Political Pressure and the Direction of Research: Evidence from China's Academia

Daron Acemoglu
David Y. Yang
Jie Zhou*

June 25, 2022

Abstract

Can China stimulate and sustain innovation with its juxtaposition of top-down emphasis on innovation and the presence of powerful leaders within academic institutions? In this paper, we investigate whether powerful actors curtail academic autonomy and freedom, and impact the direction and quality of innovation. We collect comprehensive data on the scientific publications of researchers in the leading 109 Chinese universities and the leadership changes in these universities. We use NLP methods to measure the similarity between faculty members' and their leaders' research portfolios. We find that immediately after — and not before — the leaders take office, faculty members begin to shift their research direction towards that of their leaders. Such shifts cannot be explained by the signaling of star researchers' activities, but can be attributed to leaders' political power over faculty members' career trajectories. Leaders appointed by the Communist Party exert greater influence on faculty members' research directions, and leaders' influence is stronger among disciplines and institutions that have historically or recently experienced academic persecution. We also document significant costs of leaders' influence on research quality. Below-median productivity leaders lead to even greater increases in similarity, and switches from above-median to below-median leaders is associated with sizable declines in citations. Such decline is driven by citations to papers that are most similar to new leaders.

*Acemoglu: Massachusetts Institute of Technology and NBER; email: daron@mit.edu. Yang: Harvard University and NBER; email: davidyang@fas.harvard.edu. Zhou: Massachusetts Institute of Technology; email: jiezhou@mit.edu. Junxi Liu, Michelle Han, Saier Wu, Peilin Yang, Andy Zhao, Hanyao Zhang, Xiaoyang Zhang, and Xinrui Zhou provided excellent research assistance.

1 Introduction

Fostering innovation has been a central aim of the modern Chinese regime, which recognizes the importance of technological progress for continued economic growth and has substantially increased funding for academic research during the past decades.[1] This desire notwithstanding, Chinese universities have enjoyed only limited autonomy, as exemplified by the fact that each department in every university has a Communist Party representative in charge and these Party representatives have considerable power over academic personnel decisions.

While the freedom to experiment and pursue new and sometimes unconventional ideas is considered as being fundamental to innovation (Hayek, 1978; Ridley, 2020),[2] can China stimulate and sustain innovation with its juxtaposition of top-down emphasis on innovation and the presence of powerful leaders within academic institutions? Do powerful actors curtail academic autonomy and freedom, and impact the direction and quality of innovation?

In this paper, we investigate these questions by estimating the effects of new academic leaders (e.g., deans, department heads or Communist Party representatives) on the type of research conducted by impacted faculty members across a large number of disciplines in Chinese universities. These leaders have considerable power over the careers of researchers in China, since they allocate resources, decide promotions, and can terminate faculty without meritocratic review. Do career concerns imposed by the academic leaders motivate faculty members to change their research direction to be more similar to their leaders' work? If so, does such change in research direction affect research quality? Finally, to the extent that many leaders have both political and academic objectives since they are typically appointed by the Communist Party, how much of the leaders' influence on faculty members can be attributed to politically-driven factors?

Our exploration of the linkages between political pressure and academic research is built on three new datasets. First, we collect information on organizational structures

[1] According to China's National Bureau of Statistics, fiscal spending on basic research rose from 25.8 billion RMB in 2013 to 42.3 billion RMB in 2018, a 64% increase; source: `http://www.stats.gov.cn/tjsj/ndsj/2019/indexeh.htm`.

[2] In reality, however, ample examples exist where powerful actors in political regimes, academic organizations, and commercial corporations exert influence over innovation and research activities. Such influence could take many forms: ranging from discouraging certain lines of inquiry, to explicitly blocking certain innovation or even persecuting innovators, to imposing and elevating political and non-economic considerations ahead of innovation potentials. Prominent examples of political interference include the story of the Roman Emperor Tiberius killing an inventor who had invented unbreakable glass in order to suppress his innovation, and Queen Elizabeth I of England blocking William Lee's "stocking frame" knitting machine (Acemoglu and Robinson, 2012).

and leadership transitions in the social science departments of Chinese universities since 1990. Second, we collect (close to) the universe of research publications — in both Chinese and English language journals — by faculty members and department leaders during this period. We use natural language processing (NLP) methods to construct measures of research similarity between faculty members and their department leaders. Third, we collect citation information for this sample of research publications. Using these measures of research similarity and quality, we investigate the impact on the direction and quality of faculty members' research in the event of leadership shifts.

Our baseline empirical strategy compares the (topical and textual)similarity in research output between faculty members and their respective leaders, before and after leadership switches. Identification with this strategy relies on faculty-leader pair level variations over time, exploiting the fact that each researcher is observed under multiple leaders over the course of her academic career, and leadership transitions take place at different times and with different frequencies across departments. This setup allows us to include a full set of faculty-leader pair fixed effects as well as calendar year fixed effects — thus accounting for other factors associated with the similarity between faculty members' and department leaders' research portfolios, and isolating sharp changes in faculty research direction driven by the appointment of specific leaders.

We find that, after new leaders take office, faculty members' direction of research shifts towards the portfolio of the current leader. The similarity in research output between faculty members and their respective leaders significantly and substantially increases: such change emerges almost immediately after the leadership transition, and it persists for at least four years into a new leader's tenure. Bolstering confidence in our identification strategy, we do not find analogous changes in similarity *before* a leader takes up her position. This shift in research direction after a leadership transition is robust to a range of alternative empirical specifications, including different measures of research similarity and different controls.[3] Reassuringly, in a falsification exercise, we do not find any increase in research similarity between faculty members and new leaders in *other* disciplines in the same university. Moreover, the results are not driven by differential publication strategies in Chinese and English language journals: we find very similar results when analyzing English publications alone, or when we combine English and Chinese publications together.

[3]Academic leaders can have an equally defining effect on what types of faculty they hire. We show that this channel is also important in China by documenting that departments start hiring new faculty that are more similar to new leaders. Moreover, these newly-hired faculty shift their research even more in the direction of their leader's portfolio. This hiring channel thus highlights the potentially persistent effects of leaders on the research trajectory of the departments under their control.

Shifts in research activities after a leadership transition could be driven by two related but distinct channels. First, the appointment of a new academic leader may signal to researchers promising research areas or research directions favored by the regime, thus inducing a change in research direction among faculty members. Second, and more directly related to our focus, researchers may attempt to pander to new leaders who have direct power over them. To distinguish between these two mechanisms, we estimate the effects of appointments of new leaders in the same discipline but in other institutions, who thus have no direct power over a faculty member's career, but may still signal changing priorities in Chinese academia. We also repeat this exercise for new leaders in other top institutions (and still in the same discipline) where such signaling may be particularly salient. We find that same-discipline leaders from lower-ranked universities have no significant impact on faculty members' research direction, while same-discipline leaders from higher-ranked universities have positive but much smaller effects than the impact of one's own leader. Finally, we estimate the effects of department colleagues' receiving prestigious research award, as such recognition could signal research promises. We find little effects of public recognition on colleagues' research direction. These patterns suggest that a small component of the leadership transition effects can be explained by new leaders' signaling of new research directions, and the bulk of the effects is driven by the leaders' influence through their direct control over researchers' career.

How much of the leadership influence on research direction can be attributed to career concerns arisen from political pressures? First, compared to department chairs, we find that the department's Communist Party representatives exert stronger influence on faculty's research. Tellingly, this is despite the fact that party representatives are rarely leading academics; they are directly appointed by the Chinese Communist Party and represent closer alignment to the intention and direction of the party. Second, we decompose the sources of convergence in research output into scientific and political language. We separately train text similarity models using word embeddings from the Chinese Wikipedia (scientific semantics) and the *People's Daily*, the Communist Party's official newspaper (political semantics). We observe convergence in research similarities in political language, controlling for the convergence in scientific language, suggesting a political source of leadership transmission on researchers' activities. Third, we examine whether the leadership effects become stronger after major incidences of researchers removed explicitly due to their political stances.[4] We find that after the occurrence of these incidents

[4]These incidents occur after 2013, reflecting the further decline of academic freedom since President Xi Jinping's accession to power in November 2012. This is best illustrated by the decisions of many leading universities to remove clauses related to the "freedom of thought" from their charters; source: `https://reut.rs/39RVCNx`.

in a given university, leadership transitions in the same institution (across all disciplines) exhibits substantially stronger effects on researchers' output. Fourth, we explore whether academic persecution in the past may have persistently shaped the present-day academic environment and influenced the extent to which leaders are able to influence research directions. Following Wang and Kung (2021), we measure each academic discipline's likelihood of facing top-down persecution during the Cultural Revolution. We find that disciplines that suffered worse persecution four decades ago exhibit greater impacts on research direction from leadership switches. We interpret this finding as indicating that political career concerns matter more in disciplines that suffered greater (political) persecution in the past. As such, our findings suggest that there exists effects of political interference in academic research resulting from a highly persistent legacy of past repression on academia. Furthermore, the specific mechanism through which these effects are realized is through amplified political career concerns.

Finally, we establish that shifts in research direction due to leadership transitions impose significant costs on research quality. We do this by comparing the effects of above-average and below-average leaders (defined in terms of research productivity) on the citations of researchers after leadership switches. Our results show significant deterioration in citations after switching from an above-average to a below-average leader and significant improvements after switching from a below-average to an above-average leader. Switches between leaders of the same quality have no impact on citations. More tellingly, we find that the change in citations received by a researcher after a leadership switch is driven entirely by their papers that have the greatest similarity to their new leaders' research, and that there is no change in citations received by papers that are minimally related to a new leader's agenda. Given that leaders of lower productivity exert greater influence on faculty members' research portfolios, together these results indicate that the change in research direction induced by political pressures in China have significant costs in terms of research quality. We also find some evidence that the negative effects of leadership switches on research quality are cumulative and become larger after a faculty member experiences several leadership switches.

Taken together, these results suggest that political factors have a major impact on the direction and quality of academic research in China. Political pressure induces Chinese scholars to align their personal research agenda more closely to that of their leaders, frequently resulting in lower-quality research output. Moreover, because China now accounts for a significant fraction of the world's research and innovation output, potential distortions in Chinese academia are likely to have significant consequences for global innovation.

Our paper is most closely related to the branch of existing literature on political economy investigating linkages between political factors and innovation. Much of the emphasis in this literature has been on the risk of expropriation or political interference on entrepreneurship, investment and innovation (e.g., North et al., 2009; Acemoglu and Robinson, 2012). Potential future political threats from successful entrepreneurs may also encourage elites to block innovation to preserve their political power and rents (Acemoglu and Robinson, 2006).[5] Our mechanism is rather different as it shows the effects of local political pressure in academia — though the origin of this political pressure likely comes from national institutions. In so doing, our work also joins the recent studies that examines the impact of academic persecutions, ranging from the Nazi Germany's persecution of Jewish researchers (e.g., Waldinger, 2012; Becker et al., 2021), to persecution of scholars during China's Cultural Revolution (Wang and Kung, 2021). Our focus on academic research and innovation also connects our work to the growing literature on innovation economics, specifically, the various incentives that affect research production (e.g., Azoulay et al., 2011; Manso, 2011; Akcigit et al., 2018; Hill and Stein, 2021).

Our work also contributes to the literature on innovation in China. A large literature studies the organization of the Chinese economy and the factors that drive its economic growth over the past four decades. Recent works have carefully described the innovation landscape in China (e.g., Wei et al., 2017; Bombardini et al., 2018) and its potential implications for academic research (e.g., Freeman and Huang, 2015). More closely related to our paper is the innovative paper by Jia et al. (2019), who document that academic leaders in economics departments in China's top universities tend to become more productive through co-authorship after they become leaders. This pattern — political power being used by academic leaders for their own benefit — suggests a different type of political distortion and is thus complementary to our research. We contribute to the understanding of how politics interact with innovation by providing, to the best of our knowledge, the first systematic analysis of the effects of powerful, political actors on the direction and quality of academic research.

The rest of the paper is organized as follows. Section 2 describes the data sources used for this project. Section 3 describes the key measure of research similarity and the empirical strategy; Section 4 presents the results of leadership transition on career incentives. Section 5 presents evidence of that the research similarity effects we estimate are a conse-

[5] A nascent literature documents the alignment between the autocratic institutions and private innovation, particularly in the context of China. For example, Bai et al. (2020) examine how crony capitalism combined with local governments' competition can foster growth; Beraja et al. (2021) study how provision of government data and the state's demand for AI for surveillance purposes can promote private innovation in the AI sector, due to the shareability of government data across multiple purposes.

quence of politically-motivated career concerns. Section 6 assesses the costs on research quality associated with politically-charged career incentives. Section 7 concludes, while the Appendix contains additional results and more information on data construction.

2 Data

Our empirical analysis combines three primary datasets that we collect from scratch: *(i)* the structure of Chinese universities and the leadership information in the university departments; *(ii)* the scientific publications of all affiliates in these institutions; and *(iii)* the number of citations for each of the publications. We now describe each of these datasets in turn. Figure 1 plots the overview of the data construction procedures.

2.1 University structure and department leadership

We first construct a dataset tracking the organizational structure and leadership changes in Mainland Chinese universities. We begin by examining all social science departments among the top universities in China. We focus on the 109 universities that belong to "Project 985" and "Project 211," two higher education ranking schemes that unambiguously list the top academic institutions in China.[6] Out of a total of 2,914 universities in Mainland China, these 109 top universities capture 70% of all research funding, and more than 50% of major scientific publications (Zhu, 2009; Zong and Zhang, 2019).

For each university, we collect data on organizational structure for departments within social science disciplines. We focus on the organization structure one level beneath the university's top administrative hirearchy, which corresponds to "departments" in some universities and "schools" in a few others. We will refer to these as "departments" for the rest of the paper for brevity. This is the level at which leadership has the most direct control over resources, promotions and hiring decisions. We standardize the organizational structure to make the school level definition comparable across universities.

We focus on schools and departments that are continuously active between 1990 and 2019, which is also the time window for which we collect publication records. For the schools and departments that cease to exist either due to splits or mergers, we track these changes and link schools and departments together, so that past research activities in previous academic units can be appropriately attributed to the corresponding departments

[6] "Project 985" and "Project 211" are two major projects undertook by the Chinese government to promote the development and reputation of the Chinese higher education system by founding world-class universities. The universities included in these projects are top ranked in China, and many of them have since then ranked among the top 500 universities globally; source: `https://bit.ly/3ibF8Uo`.

today. This ensures that we don't have changes in leadership that are caused mechanically by changes in school structure. Overall, there are on average 7.8 schools in a given university in the period between 1990 and 2019.

Broadly speaking, we group various schools and departments into a total of 11 disciplines: economics, management, business, and finance; political science and public management; law; education; literature and media; history; psychology; philosophy, anthropology, ethnology, and sociology; regional studies; foreign language; and Marxism. For the schools and departments that are interdisciplinary in nature, we classify them into 11 disciplines by categorizing disciplines within a school as children and the school as the parent; we then group parents into a single classification if they shared connected components. The details are described in Appendix A.

Finally, we identify school or department leaders during our sample period spanning from 1990 to 2019 from a variety of sources: official websites of universities, university yearbooks, *Baidu Baike* (a Chinese-language collaborative online encyclopedia), and various online reports that mention school leadership. We manually extract the department chairs and party secretaries for each school. For the years that we cannot locate precise leadership information, we employ several interpolation methods. [7] On average, each school experiences 2.8 leadership transitions during the three decades between 1990 and 2019. The average tenure of a given school chair is thus 5.8 years, though this varies fairly substantially across disciplines: ranging from 4.8 years in the discipline of Marxism, to 6.3 in the discipline of foreign language.

Similar to the bureaucratic structure in many organizations in China, universities and the schools within them have two parallel leadership posts: school chair in the academic track and Chinese Communist Party secretary in the political track. We primarily focus on the leadership in the academic track since those individuals are scholars and have records of academic publications, making it relevant to study the potential re-pivoting of research effort by faculty members. In contrast, party secretaries often have no academic background and are rotated in from other Communist Party organs. Nevertheless, we are able to identify a subset of school party secretaries who have academic track records and compare their influences on faculty members to the influences of academic leaders.

[7] Normally, if we find a faculty member appeared as a leader in the news in Year_1 and Year_2, we assume that this leader was holding this position from Year_1 to Year_2 if leadership information for the years in between are missing. When it is different leaders before and after the missing cell, if the missing years are no more than three, we conjecture that one of the leaders before and after the missing cell was still a leader, and we interpolate by assigning the past leader to the missing cell. This is assuming that there may be less information about leaders that are about to step down, but for the new leader who just began a position, it is more likely to obtain information about her. 16% of the missing department chairs are be solved under this assumption.

2.2 Research publications

We construct a dataset of all scientific publications by scholars in the corresponding institutions during the three decades between 1990 and 2019. The scientific publication dataset serves two primary purposes: firstly, it provides a description of research output of researchers, which we rely on to construct our primary outcomes of interest (described in detail in Section 3.1). Moreover, this publication dataset allows us to retrospectively construct a roster of scholars affiliated with each institution and school, since administrative records of faculty rosters are incomplete or absent for most schools in most years during the previous decades.

To construct this dataset, we rely on two major sources. The first is *China National Knowledge Infrastructure (CNKI)*, a full-text database covering 90% of all official published Chinese journals. The second is *Wanfang Data*, which is a comprehensive database of Chinese journals, dissertations, and academic conferences. It provides access to 8,183 journals published in China and over 43.17 million articles, including 42.89 million full-text records (as of May 2019). To the extent that the coverage of these two databases do not fully overlap, they complement each other and when combined together, provide us with close to full coverage of scientific publications in Chinese journals.

In addition to Chinese-language publications, we also collect information on scholars' publications in English-language journals. For each scholar we already collected Chinese-language publications, we searched her publications, compiled in the CNKI, from 1990 and 2019 and restrict the language of publications to be English. We verify and consolidate authors' information and consolidate duplicate English spellings of different Chinese names. In total, we found 14,105 scholars (31%) with 97,031 English papers. Appendix D provides details on cross-checks of data completeness based on information collected in Google Scholar.

For each researcher affiliated with the universities of interests described in Section 2.1, we collect all the papers she publishes between 1990 and 2019. We exclude publications on non-academic outlets such as newspaper opinion pieces. We also exclude dissertations (e.g., part of the graduate studies) and other internal school journals. This amounts to a total of 5,290,503 papers. For each paper in the collection $\mathcal{D}$, we collect information on the paper's title, authors, publication year, abstract, and citation countss.

