor with just a master's degree. We checked these cases and use the year of master's degree as the year of graduation. Controlling for the polynomial of experience or not does not qualitatively change our conclusions.

and year.

For our main analysis regarding the effects of deans taking office, we keep deans who were not yet retired by the end of the sample period (2016). To further test the role of political power, we also look at the sample of deans who retired during the sample period. These deans were already incumbent as of 2006. For these deans, we also collect information on their following positions. We classify these retired deans into "losing power" and "further promotion". Those in the "losing power" group returned to the position of an ordinary professor or leader of a laboratory. Those in the "further promotion" group moved up in the administrative hierarchy into positions such as president or vice president of the university.

Other Local Researchers

To investigate how potential spillovers of dean's political power might affect resource allocation, we construct a panel of local researchers during the same time period (2006-2016). Due to data limitations, we construct this sample using the universe of patent application records.[22] In the first step, we identify each local researcher using the combination of the institution (university) and the person's name from patent records. Many of the identified individuals, however, might be PhD students or professors entering/retiring from the university during the sample period. To address this concern, we keep local researchers in the sample using the following criteria: (1) we observe at least one patent application record for this researcher in the corresponding university during the years around the start of the sample period (2005-2008); (2) we observe at least one patent application record for this researcher in the corresponding university during the years around the end of the sample period (2013-2017); (3) we observe at least one patent application record for this researcher in the corresponding university between 2009 and 2012. We then construct the co-invention link between the local researchers and deans using patent applications. We restrict

[22] Many professors do not publish their resumes/biographical information online.

the sample to local researchers who had ever co-invented with a dean before the dean took office in order to examine the spillover effects of having a co-inventor promoted to deanship. The results reported in Table 3.14 are almost identical when we also include the local researchers who had never co-invented with a dean as a pure control group.

3.8.2 Additional Figures and Tables

Figure 3.12: OLS Estimates for the Two-way Fixed Effects Model

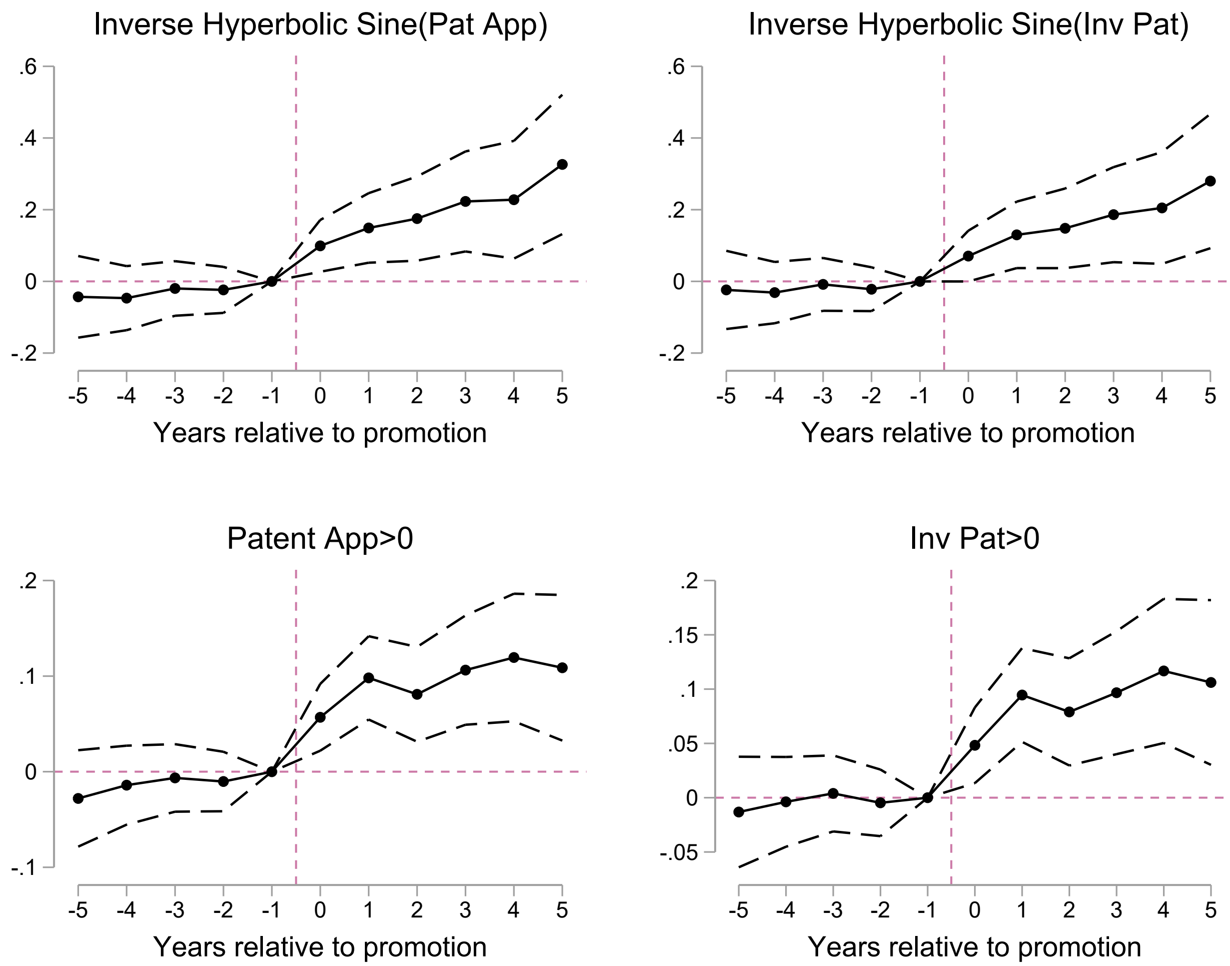

Notes: Point estimates and 95% confidence bands are plotted.

Figure 3.13: Number of Inventors of Patents before and after Promotion to Deanship

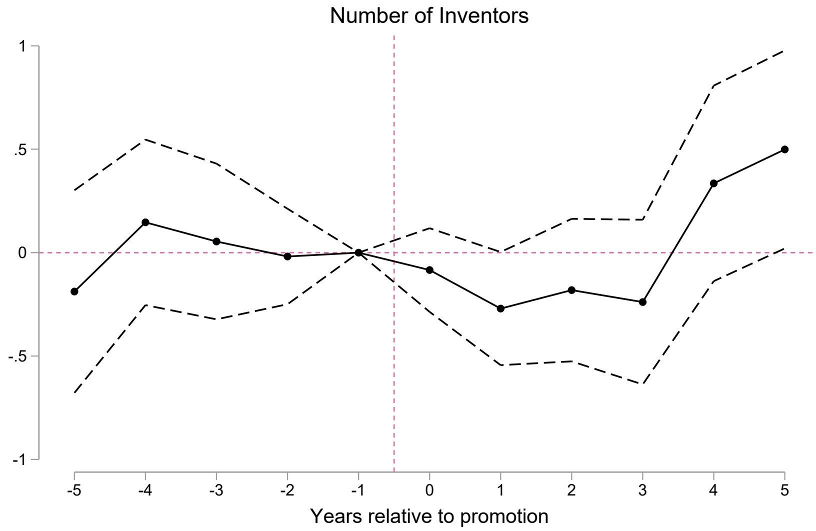

Notes: Event study estimates using the Abraham and Sun (2021) estimator and 95 percent confidence intervals are plotted. Regressions control for individual fixed effects, year fixed effects, and quadratic polynomial of potential experience.

Figure 3.14: Impact of Deanship on Patents with Varied Numbers of Inventors

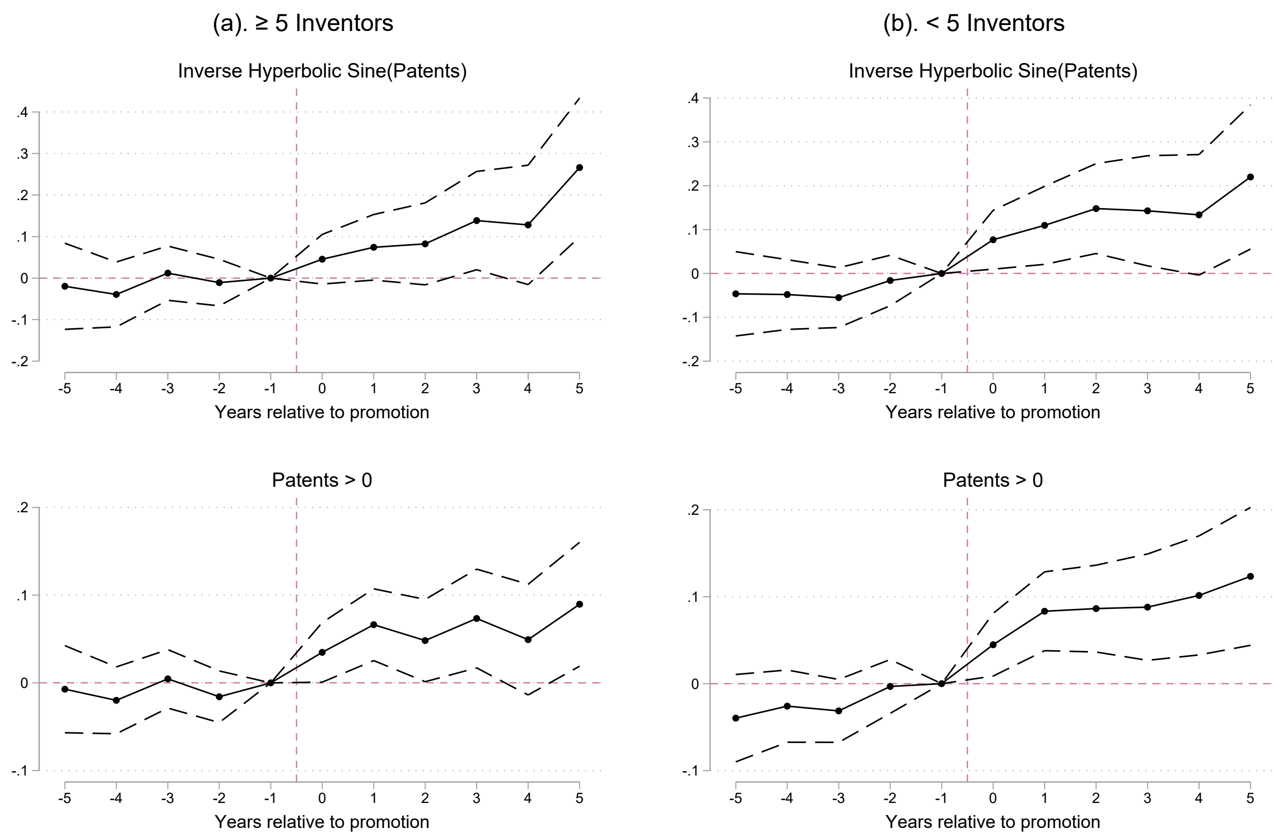

Notes: This figure reports the effects of deanship on two kinds of patents separately: the number of inventors above and below median. Event study estimates using the Abraham and Sun (2021) estimator and 95 percent confidence intervals are plotted. Regressions control for individual fixed effects, year fixed effects, and quadratic polynomial of potential experience.

Table 3.16: Deanship and Patent Applications: OLS Estimation

	(1) IHS(Patents)	(2) Patents>0	(3) IHS(Invention Patents)	(4) Invention Patents>0
5+ years before	-0.043	-0.028	-0.024	-0.013
	(0.058)	(0.026)	(0.056)	(0.026)
4 years before	-0.047	-0.014	-0.031	-0.004
	(0.046)	(0.021)	(0.044)	(0.021)
3 years before	-0.020	-0.006	-0.008	0.004
	(0.039)	(0.018)	(0.038)	(0.018)
2 years before	-0.024	-0.010	-0.022	-0.005
	(0.033)	(0.016)	(0.031)	(0.016)
Year of promotion	0.099***	0.057***	0.071*	0.048***
	(0.037)	(0.018)	(0.036)	(0.018)
1 year after	0.149***	0.098***	0.130***	0.095***
	(0.049)	(0.022)	(0.047)	(0.022)
2 years after	0.175***	0.081***	0.148***	0.079***
	(0.060)	(0.025)	(0.057)	(0.025)
3 years after	0.223***	0.106***	0.186***	0.097***
	(0.071)	(0.029)	(0.068)	(0.029)
4 years after	0.228***	0.119***	0.205**	0.117***
	(0.084)	(0.034)	(0.079)	(0.034)
5+ years after	0.326***	0.109***	0.280***	0.106***
	(0.099)	(0.039)	(0.095)	(0.039)
Average deanship effects				
	0.142***	0.080***	0.116***	0.070***
	(0.038)	(0.016)	(0.037)	(0.016)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Dep. Var. Mean	2.941†	0.425	2.507†	0.412

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are(1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, and (4) a dummy for having positive invention patent applications.Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

Table 3.17: Deanship and the Number of Co-inventors of Patents

	(1)
	Number of Inventors
5+ years before	-0.188
	(0.249)
4 years before	0.146
	(0.204)
3 years before	0.054
	(0.192)
2 years before	-0.019
	(0.118)
Year of promotion	-0.084
	(0.103)
1 year after	-0.270*
	(0.139)
2 years after	-0.181
	(0.175)
3 years after	-0.239
	(0.203)
4 years after	0.335
	(0.241)
5+ years after	0.499**
	(0.244)
Observations	29052
Individual FE	Yes
Year FE	Yes
Dep. Var. Mean	5.942

Notes: The dependent variable is the number of co-inventors of the patent associated with a dean. All regressions control for quadratic polynomial of potential experience. Event study coefficients are estimated using the interaction-weighted estimator proposed by Abraham and Sun (2021). Years more than 5 years before taking office are stacked to -5. Years more than 5 years after taking office are stacked to 5. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$.

Table 3.18: Local Anti-corruption Efforts and Deanship Effects: Alternative Data from Tencent

A. All Local Investigations

	(1)	(2)	(3)	(4)
	IHS(Patents)	Patent App>0	IHS(Invention Patents)	Inv Pat>0
Dean	0.162***	0.085***	0.134***	0.074***
	(0.051)	(0.023)	(0.049)	(0.023)
Dean*Local Investigations	-0.036*	-0.014	-0.033*	-0.014
	(0.020)	(0.009)	(0.019)	(0.009)

B. College-related Local Investigations

	(1)	(2)	(3)	(4)
	IHS(Patents)	Patent App>0	IHS(Invention Patents)	Inv Pat>0
Dean	0.174***	0.094***	0.140***	0.082***
	(0.046)	(0.020)	(0.044)	(0.020)
Dean*Local Investigations	-0.077*	-0.040**	-0.077*	-0.041**
	(0.042)	(0.020)	(0.041)	(0.020)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Individual FE*Local Investigations	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Year FE*Local Investigations	Yes	Yes	Yes	Yes
Dep. Var. Mean	2.941†	0.425	2.507†	0.412

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, and (4) a dummy for having positive invention patent applications. *Local Investigations* is the (inverse hyperbolic sine of) cumulative number of anti-corruption investigations for officials with level *Chuji* or above in the city since the start of 2013. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

Table 3.19: Anti-corruption Efforts and Patenting in Deans' Previous Fields of Expertise: Alternative Data from Tencent

A. All Local Investigations

	(1) IHS(Patents)	(2) Patent App>0	(3) IHS(Invention Patents)	(4) Inv Pat>0
Dean	-0.084***	-0.032**	-0.085***	-0.032**
	(0.026)	(0.013)	(0.026)	(0.013)
Dean*Local Investigations	0.007	0.008	0.010	0.008
	(0.012)	(0.005)	(0.012)	(0.006)

B. College-related Local Investigations

	(1) IHS(Patents)	(2) Patent App>0	(3) IHS(Invention Patents)	(4) Inv Pat>0
Dean	-0.098***	-0.027**	-0.098***	-0.026**
	(0.025)	(0.011)	(0.026)	(0.012)
Dean*Local Investigations	0.028	0.008	0.026	0.005
	(0.026)	(0.010)	(0.026)	(0.012)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Individual FE*Local Investigations	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Year FE*Local Investigations	Yes	Yes	Yes	Yes
Dep. Var. Mean	1.548†	0.238	1.347†	0.233

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are (1) the inverse hyperbolic sine of patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, and (4) a dummy for having positive invention patent applications. *Local Investigations* is the (inverse hyperbolic sine of) cumulative number of anti-corruption investigations for officials with level *Chuji* or above in the city since the start of 2013. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

Table 3.20: Anti-corruption Effects: Inspection Teams

	IHS(Patents)	Patents>0	IHS(Invention Patents)	Invention Patents>0
Dean	0.180***	0.085***	0.152***	0.078***
	(0.048)	(0.021)	(0.045)	(0.021)
Dean*Inspected	-0.148*	-0.047	-0.142*	-0.055
	(0.079)	(0.036)	(0.077)	(0.036)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Individual FE*Inspected	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Year FE*Inspected	Yes	Yes	Yes	Yes
Dep. Var. Mean	2.941†	0.425	2.507†	0.412

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, and (4) a dummy for having positive invention patent applications. *Inspected* equals 1 if a central inspection team has been sent to the province in or before the current year. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

Table 3.21: Anti-corruption Effects on Patents in Previous Fields of Study: Inspection Teams

	IHS(Patents)	Patents>0	IHS(Invention Patents)	Invention Patents>0
Dean	-0.083***	-0.029**	-0.081***	-0.027**
	(0.025)	(0.012)	(0.025)	(0.012)
Dean*Inspected	0.038	0.021	0.038	0.015
	(0.046)	(0.018)	(0.046)	(0.020)
Observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Individual FE*Inspected	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Year FE*Inspected	Yes	Yes	Yes	Yes
Dep. Var. Mean	1.548†	0.238	1.347†	0.233

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, and (4) a dummy for having positive invention patent applications. *Inspected* equals 1 if a central inspection team has been sent to the province in or before the current year. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

Table 3.22: Anti-corruption Effects: Before vs. After the Onset of the Campaign

	IHS(Patents)	Patents>0	IHS(Invention Patents)	Invention Patents>0
Dean	0.198***	0.089***	0.170***	0.081***
	(0.052)	(0.024)	(0.049)	(0.024)
Dean*Post	-0.142*	-0.033	-0.136*	-0.039
	(0.076)	(0.035)	(0.073)	(0.035)
observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Individual FE*Inspected	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Year FE*Inspected	Yes	Yes	Yes	Yes
Dep. Var. Mean	2.941†	0.425	2.507†	0.412

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, and (4) a dummy for having positive invention patent applications. *Post* equals 1 for the years after (including) 2013. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

Table 3.23: Anti-corruption Effects on Patents in Previous Fields of Study: Before vs. After the Onset of the Campaign

	IHS(Patents)	Patents>0	IHS(Invention Patents)	Invention Patents>0
Dean	-0.078***	-0.036***	-0.077***	-0.033**
	(0.026)	(0.013)	(0.026)	(0.013)
Dean*Post	0.029	0.038**	0.037	0.034*
	(0.042)	(0.018)	(0.041)	(0.019)
observations	9878	9878	9878	9878
Individual FE	Yes	Yes	Yes	Yes
Individual FE*Inspected	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Year FE*Inspected	Yes	Yes	Yes	Yes
Dep. Var. Mean	1.548†	0.238	1.347†	0.233

Notes: All regressions control for a quadratic polynomial of potential experience. Dependent variables are (1) the inverse hyperbolic sine of total patent applications, (2) a dummy for having positive patent applications, (3) the inverse hyperbolic sine of invention patent applications, and (4) a dummy for having positive invention patent applications. *Post* equals 1 for the years after (including) 2013. Standard errors in parentheses are clustered at the individual level. * $p < .1$, ** $p < .05$, *** $p < .01$. † Average number of patent applications.

Chapter 4: Anti-Corruption and Political Trust: Evidence from China

4.1 Introduction

Political trust, i.e., people's faith that the government will act in accordance with their interests,[1] has profound implications for regime legitimacy and state capacity. In a low-trust environment, citizens are less law-abiding and it is more difficult for the government to implement policy (Hetherington, 1998; Levi and Stoker, 2000; Zmerli, 2014), which can undermine institutional quality and socioeconomic development.[2] Among the most salient threats to political trust is corruption in government (Seligson, 2002; Chang and Chu, 2006; Gingerich et al., 2009), and many governments respond to corruption in order to win people's hearts and minds. As such, it is intriguing to investigate how political trust functions for institutions that aim to combat corruption.

Previous research on this interaction shows that trust facilitates the effectiveness of anti-corruption efforts (Bjørnskov, 2011; Kong et al., 2020). Related more closely to our inquiry, evidence from *democracies* suggests that revealing corruption lowers political trust, measured by citizens' weaker confidence in government and politicians' electoral losses (Guriev et al., 2021; Ferraz and Finan,

[1] Here, we use the term "government" to refer to a political authority that holds power to influence policy, including political institutions and political leaders. Throughout the paper, we use "political trust" and "trust in government" interchangeably.

[2] On political support, the pioneering work by Putnam et al. (1993) argues that trustful citizens are more law-abiding and more engaged in civic affairs. Hetherington (1998) finds that political trust relates to support for the incumbent president. State capacity refers to a government's ability to implement policy. The literature has demonstrated that political trust is highly relevant to state capacity. For instance, Sapienza and Zingales (2013) and Cullen et al. (2021) both find that trust in government affects the implementation of tax policy in the U.S. More recently, scholars have contended that political trust or identification with the nation facilitated success in combating the COVID-19 pandemic (Bollyky et al., 2022; Fukuyama, 2020; Rothstein, 2020; Van Bavel et al., 2022).

2008). This chapter attempts to investigate how anti-corruption affects political trust in *autocracies* (in this case, China), given that they often, perhaps surprisingly, have the most successful anti-corruption efforts (Carothers, 2022).

Our investigation builds upon a unique context: China's 2013 anti-corruption campaign, initiated shortly after President Xi Jinping rose to power. This campaign was in part a moral mobilization that aimed to reassert the Communist Party of China's (CPC) righteous image and earn public support (Javed, 2022), in the context of rising resentment about corruption since the earlier economic reforms (Wederman, 2004).[3] As President Xi put it during an interview with the *Wall Street Journal*, *"our Party owes its governing status to the support of the people... our people hate corruption more than anything else and we must act to allay their concerns"* (Xi, 2015). Therefore, it is useful to investigate to what extent the campaign actually influenced political trust, as it offers a useful lens to study the formation of public opinion in an autocracy like China.

