How Does Data Access Shape Science? Online Appendix

A Qualitative Evidence

A.1 Interviews

We carried out nine semi-structured interviews with FSRDC administrators and users to learn more about our empirical setting. The interviews have been crucial to provide us with institutional details on the FSRDC network and to validate our research design. Before our first interview, we consulted the Berkeley IRB to assess whether our questions would constitute human subjects research. The IRB assessed that our questions were factual and focused on the institution, thus not requiring formal IRB review. Below we summarize the main topics that emerged during our conversations.

Impact of U.S. Census Bureau's confidential data: All the people we talked to confirmed the revolu-tionary impact that U.S. Census's confidential data had on economics research. Labor, productivity, health, trade, public, and environmental economics were regarded as the fields that benefited the most. However, even more importantly, there is a sense that the linked datasets available in the FSRDCs "opened up fields of research that weren't possible before" (interview T14). One example is constituted by the work of John Haltiwanger and colleagues, whose impact transcended the traditional boundaries between labor and macroeconomics:

> "I think if you look at Haltiwanger's work, I think his work has transformed, I don't know whether you want to call it labor or macro or, you know, kind of the intersection of labor and macro. I think he's had a big impact." (interview Y79)

But besides research, our respondents felt that Census data yielded findings with an outsized policy impact. Our respondents said that this was due to two characteristics of Census data: representativeness and granularity. Without representative data, "you're only getting a sliver of the story", as Nick Bloom said on the occasion of the Stanford FSRDC opening.[1] Without granularity, "special populations get lost in the noise of aggregated data sets," as Bill Maurer remarked during the opening of the UC Irvine FSRDC.[2] In practice, confidential microdata allow to study heterogeneous effects for sub-populations, which is what policymakers usually need (interview P39).

Expansion of the FSRDC network: We inquired about the history and institutional evolution of the FSRDC network. Three key factors shaped the progressive diffusion of research data centers: the presence

[1]Source: https://news.stanford.edu/news/2010/february1/census-data-center-020210.html
[2]Source: https://news.uci.edu/2015/new-uci-center-gives-researchers-link-to-us-census-data/

of a potential user community, the goal of geographical parity in access, and the advocacy of local research leaders. Upon becoming involved in establishing new FSRDCs, the National Science Foundation (NSF) explicitly tried to achieve an equitable geographical diffusion across the United States (interview T14). The NSF was able to pursue this objective by carefully allocating the grants that help institutions set up a new data center. Obtaining the NSF grant was fairly competitive and "lots of places that have tried to get the NSF money haven't been able to" (interview D88). Given this allocation process, the likelihood of winning NSF's support often hinged on the distance from the closest FSRDC: "if you're close to an already existing location, that probably hurts your application" (interview S12). Moreover, the sheer presence of a potential user community was not enough in the absence of individuals that would advocate to open a new data center (interview S94). The following quote well summarizes the importance of vocal advocacy:

> "And he persuaded the Dean that this would be a really good thing for the School. Well, it would be a really good thing for himself, which would be a really good thing for the School, which would be really good for the University." (interview Y79)

In practice, a combination of potential users and individual advocates was key for a successful application to the NSF:

> "[...] for an RDC to get established, you're going to need at least one and preferably a few sort of well-connected and very enthusiastic people who want to push for having an RDC. I don't think there's any case where you'd find that there were sort of 40 people who kind of felt it would be nice to have an RDC, and when number 41 said they'd kind of like to do it, then it all sort of magically happens. You've got to have, and in some cases they might have just been one influential senior person who could kind of really push for it and talk it up." (interview D88)

Moreover, our interviewees revealed a surprising amount of idiosyncrasies behind each FSRDC opening. For instance, the choice of Boston for the first FSRDC was in no small part due to the presence of a Census Regional Office willing to host it, before it was eventually transferred to NBER premises (interview Y79). Another handful of FSRDCs was either opened because of the will of university administrators or because of specific collaborations between some faculty members and the Census Bureau (interview S94). In cases where the FSRDC was opened by a consortium of universities, the choice of which consortium member would host the data center was often the result of a compromise (interview S12). Taken together, this qualitative evidence suggests that both the timing and the locations of FSRDCs were strongly influenced by geographic and idiosyncratic factors, thus lending support to our identification strategy.

Career consequences: Among our respondents, there was a sense that confidential data were "incredibly good" for the careers of the researchers who had access to them (interview S94). Access to Census

data "opens up a lot of research questions" (interview S12) that translate into high-quality publications. Interestingly, our respondents also highlighted that the local availability of Census data was likely to have heterogeneous benefits for different types of scholars. Most people agreed that the lengthy and uncertain approval process to work in an FSRDC posed risks for junior faculty. The whole process requires researchers with a tenure clock to be "a little bit careful and a little bit lucky" (interview S12). However, working with these data is perceived to be an upfront investment that can offer large payoffs throughout a career, especially for researchers that started working in an FSRDC during their PhD (interviews S12, T14, S94). The following quote summarizes well the point:

> "I think the problem with being a junior professor doing that, it just takes a long time to get something done. So it'd be dangerous as a first year assistant professor to say, 'I'm going to start using an RDC and I'm going to get tenure' because that may not work out so well, but if you worked in it as a grad student and you've already got stuff going on that can be pretty, pretty effective." (interview D88)

Awareness of data available in FSRDCs: A theme that emerged from our conversations was that even relatively small geographic barriers are a major hindrance to using an FSRDC. Several administrators indicated that, in their experience, a one-hour commute is enough to discourage a researcher from ever applying to the data (interviews P39, S94, T52). However, even among researchers in proximity to a data center, a crucial factor determining the diffusion of Census data is awareness about their availability and research potential. For instance, administrators of the FSRDC often engage in outreach activities to inform prospective users that "there's all this great data" (interview S12). But the main determinant of the diffusion of confidential data seems to be exposure to research that uses them, because researchers might be inspired after hearing about their colleagues' research. One economist, who happened not to be affiliated with the economics department, told us that many colleagues from her department started using Census data because of exposure to her work (interview S94). Another FSRDC director emphasized the importance of word of mouth to spread the impact of local FSRDCs (interview D88).

Diffusion of other confidential databases: A few respondents noticed that submitting a project to the Census Bureau ultimately presented a risk-reward trade-off. In certain cases, the procedural uncertainty around the approval process might not be worth it: "the trade-off goes the other way where you say, well, I could be a little bit better with the RDC data, but do I really wanna wait nine months?" (interview S12). Moreover, there is an increasingly large number of comparable microdata from foreign countries available to U.S.-based researchers. One researcher even suggested that the hurdles in using Census data might have spurred the diffusion of foreign microdata (interview F72). In the words of one respondent:

> "I think at the moment there is this issue which people in the RDC community are sort of talking

about, which is that the U.S. data is a lot harder to work with than German data or Dutch data or whatever. [...] And there's a sense in the U.S. perhaps that [...] the data is very good and those who are working with the data can do good stuff. But there's a lot of hurdles to get to it and it's not clear that the data is really even necessarily as good or better than some of the European sorts of data. [...] If you have to look at the U.S., then it is the best data available. But if you just want to test the theory and you don't really care whether it's for the U.S. or Germany or France, the U.S. data may not be quite as good." (interview D88)

Interestingly, researchers have long feared that the United States might be losing the edge in applied research due to the limited access to administrative microdata (Card et al., 2010). Our interviews suggest that researchers might have pragmatically shifted to greater use of foreign administrative data with lower barriers to access (interviews F72, D88).

A.2 Case Study

In 2009, Chang-Tai Hsieh and Peter Klenow published the already-classic paper "Misallocation and manufacturing TFP in China and India" on the *Quarterly Journal of Economics*. At the end of 2022, the paper was already counting over 5,900 citations on Google Scholar, corresponding to over 420 cites every year. Moreover, the paper also proved extremely influential in the policy debate (Figure A1). According to the data collected by Altmetric.com, over 74 policy documents have cited this paper up to 2021.

Figure A1: Impact of the Misallocation Paper by Hsieh and Klenow (2009)

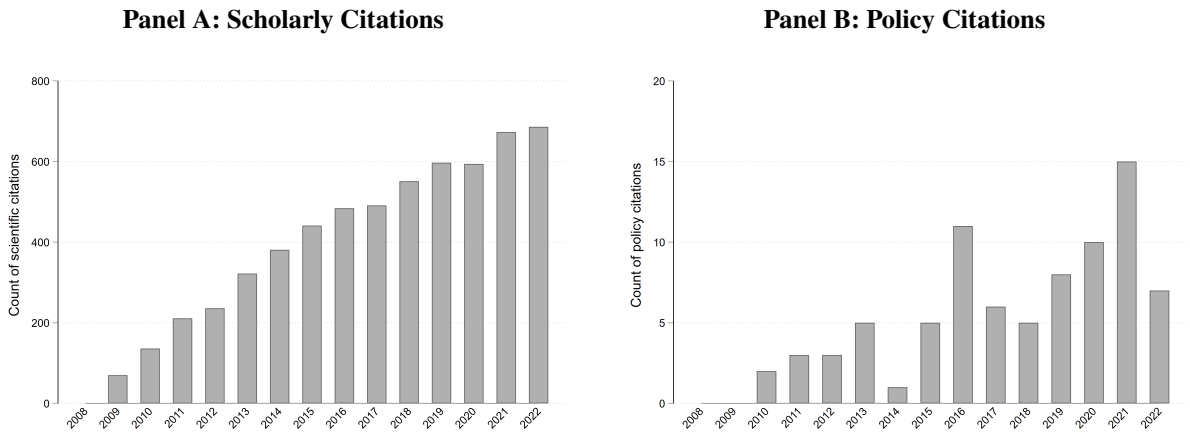

Note: This figure shows the scholarly and policy impact of the Misallocation paper by Hsieh and Klenow (2009). Panel A reports the count of scholarly citations received by the Misallocation paper, using data from Google Scholar. Panel B reports the count of policy citations received by the Misallocation paper, using data from Altmetric.com. See text for more details.

The starting point of the paper is that firm heterogeneity and the allocation of resources across firms play a key role in determining aggregate productivity. If the factors of production are not allocated to their most efficient use, that is to the most productive firms, than aggregate productivity and welfare are reduced. Hsieh

and Klenow (2009) showed that the extent of misallocation could be estimated from firm-level microdata. Their paper first develops a model of monopolistic competition with heterogeneous firms, which is then employed to measure the contribution of resource misallocation to aggregate manufacturing productivity in China and India versus the United States. Their results are striking: aggregate productivity would grow up to 50% in China and up to 60% in India should those countries re-allocate capital and labor similarly to what is observed in the U.S. Since its publication, this paper has become the backbone of a fertile strand of empirical and theoretical research (Hopenhayn, 2014).

The paper by Hsieh and Klenow (2009) was carried out using confidential microdata available only at the U.S. Census Bureau (Figure A2). We gathered additional information about this paper from U.S. Census Bureau's records. The misallocation project started in 2006 at the Berkeley FSRDC, where Hsieh was affiliated at the time. In the abstract of the project submitted to the CES, the authors delineated their objective of developing a new methodology to "help shed light on the underlying sources of productivity differences". To do so, they asked access to confidential U.S. company microdata from the Census of Manufactures (from 1963, 1967, 1972, 1977, 1982, 1987, 1992, 1997, and 2002) and the Annual Survey of Manufacturers (1973-2001). The granularity of the establishment-level data collected by the U.S. Census Bureau was crucial to study the misallocation of production factors across firms. Indeed, Hsieh told us that it would not have been possible doing the same paper without Census data (Hsieh, personal communication, 21st December 2020).

Figure A2: Footnote from the Hsieh and Klenow (2009) Misallocation paper

*We are indebted to Ryoji Hiraguchi and Romans Pancs for phenomenal research assistance, and to seminar participants, referees, and the editors for comments. We gratefully acknowledge the financial support of the Kauffman Foundation. Hsieh thanks the Alfred P. Sloan Foundation and Klenow thanks SIEPR for financial support. The research in this paper on U.S. manufacturing was conducted while the authors were Special Sworn Status researchers of the U.S. Census Bureau at the California Census Research Data Center at UC Berkeley. Research results and conclusions expressed are those of the authors and do not necessarily reflect the views of the Census Bureau. This paper has been screened to ensure that no confidential data are revealed. *chsieh@chicagobooth.edu, pete@klenow.net*.

Note: The figure reports the footnote of the paper by Hsieh and Klenow (2009), highlighting their usage of the Berkeley FSRDC.

We asked to both Hsieh and Klenow about the origins of their collaboration, and how much the need to physically access Census data shaped their work. The datawork for the Misallocation paper was mostly supervised by Hsieh, who had an easier access to the Berkeley FSRDC. Indeed, the physical distance between Berkeley and Stanford was a formidable barrier to the use of confidential Census data: as noted by Klenow, "if I hadn't had a co-author at Berkeley, I don't think I would have started the Misallocation paper" (Klenow, personal communication, 8th December 2020). This is remarkable given that the campuses of UC Berkeley and Stanford are relatively close-by, just about 40 miles apart, which was the reason that prevented Stanford

from obtaining its own FSRDC for many years.

Before starting the project, the two authors had been thinking about the idea and potential methodology, but needed data for empirical validation. Then, one day Hsieh found out about the data available at the Berkeley FSRDC from one graduate student who was using the same data for her dissertations. This detail highlights how local FSRDCs can result in knowledge spillovers that substantially alter research trajectories. Interestingly, even if the Berkeley FSRDC had been in operation since 1999, Hsieh became aware of the potential of Census data for his own research only after seeing the same data used by someone else.

In 2010, Stanford was eventually allowed to establish a branch of the California FSRDC on its campus. The new FSRDC branch drastically reduced the geographical barriers faced by Stanford researchers. In the official press release of the opening, Nick Bloom was quoted saying that "having to go to Berkeley was an immense waste of time"[3]. The same point emerged in our interviews, with a researcher remarking about the long commute between Stanford and Berkeley (interview S94).

B Additional Data Details

The main data source we leverage is *EconLit*, a proprietary database of economic scholarship curated by the American Economic Association. Our version of EconLit includes 839,513 scientific articles published in 1856 journals between 1990 and 2019 inclusive. Unfortunately, EconLit lacks unique author identifiers, preventing from aggregating the publication data into an author-year level panel. To reconstruct authors' publication records we need to disambiguate the author names and affiliations, a complex and time-consuming task that we implement in several steps summarized in Figure B1.

[3]Source: `https://news.stanford.edu/news/2010/february1/census-data-center-020210.html`

Figure B1: Steps Followed for Panel Construction

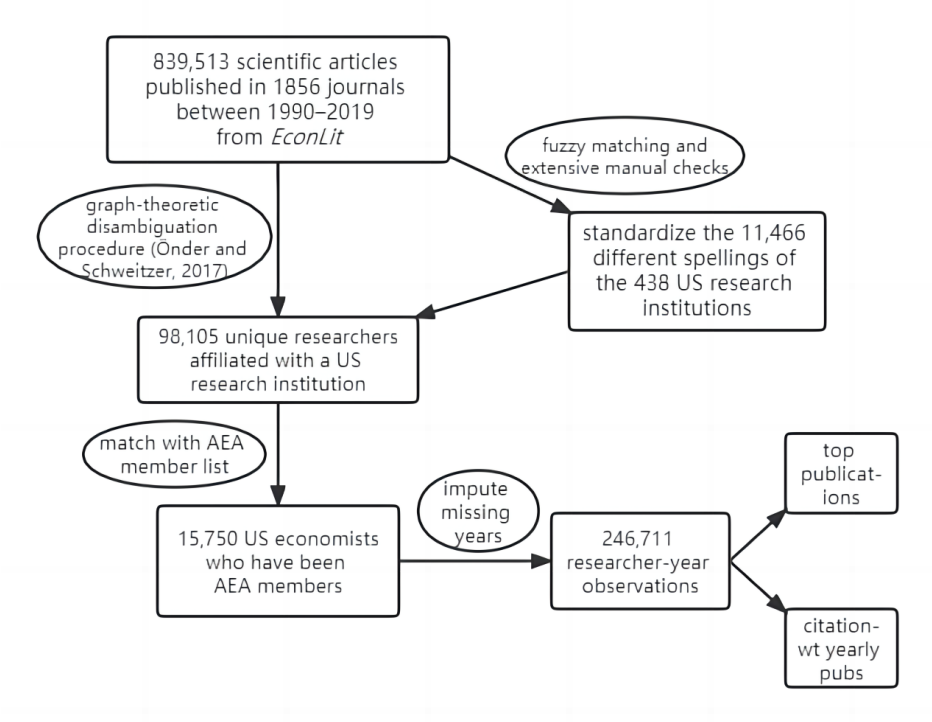

Note: This figure provides an illustration of the steps we took to construct our panel of U.S.-based publishing economists. See text for more details.

B.1 Data Disambiguation

We disambiguate the name of researchers appearing in EconLit following three steps. First, we replace all non-English characters (e.g., "ó" is replaced by "o") and transform the most common name abbreviations into standardized names (e.g., "Ted" is replaced by "Edward"). Second, we apply the disambiguation procedure developed by Önder and Schweitzer (2017). The procedure employs a graph-theoretic approach that follows a hierarchical process. After identifying the set of name entries with identical surnames, the algorithm constructs a graph of the relationships of all the corresponding first names to each other. Names associated with a certain surname can be identical, different, or subsets of each other. For example, the first name "Michael J." is identical to "Michael J.", it is different from "Tom", and it is a subset of "Michael". The algorithm classifies "Michael J." and "Michael" as the same person if no other "Michael *x.*", with *x.* different from J., appears in the data. This approach is equivalent to assuming that the combination of first, middle, and last names uniquely identify each economist (Card et al., 2022), while at the same time being conservative in assigning ambiguous names that lack a clear middle name. Third, we performed extensive manual checks and corrected several misclassifications due to either misspelling in EconLit (e.g., "Tabellini, Gudio" instead of "Tabellini, Guido") or to ambiguous names (e.g., "David Levine", that could refer to David I. Levine or David K. Levine). This three-step procedure results in a database consisting of 434,938 unique researchers from the 552,570 names originally appearing in EconLit.

Next, we standardize the names of the 178,798 affiliations appearing in EconLit. This step is necessary to pin down researchers' treatment status through their affiliation, as well as to restrict the sample to U.S.-affiliated economists who are ever at risk of being co-located to an FSRDC. We begin with a list of all research universities in the United States taken from the 2018 Carnegie Classification of Institutions of Higher Education.[4] In particular, we consider all doctoral universities (corresponding to the codes 15, 16 and 17 in the Carnegie Classification), to which we add the main institutions active in economic research (such as the IMF, RAND Corporation, World Bank, and all the regional FED offices). The result is a list of 438 universities and research centers, which we merge with the EconLit record via fuzzy string matching. This is done by employing string partial ratio similarity to consider the information included in the affiliation word ordering (e.g., to distinguish "University of Washington" from "Washington University"). We retain all matches achieving a partial ratio similarity score equal to or greater than 90 over 100, and we manually check them.[5] The result is a list of 11,466 different spellings of the 438 U.S. research institutions appearing in the Carnegie Classification.

B.2 Panel Construction

After the disambiguation of author names and affiliations, we can use bibliographic data to construct an author-year level database with an annual record for each economist who has published at least one paper in the journals included in EconLit (Moed et al., 2013). Out of the 434,938 unique researchers in our disambiguated data, we retain 98,105 scholars affiliated with a U.S. research institution for at least one year in our sample period. To further restrict our sample to academic economists, we match these names to nineteen yearly lists of members of the American Economic Association (AEA). This step allows us to avoid confounding effects arising from the inclusion of researchers working in unrelated fields but occasionally publishing in economics outlets. The result is an unbalanced panel of 15,750 publishing economists.

We use this set of authors to derive a panel of 246,711 researcher-year observations by imputing missing years between the first and the last year in which we see a researcher publishing. In practice, we interpolate researcher-year pairs that are missing by assigning a zero in the count of scientific output. As noted by Moretti (2021), this interpolation choice leads to partially conflating the extensive margin (i.e., the probability of publishing at least one paper) with the intensive margin (i.e., the number of papers published in a given year, given a positive number of publications). This issue could meaningfully change the interpretation of estimates in contexts where observable outputs are relatively rare, such as in the case of inventors obtaining a patent (Moretti, 2021). However, this is less of a concern in our case because publishing a paper is a

[4]The Carnegie Classification of Institutions of Higher Education is available online at `https://carnegieclassifications.iu.edu/`.

[5]The rate of false positives for matches achieving a score of 90/100 is around 86%, so we decided to drop affiliations with lower scores and consider them as non-U.S. universities.

more common event than patenting, hence increasing the granularity of our panel and reducing the need of imputations.

Our research design takes advantage of the requirement that U.S. Census confidential data can only be analyzed in FSRDC facilities. In practice, we can approximate data access by measuring the distance between the location where the researcher works and the closest data center. We use the information on university affiliations in scientific articles to pinpoint scholars' location and mobility over time (see Figure B2 for an example). In cases when the affiliation changes in non-consecutive years with gaps in between, we attribute the old affiliation for the first third of missing years and the new affiliation for the remaining two-thirds of years. Fortunately, EconLit has a broad coverage of economics journals, allowing us to record scientific mobility with high precision, even for less prolific authors.