We then use the publication dataset to extract rosters of faculty members (and those who ever served as school or department leaders) in each academic unit. In a nutshell, we assign an academic affiliation to each author of a paper based on publication information. Because not all papers have information on affiliations at the school level, we assign the

school level affiliation from any publication of a given author to all of her papers. In order to rule out individuals who are affiliated with a school as a student (and hence publishing sparsely) rather than a faculty member, we use the dissertation database to locate the graduation year of a given researcher and consider the post-graduation period as their faculty affiliation. We also restrict faculty members to those who publish more than 5 papers under a given affiliation and has publication records for more than 3 years, further excluding ones that may publish during a temporary position such as while a visiting scholar. Our faculty roster extraction procedure performs well when we validate it with a set of contemporaneous faculty lists that we can obtain from the school's official website (see Appendix B for details).

This procedure provides us with a list of faculty members affiliated with a particular school s, at a university u, in year t. Overall, we identify 42,395 active faculty members in social science disciplines in top universities between 1990 and 2019. On average, there are 62.2 faculty members in each school, ranging from 15.3 in the discipline of regional studies, to 103.0 in the discipline (category) of management, economics, finance, and business. Each faculty member publishes on average 1.3 papers in any given year, ranging from 0.8 paper per year in the discipline of foreign language, to 2.4 papers per year in the discipline of psychology.

We notice a general trend of increased publication by scholars across all disciplines over the sampling period: the research productivity grows from 102 papers per year in 1991 to about 32,428 papers per year in 2018, reflecting the overall growth of Chinese academic institutions and research capacity over this period. We include year fixed effects in all baseline specifications to account for the secular trend in research activities.

2.3 Citation counts

In order to measure research quality, we also collect data on the citation counts for all research publications in our sample.

Data on citation counts are constructed separately from the publication dataset previously described. From the same CNKI and Wanfang databases, citation counts data is collected and matched back to each paper in the publication dataset based on the paper's title and publication year. When CNKI and Wanfang databases report different citation counts, we take the higher count of the two as the paper's overall citation counts. This process yields citation counts for 95.6% of the papers included in analysis. The remaining papers with missing citation counts are dropped from the dataset.

On average, each paper has a mean of 13.9 citations, with leaders averaging slightly

more at 18.2 citations per paper, and faculty averaging slightly less at 13.7 per paper (see Appendix Table A.2 for summary statistics). Faculty at higher-ranked universities also average higher citation counts: at top-ranked Peking University, mean citation counts are almost double at 24 citations per paper.

3 Empirical strategy

In this section, we present our empirical strategy. The first step is the construction of our research similarity measures, which are described in the next subsection. We discuss our empirical design in in Section 3.2.

3.1 Measurement of research similarity

3.1.1 Similarity of paper pairs

To measure similarity between any given two papers, we construct similarity scores between pairs of scientific publications. For each pair of papers in the paper collection $\mathcal{D}$, we construct a variety of measures of similarity, each of which can be viewed as a mapping from pairs of research papers into a similarity score — $s : \mathcal{D} \times \mathcal{D} \mapsto \mathbb{R}^+$.

There are two broad classes of similarity measures: *(i)* non-parametric methods; and *(ii)* methods based on machine learning. Our baseline estimation uses the "term frequency inverse document frequency" (TF-IDF), which is a non-parametric method, to measure similarity. We also use several alternative text-similarity measures, which we describe in turn.[8]

Term frequency inverse document frequency (TF-IDF) TF-IDF is a statistical measure commonly used to evaluate how important a word is to a document in the context of a given corpus of documents (Biasi and Ma, 2020; Kelly et al., 2018). The importance score increases proportionally to the number of times a word appears in the document, but is offset by the frequency of the word in the corpus. Mathematically, for a word i, its TF-IDF score is:

$$\text{TF-IDF}(i, d, \mathcal{D}) = \text{tf}(i, d) \times \text{idf}(i, \mathcal{D}), \tag{1}$$

where $\text{tf}(i, d)$ is the frequency of word i in document d, and $\text{idf}(i, \mathcal{D})$ is the log of one over the share of documents containing i in the corpus $\mathcal{D}$.

[8] In this preliminary draft, we focus on results based on TF-IDF. We will introduce results using other similarity measures in a subsequent draft.

The collection of publications forms our text corpus $\mathcal{D}$, elements of which are individual papers represented by their abstracts. By adopting the bag-of-words model, each document d can be represented as a vector $\mathbf{v}^d$ based on its words, discarding specific grammar and word order. The length of $\mathbf{v}^d$ is equal to the number of words in the vocabulary of the corpus $\mathcal{D}$.[9]

Intuitively, we can let each element v_i^d be the number of times word i occurs in document d. Simply calculating the distance between the vectors of word frequencies to measure the similarity is problematic, however, because words that occur commonly in every document (often called "stop words") will introduce bias in the similarity score.

With TF-IDF, we are able to map a document d to a vector $\mathbf{v}^d$ in which each element $v_i^d = \text{TF-IDF}(i, d, \mathcal{D})$. Then for two document $f, l \in \mathcal{D}$, the similarity measure is defined as:

$$s(f, l) = \mathbf{v}^f \cdot \mathbf{v}^l. \tag{2}$$

Doc2Vec Doc2Vec (Dai et al., 2015) is an unsupervised neural network algorithm that learns the fixed-length feature vectors from variable-length documents. Doc2Vec predicts each word in document d by utilizing the current document and a sliding window around the target word. We train a Doc2Vec model by using titles and abstracts of papers in paper collection $\mathcal{D}$ as the training set. After the training stage, each document $d \in \mathcal{D}$ is mapped to a document-unique feature vector $\mathbf{v}^d$, which represents the "concept" of the document.

Training a neural network like Doc2Vec requires a large train set and computation power. Rather than directly training this neural network, however, we can make use of the training results of other researchers and institutions, based on large corpus such as Wikipedia and newspapers. This allow us to use their results as pre-trained models and further fine-tune models to fit our paper collection $\mathcal{D}$. In our setting, we utilized two pre-trained Chinese word embeddings provided by Li et al. (2018): *(i)* Word2Vec trained with People's Daily News[10], and *(ii)* Word2Vec trained with Chinese Wikipedia. More details about the training process are in Appendix C.

After constructing the vector representation of each document, similar to TF-IDF, we take *cosine* distance to measure the similarity between two documents $f, l \in \mathcal{D}$: $s(f, l) = \mathbf{v}^f \cdot \mathbf{v}^l$.

[9] With a slight abuse of notation, here $\mathcal{D}$ refers to a structured set of texts. Each text in this set is the abstract of a paper.

[10] People's Daily is an official newspaper of the Chinese Communist Party and the largest newspaper group in China.

3.1.2 Similarity score for a faculty-leader pair

Based on the similarity score between pairs of papers, we construct measures of similarity in research portfolio for each pair of faculty member and department leader.

Specifically, let $\mathcal{F}$ be the population of faculty members and $\mathcal{L}$ be the set of all leaders. For each faculty-leader pair $(F, L) \in \mathcal{F} \times \mathcal{L}$, we denote the set of papers published by the faculty member F in year t as $\mathcal{D}^F(t) = \{f_{t1}, f_{t2}, ..., f_{tn}\}$. The set of papers published by the leader $\mathcal{L}$ in year t is denoted by $\mathcal{D}^L(t) = \{l_{t1}, l_{t2}, ..., l_{tn}\}$. Finally, the similarity score of a pair of papers (f, l) is designated as $s(f, l)$.

Next, for the faculty-leader pair, we construct the research similarity score at time t based on pairs of papers belonging in the following set:

$$\mathcal{P}^{(F,L)}(t) = \{(f, l) | f \in \mathcal{D}^F(t), l \in \cup_{k \leq t} \mathcal{D}^L(k)\}.$$

Namely, we construct pairwise similarity scores by comparing all papers published by the faculty member in year t with all the papers that the leader has published up until (and including) year t.

In order to capture the pivoting of research activities for a subset of salient papers, we define the similarity score between faculty-leader pair $i = (\mathcal{F}, \mathcal{L})$ in year t as the maximum similarity score among all pairs of papers published by these two researchers during the corresponding period: $y_{i,t} = \max\{s(f, l) | (f, l) \in \mathcal{P}^{(F,L)}(t)\}$.[11]

3.2 Empirical specification

Our baseline empirical strategy follows a modified event-study design. We compare research similarity between faculty-leader pairs before and after the leader takes office. The specification controls for faculty-leader pair fixed effects as well as calendar time fixed effects. This implies that changes in research similarity are identified entirely from within faculty-leader pair variation (i.e., from variation in similarity between a faculty member and a leader over time). Our baseline specification can be written as:

$$Y_{i,t} = \sum_{l \neq -1; l = -3}^{4} \psi_l D_{i,t}^l + \alpha_i + \lambda_t + v_{i,t}, \tag{3}$$

where $Y_{i,t}$ is the similarity score for the faculty-leader pair i at time t; $D_{i,t}^l$ is an indicator for faculty-leader pair i being l periods away from initial treatment at calendar year t; α_i is a full set of faculty-leader pair fixed effects; and λ_t denotes a full set of calendar time fixed effects. For each faculty-leader pair, we focus on the time window comprising of

[11] To capture the average shifts in research portfolio, we also define a specification of the research similarity based on the average similarity scores across all pairs of papers: $y_{i,t} = \frac{1}{|\mathcal{P}^{(F,L)}(t)|} \sum_{(f,l) \in \mathcal{P}^{(F,L)}(t)} s(f, l)$.

three years before and four years after the leadership transition. Our baseline results are robust to alternative choices of time window.

By conducting the analyses at the faculty-leader pair level, we are taking advantage of the fact that academic leadership transitions are not synchronized across universities and departments. This strategy controls for any general shifts in a discipline's research priorities over time and also filters out any differences in research similarity resulting from the fact that leaders work in different areas and have different productivity levels.

3.3 Threats to identification

Our key identifying assumption is that variation in the similarity between a faculty member and a leader is orthogonal to other changes that happen at the same time as the leadership transition. Potential threats to the validity of this assumption include changes in (national) research priorities that take place at the same time as the appointment of a new leader and various types of selection determining which leaders are appointed to which departments. Our extensive fixed effects (most importantly at the faculty-leader pair level) should account for these selection-related concerns. Additionally, we bolster the plausibility of this identifying assumption in two ways. First, we examine the pretrends in research similarity between a researcher and a leader in the years leading up to the appointment of the leader. Second, we conduct a range of placebo exercises, which indicate that our results are not driven by spurious correlation between faculty and leader research styles and also enable us to distinguish the effects of new leadership appointments working via signals to all faculty within a discipline about which areas are prioritized by the Communist Party or other higher authorities.

One may also be worried that a leadership transition induces changes in faculty members' research productivity and research output quantity, thus changing the denominators of the research portfolio similarity between faculty members and leaders. We examine faculty members' productivity changes leading up to and after leadership transitions in the corresponding academic unit. Specifically, we regress the total number of academic publications per year on time relative to the date of leadership transition, controlling for faculty member fixed effects and calendar year fixed effects. This allows us to isolate the differential effects of productivity changes due to leadership transitions and shows that leadership transitions does not induce changes in researcher productivity (defined as number of publications) as depicted in Appendix Figure A.4. This pattern makes research similarity between faculty and leaders (later citation counts) more straightforward to interpret.

4 Leadership transition and direction of research

This section presents our main results on the similarity of research style between faculty members and department leaders and confirms the robustness of these results.

4.1 Baseline results

We first examine the average effects of leadership transitions on faculty members' research direction across all disciplines and all institutions over the past three decades. We pool our entire sample together and estimate the baseline specification in (A.8). Figure 2 presents the results by plotting the non-parametrically estimated ψ_l coefficients along with the corresponding 95% confidence intervals. The research portfolio similarity score between leaders and faculty members in the year *prior to* the leader taking office is normalized to zero and the timing of the leadership transition is marked by the vertical red line.

The estimates in Figure 2 show a significant increase (by approximately 7%) in research similarity between faculty members and their leaders. There is no increase in similarity before the leader takes office, and the similarity index takes off immediately after leader turnover and persists for at least four years into the new leader's tenure. This timing, with no pre-trends, is reassuring for the validity of our identification strategy. The absence of pre-trends suggests that there are no anticipation effects before new leaders take office and there is also no evidence that researchers are selected to lead departments based on the similarity of their research portfolios with the rest of the faculty members.

To the extent that faculty members' research activities and output within a department are diverse, the patterns depicted in Figure 2 suggest that, after the appointment of a new leader, faculty members pivot their heterogeneous research activities towards the same direction, getting closer to that of their leader's research portfolio. By the same token, the estimates also indicate that, after the appointment of a new leader, researchers pivot away from the research of past leaders.

4.2 Effects of leaders from other departments

Leaders may affect the research trajectories of faculty through two distinct but related channels: the local career concerns of faculty members under their direct jurisdiction, and the signals that all faculty in the discipline receive from the appointment of a leader with a particular research style and portfolio (which may also be related to career concerns since those heeding such signals may be more successful).

To isolate the effects working through the local career concerns, we estimate the baseline specification (A.8), including new leaders in the same discipline but from other universities. Because signals from lower-ranked universities may be less influential, we also distinguish between same-discipline leaders in lower-ranked and higher-ranked departments.

Specifically, to construct pairs of faculty members and leaders within the same discipline, we proceed as follows: we first locate all leaders within the same discipline at different universities, appointed within a seven year window around the leadership switch of interest. We separate the resulting pool of leaders into two groups based on whether their affiliated universities are ranked above median in the sample. From this pool of leaders, we randomly allocate a same-discipline leader to the researcher in question, separately drawing leaders from lower-ranked and higher-ranked universities. We then jointly analyze the effects of these hypothetical leaders on faculty members' research directions.

The results are presented in Figure 3. The patterns depicted in this figure are clear. There is no increase in research similarity after a leadership switch between researchers and same-discipline leaders from other higher-ranked universities, and we see significant but much smaller increases in similarity after a leadership switch between researchers and same-discipline leaders from other lower-ranked universities. These findings suggest that only a minimum component of the effects presented so far could be driven by signaling of prioritized research directions. Rather, it seems that most of our estimates reflect the political power of leaders that control resources and promotions in the focal researcher's institution.

4.3 Placebo: effects of leaders from other disciplines

We also carry out a related placebo exercise where we re-estimate our baseline specification (A.8), but for leaders in different disciplines. Once again, we focus on leaders that are appointed within a seven year window around a leadership switch and randomly allocate them to a faculty member experiencing a leadership switch as a hypothetical leader. The results of this exercise are presented in Figure 4, which again includes our baseline estimates from Figure 2 for reference. This placebo exercise shows no significant increase in similarity in research portfolios between faculty members and placebo leaders, and the point estimates are quantitatively much smaller than our baseline estimates. Overall, this placebo exercise bolsters our interpretation that the increase in research similarity detected in Figure 2 is not driven by spurious factors and reflects the causal effect of a leader switch.

4.4 Political influence and hiring decisions

Having documented that political interference shifts the research direction of faculty members towards their leaders' style, we next examine whether part of such change is driven by changes in the composition of faculty members. In particular, leaders in Chinese academia have a major impact on who is hired and this might be one channel through which changes in research similarity between faculty and leaders transpire (though we note that this could not explain all of our results, since they are exploiting within faculty-leader pair variation).

To shed light on this question, we examine whether leaders tend to hire faculty members whose research portfolio is closer to their own. Focusing on all faculty members who begin affiliation with a particular department or school between 1990 and 2019, we construct average research similarity scores between these faculty members and all leaders in the corresponding departments. The leaders who are in a leadership position at the time a faculty members' affiliation begins are the "hiring leader," under whose leadership (or potentially direct influence) the specific hiring decision is made. We investigate whether research similarity between new hires and department leaders are particularly high when the leader in question is also the hiring leader.

Table 1 presents results of the regression estimates of the research similarity between new faculty members and the leaders who hired them, relative to that between new faculty members and non-hiring leaders. We include a full set of faculty member fixed effects, which means that we are only exploiting variation across leaders. The estimates indicate that the research similarity between faculty members and hiring leaders are higher than with other leaders in the same department. This indicates that leaders tend to hire new faculty members whose research portfolios are similar to theirs.[12]

The final question we investigate in this section is: conditional on already having a similar research direction with their hiring leaders, do newly hired faculty members further shift their research portfolio towards their leader after being hired? To examine this question, we re-estimate our baseline specification on the sub-samples of existing faculty members and newly-hired members around a leadership switch. The results, which are presented in Figure 5, show a pronounced impact on newly-hired faculty. As in our results so far, the entire department shifts towards the new leader's research portfolio, but in particular, the pivot of newly-hired faculty in this direction is even larger. This pattern is consistent with politically-powerful leaders having an even more defining influence on

[12]Leaders distorting the hiring decisions of their departments towards their own priorities is another dimension of misallocation created by (potentially) top-down excessive controls, which we will explore in greater detail in Section **??**.

researchers who start their employment under their rule.

Combined with the results presented so far, this pattern suggests that new leaders tend to have a substantial impact on the research direction of the departments they control, both by changing the research direction of existing faculty and by hiring new researchers more aligned with their research style or priorities.

4.5 Robustness

Our baseline measure computed using TF-IDF may underestimate the similarity between two documents. This potential underestimation is rooted in two assumptions that this methodology utilizes: *(i)* the meaning of the words are orthogonal to each other, and *(ii)* the word order is not taken into account when measuring the similarity between two texts.

To relax the two assumptions, we calculate a vector $\mathbf{v}^d$ for each document d by training a Doc2Vec model (Dai et al., 2015), as described in Section 3.1. To further show that our results are robust to similarity scores measured under different contexts, we construct two different Doc2Vec similarity scores by combining our paper collection with Chinese Wikipedia and People's Daily respectively. Figure A.5 shows the impact of leader switches on the two types of Doc2Vec similarity scores. Consistent with the TF-IDF similarity score, before the new leader takes office, there is no rise in similarity, but once the new leader is in charge, the similarity index increases significantly and we see this effect lasting for at least four years.

Moreover, Figure A.13 shows the robustness of our results to including just leader and faculty fixed effects (rather than leader-faculty pair fixed effects), which is useful, since our analysis of political career concerns on research quality will not include leader-faculty pair fixed effects. The results are nearly identical to those in Figure 2.