Indeed, this campaign was arguably China's most massive anti-corruption drive in the post-Mao era, as attested by the large number of officials subject to corruption investigations and the series of reforms implemented to strengthen bureaucratic norms. Importantly, these anti-corruption efforts were extensively publicized, generating an influx of information about corruption to the general public. Given the discreet nature of corruption, at this point most Chinese people only had limited knowledge about corruption in government, so they may have used the new information on corruption to re-evaluate the government and update their political trust. However, it is not *ex-ante* clear in which direction this would change political trust, as different people are likely to read the same piece of information in different ways. Therefore, it remains an empirical question to identify

[3]There is debate on the anti-corruption campaign's true intentions. Some argue that the campaign serves as a tool to consolidate power. However, power consolidation is unlikely the only goal, for several reasons. First, the campaign investigated millions of officials, most of whom were low-rank officials. It is unlikely that a leader would have so many opponents, especially at the bottom tiers of government (Carothers, 2022). Second, the campaign included a range of institutional reforms, trying to straighten out previous corrupt bureaucratic norms (Ang, 2020), a task outside the scope of a pure political purge. Third, even if power consolidation were one goal of the campaign, dealing with corruption should help as it would be welcomed by the populace. All in all, the anti-corruption campaign was, to some extent, a real housecleaning.

the anti-corruption campaign's specific effects on political trust.

To guide our empirical investigation, we construct a simple model, following Dixit and Weibull (2007), to predict the heterogeneity of changes in political trust. The campaign offers *information* about corruption, from which a person can infer the government's honesty, which is linked to that person's political trust. However, there is a fundamental challenge in inference. Honest and corrupt governments can be observationally similar in terms of revealed corruption: a high level of corruption can be detected either because an honest government is willing and able to root out corruption, or because a corrupt government breeds rampant misdeeds. The ultimate judgment depends on a person's preferred *intepretation* that places different weights on two distinct inferences. Therefore, if a person is pro-government (anti-government), we predict that she will tend to read the campaign as indicating an honest (a corrupt) government, which enhances (lowers) her overall trust in government.

To empirically examine these predictions, we utilize a large, individual-level panel dataset based on the China Family Panel Study (CFPS). The panel data structure enables us to include individual fixed effects and thus track how an individual's trust varies as the anti-corruption campaign progressed. Specifically, our sample tracks political trust in 2012, 2014, and 2016, covering one pre-campaign period (2012) and two post-campaign periods (2014 and 2016). We merge the sample with data on city-level corruption investigations disclosed by anti-corruption authorities. Political trust is measured by trust in local government. We also collect rich individual and contextual characteristics to gauge the underlying mechanisms. Then, exploiting *temporal* and *regional* variation in corruption investigations, we implement a *difference-in-differences* strategy to estimate the anti-corruption campaign's effects on political trust. We present evidence in favor of our identifying assumption, which requires the trends in political trust to be similar between cities in the absence of the campaign (common trends). First, we show that conditional on province fixed effects, corruption investigations are orthogonal to predetermined factors that may drive the evolution of political trust, such as trust levels in 2012 and both the levels and growth rates of city characteristics. Second, using another survey dataset, we document a lack of differential pretrends in political trust between high- and low-investigation cities.[4]

Our main results imply that on average, political trust dropped immediately following the anti-corruption campaign. In 2014, a one-SD (= 38) increase in corruption investigations made an average individual 2.2 percentage points less likely to be trustful of the government. Although the effect was not as pronounced in 2016, further analyses indicate that the null average effects conceal intriguing heterogeneous effects. The campaign increased the level of political trust among highly educated individuals while decreasing it among those with lower levels of education. The main results are robust to a battery of robustness checks. First, they are virtually the same even if we control for trends related to previous corruption levels or to province-by-year fixed effects, suggesting the plausibility of the common trends assumption. Second, a permutation test confirms that our findings are unlikely to arise by chance. Third, we show that our results are robust to using an alternative estimator proposed by recent econometric literature on difference-in-differences designs with a continuous treatment variable (Callaway and Sant'Anna, 2021; de Chaisemartin et al., 2022b).

Beyond identifying the anti-corruption campaign's average effect on political trust, we also detect slight polarization. Significantly fewer people hold a moderate level of trust in government after the campaign. Instead, there are substantially more people with a very low level of political trust and slightly more people with a high level of political trust. These patterns echo our simple model's view that the campaign could have heterogeneous effects on political trust, since people interpret the same information brought by the campaign differently, depending on their priors about whether the government is honest or corrupt.

To corroborate this view, we first show that the campaign was indeed informative. A testable implication of informativeness is that the campaign's effects on political trust ought to be more pro-

[4]We are unable to conduct this pretrends check in the CFPS sample since we only have one pre-campaign period (2012).

nounced in scenarios where information about corruption was previously more limited. Tellingly, we show that the drop in political trust is greater for those respondents who did not pay attention to corruption news, and in provinces where corruption news was more covered up and diluted due to internet censorship and propaganda, indicating the campaign's informativeness.

Furthermore, we provide evidence that the pro- or anti-government cleavage drives different interpretations of information about corruption, which then bifurcates changes in political trust. We start by considering unpleasant experiences with government officials as direct determinants of individuals' priors, as these memories may make people develop anti-government sentiments. We find that people with unpleasant experiences indeed lower their trust to a greater extent after the campaign, indicating that they read corruption investigations as a sign of a corrupt government, confirming their negative impression of that government. We also probe the role of education as an indirect determinant of attitudes towards the government. A large state-building literature has stressed education's indoctrination function (e.g., Ramirez and Boli, 1987; Lott, 1999; Aghion et al., 2019), and particularly in an autocracy like China, education is leveraged as a tool of cultivating pro-government attitudes (Lott, 1999; Cantoni et al., 2017; Qi et al., 2022). Tellingly, our results highlight a monotonically decreasing relationship between education and declines in political trust: more educated people lower their trust to a lesser degree or even enhance it, especially those who are college-educated. These results are not driven by socioeconomic status, which is traditionally associated with education, possibly substantiating education's unique role in crafting attitudes. We supplement this interpretation by documenting that education's impacts are more pronounced in more Confucian provinces, where pro-government indoctrination could be more successful since Confucianism features similar norms (Acemoglu and Robinson, 2020, 2021b).

In addition, we rule out several alternative explanations for our findings. First, one concern is that political trust may have been lowered because anti-corruption crackdowns resulted in turmoil, undermining government performance (Wang, 2022). However, we find that the campaign does not influence people's perceptions of government performance, at least in the period under study. Second, the drop in political trust could be a consequence of changes in general trust, whether resulting from the campaign or from President Xi's other reforms that are correlated with the campaign. This appears unlikely, however, as we find that the campaign had null effects on individuals' trust in other groups (e.g., parents, strangers, and Americans). Third, one may conjecture that people would naturally report lower trust following the campaign, as they now think it is more legitimate than before to criticize the government, given that the government itself voluntarily disclosed scandals. Though we cannot fully rule out this possibility, we show that it is not the main driver of our story — the results survive, and are even more pronounced, if we exclude those who tended to see criticizing the government as taboo before (due to deference or fears) and so are more likely to lower their trust mechanically after the campaign. Taken together, our results are best explained by people updating their political trust upon receiving information about corruption provided by the campaign.

This chapter contributes to several strands of literature. First and foremost, it joins the burgeoning literature on trust in general (Arrow, 1972; Algan and Cahuc, 2010, 2014) and political trust in particular. Due to political trust's importance to a well-functioning government, voluminous studies have been devoted to understanding its formation, in which information about government performance is often considered a key factor (e.g., Chen and Yang, 2019; Saka et al., 2022; Khan et al., 2021; Shi, 2001). The link between corruption and political trust has received similar attention in this strand of literature. For instance, by analyzing a large cross-country dataset (including China), Guriev et al. (2021) show that increasing revelation of corruption scandals, induced by the expansion of 3G networks, reduces citizens' political trust on average, though interestingly they only find this effect in countries with uncensored internet. A major distinction between our papers is that in their context, citizens disclose corruption through social media, while in ours the government itself discloses corruption. We show that even *government-disclosed* information about corruption could also lead to a drop in political trust, and the drop is amplified by preexisting internet censorship, which complements Guriev et al. (2021). In addition, corruption is specifically associated with trust in crucial ways. Unsurprisingly, corruption has been documented to reduce general trust (Banerjee, 2016) and political trust (Anderson and Tverdova, 2003). But trust also affects corruption: Bjørnskov (2011) shows that legal quality is more effective at reducing corruption given a higher level of social trust. This chapter adds to the literature on the interrelationship between corruption and (political) trust by discussing the effects of unveiling information about corruption and efforts to reduce corruption on political trust, and how the effects critically depend on citizens' priors.

Another paper closely related to ours is Wang and Dickson (2022). Based on surveys conducted before and after China's anti-corruption campaign, they similarly find that the campaign reduces political trust, which they contend is because people were shocked by the great amount of corruption in government and updated their beliefs to discredit the government. We enhance these insights in two significant ways. First, we improve the identification. Wang and Dickson (2022)'s analysis relies on a repeated cross-sectional dataset, so they compare two different groups of individuals over time. If there were compositional changes in survey respondents after the campaign, this could bias their results. Also, they have to make the strong assumption that political trust measures are comparable between both sets of respondents. By contrast, we use a panel dataset to circumvent these concerns: a fixed group of individuals is studied, and by including individual fixed effects we trace how individuals' political trust evolves over time, ensuring better comparability. Second, we provide a more comprehensive view of the underlying mechanisms. Wang and Dickson (2022)'s argument is embedded in the informativeness channel of this chapter: the campaign offers information for people to re-evaluate the government. They implicitly assume that people interpret this information negatively, leading to lower political trust. By contrast, we propose and provide some evidence that interpretations could differ due to the pro- or anti-government cleavage. In this regard, we also offer, to the best of our knowledge, the first evidence of polarization in China.

Besides trust, we also add to the literature on public opinion at large. Existing studies have documented many ways in which information influences the electorate (Farzanegan and Hofmann, 2021; Enikolopov et al., 2018; Chong et al., 2015). We investigate the effects of information on public opinion in a non-electoral context, providing casual evidence on accountability in authoritarian regimes.

Last but not least, this chapter relates to the growing interest in China's anti-corruption campaign. Previous research predominantly focuses on the campaign's impacts on government officials' behaviors, such as rent seeking (Chen and Kung, 2019b; Ding et al., 2020; Kong et al., 2020), work incentives (Wang, 2022), and bureaucratic appointments (Wang, 2022). Few have examined citizens' responses, though notable exceptions include Jiang (2016) and Lai and Li (2021), who investigate the campaign's impacts on labor supply to bureaucracy. The current paper offers insights into how the campaign affects people's trust in government, a topic too important to miss given that the campaign is in part intended to garner support. Very interestingly, Kong and Qin (2021) find that the anti-corruption campaign increases entrepreneurship, and that the effects are larger in regions with a higher level of trust. Their finding points to the importance of trust in economic dynamism, particularly in the anti-corruption context. This chapter brings a new perspective to this question: anti-corruption can directly affect trust, which in turn may either amplify or weaken the joint effects of anti-corruption efforts and trust on entrepreneurship or other economic outcomes, depending on the direction of the impacts of anti-corruption on trust.

The rest of this chapter proceeds as follows. Section 4.2 introduces the background and provides a conceptual framework that guides our investigation. Section 4.3 presents the data. Section 4.4 introduces the empirical strategy. Section 4.5 reports the main results, followed by Section 4.6 discussing the underlying mechanisms. Section 4.7 concludes the paper.

4.2 Background and Conceptual Framework

In this section, we first introduce the main features of the anti-corruption campaign. Then, we build a conceptual model to illustrate how the campaign may influence political trust.

4.2.1 The Anti-Corruption Campaign

In 2013, shortly after President Xi Jinping took office, the Chinese government initiated an unprecedented anti-corruption campaign. Its onset was marked by President Xi's directive in the Second Plenary Session of the Eighteenth Central Commission for Discipline Inspection, January 2013. This campaign was arguably the greatest anti-corruption drive in the post-Mao era (Chen and Kung, 2019b; Ang, 2020; Carothers, 2022) due to several features. First, it was unusually long and in fact is still proceeding as of 2023. Past campaigns were dramatic but short. President Xi's anti-corruption campaign is undoubtedly a massive mobilization, but it is prolonged and seeks to establish a new normal. Second, the campaign has featured strict enforcement. By 2017, over 1.5 million officials had been investigated for misdeeds. Notably, many of them were high-ranking officials who often got leniency in the past.[5] For example, the campaign purged Zhou Yongkang, a former member of the Politburo Standing Committee, the most powerful body in the Chinese government. Third, the campaign was influential beyond the political arena. Existing research has documented the campaign's influence on a wide range of issues, including rent-seeking of local officials (Chen and Kung, 2019b), firm performance (Ding et al., 2020; Kong et al., 2020), and labor supply to the bureaucracy (Jiang et al., 2020; Lai and Li, 2021).

Given its high-profile nature, the campaign received a blaze of publicity. All of China's media outlets reported the campaign's achievements (e.g., the number of corruption investigations conducted, improvements in bureaucratic work ethics) and covered significant stories about some cor-

[5] The campaign was alleged to punish all corrupt officials, regardless of their seniority. In President Xi's own words, the campaign aimed to *"crack down on both tigers [high-rank officials] and flies [low-rank officials]"* (Xi, 2015).

rupt officials (Wang and Dickson, 2022; Zhuang, 2022). Notably, WeChat, China's most popular social media with 1.1 billion users in 2021, established a database that assembled all the disclosed information about government corruption, offering easy access to its users. Taken together, this campaign created an unprecedented influx of information about corruption, enabling many people to learn more about what used to be secret.

4.2.2 Conceptual Framework: Anti-Corruption and Political Trust

Political trust is people's belief that the government will act in accordance with their interests (Hetherington, 1998; Levi and Stoker, 2000; Zmerli, 2014). In forming political trust, people make judgments using available information about several aspects of government performance. Important among them is corruption, since it could harm public interests severely and since people ought to be concerned about whether their government is honest.

Therefore, as the anti-corruption campaign distributes a great amount of information about corruption to the public, people may use it to (re-)evaluate the government and update their political trust. However, it is not *ex-ante* clear how political trust responds to the anti-corruption campaign, since different people could interpret the same piece of information in different ways. To fix these ideas, we build a simple model in the spirit of Dixit and Weibull (2007) to consider the (heterogeneous) impacts of the anti-corruption campaign on political trust.[6]

Model. Suppose that individual i's political trust depends on government honesty, s.[7] Although s is not directly observed, one can infer it from x, their information about corruption, which is increased by the anti-corruption campaign. Individual i has her own priors about s and x, denoted

[6]Dixit and Weibull (2007)'s original model aims to explain why people's opinions on monetary policy polarize even though they observe the same economic conditions, e.g., inflation. They suggest that people rationally update their beliefs about the real state of the world and form policy opinions in a Bayesian fashion. However, different priors make people weight inferences from the same information differently in the process of belief updating, leading to polarization.

[7]To substantiate our focus on the impact of corruption, we abstract away from the reality that political trust also depends on other factors. Nonetheless, our subsequent empirical analysis takes into account alternative channels through which political trust is influenced (see Section 4.6.3).

by $\bar{s}_i$ and $\bar{x}_i$. Given the anti-corruption campaign's unprecedented nature (see Section 4.2.1), we assume that it reveals government corruption more thoroughly than expected, i.e., $x > \bar{x}_i$.

Inferring s from x relies on their functional relationship. We suppose that individual i perceives the following relationship:

$$x = \bar{x}_i + |s - \bar{s}_i|. \tag{4.1}$$

That is, individual i considers the difference between observed information about corruption revealed by the campaign and their priors of the level of corruption, $x - \bar{x}_i$, to be due to the deviation of unobserved government honesty from her priors, $s - \bar{s}_i$. Thus, she can extract knowledge about s upon observing x. However, it is worth noting that the relationship between x and s is not monotonic. Specifically, each x is compatible with two opposite narratives: (i) s is high — an honest government is able and willing to combat corruption, or (ii) s is low — a corrupt government breeds rampant misdeeds.

Such non-monotonicity makes it challenging for an individual to identify the underlying s. As illustrated by Figure 4.1, when an x is observed, individual i can draw two inferences about the level of government honesty, s:

$$s_i^h = \bar{s}_i + (x - \bar{x}_i) \tag{4.2}$$

$$s_i^c = \bar{s}_i - (x - \bar{x}_i), \tag{4.3}$$

where s_i^h refers to an honest government but s_i^c corresponds to a corrupt government. A Bayesian individual would draw her ultimate inference of s by weighting s_i^h and s_i^c. The weights depend on an individual's preexisting belief about whether the government is honest or corrupt. Let $p_i \in [0,1]$ denote individual i's believed probability that the government is honest, while $1 - p_i$ denotes her believed probability that the government is corrupt. Despite its authoritarian system, such a pro- or anti-government cleavage exists in modern China's ideological spectrum (Pan and Xu, 2018).[8]

[8] As noted by previous literature, the ideological spectrum may have multiple dimensions (e.g., liberalism, nation-

Figure 4.1: Relationship Between x and s

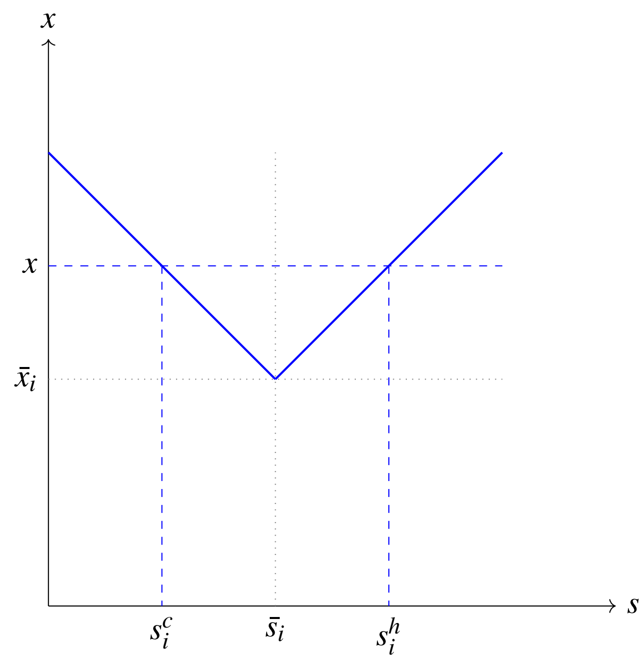

Therefore, upon receiving information about corruption, x, individual i revises her perceived government honesty and thus her political trust. This process is captured by her posterior $\tilde{s}_i$, which is determined in the following way:

$$\tilde{s}_i = p_i s_i^h + (1 - p_i) s_i^c \tag{4.4}$$

$$\Delta s_i = \tilde{s}_i - \bar{s}_i = \underbrace{(x - \bar{x}_i)}_{\text{informativeness}} \times \underbrace{(2p_i - 1)}_{\text{interpretation}} . \tag{4.5}$$

Inspecting Equation 4.5 indicates that the anti-corruption campaign, by increasing x, would influence political trust differently across individuals. How and how much the campaign changes one's political trust, i.e., the sign and magnitude of Δs_i, depends on two factors. The first one is the *informativeness* of the anti-corruption campaign, $x - \bar{x}_i$. Holding p_i constant, political trust would vary

alism, market economy, etc.) or reflect socioeconomic differences across groups. As such, the pro- or anti-government cleavage may correlate with or result from cleavages in other dimensions and differences in a range of attributes. In this regard, the pro- or anti-government cleavage should be interpreted as a reduced-form representation of underlying factors that make people's ideologies differ. In Section 4.5, we analyze several important factors.

more significantly if the campaign offers more information about corruption than what individual i had anticipated. However, it is unclear whether such informativeness would increase or decrease political trust, since different people may process the same information differently. This points to a second critical factor — an individual's *interpretation*, $2p_i - 1$. Conditional on informativeness, a relatively pro-government individual ($p_i > 1/2$) would enhance her trust, i.e., $\Delta s_i > 0$, since she overweights s_i^h, the inference aligned with her prior that the government is honest and catches many corrupt bureaucrats. By contrast, a relatively anti-government individual ($p_i < 1/2$) would lower her trust, i.e., $\Delta s_i < 0$, because she overweights s_i^c that considers a high x as confirming her prior that the government is corrupt. The following proposition summarizes the heterogeneous effects of the anti-corruption campaign on political trust.

Proposition 1. *By revealing unanticipated information about corruption, $x > \bar{x}_i$, the anti-corruption campaign increases (decreases) political trust among individuals who hold priors that the government is honest (corrupt), i.e., $p_i > 1/2$ ($p_i < 1/2$), and the increase (decrease) is greater if the campaign is more informative, i.e., $x - \bar{x}_i$ is larger.*

Remarks. In the model, we have made a simplifying assumption: the anti-corruption campaign (x) affects political trust by intervening in perceived government honesty ($\bar{s}_i$). However, it is likely that the campaign or other contemporaneous shocks operate through other channels, such as government performance and general willingness to trust. Our empirical investigations take into account alternative explanations of the campaign's impacts on political trust (see Section 4.6.3).

Proposition 1 underscores the importance of informativeness and interpretation in shaping political trust changes, advancing the previous literature's insights (Wang and Dickson, 2022) in two main directions. First, we allow for flexibility in people's interpretations of information about corruption. Wang and Dickson (2022) similarly contend that people would use information about corruption brought by an anti-corruption campaign to update political trust. However, they hypothesize that the campaign should *suppress* political trust, since people would be shocked by the many corruption scandals and become pessimistic about officials' integrity (their Hypotheses 1 and 2). Thus, they implicitly assume that people *negatively* interpret corruption disclosed by the campaign, an assumption which may not be warranted for the entire population. Second, we enrich the role that priors play in shaping the campaign's impacts. In Wang and Dickson (2022), priors govern the campaign's informativeness and then lead to heterogeneity in impacts: they argue that the campaign should have a smaller (larger) effect of reducing political trust if an individual had more (less) knowledge of government corruption previously (their Hypothesis 3). Our framework embeds this informativeness channel through the term $x - \bar{x}_i$ in Equation 4.5. Moreover, with the term $2p_i - 1$ in Equation 4.5, we consider the possibility that priors can bifurcate opinions regarding the same information. This is not rare in politics, as documented by the massive literature on public opinion and political polarization (Adena et al., 2015; Bullock, 2009; Bisgaard, 2015; Bisgaard and Slothuus, 2018; Spenkuch et al., 2021; Gerber and Green, 1999).[9]

Proposition 1 guides much of our subsequent analysis. We first investigate the anti-corruption campaign's effect on an average individual's political trust. Then, we probe into two underlying mechanisms implied by Proposition 1: informativeness and interpretation.

4.3 Data

4.3.1 Local Information about Corruption

We hypothesize that the anti-corruption campaign could affect political trust since it results in an influx of information about government corruption. To empirically examine this hypothesis, we need to measure the amount of information available to people.