Figure B2: Example of Researcher Mobility from Bibliographic Records

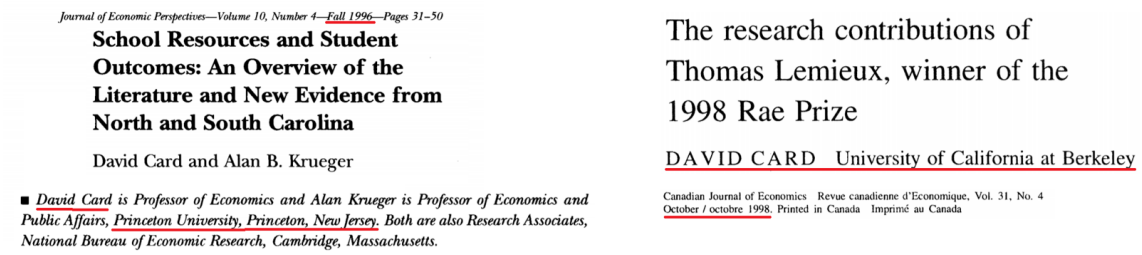
Journal of Economic Perspectives—Volume 10, Number 4—Fall 1996—Pages 31–50

School Resources and Student Outcomes: An Overview of the Literature and New Evidence from North and South Carolina

David Card and Alan B. Krueger

■ *David Card is Professor of Economics and Alan Krueger is Professor of Economics and Public Affairs, Princeton University, Princeton, New Jersey. Both are also Research Associates, National Bureau of Economic Research, Cambridge, Massachusetts.*

The research contributions of Thomas Lemieux, winner of the 1998 Rae Prize

DAVID CARD University of California at Berkeley

Canadian Journal of Economics Revue canadienne d'Economique, Vol. 31, No. 4
October / octobre 1998. Printed in Canada Imprimé au Canada

Note: This figure shows how bibliographic records can be leveraged to construct a panel of publishing economists. In particular, the affiliation details permit to detect mobility events over time and co-location to an FSRDC.

We construct productivity metrics at the researcher-year level by counting the number of yearly published articles. We consider as "top publications" all papers appearing on top five and top field journals as defined by Heckman and Moktan (2020). Figure B3 shows the number of articles recorded in EconLit for each of the main outlets. Next, we augment EconLit by merging each article with the yearly citations count extracted from SSCI/Web of Science. We are able to match 97.2% of our data with the corresponding citation records to construct citation-weighted metrics of academic productivity. We use this information to adjust yearly paper counts by the number of citations received by each article up to five years following publication.

Information on the ranking of economics departments is taken from Kalaitzidakis et al. (2003). We use this information to code seven tiers of universities that we use when estimating models that include a tier-specific time trend. We divide the 100 North American institutions reported by Kalaitzidakis et al. (2003) into five groups that have a roughly similar number of observations in our sample. The top tier include the first five U.S. universities (Harvard, Chicago, MIT, Northwestern, Penn), the second tier the following five (Yale, Princeton, Stanford, UC Berkeley, NYU), the third tier the following ten (from Columbia to Boston U), the forth tier the following eighteen (from Brown to U Virginia), while the fifth tier includes the remaining ones (from Wash U to Stony Brook). The two residual tiers are constituted, respectively, by researchers affiliated

to foreign universities or to U.S.-based research institutions not covered by the ranking.

Figure B3: Count of Articles in the Main Economics Journals (1990-2019)

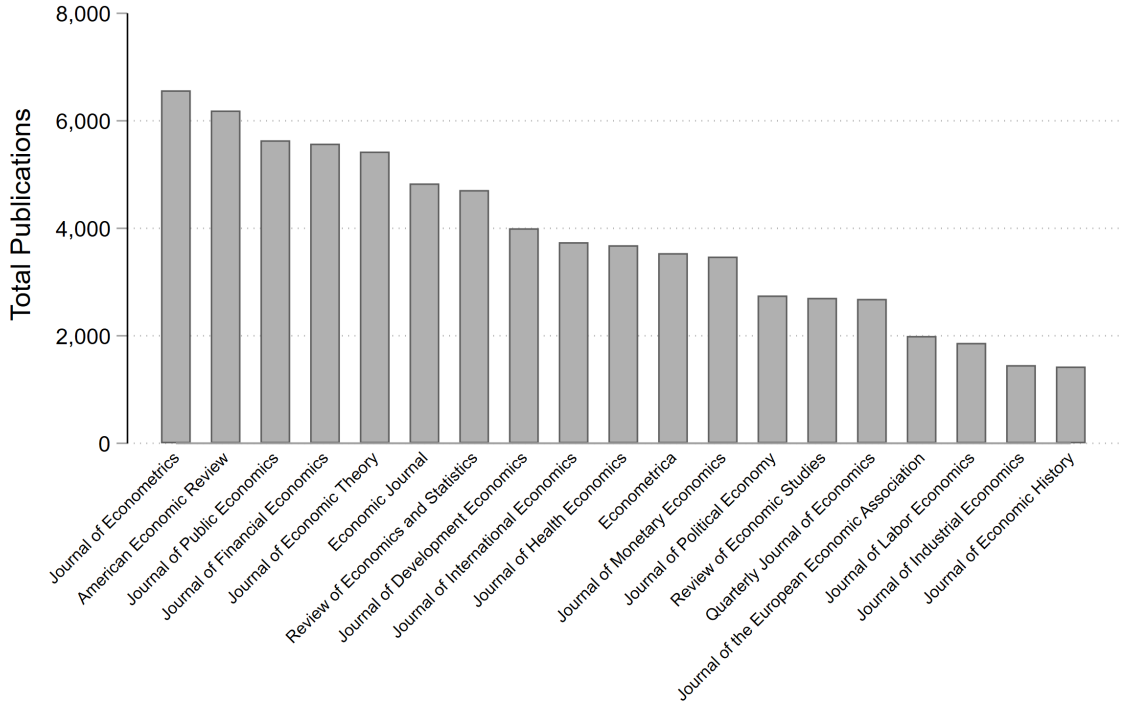

Note: This figure shows the count of papers in the most prestigious economics journals as recorded by *EconLit*. The list of top five and top field journals is from Heckman and Moktan (2020). See text for more details.

B.3 Data Checks

A possible concern of our data construction is that it might lead to measurement errors in our dependent variables if the bibliographic records employed are incomplete. To ensure that EconLit provides good coverage of published research, we randomly drew 15 researchers from our panel and manually searched their publication records. For 14 of them we could find a pdf version of their CV with a list of their scholarly work. We recorded all the publications listed in their CV as the "ground truth" against which we assessed the coverage of EconLit. These researchers collectively published 8 articles in a top five journal and 21 articles in a top field journal during our sample period. Except for one article appearing in the *Journal of Economic History*, all of the other papers were correctly recorded in our data, suggesting that EconLit offers a reliable coverage.

It is also important to note that we rely on published records to infer researchers' careers and mobility patterns. This could lead to measurement error in our independent variable for researchers who gain employment in a city with a local FSRDC, but do not immediately publish articles with the new affiliation. We expect this issue to induce, at most, a downwardly biased coefficient, since it would translate into

some researchers-year observations being wrongly coded as "non-treated" while having local access. Our estimates should be more conservative as a result of this type of measurement error.

However, downward bias could be especially larger for junior economists who might need several years to see their first papers published after completing their studies. To assess the plausibility of this concern, we merged our panel with data on researchers' careers from Sarsons et al. (2021). The result is a subset of 349 economists for which we have the precise date of PhD conferral and tenure decision alongside their publications. We find that the mean time from PhD conferral to first publication is 1.42 years (SD=2.88), with the large majority of researchers publishing their first paper within three years of graduation. Overall, this implies that while junior scholars will enter our panel with some delay, the resulting attenuation bias in our estimates should be fairly small.[6]

C FSRDC: Measuring Diffusion and Impact

This Appendix provides details on the diffusion of Federal Statistical Research Data Centers and the measurement of their research impact.

C.1 Access to FSRDCs Over Time

Table C1 shows the list of the 30 Federal Statistical Research Data Centers opened until 2019. The first was opened in Boston in the Census Bureau's Boston Regional Office (Atrostic, 2007). After that, Carnegie Mellon University and the Center for Economic Studies pioneered a new institutional model where the data center would have been located and operated by an academic institution, with the Bureau just keeping an oversight role. This institutional arrangement became the standard model followed by all subsequent FSRDCs. Interestingly enough, the FSRDC at CMU is also the only one that later closed because of low usage. This episode is consistent with the evidence from our interviews: FSRDCs usually open thanks to the leadership of a small number of researchers, but sometimes this might not translate into broader usage patterns. Appendix Figure G1 confirms qualitative evidence suggesting that the pattern of FSRDC openings did not explicitly prioritize higher-status universities.

Our main analysis explores the dynamics of individual productivity after gaining access to an FSRDC. We define as "treated" researchers working in a city with an active FSRDC. In the Appendix Table G9 we show the robustness of our main results with alternative definitions of the treatment, such as being affiliated with the same institution or being affiliated with an FSRDC consortium member. Alternatively, one could define the treatment status based on the actual distance between the institutional affiliation listed by researchers on

[6]The mean time from PhD conferral to tenure is 7.13 years (SD=2.65), similar to the estimate of Heckman and Moktan (2020). Since the first publication appears on average one and a half years after the PhD, this validates our choice of using a cutoff of 7 years after the first publication to identify economists likely to be untenured.

Table C1: Federal Statistical Research Data Centers Opened during our Sample Period

Year opened	Institution of location	City	State
1994	NBER Cambridge	Boston	Massachussets
1997*	Carnegie Mellon U	Pittsburgh	Pennsylvania
1999	UC Berkeley	Berkeley	California
1999	UCLA	Los Angeles	California
2000	Duke	Durham	North Carolina
2002	U Michigan	Ann Arbor	Michigan
2003	Chicago FED	Chicago	Illinois
2004	Cornell	Ithaca	New York
2006	Baruch College-CUNY	New York	New York
2010	U Minnesota	Minneapolis	Minnesota
2010	Stanford	Stanford	California
2011	Atlanta FED	Atlanta	Georgia
2012	Texas A&M U	College Station	Texas
2012	U Washington	Seattle	Washington
2014	UC Irvine	Irvine	California
2014	USC	Los Angeles	California
2014	Penn State	University Park	Pennsylvania
2015	U Missouri	Columbia	Missouri
2015	U Nebraska-Lincoln	Lincoln	Nebraska
2015	UW Madison	Madison	Wisconsin
2015	Yale	New Haven	Connecticut
2016	U Maryland	College Park	Maryland
2016	FED Kansas City	Kansas City	Missouri
2017	UT Austin	Austin	Texas
2017	U Colorado Boulder	Boulder	Colorado
2017	U Kentucky	Lexington	Kentucky
2017	Philadelphia FED	Philadelphia	Pennsylvania
2017	Georgetown U	Washington	DC
2018	Dallas FED	Dallas	Texas
2019	Federal Reserve Board	Washington	DC

Note: The FSRDC at Carnegie Mellon University closed in 2004. Additional FSRDCs opened at the University of Utah (2020), University of Illinois Urbana-Champaign (2020), and University of Florida (2022) are outside of our sampling period.

their papers and the location of the nearest FSRDC. Appendix Figure G8 presents elasticities of our main results at various distances, showing results consistent with our main specifications.

Importantly, researchers can gain access to Census data in two ways: either when a new data center opens in the city of their current affiliation, or when they change employer and move to a new city with an active FSRDC. Figure C1 provides a breakdown of the researchers in our sample by the modality in which they received access (see also Appendix Figure G2). We might worry that the individuals who advocated opening a local FSRDC might have an interest in using it, potentially inflating our estimates. Appendix Table G4 shows that our results are robust to the exclusion of the researchers that appeared as Principal Investigators (PIs) in the NSF grants that led to establishing the data centers. More in general, the institutions hosting the data center on their premises might be the ones that had the most to benefit from it, which would lead to upwardly biased estimates. Appendix Table G7 shows robustness of our results if we estimate our models only on the universities that gained vicarious access by being in the same city as the hosting institution.

An alternative concern is that researchers who gain access by moving to a city with an active data center might do so precisely because they want to work on confidential data. This endogenous geographical sorting could bias our estimates upward. To rule out this concern, we repeated the main analyses excluding all scholars who become co-located to a data center after changing affiliation. All the main results are robust to this test (Appendix Table G5). In addition, qualitative evidence from our interviews suggests that this concern is unlikely to hold in practice. While the presence of a local FSRDC could in principle convince a scholar to accept a job offer, this is usually true only for previous users whose research agenda crucially hinges on having an FSRDC nearby (interview S94).

Figure C1: U.S.-Based Economists by Modality of FSRDC Access

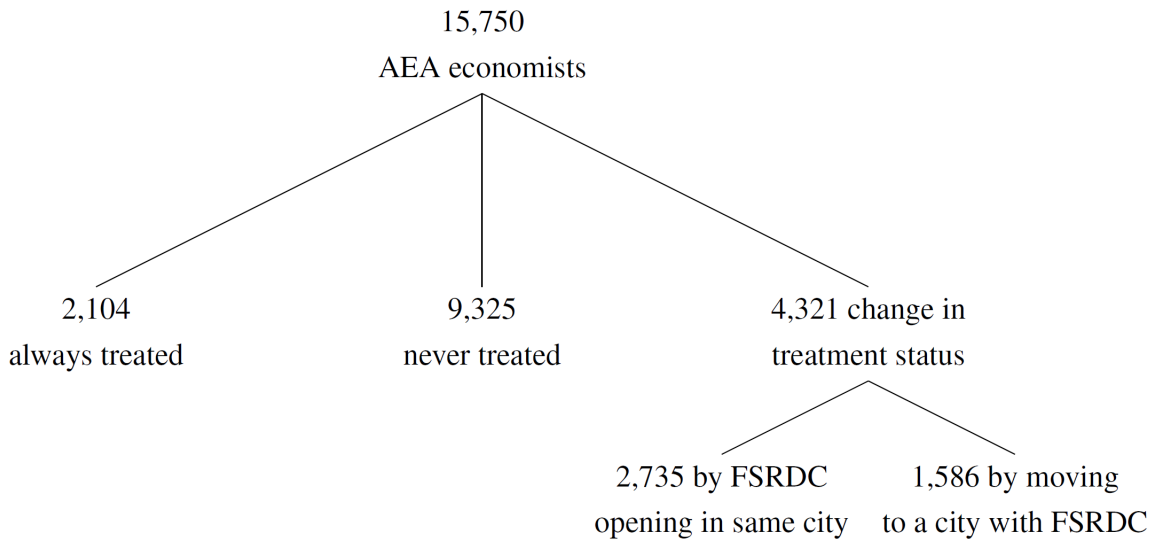

Note: This figure summarizes the treatment status of all publishing economists included in our panel. See text for more details.

C.2 Articles using Census Data

We assembled a novel dataset of all articles that *directly* employ restricted-access microdata accessible only in U.S. Census' secure facilities. Since there is no official bibliographic record of research using Census data, we carefully searched for them using several complementary strategies.

As a starting point, we exploited the fact that papers using U.S. Census confidential data are expected to indicate it clearly in the acknowledgment of the published version of the paper. We started by collecting all the most commonly used sentences appearing in the acknowledgments of a sample of FSRDC papers (such as "Census Research Data Center", "do not reflect the views of the Census Bureau", and "Special Sworn Status researchers"). Then we searched for them in the main databases of published research, Web of Science and Scopus, which recently started collecting the acknowledgment sections of journal articles. However, we found that many papers do not report the standard disclaimers required by the Census Bureau.

We tried to overcome this limitation with additional searches in databases that allow full-text searches, such as JSTOR,[7] Google Scholar, and the NBER working paper repository.

We further expanded our search by exploiting the fact that projects approved by the FSRDC are expected to submit a final working paper to the Center for Economic Studies of the U.S. Census Bureau for online publication.[8] We collected the metadata of 1,081 working papers and matched them to EconLit through a combination of fuzzy title matching and extensive manual checks. Overall, just about half of these papers were ever published, and only 455 papers could be linked to the corresponding EconLit record.[9] The high share of CES working papers that never get published suggests that output-based assessments of FSRDCs will understate the actual use of the network by academics (see Section C.3 for some alternative approaches).

Finally, we asked for access to all approved FSRDC projects via Freedom of Information Act (FOIA) request to the Census Bureau (FOIA ID No. DOC-CEN-2020-001640). Notably, the results of our request have been published online for the benefit of everyone interested in tracking the use of Census Bureau's administrative data.[10] We manually searched the publication records of each researcher whose projects were approved to be carried out in an FSRDC. In total, we were able to find a total of 861 papers published in peer-reviewed journals that could be matched with EconLit. Once we restrict our sample to U.S. researchers affiliated with the AEA, the final sample of papers employing Census data consists of 587 articles written by 509 economics researchers.[11]

C.3 Other Measures of FSRDC Impact

One potential shortcoming of our measures of FSRDC impact is that they are based on research in peer-reviewed economics journals. However, we might be understating the actual diffusion of Census data if many projects are not published. This would be a concern if papers using FSRDC have a lower propensity to be published, perhaps because they constitute riskier research with a wider variance in scientific quality. To sidestep this type of concern, we carry out robustness tests where the measure of data adoption is given by whether a scholar has a project approved to be carried out in an FSRDC. We obtained this information thanks to a FOIA request (FOIA ID No. DOC-CEN-2020-001640) that resulted in a list with details about all projects approved by the U.S. Census Bureau. We supplement these data with similar records from the

[7] The main limitations of JSTOR are that its coverage is unreliable for more recent years and that it does not encompass journals published by Elsevier.

[8] The papers are available online at the following link: `https://ideas.repec.org/s/cen/wpaper.html`.

[9] Published papers that do not appear in EconLit have either appeared in journals not covered (e.g., the *Strategic Management Journal*) or as a book chapter (mostly in NBER-edited books).

[10] `https://www.census.gov/about/adrm/fsrdc/about/ongoing-projects.html`

[11] During our data construction, we also found several articles from other disciplines, such as sociology, demography, and health policy. We excluded them because they are outside the scope of this analysis. However, it must be noted that a complete policy evaluation of the scientific impact of the FSRDC network would require expanding the focus outside economics sciences.

other agencies participating in the FSRDC program, all obtained with additional FOIA requests.[12] These data give us an upstream measure of data usage that is not dependent on the project's outcomes.

Moreover, projects carried out using Census confidential data must be submitted as working papers to the Center for Economic Studies once completed. We use the count of CES working papers up to 2019 inclusive as an additional measure of data usage. Interestingly, only 48.4% of such CES working papers were published as of June 2020. This figure is close to recent estimates that around 50-74% of NBER working papers are eventually published (Baumann and Wohlrabe, 2020; Lusher et al., 2021). The large majority of CES working papers appeared in a peer-reviewed journal included in EconLit, with just 13% of them appearing as book chapters or in journals of other disciplines. Again, this number is very close to the finding that 12% of published NBER working papers do not appear in economics journals (Lusher et al., 2021).

C.4 Spillovers from Local Access to FSRDCs

A. Articles citing FSRDC papers: Besides using confidential data available only in an FSRDC, we aim to capture how access to U.S. Census data might also affect research that does not directly use confidential Census data. Local access to Census data might be shaping the topics or questions researchers decide to work on, for instance, by increasing awareness of past research done with such data or being exposed to the findings of colleagues working in the FSRDC. We measure this type of influence by recording which articles build on the findings of the papers written using Census data. This is done by taking the list of papers using Census data (identified with the procedure of Appendix Section C.2) and downloading their references from the Social Sciences Citation Index (SSCI) curated by Web of Science. We exclude papers directly using Census data from the count since they are likely to mechanically cite other FSRDC papers (for instance, in their data section). The result is a set of papers that build directly on the results made possible by the FSRDC network, even if they do not use its confidential microdata.

B. JEL codes: Another approach to measuring spillovers from FSRDC access is to consider its impact on publications that pertain to topics commonly associated with using Census data. Researchers working in areas such as labor, productivity studies, trade, and environmental economics should benefit more from research exposure based on Census data. We capture similarity in research topics using the paper-level JEL codes recorded in EconLit. This is conceptually similar to the approach of Azoulay et al. (2019), who leverage the keyword assigned to all life-science publications by the National Library of Medicine to define research that is proximate in topical space.

[12] We obtained data from the AHRQ (FOIA ID No. 2021-00311-FOIA-PHS), the BEA (FOIA ID No. DOC-BEA-2021-000222), the BLS (FOIA ID No. 2021-F-00826) and the NCHS (FOIA ID No. 21-00102-FOIA).