Finally, we show that our results are robust when including English publications. We construct TF-IDF similarity score between Chinese and English publications following the methods illustrated in Appendix D and re-estimate the baseline specification. The yellow points/lines in Figure A.14 shows the impact of leader switches on both Chinese and English publications, which are very close to the baseline leader effects (in grey).

4.6 Heterogeneous effects

Anticipating our results on the implications of faculty members' career concerns on their research quality, we now show that the increase in similarity to current leaders' research

is greater in lower-ranked departments and for leaders who are academically less accomplished.

First, we re-estimate our baseline specification, but allow the leadership transition effects to differ across three sub-samples, divided according to the department's research output relative to other departments in the same discipline — those in the top 10% of this relative ranking, those in the range 10%-40%, and finally those in 40-70%.[13] Note that even the relatively lowered-ranked departments in our sample are among the top academic institutions in China, since we are focusing on the top 109 universities in the country. Figure 6 presents the results. While we observe greater research similarity between faculty members and leaders after the leadership transition across all ranking groups, the effects are noticeably larger for schools ranked below the 50th percentile and smallest for the top 10% group. Reassuringly, there are no statistically significant pre-trends for any of the three groups. That our effects should be stronger in lower-ranked departments is plausible. Reflecting the top-down nature of Chinese academic institutions, leaders tend to have substantial power over promotion and dismissal decisions, as we noted in the Introduction, and this power is even greater in lower-ranked departments.

Second, we separate leaders between those that are above-median and below-median based on their publication record. Figure 7 shows that the increase in similarity is larger for below-average leaders.

Third, we compare the effect of star researchers on research similarity to the baseline effect of department heads. We define star researchers as researchers under Changjiang Scholar Program from 1999 to 2018. And we use the year when they earn the award to determine the time of treatment and considered it as a signal of promising researcher areas. Then we follow the same procedure as department heads to investigate how researchers respond before and after the award winning. Figure 8 shows the results. Contrary to department leaders, however, the similarity between star researchers and their colleagues decreases after the award winning.

Taken together, these three findings imply that the influence of leaders on the research direction of faculty is driven more by leaders in charge of lower-ranked universities and that have below-average academic achievements, both of which suggest that this change in research direction may be associated with significant distortions and might even lead to substantively lower-quality research. This is what we will explore more systematically in Section 6.

[13] These are rankings according to China's Ministry of Education, including all Chinese universities, and the bottom 30% departments are not in our sample.

5 Politically-charged career incentives

There is politics and career concerns in every academic institution. Is academia in China different? In the next subsection, we undertake several complementary exercises to argue that the answer is likely yes. We then explore the long term effects of past academic persecution on politically-charged career incentives.

5.1 Academic leadership from the Communist Party

First, we examine whether the Communist Party representatives in academic departments exert as strong, or even stronger, influence on the faculty members' direction of research as academic leaders. Similar to the bureaucratic structure of many organizations in China, two parallel leadership posts co-exist in each department: academic leaders (who are department heads or deans) and the Communist Party representative or secretary. Party secretaries often have little academic background and are directly appointed within the party organization. Typically they are rotated from other Communist Party organs. For our exercise, we focus on the subset of party secretaries who have academic track records but have embarked on a political track within academic leadership.

We re-estimate our baseline specification to explore the effects of party secretaries on the research direction of faculty members in their departments. Figure 9, Panel A, presents the estimates. Reassuringly, there are no significant pre-trends prior to the party secretary taking office, but the research similarity between faculty members and their party secretaries starts increasing immediately thereafter. Although the estimates are noisy (which is inevitable given the smaller sample size of academically-active party secretaries), their magnitude is, on average, about three times the magnitude of the effect for academic leaders. This pattern confirms the political nature of the career concerns we have documented in the previous section.

Next, we examine whether attempts to introduce a type of tenure-track system in some Chinese universities changes the underlying impact of leadership transition on faculty members' research directions. Although the aim of the reform was to make Chinese academic institutions more similar to their US or European counterparts by encouraging autonomy and higher-quality research, these reforms may have also increased the power of leaders, who were enabled to make more decisions about the careers of the faculty under them.[14] While several departments (such as the Department of Physics at Tsinghua

[14] Albeit an extreme example, a recent murder case of a department Party Secretary in Fudan University due to his denial of tenure to a faculty member demonstrates the controls leaders can exert during the tenure promotion process (see `https://bit.ly/3qq3hKG` for further details).

University) introduced tenure track system in the early 2000s, the vast majority of departments in the universities that we examine formally began tenure track reform on or after 2013. Specifically for 20 social science departments at Peking University, we collected the starting date of the tenure-track system. Because the tenure-track reform only applied to new faculty members hired after the start of the reform, we classify faculty-leader pairs as falling under the new tenure-track system if the faculty in question was hired after the reform. These restrictions yield a sample of 5,770 faculty-leader pairs, of which 911 were under tenure-track. We then jointly estimate the effect of a leader switch on similarity scores for faculty-leader pairs that are either unaffected or affected by the reform. The results, presented in Appendix Figure A.11, show that there is an even greater impact of leaders on research direction after tenure-track reform. This finding suggests that, as hypothesized, tenure-track reform within the institutional setup of Chinese academia has exacerbated the control of academic leaders and amplified the politically-charged career concerns of researchers.

5.2 Effects under political semantics

Word embeddings in NLP frequently reflect social biases (Garg et al., 2018) and different political sentiments (Yang and Roberts, 2021). More specifically, Yang and Roberts (2021) show that, compared with Chinese language Wikipedia which is free of censorship, word embeddings trained on corpora that are monitored by the Chinese government can have very different semantic meaning for a variety of concepts such as democracy, freedom, collective action, and equality.

Will the similarity score be different when semantics reflected by word embeddings are different? To see this, we use two sets of word embeddings from Word2Vec models pre-trained with either *People's Daily Newspaper* (political semantics) or *Chinese Wikipedia* (scientific semantics). We weighted the two types of word embeddings by TF-IDF scores, and construct two separate similarity scores. Table A.15 gives an example that using word embeddings trained with political and scientific language can generate different results when evaluating the similarity between paper pairs.

Using the similarity score measured with political language as the outcome variable, we then estimate the leadership effect with the baseline specification and control the similarity score measured with scientific language. We note the significant increase in the political similarity score after leadership transmission in Figure A.12. This suggests politics as an important mechanism through which leaders can influence faculty's research direction.

5.3 Persistence of past academic persecution

Based on a record of academic persecution, which includes 38 events at the department level after 2013, we create a university-year-specific shock and see if it increases the main effect we identify after the shock. These events are about people who publicly publish an article or express opinions against the government's political agenda. These people were removed from their positions.

Based on this record, we create a university-year-specific shock and see if it increases the main effect we identify after the shock. After hand-coding these events and matching them with the 108 universities in our sample, we find that:

1. Most of the academic persecutions happen at lower-level universities. Only 11 events are about the schools in our sample since we mainly consider high-ranked universities in China;

2. Social sciences are more vulnerable to academic persecution. There are very few events happen to natural science. Only 3 of the 38 events happen to professors in the natural science department.

Figure A.6 Panel A shows that the leader effect on research similarity is more salient after the persecution.

The extent to which academic leaders are able to influence research directions of the faculty members is linked to their control over them, which has its roots in the broader institutional structure of modern China. If so, we may also expect that these roots are persistent and may be linked to prior episodes of top-down actions to control academia by the Chinese Communist Party. To explore this issue, we investigate whether experiences of past academic persecution initiated by the Communist Party have persistent effects and are linked to the political pressures faced by contemporary Chinese academics. For this purpose, we follow Wang and Kung (2021) and measure the academic disciplines' likelihood of facing top-down persecution during the Cultural Revolution. Wang and Kung (2021) show that academics in disciplines with higher ideological dissension, particularly those in humanities and social sciences, faced disproportionately higher rates of intellectual persecution during the Cultural Revolution than their peers in sciences, applied sciences and engineering.

To study the influence of past academic persecution from the Communist party on political pressure among contemporary Chinese academia, we proxy the severity of persecution during the Cultural Revolution with the extent of ideological dissent within a particular discipline and then explore whether the similarity effects we estimate our more

pronounced in such disciplines. Specifically, we follow Wang and Kung (2021) in using the Discipline Classification and Code (*xueke fenlei yu daima*) developed by the National Standardization Management Committee (*guojia biaozhunhua guanli weiyuanhui*) to construct a measure of ideological dissent for each academic discipline. Under this classification scheme, each academic discipline is ranked based on the level of "consensus" and "paradigmatic development" in the discipline.[15] Consensus within a discipline around general accepted theories, laws, frameworks, methods and beliefs serves as a proxy for how much ideological dissent vs. general academic co-optation there is within that discipline. Disciplines that exhibit lower consensus (and thus higher ideological dissent) are those in humanities and social sciences, while natural sciences show higher consensus and thus lower ideological dissent. Based on the ranking of second-level classification codes, we construct a normalized index which assigns a value of zero to the discipline with the lowest level of ideological dissent and one to the discipline with the highest dissent.

We then re-estimate the baseline specification in (A.8) separately by discipline and plot in Figure 10 the estimated leadership transition effects against the ideological dissent rank. The pattern we find is highly suggestive.[16] There is a positive relationship between the severity of persecution during the Cultural Revolution and the impact of leaders on the research direction of the faculty under their control. The implied relationship is significant at 10% (*p-value* = 0.0966) and suggests that more severely-persecuted disciplines (during the Cultural Revolution) exhibit greater responsiveness to leader switches today. We interpret this finding as showing a persistent legacy of political pressure and deprivation of academic freedom that originated during the Cultural Revolution. It also raises the intriguing possibility that lack of academic freedom can have persistent effects holding back independent research initiatives in present day, potentially resulting in substantial long-term cost to research quality (as we explore next).

[15]The classification scheme assigns a unique code that identifies up to three levels of classification for each academic discipline: 62 first-level discipline groups (e.g., economics) are divided into 676 second-level disciplines (e.g., labor economics), and then further subdivided into 2,382 third-level disciplines (e.g., labor economic history). As in Wang and Kung (2021), we use the second-level classification code for analysis. A faculty member in the discipline of economics, for example, is assigned a unique second-level classification code based on whether they specialized in political economy, labor economics, development economics, business economics, economic history, etc.

[16]Appendix Figure A.7 presents the estimated baseline effects of leadership transition, discipline by discipline. In this exercise, the largest effects are in education, Marxism, management, economics, finance, business, law, media, philosophy, anthropology, ethnology, and sociology, while the effects are muted in political science, public management, and foreign language, and even negative, though imprecise, in history and psychology.

6 Implications for research quality

Do politically-charged incentives impact research quality? We already saw that the increase in similarity after a leadership switch is greater in lower-ranked universities and for leaders who are themselves below-average in terms of their research output. We now directly look at whether attempts to curry favor with leaders results in lower-quality research.

Our first strategy is to look at whether the quality of a leader has an impact on the citation counts received by faculty research papers published by the faculty member that follow her appointment.[17] Specifically, we define high (low) productivity leaders as those who have produced above (below) median numbers of research publications prior to their leadership appointment. The median is computed for the sample of other leaders in the same discipline in similarly-ranked universities.

We estimate the effects on citation counts of future research papers following leader switches from below-average to above-average leaders, and vice versa. For comparison, we show the effects of leadership transitions with no change in leader quality as well. In all specifications, we control for a full set of faculty member fixed effects and year fixed effects. Finally, we also look separately at average citation counts, citations to papers that are most similar to the new leaders (defined as the paper with the highest similarity score with the current leader for each researcher) and citations to papers least similar to leaders (defined as as the paper with the lowest similarity score with the current leader for each researcher).

The results are presented in Figure 11. We find that a switch from a below-median to an above-median leader is associated with greater citation counts (on average a 7.6% increase) and a switch from an above-average leader to a below-average leader is associated with significantly fewer citation counts (on average a 19.7% decrease). There are no effects from switches that involve leaders in the same quality category. Notably, we also see that these results are entirely driven by citation counts of papers that are most similar to the leaders' research. In the bottom two panels, when we look at papers that have little similarity to leaders, there are precise zeros, which suggests that our estimates are not driven by spurious factors. In addition, in none of these cases do we see any effect before

[17] Since the faculty members are always exposed to academic leaders, we cannot identify the main effect of political influence by leaders on research quality — for example, if the political career concerns caused by leaders in Chinese academia leads to, say, a 10% lower citation counts per paper, this will be the same both before and after a leader switch. (Indeed, Appendix Figure A.8, discussed below, shows that a first leadership transition has no impact on research quality.) This motivates our empirical design where we look at the differential effects of below-average leaders.

the relevant switch, which is comforting regarding the validity of this empirical design.

Overall, we interpret these estimates as capturing the significant costs of politically-charged incentives in Chinese academia for research quality. Combined with our earlier results that showed an oversized influence of below-average leaders and Communist Party representatives, these results suggest that Chinese researchers are often incentivized to change their research style to suit the preferences of academically-undistinguished leaders, and indeed, this is often associated with low-citation research papers.

In Appendix Figure A.8, we add to this evidence by separately estimating leadership transitions' impact on faculty members' citation counts, depending on how many leadership transitions a faculty member has experienced since joining their current department. It is worth noting that here we are not distinguishing between below-average vs. above-average leaders, and thus these estimates should be read as the impact of additional leadership transitions, regardless of the characteristics of leaders. We find that while the first and second leadership transitions have relatively small effects, further transitions have significant negative effects. This evidence suggests that there may be a cumulative negative impact of leadership transitions on research quality — perhaps because several changes in research direction intended to curry favor with political leaders start having cumulative costs. Figure A.6 Panel B shows the leader effect on citation before and after academic persecution. Since the leader switch effects on citation are in different directions for transitions to low and high-quality leaders, we change the outcome variable to negative for transitions to high-quality leaders. The negative leader effect on citation increases after universities experience academic persecution.

Taken together, the evidence presented in this section indicates that the pivot of research direction towards a new leader's style tends to move faculty members away from their academic strength and comes at the cost of producing lower quality research.

7 Conclusion

Throughout history, most authoritarian regimes have been suspicious of innovation, research and new technologies, and have often discouraged or even sometimes blocked them (Mokyr, 1992; Acemoglu and Robinson, 2012). Even Soviet Russia, which poured huge resources into military and nuclear technologies and cultivated top-quality research in chemistry, physics and mathematics, was opposed to new technologies that were deemed to be destabilizing (Fitzpatrick, 1999). In this light, modern-day China may be viewed as an almost unique case of an authoritarian regime deeply committed to innovation. But is this enough for producing high quality research? Or do the authoritarian political sys-

tem and its reverberations throughout Chinese bureaucracy and society still distort the direction of research and suppress its quality? These questions are central not just for the future of China's growth, but also for global innovation, especially given China's growing role therein. Nevertheless, we are not aware of any systematic investigation of the impact of political factors in the direction and quality of research and innovation.

In this paper, we undertake such a study. We exploit the appointment of new department leaders in the 109 top Chinese universities, who typically have extensive powers for resource allocation, promotion and termination. These leaders and their powers at least partly reflect the authoritarian nature of Chinese political system and the organizational structure of Chinese academia. The main question we explore is whether the appointment of new leaders leads to a change in the research portfolio and style of faculty members under their jurisdiction, and whether this comes at a significant cost in terms of research quality. We build a data set comprising the academic publications of all leaders and faculty members in these universities. Using NLP methods, we construct measures of similarity between leaders' and faculty members' research output. We combine these data together with data on changes in leadership switches (where leaders comprise of department heads, deans and Chinese Communist Party representatives and departments) and data on citation counts.

Our main finding is strong increase in research similarity between a leader and the faculty under her jurisdiction. Reassuringly, there is no pre-trend — the increase in similarity starts after the new leader takes up office. We also show that after a leadership transition, leaders in the same discipline but in other universities as well as leaders in other disciplines do not have similar effects. We interpret these results as being due to politically-charged career concerns in Chinese academia that primarily impact faculty via local pressures exerted by (or implicitly felt from) the leaders under whose jurisdiction they are.

Career concerns are not confined to Chinese academia or authoritarian settings, however. It is plausible to presume that analogous changes in research strategy may happen in academic systems with greater autonomy. To build the case that the patterns we describe go beyond what would happen in situations where there is greater autonomy, less political interference and better institutional safeguards for meritocratic promotions and external review, we adopt a number of complementary strategies. First, the effects of new leaders on their faculty is more pronounced in lower-ranked departments, which typically lack procedures for external review. Second, we document that they have grown in importance after tenure-track reforms were introduced, which in practice increased the power of academic leaders over the careers of faculty under their rule.

Do political pressures affect the quality as well as the direction of research? We show that the answer is likely yes. When there is a switch from a below-average leader to an above-average leader (in terms of publication output), there is an increase in the quality of research as measured by citation counts. Conversely, a switch from an above-average to a below-average leader is associated with significant costs in terms of citation counts. We document that these results are driven entirely by papers that are most similar to new leaders' research portfolio, thus bolstering the case that research is being redirected away from areas in which faculty have expertise towards less highly cited, lower-quality areas intended to curry favor with new leaders. We also present results suggesting that leadership switches have negative cumulative effects — faculty experiencing several leadership switches become less productive.

We view our paper as a first step in a research agenda that explores the relationship between political factors and the direction and quality of innovation, originating both from national institutions and local organizations. This agenda is important for several reasons. As new technologies such as AI, nanotechnology and new materials become increasingly important, the quality and direction of research and innovation for global prosperity are likely to grow. However, how different political systems and local and global incentives coming from political or other considerations, impact the direction of this research is unclear. This is critical for the future of Chinese growth, which can be seen as a unique historical experiment in combining an authoritarian political system with a relentless focus on innovation. It is also central for understanding the forces impacting the nature of academic research in many other contexts, and as input into designing better academic institutions under democratic institutions.

In this light, there are several interesting research areas that can be further explored. First, it is important to conduct similar studies in other contexts, which would enable a valuable comparative perspective — in particular to see whether similar politically-charged career concerns are present in the academia of less authoritarian countries. Second, in more data-rich environments, it may be possible to look at other characteristics of leaders, such as where they have obtained their degree and how they have risen in the academic hierarchy. Last but not least, a similar analysis in the context of corporate innovation, for example, linking the nature of patents to the priorities and organizational structure of the firms under which the research is being conducted, would be a very fruitful area.

References

Acemoglu, Daron and James A Robinson, "Economic Backwardness in Political Perspective," *American Political Science Review*, February 2006, *100* (1), 1–17.