We first measure the regional variation in information about corruption, using a comprehensive database of virtually *all* the corruption investigations disclosed by the anti-corruption authorities

[9]For instance, Adena et al. (2015) document that Germans with high (low) anti-Semitic predispositions were persuaded (dissuaded) by Nazi propaganda. Gerber and Green (1999) argue that "observers with different preconceptions interpret the same piece of evidence in ways that conform to their *initial views*".

between 2011 and 2016 (Wang and Dickson, 2022).[10] The database was developed by China's internet tycoon, Tencent, and it was widely circulated over Tencent's WeChat, the most popular social media in China, with over 1.1 billion users in 2021. In the database, people can easily check which officials have been investigated in their cities and access related stories. Therefore, people can be well exposed to information brought in by the anti-corruption campaign. More importantly, by building upon official sources, the database arguably includes the majority of publicly available information about corruption in Chinese society, across possible transmission channels (e.g., news reports, internet, and word of mouth).[11]

We compute the cumulative number of corruption investigations for each city p as of time τ, starting from 2011, denoted by D_p^τ. Due to the campaign's publicity and the database's popularity and coverage, we consider D_p^τ a proxy for the amount of information about corruption received by local people.

Figure 4.2 presents a monthly series of cumulative corruption investigations at the national level. There were barely any corruption investigations disclosed before President Xi's anti-corruption campaign. Immediately after the campaign's onset in January 2013, however, corruption investigations sharply increased, and the increase halted in 2016. This trend confirms that the campaign did lead to the disclosure of a great amount of information about corruption.

4.3.2 China Family Panel Study Sample

Sample Construction

We conduct an analysis relying on the China Family Panel Study, a nationally representative biennial household survey starting in 2010. We construct a *balanced panel dataset* using CFPS

[10] See https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/9QZRAD for further details about the dataset (last accessed on May 17, 2020).

[11] The information circulated through other channels can be ultimately reflected by official sources that are then collected by the database.

Figure 4.2: Time Series of Cumulative Investigations

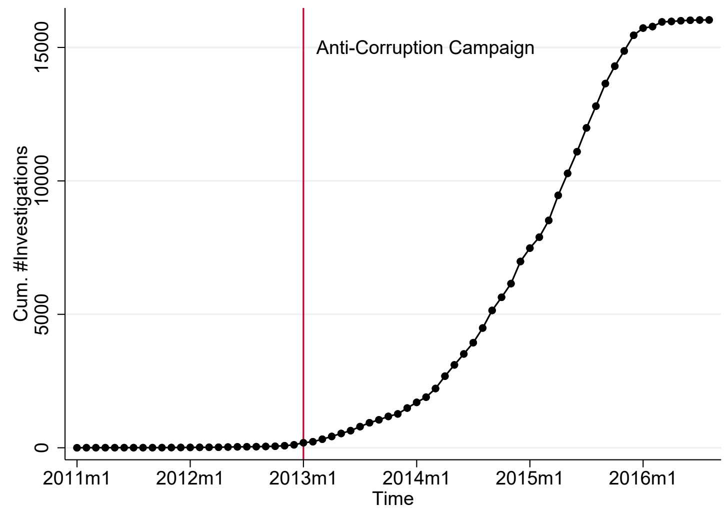

Note: Data from the Tencent database (Wang and Dickson, 2022). The vertical line marks the onset of the anti-corruption campaign (January 2013).

data from 2012, 2014, and 2016, as starting in 2012 the survey included questions on political attitudes. We elicit some measures from the 2010 survey. Our sample only includes individuals who (i) responded to all three waves of surveys, (ii) were born between 1950 and 1990, and (iii) never migrated between 2012 and 2016. These restrictions allow us to concentrate on a fixed group of individuals who are mature enough to form meaningful political attitudes, and they also enable us to correctly match city-level information about corruption. We end up with a balanced panel dataset with 11,870 individuals from 121 cities, all of whom were consecutively surveyed in 2012, 2014, and 2016. We discuss the main variables below. Summary statistics are reported in Table 4.1.

Political Trust

Measurement. CFPS elicits political trust based on the following question:

Please rate to what extent you trust the local government cadres. Answers range from 0–10 (0 = lowest trust, 10 = highest trust).

Figure 4.3 displays the distribution of the political trust score by year. The distribution is bell-shaped: responses pool in the middle of the scale (mean, median ≈ 5), which is not uncommon in the literature on trust or public opinion more broadly (Gaziano and McGrath, 1986; Petty and Krosnick, 2014; Allen and Birch, 2015).[12] To break ties in judgment and to capture unambiguous political trust, we create a dummy variable for high-level political trust, which equals one if the reported score is greater than or equal to 5. Note that this formulation adopts a relatively conservative definition of *distrust* (for a score ≤ 4), creating a more powerful test for whether the anti-corruption campaign has lowered political trust.

Figure 4.3: Distribution of Political Trust

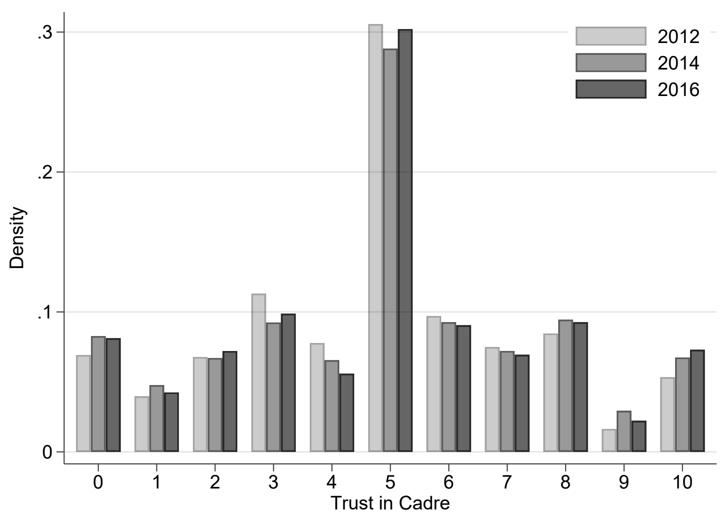

Note: Data from CFPS of 2012, 2014, and 2016.

Scope. What does our political trust measure capture? According to the wording of the question, it specifically measures people's trust in local government. CFPS does not ask for people's trust in central government or other (higher-level) sections of the Chinese state like the legislature, the military, or even the ruling party. However, Chen and Yang (2019) find a strong correlation between trust in local and central governments, suggesting that to some extent, examining how trust

[12]There are several possible reasons for this "overstuffed middle" problem (Allen and Birch, 2015). It could mean that respondents are indeed centrist, that they do not have the information or knowledge to make a deterministic judgment, or that they are ambivalent.

in local government changes might inform us about changes in general political trust. Moreover, even if trust in higher-level sections of the Chinese state were surveyed, trust in local government might arguably provide the most useful data, since it is more tolerated to criticize the local government than higher-level sections (Lorentzen, 2014; King et al., 2013; Qin et al., 2017).[13]

Validity. Due to the self-reporting nature of the CFPS, there may be concerns about the validity of our political trust measure. People may still be reluctant to report their political trust truthfully, despite the fact that the trustee is local government and that it is legitimate for respondents to be more outspoken. Were such misreporting salient and associated with underlying determinants of corruption investigations, our results could pick up a spurious impact of the anti-corruption campaign on political trust. However, our design and results try to alleviate this self-censorship concern in several ways.

First, the item response rate for the question on political trust is high (e.g., 96.75% in 2012), indicating that it is unlikely for people to be intimidated into silence. In addition, Figure 4.3 shows that the distribution of political trust is not skewed towards "politically correct" high trust, and many respondents reported low trust.

Second, our political trust measure exhibits reasonable patterns with high internal consistencies. Appendix Figure 4.10 correlates pre-campaign political trust (measured in 2012) with *predetermined* negative experiences with local governments (measured in 2010). As expected, political trust is lower for respondents who had been unfairly treated by local cadres, had conflicts with cadres, encountered slack cadres, or had been asked for bribes. Moreover, Figure 4.11 shows that higher political trust is strongly associated with more (peaceful) political engagement. Specifically,

[13] One piece of anecdotal evidence is from the China General Social Survey (CGSS), which did elicit trust in both local and central governments between 2010 and 2012. Only 3.71 percent of respondents reported distrust in the central government, marking a sharp contrast to the 15.53 percent of respondents were distrustful of the local government. In addition, the response rate for trust in central government was 20 percentage points lower than that for trust in local government. Taken together, it is likely that people are more outspoken when judging the local government than when judging the central government. However, one needs to take this assertion with caution: perhaps the central government is genuinely more trustworthy. We do not use CGSS our analysis since due to unknown reasons, it stopped eliciting political trust after 2012, around the time of the anti-corruption campaign.

a trustful individual is more likely to vote in grassroots elections (in line with Tao et al., 2011), and she is more prone to resolve dissent (if any) via petitions rather than protests. Figure 4.11 also shows that higher political trust has led to more optimistic evaluations of China's social governance in terms of handling challenges from corruption, environmental issues, inequality, etc.

Third, the panel data structure enables us to include individual fixed effects, which removes any individual-invariant heterogeneity. This would largely purge misreporting if it is relatively stable, which may be plausible given that we have a six-year short panel. In our empirical investigations we also flexibly control for differential trends in political trust, to absorb possible time-varying reporting patterns (see Section 4.4).

Besides self-censorship, another concern is about the comparability of our political trust measure. This is especially concerning in studies using cross-sectional or repeated cross-sectional data, which have to assume that the trust measures are comparable between different individuals and/or times. However, the panel structure enables us to get around this problem. First, we are able to trace changes in political trust within the *same* individual. Second, we can avoid disturbances from compositional changes. Repeated cross-sectional studies have to contrast different (and likely incomparable due to selection into survey response) individuals over time. But if people self-select into and out of response groups because of the anti-corruption campaign, it is unclear how the campaign causes a change in political trust. By contrast, the panel data concentrate on a fixed group of individuals.

Additional Variables

Besides corruption investigations and political trust, Table 4.1 presents rich variables that we collect from CFPS as well as other sources. We briefly describe them below and will introduce them in greater detail when they become pertinent to our analysis.

Attitudinal Variables. As shown in Panel (B), we include trust in other groups (parents, strangers, and Americans) and perceived government performance, which helps us pin down accurate interpretations of our results (Section 4.6.3).

Experiences with the Government. To investigate the underlying mechanisms (Section 4.6), we exploit information about individual experiences with the government: whether respondents have been unfairly treated by local cadres, had conflicts with cadres, encountered slack cadres, and have been asked for bribes (see Panel (C)).

Covariates. Our sample also contains a range of individual characteristics (see Panel (D)) including birth cohort, gender, Han ethnicity, *hukou* status, Communist Party membership, educational attainment, employment in state sectors, parental educational attainment, and parental Communist Party membership. All of them can be conducive to the formation of political trust.

Other Variables. Panel (E) presents several variables we use for robustness checks and for disentangling different mechanisms. We provide a description of them below in order of Panel (E).

Past Corruption. To measure a city's past level of corruption, we use the ETC index proposed by Cai et al. (2011). ETC refers to Chinese firms' spending on entertainment and travel costs, amenities which are commonly used to bribe government officials. Therefore, the literature has used ETC as a proxy for local corruption in China (e.g., Fang et al., 2019; Jiang, 2016). Cai et al. (2011)'s ETC indices (share of ETC) available between 2002 and 2004. We take the average as a measure of a city's past level of corruption.

Attention to Corruption News. The CFPS 2010 survey asks individuals if they have ever paid attention to news about corruption. We code people's answers into a dummy variable.

Censorship and Propaganda. The anti-corruption campaign offers information about corruption which used to be unavailable to the public. To capture to what extent such information has been blocked, we use Qin et al. (2017)'s measures of censorship and propaganda at the provincial level: the share of deleted Weibo posts (Weibo is "Chinese Twitter") and the share of government users on Weibo.

Confucianism. Confucian philosophy has enduring influences on China's political traditions

(Bell, 2010; Jiang, 2016; Pan and Xu, 2018). To gauge its implications for political trust, we use the number of Confucian temples (in log form) to capture the city-level Confucian norms, following previous literature (Kung and Ma, 2014; Chen et al., 2020; Alm et al., 2022).

Special Background. Some special backgrounds may have unique impacts on individuals' political attitudes and behavior. We take into account three types of backgrounds. First, military services could foster loyalty to the state, which may spill over to others in the family. So, we code an individual to have such a background if anyone in her family has ever served in the military. Second, people may see criticizing the government as a taboo if they or a family member were purged in the Communist Revolution (1950s) or the Cultural Revolution (CR, 1966–76). Based on experiences elicited in the CFPS 2010 survey, we create a dummy variable that equals one if one's family was purged in these Revolutions.[14] Third, witnessing intense state violence can also lead to fears of criticizing the government. The Cultural Revolution was the most violent episode in China's modern history. Therefore, we create a dummy variable that equals one if the respondent is from a city with above-median CR casualties (data from Walder, 2014), or if she experienced the CR during the "impressionable years."[15]

[14]The CFPS 2010 survey directly asks if one's family was assigned a bad class label (landlords, rich peasants, and capitalists), was sent to the May Seventh Cadre School (where intellectuals were re-educated through labor in the Cultural Revolution), or experienced the Sent-Down Youth Movement. Intentional misreporting is not very likely. Also using CFPS data, Alesina et al. (2020) document that for nearly 94.3% of Chinese households, all household members report identical class labels.

[15]The impressionable years hypothesis (Alwin and Krosnick, 1991; Giuliano and Spilimbergo, 2014; Cotofan et al., 2020) suggests that the period between the ages of 18 and 25 is a critical period for the formation of political attitudes. For instance, Giuliano and Spilimbergo (2014) find that people are more supportive of redistribution policies if they have experienced recessions in impressionable years.

Table 4.1: Summary Statistics

	Measuring Time	Source	Obs.	Mean	SD	Min	Max
Panel (A): Anti-Corruption							
# Corruption investigations	2012, 14, 16	0	35610	28.652	37.857	0	262
Panel (B): Attitudes							
Trust: cadres	2012, 14, 16	1	35610	4.879	2.623	0	10
Trust: parents	2012, 14, 16	1	35536	9.313	1.475	0	10
Trust: strangers	2012, 14, 16	1	35528	1.902	2.066	0	10
Trust: Americans	2012, 14, 16	1	34858	2.126	2.402	0	10
Government performance	2012, 14, 16	1	34628	3.411	0.909	1	5
Panel (C): Experiences							
Experience: unfairly treated by cadres	2012	1	35466	0.089	0.285	0	1
Experience: having conflicts w/cadres	2012	1	35520	0.035	0.184	0	1
Experience: slack cadres	2012	1	35463	0.125	0.331	0	1
Experience: asked for bribes	2012	1	35478	0.066	0.249	0	1
Panel (D): Covariates							
Birth cohort	2012	1	35610	1967.386	10.539	1950	1990
Male	2012	1	35610	0.467	0.499	0	1
Han ethnicity	2012	1	35610	0.925	0.263	0	1
Urban	2012	1	35610	0.463	0.499	0	1
Years of schooling completed:2012	2012	1	35610	6.976	4.729	0	22
Communist Party member	2012	1	35610	0.075	0.263	0	1
State sector employee	2012	1	35610	0.080	0.272	0	1
Degree completed, father	2010	1	35610	4.512	4.233	0.460	16
Degree completed, mother	2010	1	35610	2.594	3.496	0.460	16
Communist Party member, father	2010	1	35610	0.161	0.367	0	1
Communist Party member, mother	2010	1	35610	0.023	0.149	0	1
Panel (E): Other Variables							
ETC index	2002–04	2	18810	0.012	0.005	0.003	0.025
Attention to corruption news	2010	1	35610	0.231	0.422	0	1
Share of Delted Weibo posts	2009–13	3	33525	0.182	0.048	0.120	0.285
Share of Govt. Weibo users	2009–13	3	33525	0.041	0.010	0.025	0.061
ln(# Confucian temples)	Ming-Qing	4	29910	4.584	1.526	0	8.677
Military Family	2010	1	35610	0.063	0.243	0	1
Family purged in Revolutions	2010	1	35610	0.127	0.333	0	1
Witnessed violent CR	2010	5	35610	0.626	0.484	0	1

Source: 0 = Wang and Dickson (2022), 1 = CFPS, 2 = Cai et al. (2011), 3 = Qin et al. (2017), 4 = Chen and Kung (2019b), 5 = defined by authors based on Walder (2014).

4.3.3 Selection Into Balanced Panel

We construct a balanced panel from the original unbalanced panel of CFPS, which enables us to track the political trust of a given set of individuals overtime. However, individuals in the balanced sample, who responded to all the waves of CFPS and responded to all the specific questions about political trust, may have different characteristics from the original representative sample. This could undermine internal validity if non-response was not orthogonal to treatment or political trust and affect external validity if the balanced panel is not representative of the population. To investigate the potential problem of sample selection, we test the balance of individual characteristics and regional features between selected and non-selected individuals. The summary statistics by group are reported in Table 4.2.

Reassuringly, the analysis indicates that there is no systematic difference in baseline political trust or the proxy of anti-corruption efforts between selected and unselected individuals. The cumulative number of anti-corruption investigations in 2014 is higher for the selected sample, but the difference seems to be economically insignificant. However, selected individuals differ from the unselected sample in several demographic and political characteristics. First, selected individuals tend to be younger, have higher levels of education and come from higher socio-economic backgrounds. Second, selected individuals are also more likely to be pro-government, as evidenced by higher rates of employment in the state sector and Communist Party membership. Selected individuals also tend to live in regions with more Confucian temples. Third, selected individuals are more likely to have paid attention to corruption news before the anti-corruption campaign. In addition, selected individuals come from regions with weaker government censorship and propaganda (measured by the share of deleted Weibo posts and the share of government users on Weibo). However, these differences are not economically significant.

Overall, the balanced panel used for our analysis may not be entirely representative of the CFPS or Chinese individuals. However, we observed no direct selection problem concerning treatment status or political trust since treatment intensity and baseline political trust were similar between the balanced panel and the unselected sample. To probe into the remaining concern about unbalanced covariates, we also examine the correlation between political trust in 2012 and the unbalanced covariates, as reported in Appendix Table 4.9. We found that while some variables predicted political trust, there was no uniform direction of the potential bias. For example, while the selected individuals were more likely to be members of the Communist Party, associated with higher political trust, they were also more likely to be urban residents, negatively correlated with political trust. These patterns suggest that political attitudes are not the main reason for the observed difference between the selected and the unselected sample. That said, we should still interpret our findings with caution. For instance, we find that a higher fraction of the selected individuals paid attention to corruption news even before the anti-corruption campaign. As we will show later, those who did not pay attention to corruption news before the campaign were more negatively impacted by the campaign in terms of political trust (Figure 4.6).

Table 4.2: Comparing the Balanced Panel Sample and the Unselected Sample

	(1)	(2)	(3)	(4)	(5)	(6)
				Orthogonality Test		
	Full Sample	Selected	Unselected	Mean Diff.	P-value	Obs.
Panel (A): Anti-Corruption						
# Corruption investigations, 2014	27.838	28.136	26.054	2.082	0.001	13854
	(0.219)	(0.240)	(0.526)	(0.626)		
# Corruption investigations, 2016	57.137	57.116	57.262	-0.146	0.893	13854
	(0.380)	(0.413)	(0.974)	(1.085)		
Panel (B): Attitudes in 2012						
Trust: cadres	4.829	4.828	4.840	-0.012	0.864	13405
	(0.022)	(0.023)	(0.065)	(0.068)		
Trust: parents	9.095	9.131	8.827	0.304	0.000	13413
	(0.014)	(0.015)	(0.049)	(0.045)		
Trust: strangers	2.041	2.049	1.980	0.069	0.222	13399
	(0.018)	(0.019)	(0.054)	(0.056)		
Trust: Americans	2.365	2.369	2.335	0.034	0.614	12994
	(0.021)	(0.023)	(0.064)	(0.068)		
Government performance	3.448	3.446	3.465	-0.019	0.465	13058
	(0.008)	(0.009)	(0.024)	(0.025)		
Panel (C): Experiences by 2012						
Experience: unfairly treated by cadres	0.091	0.089	0.101	-0.012	0.114	13403
	(0.002)	(0.003)	(0.008)	(0.008)		
Experience: having conflicts w/cadres	0.037	0.035	0.049	-0.014	0.006	13432
	(0.002)	(0.002)	(0.005)	(0.005)		
Experience: slack cadres	0.125	0.125	0.126	-0.000	0.969	13404
	(0.003)	(0.003)	(0.008)	(0.009)		
Experience: asked for bribes	0.067	0.066	0.071	-0.005	0.488	13407
	(0.002)	(0.002)	(0.006)	(0.007)		
Panel (D): Individual and Family Covaraites						
Birth cohort	1967.147	1967.386	1965.717	1.669	0.000	13854
	(0.090)	(0.097)	(0.250)	(0.258)		
Male	0.466	0.467	0.461	0.006	0.633	13854
	(0.004)	(0.005)	(0.011)	(0.012)		
Han ethnicity	0.921	0.925	0.896	0.029	0.000	13854
	(0.002)	(0.002)	(0.007)	(0.007)		
Urban	0.429	0.444	0.340	0.104	0.000	13843
	(0.004)	(0.005)	(0.011)	(0.012)		
Years of educ.	6.531	6.739	5.290	1.449	0.000	13854
	(0.041)	(0.044)	(0.105)	(0.116)		
Total Observations	13854	11870	1984			

Note: This table summarizes individual and regional characteristics at the individual level. Individuals that are selected into the balanced panel for the main analysis are compared to the original CFPS sample and individuals who are not selected into the balanced panel. Columns (4)–(6) test the balance of covariates between the selected and unselected individuals. Standard errors are reported in parentheses.

Table 4:2: Comparing the Balanced Panel Sample and the Unselected Sample, Continued

	(1)	(2)	(3)	(4)	(5)	(6)
				Orthogonality Test		
	Full Sample	Selected	Unselected	Mean Diff.	P-value	Obs.
Panel (D): Individual and Family Covaraites, Continued						
Communist Party member	0.070	0.075	0.037	0.038	0.000	13854
	(0.002)	(0.002)	(0.004)	(0.006)		
State sector employee	0.075	0.080	0.041	0.039	0.000	13854
	(0.002)	(0.002)	(0.004)	(0.006)		
Years of educ., father	4.486	4.512	4.221	0.291	0.026	13012
	(0.037)	(0.039)	(0.123)	(0.131)		
Years of educ., mother	2.503	2.594	1.817	0.776	0.000	13442
	(0.030)	(0.032)	(0.075)	(0.092)		
Communist Party member, father	0.152	0.161	0.095	0.066	0.000	13854
	(0.003)	(0.003)	(0.007)	(0.009)		
Communist Party member, mother	0.022	0.023	0.015	0.007	0.034	13854
	(0.001)	(0.001)	(0.003)	(0.004)		
Panel (E): Other Variables						
ETC index	0.012	0.012	0.012	0.000	0.647	7144
	(0.000)	(0.000)	(0.000)	(0.000)		
Attention to corruption news	0.223	0.231	0.171	0.060	0.000	13854
	(0.004)	(0.004)	(0.008)	(0.010)		
Share of Delted Weibo posts	0.183	0.182	0.190	-0.008	0.000	13095
	(0.000)	(0.000)	(0.001)	(0.001)		
Share of Govt. Weibo users	0.041	0.041	0.042	-0.001	0.015	13095
	(0.000)	(0.000)	(0.000)	(0.000)		
ln(# Confucian temples)	4.557	4.584	4.406	0.178	0.000	11751
	(0.014)	(0.015)	(0.037)	(0.039)		
Military Family	0.063	0.063	0.064	-0.001	0.866	13854
	(0.002)	(0.002)	(0.005)	(0.006)		
Family purged in Revolutions	0.124	0.127	0.109	0.018	0.026	13854
	(0.003)	(0.003)	(0.007)	(0.008)		
Witnessed violent CR	0.630	0.626	0.656	-0.030	0.010	13854
	(0.004)	(0.004)	(0.011)	(0.012)		
Total Observations	13854	11870	1984			

Note: This table summarizes individual and regional characteristics at the individual level. Individuals that are selected into the balanced panel for the main analysis are compared to the original CFPS sample and individuals who are not selected into the balanced panel. Columns (4)–(6) test the balance of covariates between the selected and unselected individuals. Standard errors are reported in parentheses.

