The Effect of Funding Delays on the Research Workforce: Evidence From Tax Records*

Preliminary: Please do not share without permission

Stephanie Cheng† Elisabeth Perlman‡ Joseph Staudt§
Wei Yang Tham¶

2022-07-15

Abstract

We study how an interruption in the flow of research funding to a major NIH grant — the R01 — affects the career outcomes of research personnel using comprehensive wage and tax records that have been linked to university transaction data. Using a difference-in-differences design, we find that for employees who work for labs with fewer grants, an interruption of more than 30 days has a substantial effect on job placement, including a 2.5 pp increase in the probability of no longer working in the US. The effects are strongest among postdocs and graduate students. We also find that interrupted employees who stay in the US labor force earn 20% less than uninterrupted employees who also remain in the US.

*This paper uses data from the U.S. Census Bureau. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this data product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. (DRB Approval Number: CBDRB-FY21-CES009-002).

†Edgeworth Economics

‡US Census Bureau

§US Census Bureau

¶Harvard University, wtham@hbs.edu

"My current job started 8 years ago when my boss told me he had 6 months of guaranteed funding. I worked for him full-time for 4 years, my salary cobbled together from a half-dozen grants over that time. . . (W)hile my skills are undoubtedly valuable to a research lab, it is incredibly difficult for someone like me to find a stable job because of the funding issues and lack of recognition of the value of a supertech position." - Anonymous lab technician/manager (Guzey 2019)

1 Introduction

In 1945, Vannevar Bush – director of the Office of Scientific Research and Development and a key figure in mobilizing science for World War 2 – submitted a report, *Science: The Endless Frontier*, to the President of the United States. The report contained Bush's recommendations for how the federal government should support research. Guided by the premise that basic research takes time to yield rewards, Bush called for the creation of a federal science agency that was provided with "stability of funds so that long-range programs may be undertaken." One vital use of these funds was to "strengthen" universities, which he viewed as essential to the production of basic research because they were "least under pressure for immediate, tangible results." While *The Endless Frontier* served as a blueprint for science policy in the post-WW2 era, the research environment today has evolved in a way that compromises on Bush's vision of stable funding for universities.

Instead, universities have become highly dependent on federal funding, leaving them exposed to periods of instability such as the boom and bust in NIH funding from 1998 to 2003 (Stephan 2013; Freeman and Van Reenen 2009). However, the problem is not only a matter of changes in funding levels. Federal science agencies need the US government to allocate them a budget every year, and with that comes uncertainty over when Congress will pass the federal budget (not to mention the non-negligible threat of a government shutdown). The unpredictability of when a given year's budget will be passed is such a regular occurrence that it might even be considered a permanent feature of the scientific funding landscape.[2]

Naturally, researchers are concerned that budgetary delays at the federal level may lead to delays in the disbursement of grant funding (Fikes 2018; DrugMonkey 2009, n.d.; Mervis and Marshall 1996), even if they do receive the money eventually. There could be an effect on Principal Investigators (PIs) if they raise their salary through grants or on the personnel they hire, from trainees (e.g. graduate students, postdocs) to staff (e.g. research scientists, lab managers). However, there are several challenges to quantifying these effects. One is finding "micro" variation for an aggregate-level shock (i.e. everyone is affected by the federal budgeting process). Another is defining and measuring a delay. The notion of a delay means that funding arrived later than it "should have," but without access to data that is usually confidential, we are more likely to observe when funding *did arrive* instead of when it "should have" arrived. Finally, even if those challenges are solved, there is the task of linking delays to the characteristics and outcomes of the people they might affect, beyond the PIs of grants.

This paper attempts to overcome these challenges and quantify the effects of funding delays. We address the first two challenges (finding variation and defining delays) by focusing on a major NIH grant – the "R01." R01s are generally regarded as being necessary to establish an independent research lab in the biomedical sciences. They are usually granted for four to five years, after which the Principal Investigator of the R01 has to apply to renew it for another term. Renewal is not guaranteed, but even if the PI's application is successful, there can be a delay between the expiration of funds pre-renewal and the disbursement of funds upon renewal. Using public grant data (NIH ExPORTER), university administrative data on grant transactions (UMETRICS), tax records (W-2 and 1099) and unemployment insurance records that have been linked, we are able to connect the length of delays in R01 renewals to employees who were in labs undergoing an R01 renewal, observing their earnings and *all* job placement within the US.

[2]For example, the National Institute of Allergy and Infectious Diseases (NIAID), which accounted for 14% of the NIH's budget in FY 2020, explicitly addresses this issue in an online guide to the grant application process, stating that it is "assiduous about issuing awards using funds from the CR." "CR" stands for continuing resolution.