_ **and** _ , *Why Nations Fail* The Origins of Power, Prosperity, and Poverty, New York: Crown Business, August 2012.

Akcigit, Ufuk, Santiago Caicedo, Ernest Miguelez, Stefanie Stantcheva, and Valerio Sterzi, "Dancing with the Stars: Innovation Through Interactions," *NBER Working Paper*, March 2018, pp. 1–62.

Azoulay, Pierre, Joshua S Graff Zivin, and Gustavo Manso, "Incentives and creativity: evidence from the academic life sciences," *The RAND Journal of Economics*, September 2011, *42* (3), 527–554.

Bai, Chong-En, Chang-Tai Hsieh, and Zheng Song, "Special Deals with Chinese Characteristics," *NBER Macroeconomics Annual*, January 2020, *34*, 341–379.

Becker, Sascha O, Volker Lindenthal, Sharun Mukand, and Fabian Waldinger, "Scholars At Risk: Academic Networks and High-Skilled Emigration from Nazi Germany," 2021.

Beraja, Martin, David Y. Yang, and Noam Yuchtman, "Data-intensive Innovation and the State: Evidence from AI Firms in China," *NBER Working Paper*, January 2021, pp. 1–84.

Biasi, Barbara and Song Ma, "The Education-Innovation Gap," Technical Report, working paper, Yale 2020.

Bombardini, Matilde, Bingjing Li, and Ruoying Wang, "Import Competition and Innovation: Evidence from China," *Working Paper*, January 2018, pp. 1–44.

Dai, Andrew M, Christopher Olah, and Quoc V Le, "Document embedding with paragraph vectors," *arXiv preprint arXiv:1507.07998*, 2015.

Fitzpatrick, Sheila, *Everyday Stalinism* Ordinary Life in Extraordinary Times: Soviet Russia in the 1930s, Oxford University Press, March 1999.

Freeman, Richard B and Wei Huang, "China's "Great Leap Forward" in Science and Engineering," *NBER Working Paper*, April 2015, pp. 1–31.

Garg, Nikhil, Londa Schiebinger, Dan Jurafsky, and James Zou, "Word embeddings quantify 100 years of gender and ethnic stereotypes," *Proceedings of the National Academy of Sciences*, 2018, *115* (16), E3635–E3644.

Hayek, Friedrich August, *The Constitution of Liberty*, University of Chicago Press, 1978.

Hill, Ryan and Carolyn Stein, "Race to the Bottom: Competition and Quality in Science," *Working Paper*, 2021.

Jia, Ruixue, Huihua Nie, and Wei Xiao, "Power and publications in Chinese academia," *Journal of Comparative Economics*, 2019, *47* (4), 792–805.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, "Measuring technological innovation over the long run," Technical Report, National Bureau of Economic Research 2018.

Li, Shen, Zhe Zhao, Renfen Hu, Wensi Li, Tao Liu, and Xiaoyong Du, "Analogical Reasoning on Chinese Morphological and Semantic Relations," in "Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)" Association for Computational Linguistics 2018, pp. 138–143.

Manso, Gustavo, "Motivating Innovation," *The Journal of Finance*, September 2011, *66* (5), 1823–1860.

Mation, Lucas Ferreira and Aguinaldo Maciente, "GROUP_TWOWAY: Stata module to group observations by the connected components of two variables," 2014.

Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.

Mokyr, Joel, *The Lever of Riches* Technological Creativity and Economic Progress, Oxford University Press, April 1992.

North, Douglass C, John Joseph Wallis, and Barry R Weingast, *Violence and Social Orders* A Conceptual Framework for Interpreting Recorded Human History, Cambridge: Cambridge University Press, February 2009.

Ridley, Matt, *How Innovation Works* And Why It Flourishes in Freedom, HarperCollins, May 2020.

Waldinger, Fabian, "Peer effects in science: Evidence from the dismissal of scientists in Nazi Germany," *The review of economic studies*, 2012, *79* (2), 838–861.

Wang, Alina Yue and James Kai-Sing Kung, "Silence is Golden: Ideological Dissension and Selective Persecution in China's Cultural Revolution," *Working Paper*, 03 2021, pp. 1 – 57.

Wei, Shang-Jin, Zhuan Xie, and Xiaobo Zhang, "From "Made in China" to "Innovated in China": Necessity, Prospect, and Challenges," *Journal of Economic Perspectives*, February 2017, *31* (1), 49–70.

Yang, Eddie and Margaret E Roberts, "Censorship of online encyclopedias: Implications for nlp models," in "Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency" 2021, pp. 537–548.

Zhu, Junwen, " SCIE : 1978-2007." PhD dissertation, : 2009.

Zong, Xiaohua and Wei Zhang, "Establishing world-class universities in China: deploying a quasi-experimental design to evaluate the net effects of Project 985," *Studies in Higher Education*, 2019, *44* (3), 417–431.

Figures and Tables

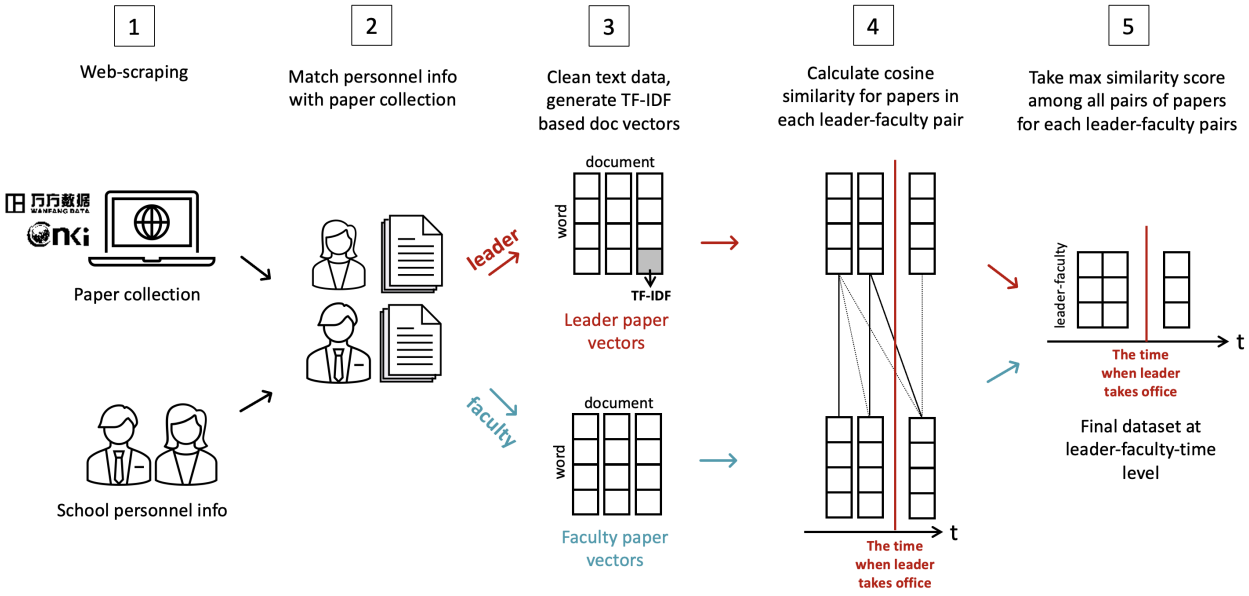


Figure 1: The data construction procedure.

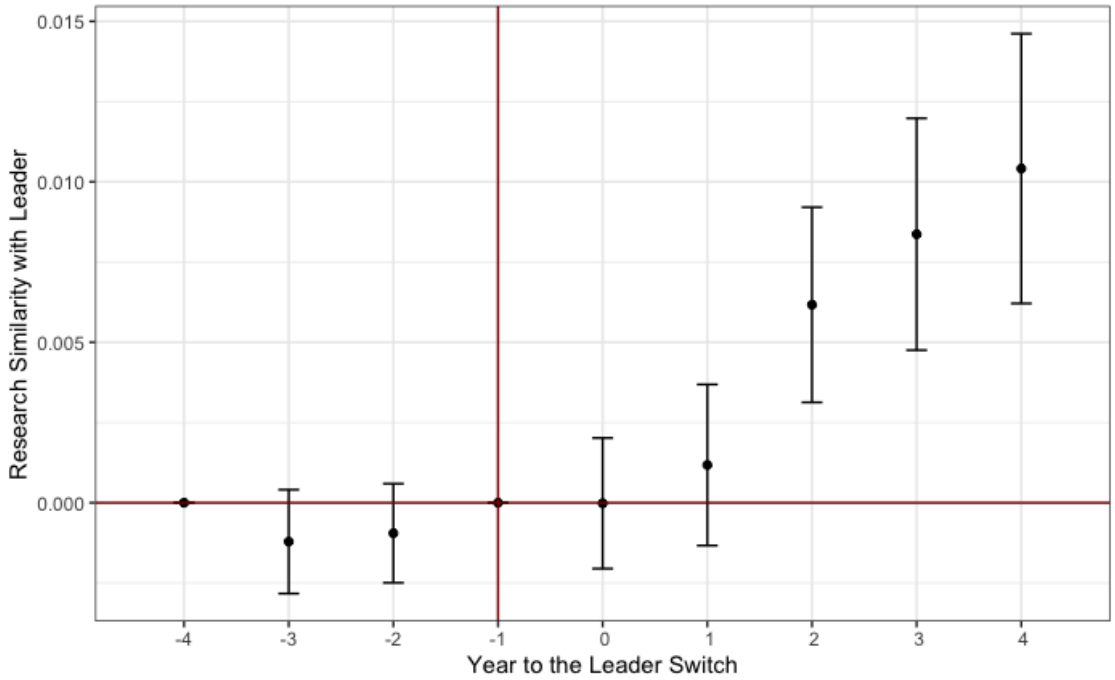


Figure 2: Impact of leader switch on the faculty-leader similarity score. The points in the figure represent the estimated effects of event time relative to leadership switch, controlling for the full sets of faculty-leader pair fixed effects and time fixed effects (i.e., the ψ_l from the nonparametric event study in equation A.8, $Y_{i,t} = \sum_{l \neq -1; l=-3}^{4} \psi_l D_{i,t}^{l} + \alpha_i + \lambda_t + v_{i,t}$). The error bars represent the 95% confidence intervals.

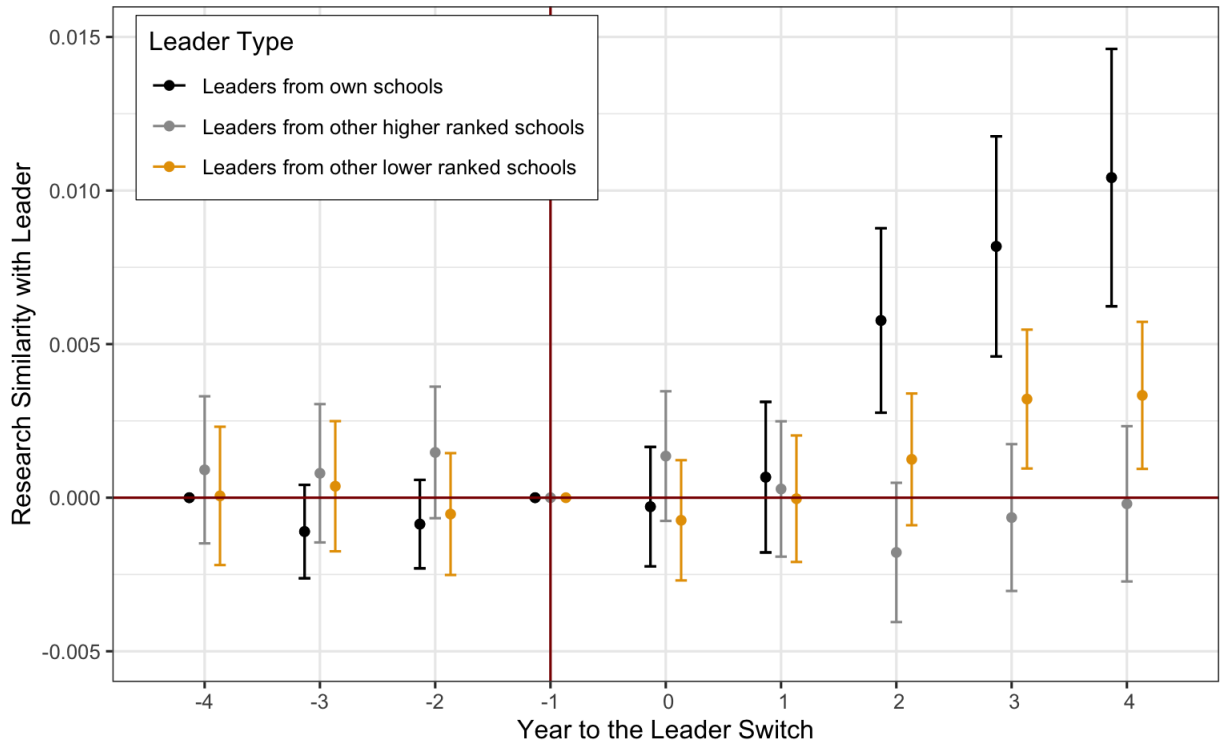


Figure 3: Impact of leaders from other higher ranked schools vs. leaders from other lower ranked schools on the similarity score. This panel uses the faculty-leader pairs in which leaders from other higher ranked schools and leaders from other lower ranked schools. We estimate the effect of the two type of leaders simultaneously in regression $Y_{i,t,c} = \sum_{l \neq -1; l=-3}^{4} \mu_l D_{i,t,c}^{l} \times L_i + \sum_{t \neq -1; t=-3}^{4} \psi_l D_{i,t,c}^{l} + \alpha_i + \lambda_t + v_{i,t,c}$, where L_i is the indicator for whether the leader in pair i is from other higher ranked schools (=1) or not (=0). The grey lines/markers represent the estimated effects of leaders from other higher ranked schools (i.e., the ψ_l in the regression). The yellow lines/markers represent the estimated effects of leaders from other lower ranked schools (i.e., the $\mu_l + \psi_l$ in the regression). The black lines/markers replicate our baseline results in Figure 2.

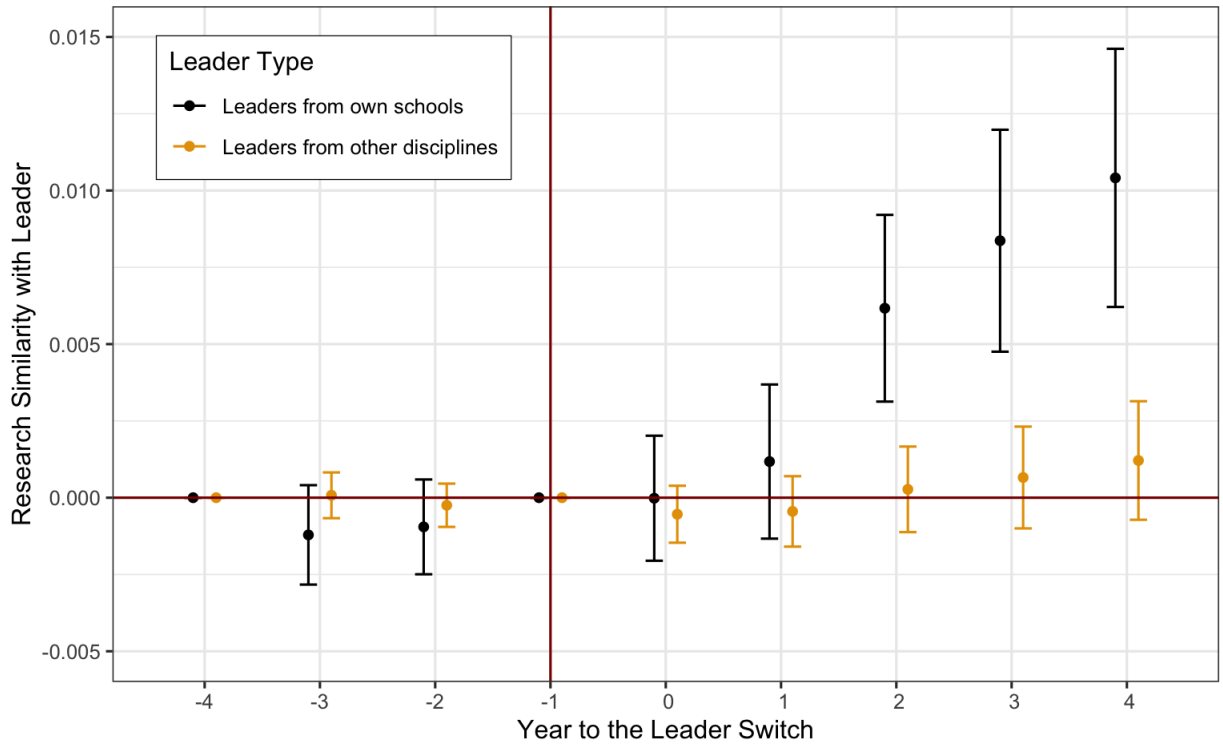


Figure 4: Impact of leaders from other disciplines on the similarity score. This panel uses the faculty-leader pairs in which leaders from other disciplines. The yellow lines/markers in the figure the estimated effects of event time relative to leadership switch, controlling for the full sets of faculty-leader pair fixed effects and time fixed effects (i.e., the ψ_l from the nonparametric event study in equation A.8, $Y_{i,t} = \sum_{l \neq -1; l=-3}^{4} \psi_l D_{i,t}^{l} + \alpha_i + \lambda_t + v_{i,t}$). The black lines/markers replicate our baseline results in Figure 2.