4.4 Empirical Strategy

4.4.1 Econometric Model

The combination of regional and temporal variations in corruption investigations motivates a *difference-in-differences* strategy. Thus, to assess the anti-corruption campaign's impact on political trust, we rely on the following specification:

$$y_{ipt} = \alpha + \beta_1 \left(D_p^{14} \times T_t^{14}\right) + \beta_2 \left(D_p^{16} \times T_t^{16}\right) + (X_i' \cdot \mu_t)\gamma + \lambda_i + \mu_t + \varepsilon_{ipt}. \tag{4.6}$$

Subscript i indexes individuals, p indexes cities, and t indexes years (2012, 2014, and 2016). We include individual fixed effects λ_i and year fixed effects μ_t. y_{ipt} is the political trust measure. D_p^{14} and D_p^{16} are city p's cumulative corruption investigations as of 2014 and 2016, respectively. T_t^{14} and T_t^{16} are the dummy variables for years 2014 and 2016. X_i is a set of individual characteristics, including indicators of birth cohort, gender, Han ethnicity, *hukou* status, Communist Party membership, educational attainment, employment in state sectors, parental educational attainment, and parental Communist Party membership. As they are mostly invariant over time, we interact these characteristics with λ_i to allow for differential impacts on political trust (or reporting of trust, as mentioned in Section 4.3.2) over time. ε_{ipt} is the error term. We subject standard errors to clustering at the city level.

4.4.2 Identifying Assumption

Equation 4.6 makes full use of our three-period panel data to trace how an individual's political trust varies with the anti-corruption campaign. β_1 and β_2 are the parameters of interest, capturing how an increase in cumulative investigations is associated with an average individual's political trust in 2014 and 2016, respectively.

The difference-in-differences design compares the trends of political trust between individuals in high- and low-investigation cities. To attribute the trend differences to the gap in corruption investigations, i.e., to causally interpret estimated β_1 and β_2, the common trends assumption needs to be met — if corruption investigations were at the same level, individuals would share similar trends of political trust between cities, conditional on the controls.

The major concern is that if political trust were already on a different trend in high-investigation cities than in low-investigation cities, our estimates would be biased. However, the bias may be limited since the trends may not be very distinct, depending on corruption investigations. On one hand, more corruption investigations can be associated with *declining* political trust, as they may reflect the severity of preexisting corruption. On the other hand, to the extent that corruption is a byproduct of economic growth, more investigations may be associated with *rising* political trust, as people give credit to the developmental government despite the revealed corruption (Ang, 2020).[16] Taken together, these two competing narratives could counteract each other, mitigating differential trends across cities.

We conduct a battery of checks to ensure that the common trends assumption is plausible. First, in Table 4.10, we show that once conditioned on province fixed effects (embedded in individual fixed effects), a city's cumulative investigations (D_p^{14} and D_p^{16}) are not correlated with its pre-campaign (2012) political trust level or with other factors that could affect the evolution of political trust, including the predetermined *levels* and *growth rates* of public sector size, private sector size, GDP per capita, tax revenue per capita, and wage rate. This indicates that the cumulative investigations may be conditionally idiosyncratic, favoring the common trends assumption.[17]

Second, in Equation 4.6, we flexibly control for possible differential trends by including the interactions of individual covariates and year dummies. As robustness checks, we also include past

[16]Note that this would attenuate negative estimates (see Section 4.5.1), suggesting that our results are still informative in the sense of providing a lower bound.

[17]This might not be surprising, as most variations in corruption come from time-invariant factors such as resource endowments, culture, social networks, and so on, which are absorbed by province fixed effects. The remaining variations are due to the anti-corruption campaign's idiosyncratic enforcement.

corruption levels, interacted with year dummies and province-by-year fixed effects, and the results persist (see Section 4.5.2).

Third, although we are unable to test for pretrends using the usual event-study exercise because we have only one pre-campaign period (2012) in the CFPS sample, we show evidence in the same vein using another dataset from the China Social General Social Survey (CGSS). CGSS elicited political trust (in local government) between 2010 and 2012. We correlate political trust in these years with upcoming corruption investigations in 2014 and 2016.[18] Figure 4.13 shows a lack of differential trends in political trust prior to the campaign, also lending support to the common trends assumption.

4.5 Effects of Corruption Investigations on Political Trust

4.5.1 Main Results

Average Effect. Based on Equation 4.6, Table 4.3 presents the impact of corruption investigations on political trust. We scale the estimates to reflect how political trust is associated with a one standard deviation change in the number of cumulative investigations that people have been exposed to (SD = 38). The dependent variable is the political trust dummy (= 1 if the score ≥ 5). As mentioned earlier, this formulation defines distrust conservatively (score ≤ 4) and so enhances the power of the test for whether the anti-corruption campaign *reduces* political trust. In Table 4.3, all the estimates imply that on average, the corruption investigations brought by the anti-corruption campaign have lowered political trust. This decline is pronounced in 2014 but not in 2016.[19]

Specifically, Column (1) displays results from a minimum specification, where only individual and year fixed effects are controlled. In the rest of the columns, we stepwise add covariates. We include birth cohort-by-year fixed effects in Column (2), individual characteristics (gender, edu-

[18] Here we measure corruption investigations at the province level because CGSS only provides a province identifier.

[19] Implications are similar if we use the political trust scale as the dependent variable (see Table 4.13).

Table 4.3: Effect of Anti-Corruption on Political Trust

	(1)	(2)	(3)	(4)
$D^{14} \times T^{14}$	-0.025***	-0.025***	-0.022***	-0.022***
	(0.008)	(0.008)	(0.008)	(0.008)
$D^{16} \times T^{16}$	-0.006	-0.006	-0.009	-0.009
	(0.009)	(0.009)	(0.009)	(0.009)
D.V. mean, pre-campaign	0.633	0.633	0.633	0.633
Individual FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Cohort × year FE		✓	✓	✓
Indiv. char. × year FE			✓	✓
Fam. bkgd. × year FE				✓
Obs.	35610	35610	35610	35610
R^2	0.505	0.507	0.508	0.509

Note: The dependent variable is the political trust dummy. The number of investigations (D) is standardized. Individual characteristics include gender, indicators of educational attainment, *hukou* status, Han ethnicity, Communist Party membership, and state sector employment. Family background includes parents' educational attainment and their Communist Party membership. Robust standard errors, clustered at the city level, are reported in parentheses.
* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

cation, ethnicity, party membership, and state sector employment) interacted with year dummies in Column (3), and family background (parental education and party membership) interacted with year dummies in Column (4). The estimates are remarkably stable with the inclusion of these controls, which indicates that corruption investigations are orthogonal to local conditions. This result lends support to the identifying assumption that requires the paths of political trust to be similar between cities if corruption investigations were at the same level.

In our preferred specification (Column (4)), the estimate shows that in 2014, a one SD increase in corruption investigations on average reduces the likelihood of trusting the government by 2.2 percentage points.[20] Such an effect is sizeable. For example, in the baseline year of 2012, this magnitude is more than two-thirds of the gap in political trust between those who paid attention to corruption news and those who did not (Column (2) of Appendix Table 4.9). It is also one-fifth of

[20] One SD increase is also very similar to the difference between the third quartile and the first quartile of corruption investigations in both 2014 and 2016.

the gap in political trust between Communist Party members and other individuals.

Although the anti-corruption campaign had a negative effect on political trust in 2014, our estimates suggest that by 2016, the effect had become insignificant. However, the average effects mask heterogeneous effects across different groups. Notably, the anti-corruption campaign had varying effects on individuals with different education levels. As we will demonstrate in greater detail later, the campaign decreased political trust among individuals with lower levels of education, while individuals with a college degree experienced an increase in political trust (see Table 4.6 and Figure 4.8). The results show interesting distributional effects of the campaign on political trust, and point to potential polarization. The heterogeneity in the treatment effects also provides insight into the mechanisms behind the change of individuals' attitudes. Later in the paper, we examine the interplay of education and anti-corruption more closely (see Section 4.6.2).

Effects at Different Margins. Apart from identifying the average effect, we also zoom into corruption investigations' effects at different margins of political trust. We create dummy variables for each level of political trust scores, ranging from 0 to 10, and use them as the dependent variable of Equation 4.6.

Figure 4.4a presents the changes in 2014. We see that corruption investigations have reduced the proportion of people who hold a moderate level of political trust (score = 4–6). Such an effect is notable — these people may have been ambivalent about judging the government; however, the anti-corruption campaign appears to have pushed them to make a more definite judgment. The changes in judgment are slightly *polarized*: most people adjust their political trust to low levels (1–3), but there is also a significant increase in the size of people who are highly trustful of the government (score = 9), which renders the negative average effect observed in Table 4.3. Figure 4.4b shows the effects of corruption investigations on political trust by margin in 2016, where the patterns are similar to those in 2014 but are far less pronounced.

Figure 4.4: Effects at Different Margins of Political Trust

(a) Effects in 2014 (b) Effects in 2016

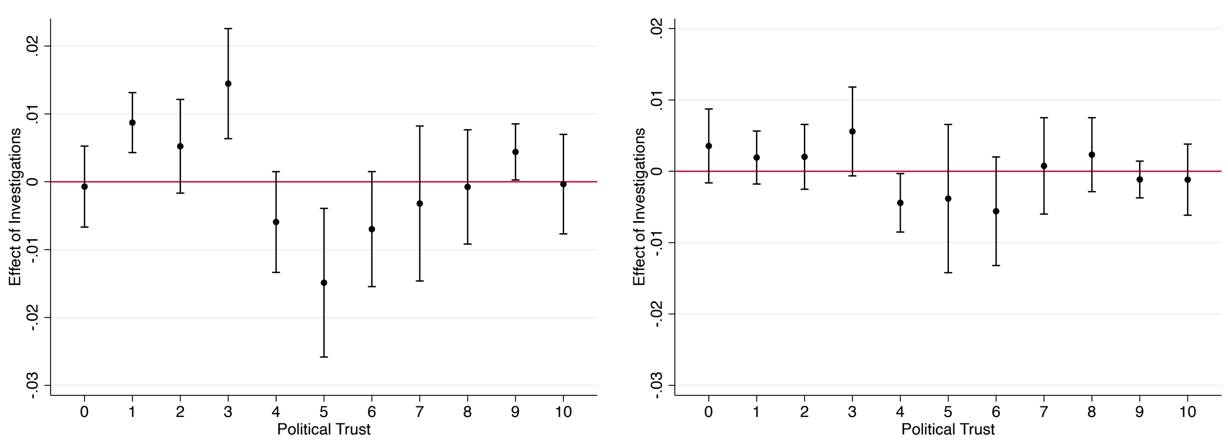

Note: We estimate Equation 4.6 using the dummy variable for each level of political trust as the independent variable. The estimated coefficients on the 2014 investigations are plotted in (a), and the estimated coefficients on the 2016 investigations are plotted in (b). The solid points are point estimates and the caps are 90 percent confidence intervals.

Within-individual Change in Political Trust. Our fixed-effects models exploit variation in political trust of individuals that changed their attitudes during the sample period. As most people's attitudes tend to be stable, our identification of the causal effects of the anti-corruption campaign is based on a small fraction of individuals whose political trust changed over time. Investigating this subset of individuals provides insight into our empirical strategy and can potentially shed light on the underlying mechanisms. Figure 4.5 displays the distribution of within-individual changes in political trust over time, showing that a majority of individuals did not change their political trust. Relative to 2012, more than 60% of individuals did not change their political trust in 2014 or 2016.[21] We also examine the distribution of the change in the original reported political trust score in Appendix Figure 4.12, showing similar patterns. Notably, we do not observe more people reducing political trust than those increasing political trust, which could be attributed to other events that occurred during the same period, such as economic growth and more intensive propaganda.

[21] Actually, the political trust dummy stayed unchanged for 49% individuals from 2012 through 2016.

Therefore, the DID strategy (with time fixed effects) is necessary for identifying the causal effects of the anti-corruption campaign on political trust. Moreover, as we will discuss later, the campaign had heterogeneous effects, including increasing the political trust of highly educated individuals.

Are there unique features that distinguish individuals who changed their political trust from those who did not? In Appendix Table 4.11 and Appendix Table 4.12, we check the balance of covariates between different types of individuals defined based on their change in political trust over time. In most cases, the difference between the groups appear statistically or economically insignificant. However, there are some interesting exceptions. First, the baseline level of political trust in 2012 is much lower for individuals whose political trust increased subsequent years, likely due to mechanical reversion to the mean. Second, individuals whose political trust eventually increased in 2016 are more likely to come from regions with a stronger Confucian ideology, as measured by the number of Confucian temples. As we will discuss later, Confucian ideology likely played a key role in driving the positive response of highly educated individuals to the campaign. Still, we need to interpret these patterns with caution as they are based only on inter-temporal variation before and after the onset of the campaign.

Figure 4.5: Distribution of Change in Political Trust Dummy Within Individuals

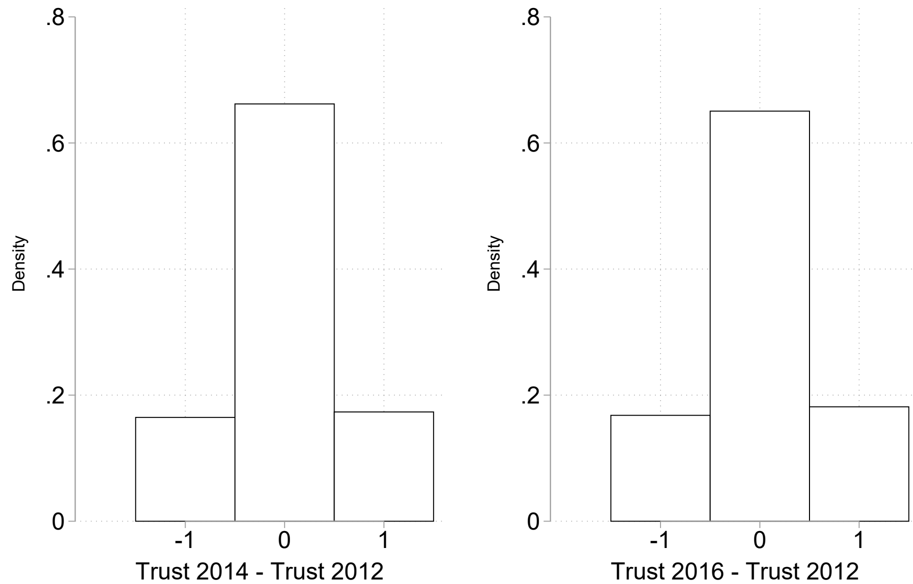

Note: This figure reports the distribution of within-individual change in the political trust dummy relative to 2012 in the balanced panel sample.

Remarks. Taken together, our results show that corruption investigations brought by China's anti-corruption campaign lower the average individual's political trust. But this masks important heterogeneity: there is a small group of people who actually enhance their trust, despite the majority reacting negatively upon learning of a great deal of corruption in government. In light of our theoretical prediction, Proposition 1, our results suggest that the average individual interprets corruption revelation from the campaign negatively, but there exists a significant cleavage in beliefs about whether the government is honest or corrupt, making the same information about corruption interpreted differently and ultimately bifurcating political trust. We supplement this view with more heterogeneity analyses and discussions of alternative explanations in Section 4.6. Meanwhile, in the rest of this section, we provide several robustness checks for our results.

4.5.2 Robustness Checks

Further Controls for Differential Trends. The key assumption to be met for causal interpretations of our estimates is the common trends assumption: were corruption investigations at the same level, the trends of political trust would be similar between cities (see Section 4.4.2). This assumption is plausible as corruption investigations are conditionally idiosyncratic, i.e., they are conditionally orthogonal to a variety of factors in political trust's evolution (see Table 4.10) and are not associated with pretrends in political trust (see Figure 4.13).

Even so, to shed more light on the common trends assumption, we explicitly add differential trends of political trust in two ways. First, corruption investigations may relate to past levels of corruption, which shape the long-term trend of political trust, so we include interactions of past corruption levels and year dummies in Equation 4.6 to further purge any differential trends. Past corruption levels are measured using firms' entertainment and travel costs (ETC), which are often spent on bribing government officials (Cai et al., 2011). Cai et al. (2011)'s ETC indices from 2002–2004 cover half of the cities in our sample (60). We take the three-year average. Columns (1) and (2) of Table 4.4 display the results of controlling for past corruption-related paths. Since ETC data only cover a subset of the cities in our sample, to aid in comparison we re-estimate Equation 4.6 in Column (1), using the subsample where ETC is available. Reassuringly, corruption investigations still reduce political trust. Column (2) shows that the inclusion of average ETC interacted with year dummies does not materially change the estimates. Second, we examine the robustness of our results by including province-by-year fixed effects, as province-invariant factors may result in differential trends of political trust. Columns (3) and (4) of Table 4.4 show that the estimates with and without including province-by-year fixed effects deliver the same implications. Column (3) replicates Column (4) of Table 4.3. After further controlling for province-by-year fixed effects in Column (4), we find that the estimated effect is very similar to our baseline results for the year of 2014. The estimated effect on political trust in 2016 is also negative and statistically significant with the inclusion of province-by-year fixed effects. Taken together, our results were most likely not confounded by differential trends of political trust.

Table 4.4: Robustness Checks: Further Controls for Differential Trends

	+ ETC × Year FE		+ Province × Year FE	
	(1)	(2)	(3)	(4)
$D^{14} \times T^{14}$	-0.027***	-0.026***	-0.022***	-0.026**
	(0.008)	(0.008)	(0.008)	(0.012)
$D^{16} \times T^{16}$	-0.005	-0.004	-0.009	-0.019**
	(0.011)	(0.010)	(0.009)	(0.008)
D.V. mean, pre-campaign	0.624	0.624	0.633	0.633
Individual FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Covariates	✓	✓	✓	✓
ETC × year FE		✓		
Province × year FE				✓
Obs.	18810	18810	35610	35610
R^2	0.523	0.524	0.509	0.511

Note: The dependent variable is the political trust dummy. The number of investigations (D) is standardized. Covariates include gender, educational attainment, *hukou* status, Han ethnicity, Communist Party membership, state sector employment, parental educational attainment, and parental Communist Party membership, all of which are interacted with year dummies. The ETC index is from Cai et al. (2011). Column (1) replicates Column (4) of Table 4.3 using the subsample for which the ETC information is available. Column (3) is the same as Column (4) of Table 4.3. Robust standard errors, clustered at the city level, are reported in parentheses.

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

Permutation Test. We also conduct a permutation test to ensure that the results are not simply due to chance. We permute corruption investigations across cities and re-estimate Equation 4.6 to derive a counterfactual effect of corruption investigations on political trust. Figure 4.14 displays the distribution of counterfactual effects in 2014, derived from 1,000 permutations. The vertical line is the true effect estimated using the actual sample. As shown, the counterfactual effects are centered around zero, while the true effect is at the distribution's tail and is statistically significant (p-value = 0.066), indicating that the true effect is not coincidental.[22]

Aggregation of Heterogeneous Effects. Recent econometric literature on difference-in-differences designs with continuous treatments (Callaway and Sant'Anna, 2021; de Chaisemartin et al., 2022b) suggests that the two-way fixed effects (TWFE) estimator may place weights that are not sensible when aggregating treatment effects, even if there is no variation in treatment timing.[23] What is concerning is that if the treatment effects are highly heterogeneous and there are negative weights, then our estimate probably recovers a causal parameter with a sign opposite to conventional causal parameters of interest (e.g., average treatment effect), leading to interpretation challenges. To alleviate this concern, we implement de Chaisemartin et al. (2022b)'s heterogeneity-robust estimator, which confirms the main findings that corruption investigations reduced political trust in 2014 (see Table 4.14).[24] That is to say, our results are not due to incorrect aggregation of treatment effects.

4.6 Mechanisms and Discussions

Thus far, our results have provided robust evidence that on average, corruption investigations brought by the anti-corruption campaign lowered political trust. When we zoom into the effects

[22] Put another way, this test rejects the sharp null that the campaign had no effect on political trust in any city at a significance level of 0.066.

[23] In a difference-in-difference design with binary treatment, the aggregation problem occurs when there are variations in treatment timing (Goodman-Bacon, 2021).

[24] To implement de Chaisemartin et al. (2022b), a group consisting of individuals facing low corruption investigations needs to be defined for comparison. The low level was chosen by the researchers. We attempt both 25th and 50th percentiles as cutoffs, and the results are qualitatively similar.

at each margin of political trust, we see slight polarization. What has driven these changes? The patterns emerging in Section 4.5 seem to echo predictions in Proposition 1: the anti-corruption campaign has heterogeneous impacts on political trust, since people interpret new information about corruption based upon their distinct priors.

To shed more light on this view, we corroborate Proposition 1's two core mechanisms. First, *informativeness* — the anti-corruption campaign ought to provide information about corruption in government, and then people use that information in the evaluative process. Second, *interpretation* — conditional on informativeness, how to infer government honesty from the same piece of information depends on an individual's pro- or anti-government slant. We examine these two mechanisms in Section 4.6.1 and Section 4.6.2, respectively. Additionally, we discuss other alternative explanations to further ascertain Proposition 1's implications.

4.6.1 Informativeness

If informativeness is at work, then a testable implication is that the anti-corruption campaign should have a more pronounced impact on political trust among the group that had less information about corruption before (low $\bar{x}_i$), *ceteris paribus*. In light of Equation 4.5, the change in political trust, $\Delta s_i = (2p_i - 1)(x - \bar{x}_i)$, can be amplified by informativeness, $x - \bar{x}_i$.