Our research design builds on Tham (2021), which focuses on how funding delays affect spending patterns and overall lab productivity. Following their analysis, we define R01 grants that were renewed after more than 30 days as "interrupted" and R01 grants that were renewed within 30 days as "continuously funded" or "uninterrupted."[3] We use a difference-in-differences design that compares the outcomes of employees of a lab with an *interrupted* R01 with the outcomes of employees of a lab with a *continuously funded* R01 before and after R01 renewal occurred. Our estimation procedure combines "stacking" (A. C. Baker, Larcker, and Wang 2022; Cengiz et al. 2019) and the estimator from Callaway and Sant'Anna (2020).

To take into account that PIs with more grants may use their additional grant money to make up for the delayed funding, we split the sample by whether an employee is working with a PI with only one R01 and a PI with multiple R01s. This is also motivated by the results in Tham (2021), which shows that when PIs of an R01 experience an interruption, PI spending across all their NIH grants decreases substantially until the grant is renewed, but the decrease in spending is concentrated among PIs with only one R01, while spending decreases much less for PIs with multiple R01s.

We first describe the results for employees working with a single-R01 lab (i.e. more funding constrained). We find that funding interruptions do have a long-term effect on an employee's job placement. We define three mutually exclusive categories of job placement: (1) working at a US university, (2) working in industry in the US, or (3) not working at a US entity (we clarify how we define these categories in the Results section). Employees are 5 percentage points less likely to work at a university for the first two years after an interruption. About half (2.5 percentage points) of the departure from US universities is due to employees working in US industry and the remaining half is due to employees leaving the US labor force. The half of employees who leave for industry move back to a university job after 2 years. However, the departure of employees from the US labor force persists at 2.5 percentage points up to 6 years after interruption.

[3] Our choice of 30 calendar days is meant to approximate a month – grants are usually funded on the first of the month, thus the arrival of new grant funding can be thought of as occurring on a monthly basis.

Given the importance of immigrants to the scientific enterprise, one question of interest is whether non-US citizens are more likely to be affected by funding interruptions due to the requirements of maintaining their immigration status in the US, such as being employed.[4] Since we do not observe immigration status, we instead estimate results separately for US-born and foreign-born employees.[5] Surprisingly, we find that US-born employees are driving the job placement results. One possible explanation may be that US citizens are more willing to take jobs outside of the US since it would be easier for them to return while non-US citizens try harder to find (possibly lower quality) jobs in the US.

We also explore whether these results are driven by certain employees in certain occupations. We organize employees into three aggregated occupation categories: faculty, postdocs/graduate students, and "others" (which includes occupations such as staff, research scientists, and undergrads). Our results on placement changes are mostly driven by non-faculty employees, particularly postdocs and graduate students. Specifically, postdocs/grad students leave academia, and instead of moving to a non-university job in the US, they exit the US labor force entirely (possibly taking a job abroad).

Next, we look into whether interruptions affect the earnings of employees. Interruptions may have a negative effect on the earnings of employees if they force employees to find a new job or delay career progression (e.g. from postdoc to faculty). However, earnings may also increase in some cases. For example, a graduate student forced to graduate early may take a job in the private sector instead of academia and have higher earnings than if they had stayed in academia. We find that on average, the earnings of interrupted employees in single-R01 labs decrease substantially. However, though always negative, the magnitude of our estimates depends on how we define the sample.

Our data *does not* include earnings outside the US. Thus, if an employee leaves their US job to work in a non-US job, then their earnings appear in our data as zero. While this is accurate in a narrow sense, it is clear that what we are really interested in is the effect of interruptions on earnings regardless of country. To get around this limitation

[4] For example, foreign-born individuals comprised more than 45% of PhDs across Science & Engineering occupations except the Social Sciences Foundation (2018)

[5] Although "non-US-born" is more accurate, we use the term "foreign-born" for easier reading.

of our data, we estimate the effect of interruptions on earnings for a subsample of "fully attached" employees who report earnings in the US every year starting from one year before interruption to five years after interruption, (implying that they reported earnings in the US for at least seven years consecutively). We find that interrupted employees who continue to work in the US experience a 20% decrease in earnings.

We repeat the same analyses on employees working for a PI with multiple R01s and find null effects. We interpret this as evidence for the following mechanism: interruptions constrain a PI's funding *and* for some PIs, other sources of funding (e.g. university-provided bridge funding) are not enough to make up for that loss in funding, which in turn leads them to reduce hiring. This mechanism makes it unlikely that the effects we observe for interrupted single-R01 employees are actually driven by interruptions rather than being driven by time-varying confounders that happen to affect both the likelihood of an interruption and employee outcomes.[6]

Overall, we find that funding interruptions have a sizable impact on the earnings and job placements of members of the scientific workforce. These impacts are concentrated on non-faculty, particularly postdocs and graduate students, and even lead to departures from the US entirely. Even though all labs in our sample are eventually funded, even differences in *timing* can have a meaningful impact on research personnel.