Figure C2: List of JEL Codes Associated to Papers using Confidential Census Data

(i) JELs most representative of FSRDC research

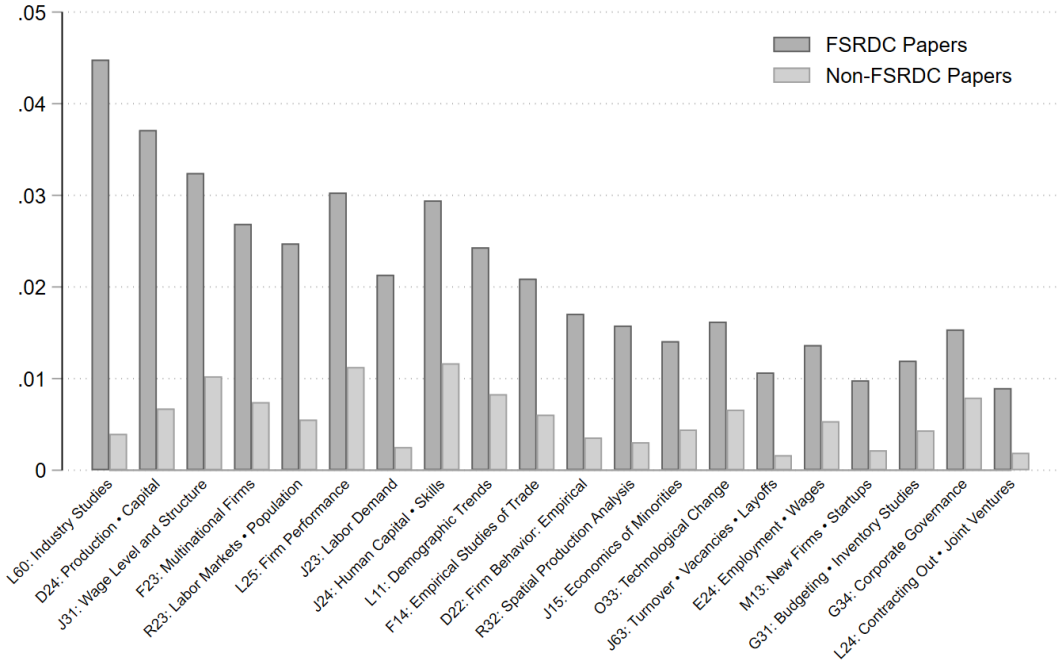

(ii) JELs least representative of FSRDC research

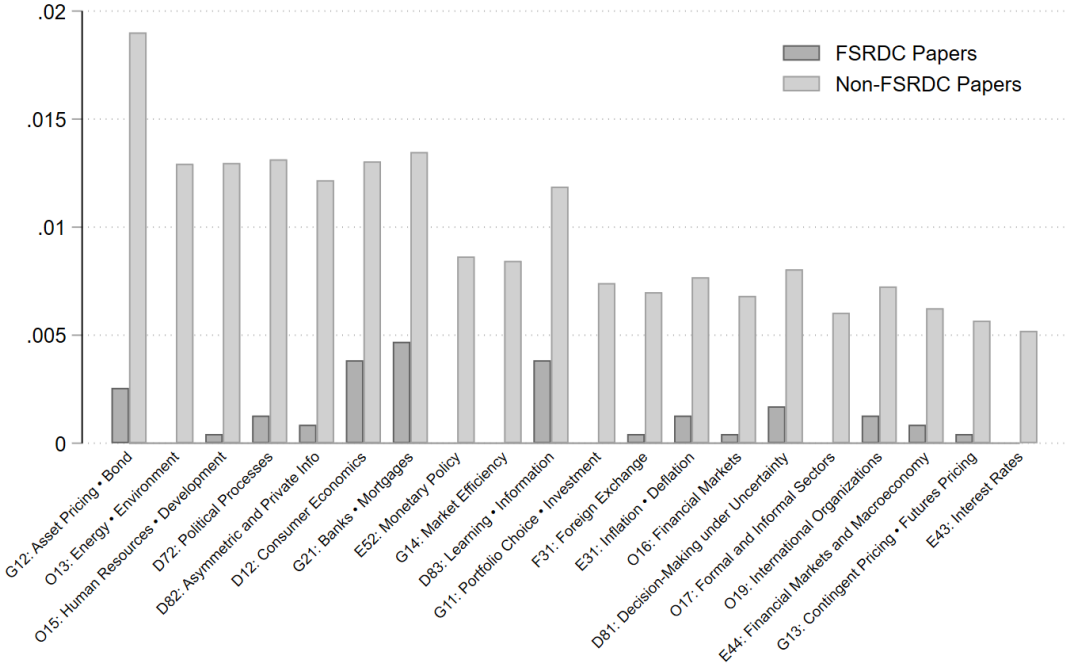

Note: This figure shows the JEL codes that are most distinctive of research carried out with and without confidential Census data. Each bar represents the share of papers that include that JEL code for each of the two groups of papers. Panel (a) reports the 20 JEL codes with the largest differences in frequency between papers written in an FSRDC and the rest of the sample. Panel (b) reports the 20 JEL codes with the largest negative differences in frequency between papers written in an FSRDC and the rest of the sample. JEL codes are sorted from largest to smallest difference in frequency. See text for more details.

We use the list of FSRDC papers to find the JEL codes most likely to be associated with confidential Census data. Figure C2 displays in Panel (a) the 20 JEL codes with the largest differences in frequency between papers written in an FSRDC and the rest of the sample. The figure confirms anecdotal evidence that underscores the considerable potential of these data for firm-level analyses. Other common topics include studying labor markets, demographic trends, and international trade. Likewise, Panel (b) of Figure C2 shows that research carried out in Federal Statistical Research Data Centers is much less likely to encompass topics like financial markets, development, and monetary economics, as well as theoretical investigations of market efficiency and uncertainty. Appendix Table G12 shows the robustness of our main results to alternative definitions of which JEL codes are the most representative of research enabled by access to Census data.

C. Topical focus of research: The main shortcoming of relying on JEL codes is that they are assigned by the authors of the paper. Since there are no explicit rules on defining the proper JEL codes, differences in researchers' approach when choosing JEL codes might introduce measurement errors. We sidestep this issue by developing a data-driven categorization of research fields to characterize researchers' research trajectories (Furman and Teodoridis, 2020). Our approach consists in using topic modeling, a popular method for text analysis (Gentzkow et al., 2019). The main advantage of this choice is its ability to identify similarities between bodies of text without being influenced by authors' (potentially strategic) choice of JEL codes.

More in detail, we merge titles and abstracts of all the papers written by a researcher in a given year to generate researcher-year level bodies of text. Then, we follow common practices and pre-process the text using the "term frequency–inverse document frequency" (tf–idf) approach. In this way, we can downplay the impact of common words while giving more weight to rare words when computing similarities between

Table C2: Latent Dirichlet Allocation (LDA) Topics and Their Most Salient Words

Topic Label	Most Salient Words				
topic 1 (innovation)	network	innovation	research	firm	social
topic 2 (institutions)	patent	r&d	election	voter	vote
topic 3 (trade)	trade	growth	export	firm	country
topic 4 (energy)	energy	electricity	policy	cost	efficiency
topic 5 (migrations)	quantile	immigrant	exit	immigration	alliance
topic 6 (corporate governance)	governance	seller	buyer	culture	girl
topic 7 (finance)	firm	bank	risk	stock	investor
topic 8 (welfare)	food	jump	maternal	nutrition	childhood
topic 9 (corporate finance)	merger	ambiguity	cultural	director	conservation
topic 10 (geography)	friction	venture	gold	building	corruption
topic 11 (econometrics)	model	forecast	method	estimator	test
topic 12 (health/education)	health	child	household	woman	income
topic 13 (natural resources)	renewable	sovereign	nuclear	screening	fine
topic 14 (macroeconomics)	rate	monetary	shock	tax	price
topic 15 (labor)	wage	worker	job	employment	labor
topic 16 (statistics)	tail	dependence	nonparametric	bootstrap	covariate
topic 17 (environmental)	emission	oil	carbon	climate	fuel
topic 18 (microeconomics)	game	equilibrium	preference	agent	information
topic 19 (accounting)	audit	dataset	reporting	auditor	transparency
topic 20 (healthcare)	patient	hospital	ecosystem	rating	student

Note: The table reports the five words most closely related to each LDA topic. After manual examination, we coded a plausible label to capture the content of each topic. See text for more details.

Table C3: Keywords Used to Tag Articles' Research Design and Data

Category	Keywords (from Currie et al., 2020)
Administrative data	"administrative data" "admin data" "administrative-data" "admin-data" "administrative record" "admin record" "administrative regist" "admin regist" "registry data" "register data" "confidential data" "restricted data"
Survey data	"survey data" "phone survey" "survey administered" "household survey"
Quasi-experimental methods	"difference in diff" "differenceindiff" "differences in diff" "differencesindiff" "d-in-d" "diff in diff" "diffindiff" "event stud" "eventstud" "staggered adoption" "regression discontinuit" "regressiondiscontinuit" "regression kink" "regressionkink" "rd resign" "rdresign" "rd estimat" "rdestimat" "rd model" "rdmodel" "rd regression" "rdregression" "rd coefficient" "rdcoefficient" "rk design" "rkdesign" "rdd" "rkd" "instrumental variable" "instrumentalvariable" "two stage least squares" "twostage least squares" "2sls" "tsls" "valid instrument" "exogenous instrument" "iv estimat" "ivestimat" "iv regression" "ivregression" "iv strateg" "ivstrateg" "we instrument" "i instrument" "exclusion restriction" "paper instruments" "weak first stage" "simulated instrument"
Experimental methods	"randomized controlled trial" "randomized field experiment" "randomized controlled experiment" "randomised controlled trial" "randomised control trial" "randomised field experiment" "randomised controlled experiment" "social experiment" "rct" "laboratory experiment" "lab experiment" "public good game" "public goods game" "ztree" "orsee" "showup fee" "laboratory participant" "lab participant" "pre analysis plan" "preanalysis plan" "pre registered" "preregistered" "preregistration" "dictator game" "ultimatum game" "trust game"

Note: This table presents the keywords used to tag each paper as belonging to one of the following non-mutually exclusive categories: administrative data, survey data, quasi-experimental methods, experimental methods. The list of keywords is taken from Currie et al. (2020).

bodies of text. Next, we use an off-the-shelf algorithm known as latent Dirichlet allocation (LDA) to inductively derive 20 topics made of words that tend to co-appear in abstracts and titles. Each term is assigned a probability distribution over these clusters of words, with probabilities being higher when the term is more representative of that topic. Manual audits of the terms most representative of each data-driven cluster confirm that the procedure defines meaningful research topics (Table C2).

D. Text-based measures of research design: Finally, we compile data to assess potential spillovers on the adoption of empirical methods and administrative data from other sources. Following Currie et al. (2020), we use a series of regular expression searches to find keywords in the titles and abstracts that identify the use of a certain method (e.g., difference-in-differences) or type of data (e.g., survey data). The list of n-grams we use is presented in Table C3 and directly taken from Appendix A of Currie et al. (2020). These keywords allow us to tag each paper as belonging to one of the following non-mutually exclusive categories: administrative data, survey data, quasi-experimental methods, and experimental methods.

Table C4 explores what keywords are more likely to be associated with papers using Census data. Confirming our priors, these papers are eight times more likely to mention the use of confidential data or administrative records, while the same is not true for words denoting the use of traditional surveys compiled for research purposes. Further, Table C4 matches our intuition that research carried out in an FSRDC does not employ experimental research design (such as RCTs or behavioral laboratory experiments). The inclusion of year fixed effects helps alleviating the concern that these findings are just the reflection of a wider "credibility revolution" in the years of FSRDC activity.

Table C4: Characteristics of Papers Using Census data

	Administrative data	Survey data	Quasi-exp. methods	Experimental methods
FSRDC Paper	0.0384***	0.0105	0.00126	-0.0130***
	(0.00961)	(0.01117)	(0.00798)	(0.00157)
Year FE	YES	YES	YES	YES
N	188,181	188,181	188,181	188,181
Mean of the Dep Var	0.0046	0.0138	0.0232	0.0111

Note: This table presents estimates from OLS models evaluating the likelihood that papers using Census data mention the keywords listed in Table C3. *, **,*** denote significance at 5%, 1% and 0.1% level respectively. Robust standard errors reported.

D Classification of Research by Methodological Orientation

In this Appendix, we explain the machine learning procedure followed in classifying the methodological orientation of each EconLit record. We then discuss how we infer the specialization of each researcher from her published work and explore the robustness of our results.

D.1 Classification Algorithm

We employ the machine learning classification procedure developed by Angrist et al. (2020) to classify each economics paper as either empirical or theoretical. The algorithm aims to separate research that produces data-based estimates of economically meaningful parameters from purely theoretical or methodological papers. We also classify papers that develop new methods or models but apply them to produce substantively meaningful estimates as empirical. The classifier developed by Angrist et al. (2020) is a logistic ridge regression that fits a dummy variable indicating empirical papers using article titles, journals, JEL codes, publication years, and abstracts as inputs.[13] Angrist et al. (2020) show that the algorithm has an 87% accuracy for entries with the abstract, so we scraped the internet to fetch the abstracts of all articles that do not have it reported in EconLit.

[13]The algorithm is trained on a set of 5,469 hand-classified papers.

Figure D1: Example of Article Classification using Machine Learning

Panel A: Empirical Article

DO ENERGY EFFICIENCY INVESTMENTS DELIVER? EVIDENCE FROM THE WEATHERIZATION ASSISTANCE PROGRAM*

MEREDITH FOWLIE
MICHAEL GREENSTONE
CATHERINE WOLFRAM

A growing number of policies and programs aim to increase investment in energy efficiency, because conventional wisdom suggests that people fail to take up these investments even though they have positive private returns and generate environmental benefits. Many explanations for this energy efficiency gap have been put forward, but there has been surprisingly little field testing of whether the conventional wisdom is correct. This article reports on the results of an experimental evaluation of the nation's largest residential energy efficiency program—the Weatherization Assistance Program—conducted on a sample of approximately 30,000 households in Michigan. The findings suggest that the upfront investment costs are about twice the actual energy savings. Furthermore, the model-projected savings are more than three times the actual savings. Although this might be attributed to the "rebound" effect—when demand for energy end uses increases as a result of greater efficiency—the article fails to find evidence of significantly higher indoor temperatures at weatherized homes. Even when accounting for the broader societal benefits derived from emissions reductions, the costs still substantially outweigh the benefits; the average rate of return is approximately −7.8% annually. *JEL Codes:* Q4, Q48, Q5.

P(Empirical) = 96.95%

Panel B: Theoretical Article

Econometric Theory, **13**, 1997, 467–505. Printed in the United States of America.

GAUSSIAN ESTIMATION OF MIXED-ORDER CONTINUOUS-TIME DYNAMIC MODELS WITH UNOBSERVABLE STOCHASTIC TRENDS FROM MIXED STOCK AND FLOW DATA

A.R. BERGSTROM
University of Essex

This paper develops an algorithm for the exact Gaussian estimation of a mixed-order continuous-time dynamic model, with unobservable stochastic trends, from a sample of mixed stock and flow data. Its application yields exact maximum likelihood estimates when the innovations are Brownian motion and either the model is closed or the exogenous variables are polynomials in time of degree not exceeding two, and it can be expected to yield very good estimates under much more general circumstances. The paper includes detailed formulae for the implementation of the algorithm, when the model comprises a mixture of first- and second-order differential equations and both the endogenous and exogenous variables are a mixture of stocks and flows.

P(Theoretical) = 83.6%

Note: This figure shows the results from the machine learning classification algorithm of Angrist et al. (2020) for two papers in our sample. See text for more details.

This procedure results in a paper-level score (between 0 and 1) that captures the probability that an article is empirical. Figure D1 shows an example of two articles and the results from the machine learning classification algorithm. The algorithm has a remarkable performance also in cases where theoretical papers mention in their abstract or title keywords usually associated with empirical papers, such as "data" or "estimates". We classify as empirical all papers with a predicted score larger than 0.5.[14] Our data show an increasing share of empirical publications over time, from 52.3% in 1990 to 71.3% in 2019 (Figure D2). Relative to the findings of Angrist et al. (2020), we document a slightly larger share of empirical research. Possible reasons for this discrepancy are the better coverage of abstracts in our data (that allow for a more precise classifications) and the fact that we consider in our analyses a broader set of journals.

D.2 Robustness Checks

We define as empiricist all researchers with more than half of their publication output classified as empirical (excluding any paper using FSRDC data). This measure has the advantage of being available for every researcher in our sample of publishing economists. In total, we classify as empiricists 73% of U.S.-based economists in our sample. Figure D3 shows the distribution of the share of empirical articles for the researchers in our sample.

We collected additional data to show the face validity of our classification. First, we adopted a case-control

[14] Unlike Angrist et al. (2020), we do not separately classify econometrics articles into an ad-hoc category.

Figure D2: Publications and Citations to Economics Papers by Methodological Orientation

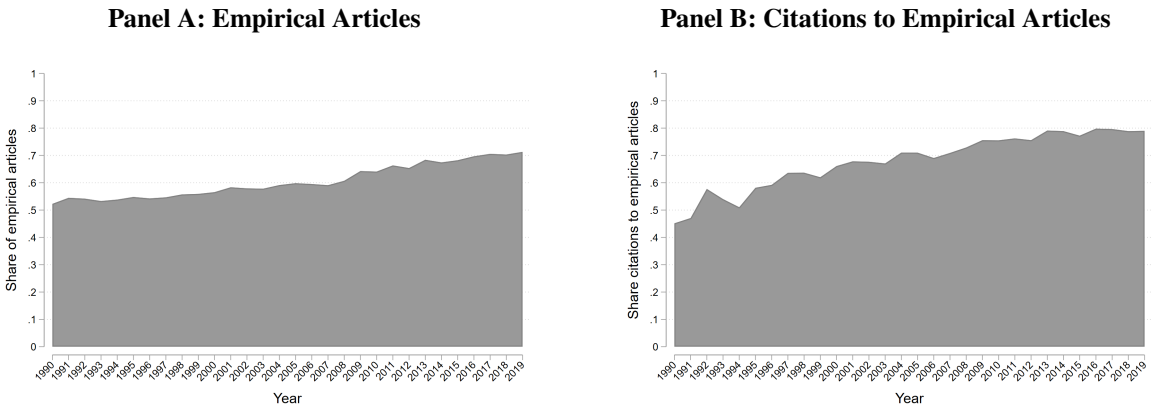

Note: This figure shows the share of empirical articles (Panel A) and citations received by empirical articles (Panel B) over time. See text for more details.

approach and checked the results of our classification for the Editorial Boards of some high-profile economics journals. In particular, we collected data for five journals that publish a broad spectrum of work, from mostly empirical (e.g., *AEJ: Economic Policy*) to mostly theoretical (e.g., *Journal of Economic Theory*). Table D1 presents the total number of Editors reported on the website of each of these journals as of November 2022. We matched these Editors with our sample of U.S. affiliated economists and checked how many of them our procedure classified as applied researchers. Confirming our priors, our procedure classified 98% of *Journal of Economic Theory*'s Editors as theoretical researchers and most components of the Editorial Boards of *AEJ: Applied Economics* and *AEJ: Economic Policy* as empirical researchers. Other journals encompass a greater variety of expertise, with *AEJ: Microeconomics* tilted towards theory-minded Editors and *AEJ: Macroeconomics* showing a preponderance of empirical Editors.

Figure D3: Empirical Articles as a Share of Total Published Works for Researchers in our Sample

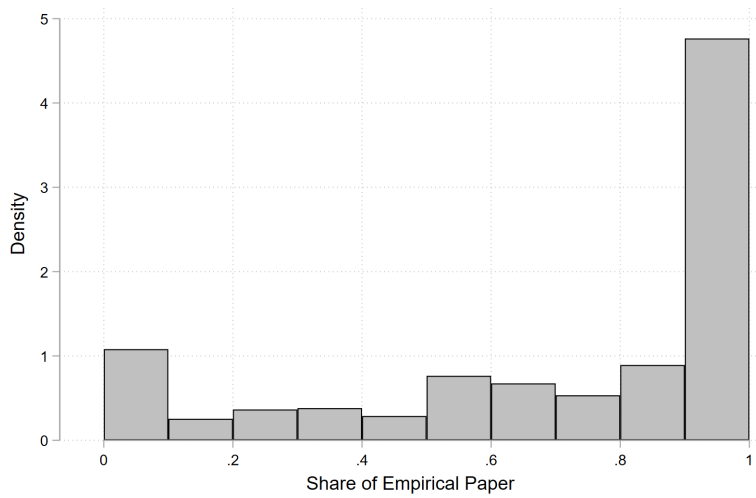

Note: The figure reports the distribution of the share of empirical published work for researchers in our sample. In our main analyses, we classified as empiricist all scholars with a share > 0.5. See text for more details.

Table D1: Methodological Classification for the Editorial Board of Selected Economics Journals

Journal	Number of Editors	Ever U.S.-affiliated	Share Empirical
Journal of Economic Theory	62	47	2.13%
AEJ: Microeconomics	21	19	31.58%
AEJ: Macroeconomics	19	18	66.67%
AEJ: Applied Economics	36	33	96.97%
AEJ: Economic Policy	37	35	97.14%

Note: The Table reports the composition of the Editorial Boards of five economics journals (as of November 2022) and the share of them that our machine learning procedure classified as empirical researchers.

Second, we digitized the lists of North American Ph.D. graduates published yearly in the last issue of the *Journal of Economic Literature*. Figure D4 shows the topics of the doctoral dissertation of the 5,209 researchers we could match to our data, divided by whether we classified the researcher as an empiricist or a theorist. The figure shows that theorists are most likely to have written their doctoral thesis in Microeconomics, Macroeconomics, or Mathematical Methods, while the most popular fields for empiricists are Labor, Health, and Development Economics.