Our first test is to investigate how the campaign's effects vary with previous exposure to corruption news. For those who had been unmindful of corruption news, the anti-corruption campaign may be relatively more informative, i.e., $x - \bar{x}_i$ is larger. Consequently, we expect the campaign to have a more discernible effect among this group. The CFPS 2010 survey asked respondents whether they had paid attention to corruption news. Thus, we estimate Equation 4.6 separately for individuals with and without attention to corruption news. Figure 4.6 compares the estimates emerging from this subsample analysis. Tellingly, the drop in political trust appears to be driven by the campaign's influences on those who had paid little attention to news about corruption, and a test strongly rejects that the effects are equal between the two subsamples (p-value for the 2014 difference = 0.049, p-value for the 2016 difference < 0.001), which is consistent with our hypothesis.

The first test exploits variations in people's knowledge about corruption from a *demand* perspective. By contrast, our second test leverages variations driven by *supply*-side factors. For many people the Internet, especially social media, is a prominent source (if not the only source) to learn about corruption in government, and not only in China (Guriev et al., 2021; Qin et al., 2017). However, internet censorship could suppress the supply of information about corruption. As a result, in highly censored regions, the anti-corruption campaign would ironically be more informative as it reveals corruption that tended to be covered up before, enlarging its impact on political trust. To test this hypothesis, we use the share of posts deleted on Weibo ("Chinese Twitter") provided by Qin et al. (2017) to measure the censorship intensity in each province. Then, we divide our sample by quintiles of censorship intensity and estimate Equation 4.6 separately using each subsample. Figure 4.7a shows patterns in line with our hypothesis — the campaign induced a distinctive drop in political trust in 2014 in the most censored provinces, which is statistically distinguishable from the impacts in other less censored provinces.[25]

In addition, the Internet's ability to provide information about corruption can be rather restricted due to propaganda. Government users on Weibo can disseminate "neutral or positive" messages to distract the public from scandals (King et al., 2017; Qin et al., 2017), making the anti-corruption campaign more informative and its impact on political trust more pronounced in regions subject to intense propaganda. To examine this hypothesis, we measure propaganda intensity using the share of government users among a province's Weibo users, again provided by (Qin et al., 2017), and perform the same heterogeneity exercise as before. Reassuringly, Figure 4.7b shows that the greatest drop in political trust occurs in the top quintile of propaganda intensity.[26]

[25]Testing the respective Q1-Q5, Q2-Q5, Q3-Q5, and Q4-Q5 differences in 2014 estimates yields p-values of 0.039, 0.003, 0.013, and 0.022 (Q means quintile).

[26]Testing the respective Q1-Q5, Q2-Q5, Q3-Q5, and Q4-Q5 differences in 2014 estimates yields p-values of 0.050,

Overall, exploiting variations in people's knowledge about corruption from both the demand and supply sides, our results support the informative nature of the anti-corruption campaign.[27] However, they do not yet clarify how different people process such information. In the following subsection, we probe this issue by investigating the role of priors.

Figure 4.6: Effects by Attention to Corruption News

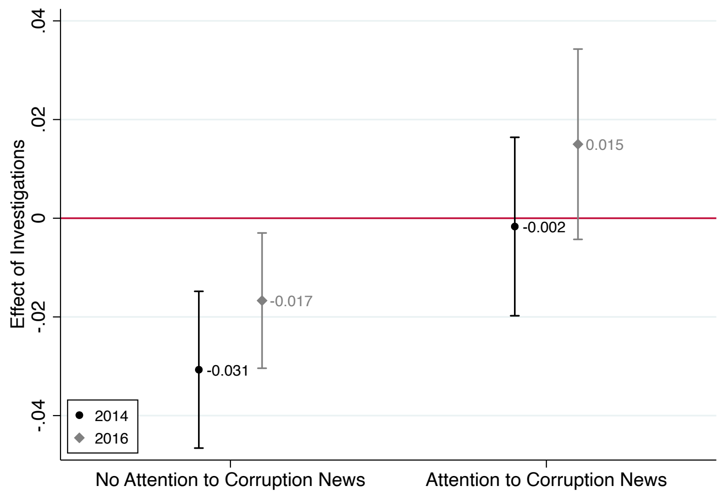

Note: The sample is divided into subsamples by individual attention to corruption news (elicited in the CFPS 2012 survey), and Equation 4.6 is estimated separately in each subsample. The solid points are point estimates and the caps are 90 percent confidence intervals.

0.034, 0.033, and 0.038. Though the 2016 estimate in Q5 has a large magnitude, it is not statistically distinguishable from its counterparts in other quintiles due to the wide confidence interval.

[27] Please note this caveat of doing *ceteris paribus* heterogeneity exercises: apart from capturing the campaign's high informativeness (high $x - \bar{x}_i$ due to low $\bar{x}_i$), the measurements we use may relate to relevant groups' negative interpretations (negative $2p_i - 1$ due to low p_i). For instance, one may be interested in corruption news since she is suspicious about government honesty, and a region often has more intense censorship and propaganda due to concerns of political instability, as people there tend to be more rebellious. However, this would not reject the conclusion that the anti-corruption campaign is informative. As the interpretation channel only operates when there is some information provided by the campaign ($x - \bar{x}_i > 0$), the results of the above heterogeneity exercises indicate existence of the campaign's informativeness, though we may not be able to disentangle how much heterogeneous effect is purely due to informativeness and how much is due to interpretations.

Figure 4.7: Effects by Censorship and Propaganda

(a) Share of Deleted Weibo Posts (b) Share of Government Accounts on Weibo

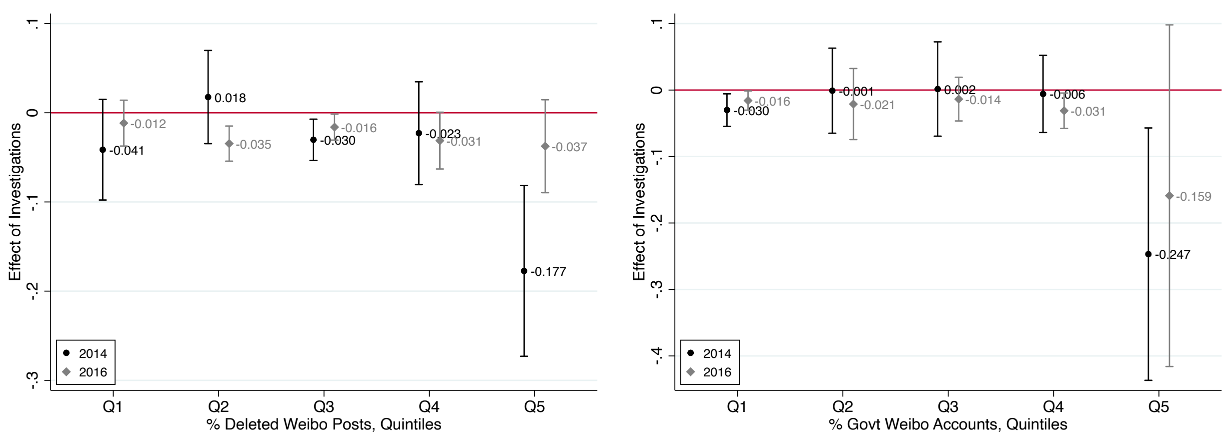

Note: In (a) and (b), the sample is divided into five subsamples by quintiles of deleted Weibo posts or government Weibo users, and Equation 4.6 is estimated separately in each subsample. The solid points are point estimates and the caps are 90 percent confidence intervals.

4.6.2 Interpretation

As we highlight in Proposition 1, the same piece of information is compatible with different interpretations, and an individual would overweight the interpretation more aligned with her pro- or anti-government priors, resulting in different impacts of the anti-corruption campaign on political trust. To shed light on the role of prior-driven interpretations, we first concentrate on factors that shape priors directly: unpleasant experiences with the government. Then, we examine a more indirect determinant: education. This choice is motivated by three factors. First, education has important influences on political behavior and political attitudes, as is well-documented by a large body of literature (Almond and Verba, 1963; Putnam et al., 1993; Dee, 2004; Sondheimer and Green, 2010; Campante and Chor, 2012; Croke et al., 2016). Second, education has been an important component in state building — a regime devises their education policy to cultivate citizens that are supportive of their very regime (Weber, 1976; Ramirez and Boli, 1987; Lott, 1999; Aghion et al., 2019; Alesina et al., 2021; Bandiera et al., 2019; Cantoni et al., 2017). Third, there has been some evidence and observations showing that educated Chinese people exhibit stronger pro-government or nationalistic sentiments (Cantoni et al., 2017; Qi et al., 2022). For instance, Cantoni et al. (2017) show that China's textbook reform enhances elite students' support for Chinese institutions. Qi et al. (2022) find that education is positively associated with support for the armed unification of Taiwan.

Role of Experiences with the Government

People may have formed their priors about the government in the course of interactions with government officials. Negative experiences can damage the government's image, making people interpret corruption investigations in a way that discredits the government. Based on the CFPS 2012 survey, we examine four types of self-reported negative experiences: (i) being unfairly treated by officials, (ii) having conflicts with officials, (iii) encountering lazy officials, and (iv) being asked for bribes. Table 4.5 compares the anti-corruption campaign's effects on political trust between people with and without these experiences. We see that people who have had unpleasant experiences incur a larger decline in trust, suggesting that the campaign may have provoked more negative sentiments among these groups due to their negative priors.

Table 4.5: Experiences with the Government and Political Trust

	Unfarily Treated		Having Conflicts		Lazy Cadres		Asked for Bribes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	No	Yes	No	Yes	No	Yes	No	Yes
$D^{14} \times T^{14}$	-0.015*	-0.073***	-0.021**	-0.074*	-0.015*	-0.060**	-0.017**	-0.113***
	(0.009)	(0.020)	(0.008)	(0.043)	(0.009)	(0.023)	(0.008)	(0.026)
$D^{16} \times T^{16}$	-0.007	-0.016	-0.008	-0.042	-0.007	-0.013	-0.006	-0.041*
	(0.009)	(0.015)	(0.009)	(0.028)	(0.010)	(0.013)	(0.009)	(0.021)
D.V. mean, pre-campaign	0.657	0.465	0.647	0.456	0.663	0.485	0.653	0.470
p-value, 2014 diff.		0.012		0.182		0.069		0.000
p-value, 2016 diff.		0.547		0.170		0.705		0.057
Individual FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Covariates	✓	✓	✓	✓	✓	✓	✓	✓
Obs.	32307	3159	34269	1251	31017	4446	33129	2349
R^2	0.497	0.580	0.506	0.577	0.497	0.554	0.501	0.580

Note: The dependent variable is the political trust dummy. The number of investigations (D) is standardized. Covariates include gender, educational attainment, *hukou* status, Han ethnicity, Communist Party membership, state sector employment, parental educational attainment, and parental Communist Party membership, all of which are interacted with year dummies. Robust standard errors, clustered at the city level, are reported in parentheses.

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

Role of Education

Table 4.6 takes a first look at how the anti-corruption campaign influences political trust differently by education levels. As Column (1) shows, the decline in political trust is mitigated by educational attainment. To be more specific, if one has completed high school, i.e., 12 years of schooling, the campaign has a virtually null effect on political trust (e.g., $-0.062+12\times 0.005=0.002$ in 2014); and if one has had some college education, then the campaign turns out to enhance political trust. Figure 4.8 uses a non-parametric approach to examine the heterogeneity by education, confirming that education mitigates the drop in political trust and identifying college education as a turning point.

Table 4.6: Education, Confucianism, and Political Trust

	(1)	(2)
$D^{14} \times T^{14}$	-0.062***	-0.066***
	(0.015)	(0.015)
$D^{16} \times T^{16}$	-0.041***	-0.039***
	(0.008)	(0.009)
$D^{14} \times T^{14} \times$ Schooling	0.005***	-0.001
	(0.001)	(0.004)
$D^{16} \times T^{16} \times$ Schooling	0.004***	-0.000
	(0.001)	(0.002)
$D^{14} \times T^{14} \times$ Schooling $\times$ ln(Conf. temples)		0.001*
		(0.001)
$D^{16} \times T^{16} \times$ Schooling $\times$ ln(Conf. temples)		0.001**
		(0.000)
Schooling mean	6.976	6.822
Confucianism mean	4.584	4.584
Individual FE	✓	✓
Year FE	✓	✓
Covariates	✓	✓
Obs.	35610	29910
R^2	0.509	0.509

Note: The dependent variable is the political trust dummy. The number of investigations (D) is standardized. Covariates include gender, educational attainment, *hukou* status, Han ethnicity, Communist Party membership, state sector employment, parental educational attainment, and parental Communist Party membership, all of which are interacted with year dummies. Robust standard errors, clustered at the city level, are reported in parentheses.

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

Figure 4.8: Effects by Educational Attainment

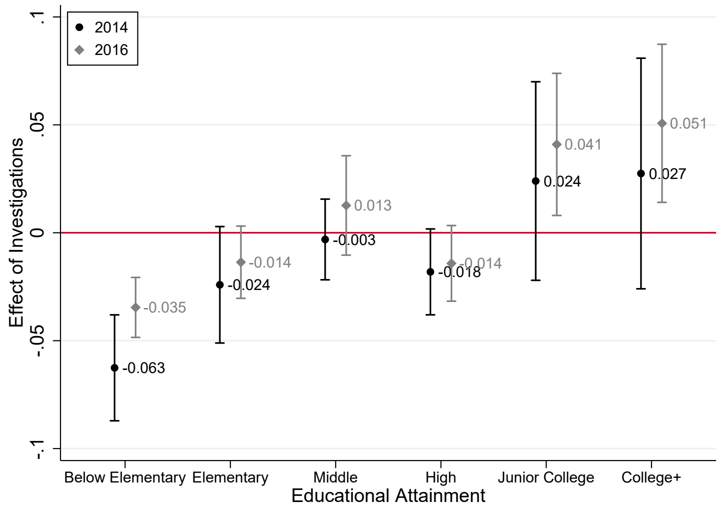

Note: The sample is divided into subsamples by individual educational attainment, and Equation 4.6 is estimated separately in each subsample. The solid points are point estimates and the caps are 90 percent confidence intervals.

These patterns are intriguing. They suggest that education may have shaped pro-government priors, making the educated interpret the anti-corruption campaign in a positive way. It is worth noting that educated people also tend to be better informed about corruption in government, making the anti-corruption campaign less informative to them. The increase in political trust among college-educated respondents underscores the distinctive interpretation associated with education: if the college-educated have the same interpretation as the average person, then they would have a weaker decrease or even no change in political trust, rather than an increase.

We further confirm that our results reflect a unique, education-induced, pro-government prior. In Table 4.15, we horse race education and individual attributes that are associated with education and may also foster pro-government, such as urban *hukou*, Communist Party membership, and state sector employment. The results show that education's impacts are remarkably stable.

So what enables more education to shape stronger pro-government attitudes? We suggest China's Confucianism as a driving force, for two reasons. First of all, as China's traditional political philosophy, Confucianism heavily influences China's political traditions (Pan and Xu, 2018; Bell, 2010; Perry, 2008; Jiang, 2016; Economist, 2021; Page, 2015).[28] Importantly, it features a "benevolent dictator model" that encourages people to be faithful to the ruling body and discourages critiques (e.g., Acemoglu and Robinson, 2020). As Confucius himself put it, *"commoners do not debate matters of government"* (Confucius, 2003). In this vein, some scholars even invoke Confucianism as an explanation of the long and consistent autocratic history of China (Huntington, 1991; Acemoglu and Robinson, 2020, 2021a). Second, Confucian doctrines are well integrated into China's education system (Jiang, 2016). For instance, they are taught in Chinese literature and history classes and even tested in the college entrance exam in some regions. Accordingly, it could be easier for people to accept the pro-government doctrines embedded in education in highly Confucian provinces, where the local norms have been historically more pro-government.

Therefore, we expect an interplay between education and Confucianism in fostering a pro-government prior and thus enhancing political trust. To test this hypothesis, we measure Confucianism using the number of Confucian temples (Chen et al., 2020), which are historical sites for Confucian teachings and so capture the spread of Confucianism.[29] Relying on this measurement, a further heterogeneity exercise in Column (2) of Table 4.6 favors our hypothesis: education's heterogeneous effect is completely driven by Confucianism.

To take a granular look at the interplay between education and Confucianism, we partition our sample into four groups by college completion: (i) below college and below-median Confucianism, (ii) below college and above-median Confucianism, (iii) above college and below-median Confucianism, and (iv) above college and above-median Confucianism. Then, we conduct subsample

[28] For the first three decades of its reign in China, the Communist Party tried to extirpate Confucianism that it saw as feudal and backward. However, many of its practices were nonetheless implicitly shaped by Confucian values (Bell, 2010; Perry, 2008). Moreover, in the post-Mao era, the Party rehabilitated and promoted Confucianism, labeling it China's homegrown political philosophy (Jiang, 2016; Economist, 2021).

[29] Using the same measure, previous literature has provided evidence that Confucianism may lead to conformity to the government. For instance, Kung and Ma (2014) find that peasant rebellions were less likely to occur in areas where Confucian culture was stronger. Alm et al. (2022) find that people were more likely to conform to housing market regulations in cities with a stronger Confucian culture.

analysis based on Equation 4.6. Figure 4.9 presents the *gap* in the anti-corruption campaign's effects on political trust between high- and low-Confucianism provinces, separately by college completion. It delivers two messages. First, all the gaps are positive, implying that regardless of college attainment, the drop in political trust induced by the campaign is attenuated or even reverted to an increase in more Confucian provinces. This suggests that Confucianism carries, as it advocates, pro-government attitudes. Second, the positive gaps are much larger for the college-educated group, suggesting that Confucianism facilitates pro-government indoctrination inherent in China's education system. In Appendix Figure 4.15, we provide a more granular analysis by dividing the full sample into subgroups based on education and Confucianism. We then estimate the effects of the campaign on political trust separately for each subgroup. Our findings confirm that the observed heterogeneous effects by education are driven by regions with a stronger Confucianism ideology.

Taking these results together, we find that the campaign's effect on political trust varies dramatically with educational attainment, and that the interplay between education and Confucianism drives it. These findings shed light on the interpretation mechanism through education: education shapes a pro-government prior, which leads to different interpretations of corruption information and political trusting behavior.

Figure 4.9: Effects of Education Compared: High Confucianism versus Low Confucianism

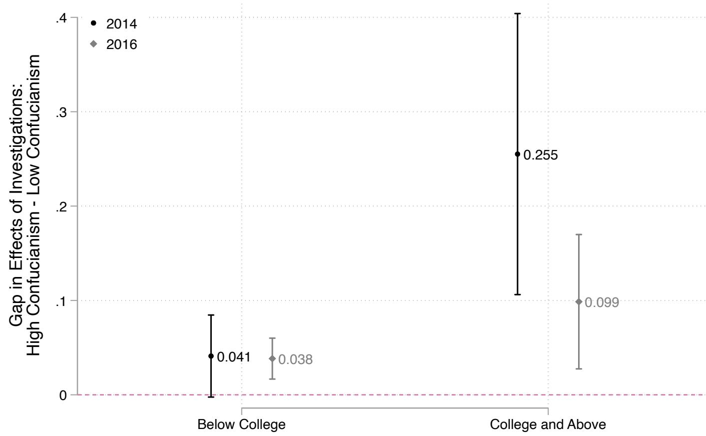

Note: The sample is divided into two groups: the below-college-educated and the above-college-educated. Within each group, we further divide individuals by residence in high and low Confucianism cities (Confucian temples above or below the median). Then, we estimate Equation 4.6 in each subsample. The solid points are the differences in effects between high and low Confucianism cities, and the caps are 90 percent confidence intervals.

4.6.3 Discussions: Alternative Explanations

We read our results as the anti-corruption campaign affecting political trust by intervening in people's perceptions of underlying government honesty. However, as we have noted in Section 4.2.2, changes in political trust may be rendered through other channels related to the anti-corruption campaign. To pin down the interpretations, we consider two possibilities below. We do not find strong evidence that they significantly threaten our interpretations.

Government Performance. Besides offering information about corruption for people to infer underlying government honesty, the anti-corruption campaign may affect political trust by influencing government performance. We do not think this would threaten our results much, since if anything, stricter monitoring is expected to deter misconduct and improve government performance, making the findings of a negative effect on political trust even more surprising. However, the campaign could result in a backlash. The performance deteriorates due to a chilling effect: officials shirk their duties to avoid unconsciously making mistakes that would be targeted by the campaign (Wang, 2022).

To address this concern, we investigate the role of government performance in Columns (1)–(3) of Table 4.7. Column (1) shows that corruption investigations do not significantly affect people's assessment of government performance (measured on a scale from 1–5, the higher the better). This may be reasonable, since it is not practical for government performance to improve much in the short run. As expected, Columns (2) and (3) show that controlling for government performance does not markedly change the effects of corruption investigations on political trust. Therefore, our findings cannot be explained by the campaign's impacts on government performance.

General Trust. Chinese society has witnessed many changes under President Xi's administration, and the anti-corruption campaign was just the tip of the iceberg. One may be concerned that instead of speaking to people's updated views on the government, the negative relationship between political trust and corruption investigations we uncover merely reflects changes in general trust (in any entity, not just in government) due to the anti-corruption campaign or other contemporaneous shocks that correlated with it.

To examine if our findings are just a manifestation of changes in general trust, we conduct a couple of placebo tests that investigate whether trust in other groups is affected by the anti-corruption campaign (or shocks correlated with it). We expect to see null effects in these tests if changes in political trust are not driven by changes in general trust. As expected, Columns (4)–(6) of Table 4.7 show that corruption investigations have no impact on respondents' trust in parents, strangers, and Americans. In other words, our results are specific to changes in people's perceptions about the government rather than about changes in overall willingness to trust.

Signaling. Another alternative interpretation is that by voluntarily disclosing corruption, the government may have sent out a (credible) signal to people through the anti-corruption campaign: it is now legitimate to criticize the government and lower political trust, which used to be politically

Table 4.7: Alternative Interpretations: Government Performance and General Trust

	Govt. Performance			Trust in Other Groups		
	(1) Performance	(2) Political Trust	(3) Political Trust	(4) Parents	(5) Strangers	(6) Americans
$D^{14} \times T^{14}$	0.035	-0.018**	-0.020***	-0.003	0.009	-0.000
	(0.028)	(0.008)	(0.008)	(0.002)	(0.010)	(0.016)
$D^{16} \times T^{16}$	0.019	-0.008	-0.009	-0.000	0.003	0.006
	(0.020)	(0.010)	(0.009)	(0.001)	(0.005)	(0.006)
Performance			0.062***			
			(0.004)			
D.V. mean, pre-campaign	3.412	0.632	0.632	0.981	0.183	0.260
Individual FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Covariates	✓	✓	✓	✓	✓	✓
Obs.	33078	33078	33078	33017	33020	32524
R^2	0.511	0.510	0.516	0.400	0.469	0.488

Note: The number of investigations (D) is standardized. Covariates include gender, educational attainment, *hukou* status, Han ethnicity, Communist Party membership, state sector employment, parental educational attainment, and parental Communist Party membership, all of which are interacted with year dummies. Robust standard errors, clustered at the city level, are reported in parentheses.
* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

incorrect or even prohibited, as the government portrayed itself as impeccable.[30]

We cannot completely rule out this possibility, which requires knowledge about each individual's take on the campaign's nature. Nonetheless, we provide some evidence to show that the signaling story is unable to fully drive our findings. First of all, were signaling the only mechanism at play, the anti-corruption campaign should have universally reduced political trust. However, some groups actually increased their political trust following the campaign (see Figure 4.4a and Section 4.6.2).