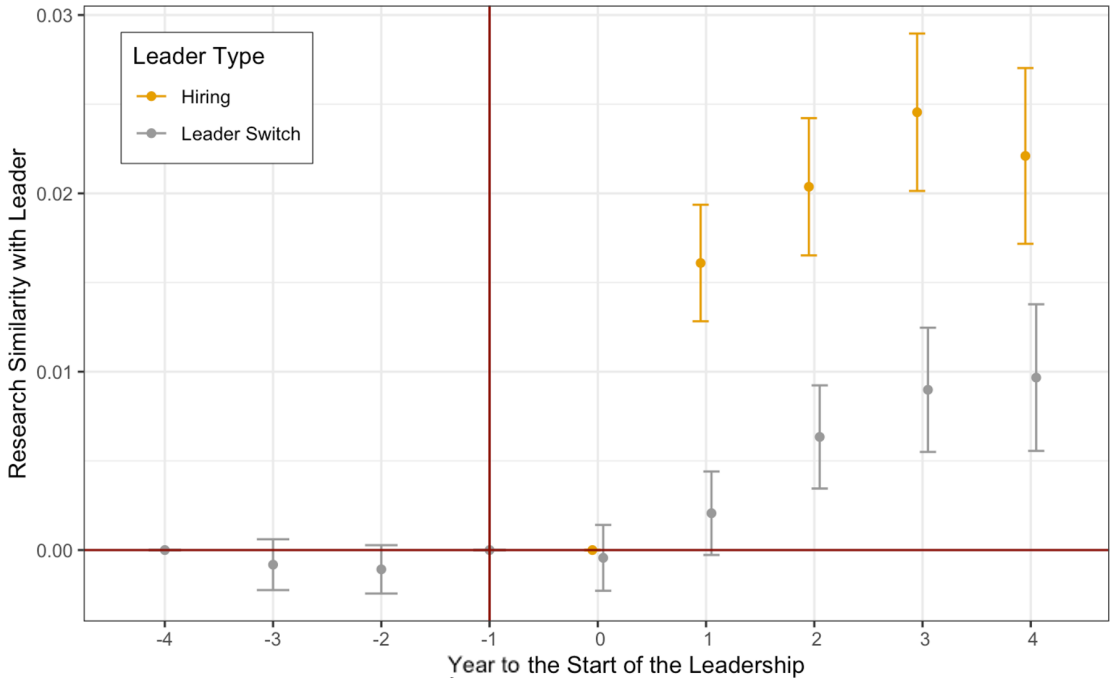


Figure 5: Impact of leadership transition for hiring leaders versus department heads. The yellow lines/marks represent the estimators for the effect of hiring leaders. The grey lines/marks represent the effect of leader switch. We add faculty members that are newly hired to the baseline sample: (1) similarity scores will the similarity between the faculty member and the leader hired her; (2) the treatment year is the year when the faculty member is recruited by the school.

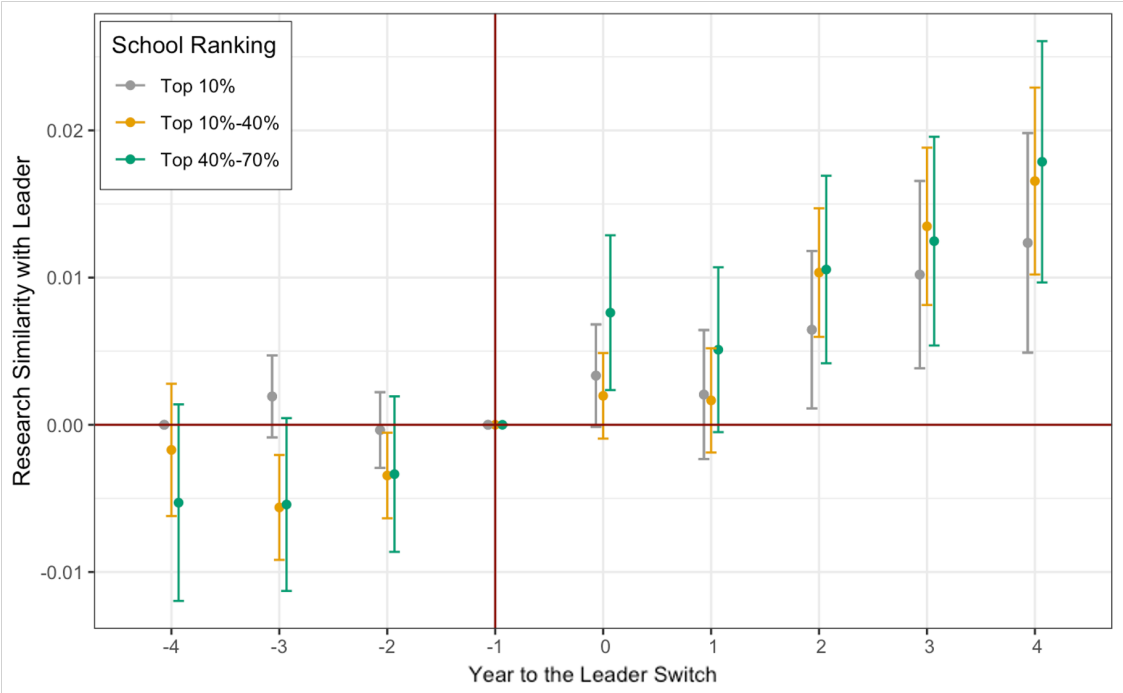


Figure 6: Heterogeneous effect of leader switch by school ranking. We estimate the effects simultaneously in regression $Y_{i,t} = \sum_k \sum_{l \neq -1; l=-3}^{4} \mu_l^k D_{i,t}^l \times R_i^k + \sum_{t \neq -1; t=-3}^{4} \psi_l D_{i,t}^l + \alpha_i + \lambda_t + v_{i,t}$, where R_i^k is the indicator for the rank of the school of pair i. The grey lines/markers represent the estimated effects of leaders from schools ranked top 10%. The yellow lines/markers represent the estimated effects of leaders from schools ranked 10%-40%. And the green lines/markers represent the estimated effects of leaders from schools ranked 40%-70%.

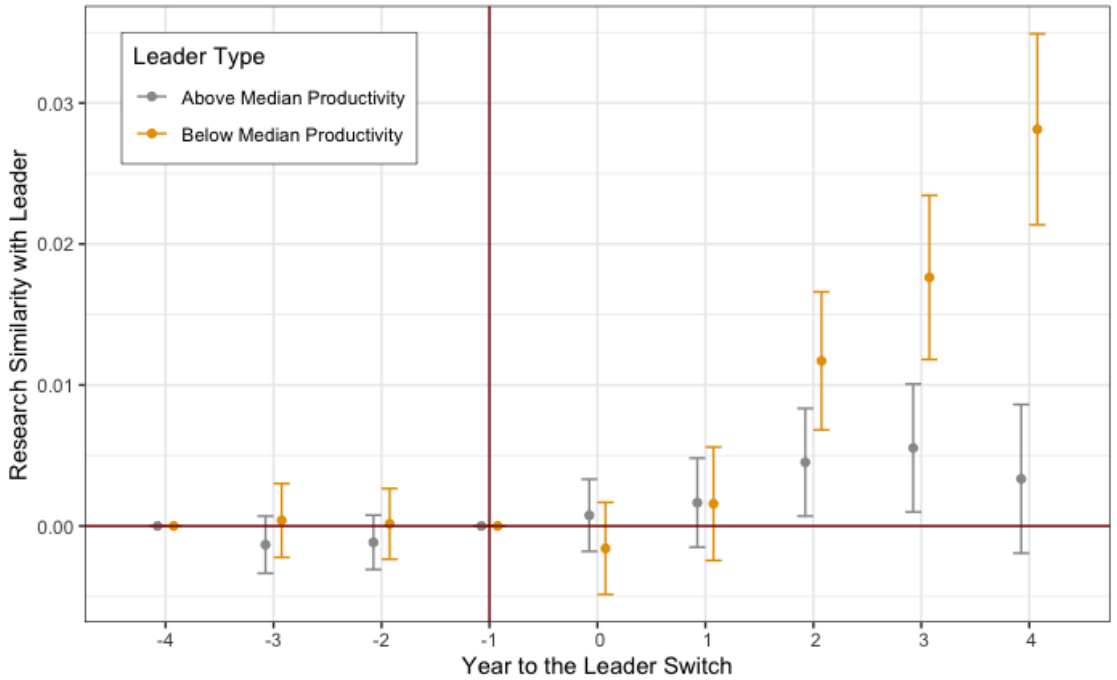


Figure 7: Heterogeneous effect of leader switch by the productivity of leaders. We define the productivity of a leader as the average number of publications of the leader in the 3 years before she is in office. Then we take the median of leaders' productivity for each discipline. A leader will be categorized as "above mean productivity" if her productivity is above the median number of her discipline. Otherwise, she will be categorized as "below median productivity". We estimate the effect of leader with different productivity level separately with our baseline regression (i.e., we assume that pairs with different leader productivity have different calendar year fixed effects). The yellow lines/marks represent the estimators for the effect of leaders that are below the median productivity. The grey lines/marks represent the estimators for those above the median productivity.

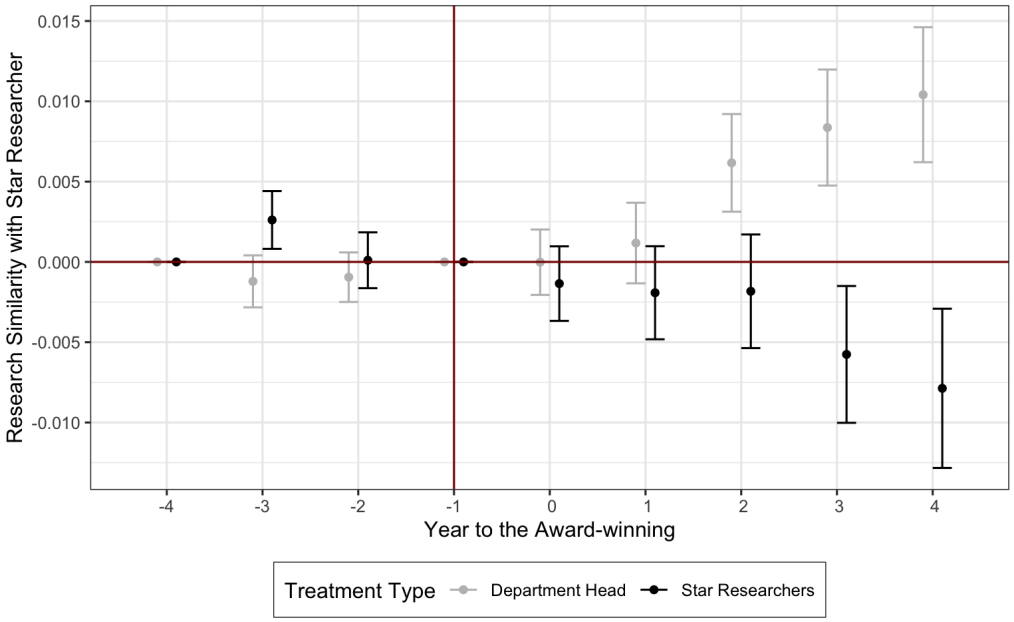


Figure 8: The impact of star researchers on the faculty-star researcher similarity score versus the impact of department heads on faculty-leader similarity score. The black dots/lines show the effect of winning the award on research similarity. Time 0 is the year when the star researchers earn the award from Changjiang Scholar Program. The grey dots/lines show our baseline result – the effect of leader switch on research similarity. Time 0 is the year when the department head starts her position. The error bars represent the 95% confidence intervals.

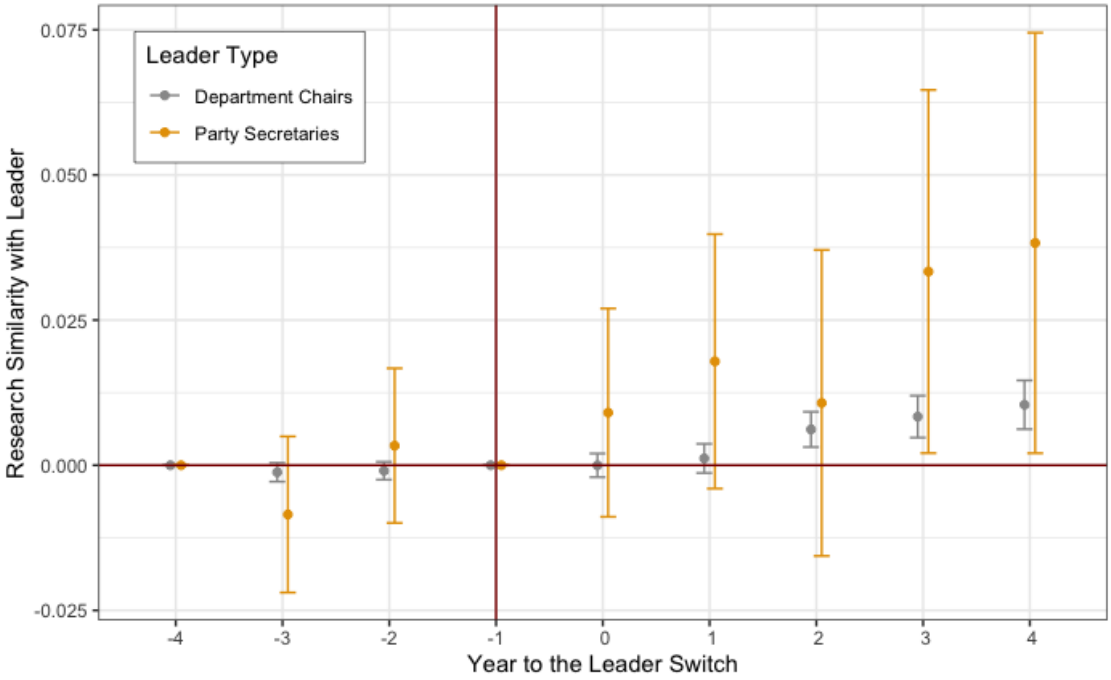


Figure 9: Heterogeneous effect of department chairs and party secretaries. The yellow lines/marks represent the estimators for party secretaries, which is also based on the baseline regression. Since many party secretaries are non-academic personnel, we restrict our sample to party secretaries whose productivity is above the median of department chairs to ensure that we are focusing on academic party secretaries. The grey lines/markers represent the estimated effects of leaders department heads.

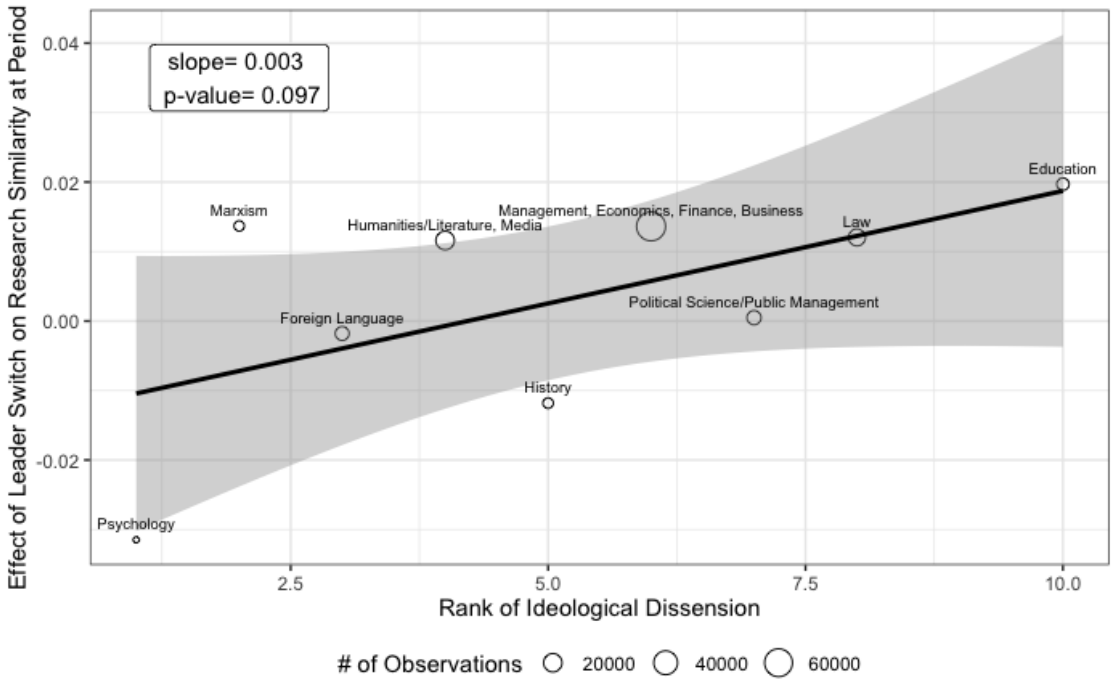


Figure 10: The correlation between the effect of leadership switch and the measure of ideological dissension. The x-axis shows the ranking of "ideological dissension." There is less ideological dissension in the discipline when the ranking is lower. The y-axis is the effect of leader switch at t=4. Each point represents a discipline.