Figure D4: Distribution of PhD Dissertation Fields for Researchers Classified as Empiricists or as Theorists

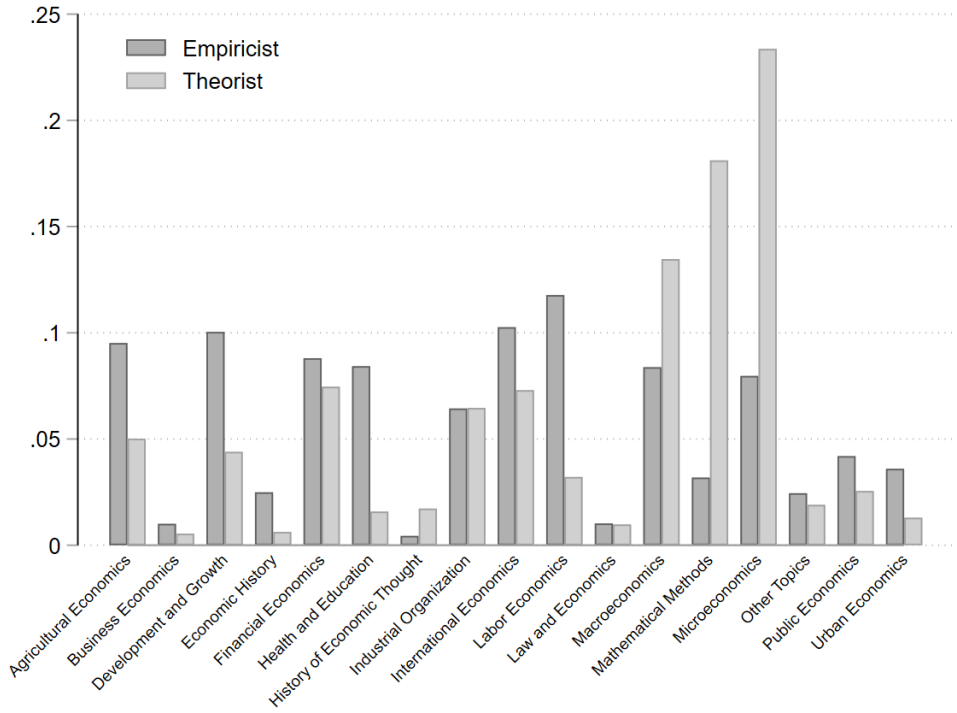

Note: The figure reports the distribution of dissertation topics for each researcher in our sample that we could match to the *Journal of Economic Literature* (JEL) yearly lists of graduates. In total, we matched 5,209 economists, of which 3,876 classified as empiricists and 1,333 as theorists. See text for more details.

Our classification is based on the lifetime count of individual publications. As an alternative, we repeat it considering only articles published before being co-located to an FSRDC. The concordance between the

two approaches is 92.8%, and grows to 95% when looking at researchers for whom we observe at least three publications before they have access to a local FSRDC. This confirms our intuition that researchers rarely change methodological orientation during their career, and, if so, hardly because of gaining access to an FSRDC. We also performed robustness checks to show the stability of our results to different thresholds in defining empirical researchers. In particular, we ran our main regressions with increasingly stringent definitions of empirical researchers. Figure G9 plots the coefficients of these additional regressions. Our results are driven by the most empirically-minded researchers in our sample, with the magnitude of the coefficients monotonically growing when using more stringent thresholds to define empiricists.

E Measuring the Policy Impact of Academic Research

E.1 Altmetric.com Data

Citations from policy documents (or "policy citations" for brevity) are helpful in recording the influence of academic research in sectors outside the ivory tower. Yet, the use of this type of bibliometric data in academic studies is still rare (with a few notable exceptions, such as Yin et al., 2022). For a long time, the major obstacle has been the paucity of databases that reliably record citations from unstructured policy documents and merge them with unique identifiers for scientific articles. More recently, the increasing availability of open data on the scientific bibliome and the possibility of large-scale scraping of policy sources from the internet has enabled the first systematic collections of policy citations.

In this paper, we measure the policy impact of economics articles using the data from Altmetric.com, a company focused on providing evidence of the scientific impact that goes beyond the traditional reliance on paper-to-paper citations. In particular, we leverage the fact that since 2014 Altmetric has been collecting policy documents and merging them with the research output they cite. Altmetric.com scrapes data within policy documents going as back as the 1920s, thus capturing impact that might take several years to materialize. The type of documents tracked range from government documents to white papers of international development organizations, from research institutes to think tanks' reports. To the best of our knowledge, very few papers have used this source of policy document citations to date (Haunschild and Bornmann, 2017).

We merge our publication-level dataset to the Altmetric.com data using the articles' DOI. In total, we find 83,640 citations from policy documents to 28,149 economics articles authored by U.S.-based economists (14.96% of our sample).[15] Figure E1 shows the distribution of the publication year of the articles cited (left panel) and the documents citing them (right panel). Starting with the latter, one notices that the coverage of

[15]We exclude from the Altmetric.com data citations from NBER working papers since they are the working paper version of academic articles and not policy documents.

Figure E1: Year of Publication of the Economics Articles with Policy Citations and of the Citing Policy Documents

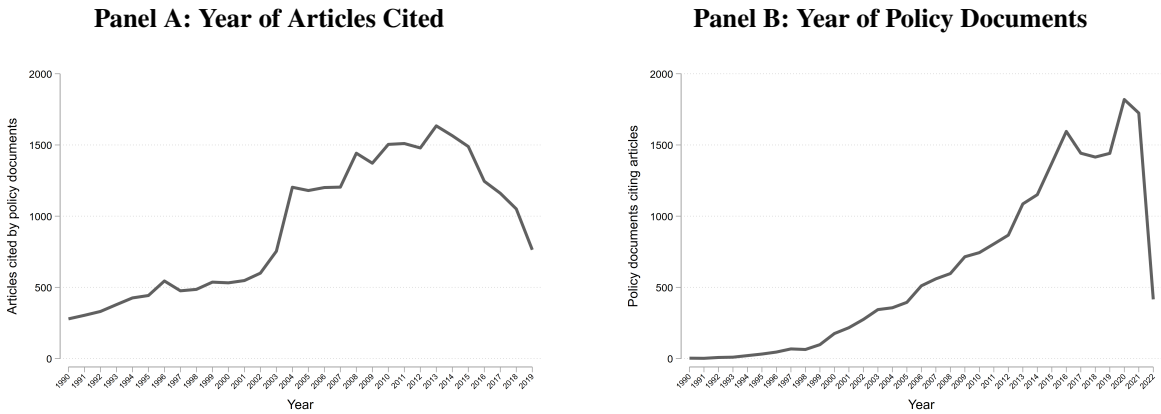

Note: This figure shows the publication year of the economics articles that received at least one policy citation (Panel A) and the publication year of the policy sources that cite them (Panel B). See text for more details.

Altmetric.com is skewed towards recent policy sources. This is likely to explain why the left panel shows that scientific papers published in the mid-2000s have a much higher likelihood of being cited by policy documents. However, insofar as the recording of policy references is not systematically biased within the year of article publication, the inclusion of year fixed effects should take into account the time-varying propensity of receiving a policy cite.

The geographic origin of policy documents is heavily skewed towards Western Countries, as shown by Figure E2. The United States alone account for almost 40% of all the policy citations received by the economics articles in our sample. When focusing on individual sources, one sees many government bodies or international organizations well-known for their policy advocacy (Figure E2). Interestingly, almost a fifth of the cites comes from documents produced by the World Bank. Unreported robustness checks confirm that all our results are unchanged if we drop them from the sample. Figure E3 shows that articles appearing

Figure E2: Origin of Policy Citations to Economic Scholarship

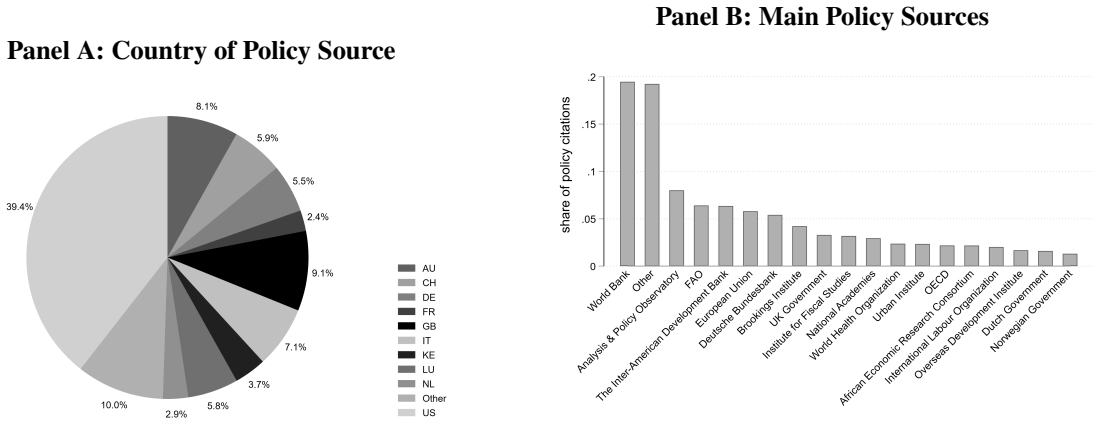

Note: This figure shows the countries of origin of the policy documents citing economics research (Panel A) and their main sources (Panel B). See text for more details.

Figure E3: Share of Articles Receiving at Least One Citation from Policy Documents by Journal

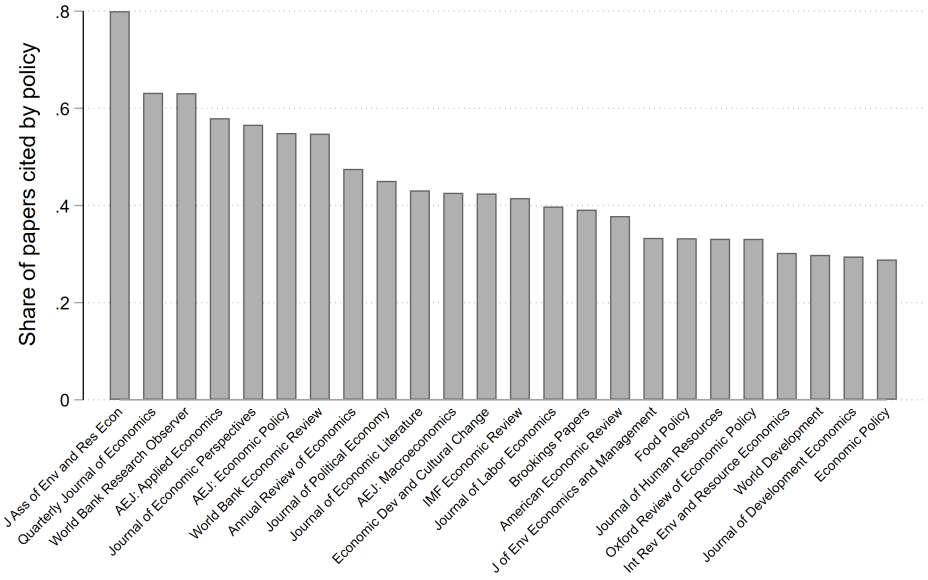

Note: The figure reports the share of published articles that received at least one policy citation by journal. Only the journals with the highest shares are showed in the graph. For the average economics journal in our sample, only 14.96% articles receive mentions from policy sources.

in more prestigious journals (e.g., the *Quarterly Journal of Economics* or the *American Economic Review*) or in more policy-oriented outlets (e.g., the *AEJ: Economic Policy* or the *World Bank Research Observer*) have a higher share of articles that receive policy cites.

We empirically explore the policy impact of applied work compared to theoretical economics articles. Table E1 shows that articles we classified as empirical receive a much larger number of cites from policy documents, an effect proportionally larger when considering only U.S. sources. This effect is robust to the inclusion of journal and year fixed effects, thus effectively capturing the higher policy relevance of empirical papers. Interestingly, we find that the effect increases substantially in the case of papers that employ confidential data available only in FSRDCs: the average number of policy citations from U.S. sources grows by more than four times. The fact that the increase of citations from outside the United States is much smaller fits well the intuition that evidence deriving from Census data should be particularly salient to inform policy-making in the U.S. (Einav and Levin, 2014b). Columns 5 and 6 of Table E1 show that articles citing FSRDC papers are also more cited by policy, but the effect is much smaller.

We find that scientific articles with a better research design tend to receive more attention from policy sources. Table E1 highlights that papers mentioning the use of administrative data in their abstract receive, on average, twice as many policy citations. The last two columns of Table E1 confirm that empirical studies adopting quasi-experimental methods receive more attention from policy. Figure E4 shows the considerable heterogeneity in the consumption of economic research by field: papers in development, labor, and urban

Table E1: Statistical Association Between Research Design and Policy Impact of Economics Papers

	US cites (1)	Non-US cites (2)	US cites (3)	Non-US cites (4)	US cites (5)	Non-US cites (6)	US cites (7)	Non-US cites (8)	US cites (9)	Non-US cites (10)
Empirical (0/1)	0.1627***	0.1541***								
	(0.00939)	(0.01275)								
FSRDC use (0/1)			0.8086***	0.3240**						
			(0.13732)	(0.09898)						
FSRDC cite (0/1)					0.1835***	0.1871***				
					(0.03096)	(0.03798)				
Administrative data (0/1)							0.1805*	0.1641**		
							(0.06879)	(0.04999)		
Causal method (0/1)									0.1130**	0.1143**
									(0.03041)	(0.03281)
Year FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Journal FE	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
N	188,181	188,181	188,181	188,181	188,181	188,181	188,181	188,181	188,181	188,181
Mean of DV	0.1743	0.2676	0.1743	0.2676	0.1743	0.2676	0.1743	0.2676	0.1743	0.2676

Note: This table presents estimates from OLS models evaluating the average increase in policy citations for paper that are empirical (columns 1 and 2), use FSRDC data (columns 3 and 4), cite other FSRDC papers (columns 5 and 6), mention administrative data in their abstract (columns 7 and 8), or mention quasi-experimental methods in their abstract (columns 9 and 10). Papers in the last four columns are tagged using the keywords listed in Table C3. *, **,*** denote significance at 5%, 1% and 0.1% level respectively. Standard errors are clustered by year of publication.

tend to be very influential in the policy debate, unlike fields such as history or microeconomic theory. Finally, the Pearson correlation coefficient between the count of citations from policy sources and academic papers is 0.42, a level of association consistent with past work (Yin et al., 2022). In unreported analyses, we confirm that all the results of this Appendix are robust to explicitly controlling for academic citations.

E.2 Example: Greenstone et al. (2010)

To exemplify the type of policy citations captured by the Altmetric.com data, consider the paper "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings" by Michael Greenstone, Richard Hornbeck, and Enrico Moretti (2010). In this influential paper (1,216 citations on Google Scholar as of 2022), the authors quantify the extent of localised spillovers on productivity. The research design exploits the opening of large manufacturing plants to compare changes in total factor productivity of incumbent plants in "winning" and "losing" counties. To carry out their analyses, the authors accessed several confidential dataset hosted in the FSRDC: the Standard Statistical Establishment List (SSEL), the Annual Survey of Manufactures (ASM), and the Census of Manufactures (CM). Using these data sources, the authors find that incumbents' productivity is 12 percent higher in winning counties five years after the plant openings. Such findings have important implications for policies aimed at local economic development, as remarked by the authors themselves in the introduction of the paper.

Indeed, the paper by Greenstone et al. (2010) received fourteen citations from policy documents in the Altmetric.com data. The documents citing the paper come from a variety of countries: United States (9),

United Kingdom (2), France (1), Belgium (1), and Australia (1). The first citation came in 2012, two years after the publication of the paper, in a policy report prepared by London Economics (a policy consultancy) for the Department for Business, Innovation and Skills of the U.K. Government.[16] The document contains a literature review of the evidence regarding the productivity spillovers of investment in intangible assets. In particular, the report cites Greenstone et al. (2010) as one of the few studies quantifying the magnitude of agglomerations spillovers that does not depend on the usage of patent citations. Similarly, the "5 Year Productivity Review" compiled by the Australian Productivity Commission in 2017 lists this paper as evidence for agglomeration economies.[17]

Some policy documents appear to cite the paper to extrapolate the likely effect of place-based policies aimed at firms, and less for its academic contribution in causally showing the existence of agglomeration spillovers. The 2019 report "How to Solve the Investment Promotion Puzzle: A Mapping of Investment Promotion Agencies in Latin America and the Caribbean and OECD Countries" by the Inter-American Development Bank cites the Million Dollar Plant paper as showing the efficacy of place-based policy aimed at attracting manufacturing investments.[18] Other documents using the findings of Greenstone et al. (2010) as evidence of place-based policy effectiveness include books from Bruegel,[19] policy briefs from the Brooking Institute,[20]

[16]Source: https://gov.uk/government/uploads/12-793-investment-intangible-assets-on-productivity-spillovers.pdf
[17]Source: https://apo.org.au/sites/2017-10/apo-nid115951_1.pdf
[18]Source: https://publications.iadb.org/how-solve-investment-promotion-puzzle
[19]Source: https://www.bruegel.org/sites/2016/01/Blueprint-XXIV.pdf
[20]Source: https://www.brookings.edu/2018/ES_THP_CompetitionFacts.pdf

Figure E4: Share of Articles Receiving at Least One Citation from Policy Documents by Field

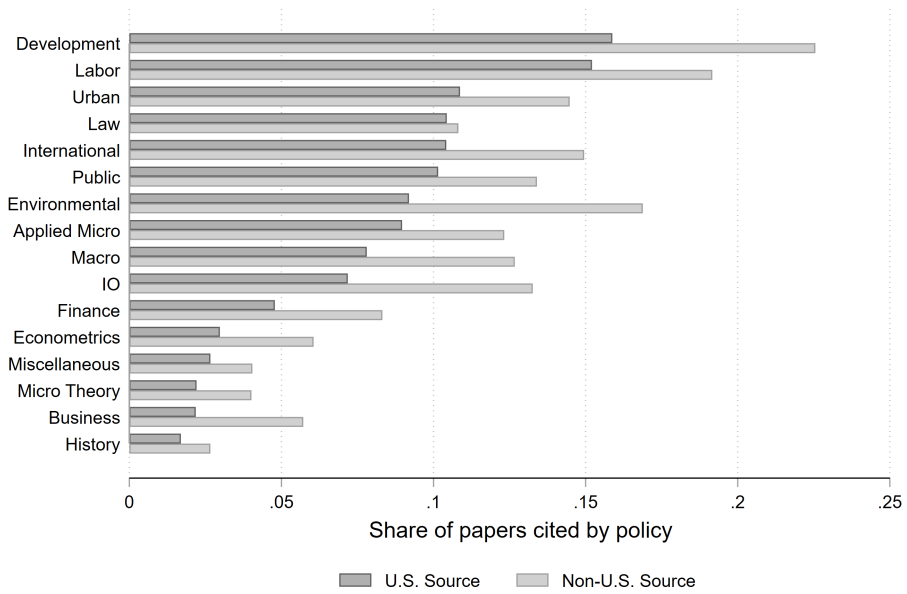

Note: The figure reports the share of published articles that received at least one policy citation by field. Articles are classified into fields using the method of Angrist et al. (2020). For the average economics journal in our sample, only 14.96% articles receive mentions from policy sources.

and policy reports from the OECD.[21]

F Evidence on Articles using FSRDC Datasets

This Appendix briefly describes the nature and expansion of datasets available to researchers in FSRDCs. Next, we show descriptive evidence on the characteristics of the economic articles that directly use confidential Census data. Using manually coded information on the specific datasets used in each paper, we provide suggestive evidence on the lifecycle of research data sources.

F.1 The Development of New Datasets at the FSRDC

Cognizant of the immense research potential of its microdata, in 1982, the U.S. Census Bureau established the Center for Economic Studies (CES) to make these resources accessible to researchers in economics (CES, 2017). The initial focus of the CES was the creation of data resources on the manufacturing sector (McGuckin, 1995). The first matched dataset, developed in 1984, was called Longitudinal Establishment Database (LED) and consisted in pooling the manufacturing data from 1972 to 1981. The significant contribution of that database was the creation of Permanent Plant Numbers (PPNs) that enabled merging survey waves from different years (Atrostic, 2007). Following this first step, CES staff members continued expanding the LED with data from the Economic Censuses and the Annual Survey of Manufactures, eventually creating what is now known as the Longitudinal Research Database (LRD) (McGuckin et al., 1993). The LRD became a favourite of academic researchers, allowing them to conduct pathbreaking empirical research on business dynamics and business demographics (Davis et al., 1998)

In the late 1990s, CES began developing a new database, later called the Longitudinal Business Database (LBD) (Jarmin and Miranda, 2002). According to Chow et al. (2021), creating the LBD was prompted by the desire to test if the results based on the manufacturing data of the LRD also applied to other sectors. The LBD was created by merging the Standard Statistical Establishment List with the Economic Census, thus creating a source that contains basic information on the universe of all U.S. business establishments. This impressive data effort was joined by the concurrent development of the Longitudinal Employer-Household Dynamics (LEHD) data (Abowd et al., 2004). The LEHD merges worker and employer records from Census Bureau surveys with state unemployment insurance claims to create matched employer-employee panel data. Together, the LBD and LEHD constitute a unique research tool to investigate the dynamics of the U.S. economy.