Second, we show that our findings survive excluding those respondents who had been reticent due to deference or fears, and so would be most prone to express distrust after the campaign's signaling. Table 4.8 presents this exercise. In Column (1), we exclude members of the Communist

[30] Newman et al. (2021) tells a similar story in the US context by investigating the effects of Donald Trump's campaigns on demonstrations of racial prejudice. Prejudiced citizens usually tend to constrain the expression of their prejudice. However, they are emboldened to express and act upon their prejudices if there are political elites (e.g., Trump) doing so.

Party of China (CPC) and those from military families. Party disciplines and indoctrination may cultivate their loyalty to the state, making them see criticizing or reporting distrust in government as taboo. In Column (2), we exclude those from families who were persecuted in the Communist Revolution (1950s) and the Cultural Revolution (1966–76). Persecutions are measured by government-assigned bad class labels (e.g., landlords, rich peasants, and capitalists) and experiences of the Sent-Down Youth Movement and the May Seventh Cadre School, which were elicited in the CFPS 2010 survey. State repression can credibly make people frightened of criticizing the government unless they are allowed to do so (in fact, they can be highly motivated to do so). In the same avenue, Column (3) further excludes those who witnessed the violent Cultural Revolution — they either came of age (reached impressionable years, 18–25) during the CR or were from cities with a large share of the population afflicted (above median) (Alwin and Krosnick, 1991; Walder, 2014). In Column (4), all three groups are excluded.

Patterns in Table 4.8 suggest that the signaling story may have played a minor role. We see that after excluding each respective group, the mean pre-campaign trust is not dramatically lower than the full sample mean (0.633). Thus the excluded individuals, who are supposedly reticent because of deference or fears, in fact do *not* report significantly lower trust. In addition, notwithstanding the exclusion of these groups, Table 4.8 shows that corruption investigations reduce political trust, implying that our results are not entirely driven by the signaling story. Notably, the subsample estimates reported in Columns (3) and (4) are larger than the full sample estimates (cf. Table 4.3), suggesting the role of signaling may not be the most prominent.

4.7 Conclusion

This chapter studies the impacts of China's recent anti-corruption campaign on political trust. Using individual panel data to trace the evolution of political trust, we find that the campaign, on average, has reduced political trust. We provide suggestive evidence for two (interrelated)

Table 4.8: Alternative Interpretations: Signaling

	(1)	(2)	(3)	(4)
$D^{14} \times T^{14}$	-0.019**	-0.018**	-0.054**	-0.045*
	(0.008)	(0.008)	(0.026)	(0.026)
$D^{16} \times T^{16}$	-0.010	-0.008	-0.044***	-0.042***
	(0.009)	(0.009)	(0.011)	(0.010)
Excluded	CPC & Military	Hit by Revolutions	Witnessed Violent CR	All 3 Groups
D.V. mean, pre-campaign	0.628	0.640	0.627	0.630
Individual FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Covariates	✓	✓	✓	✓
Obs.	31242	31098	13122	10551
R^2	0.506	0.504	0.510	0.505

Note: The dependent variable is the political trust dummy. The number of investigations (D) is standardized. Covariates include gender, educational attainment, *hukou* status, Han ethnicity, Communist Party membership, state sector employment, parental educational attainment, and parental Communist Party membership, all of which are interacted with year dummies. Robust standard errors, clustered at the city level, are reported in parentheses.
* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

mechanisms. First, the campaign may have functioned as an information treatment by improving people's limited knowledge about corruption. Indeed, the drop in political trust is more pronounced among less informed groups. Moreover, we uncover strong heterogeneity driven by the pro- and anti-government cleavage. Political trust is reduced to a greater extent for those who have had unpleasant experiences with the government. Education mitigates this reduction and even reverses it to an increase, possibly by forging pro-government sentiments, as exemplified by the strong interplay between education and Confucianism.

Our results suggest that government officials face a dilemma when attempting to advertise a seemingly popular anti-corruption reform. Such a reform inevitably reveals the government's downsides, and its influence, at least in the short run, depends on people's priors about whether the government is good or not, providing space for polarization and backlash should the cleavage in priors be significant. Moreover, it could be hard to alter these effects in the short run. As we have shown, they are rooted in the long-term practices of the government (e.g., censorship, propaganda, governance), as well as in cultural norms (e.g., Confucianism).

We close this chapter by noting one limitation of our results. We, at best, speak to the short-run effects of the anti-corruption campaign. However, people may update their political trust in the long term, changing the implications of the anti-corruption campaign on political trust. In addition, although we do not find evidence that the effect operated through the channel of government performance, we conjecture that in the long run, government performance could matter since the anti-corruption campaign has tremendous economic effects (e.g., Kong et al., 2020; Kong and Qin, 2021; Chen and Kung, 2019b; Xu and Yano, 2017b; Xu et al., 2021). It will be interesting for future research to examine the long-run effect, given that the campaign is still proceeding and the anti-corruption practices in the campaign tend to be a regular part of the institution.

4.8 Appendix

4.8.1 Additional Figures and Tables

Figure 4.10: Correlates of Political Trust

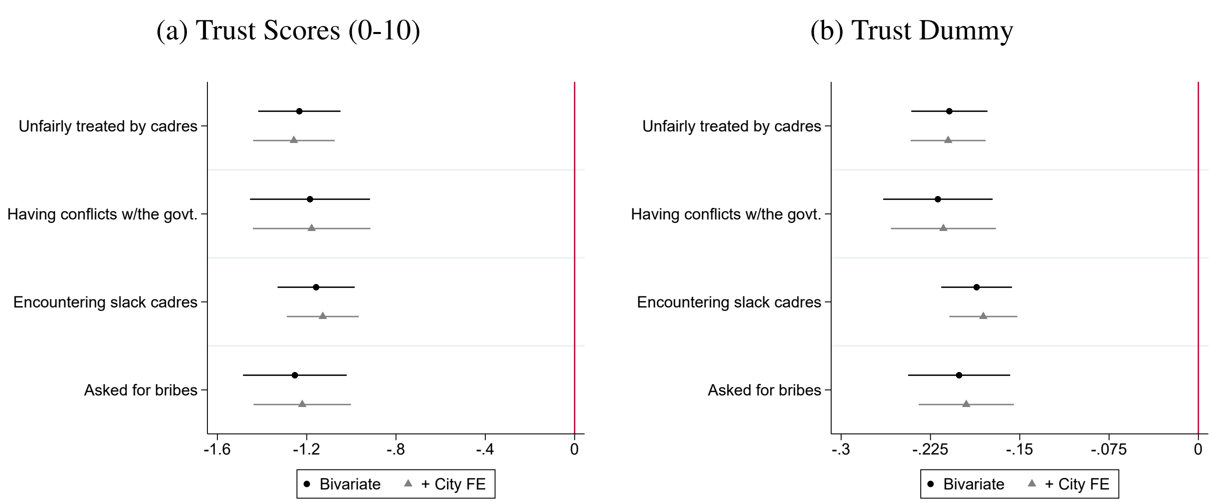

Note: In this figure, political trust is regressed on unpleasant experiences with the government. The solid dots are point estimates of coefficients on experiences and the caps are 90 percent confidence intervals. Robust standard errors are clustered at the city level.

Figure 4.11: Outcomes of Political Trust

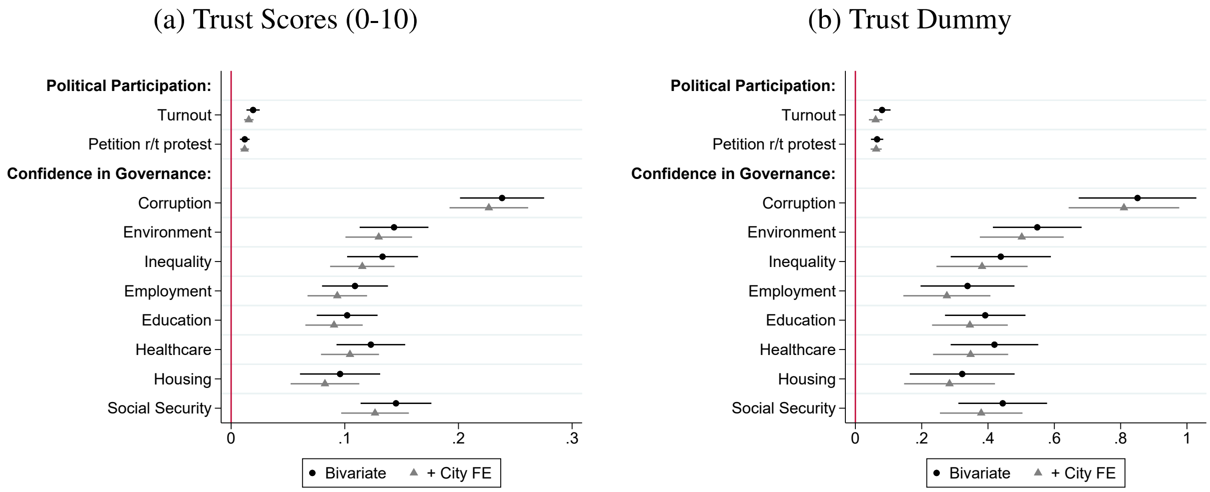

Note: In this figure, political outcomes are regressed political trust levels (measured on a 0–10 scale or using a dummy). There are two categories of behavioral outcomes: (1) political participation (voting in grassroots elections, petitioning rather than protesting if there is dissent) and (2) confidence in the governance of various issues. The solid dots are point estimates of coefficients on political trust and the caps are 90 percent confidence intervals. Robust standard errors are clustered at the city level.

Figure 4.12: Distribution of Change in Political Trust Dummy Within Individuals

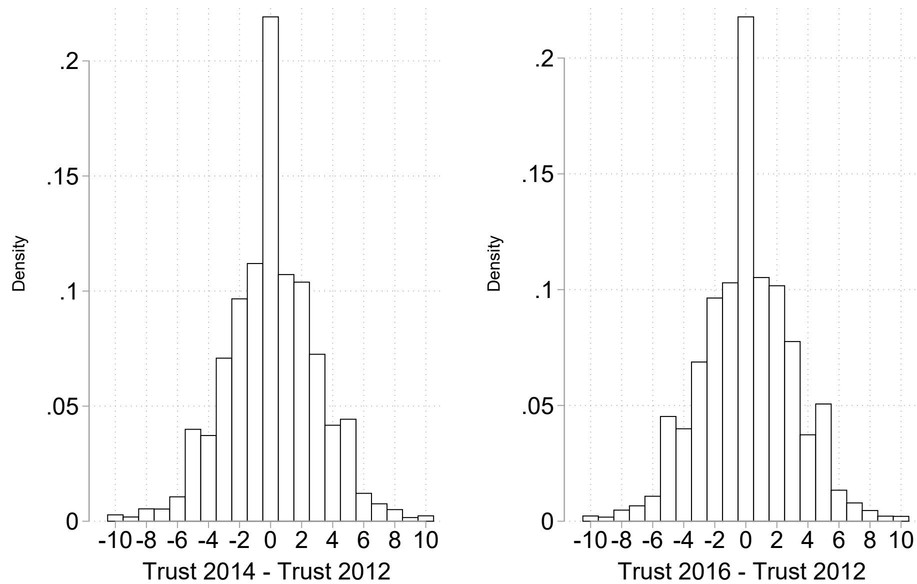

Note: This figure reports the distribution of within-individual change in reported political trust (scale 0–10) relative to 2012 in the balanced panel sample.

Figure 4.13: Effects of Investigations on Political Trust Before 2012

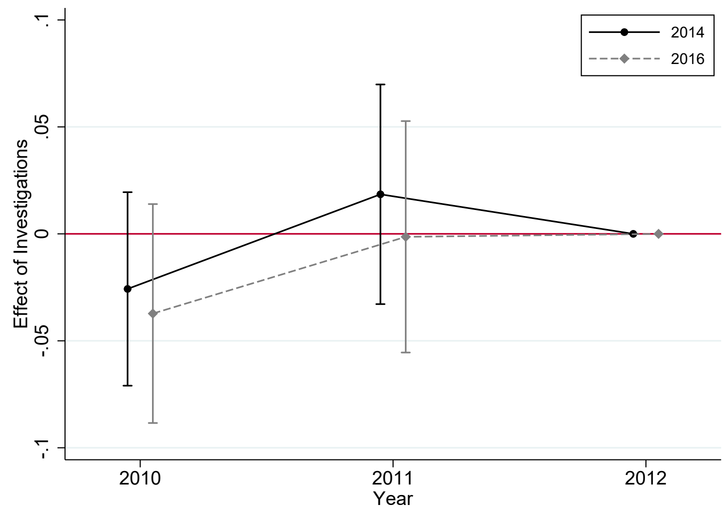

Note: The China General Social Survey (CGSS) elicited data on political trust (in local government) in 2010, 2011, and 2012 (before the campaign). In this figure, we regress political trust on upcoming investigations in 2014 and 2016 (province level, as CGSS only provides a province identifier), interacted with dummies for 2010, 2011, and 2012, and controlling for province and year fixed effects. The number of investigations (*D*) is standardized. The solid dots are point estimates of coefficients on investigations and the caps are 90 percent confidence intervals. Robust standard errors are clustered at the province level.

Figure 4.14: Placebo Test

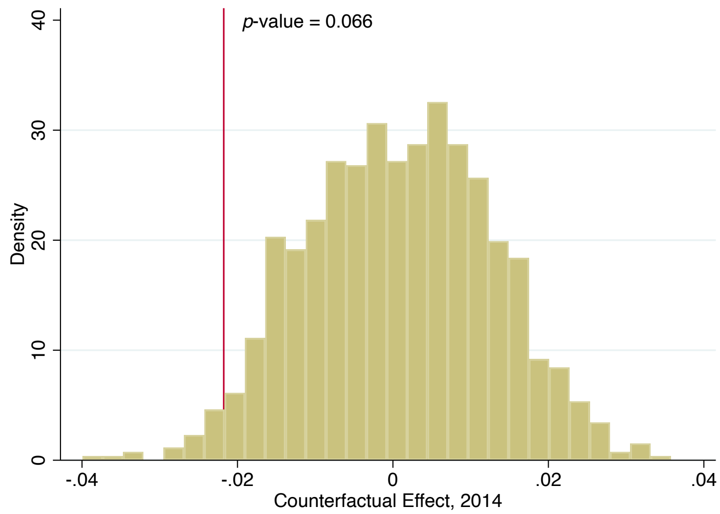

Note: This figure is derived from 1,000 permutations. The bars display the distribution of counterfactual estimates. The vertical line marks the true estimate.

Figure 4.15: Heterogeneous Effects by Education and Confucianism

(a) Effects in 2014 (b) Effects in 2016

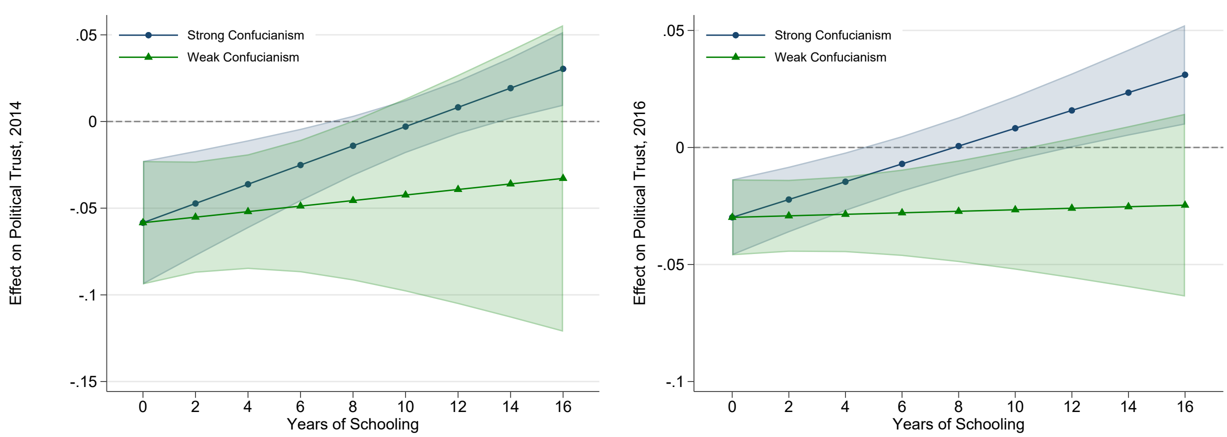

Note: This figure displays the estimated DID coefficients in Equation 4.6 estimated using subsamples of different levels of education and for regions with strong and weak Confucianism (Confucian temples above or below the median), respectively. Point estimates and 95% CI are plotted.

Table 4.9: Correlates Between Political Trust and Individual Characteristics

	(1)	(2)	(3)	(4)
	Political Trust Dummy		Political Trust (Scale 0--10)	
Age	0.003***	0.004***	0.020***	0.026***
	(0.000)	(0.001)	(0.002)	(0.003)
Male	-0.003	0.007	-0.079*	-0.012
	(0.009)	(0.010)	(0.045)	(0.052)
Han ethnicity	-0.076***	-0.085***	-0.524***	-0.583***
	(0.015)	(0.017)	(0.080)	(0.085)
Urban	-0.063***	-0.067***	-0.504***	-0.495***
	(0.009)	(0.011)	(0.048)	(0.055)
Communist Party member	0.108***	0.111***	0.613***	0.615***
	(0.017)	(0.019)	(0.083)	(0.094)
State sector employee	-0.028	-0.027	-0.076	-0.069
	(0.018)	(0.021)	(0.083)	(0.099)
Years of educ.	-0.001	-0.000	-0.013**	-0.005
	(0.001)	(0.001)	(0.006)	(0.007)
Degree completed, father	-0.008*	-0.008	-0.032	-0.023
	(0.005)	(0.006)	(0.026)	(0.029)
Degree completed, mother	-0.003	0.006	-0.067**	-0.031
	(0.007)	(0.008)	(0.033)	(0.039)
Communist Party member, father	-0.022*	-0.019	-0.036	-0.064
	(0.012)	(0.014)	(0.062)	(0.071)
Communist Party member, mother	0.041	0.069**	0.111	0.170
	(0.029)	(0.033)	(0.141)	(0.161)
Attention to corruption news		-0.032**		-0.257***
		(0.013)		(0.063)
ln(# Confucian temples)		0.001		-0.024
		(0.003)		(0.018)
Share of Govt. Weibo users		-0.266		-3.721
		(0.715)		(3.675)
Share of Delted Weibo posts		-0.285*		-1.325*
		(0.152)		(0.780)
Family purged in Revolutions		-0.034**		-0.236***
		(0.016)		(0.079)
Witnessed violent CR		-0.024**		-0.102*
		(0.011)		(0.056)
observations	12521	9840	12521	9840
R^2	0.018	0.021	0.034	0.037

Notes: Robust standard errors are reported in parentheses. * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

Table 4.10: Correlates of Cumulative Investigations

	Cum. Investigations 2014		Cum. Investigations 2016	
	(1)	(2)	(3)	(4)
Political trust dummy (2012)	0.130*	0.023	0.156**	0.026
	(0.069)	(0.074)	(0.071)	(0.065)
% Public employment	-0.076**	-0.011	-0.108***	-0.016
	(0.033)	(0.042)	(0.033)	(0.037)
% Private employment	-0.313***	-0.009	-0.317***	-0.101
	(0.107)	(0.213)	(0.100)	(0.228)
ln(GDP p.c.)	0.163	0.131	-0.098	-0.058
	(0.195)	(0.245)	(0.190)	(0.201)
ln(tax p.c.)	-0.286	-0.479*	-0.128	-0.340*
	(0.213)	(0.247)	(0.181)	(0.191)
ln(wage rate)	0.575***	0.832**	0.614***	0.821***
	(0.193)	(0.410)	(0.170)	(0.296)
GR. % public employment	-0.059	-0.081	-0.116*	-0.127*
	(0.061)	(0.084)	(0.064)	(0.076)
GR. % private employment	-0.195**	-0.248	-0.236**	-0.302
	(0.095)	(0.205)	(0.112)	(0.204)
GR. ln(GDP p.c.)	-0.010	-0.038	-0.017	-0.050
	(0.041)	(0.073)	(0.050)	(0.069)
GR. ln(tax p.c.)	0.167	0.228*	0.092	0.189*
	(0.110)	(0.134)	(0.088)	(0.099)
GR. ln(wage rate)	-0.112***	-0.076	-0.131***	-0.070
	(0.033)	(0.061)	(0.036)	(0.049)
Province FE		✓		✓
F stat.	7.865	0.797	8.779	1.020
F-test p-value	0.000	0.643	0.000	0.438
Obs.	113	113	113	113
R^2	0.351	0.562	0.310	0.637

Note: All dependents and independents are standardized. "GR." = growth rate. Economic variables are an average of values from 2000 to 2010. Robust standard errors are reported in parentheses. The null for the F test is that coefficients on all independents are zero.