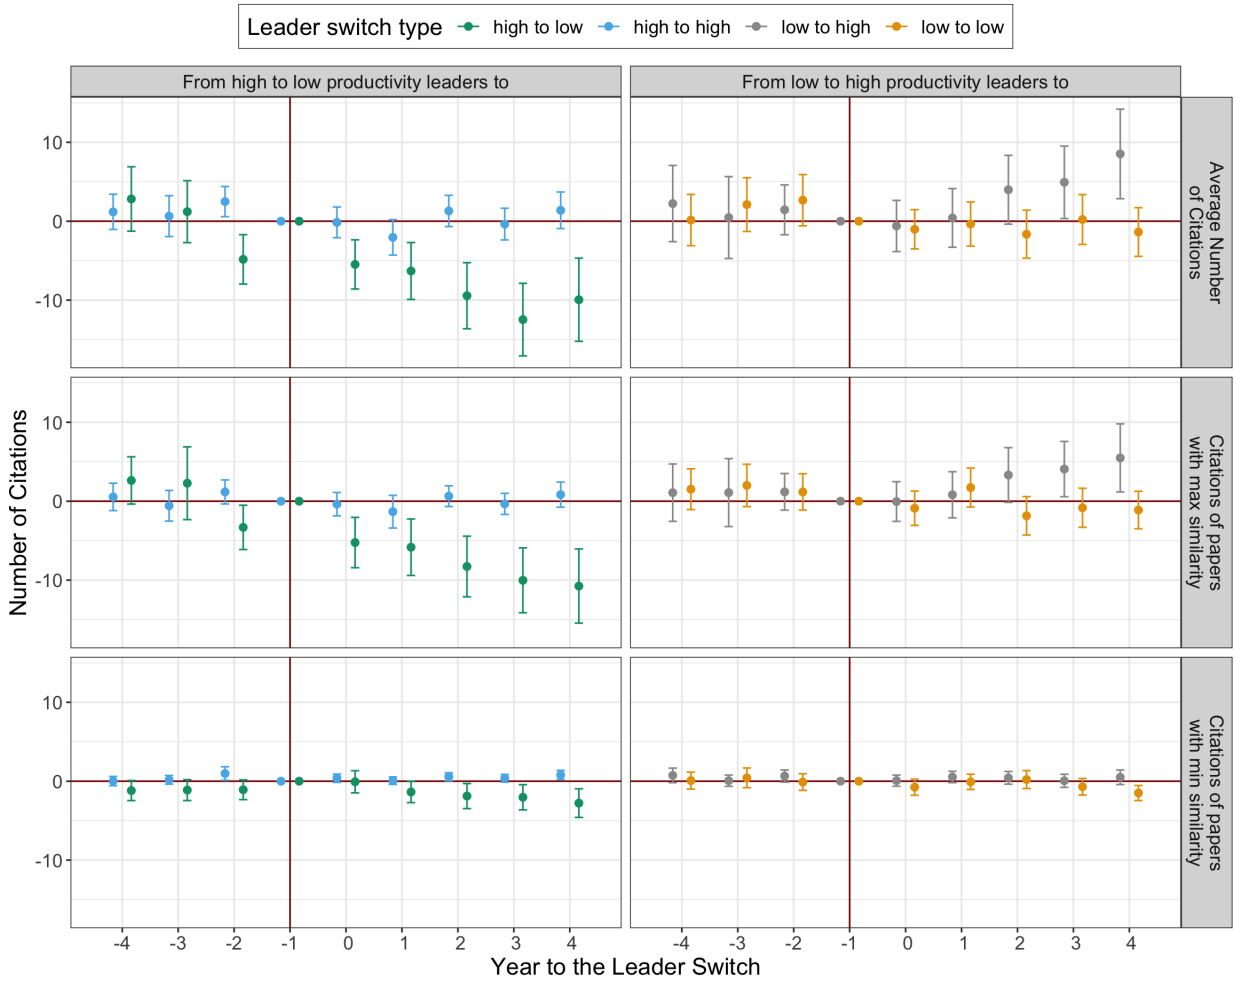


Figure 11: Impact of leaders on citation counts. The points in the figure represent the estimated effects of event time in the following regression: $Y_{i,t} = \sum_{l \neq -1; l=-4}^{4} \psi_l D_{i,t}^l + \alpha_i + \lambda_t + v_{i,t}$, where $D_{i,t}^l$ is an indicator for faculty i being l periods away from initial treatment at calendar year t; α_i is a full set of the faculty fixed effects; and λ_t is a full set of calendar time fixed effects. The error bars represent the 95% confidence intervals. The grey lines/markers represent the estimated effects using the sample of low-to-high leader switches. The yellow lines/markers represent the estimated effects using the sample of low-to-low leader switches. The green lines/markers represent the estimated effects of leaders on citation counts using the sample of high-to-low leader switches. And the blue lines/markers represent the estimated effects of leaders on citation counts using the sample of high-to-high leader switches

Table 1: Hiring leaders vs. other leaders

	Similarity Score		
	Pooled (1)	First year (2)	Pooled (3)
Panel A: TF-IDF Mean			
Dummy for hiring leader	0.00019	0.00015	0.00014
	(0.00021)	(0.00032)	(0.00021)
Panel B: TF-IDF Max			
Dummy for hiring leader	0.01097	-0.00226	0.00893
	(0.00137)	(0.00197)	(0.00139)
Faculty FE	Yes	Yes	Yes
Calendar Year FE	Yes	Yes	Yes
Control for Event Time	No	No	Yes
Number of obs	207852	61947	207852

Note: (1) Use the sample where the relative year is between 0 and 4. Regression here is: $Y_{ijt} = \beta H_{ijt} + \alpha_i + \gamma_t + \epsilon_{ijt}$, where Y_{ijt} is the similarity score between faculty i and leader j at calendar year t, H_{ijt} is an indicator for whether leader j is the hiring leader of faculty i at year t, α_i and γ_t control for the faculty member and year fixed effect.
(2) Only use the sample where the relative year equals to 0. Regression here is: $Y_{ijt} = \beta H_{ijt} + \alpha_i + \gamma_t + \epsilon_{ijt}$.
(3) Use the sample where the relative year is between 0 and 4. Regression here is: $Y_{ijt} = \beta H_{ijt} + \sum_l \psi_l D^l_{ijt} + \alpha_i + \gamma_t + \epsilon_{ijt}$, where D^l_{ijt} is an indicator for faculty i and leader j being l periods away from initial treatment at calendar year t.

APPENDIX

A Categorizing multidisciplinary schools

The categorization is done with the following steps:

Step 1 For each school, we extract disciplines that are (i) listed in school level code, and (ii) listed in the name of departments that are under the school. We need (ii) to improve accuracy, since some schools can cover disciplines that are not indicated by their school-level names. For example, the Guanghua School of Management of Peking University (北京大学光华管理学院) only has the discipline "management" in its name, but actually covers other disciplines including business, finance and economics.

At the end of this step, for each school, we get an array of disciplines in the school. For example, the Guanghua School of Management of Peking University has the following disciplines: (Management, Management, Management, Economics, Economics, Finance, Finance, Business).

Step 2 Within each school, we drop disciplines which only take less than 25%. The threshold is chosen ad hoc by checking if the final categorization makes sense. Notice that there are 10 out of 787 schools are dropped in this process since there is no discipline in the school is more than 25%. These schools are usually called "School of Social Science" (人文社会科学院) which are a mixture of all social sciences. If we include the 10 schools, all disciplines need to be combined into one.

After this step, one school has at most 3 kinds of disciplines. For example, the Guanghua School of Management of Peking University has 3 disciplines: Management, Economics, Finance.

Step 3 Within each school, disciplines are ranked by percentage. So the first discipline will be the major discipline of the school. Given the first discipline, we need to check what other disciplines are usually linked to it. Some links are very rare. For example, most of the Marxism schools are independent of other disciplines. But Northeastern University (东北大学) combines Marxism and Philosophy, and China Agricultural University (中国农业大学) combines Marxism and History. These links will mess up the categorization. So I dropped these rare links and categorize the schools with their first discipline. The criteria that we use is: given the first discipline, drop if this type of links only takes less than 15%.)

Step 4 Finally we group observations by the connected disciplines by using *group_twoway* by Mation and Maciente (2014). The final categorization is as follows:

- Marxism
- Political Science, Public Management
- Law
- Management, Economics, Finance, Business
- Education

- Foreign Language
- Humanities, Literature, Media
- History
- Psychology
- Philosophy, Anthropology, Ethnology, Sociology
- Regional Studies

B Identification of faculty members

In this section we discuss the strategy we use to identify faculty members from students or other unaffiliated researchers for a given department.

The difficulty for getting a full set of faculty from 1990 and 2019 is that most of the universities don't have good records of faculty at department level. We utilize the scientific publications of all affiliates in the 109 universities and extract faculty lists based on authors and affiliations to pin down the list of faculty and assign them to schools they are affiliated.

The general process is: *(i)* Manually find keywords for identifying department; *(ii)* Filter faculty members with certain criteria.

Manually find the "keys" for identifying department This is a key step for identifying faculty members. The affiliation entries in our publication data are typically very messy. What makes thing worse, some people don't use the full name of their departments/schools to put it in the affiliation. For example, a professor affiliated to the Department of Applied Economics of Guanghua School of Management at Peking University (北京大学光华管理学院应用经济系) could possibly put something like "Guanghua School at Peking University" (北京大学光华学院) or "DAE of Guanghua School of Management at Peking University" (北京大学光华管理学院应经系) to the affiliation of his paper. Therefore, we must manually extract some "keyword" to match the affiliation. The protocol for adding the searching keywords is:

1. University names are automatically added to the set of keywords for all the affiliated departments;
2. We make sure that each set of keywords can uniquely identify one department. One should be very careful when trying to use generic keywords to identify a department. For example, when we try to use "Fianace" (金融) to extract papers and faculty members for the Department of Finance of School of Economics at Peking University (北京大学经济学院金融系), which will be contaminated by the Department of Finance of Guanghua School of Management at Peking University (北京大学光华学院金融系) and the Department of Financial Mathematics of School of Mathematics at Peking University (北京大学数学学院金融数学系). Therefore a possible key in this case to uniquely identify the Department of Finance of School of Economics at Peking University (北京大学经济学院金融系) could be "Peking University" (北京大学), "Econ" (经), and "Finance" (金融).
3. Most affiliations in the papers are precise only to the school level (just like HBS and Harvard SEAS), not to department level (in US equivalent, Harvard Econ department and Political Science department). This is because, most Chinese "schools" are US "departments" equivalent, and Chinese "departments" are equivalent to something like the macroeconomics group at Harvard econ department, although there are typically still bureaucracy structures and CCP establishments in this very disaggregated level. We try to identify faculty members by their department (in the Chinese sense), however in many times we are not able to do that.

Filter faculty members with certain criteria After we can identify department/school from papers, we use these to identify faculty members. Our current criteria for a faculty member are: *(i)* having more than 3 years of publication span; *(ii)* having 5 or more papers. The first is to exclude PhD students who typically publish papers in 2-3 years, and the second is to guarantee enough variations for us to exploit in regressions.

Validation with one university To validate the method that we use to extract faculty members, we compare the faculty we extracted with the faculty list that we can obtain from the the official website of the School of Economics at Sun Yat-sen University. Table A.1 shows the validation result.

Table A.1: Compassion between extracted faculty and the actual personnel

Year	Actual Number of Faculty	Number of people extract from raw data	Number of Faculty after filtering
2019	44	353	23
2018	41	327	31
2017	37	305	35
2016	35	301	39
2015	33	297	39
2014	31	296	40
2013	29	337	44
2012	28	335	46
2011	24	340	52

C Similarity score based on Doc2Vec methods

Possible measurement error in TF-IDF The biggest issue for TF-IDF is that it can severely underestimate the similarity between two documents, which might be a major measurement error for our baseline regression. Consider comparing the similarity between the following two documents:

- Document A is "The car is driven on the road."
- Document B is "The truck is driven on the highway."

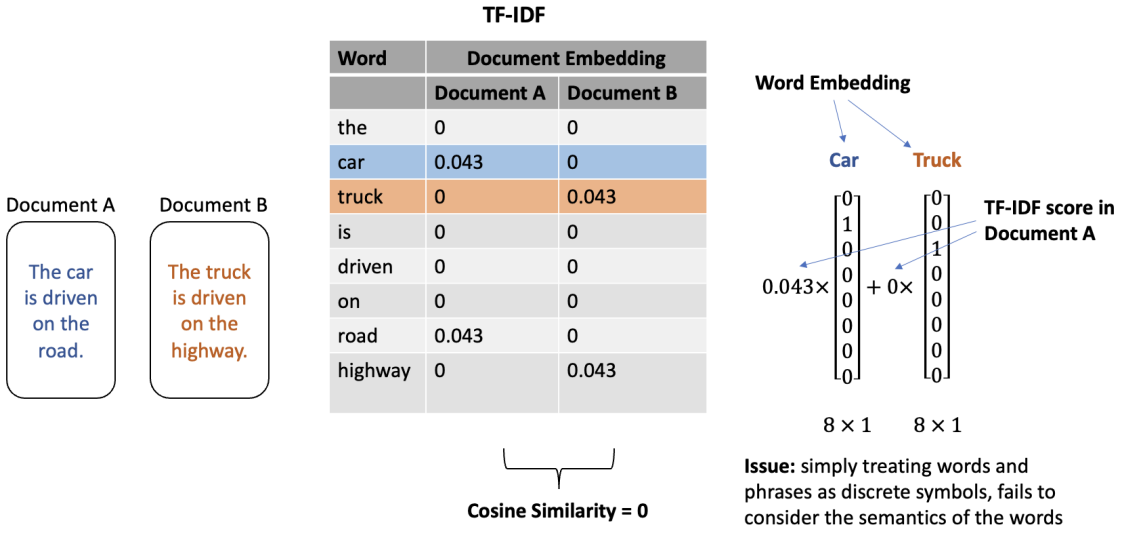


Figure A.1: A illustrative example of TF-IDF

The table in the middle shows the document embeddings (vectors) of the two documents. The document embedding is an 8x1 vector, where 8 is the size of the vocabulary. Each element in the vector is the TF-IDF score of the corresponding vector. After calculating the cosine similarity score of the two documents, we get the similarity score between the two documents equals 0. For each document embedding, we can view it as a weighted average over all of the word embeddings in the document. As shown on the right of Figure A.1, the weights are TF-IDF scores, and the word embeddings are standard unit vectors.

Representing words as standard unit vectors treats words as discrete symbols. We assume "car" and "truck", "road" and "highway" are totally different, which leads to the result that the similarity score is 0.

Taking semantics of the words into consideration Replacing word embeddings with Word2Vec (Mikolov et al., 2013) will relax the assumption of TF-IDF that words are orthogonal.

The Word2Vec algorithm uses a neural network model to learn word associations from a large corpus of text. One of the Word2Vec structures – continuous bag of words (CBOW) – forms a sliding window around the current word and allows it to be predicted using the

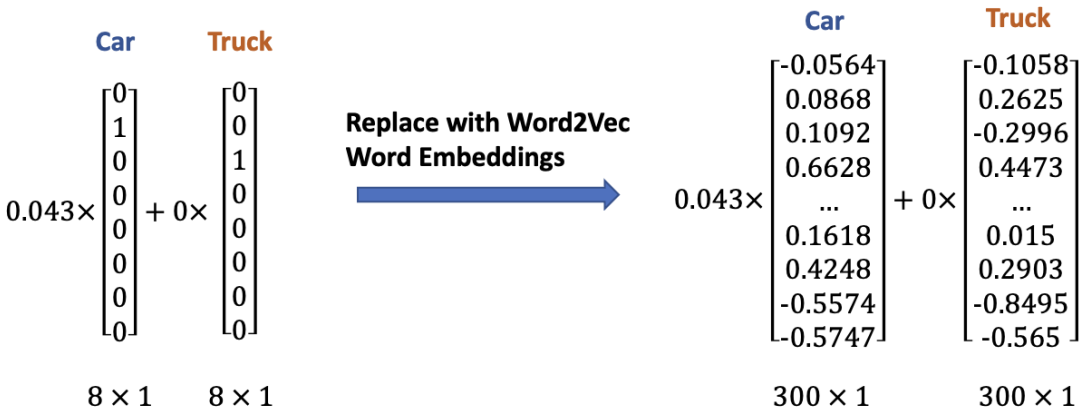


Figure A.2: Change from TF-IDF to Weighted Word2Vec

words around it. A feature vector is used to represent each word. With this process, each word will be mapped to a $n \times 1$ vector and can represent the semantic similarity between words[1]. If we reconsider our previous example, using word embeddings from Word2Vec will increase the similarity score between Document A and B. The cosine similarity between "car" and "truck" is 0.81 rather than 0 in the TF-IDF case. With the new weighted average over all of the word embeddings in the document, the similarity score between Document A and B is 0.8442.

While training an informative Word2Vec model requires large train set and computation power, many researchers and institutions publicize their training results based on large corpus such as Wikipedia, newspapers, and twitter. One can use their results as pre-trained word embeddings and further finetune models to fit her own target corpus.

In our setting, we utilized two pre-trained Chinese word embeddings provided by Li et al. (2018): *(i)* Word2Vec trained with People's Daily News, and *(ii)* Word2Vec trained with Chinese Wikipedia.

Taking the ordering of the words into consideration Although weighted Word2Vec relaxes the orthogonality of semantics, another assumption is still concerning when numerically representing paragraphs – it ignores the syntax within each paragraph. Doc2Vec is developed to further tackle this issue and give paragraph embeddings.

Instead of directly using TF-IDF weights, Doc2Vec puts more structure on the relationship between words and documents to estimate the weights. Based on CBOW which utilizes words to anticipate the following word, Doc2Vec also includes a document-unique feature vector. As a result, when the word vectors are trained, the document vector is also trained which can be used as the numeric representation of the document.

[1] n will be a hyperparameter in the Word2Vec model. In the example here, we set $n = 300$.

D English Publications

Data Validation for English Publications To provide a benchmark for the completeness of English publications, we use the number of papers that we previously manually collected based on the search results of Google scholar for 100 randomly selected scholars. In Figure A.3, each dot in the figure represents a scholar. We put the number of papers collected from Google Scholar on the x-axis and the number from CNKI on the y-axis. The figure shows that our scraping results are well aligned with the 45-degree and correlation coefficient between Google Scholar and CNKI equals 0.75.

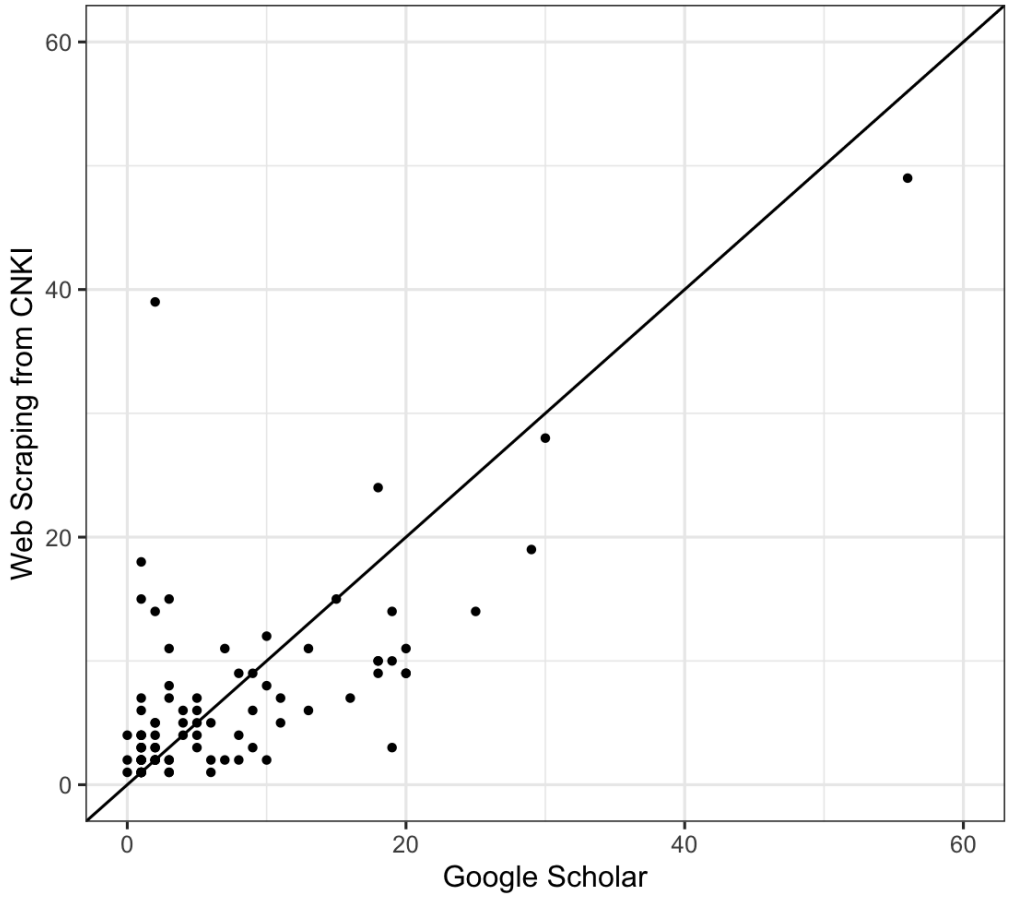


Figure A.3: Each dot represents a scholar, for which the x-axis shows her paper number in Google scholar. The left y-axis (red) is the number of the current round, and the right (blue) is the number of our previous round. The black line is a 45-degree equivalence line.