Over time, the range of confidential datasets offered by FSRDC has further expanded. Several other statistical agencies have started making their microdata available through the FSRDC network, most notably

[21] Source: https://www.oecd/OECD-Trade-Policy-Paper-242

Table F1: Summary Statistics of FSRDC Papers

Panel A: FSRDC Papers

	N	Mean	Std. Dev.	Median	Min	Max
Top Field (0/1)	632	0.354	0.48	0	0	1
Top Five (0/1)	632	0.138	0.34	0	0	1
Average Cites	632	7.010	10.63	4	0	92
Top 5% Cited (0/1)	632	0.157	0.36	0	0	1
N Authors	632	2.328	1.03	2	1	7
Year	632	2009.672	7.44	2011	1990	2019

Panel B: Non-FSRDC Papers

	N	Mean	Std. Dev.	Median	Min	Max
Top Field (0/1)	182,316	0.195	0.40	0	0	1
Top Five (0/1)	182,316	0.046	0.21	0	0	1
Average Cites	182,316	3.414	7.26	2	0	402
Top 5% Cited (0/1)	182,316	0.057	0.23	0	0	1
Number of Authors	182,316	2.039	0.95	2	1	6
Year	182,316	2007.594	8.08	2009	1987	2019

Note: This table lists summary statistics at the paper level for 632 FSRDC papers (Panel A) and the remainder 182,316 non-FSRDC papers in our dataset (Panel B). Top Field: 0/1 = 1 for papers in top field economics journals. Top 5 Papers: 0/1 = 1 for papers in a top five journal. Average Cites: count of average yearly citations receives by each published article. Top 5% Cited: 0/1 = 1 for papers in the top 95th percentile of the citation distribution by year of publication. Number of Authors: number of authors of each published paper. Year: year of publication. See text for details.

the National Center for Health Statistics (NCHS), the Bureau of Labor Statistics (BLS), and the Agency for Healthcare Research and Quality (AHRQ). The advantages of doing so are clear since they can leverage the existing system of research data centres without having to muster the resources for creating a similar

Figure F1: Count of Yearly FSRDC Articles over Time

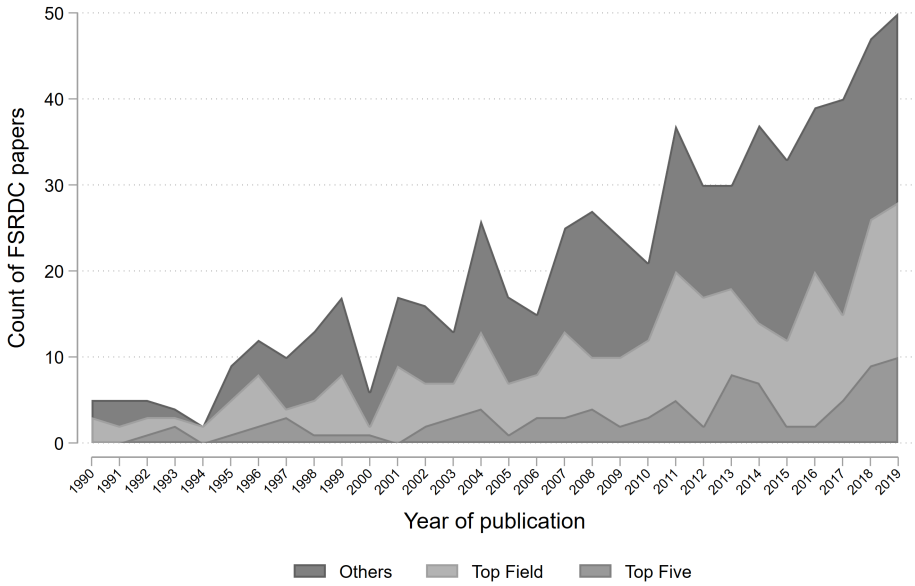

Note: The figure reports the yearly count of economics articles using FSRDC data. The area of the plot is colored in different shades of grey to indicate articles published in a top field journal or in one of the top five journals.

infrastructure (CES, 2017). Moreover, one farsighted feature of the FSRDC is leaving the possibility for approved users to merge existing microdata and to collaborate creating new databases. For example, the recent development of the Management and Organizational Practices Survey (MOPS) has been spearheaded by academic researchers, who have collaborated with the U.S. Census Bureau to design a novel survey instrument (Bloom et al., 2019). As a result, survey and administrative microdata available in FSRDC are increasingly powerful in investigating a broad swath of economic, social, health, and policy questions.

F.2 Characteristics of FSRDC Papers

We begin by assessing the characteristics of papers that employ confidential Census microdata. Table F1 reports the descriptive statistics of the 632 FSRDC papers compared to the remainder 182,316 papers in our sample. Across the board, one notices that research using confidential Census data is much more likely to be published in highly prestigious outlets, especially the so-called top five generalists economics journals. Regarding impact, the average FSRDC paper receives twice as many citations, and it is almost three times more likely to be in the right tail of highly cited papers. Interestingly, FSRDC papers do not differ substantially when looking at the size of their authorship teams. Table F2 shows the strength of the statistical association between the use of confidential Census data and the four proxies of research quality after controlling for the year of publication.

Table F2: Statistical Association Between Proxies of Research Quality and Use of FSRDC Data

	Top Field	Top Five	Average Cites	Top 5% Cites
FSRDC Paper	0.1658***	0.0954***	3.5539***	0.1010***
	(0.01899)	(0.01453)	(0.37347)	(0.01267)
Year FE	YES	YES	YES	YES
N	188,181	188,181	182,948	182,948
Mean of DV	0.1987	0.0483	3.4362	0.0568

Note: This table presents estimates from OLS models evaluating the association between proxies of research quality and the use of confidential Census data. *, **,*** denote significance at 5%, 1% and 0.1% level respectively. Standard errors are clustered by year of publication.

Figure F1 illustrates the steady growth of FSRDC articles published every year. The growing diffusion of FSRDC data in economics closely mimics the expansion of the FSRDC network exploited in the primary analyses of this paper. Yet, a potential downside of lowering access costs is that it could increase the quantity of research produced at the expense of its quality. When confidential data are hard to access, researchers might be using them only for ideas of very high quality, while easier access could increase scientific production mainly in the lower tail of the quality distribution. However, the evidence we present is inconsistent with this line of reasoning. Figure F1 shows that the share of those articles appearing in the top field or top five journals is relatively constant over time. This suggests that the expansion in data access

is not leading to articles of lower quality, at least as proxied by the prestige of the journals where they get accepted.

Another way in which this potential concern can be tested is by exploiting the distance between authors' affiliation and the closest FSRDC (or Census headquarters) to them. Assuming that the distance to the FSRDC is a synthetic proxy of the costs involved in accessing it to perform data analyses, we can test whether articles written by researchers with local access are of lower quality. Table F3 presents the results, where we consider only the distance of the author closest to an FSRDC in the cases where the paper has multiple actors. Overall, we find no relationship between article quality and distance to the closest FSRDC, suggesting once more that expanding data access does not come at the cost of lower-quality science.

Table F3: Statistical Association Between Proxies of Research Quality and Distance from the Nearest FSRDC for Articles Using FSRDC Data

	Top Field		Top Five		Average Cites		Top 5% Cites	
Below 10 miles	-0.0019		0.0084		-0.1230		-0.0058	
	(0.03711)		(0.02554)		(0.91566)		(0.02079)	
		-0.0093		0.0125		0.6783		0.0150
Below 50 miles		(0.03973)		(0.02329)		(0.90897)		(0.02502)
Year FE	YES	YES	YES	YES	YES	YES	YES	YES
N	632	632	632	632	632	632	632	632
Mean of DV	0.3544	0.3544	0.1377	0.1377	7.0099	7.0099	0.1566	0.1566

Note: This table presents estimates from OLS models evaluating whether FSRDC articles written by researchers closer to an FSRDC are of lower quality. Below 10 miles: 0/1 = 1 if at least one author of the paper is affiliated to an institution located within 10 miles of an FSRDC active at the time of publication. Below 50 miles: 0/1 = 1 if at least one author of the paper is affiliated to an institution located within 50 miles of an FSRDC active at the time of publication. *, **,*** denote significance at 5%, 1% and 0.1% level respectively. Standard errors are clustered by year of publication.

F.3 Descriptive Evidence at the Dataset-level

Our main analyses exploit the staggered roll-out of FSRDC across the United States to estimate their causal impact on the rate and direction of academic research in economics. However, one might wonder how different the impact of an FSRDC opened in the 1990s could be relative to those opened in the 2010s. On the one hand, researchers getting access earlier can easily exploit the data to their full potential, unlike those that get access later. This logic would imply diminishing returns to additional FSRDCs because many of the research lines that could be explored with confidential Census data will be exhausted over time. On the other hand, researchers that gain access later might be able to build off a larger pool of metadata and tacit knowledge about Census datasets. Moreover, one of the objectives of the FSRDC program is to enable the creation of new research data, which expands the menu of datasets available over time. Therefore, it is not clear that the value of FSRDC access should be decreasing over time.

Figure F2: Confidential Census Datasets Most Frequently Used in Economics Research

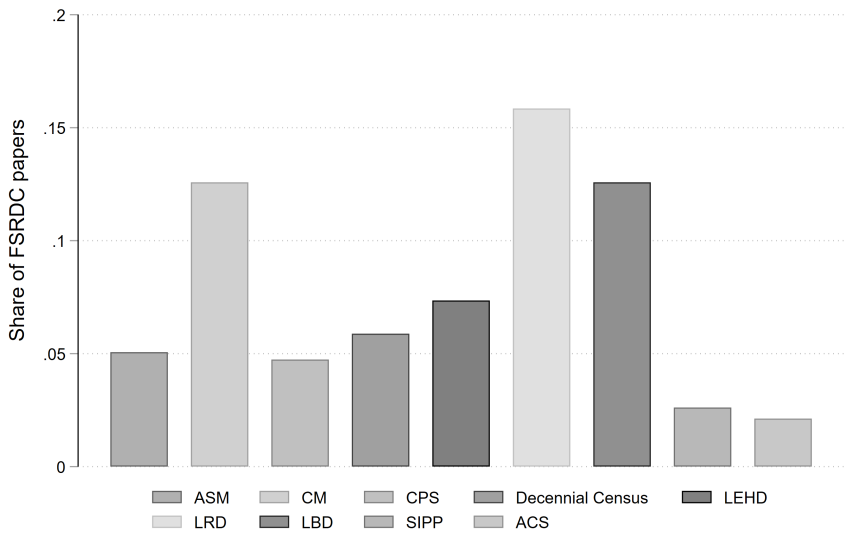

Note: The figure reports the share of articles using FSRDC data that use each of the most common datasets. The dataset showed in the figure include the Annual Survey of Manufactures (ASM), the Census of Manufactures (CM), the Current Population Survey (CPS), the Decennial Census, the Longitudinal Employer-Household Dynamics dataset (LEHD), the Longitudinal Research Database (LRD), Longitudinal Business Database (LBD), the Survey of Income and Program Participation (SIPP), and the American Community Survey (ACS). For the purposes of this figure, we code the article as using only the LRD or LBD when the CM or ASM data are used in conjunction with the LRD or the LBD, respectively.

We painstakingly coded all the specific datasets employed in each FSRDC article to shed light on these alternative explanations.[22] The result of this manual effort is a unique opportunity to explore the "lifecycle" of individual confidential datasets and provide novel information on which Census datasets are the most commonly used in economic research. Figure F2 shows the most popular dataset used in our sample of FSRDC articles. One notices immediately that the LBD and LRD account for almost 30% of all economics papers written in FSRDCs, suggesting that the CES' efforts in creating these resources for research have been highly successful. In general, data on manufacturing plants are among the most popular in economic research. Datasets that provide details on demographic trends, such as the Decennial Census or the Current Population Survey, are used in a smaller number of papers but are likely to be among the most relevant for other disciplines, such as sociology and demography.

Figure F3 breaks down the time series of FSRDC papers by the type of dataset used. A few trends emerge clearly. First, the LRD constituted the main source of data for the first two decades of our sample period, consistent with the accounts of Atrostic (2007). Second, in recent years, the LBD has replaced the LRD, which is now the most popular confidential dataset provided by the Census Bureau. Third, the LEHD is seeing growing popularity in recent years, but much lower than the LBD. Potential reasons for

[22] We thank Buyi Geng and Jiamei (Jasmine) Xu for their precious help in manually coding the confidential dataset used by each FSRDC paper.

Figure F3: Usage of Specific Confidential Census Datasets over Time

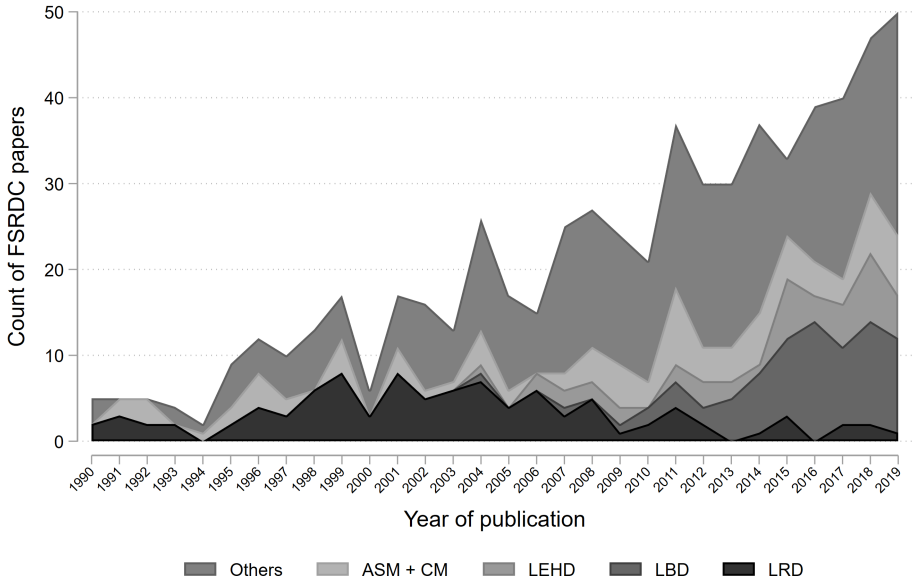

Note: The figure reports the yearly count of economics articles using FSRDC data, divided by the specific confidential dataset used. The figure shows the most common datasets on the manufacturing sector, namely the Longitudinal Research Database (LRD), Longitudinal Business Database (LBD), Longitudinal Employer-Household Dynamics dataset (LEHD), the Annual Survey of Manufactures (ASM) and the Census of Manufactures (CM). For the purposes of this figure, we code the article as using only the LRD or LBD when the CM or ASM data are used in conjunction with the LRD or the LBD, respectively.

this are its partial coverage of U.S. states and the availability of similar matched employer-employee data from Scandinavian countries. Other sources, such as the Census of Manufactures and Annual Survey of Manufactures, are primarily used in conjunction with the LRD and LBD.

Finally, we directly assess whether there are decreasing returns to using confidential Census data over time. Panel A of Figure F4 plots the likelihood that a paper will be published in a top five journal over time,

Figure F4: Likelihood of Being Published on a Top Five Journal, FSRDC Papers vs. PSID Papers

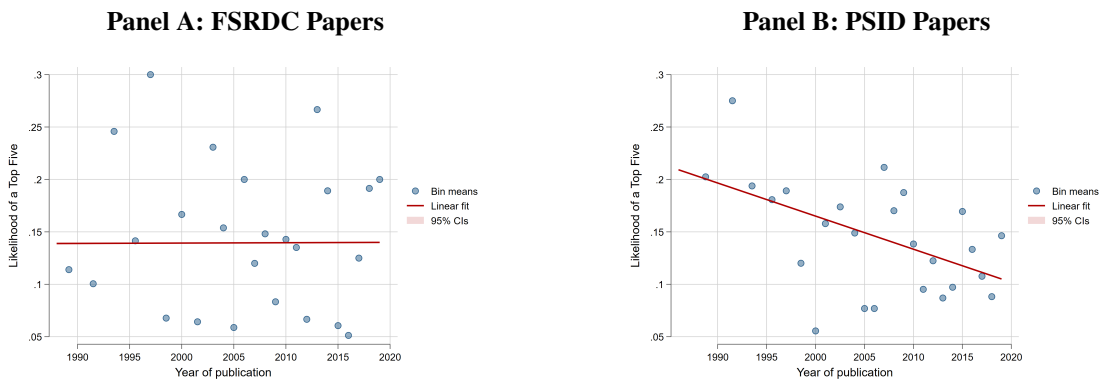

Note: This figure shows a binned scatterplot of the likelihood that an article appears in a top five economics journal over time, separately for papers using confidential Census data (Panel A) and the Panel Study of Income Dynamics (or PSID) data (Panel B). See text for more details.

conditional on using data accessible only in an FSRDC. The figure shows some year-to-year variation but no specific time trends. As a comparison, Panel B shows the same for articles that employ data from the Panel Study of Income Dynamics (PSID). In this case, we can see a clear downward trend, suggesting that the same dataset is less likely conducive to top publication over time. How can this difference be explained? One possibility is that adding new confidential datasets in the FSRDCs counteracts any given data source's natural decline in productivity. Figure F5 shows some suggestive evidence consistent with this supposition. Considering only the LRD, we see a similar pattern to the PSID. However, the creation of the LBD in 2002 opened up new possibilities for scholars, offering an alternative to the progressive exhaustion of LRD's potential. In sum, the progressive addition of new confidential datasets offers a potential explanation for the continuing research value of access to a local FSRDC.

Figure F5: Likelihood of Being Published on a Top Five Journal Conditional on Using FSRDC Data (LRD vs. LBD)

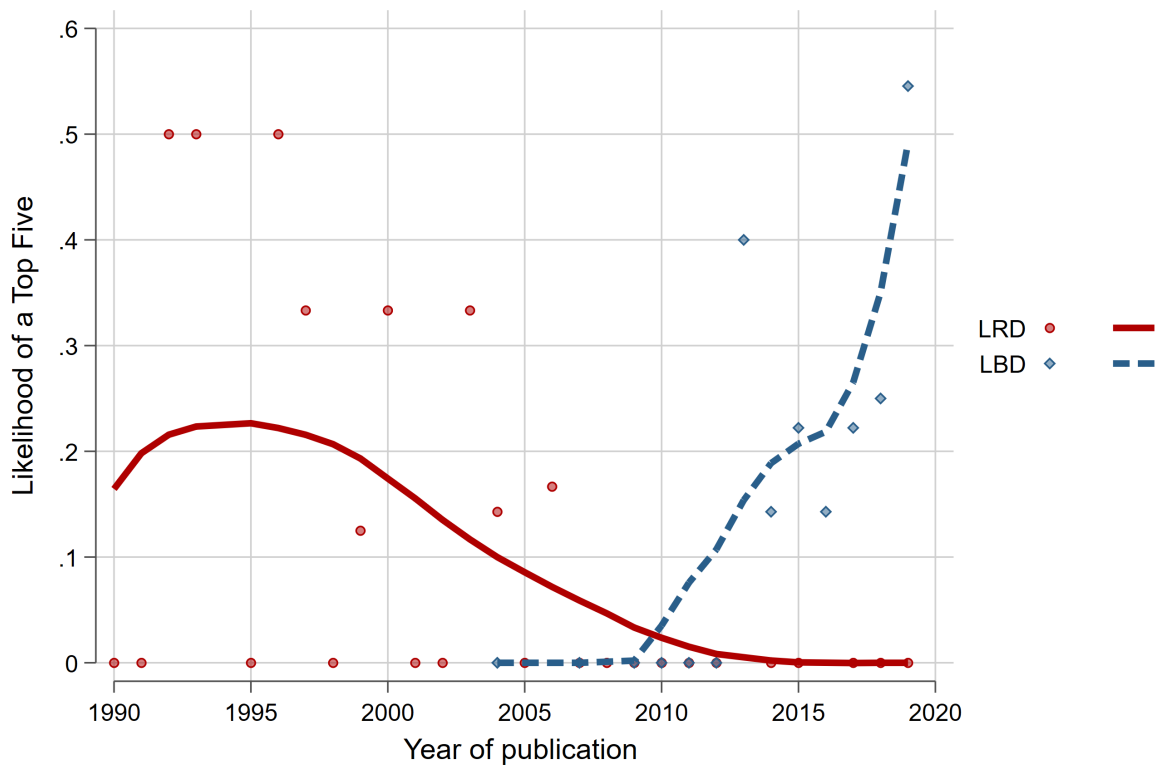

Note: This figure shows a non-parametric binned scatterplot of the likelihood that an article using confidential Census data appears in a top five economics journal over time, separately for papers using the Longitudinal Research Database (LRD) and the Longitudinal Business Database (LBD). See text for more details.

G Appendix Tables and Figures

Table G1: Effect of FSRDC Access on Projects and Working Papers

	FSRDC Project Approval (0/1)		CES Working Paper (0/1)	
	(1)	(2)	(3)	(4)
Post-FSRDC	-0.000996	-0.00178	-0.000883	-0.000448
	(0.00052)	(0.00098)	(0.00052)	(0.00093)
Post-FSRDC × Empiricist	0.00457***	0.00460***	0.00200**	0.00182*
	(0.00079)	(0.00080)	(0.00074)	(0.00078)
Researcher FE	Yes	Yes	Yes	Yes
University Tier × Year FE	Yes	No	Yes	No
University × Year FE	No	Yes	No	Yes
N	246532	245556	246532	245556

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on additional measures of administrative data adoption. Columns (1) and (2) report results from OLS models, where the dependent variable is an indicator of whether the researcher has any new FSRDC project approved. Columns (3) and (4) report results from OLS models, where the dependent variable is the number of CES working papers published. Post-FSRDC equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects. Columns (1) and (3) further include year fixed effects interacted with with university-tier dummies, and columns (2) and (4) further include year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively. See details in Appendix C.3.

Table G2: Effect of FSRDC Access on Research Output Using the Callaway-Sant'Anna Doubly-Robust DID Estimator.