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

Table 4.11: Covariates Balance by Change in Political Trust Dummy, 2012-2014

	(1)	(2)	(3)	(4)	(5)	(6)
ΔPolitical trust dummy	**-1**	**0**	**+1**	(1) vs. (2)	(1) vs. (3)	(2) vs. (3)
Panel (A): Anti-Corruption						
# Corruption investigations, 2014	28.544	28.366	26.869	0.178	1.675**	1.497**
	(0.589)	(0.297)	(0.562)	(0.665)	(0.814)	(0.648)
Panel (B): Attitudes 2012						
Trust: cadres	5.913	5.173	2.482	0.740***	3.431***	2.691***
	(0.033)	(0.029)	(0.029)	(0.060)	(0.044)	(0.058)
Trust: parents	9.239	9.190	8.801	0.048	0.438***	0.389***
	(0.035)	(0.017)	(0.045)	(0.039)	(0.057)	(0.041)
Trust: strangers	2.296	2.117	1.554	0.179***	0.742***	0.563***
	(0.049)	(0.024)	(0.037)	(0.054)	(0.061)	(0.051)
Trust: Americans	2.487	2.460	1.912	0.027	0.575***	0.548***
	(0.057)	(0.029)	(0.048)	(0.064)	(0.074)	(0.061)
Government performance	3.468	3.479	3.302	-0.011	0.166***	0.177***
	(0.020)	(0.011)	(0.021)	(0.023)	(0.029)	(0.023)
Panel (C): Experiences						
Experience: unfairly treated by cadres	0.074	0.092	0.091	-0.018**	-0.017*	0.001
	(0.006)	(0.003)	(0.006)	(0.007)	(0.009)	(0.007)
Experience: having conflicts w/cadres	0.031	0.036	0.037	-0.005	-0.006	-0.001
	(0.004)	(0.002)	(0.004)	(0.005)	(0.006)	(0.005)
Experience: slack cadres	0.122	0.128	0.118	-0.006	0.004	0.010
	(0.007)	(0.004)	(0.007)	(0.008)	(0.010)	(0.008)
Experience: asked for bribes	0.061	0.068	0.062	-0.007	-0.001	0.006
	(0.005)	(0.003)	(0.005)	(0.006)	(0.008)	(0.006)
Panel (D): Individual and Family Covaraites						
Birth cohort	1968.089	1967.159	1967.584	0.930***	0.506	-0.424
	(0.239)	(0.119)	(0.233)	(0.266)	(0.334)	(0.261)
Male	0.471	0.472	0.446	-0.001	0.025	0.026**
	(0.011)	(0.006)	(0.011)	(0.013)	(0.016)	(0.012)
Han ethnicity	0.923	0.922	0.939	0.001	-0.016**	-0.017***
	(0.006)	(0.003)	(0.005)	(0.007)	(0.008)	(0.006)
Urban	0.475	0.438	0.438	0.037***	0.037**	0.000
	(0.011)	(0.006)	(0.011)	(0.013)	(0.016)	(0.012)
Observations	1955	7857	2058	9812	4013	9915

Note: Column (1)–(3) report summary statistics of different variables by the change of political trust dummy from 2012 to 2014. Column (4)–(6) report t-test results of pairwise comparison. Standard errors are in parentheses. * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

Table 4.11: Covariates Balance by Change in Political Trust Dummy, 2012-2014 (Continued)

	(1)	(2)	(3)	(4)	(5)	(6)
ΔPolitical trust dummy	**-1**	**0**	**+1**	(1) vs. (2)	(1) vs. (3)	(2) vs. (3)
Panel (D): Individual and Family Covaraites, Continued						
Years of educ.	6.855	6.802	6.389	0.053	0.466***	0.413***
	(0.109)	(0.054)	(0.106)	(0.122)	(0.152)	(0.119)
Communist Party member	0.068	0.083	0.052	-0.015**	0.015**	0.030***
	(0.006)	(0.003)	(0.005)	(0.007)	(0.007)	(0.007)
State sector employee	0.073	0.085	0.072	-0.011	0.001	0.013*
	(0.006)	(0.003)	(0.006)	(0.007)	(0.008)	(0.007)
Years of educ., father	4.641	4.476	4.527	0.166	0.114	-0.052
	(0.097)	(0.047)	(0.094)	(0.107)	(0.135)	(0.105)
Years of educ., mother	2.722	2.581	2.519	0.141	0.203*	0.062
	(0.080)	(0.039)	(0.076)	(0.089)	(0.110)	(0.086)
Communist Party member, father	0.152	0.163	0.160	-0.011	-0.008	0.003
	(0.008)	(0.004)	(0.008)	(0.009)	(0.011)	(0.009)
Communist Party member, mother	0.030	0.021	0.020	0.008**	0.009*	0.001
	(0.004)	(0.002)	(0.003)	(0.004)	(0.005)	(0.004)
Panel (E): Other Variables						
ETC index	0.013	0.012	0.012	0.000**	0.000*	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Attention to corruption news	0.239	0.231	0.225	0.009	0.014	0.006
	(0.010)	(0.005)	(0.009)	(0.011)	(0.013)	(0.010)
Share of Delted Weibo posts	0.181	0.181	0.183	0.000	-0.002	-0.002*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Share of Govt. Weibo users	0.041	0.041	0.042	0.000*	0.000	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
ln(# Confucian temples)	4.548	4.606	4.533	-0.057	0.015	0.073*
	(0.039)	(0.019)	(0.037)	(0.042)	(0.054)	(0.041)
Military Family	0.063	0.062	0.068	0.002	-0.005	-0.006
	(0.006)	(0.003)	(0.006)	(0.006)	(0.008)	(0.006)
Family purged in Revolutions	0.119	0.131	0.117	-0.012	0.002	0.014*
	(0.007)	(0.004)	(0.007)	(0.008)	(0.010)	(0.008)
Witnessed violent CR	0.620	0.626	0.633	-0.005	-0.012	-0.007
	(0.011)	(0.005)	(0.011)	(0.012)	(0.015)	(0.012)
Observations	1955	7857	2058	9812	4013	9915

Note: Column (1)–(3) report summary statistics of different variables by the change of political trust dummy from 2012 to 2014. Column (4)–(6) report t-test results of pairwise comparison. Standard errors are in parentheses. * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

Table 4.12: Covariates Balance by Change in Political Trust Dummy, 2012-2016

	(1)	(2)	(3)	(4)	(5)	(6)
ΔPolitical trust dummy	**-1**	**0**	**+1**	(1) vs. (2)	(1) vs. (3)	(2) vs. (3)
Panel (A): Anti-Corruption						
# Corruption investigations, 2016	56.345	57.755	55.540	-1.411	0.804	2.215**
	(0.978)	(0.516)	(0.962)	(1.131)	(1.373)	(1.101)
Panel (B): Attitudes 2012						
Trust: cadres	5.944	5.210	2.426	0.733***	3.518***	2.784***
	(0.032)	(0.029)	(0.029)	(0.059)	(0.043)	(0.057)
Trust: parents	9.191	9.176	8.911	0.014	0.279***	0.265***
	(0.035)	(0.018)	(0.041)	(0.039)	(0.054)	(0.040)
Trust: strangers	2.307	2.107	1.604	0.200***	0.703***	0.503***
	(0.049)	(0.024)	(0.037)	(0.054)	(0.061)	(0.050)
Trust: Americans	2.542	2.440	1.955	0.102	0.586***	0.484***
	(0.057)	(0.029)	(0.047)	(0.064)	(0.074)	(0.060)
Government performance	3.468	3.480	3.305	-0.011	0.164***	0.175***
	(0.020)	(0.011)	(0.021)	(0.023)	(0.029)	(0.023)
Panel (C): Experiences						
Experience: unfairly treated by cadres	0.085	0.089	0.092	-0.005	-0.007	-0.002
	(0.006)	(0.003)	(0.006)	(0.007)	(0.009)	(0.007)
Experience: having conflicts w/cadres	0.035	0.032	0.045	0.003	-0.010	-0.013***
	(0.004)	(0.002)	(0.004)	(0.004)	(0.006)	(0.004)
Experience: slack cadres	0.125	0.120	0.144	0.004	-0.020*	-0.024***
	(0.007)	(0.004)	(0.008)	(0.008)	(0.011)	(0.008)
Experience: asked for bribes	0.061	0.064	0.079	-0.003	-0.018**	-0.015**
	(0.005)	(0.003)	(0.006)	(0.006)	(0.008)	(0.006)
Panel (D): Individual and Family Covaraites						
Birth cohort	1968.017	1967.146	1967.665	0.871***	0.353	-0.519**
	(0.237)	(0.120)	(0.227)	(0.265)	(0.328)	(0.257)
Male	0.461	0.475	0.445	-0.014	0.016	0.030**
	(0.011)	(0.006)	(0.011)	(0.013)	(0.015)	(0.012)
Han ethnicity	0.939	0.921	0.928	0.018***	0.012	-0.007
	(0.005)	(0.003)	(0.006)	(0.007)	(0.008)	(0.007)
Urban	0.435	0.442	0.459	-0.007	-0.024	-0.017
	(0.011)	(0.006)	(0.011)	(0.012)	(0.015)	(0.012)
Observations	1955	7857	2058	9812	4013	9915

Note: Column (1)–(3) report summary statistics of different variables by the change of political trust dummy from 2012 to 2016. Column (4)–(6) report t-test results of pairwise comparison. Standard errors are in parentheses. * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

Table 4.12: Covariates Balance by Change in Political Trust Dummy, 2012-2016 (Continued)

	(1)	(2)	(3)	(4)	(5)	(6)
ΔPolitical trust dummy	**-1**	**0**	**+1**	(1) vs. (2)	(1) vs. (3)	(2) vs. (3)
Panel (D): Individual and Family Covaraites, Continued						
Years of educ.	6.434	6.813	6.754	-0.379***	-0.320**	0.059
	(0.108)	(0.055)	(0.105)	(0.121)	(0.150)	(0.117)
Communist Party member	0.055	0.084	0.063	-0.029***	-0.008	0.021***
	(0.005)	(0.003)	(0.005)	(0.007)	(0.007)	(0.007)
State sector employee	0.079	0.080	0.084	-0.001	-0.005	-0.004
	(0.006)	(0.003)	(0.006)	(0.007)	(0.009)	(0.007)
Years of educ., father	4.485	4.500	4.578	-0.016	-0.094	-0.078
	(0.095)	(0.048)	(0.093)	(0.106)	(0.133)	(0.103)
Years of educ., mother	2.624	2.589	2.582	0.035	0.042	0.007
	(0.079)	(0.040)	(0.075)	(0.088)	(0.109)	(0.085)
Communist Party member, father	0.158	0.159	0.169	-0.001	-0.011	-0.010
	(0.008)	(0.004)	(0.008)	(0.009)	(0.012)	(0.009)
Communist Party member, mother	0.022	0.023	0.021	-0.002	0.001	0.002
	(0.003)	(0.002)	(0.003)	(0.004)	(0.004)	(0.004)
Panel (E): Other Variables						
ETC index	0.012	0.012	0.012	0.000	0.000	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Attention to corruption news	0.233	0.230	0.233	0.002	0.000	-0.003
	(0.009)	(0.005)	(0.009)	(0.011)	(0.013)	(0.010)
Share of Delted Weibo posts	0.186	0.180	0.183	0.006***	0.003	-0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Share of Govt. Weibo users	0.042	0.041	0.041	0.001***	0.000	-0.001**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
ln(# Confucian temples)	4.459	4.613	4.595	-0.154***	-0.136***	0.018
	(0.038)	(0.019)	(0.036)	(0.042)	(0.052)	(0.041)
Military Family	0.055	0.063	0.069	-0.008	-0.014*	-0.005
	(0.005)	(0.003)	(0.005)	(0.006)	(0.008)	(0.006)
Family purged in Revolutions	0.114	0.129	0.130	-0.015*	-0.016	-0.002
	(0.007)	(0.004)	(0.007)	(0.008)	(0.010)	(0.008)
Witnessed violent CR	0.606	0.632	0.623	-0.026**	-0.016	0.009
	(0.011)	(0.005)	(0.010)	(0.012)	(0.015)	(0.012)
Observations	1994	7721	2155	9715	4149	9876

Note: Column (1)–(3) report summary statistics of different variables by the change of political trust dummy from 2012 to 2016. Column (4)–(6) report t-test results of pairwise comparison. Standard errors are in parentheses. * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

Table 4.13: Effect of Anti-Corruption on Political Trust (Scale 0–10)

	(1)	(2)	(3)	(4)
$D^{14} \times T^{14}$	-0.095**	-0.093**	-0.073*	-0.070*
	(0.041)	(0.042)	(0.040)	(0.039)
$D^{16} \times T^{16}$	-0.031	-0.031	-0.045	-0.046
	(0.042)	(0.041)	(0.040)	(0.040)
D.V. mean, pre-campaign	4.828	4.828	4.828	4.828
Individual FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Cohort × year FE		✓	✓	✓
Indiv. char. × year FE			✓	✓
Fam. bkgd. × year FE				✓
Obs.	35610	35610	35610	35610
R^2	0.568	0.570	0.571	0.572

Note: The dependent variable is the political trust scale (0–10). The number of investigations (D) is standardized. Individual characteristics include gender, indicators of educational attainment, *hukou* status, Han ethnicity, Communist Party membership, and state sector employment. Family background includes parents' educational attainment and their Communist Party membership. Robust standard errors, clustered at the city level, are reported in parentheses.

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

Table 4.14: Robustness Check: Heterogeneity-Robust Estimator

	β	SE	p-value
2014, 25%	-0.040	0.023	0.079
2014, 50%	-0.036	0.019	0.061
2016, 25%	-0.015	0.014	0.283
2016, 50%	-0.008	0.012	0.485

Note: The dependent variable is the political trust dummy. de Chaisemartin et al. (2022b)'s heterogeneity-robust estimator is implemented. For implementation, a low-intensity group needs to be defined for comparison. We define it as being below the first quartile or the median.

Table 4.15: Education, Socioeconomic Status, and Political Trust

	(1)	(2)	(3)	(4)	(5)
$D^{14} \times T^{14}$	-0.062***	-0.057***	-0.063***	-0.063***	-0.059***
	(0.015)	(0.017)	(0.015)	(0.015)	(0.017)
$D^{16} \times T^{16}$	-0.041***	-0.047***	-0.042***	-0.042***	-0.048***
	(0.008)	(0.009)	(0.008)	(0.008)	(0.009)
$D^{14} \times T^{14}$ × Schooling	0.005***	0.005***	0.006***	0.006***	0.006***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$D^{16} \times T^{16}$ × Schooling	0.004***	0.004***	0.005***	0.005***	0.004***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$D^{14} \times T^{14}$ × Urban *hukou*		-0.007			-0.006
		(0.017)			(0.017)
$D^{16} \times T^{16}$ × Urban *hukou*		0.017			0.017
		(0.011)			(0.012)
$D^{14} \times T^{14}$ × CPC member			-0.042**		-0.039**
			(0.019)		(0.020)
$D^{16} \times T^{16}$ × CPC member			-0.011		-0.010
			(0.014)		(0.014)
$D^{14} \times T^{14}$ × State employee				-0.022	-0.018
				(0.015)	(0.015)
$D^{16} \times T^{16}$ × State employee				-0.007	-0.008
				(0.013)	(0.013)
D.V. mean, pre-campaign	0.633	0.633	0.633	0.633	0.633
Individual FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Covariates	✓	✓	✓	✓	✓
Obs.	35610	35610	35610	35610	35610
R^2	0.509	0.509	0.509	0.509	0.509

Note: The dependent variable is the political trust dummy. The number of investigations (D) is standardized. Individual characteristics include gender, indicators of educational attainment, *hukou* status, Han ethnicity, Communist Party membership, and state sector employment. Family background includes parents' educational attainment and their Communist Party membership. Robust standard errors, clustered at the city level, are reported in parentheses.

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

Bibliography

Abraham, Sarah, and Liyang Sun. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*.

Acemoglu, Daron, and James A Robinson. 2020. *The narrow corridor: States, societies, and the fate of liberty*. Penguin Books.

Acemoglu, Daron, and James A Robinson. 2021a. "Culture, Institutions and Social Equilibria: A Framework."Technical report, National Bureau of Economic Research.

Acemoglu, Daron, and James A Robinson. 2021b. "Non-Modernization: Power-Culture Trajectories and the Dynamics of Political Institutions."Technical report, National Bureau of Economic Research.

Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa, and Ekaterina Zhuravskaya. 2015. "Radio and the Rise of the Nazis in Prewar Germany." *The Quarterly Journal of Economics* 130 (4): 1885–1939.

Aghion, Philippe, Xavier Jaravel, Torsten Persson, and Dorothée Rouzet. 2019. "Education and military rivalry." *Journal of the European Economic Association* 17 (2): 376–412.

Alesina, Alberto F, Marlon Seror, David Y Yang, Yang You, Weihong Zeng et al. 2020. *Persistence through revolutions*. National Bureau of Economic Research.

Alesina, Alberto, Paola Giuliano, and Bryony Reich. 2021. "Nation-building and education." *The Economic Journal* 131 (638): 2273–2303.

Algan, Yann, and Pierre Cahuc. 2010. "Inherited trust and growth." *American Economic Review* 100 (5): 2060–92.

Algan, Yann, and Pierre Cahuc. 2014. "Trust, growth, and well-being: New evidence and policy implications." In *Handbook of Economic Growth*, Volume 2. 49–120, Elsevier.

Allen, Nicholas, and Sarah Birch. 2015. "Process preferences and British public opinion: Citizens' judgements about government in an era of anti-politics." *Political Studies* 63 (2): 390–411.

Alm, James, Weizheng Lai, and Xun Li. 2022. "Housing market regulations and strategic divorce propensity in China." *Journal of Population Economics* 35 (3): 1103–1131.

Almond, Gabriel, and Sidney Verba. 1963. *The Civic Culture:Political Attitudes and Democracy in Five Nations*. Princeton, NJ: Princeton University Press.

Alwin, Duane F, and Jon A Krosnick. 1991. "Aging, cohorts, and the stability of sociopolitical orientations over the life span." *American Journal of Sociology* 97 (1): 169–195.

Anderson, Christopher J, and Yuliya V Tverdova. 2003. "Corruption, political allegiances, and attitudes toward government in contemporary democracies." *American Journal of Political Science* 47 (1): 91–109.

André, Pierre, and Yannick Dupraz. 2019. "Education and hypergamy: evidence from Cameroon."

Ang, Yuen Yuen. 2020. *China's gilded age: The paradox of economic boom and vast corruption*. Cambridge University Press.

Arrow, Kenneth J. 1972. "Gifts and Exchanges." *Philosophy & Public Affairs* 343–362.

Artmann, Elisabeth, Nadine Ketel, Hessel Oosterbeek, and Bas van der Klaauw. 2021. "Field of study and partner choice." *Economics of Education Review* 84 102149.

Attanasio, Orazio P, and Katja M Kaufmann. 2017. "Education choices and returns on the labor and marriage markets: Evidence from data on subjective expectations." *Journal of Economic Behavior & Organization* 140 35–55.

Azoulay, Pierre, Joshua S Graff Zivin, and Jialan Wang. 2010. "Superstar extinction." *The Quarterly Journal of Economics* 125 (2): 549–589.

Bandiera, Oriana, Myra Mohnen, Imran Rasul, and Martina Viarengo. 2019. "Nation-building through compulsory schooling during the age of mass migration." *The Economic Journal* 129 (617): 62–109.

Banerjee, Abhijit, Rema Hanna, and Sendhil Mullainathan. 2013. "Corruption." In *The Handbook of Organizational Economics*, edited by Gibbons, Robert, and John Roberts, Chap. 27 1109–1147, Princeton University Press.

Banerjee, Ritwik. 2016. "Corruption, norm violation and decay in social capital." *Journal of Public Economics* 137 14–27.

Becker, Gary S. 1973. "A theory of marriage: Part I." *Journal of Political economy* 81 (4): 813–846.

Bell, Daniel A. 2010. "China's new Confucianism." In *China's New Confucianism*, Princeton University Press.

Bellemare, Marc F, and Casey J Wichman. 2020. "Elasticities and the inverse hyperbolic sine transformation." *Oxford Bulletin of Economics and Statistics* 82 (1): 50–61.

Biagioli, Mario, Martin Kenney, Ben R Martin, and John P Walsh. 2019. "Academic misconduct, misrepresentation and gaming: A reassessment." *Research Policy* 48 (2): 401–413.

Bisgaard, Martin. 2015. "Bias will find a way: Economic perceptions, attributions of blame, and partisan-motivated reasoning during crisis." *The Journal of Politics* 77 (3): 849–860.

Bisgaard, Martin, and Rune Slothuus. 2018. "Partisan elites as culprits? How party cues shape partisan perceptual gaps." *American Journal of Political Science* 62 (2): 456–469.

Bjørnskov, Christian. 2011. "Combating corruption: On the interplay between institutional quality and social trust." *The Journal of Law and Economics* 54 (1): 135–159.

Bollyky, Thomas J, Erin N Hulland, Ryan M Barber et al. 2022. "Pandemic preparedness and COVID-19: an exploratory analysis of infection and fatality rates, and contextual factors associated with preparedness in 177 countries, from Jan 1, 2020, to Sept 30, 2021." *The Lancet* 399 (10334): 1489–1512.

Borjas, George J, and Kirk B Doran. 2015a. "Cognitive mobility: Labor market responses to supply shocks in the space of ideas." *Journal of Labor Economics* 33 (S1): S109–S145.

Borjas, George J, and Kirk B Doran. 2015b. "Prizes and productivity how winning the fields medal affects scientific output." *Journal of human resources* 50 (3): 728–758.

Borjas, George J, Kirk B Doran, and Ying Shen. 2018. "Ethnic complementarities after the opening of China how chinese graduate students affected the productivity of their advisors." *Journal of Human Resources* 53 (1): 1–31.

Bruze, Gustaf. 2015. "Male and female marriage returns to schooling." *International Economic Review* 56 (1): 207–234.

Bullock, John G. 2009. "Partisan bias and the Bayesian ideal in the study of public opinion." *The Journal of Politics* 71 (3): 1109–1124.

Burgess, Simon, Carol Propper, and Arnstein Aassve. 2003. "The role of income in marriage and divorce transitions among young Americans." *Journal of Population Economics* 16 (3): 455–475.

Cai, Hongbin, Hanming Fang, and Lixin Colin Xu. 2011. "Eat, drink, firms, government: An investigation of corruption from the entertainment and travel costs of Chinese firms." *The Journal of Law and Economics* 54 (1): 55–78.

Callaway, Brantly, and Pedro HC Sant'Anna. 2021. "Difference-in-differences with multiple time periods." *Journal of Econometrics* 225 (2): 200–230.

Campante, Filipe R, and Davin Chor. 2012. "Schooling, political participation, and the economy." *Review of Economics and Statistics* 94 (4): 841–859.

Cantoni, Davide, Yuyu Chen, David Y Yang, Noam Yuchtman, and Y Jane Zhang. 2017. "Curriculum and ideology." *Journal of Political Economy* 125 (2): 338–392.

Cantoni, Davide, and Noam Yuchtman. 2014. " Medieval Universities, Legal Institutions, and the Commercial Revolution ." *The Quarterly Journal of Economics* 129 (2): 823–887. 10.1093/qje/qju007.

Carothers, Christopher. 2022. *Corruption Control in Authoritarian Regimes: Lessons from East*

Asia. Cambridge University Press.

de Chaisemartin, Clément, Xavier D'Haultfoeuille, Félix Pasquier, and Gonzalo Vazquez-Bare. 2022a. "Difference-in-Differences Estimators for Treatments Continuously Distributed at Every Period." *arXiv preprint arXiv:2201.06898*.

de Chaisemartin, Clément, Xavier D'Haultfœuille, Félix Pasquier, and Gonzalo Vazquez-Bare. 2022b. "Difference-in-Differences Estimators for Treatments Continuously Distributed at Every Period." *Available at SSRN*.

de Chaisemartin, Clément, Xavier D'Haultfœuille, and Yannick Guyonvarch. 2019. "Fuzzy differences-in-differences with stata." *The Stata Journal* 19 (2): 435–458.