For papers that are on Google but cannot be found through CNKI, most of the missing papers are because of CNKI's journal collection: (1) some small journals are not collected by CNK, (2) some journals on CNKI don't have author info listed. Concerns might be that these journals are censored in China and scholars will publish papers with different opinions on these journals. Based on the 10 cases we looked into, we didn't find it is the case.

Google scholar also includes papers that are (1) not published, such as NBER working paper, or (2) conference papers. These papers are not under our consideration.

NLP Analysis for English Publications To compare the similarity between Chinese and Enlgish papers, we first translate English papers to Chinese with Azure Translator. And then, we calculate similarity scores based on the Chinese translation.

We first select 3 papers with both Chinese and English abstracts. The English abstracts are translated by the authors themselves. We select the 3 papers such that P2 is very similar to P3, and P3 is less relevant to P1. Denote the three papers as P1, P2, and P3. $P1_{eng}$ will refer to the English abstract of P1, and $P1_{zh}$ is the Chinese abstract of P1. $Sim(P1_{eng}, P1_{zh})$ then represents the similarity score between the Chinese and English abstracts of P1. Then the similarity score works well if:

1. $1 = Sim(P1_{eng}, P1_{zh})$
2. $Sim(P1_{eng}, P1_{zh}) > Sim(P2_{eng}, P1_{zh}) > Sim(P3_{eng}, P1_{zh})$

The following table shows the similarity scores calculated with the method:

$Sim(P1_{eng}, P1_{zh})$	0.988960176
$Sim(P2_{eng}, P1_{zh})$	0.980958178
$Sim(P3_{eng}, P1_{zh})$	0.971544729

E Additional figures and tables

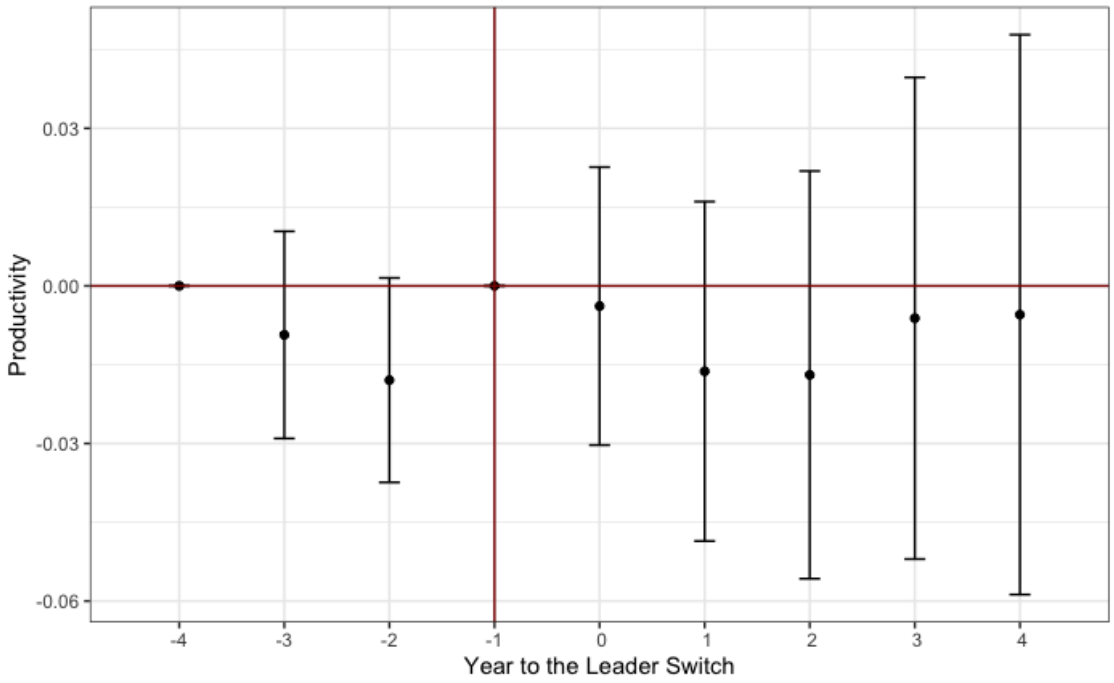


Figure A.4: Impact of leader switch on the productivity of faculty. Using the number of publications per year as the outcome variable, we re-estimate the baseline specification (equation A.8) restricting the sample to a balanced panel of faculty-leader pairs.

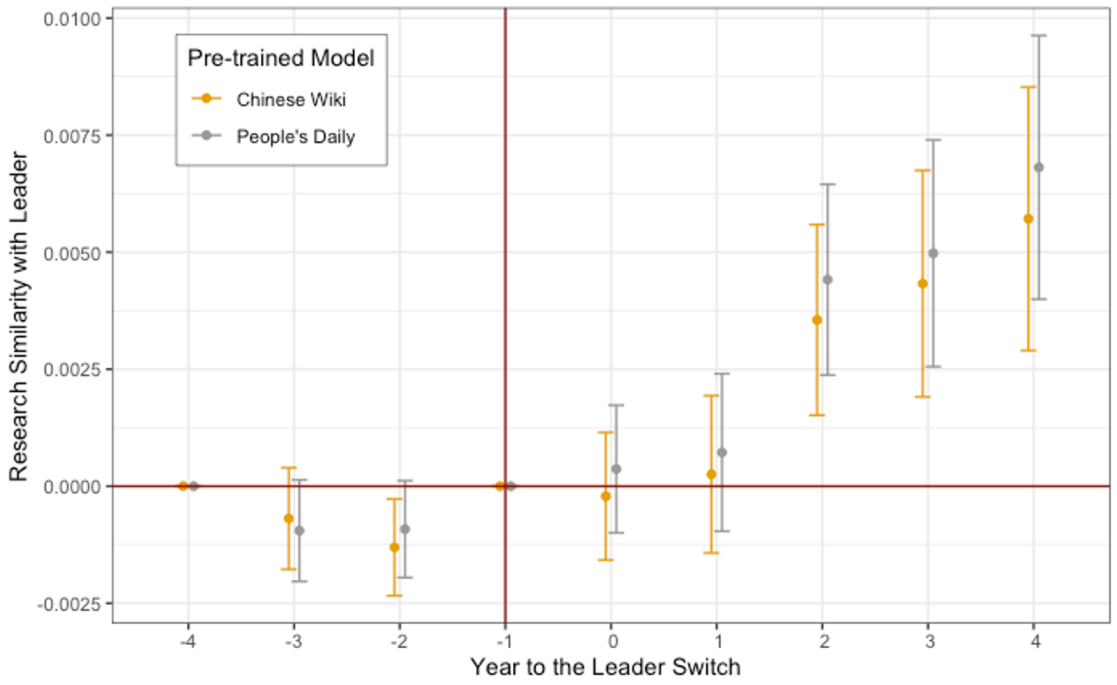
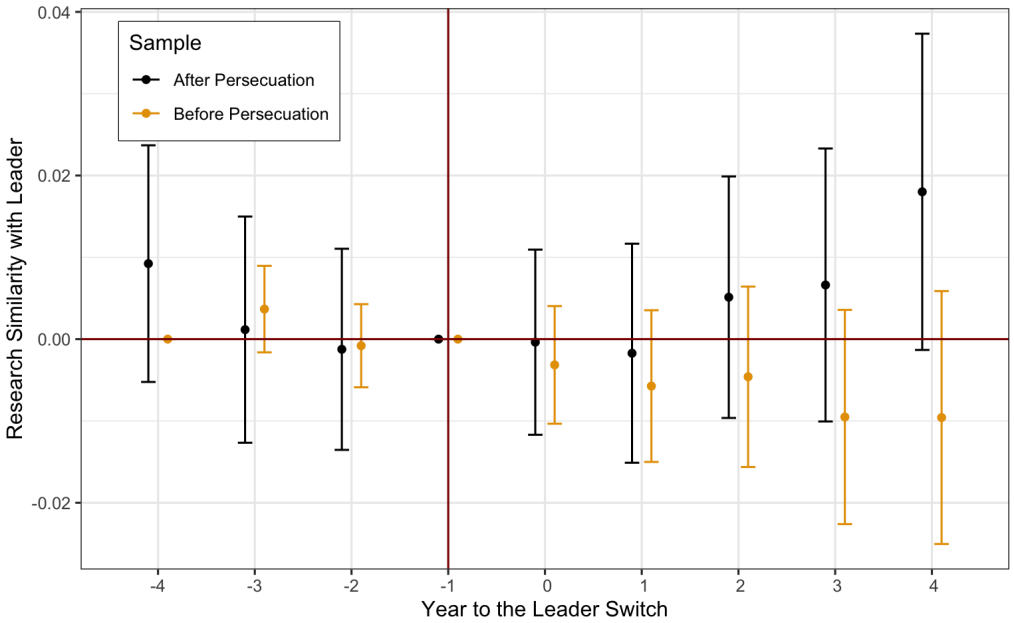
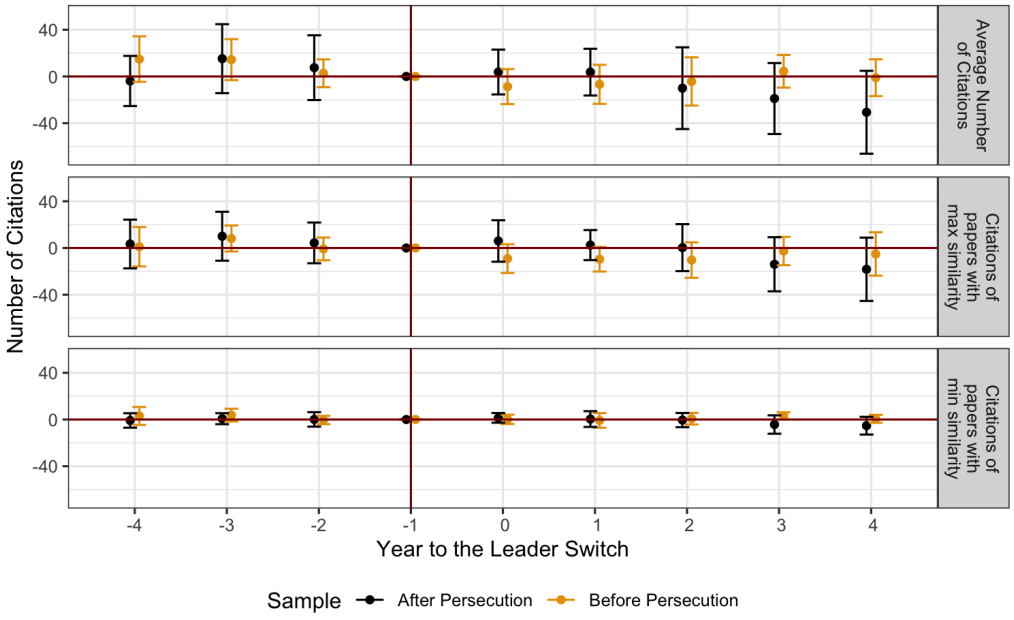


Figure A.5: Robustness check: impact of leader switch on the faculty-leader Doc2Vec similarity score. The points in the figure represent the estimated effects of event time (i.e., the ψ_l from the nonparametric event study in Equation A.8). The error bars represent the 95% confidence intervals.



Panel A: Effects of leader switch on research similarity



Panel B: Effects of leader switch on citation

Figure A.6: Heterogeneous effect of leader switch before and after academic persecution. We estimate the effects simultaneously in regression $Y_{i,t} = \sum_{l \neq -1; l=-3}^{4} \mu_l D_{i,t}^l \times T_i + \sum_{t \neq -1; t=-3}^{4} \psi_l D_{i,t}^l + \alpha_i + \lambda_t + v_{i,t}$, where T_i is the indicator for whether the leader in pair i is after the university has an academic persecution event (=1) or before the event (=0). The yellow lines/markers represent the estimated effects of leaders before the persecution (i.e., the ψ_l in the regression). The black lines/markers represent the estimated effects of leaders after the persecution (i.e., the $\mu_l + \psi_l$ in the regression).

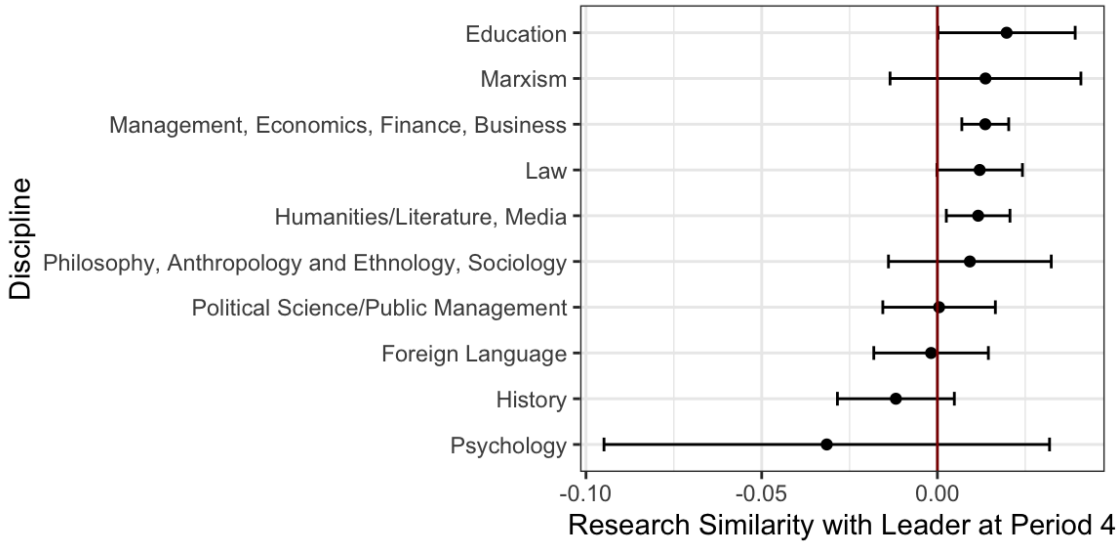


Figure A.7: Heterogeneous effect of leader switch by discipline. We estimate equation A.8 for each discipline. Each line/mark represent the point estimate of the 5th year for each discipline, ranked by size of the estimates. We classify schools into 10 categories by taking disciplines as the connected components of schools that share the same classification. The details are described in Appendix A.

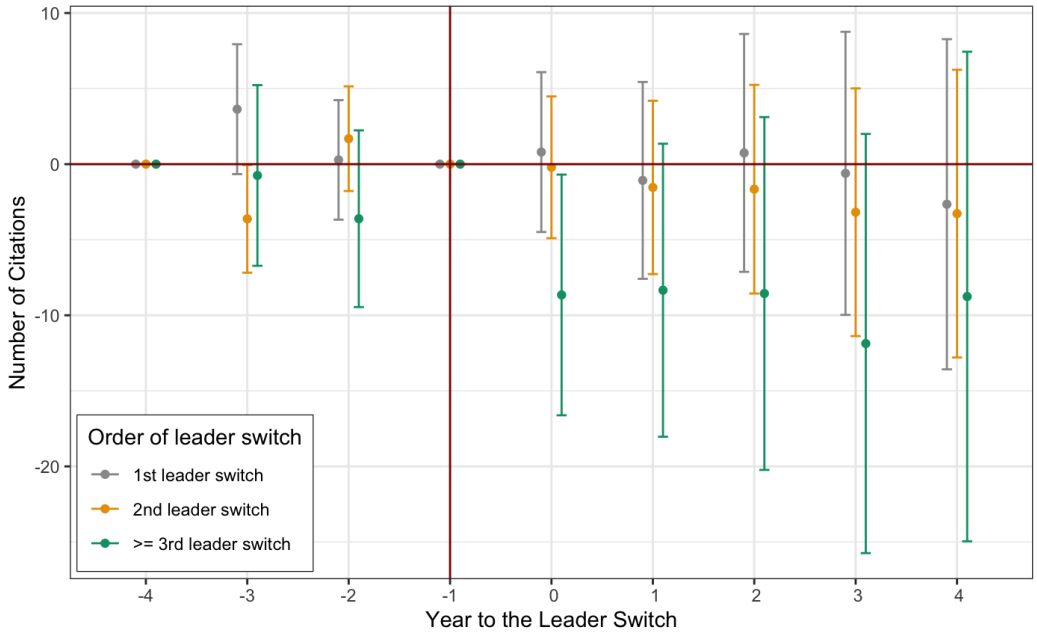


Figure A.8: Heterogeneous effect by the order of leadership transaction. We separately estimate the effect of leader switches on citation counts for different order of leadership transition within a faculty member. The points in the figure represent the estimated effects of event time relative to leadership switch on the faculty-leader similarity score, controlling for the full sets of faculty-leader pair fixed effects and time fixed effects (i.e., the ψ_l from the nonparametric event study in equation A.8, $Y_{i,t} = \sum_{l \neq -1; l=-3}^{4} \psi_l D_{i,t}^{l} + \alpha_i + \lambda_t + v_{i,t},$). The grey markers represent the estimated effects for leader switches that are firstly experienced by faculty. The yellow lines/markers represent the estimated effects of the second leader switch experienced by faculty. And the green lines/markers combine the effect of the rest of the switches.