	Use FSRDC (1)	Cite FSRDC (2)	Top Pubs (3)	Cite-Pubs (4)
ATT	0.00598***	0.0209***	0.0259*	2.235***
	(0.001)	(0.004)	(0.011)	(0.460)
Observations	230,348	230,348	230,348	230,348

Note: This table presents ATT estimates from models evaluating the impact of FSRDC access on measures of administrative data adoption and research output using the doubly-robust difference-indifference estimator developed by Callaway and Sant'Anna (2021). Column (1) reports result from models where the dependent variable is the count of articles that directly use confidential microdata accessible only in an FSRDC. Columns (2) reports result from models where the dependent variable is the count of articles that cite a paper using confidential data accessible only in an FSRDC (excluding papers that directly use confidential Census data). Column (3) reports result from models where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan, 2020). Column (4) reports result from models where the dependent variable is the count of articles weighted by the number of citations received up to 5 years following their publication. Post-FSRDC equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects and year fixed effects. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G3: Effect of FSRDC Access for Universities

	Use FSRDC (1)	Cite FSRDC (2)	Top Pubs (3)	Cite-Pubs (4)
Post-FSRDC	0.156***	0.931**	1.927**	105.2**
	(0.043)	(0.290)	(0.603)	(33.248)
University FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
N	9094	9094	9094	9094

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of administrative data adoption and research output at the university level. Column (1) reports result from OLS models, where the dependent variable is the count of articles that directly use confidential microdata accessible only in an FSRDC. Column (2) reports result from OLS models, where the dependent variable is the count of articles that cite a paper using confidential data accessible only in an FSRDC (excluding papers that directly use confidential Census data). Column (3) reports result from OLS models, where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan, 2020). Column (4) reports result from OLS models, where the dependent variable is the count of articles weighted by the number of citations received up to 5 years following their publication. Post-FSRDC equals one in all years after there is an operating FSRDC in the city of the university. All models include university fixed effects and year fixed effects. Standard errors are in parentheses, clustered at the university level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G4: Effect of FSRDC Access Excluding NSF Grant Applicants

	FSRDC Use	FSRDC Cite	Top Pubs	Cite-Pubs
	(1)	(2)	(3)	(4)
Post-FSRDC	0.00224*	-0.00196	-0.00422	-0.123
	(0.001)	(0.005)	(0.011)	(0.436)
Post-FSRDC × Empiricist	0.00323**	0.0161***	0.0316**	1.664***
	(0.001)	(0.004)	(0.011)	(0.445)
Researcher FE	Yes	Yes	Yes	Yes
University × Year FE	Yes	Yes	Yes	Yes
N	244850	244850	244850	244850

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of administrative data adoption and research output excluding researchers who were the applicants of the NSF grant leading to a FSRDC center. Column (1) reports result from OLS models, where the dependent variable is the count of articles that directly use confidential microdata accessible only in an FSRDC. Columns (2) reports result from OLS models, where the dependent variable is the count of articles that cite a paper using confidential data accessible only in an FSRDC (excluding papers that directly use confidential Census data). Column (3) reports result from OLS models, where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan 2020). Column (4) reports result from OLS models, where the dependent variable is the count of articles weighted by the number of citations received up to 5 years following their publication. Post-FSRDC equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects and year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G5: Effect of FSRDC Access Excluding Authors Treated After Mobility Events

	FSRDC Use	FSRDC Cite	Top Pubs	Cite-Pubs
	(1)	(2)	(3)	(4)
Post-FSRDC	0.00274	-0.00324	-0.0313*	-0.928
	(0.001)	(0.005)	(0.015)	(0.542)
Post-FSRDC × Empiricist	0.00347*	0.0182***	0.0442***	2.141***
	(0.001)	(0.005)	(0.013)	(0.560)
Researcher FE	Yes	Yes	Yes	Yes
University × Year FE	Yes	Yes	Yes	Yes
N	219941	219941	219941	219941

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of administrative data adoption, research output, and research direction excluding researchers who got treated because of moving to an institution with an operating FSRDC. Column (1) reports result from OLS models, where the dependent variable is the count of articles that directly use confidential microdata accessible only in an FSRDC. Columns (2) reports result from OLS models, where the dependent variable is the count of articles that cite a paper using confidential data accessible only in an FSRDC (excluding papers that directly use confidential Census data). Column (3) reports result from OLS models, where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan 2020). Column (4) reports result from OLS models, where the dependent variable is the count of articles weighted by the number of citations received during the 5 years following their publication. Post-FSRDC equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects and year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G6: Effect of FSRDC Access Excluding Authors Losing Local Access After Mobility Events

	FSRDC Use (1)	FSRDC Cite (2)	Top Pubs (3)	Cite-Pubs (4)
Post-FSRDC × Empiricist	0.00375*	0.0153**	0.0454***	2.201***
	(0.002)	(0.006)	(0.013)	(0.563)
Researcher FE	Yes	Yes	Yes	Yes
University × Year FE	Yes	Yes	Yes	Yes
N	206071	206071	206071	206071

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of administrative data adoption, research output, and research direction excluding researchers who lost local access after moving to an institution without a local operating FSRDC. The main effect of $PostFSRDC_{j,t}$ is no longer identified since the treatment status is coded as an absorbing state. Column (1) reports result from OLS models, where the dependent variable is the count of articles that directly use confidential microdata accessible only in an FSRDC. Columns (2) reports result from OLS models, where the dependent variable is the count of articles that cite a paper using confidential data accessible only in an FSRDC (excluding papers that directly use confidential Census data). Column (3) reports result from OLS models, where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan 2020). Column (4) reports result from OLS models, where the dependent variable is the count of articles weighted by the number of citations received during the 5 years following their publication. Post-FSRDC equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects and year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G7: Effect of FSRDC Access Excluding Authors in Institutions Hosting FSRDCs

	FSRDC Use	FSRDC Cite	Top Pubs	Cite-Pubs
	(1)	(2)	(3)	(4)
Post-FSRDC	0.00249	-0.00400	-0.00573	-0.621
	(0.002)	(0.007)	(0.016)	(0.636)
Post-FSRDC × Empiricist	0.00325	0.0129*	0.0315*	2.024**
	(0.002)	(0.006)	(0.015)	(0.654)
Researcher FE	Yes	Yes	Yes	Yes
University × Year FE	Yes	Yes	Yes	Yes
N	194110	194110	194110	194110

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of administrative data adoption, research output, and research direction excluding researchers who affiliate with an institution with an operating FSRDC center. Column (1) reports result from OLS models, where the dependent variable is the count of articles that directly use confidential microdata accessible only in an FSRDC. Columns (2) reports result from OLS models, where the dependent variable is the count of articles that cite a paper using confidential data accessible only in an FSRDC (excluding papers that directly use confidential Census data). Column (3) reports result from OLS models, where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan 2020). Column (4) reports result from OLS models, where the dependent variable is the count of articles weighted by the number of citations received during the 5 years following their publication. Post-FSRDC equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects and year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G8: Effect of FSRDC Access by Type of Data Access

	Use FSRDC		Cite FSRDC		Top Pubs		Cite Pubs	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
City	0.00181	0.00199	0.00302	0.00609	-0.0155	-0.0112	-1.105	-0.328
	(0.00167)	(0.00203)	(0.00490)	(0.00622)	(0.01689)	(0.01901)	(0.75919)	(0.89523)
Consortium	-0.000898	-0.000122	-0.0107*	-0.00629	-0.000615	0.0307	-0.192	0.0676
	(0.00079)	(0.00162)	(0.00514)	(0.00716)	(0.01209)	(0.01580)	(0.41284)	(0.63209)
City+Consortium	-0.0000178	0.00222	-0.0101	-0.00579	-0.00403	0.0114	-0.616	-1.379
	(0.00133)	(0.00200)	(0.00540)	(0.00746)	(0.01483)	(0.01782)	(0.62231)	(0.86739)
University	-0.000956	0.00403*	-0.0123**	0.00166	-0.0195	-0.00177	-0.373	1.391
	(0.00103)	(0.00177)	(0.00454)	(0.00646)	(0.01284)	(0.01579)	(0.47862)	(0.71213)
City × Empiricist	-0.00270	-0.00307	-0.00622	-0.00837	0.0254	0.0218	1.391	1.190
	(0.00234)	(0.00241)	(0.00658)	(0.00686)	(0.01898)	(0.01953)	(0.91808)	(0.96321)
Consortium × Empricist	0.00165	0.00102	0.0174**	0.0167**	0.00291	0.00196	0.629	0.585
	(0.00103)	(0.00110)	(0.00613)	(0.00628)	(0.01354)	(0.01373)	(0.51243)	(0.54375)
City+Consortium × Empiricist	0.00322	0.00244	0.0212**	0.0185**	0.0176	0.0177	1.734*	1.857*
	(0.00204)	(0.00212)	(0.00685)	(0.00716)	(0.01683)	(0.01684)	(0.79447)	(0.82103)
University × Empiricist	0.00741***	0.00678***	0.0298***	0.0292***	0.0540***	0.0554***	1.931**	1.932**
	(0.00169)	(0.00171)	(0.00631)	(0.00654)	(0.01431)	(0.01449)	(0.62555)	(0.63957)
Researcher FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
University Tier × Year FE	Yes	No	Yes	No	Yes	No	Yes	No
University × Year FE	No	Yes	No	Yes	No	Yes	No	Yes
N	246532	245556	246532	245556	246532	245556	246532	245556

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of administrative data adoption and research output by type of researchers' access. City equals one in all years after a researcher has been affiliated to a research institution located in a city (but not same university or consortium) with an operating FSRDC. Consortium equals one in all years after a researcher has been affiliated to a research institution part of an FSRDC consortium (but not in the same university or city). City+Consortium equals one in all years after a researcher has been affiliated to a research institution member of the FSRDC consortium and located in the same city of the FSRDC (but not same university). University equals one in all years after a researcher has been affiliated to a research institution with an operating FSRDC. The omitted category are researchers affiliated to a research institution without an FSRDC that is neither in the same city nor part of a consortium operating FSRDC. Columns (1) and (2) report results from OLS models, where the dependent variable is the count of articles that directly use confidential microdata accessible only in an FSRDC. Columns (3) and (4) report results from OLS models, where the dependent variable is the count of articles that cite a paper using confidential data accessible only in an FSRDC (excluding papers that directly use confidential Census data). Columns (5) and (5) report results from OLS models, where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan, 2020). Columns (7) and (8) report results from OLS models, where the dependent variable is the count of articles weighted by the number of citations received up to 5 years following their publication. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. Columns (1), (3), (5), and (7) further include year fixed effects interacted with university-tier dummies and columns (2), (4), (6) and (8) further include year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G9: Alternative Methods to Measure FSRDC Exposure

	Top Pubs (1)	Cite-Pubs (2)	Top-Pubs (3)	Cite-Pubs (4)	Top-Pubs (5)	Cite-Pubs (6)
Post-FSRDC (Institution)	-0.00913	1.466				
	(0.016)	(0.778)				
Post-FSRDC (Institution) × Empiricist	0.0525***	1.700**				
	(0.014)	(0.645)				
Post-FSRDC (City)			-0.00712	-0.170		
			(0.012)	(0.435)		
Post-FSRDC (City) × Empiricist			0.0352**	1.737***		
			(0.011)	(0.445)		
Post-FSRDC (Consortium)					0.0127	0.00776
					(0.010)	(0.462)
Post-FSRDC (Consortium) × Empiricist					0.0246**	1.395**
					(0.009)	(0.442)
Researcher FE	Yes	Yes	Yes	Yes	Yes	Yes
University × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	245556	245556	245556	245556	245556	245556

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of research output using alternative measures of exposure. Columns (1), (3), and (5) report results from OLS models, where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan, 2020). Columns (2), (4) and (6) report results from OLS models, where the dependent variable is the count of articles weighted by the number of citations received up to 5 years following their publication. Post-FSRDC (Institution) equals one in all years after a researcher has been affiliated to a research institution hosting an operating FSRDC. Post-FSRDC (City) equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. And Post-FSRDC (Consortium) equals one in all years after a researcher has been affiliated to a research institution belonging to a consortium operating an FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects and year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G10: Effect of FSRDC Access Controlling for Broader Time Trends in Empirical Research

	Use FSRDC (1)	Cite FSRDC (2)	Top Pubs (3)	Cite-Pubs (4)
Post-FSRDC	0.00269*	0.00653	0.00587	0.186
	(0.001)	(0.005)	(0.012)	(0.433)
Post-FSRDC × Empiricist	0.00254*	0.00408	0.0159	1.204**
	(0.001)	(0.004)	(0.011)	(0.449)
Researcher FE	Yes	Yes	Yes	Yes
University × Year FE	Yes	Yes	Yes	Yes
Empiricist × Year Bin FE	Yes	Yes	Yes	Yes
N	245556	245556	245556	245556

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of administrative data adoption and research output. Column (1) reports results from an OLS model, where the dependent variable is the count of articles that directly use confidential microdata accessible only in an FSRDC. Column (2) reports results from an OLS model, where the dependent variable is the count of articles that cite a paper using confidential data accessible only in an FSRDC (excluding papers that directly use confidential Census data). Column (3) reports results from an OLS model, where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan, 2020). Column (4) reports results from an OLS model, where the dependent variable is the count of articles weighted by the number of citations received up to 5 years following their publication. Post-FSRDC equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects, year fixed effects interacted with university dummies, and empiricist fixed effects interacted with time dummies (in bins of five years each). Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G11: Effect of FSRDC Access Excluding Direct FSRDC Users

	FSRDC Cite		Top Pubs		Cite-Pubs	
	(1)	(2)	(3)	(4)	(5)	(6)
Post-FSRDC	-0.00674**	-0.00440	-0.0147	-0.00834	-0.510	-0.331
	(0.003)	(0.004)	(0.009)	(0.011)	(0.313)	(0.392)
Post-FSRDC × Empiricist	0.0107***	0.0107**	0.0341**	0.0337**	1.278***	1.318***
	(0.003)	(0.003)	(0.010)	(0.010)	(0.372)	(0.391)
Researcher FE	Yes	Yes	Yes	Yes	Yes	Yes
University Tier × Year FE	Yes	No	Yes	No	Yes	No
University × Year FE	No	No	No	No	No	No
N	238710	237747	238710	237747	238710	237747

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of administrative data adoption and research output excluding researchers who ever used FSRDC. Columns (1) and (2) report results from OLS models, where the dependent variable is the count of articles that cite a paper using confidential data accessible only in an FSRDC (excluding papers that directly use confidential Census data). Columns (3) and (4) report result from OLS models, where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan 2020). Columns (5) and (6) report results from OLS models, where the dependent variable is the count of articles weighted by the number of citations received up to 5 years following their publication. Post-FSRDC equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects. Columns (1), (3) and (5) further include year fixed effects interacted with university-tier dummies, and columns (2), (4) and (6) further include year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G12: Effect of FSRDC Access on Alternative Measures of Research Direction

Panel A: Papers Using JEL Codes (dummy)

	Definition 1		Definition 2		Definition 3	
	FSRDC	Non-FSRDC	FSRDC	Non-FSRDC	FSRDC	Non-FSRDC
	(1)	(2)	(3)	(4)	(5)	(6)
Post-FSRDC	-0.00332	-0.0125	-0.0107	-0.0105	-0.00817	-0.0127
	(0.007)	(0.010)	(0.006)	(0.010)	(0.008)	(0.010)
Post-FSRDC × Empiricist	0.0175**	0.0278**	0.0278***	0.0263**	0.0355***	0.0294**
	(0.006)	(0.009)	(0.006)	(0.009)	(0.007)	(0.009)
Dependent Variable Mean	0.098	0.320	0.099	0.320	0.160	0.311
Researcher FE	Yes	Yes	Yes	Yes	Yes	Yes
University × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	245556	245556	245556	245556	245556	245556

Panel B: Papers Using JEL Codes (count)

	Definition 1		Definition 2		Definition 3	
	FSRDC	Non-FSRDC	FSRDC	Non-FSRDC	FSRDC	Non-FSRDC
	(1)	(2)	(3)	(4)	(5)	(6)
Post-FSRDC	-0.00837	-0.0180	-0.0188*	-0.0159	-0.0135	-0.0160
	(0.009)	(0.019)	(0.008)	(0.020)	(0.011)	(0.019)
Post-FSRDC × Empiricist	0.0268**	0.0668***	0.0412***	0.0646***	0.0576***	0.0676***
	(0.008)	(0.018)	(0.008)	(0.018)	(0.011)	(0.018)
Dependent Variable Mean	0.116	0.450	0.116	0.451	0.197	0.434
Researcher FE	Yes	Yes	Yes	Yes	Yes	Yes
University × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	245556	245556	245556	245556	245556	245556

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of research directions. We experiment with alternative coding of what JEL codes are the most representative of FSRDC research. Definition 1: JEL codes with the largest frequency difference between FSRDC and non-FSRDC papers (this reproduces the main analyses). Definition 2: JEL codes that are among the twenty most frequent for FSRDC papers, but not among the twenty most frequent in non-FSRDC papers. Definition 3: JEL codes with the frequency difference between FSRDC and non-FSRDC papers larger than 0.03. Post-FSRDC equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects and year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G13: Effect of FSRDC Access on Additional Proxies of Research Design

	Big Data	Data	Natural Exp.
	(1)	(2)	(3)
Post-FSRDC	-0.000457*	0.0204*	-0.00139
	(0.000)	(0.009)	(0.001)
Post-FSRDC × Empiricist	0.0000839	-0.00139	0.00335**
	(0.000)	(0.008)	(0.001)
Researcher FE	Yes	Yes	Yes
University × Year FE	Yes	Yes	Yes
N	245556	245556	245556

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on additional text-based proxies of research design. Columns (1) and (2) report results from OLS models, where the dependent variable is the number of published articles that mention in their title or abstract keywords related to the use of big data or just data, respectively. Column (3) reports results from an OLS model, where the dependent variable is the number of published articles that mention in their title or abstract keywords related to the use of natural experiments. Post-FSRDC equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects and year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G14: Effect of FSRDC Access: Timing of NSF Grant

	FSRDC Use	FSRDC Cite	Top Pubs	Cite-Pubs
	(1)	(2)	(3)	(4)
Post-Grant	0.000894	-0.00810*	-0.0143	-0.558
	(0.001)	(0.003)	(0.011)	(0.362)
Post-Grant × Empiricist	0.00152	0.0138***	0.0296**	1.502***
	(0.001)	(0.004)	(0.011)	(0.411)
Researcher FE	Yes	Yes	Yes	Yes
University Year × Year FE	Yes	Yes	Yes	Yes
N	229756	229756	229756	229756

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of administrative data adoption and research output using alternative measures of exposure. Column (1) reports result from OLS models, where the dependent variable is the count of articles that directly use confidential microdata accessible only in an FSRDC. Columns (2) reports result from OLS models, where the dependent variable is the count of articles that cite a paper using confidential data accessible only in an FSRDC (excluding papers that directly use confidential Census data). Columns (3) reports result from OLS models, where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan 2020). Column (4) reports result from OLS models, where the dependent variable is the count of articles weighted by the number of citations received up to 5 years following their publication. Post-Grant equals one in all years after a researcher has been affiliated to a research institution that received an NSF grant to establish a local FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects and year fixed effects interacted with university fixed effects. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Table G15: Descriptive Statistics of the Policy Citation Data from Altmetric.com

	N	Mean	Std. Dev.	Median	Min	Max
Policy Cites	246711	0.376	2.48	0	0	136
Policy Cites (U.S. only)	246711	0.178	1.31	0	0	79
Policy Cites (Non-U.S. only)	246711	0.198	1.37	0	0	83
Papers Cited by Policy	246711	0.105	0.37	0	0	8
Papers Cited by Policy (U.S. only)	246711	0.065	0.29	0	0	7
Papers Cited by Policy (Non-U.S. only)	246711	0.075	0.31	0	0	8

Note: This table presents summary statistics at the researcher-year level for an unbalanced panel of 246,711 observations. Policy cites: count of citations received by all articles published in a given year. Policy cites (U.S. only): count of citations received by all articles published in a given year from U.S. policy sources. Policy cites (Non-U.S. only): count of citations received by all articles published in a given year from non-U.S. policy sources. Papers cited by policy: count of all articles published in a given year that received at least one policy citation. Papers cited by policy (U.S. only): count of all articles published in a given year that received at least one policy citation from a U.S. policy source. Papers cited by policy (non-U.S. only): count of all articles published in a given year that received at least one policy citation from a non-U.S. policy source. See text for details.

Table G16: Effect of FSRDC Access on the Policy Orientation of Empirical Researchers

	Mentions of Policy (1)	Policy JELs (all) (2)	Policy JELs (subset) (3)
Post-FSRDC	-0.00974	-0.000619	0.00379**
	(0.008)	(0.003)	(0.001)
Post-FSRDC × Empiricist	0.0118	0.00245	-0.00166
	(0.007)	(0.002)	(0.001)
Researcher FE	Yes	Yes	Yes
University × Year FE	Yes	Yes	Yes
N	245556	245556	245556

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of policy orientation of empirical researchers. Column (1) reports result from OLS models, where the dependent variable is the count of articles that mention the words "policy" or "policies" in the title or abstract. Columns (2) reports result from OLS models, where the dependent variable is the count of articles that report JEL codes including the word "policy". Columns (3) reports result from OLS models, where the dependent variable is the count of articles that report JEL codes most associated with government, labor, and public policies. Post-FSRDC equals one in all years after a researcher has been affiliated to a research institution located in a city with an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects and year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively

Table G17: Effect of FSRDC Access Suppressing Changes in Affiliation Over Time

	Top Pubs (1)	Cite-Pubs (2)	Top Five (3)	Top 5% Cite (4)
Post-FSRDC	-0.0125	-0.955*	-0.0156*	-0.0181***
	(0.010)	(0.378)	(0.006)	(0.005)
Post-FSRDC × Empiricist	0.0192	1.435***	0.0213**	0.0212***
	(0.011)	(0.433)	(0.007)	(0.006)
Researcher FE	Yes	Yes	Yes	Yes
University × Year FE	Yes	Yes	Yes	Yes
N	245556	245556	245556	245556

Note: This table presents estimates from OLS models evaluating the impact of FSRDC access on measures of research output. The data are coded to suppress endogenous changes in affiliation over time, meaning that each researchers is considered as if they spent their entire career in their first placement. Column (1) reports result from OLS models, where the dependent variable is the count of articles published in the most prestigious economics journals (the list of Top Five and Top Field journals is from Heckman and Moktan 2020). Column (2) reports result from OLS models, where the dependent variable is the count of articles weighted by the number of citations received during the 5 years following their publication. Column (3) reports result from OLS models, where the dependent variable is the count of articles published in the Top Five generalists economics journals. Column (4) reports result from OLS models, where the dependent variable is the number of publications whose number of citations is in the top 5% of the citations distribution for the year in which they were published. Post-FSRDC equals one in all years after the city of the first affiliation of the researcher has received an operating FSRDC. Empiricist equals one for those researchers whose lifetime publications are mostly empirical in nature, as assessed by our machine-learning based paper classifier. All models include individual fixed effects and year fixed effects interacted with university dummies. Standard errors are in parentheses, clustered at the researcher level. *, **,*** denote significance at 5%, 1% and 0.1% level respectively.