Chang, Eric CC, and Yun-han Chu. 2006. "Corruption and trust: exceptionalism in Asian democracies?" *The Journal of Politics* 68 (2): 259–271.

Che, Yi, and Lei Zhang. 2018. "Human capital, technology adoption and firm performance: Impacts of China's higher education expansion in the late 1990s." *The Economic Journal* 128 (614): 2282–2320.

Chen, Ting, Yizhen Gu, and Ben Zou. 2022. "Delineating China's Metropolitan Areas Using Commuting Flow Data." *Available at SSRN*.

Chen, Ting, and James Kai-sing Kung. 2019a. "Busting the "Princelings": The campaign against corruption in China's primary land market." *The Quarterly Journal of Economics* 134 (1): 185–226.

Chen, Ting, and James Kai-sing Kung. 2019b. "Busting the "Princelings": The campaign against corruption in China's primary land market." *The Quarterly Journal of Economics* 134 (1): 185–226.

Chen, Ting, James Kai-sing Kung, and Chicheng Ma. 2020. "Long live Keju! The persistent effects of China's civil examination system." *The Economic Journal* 130 (631): 2030–2064.

Chen, Yuyu, and David Y Yang. 2019. "Historical Traumas and the Roots of Political Distrust: Political Inference from the Great Chinese Famine."

Chiappori, Pierre-André. 2017. *Matching with transfers: The economics of love and marriage*. Princeton University Press.

Chiappori, Pierre-Andre, Monica Costa Dias, and Costas Meghir. 2021. "The measuring of assortativeness in marriage: a comment."

Chiappori, Pierre-André, Monica Costa Dias, and Costas Meghir. 2018. "The marriage market, labor supply, and education choice." *Journal of Political Economy* 126 (S1): S26–S72.

Chiappori, Pierre-André, Monica Costa Dias, and Costas Meghir. 2020. "Changes in assortative matching: Theory and evidence for the us."

Chiappori, Pierre-Andre, Murat Iyigun, and Yoram Weiss. 2009. "Investment in schooling and the marriage market." *American Economic Review* 99 (5): 1689–1713.

Chiappori, Pierre-André, Bernard Salanié, and Yoram Weiss. 2017. "Partner choice, investment in children, and the marital college premium." *American Economic Review* 107 (8): 2109–67.

Chong, Alberto, Ana L De La O, Dean Karlan, and Leonard Wantchekon. 2015. "Does corruption information inspire the fight or quash the hope? A field experiment in Mexico on voter turnout, choice, and party identification." *The Journal of Politics* 77 (1): 55–71.

Choo, Eugene, and Aloysius Siow. 2006. "Who marries whom and why." *Journal of political Economy* 114 (1): 175–201.

Chu, Junhong, Haoming Liu, and Ivan PL Png. 2018. "Nonlabor income and age at marriage: evidence from China's heating policy." *Demography* 55 (6): 2345–2370.

Ciscato, Edoardo, and Simon Weber. 2020. "The role of evolving marital preferences in growing income inequality." *Journal of Population Economics* 33 (1): 307–347.

Confucius. 2003. *The Analects*. Indianapolis: Hackett Publishing, , 1st edition 193.

Cotofan, Maria, Lea Cassar, Robert Dur, and Stephan Meier. 2020. "Macroeconomic Conditions When Young Shape Job Preferences for Life." *Review of Economics and Statistics* 1–20.

Croke, Kevin, Guy Grossman, Horacio A Larreguy, and John Marshall. 2016. "Deliberate disengagement: How education can decrease political participation in electoral authoritarian regimes." *American Political Science Review* 110 (3): 579–600.

Cullen, Julie Berry, Nicholas Turner, and Ebonya Washington. 2021. "Political Alignment, Attitudes Toward Government, and Tax Evasion." *American Economic Journal: Economic Policy* 13 (3): 135–66.

Dee, Thomas S. 2004. "Are there civic returns to education?" *Journal of Public Economics* 88 (9-10): 1697–1720.

Ding, Haoyuan, Hanming Fang, Shu Lin, and Kang Shi. 2020. "Equilibrium consequences of corruption on firms: Evidence from china's anti-corruption campaign."Technical report, National Bureau of Economic Research.

Dixit, Avinash K, and Jörgen W Weibull. 2007. "Political polarization." *Proceedings of the National Academy of Sciences* 104 (18): 7351–7356.

Ebenstein, Avraham Y, and Ethan Jennings Sharygin. 2009. "The consequences of the "missing girls" of China." *The world bank economic review* 23 (3): 399–425.

Economist. 2021. "How did Confucianism win back the Chinese Communist Party?".

Eika, Lasse, Magne Mogstad, and Basit Zafar. 2019. "Educational assortative mating and household income inequality." *Journal of Political Economy* 127 (6): 2795–2835.

Else, Holly. 2018. "Does Science Have A Bullying Problem." *Nature* 563.

Enikolopov, Ruben, Maria Petrova, and Konstantin Sonin. 2018. "Social media and corruption." *American Economic Journal: Applied Economics* 10 (1): 150–74.

Fang, Ferric C, R Grant Steen, and Arturo Casadevall. 2012. "Misconduct accounts for the majority of retracted scientific publications." *Proceedings of the National Academy of Sciences* 109 (42): 17028–17033.

Fang, Hanming, Quanlin Gu, and Li-An Zhou. 2019. "The gradients of power: Evidence from the Chinese housing market." *Journal of Public Economics* 176 32–52.

Farzanegan, Mohammad Reza, and Hans Philipp Hofmann. 2021. "A Matter of Trust? Political Trust and the Covid-19 Pandemic."

Feng, Shuaizhang, Xiaoyu Xia et al. 2018. "Endogenous skill-biased technology adoption: Evidence from China's college enrollment expansion program."

Feng, Yuan. 1999. "National college entrance examinations: The dynamics of political centralism in China's elite education." *Journal of education* 181 (1): 39–57.

Fernández, Raquel, and Richard Rogerson. 2001. "Sorting and long-run inequality." *The Quarterly Journal of Economics* 116 (4): 1305–1341.

Ferraz, Claudio, and Frederico Finan. 2008. "Exposing corrupt politicians: the effects of Brazil's publicly released audits on electoral outcomes." *The Quarterly Journal of Economics* 123 (2): 703–745.

Fincher, Leta Hong. 2016. *Leftover women: The resurgence of gender inequality in China*. Bloomsbury Publishing.

Fisman, Raymond, Jing Shi, Yongxiang Wang, and Rong Xu. 2018. "Social ties and favoritism in Chinese science." *Journal of Political Economy* 126 (3): 1134–1171.

Fukuyama, Francis. 2020. "The thing that determines a country's resistance to the coronavirus." *The Atlantic* 30.

Gaziano, Cecilie, and Kristin McGrath. 1986. "Measuring the concept of credibility." *Journalism Quarterly* 63 (3): 451–462.

Ge, Run, and Jialin Huang. 2020. "Has the college expansion affected marriage and fertility?" *China Journal of Economics (in Chinese)* 3 168–201.

Ge, Suqin. 2011. "Women's college decisions: how much does marriage matter?" *Journal of Labor Economics* 29 (4): 773–818.

Gerber, Alan, and Donald Green. 1999. "Misperceptions about perceptual bias." *Annual Review of Political Science* 2 (1): 189–210.

Gewin, Virginia. 2021. "How to blow the whistle on an academic bully." *Nature* 593 (7858): 299–301.

Gingerich, Daniel W et al. 2009. "Corruption and political decay: Evidence from Bolivia." *Quarterly Journal of Political Science* 4 (1): 1–34.

Giuliano, Paola, and Antonio Spilimbergo. 2014. "Growing up in a Recession." *Review of Economic Studies* 81 (2): 787–817.

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics* 225 (2): 254–277.

Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos. 2014. "Marry your like: Assortative mating and income inequality." *American Economic Review* 104 (5): 348–53.

Guriev, Sergei, Nikita Melnikov, and Ekaterina Zhuravskaya. 2021. "3g internet and confidence in government." *The Quarterly Journal of Economics* 136 (4): 2533–2613.

Hankins, Scott, and Mark Hoekstra. 2011. "Lucky in life, unlucky in love? The effect of random income shocks on marriage and divorce." *Journal of Human Resources* 46 (2): 403–426.

Hausman. 2012. "University Innovation, Local Economic Growth, and Entrepreneurship." (12-10): , https://ideas.repec.org/p/cen/wpaper/12-10.html.

He, Minghong, Han Yu, Zhiyong Han et al. 2018. "NSFC Annual Report 2017."Technical report, The National Natural Science Foundation of China.

Hener, Timo, and Tanya Wilson. 2018. "Marital age gaps and educational homogamy-evidence from a compulsory schooling reform in the uk."

Hetherington, Marc J. 1998. "The political relevance of political trust." *American Political Science Review* 92 (4): 791–808.

Hsieh, Chang-Tai, Erik Hurst, Charles I Jones, and Peter J Klenow. 2019. "The allocation of talent and us economic growth." *Econometrica* 87 (5): 1439–1474.

Hsieh, Chang-Tai, and Peter J Klenow. 2009. "Misallocation and manufacturing TFP in China and India." *The Quarterly journal of economics* 124 (4): 1403–1448.

Huntington, Samuel P. 1991. "Democracy's third wave." *Journal of Democracy* 2 (2): 12–34.

Hvide, Hans K, and Benjamin F Jones. 2018. "University innovation and the professor's privilege." *American Economic Review* 108 (7): 1860–98.

Jaravel, Xavier, Neviana Petkova, and Alex Bell. 2018. "Team-specific capital and innovation." *American Economic Review* 108 (4-5): 1034–73.

Javed, Jeffrey A. 2022. *Righteous Revolutionaries: Morality, Mobilization, and Violence in the Making of the Chinese State*. University of Michigan Press.

Jia, Ruixue, Huihua Nie, and Wei Xiao. 2019. "Power and publications in Chinese academia." *Journal of Comparative Economics* 47 (4): 792–805.

Jiang, Junyan, Zijie Shao, and Zhiyuan Zhang. 2020. "The Price of Probity: Anticorruption and Adverse Selection in the Chinese Bureaucracy." *British Journal of Political Science* 1–24.

Jiang, Qing. 2016. *A Confucian Constitutional Order: How China's Ancient Past Can Shape Its Political Future*. Volume 4. Princeton University Press.

Kai-Ming, Cheng, Jin Xinhuo, and Gu Xiaobo. 1999. "From training to education. Lifelong learning in China." *Comparative Education* 35 (2): 119–129.

Khan, Adnan, Sanval Nasim, Mahvish Shaukat, and Andreas Stegmann. 2021. "Building trust in the state with information: Evidence from urban Punjab." *Journal of Public Economics* 202 104494.

King, Gary, Jennifer Pan, and Margaret E Roberts. 2013. "How censorship in China allows government criticism but silences collective expression." *American political science Review* 107 (2): 326–343.

King, Gary, Jennifer Pan, and Margaret E Roberts. 2017. "How the Chinese government fabricates social media posts for strategic distraction, not engaged argument." *American political science review* 111 (3): 484–501.

Kirkebøen, Lars, Edwin Leuven, and Magne Mogstad. 2021. "College as a Marriage Market."

Kong, Dongmin, and Ni Qin. 2019. "China's Anti-Corruption Campaign and Entrepreneurship." *Journal of Law and Economics, forthcoming*.

Kong, Dongmin, and Ni Qin. 2021. "China's Anticorruption Campaign and Entrepreneurship."

The Journal of Law and Economics 64 (1): 153–180.

Kong, Dongmin, Yunqing Tao, and Yanan Wang. 2020. “China’s anti-corruption campaign and firm productivity: Evidence from a quasi-natural experiment.” *China Economic Review* 63 101535.

Kung, James Kai-sing, and Chicheng Ma. 2014. “Can cultural norms reduce conflicts? Confucianism and peasant rebellions in Qing China.” *Journal of Development Economics* 111 132–149.

Lafortune, Jeanne. 2013. “Making yourself attractive: Pre-marital investments and the returns to education in the marriage market.” *American Economic Journal: Applied Economics* 5 (2): 151–78.

Lai, Weizheng, and Xun Li. 2021. “China’s Anti-Corruption Campaign and Civil Servant Fever.” *Working Paper*.

Levi, Margaret, and Laura Stoker. 2000. “Political trust and trustworthiness.” *Annual Review of Political Science* 3 (1): 475–507.

Li, Hongbin, Yueyuan Ma, Lingsheng Meng, Xue Qiao, and Xinzheng Shi. 2017. “Skill complementarities and returns to higher education: Evidence from college enrollment expansion in China.” *China Economic Review* 46 10–26.

Liu, Haoming, and Jingfeng Lu. 2006. “Measuring the degree of assortative mating.” *Economics Letters* 92 (3): 317–322.

Lorentzen, Peter. 2014. “China’s strategic censorship.” *American Journal of Political Science* 58 (2): 402–414.

Lott, John R, Jr. 1999. “Public schooling, indoctrination, and totalitarianism.” *Journal of Political Economy* 107 (S6): S127–S157.

Ma, Xiao. 2020. “College Expansion, Trade, and Innovation: Evidence from China.”

Magistad, Mary Kay. 2013. “China’s ‘leftover women’, unmarried at 27.” *BBC News* 21.

Mahmoudi, Morteza. 2019. “Academic bullies leave no trace.” *BioImpacts: BI* 9 (3): 129.

Marioulas, Julian. 2017. “China: A world leader in graduation rates.” *International Higher Education* (90): 28–29.

McFadden, Daniel et al. 1973. “Conditional logit analysis of qualitative choice behavior.”

Murphy, Kevin M, Andrei Shleifer, and Robert W Vishny. 1991. “The allocation of talent: Implications for growth.” *The quarterly journal of economics* 106 (2): 503–530.

Newman, Benjamin, Jennifer L Merolla, Sono Shah, Danielle Casarez Lemi, Loren Collingwood, and S Karthick Ramakrishnan. 2021. “The Trump effect: An experimental investigation of the emboldening effect of racially inflammatory elite communication.” *British Journal of Political Science* 51 (3): 1138–1159.

Nybom, Martin, Erik Plug, Bas Van der Klaauw, and Lennart Ziegler. 2017. “Skills, Parental Sorting, and Child Inequality.”Technical report, Working Paper.

Oransky, Ivan. 2020. “Retraction watch: What we’ve learned and how metrics play a role.” *Gaming the Metrics* 141.

Page, Jeremy. 2015. “Why China is turning back to Confucius.” *The Wall Street Journal* 20.

Pan, Jennifer, and Yiqing Xu. 2018. “China’s ideological spectrum.” *The Journal of Politics* 80 (1): 254–273.

Perry, Elizabeth J. 2008. “Chinese conceptions of “rights”: from Mencius to Mao—and now.” *Perspectives on politics* 6 (1): 37–50.

Petty, Richard E, and Jon A Krosnick. 2014. *Attitude strength: Antecedents and consequences*. Psychology Press.

Prendergast, Canice. 1993. “A theory of” yes men”.” *The American Economic Review* 757–770.

Prendergast, Canice, and Robert H Topel. 1996. “Favoritism in organizations.” *Journal of Political Economy* 104 (5): 958–978.

Putnam, Robert D, Robert Leonardi, and Raffaella Y Nanetti. 1993. *Making democracy work: Civic traditions in modern Italy*. Princeton University Press Princeton, NJ, USA.

Qi, Dongtao, Suixin Zhang, and Shengqiao Lin. 2022. “Urban Chinese Support for Armed

Unification with Taiwan: Social Status, National Pride and Understanding of Taiwan." *Journal of Contemporary China* 1–18.

Qin, Bei, David Strömberg, and Yanhui Wu. 2017. "Why does China allow freer social media? Protests versus surveillance and propaganda." *Journal of Economic Perspectives* 31 (1): 117–40.

Ramirez, Francisco O, and John Boli. 1987. "The political construction of mass schooling: European origins and worldwide institutionalization." *Sociology of Education* 2–17.

Rothstein, Bo. 2020. "Trust is the key to fighting the pandemic." *Scientific American* 24 1–7.

Saka, O, B Eichengreen, and C Aksoy. 2022. "The Political Scar of Epidemics." *The Economic Journal*.

Salcher, Isabelle. 2020. "Schooling and Marriage Markets How an Education Reform Impacted Matching Patterns: Evidence from Zimbabwe."

Sapienza, Paola, and Luigi Zingales. 2013. "Economic experts versus average Americans." *American Economic Review* 103 (3): 636–42.

Schofer, Evan, and John W Meyer. 2005. "The worldwide expansion of higher education in the twentieth century." *American sociological review* 70 (6): 898–920.

Seligson, Mitchell A. 2002. "The impact of corruption on regime legitimacy: A comparative study of four Latin American countries." *The Journal of Politics* 64 (2): 408–433.

Shellenbarger, Sue. 2019. "Hey, That's My Idea! When Your Boss Steals Your Work." *The Wall Street Journal*, https://www.wsj.com/articles/hey-thats-my-idea-when-your-boss-steals-your-work-1543329410.

Shi, Tianjian. 2001. "Cultural values and political trust: a comparison of the People's Republic of China and Taiwan." *Comparative Politics* 401–419.

Shi, Yigong, and Yi Rao. 2010. "China's research culture." *Science* 329 (5996): 1128–1128.

Siow, Aloysius. 2015. "Testing Becker's theory of positive assortative matching." *Journal of Labor Economics* 33 (2): 409–441.

Smith, Jacquelyn. 2013. "9 Things You Can Do When The Boss Takes Credit For Your

Work." *Forbes*, https://www.forbes.com/sites/jacquelynsmith/2013/04/30/9-things-you-can-do-when-the-boss-takes-credit-for-your-work/?sh=66cdeeb77951.

Sondheimer, Rachel Milstein, and Donald P Green. 2010. "Using experiments to estimate the effects of education on voter turnout." *American Journal of Political Science* 54 (1): 174–189.

Spenkuch, Jorg L, Edoardo Teso, and Guo Xu. 2021. "Ideology and performance in public organizations."Technical report, National Bureau of Economic Research.

Su, Puyu, Jiahu Hao, Zhaohui Huang, Xiao Limin, and Fangbiao Tao. 2011. "A study on intimate partner violence based on 2575 college students (in Chinese)." *Chinese Journal of Epidemiology* 32 (4): 346–351.

Tao, Ran, Fubing Su, Xin Sun, and Xi Lu. 2011. "Political trust as rational belief: Evidence from Chinese village elections." *Journal of Comparative Economics* 39 (1): 108–121.

Tirole, Jean. 1986. "Hierarchies and bureaucracies: On the role of collusion in organizations." *JL Econ. & Org.* 2 181.

To, Sandy. 2015. *China's Leftover Women: Late marriage among professional women and its consequences*. Routledge.

Van Bavel, Jay J, Aleksandra Cichocka, Valerio Capraro et al. 2022. "National identity predicts public health support during a global pandemic." *Nature Communications* 13 (1): 1–14.

Walder, Andrew G. 2014. "Rebellion and repression in China, 1966–1971." *Social Science History* 38 (3-4): 513–539.

Wang, Erik H. 2022. "Frightened mandarins: the adverse effects of fighting corruption on local bureaucracy." *Comparative Political Studies* 55 (11): 1807–1843.

Wang, Naixia. 2011. "Adult higher education in China: problems and potential." In *Higher Education Reform in China*, 44–61, Routledge.

Wang, Qinghua. 2014. "Crisis management, regime survival and "guerrilla-style" policy-making: the June 1999 decision to radically expand higher education in China." *The China Journal* (71): 132–152.

Wang, Shumei, and Likang Wang. 2000. "A study on health education about the HIV based on 2974 college students in Shanghai (in Chinese)." *Chinese Journal of School Doctor* 14 (1): 20–21.

Wang, Xiaonan. 2022. *The Politics of Appointing Insiders and Outsiders: Appointing Provincial Government Agency Heads in China*. Ph.D. dissertation.

Wang, Yuhua, and Bruce J Dickson. 2022. "How corruption investigations undermine regime support: Evidence from China." *Political Science Research and Methods* 10 (1): 33–48.

Weber, Eugen. 1976. *Peasants into Frenchmen: the modernization of rural France, 1870-1914*. Stanford University Press.

Wederman, Andrew. 2004. "The intensification of corruption in China." *The China Quarterly* 180 895–921.

Wei, Shang-Jin, and Xiaobo Zhang. 2011. "The competitive saving motive: Evidence from rising sex ratios and savings rates in China." *Journal of political Economy* 119 (3): 511–564.

World Bank. 2017. *World development report 2018: Learning to realize education's promise*. The World Bank.

Wu, Daguang, Manman Teng, and Duanmiao Li. 2016. "College Graduation Rate and the Quality of Higher Education: A Comparative Study Based on China and the US (in Chinese)." *Journa of Higher Education* 12.

Xi, Jinping. 2015. "Full transcript: Interview with Chinese President Xi Jinping." *Wall Street Journal* 22.

Xu, Gang, Xue Wang, Ruiting Wang, Go Yano, and Rong Zou. 2021. "Anti-corruption, safety compliance and coal mine deaths: Evidence from China." *Journal of Economic Behavior & Organization* 188 458–488.

Xu, Gang, and Go Yano. 2017a. "How does anti-corruption affect corporate innovation? Evidence from recent anti-corruption efforts in China." *Journal of Comparative Economics* 45 (3): 498–519.

Xu, Gang, and Go Yano. 2017b. “How does anti-corruption affect corporate innovation? Evidence from recent anti-corruption efforts in China.” *Journal of Comparative Economics* 45 (3): 498–519.

Yang, Rui. 2005. “Corruption in china’s higher education: a malignant tumor.” *International Higher Education* (39): .

Ytsma, Erina. 2019. “Effort and Selection Effects of Performance Pay in Knowledge Creation.”

Zha, Danyan. 2022. “Schooling Expansion and the Female Marriage Age: Evidence from Indonesia.”

Zhang, Hanzhe. 2021. “An investment-and-marriage model with differential fecundity: On the college gender gap.” *Journal of Political Economy* 129 (5): 1464–1486.

Zhong, Hai. 2011. “Returns to higher education in China: What is the role of college quality?” *China Economic Review* 22 (2): 260–275.

Zhuang, Maiting. 2022. “Intergovernmental Conflict and Censorship: Evidence from China’s Anti-Corruption Campaign.” *Journal of the European Economic Association*. 10.1093/jeea/jvac036, jvac036.

Zmerli, Sonja. 2014. *Political Trust*. 4887–4889, Dordrecht: Springer Netherlands, . 10.1007/978-94-007-0753-5_2202.

Zuckerman, Harriet. 2020. “Is “the time ripe” for quantitative research on misconduct in science?” *Quantitative Science Studies* 1 (3): 945–958.