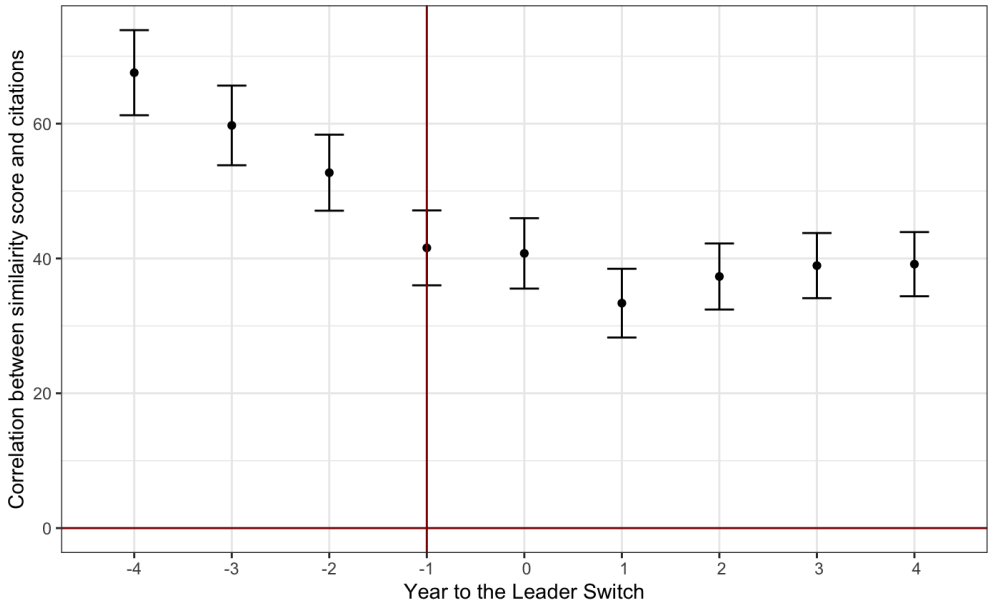


Figure A.9: Correlation between Similarity score and the number of citation counts. The points in the figure represent the coefficients ψ_l in the equation Citation $_{i,t} = \sum_{l \neq -1; l=-4}^{4} \psi_l D_{i,t}^{l} \times$ Similarity $_{i,t} + \alpha_i + v_{i,t}$, where Citation $_{i,t}$ is the average number of citations for faculty i at year t, Similarity $_{i,t}$ is the TF-IDF Max similarity score between faculty i and her leader at year t, α_i is faculty FEs.

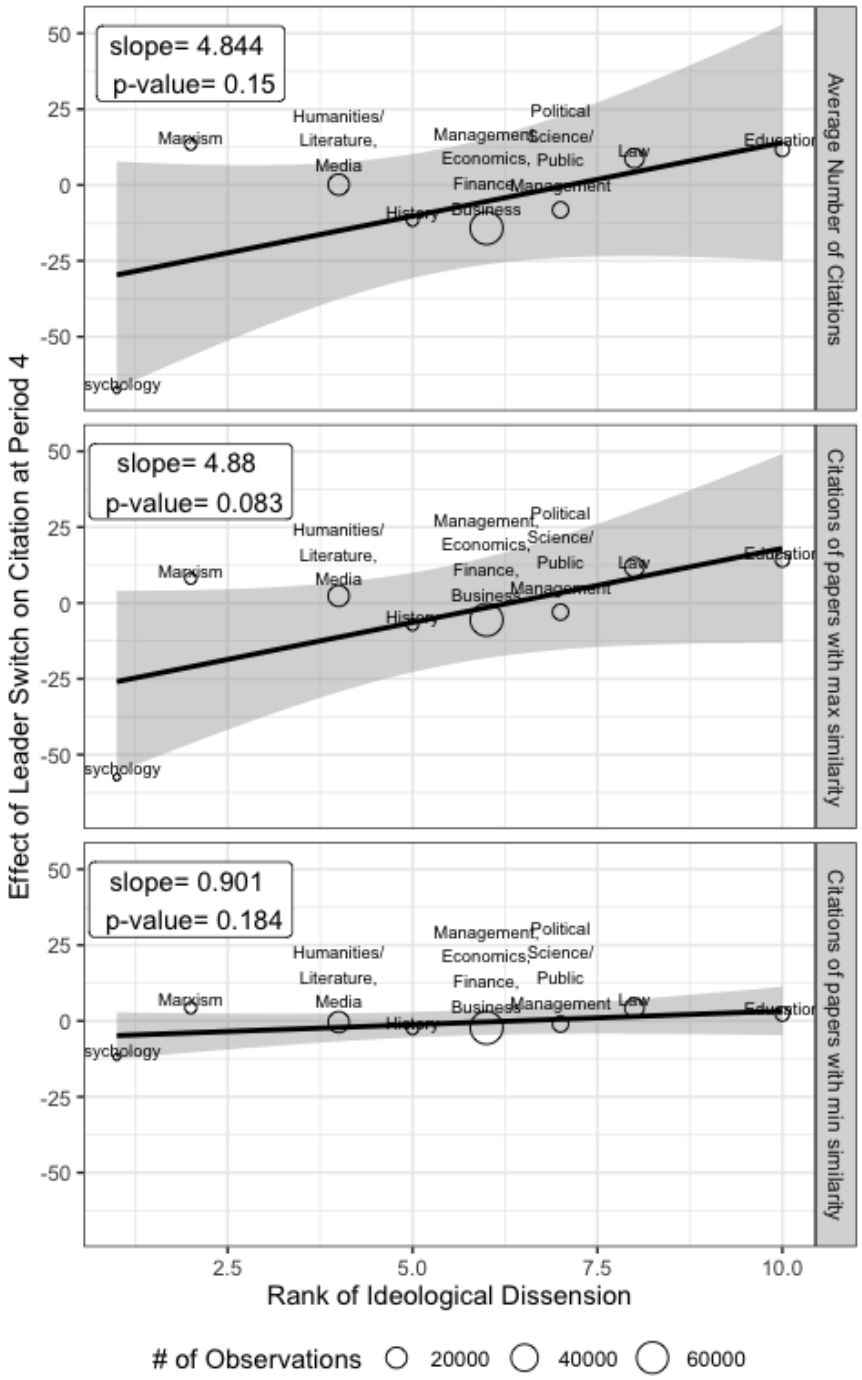


Figure A.10: The correlation between the effect of leadership switch on citation counts and the measure of ideological dissension. The sample only include high-to-low and low-to-high leader switches. We flip the sign of outcome variables for low-to-high leader switches. The panels use different measures for citation counts which are labeled on the right hand side. The x-axis shows the ranking of "ideological dissension." There is less ideological dissension in the discipline when the ranking is lower. The y-axis is the effect of leader switch at t=4. Each point represents a discipline.

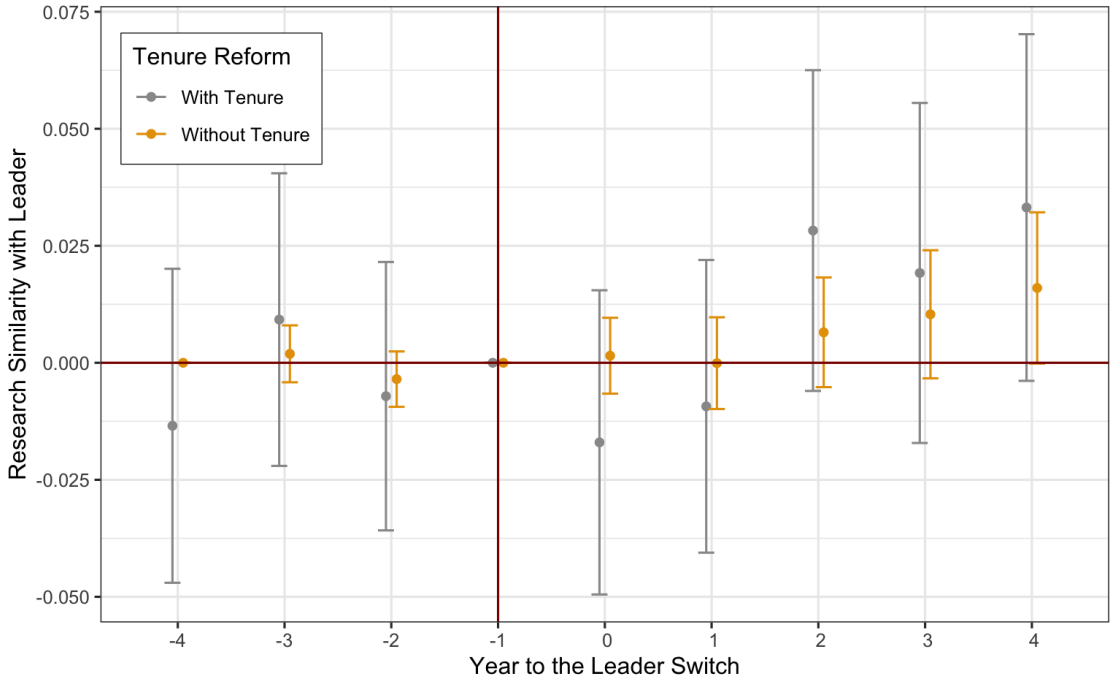


Figure A.11: Heterogeneous effect before and after tenure track reform at Peking University. The points in the figure represent the estimated effects of event time (i.e., the ψ_l from the non-parametric event study in Equation A.8), using three measures of citation counts. The error bars represent the 95% confidence intervals. The grey (yellow) lines/markers represent the estimated effects for leader switches for faculty-leader pairs that are affected (unaffected) by the tenure track system reform.

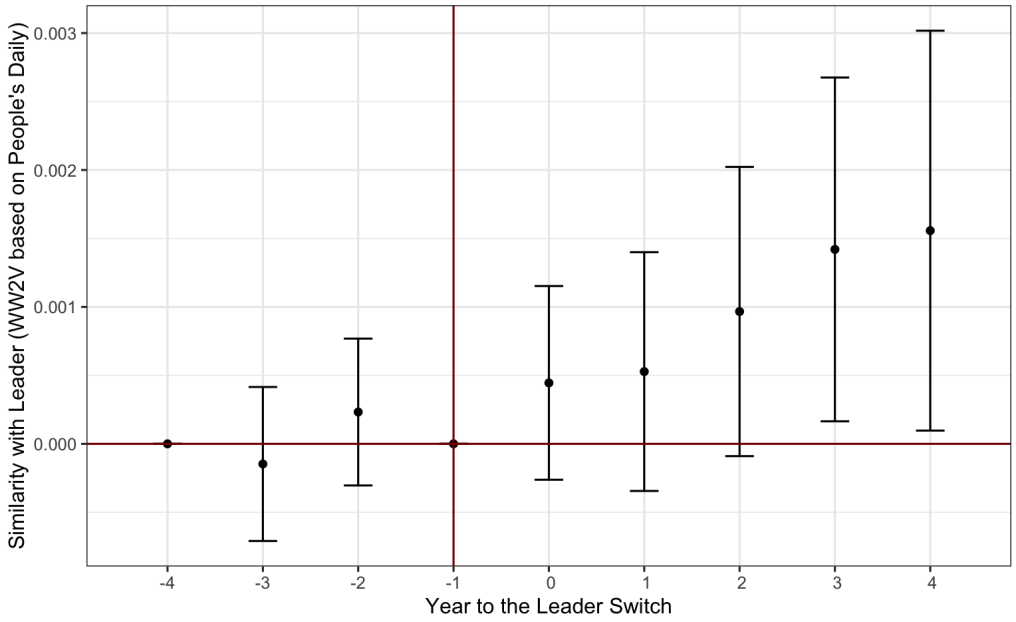


Figure A.12: Impact of department heads on faculty-leader similarity score under political semantics, while controlling scientific semantics. We estimate the leadership effect using WW2V People's Daily similarity score as outcome variable, controlling for WW2V Chinese Wiki similarity score.

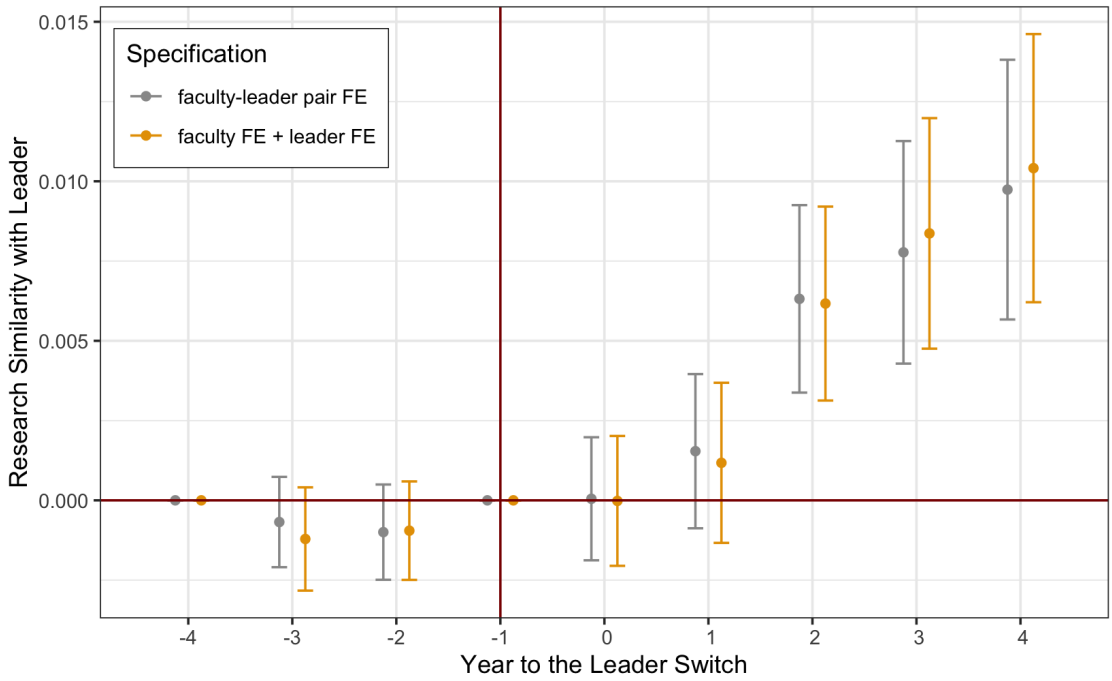


Figure A.13: Comparison between baseline and alternative specifications. faculty fixed effects (rather than leader-faculty pair fixed effects), which is useful, since our analysis of political career concerns on research quality will not include leader-faculty pair fixed effects. The grey points/lines in the figure represent the estimated effects of event time in our baseline equation A.8), controlling year fixed effect and faculty-leader pair fixed effect. The yellow points/lines show the estimated effects of event time in the following regression: $Y_{i,j,t} = \sum_{l \neq -1; l=-3}^{4} \psi_l D^l_{i,j,t} + \alpha_i + \beta_j + \lambda_t + v_{i,j,t}$,. Here $Y_{i,j,t}$ represents the similarity score between faculty i and leader j at time t. $D^l_{i,j,t}$ is an indicator for faculty i and leader j being l periods away from initial treatment at calendar year t. And here we control year fixed effect λ_t and faculty fixed effect α_i, and leader fixed effect β_j.

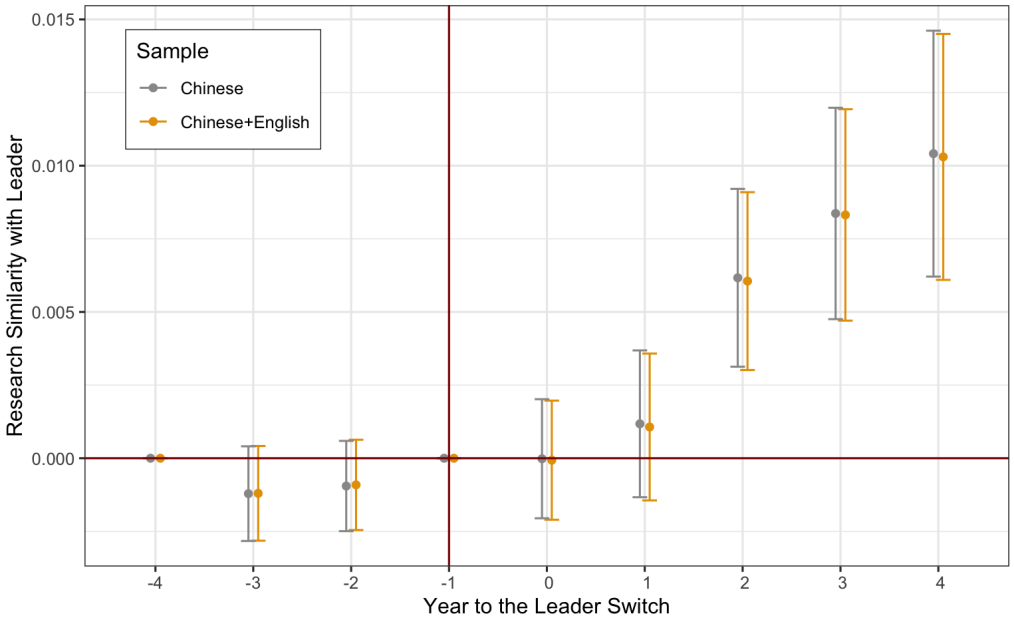


Figure A.14: Impact of leader switch on the faculty-leader similarity score, including English publications. The yellow points in the figure represent the estimated effects of event time relative to leadership switch using *both Chinese and English* publications, controlling for the full sets of faculty-leader pair fixed effects and time fixed effects. The error bars represent the 95% confidence intervals.The grey points/lines represent our baseline leader switch effects, which include Chinese publications only.

Leader's Paper	Faculty's Paper	People's Daily	Chinese Wiki	Difference
From Theoretical Research to Professional Action - a New Model of Public Welfare Projects for Precise Educational Support in Poor Areas (从理论研究到专业行动——贫困地区教育精准帮扶公益项目新模式)	Computational Skills of Junior High School Students and Their Development under the New Curriculum (新课程下初中生运算能力及其培养)	0.605416	0.605447	-0.000031
	From Education for Proletarian Politics to Education for Socialist Construction--Learning the Spirit of the Third Plenary Session of the Eleventh Central Committee of the Communist Party of China (从教育为无产阶级政治服务向教育为社会主义建设服务转变—学习中共十一届三中全会决议精神体会)	0.715717	0.518635	0.197082
	A Discussion on Teachers' Right to Education (论教师的教育权)	0.560299	0.736867	-0.176568

Figure A.15

Table A.2: Summary statistics on citation counts

	Citations per Paper		
	All Papers	Faculty Only	Leaders Only
Mean	13.9	13.7	18.2
25th Percentile	1.0	1.0	1.0
50th Percentile	4.0	4.0	5.0
75th Percentile	13.0	13.0	16.0
Total Papers	736,756	725,455	115,235