Figure G1: Rank of Universities Treated by FSRC Openings (average)

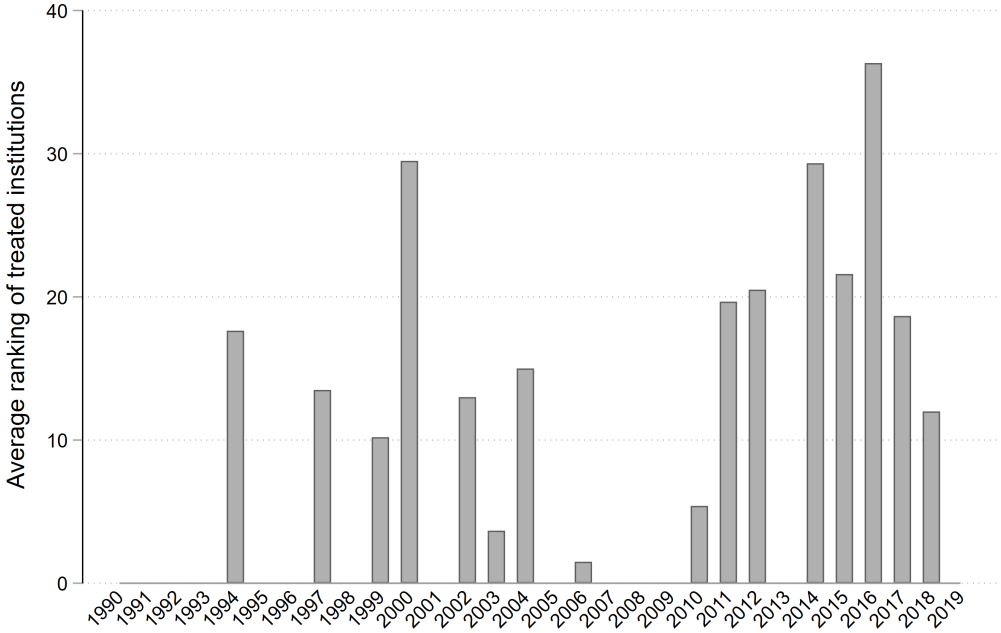

Note: The figure reports the average rank of universities treated by the opening of a new FSRDC in their city. Informations on the ranking of economics departments is taken from Kalaitzidakis et al. (2003). Lower numbers correspond to departments ranked higher. See text for more details.

Figure G2: Cumulative Number of Researchers Gaining FSRDC Access over Time

(i) By Modality of Access

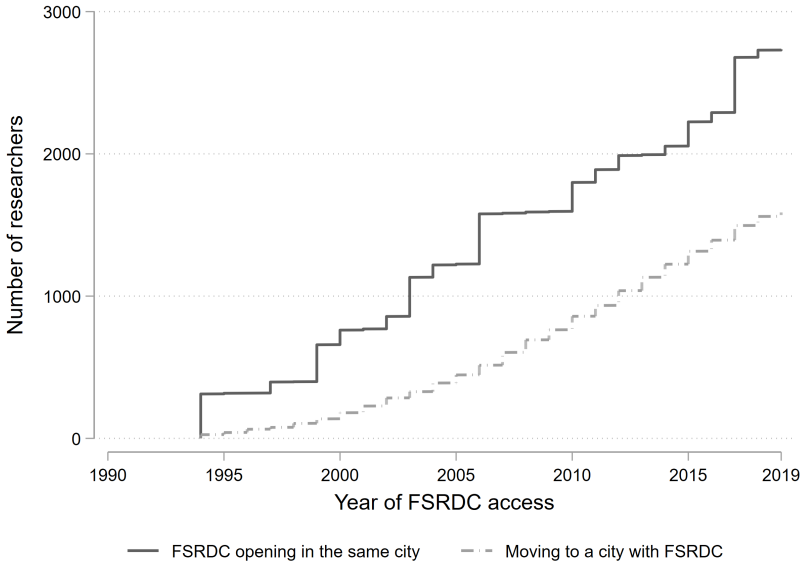

(ii) By Ranking of Institution

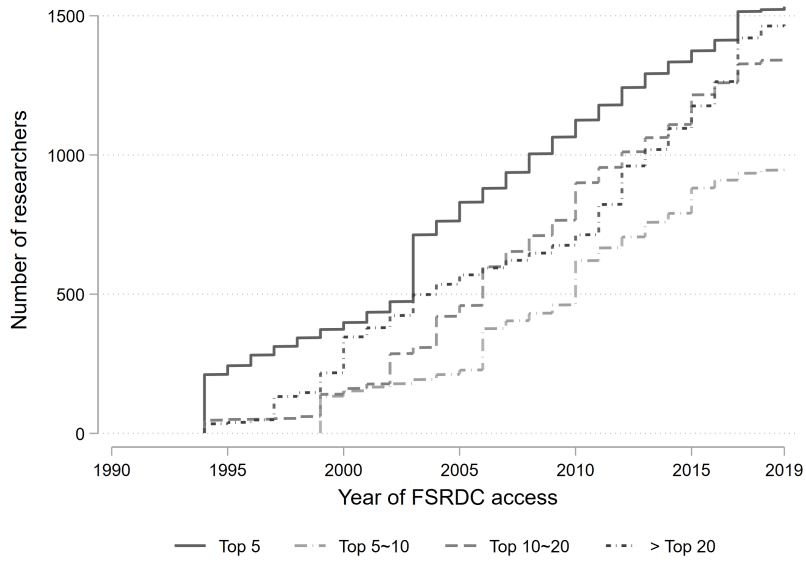

Note: This figure plots the cumulative number of researchers that gain access to an FSRDC in the city where they are located over time. The cumulative frequencies are split by the modality in which researchers become co-located to an FSRDC (panel (i)) and the ranking of their institution of affiliation (panel (ii)). See text for more details.

Figure G3: Yearly Papers using Census Data (per capita averages)

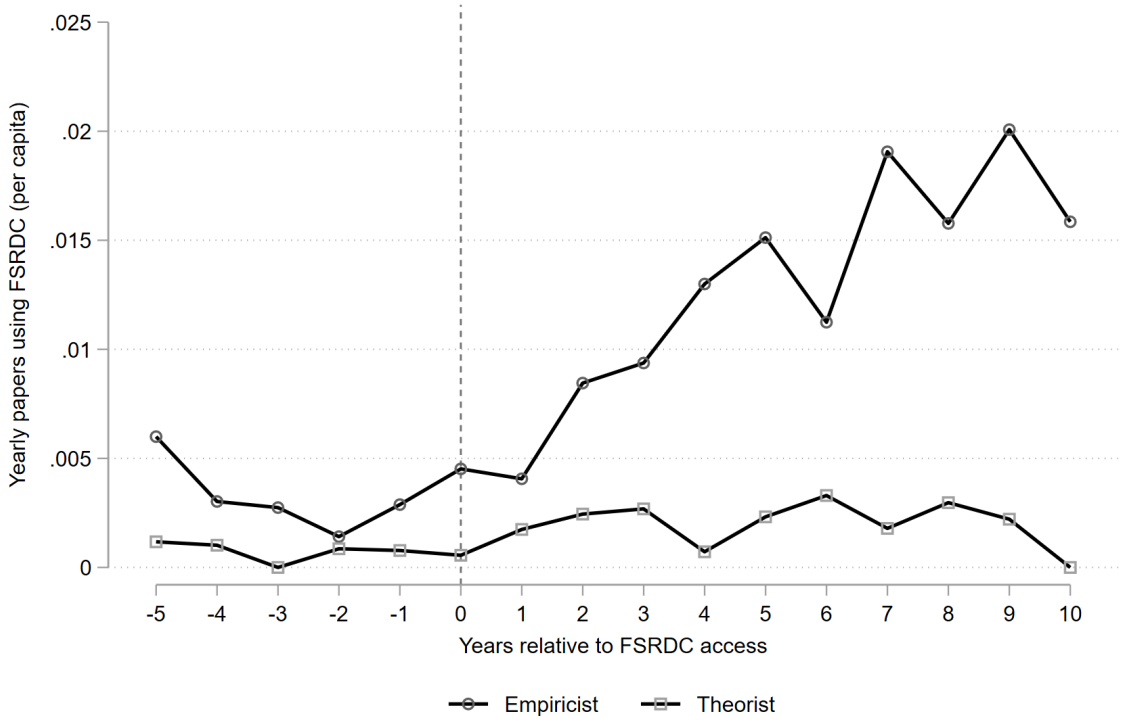

Note: The figure reports the average per capita number of yearly articles written using Census data. The yearly individual means are computed relative to the time of first access to a FSRDC. See text for more details.

Figure G4: Time-Varying Estimates of the Impact of FSRDCs on Research Output of Applied Scholars, Obtained with the Callaway-Sant'Anna Doubly-Robust DID Estimator.

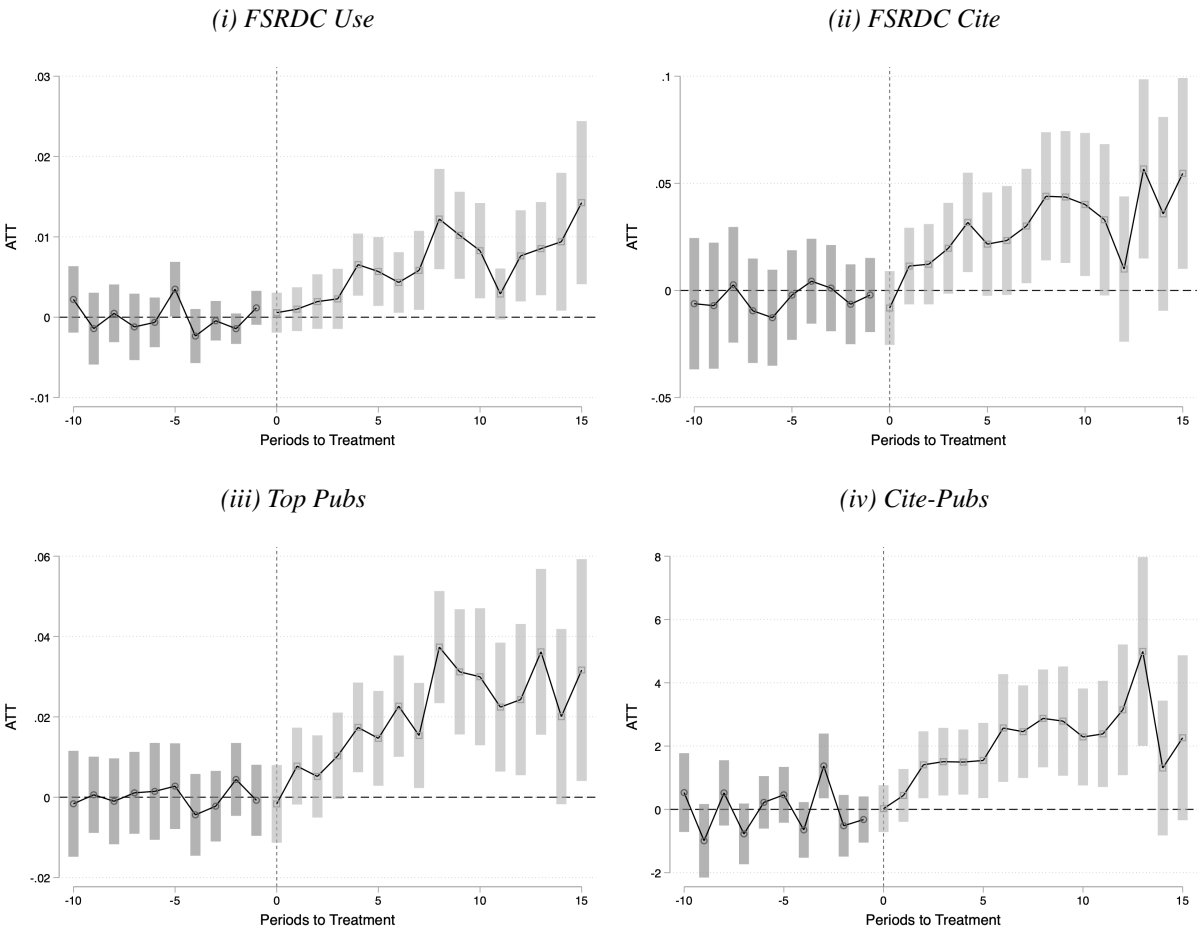

Note: This figure provides visual illustrations of the event study version of the main regressions evaluating the impacts of FSRDC access on measures of research output, estimated using the doubly-robust difference-indifference estimator developed by Callaway and Sant'Anna (2021). The main dependent variables are the number of papers written using Census data (panel (i)), the number of papers that cite FSRDC papers (panel (ii)), the number of top publications (Panel (iii)), and the citation-weighted number of publications (Panel (iv)). The charts plot values of β for different lags before and after obtaining access to an FSRDC with 95% confidence intervals. Regressions include researcher and university $\times$ year fixed effects. Standard errors are clustered at the researcher level.

Figure G5: Effect of FSRDC Access Using Split Sample Regressions (Theorists vs. Empiricists)

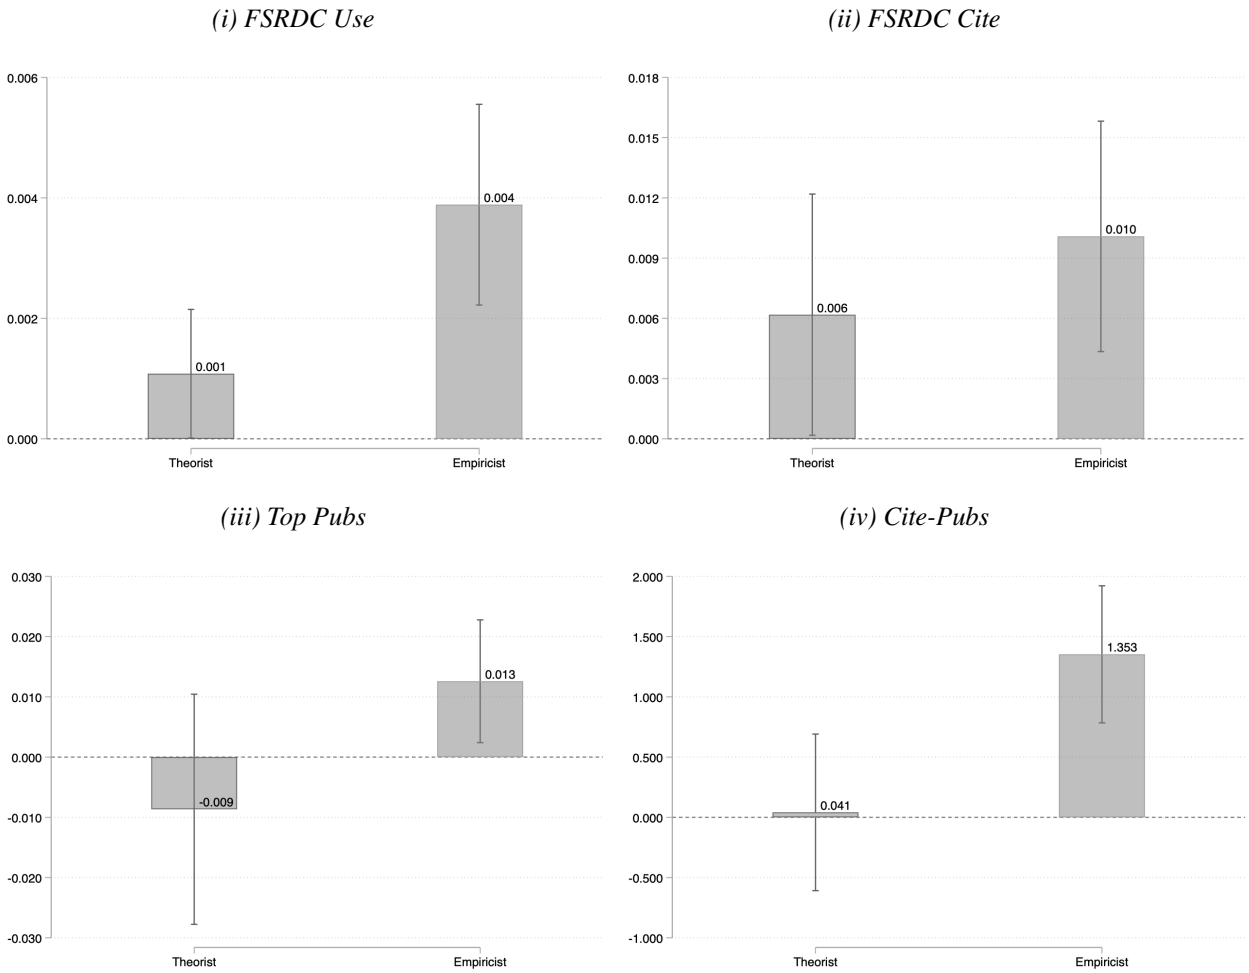

Note: This figure provides visual illustrations of the impacts of FSRDC access on measures of research output for different types of researchers using split-sample regressions. Each panel shows the main effect of FSRDC access estimated from a separate regression restricted to theorists only (first coefficient) and empiricists only (second coefficient), thus exploiting only variation in treatment time across universities. The main dependent variables are the count of publications using Census data (panel (i)), the count of publications that cite papers using Census data (panel (ii)), the number of top publications (panel (iii)), and the citation-weighted number of publications (panel (iv)). Regressions include researcher, university, and year fixed effects. Standard errors are clustered at the researcher level.

Figure G6: Heterogeneous Effect of FSRDC Access

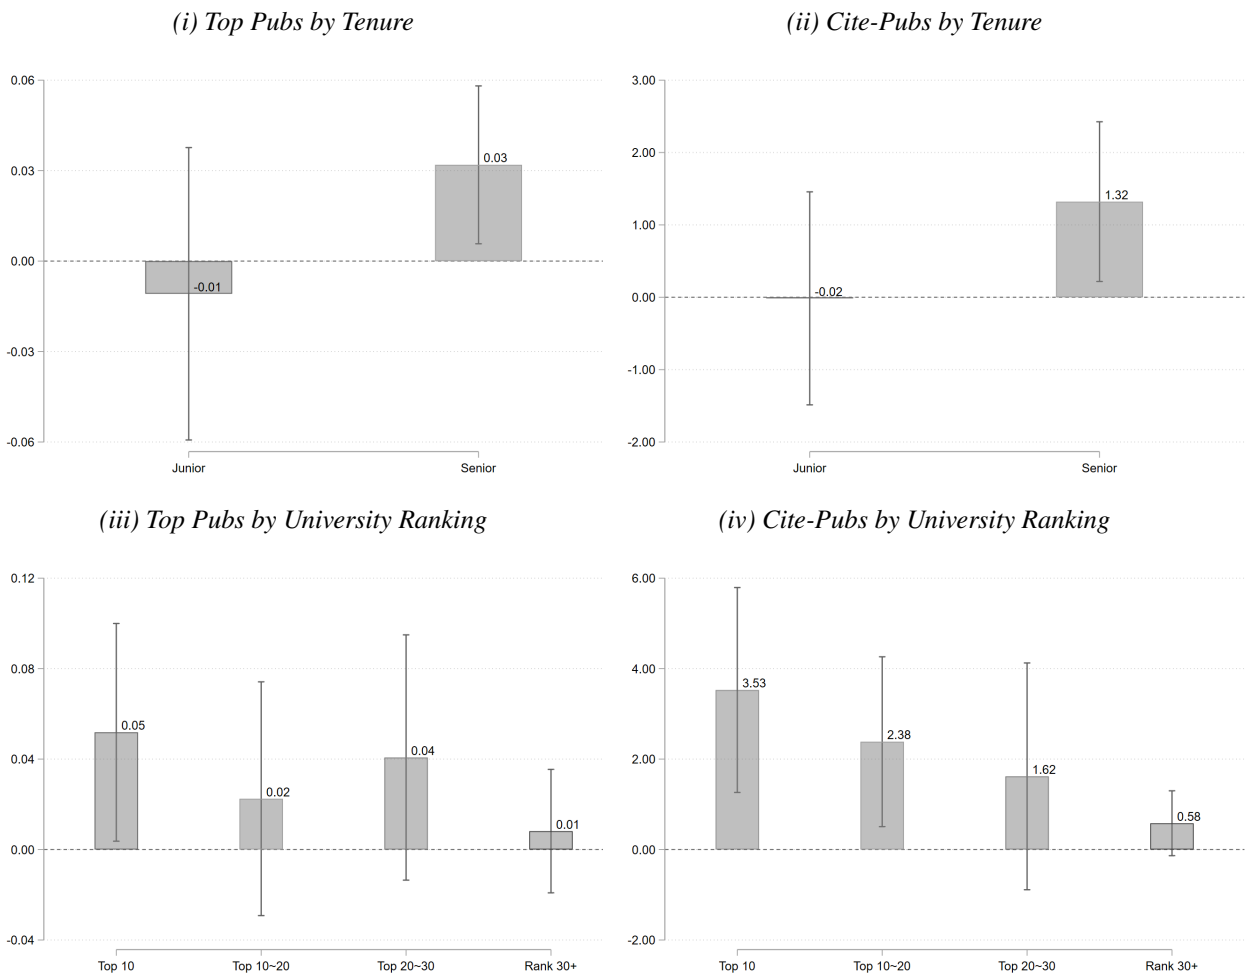

Note: This figure provides visual illustrations of the estimates from OLS models evaluating the impact of FSRDC access on research output. The main dependent variables are the number of top publications (Panel (i) and (iii)), and the citation-weighted number of publications (Panel (ii) and (iv)). The coefficients displayed are from split sample regressions. Junior: 0/1=1 for researchers that have started publishing less than seven years before. Information on the ranking of economics departments is taken from Kalaitzidakis et al. (2003). Regressions include researcher and university $\times$ year fixed effects. Standard errors are clustered at the researcher level. See text for more details.

Figure G7: University-Level Time-Varying Estimates of the Impact of FSRDCs

(i) Top Publications

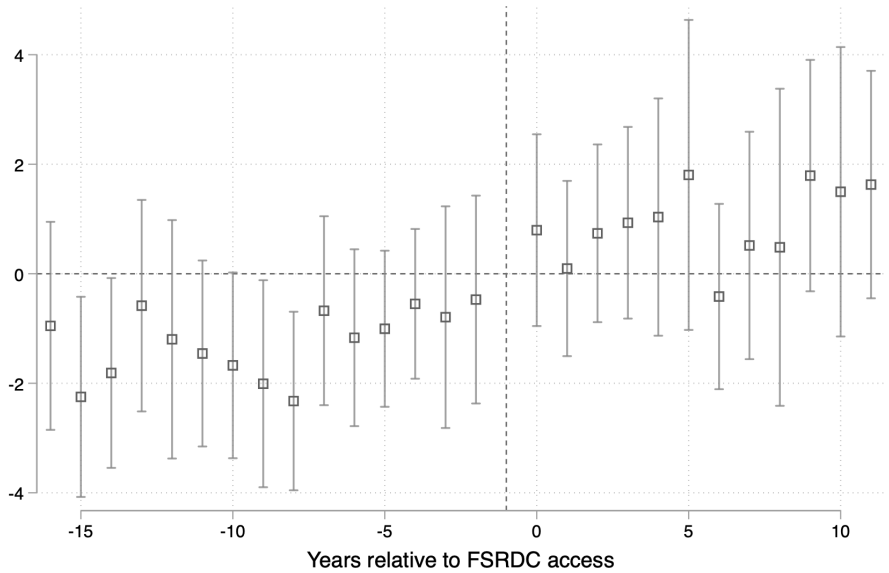

(ii) Cite-weighted Publications

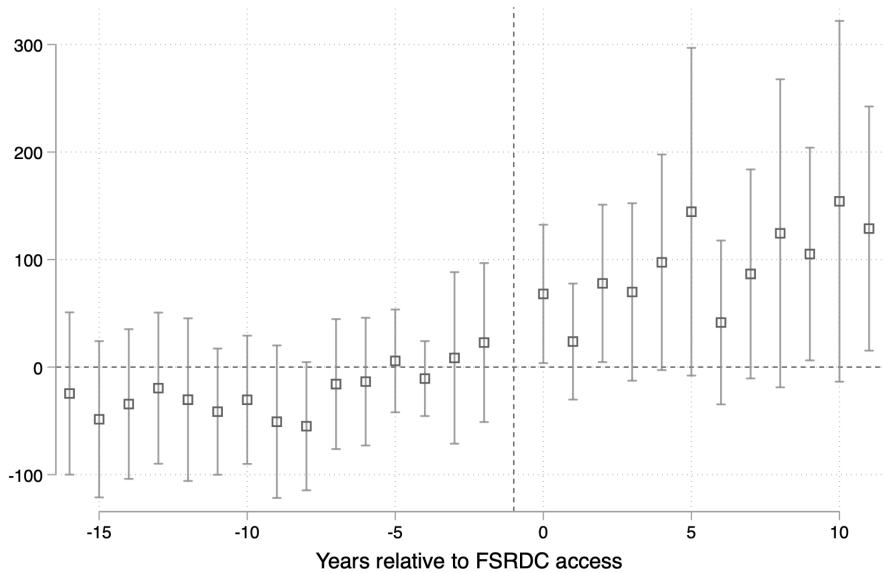

Note: This figure provides visual illustrations of the event study version of the regression evaluating the impact of FSRDC access on measures of research output at the university level. The main dependent variables are the count of articles published in the most prestigious economics journals (panel (i)) or the the count of articles weighted by the number of citations received up to 5 years following their publication (panel (ii)). The chart plots values of β for different lags before and after obtaining access to an FSRDC with 95% confidence intervals. All models include university fixed effects and year fixed effects. Standard errors are clustered at the university level.

Figure G8: Effect of FSRDC Access by Distance

(i) FSRDC Use

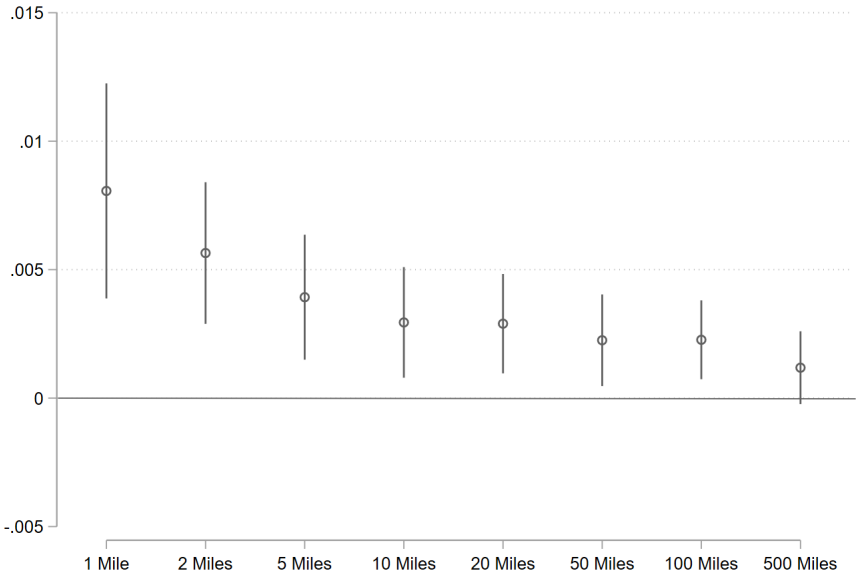

(ii) FSRDC Cite

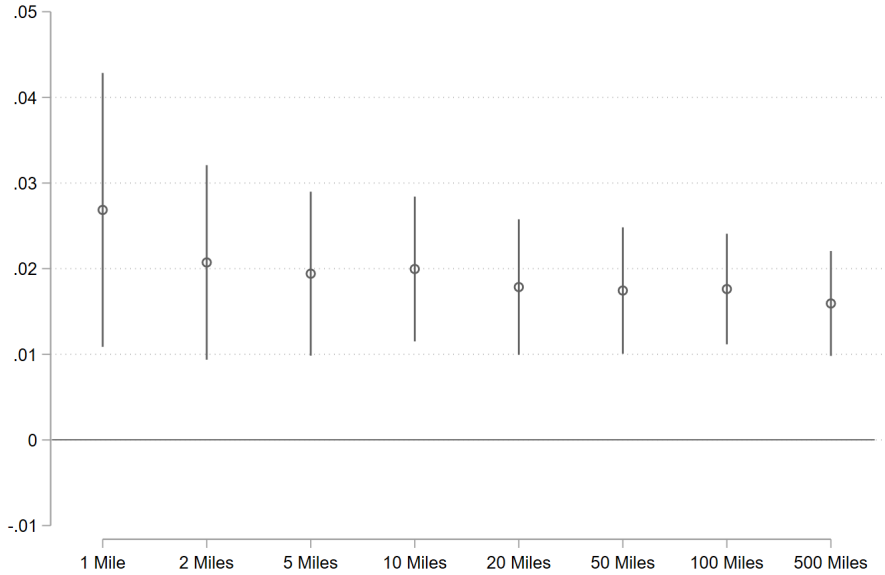

Note: This figure provides visual illustrations of the effect of data access at different distances. Distance is the geometric distance between the institution of the researcher and the closest operating FSRDC. The main dependent variables are the number of papers written using Census data (panel (i)) and the number of papers that cite FSRDC papers (panel (ii)). The chart plots values of β from different regressions where researchers are considered as treated if they are affiliated to institutions within a progressively larger radius from an FSRDC. Regressions include researcher and university $\times$ year fixed effects. Standard errors are clustered at the researcher level.

Figure G9: Effect of FSRDC Access by Share of Empirical Work

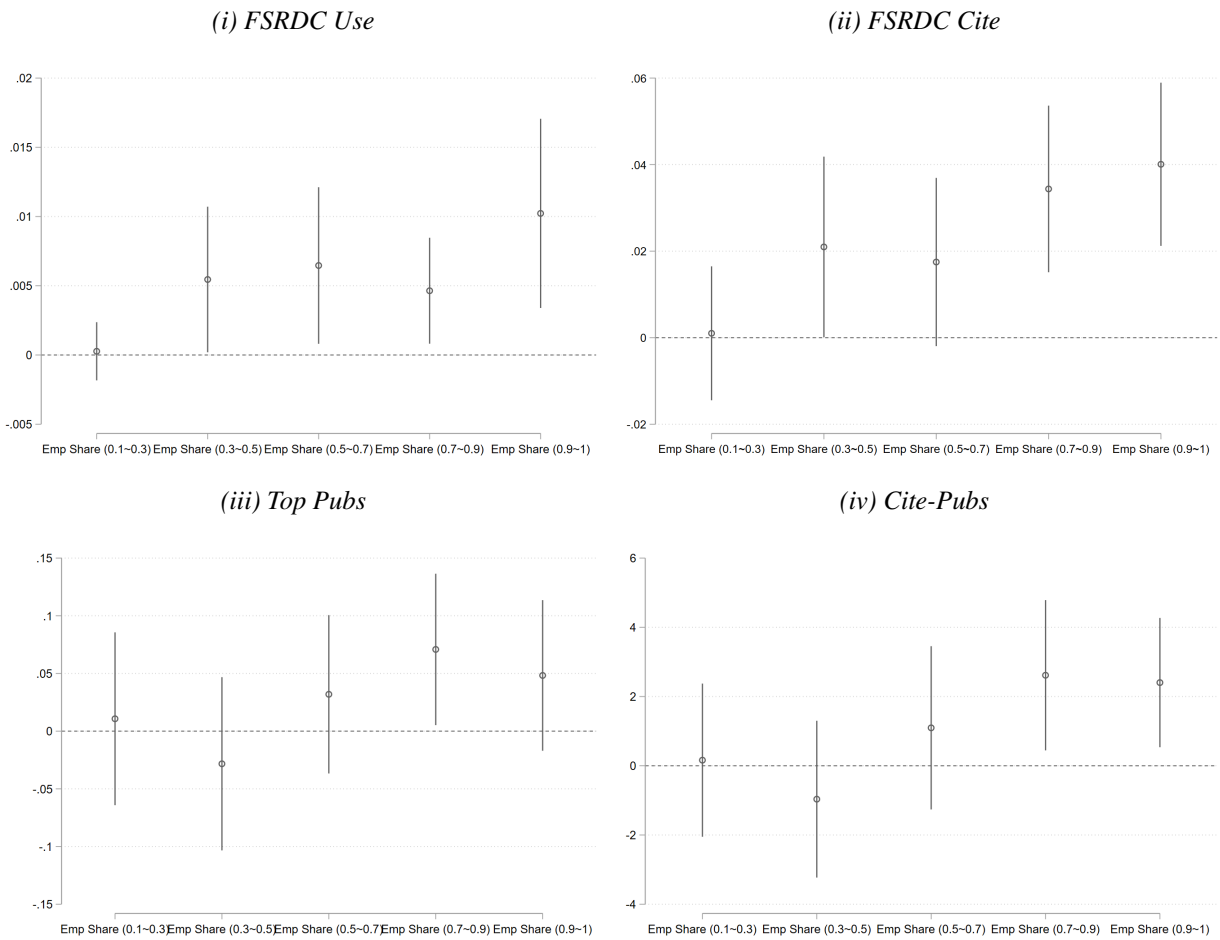

Note: This figure provides visual illustrations of the effect of data access for researchers with different methodological orientations. Empirical share is the proportion of a researcher's scholarship that is empirical in nature. The main dependent variables are the number of papers written using Census data (panel (i)), the number of papers that cite FSRDC papers (panel (ii)), the number of top publications (Panel (iii)), and the citation-weighted number of publications (Panel (iv)). The chart plots coefficients from a regression where researchers are classified into six mutually exclusive categories, depending on the share of their publications that is empirical in nature. The excluded category are researchers with no empirical papers. Standard errors are clustered at the researcher level.

Figure G10: Time-Varying Estimates of the Impact of FSRDCs on Research Output of Applied Scholars

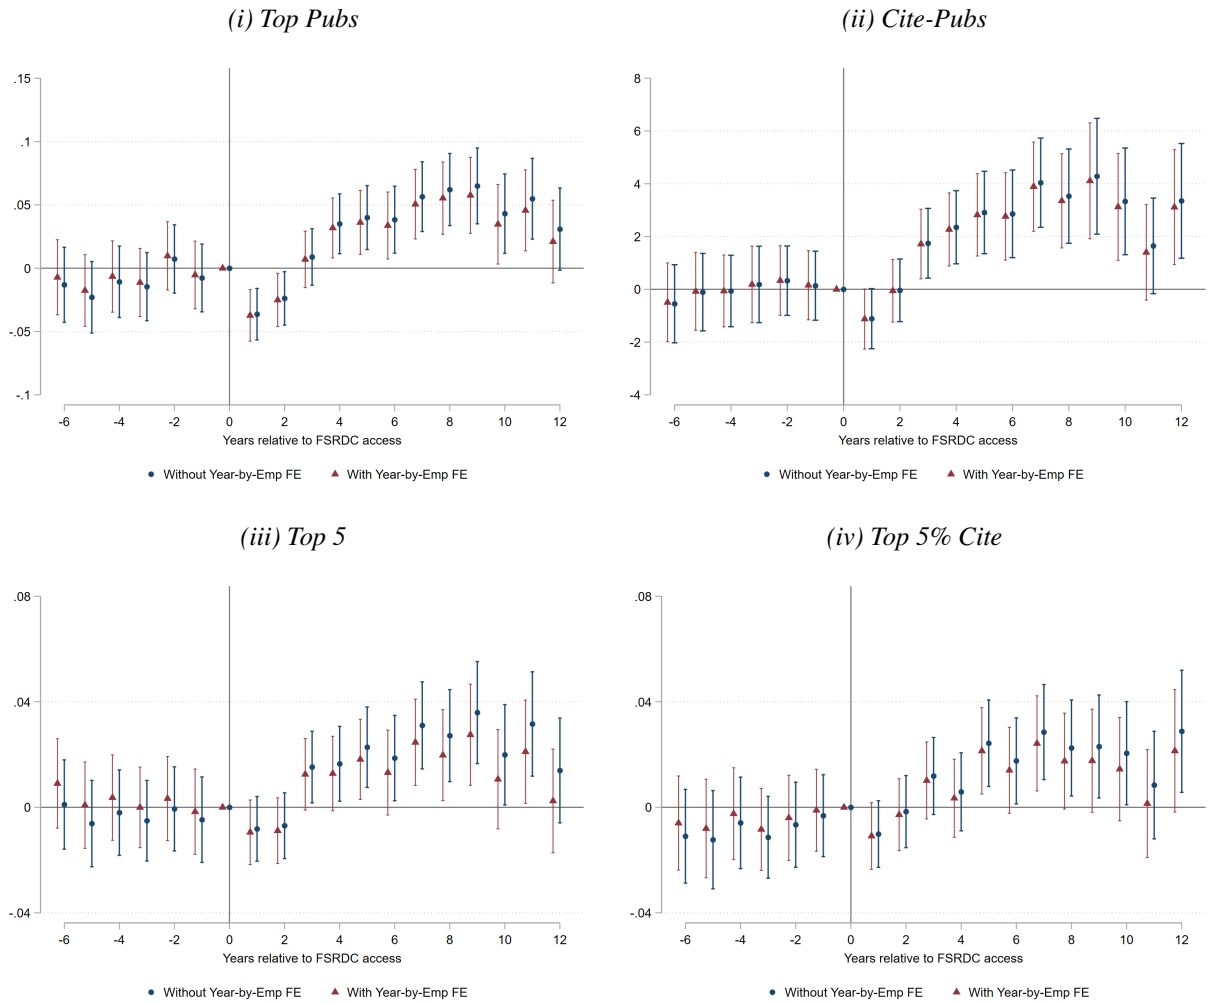

Note: This figure provides visual illustrations of the event study version of the main regressions evaluating the impacts of FSRDC access on measures of research output. The main dependent variables are the count of papers published in the main economics journals (panel (i)), the count of papers weighted by the number of citations received up to 5 years following publication (panel (ii)), count of papers published in Top Five journals (panel (iii)), and count of papers that are in the top 5% of the citations distribution (panel (iv)). The charts plot values of β for different lags before and after obtaining access to an FSRDC with 95% confidence intervals. Regressions include researcher and university × year fixed effects. Coefficients depicted with a triangle are from the same model estimated including additional time fixed effects for empirical researchers in bins of five years each. Standard errors are clustered at the researcher level.

Figure G11: Impact of FSRDCs on Policy Relevance of Academic Research

(i) Policy cites (raw means)

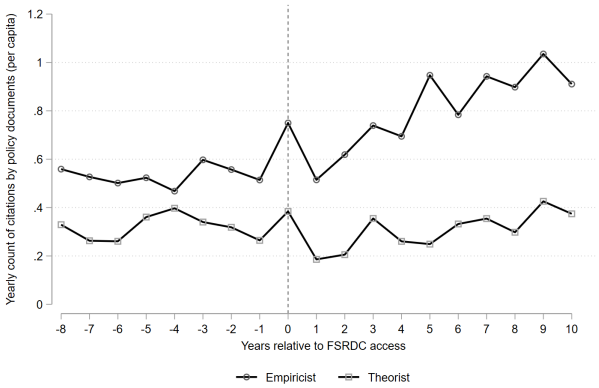

(ii) Papers cited by policy (raw means)

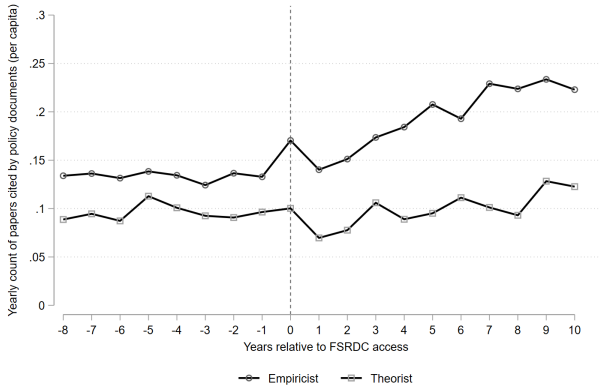

(iii) Policy cites (event study)

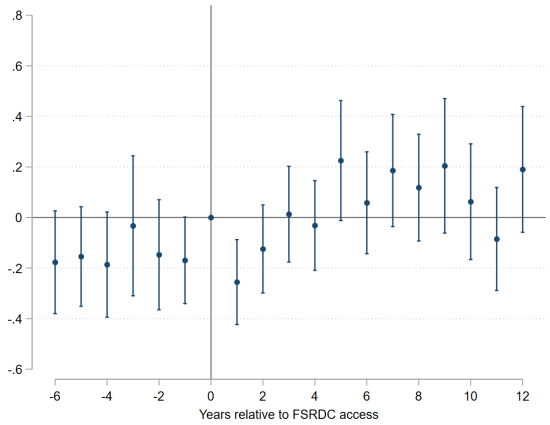

(iv) Papers cited by policy (event study)

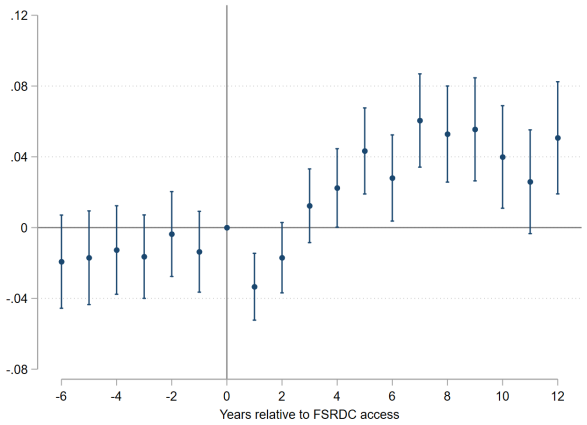

Note: This figure provides visual illustrations of the effect of access to administrative data on the policy relevance of scientific output. Panel (i) shows the raw means of the policy citations received by papers written by empiricists and theorists around the opening of an FSRDC. Panel (ii) shows the raw means of the number of papers with at least one policy citation written by empiricists and theorists around the opening of an FSRDC. Panel (iii) shows the event study version of the figure in Panel (i), while Panel (iv) does the same for the figure in Panel (ii). The charts in Panels (iii) and (iv) plot values of β for different lags before and after obtaining access to an FSRDC with 95% confidence intervals. Regressions include researcher and university × year fixed effects. Standard errors are clustered at the researcher level.