Our work contributes to a fuller picture of the relationship between uncertainty and knowledge production. Research and policy solutions have tended to focus on whether faculty or principal investigators would do riskier science in less uncertain environments (e.g. Azoulay, Graff-Zivin, and Manso (2011); "What We Learned Doing Fast Grants" (n.d.)). However, the importance of team science (Wuchty, Jones, and Uzzi 2007) means that these concerns extend beyond faculty to the careers of non-faculty personnel supported by research grants. There is some evidence that scientific careers can be affected meaningfully by events that initially seem small (Azoulay, Greenblatt, and Heggeness 2021; Hill 2019) and our results indicate that interruptions are one such event. Given the yearly recurrence

[6]This is a similar idea to a placebo test, although ex-ante we do not know that the effects of interruptions on the multiple-R01 sample are necessarily zero.

of uncertainty about the federal budget, the *how* – and not just *how much* – of science funding has important ramifications for the scientific labor force.

The rest of the paper consists of the following subsections, which the reader can jump to by clicking on the following links: Background, Data, Estimation, Results, Conclusion

2 Background[7]

2.1 NIH funding

The NIH is responsible for an annual budget of about US$40 billion, much of which is disbursed through research grants. A core part of the NIH's mission is funding basic science to generate fundamental knowledge that tends to have long-term rather than immediate impact.

The NIH is funded every fiscal year by congressional appropriation. This is part of a broader process whereby the US Congress passes regular appropriations bills to fund a wide range of government operations.[8] If appropriations have not been made by the beginning of the fiscal year, Congress can enact a "continuing resolution" to provide temporary funding. If a continuing resolution is not enacted and a "funding gap" occurs, then federal agencies have to begin a "shutdown" of projects and activities that rely on federal funds.

It is typically taken as given that regular appropriations will not have been made by the beginning of the fiscal year on 1 October, and that federal agencies will have to operate under a continuing resolution for at least some portion of the year. Under a continuing resolution, the NIH continues to fund existing projects, albeit at a reduced rate initially. However, it might also choose to delay funding for new or renewed projects in response to uncertainty about the size of the NIH's budget for the fiscal year.

To illustrate, suppose that at the beginning of the fiscal year, the NIH knows (1) its budget

[7] This section of the paper borrows heavily from Tham (2021).

[8] A fiscal year is identified by the year in which it ends. E.g. FY 2001 started on 1 October 2000 and ended on 30 September 2001.

and (2) its own ranking of projects available to be funded (rank could be based on project quality but also other factors such as NIH priorities). In this scenario, the NIH knows which projects it wishes to fund *and* whether it can fund them before the projects are set to run out of funding. Thus, there are no funding interruptions.

The scenario above illustrates that funding interruptions arise from uncertainty about either (1) the NIH's budget or (2) the quantity and quality of projects that need funding that fiscal year, or both. Some uncertainty over projects is built-in as there are three review cycles throughout the fiscal year.

2.2 R01 grants

The R01 is the largest grant mechanism through which the NIH funds extramural research. It is designed to provide enough funding to establish an independent research career. An R01 *project period* lasts for 4-5 years, after which it must be renewed in order to receive additional funding for a subsequent project period.[9] Thus, the same *project* can last for multiple *project periods*.

Principal Investigators (PIs) generally want to maintain R01 funding for as long as possible, so as their current project period ends, they have to apply to renew their project for another 4-5 year project period.[10] In order to avoid lapses in funding between two project periods, PIs usually start to apply for renewal 1-2 years before a project period ends. This allows time to prepare the renewal application itself as well as time to resubmit an application that is rejected the first time.

The focus of our paper is on the effects of lapses in funding between two project periods. Thus, we exclusively analyze projects that are successfully renewed at least once and so span multiple project periods, even though it is possible that a project is not renewed upon expiry.

[9] They can also be shorter (1-3 years), but this is uncommon.

[10] R01 renewal is sometimes even listed as a criterion for receiving tenure e.g. Medicine (2020).

2.3 Where do funding interruptions come from?

Suppose that at the beginning of the fiscal year, the NIH knows (1) its budget and (2) its own ranking of projects available to be funded (rank could be based on project quality but also other factors such as NIH priorities). In this scenario, the NIH knows which projects it wishes to fund *and* whether it can fund them before the projects are set to run out of funding. Thus, there are no funding interruptions.

The scenario above illustrates that funding interruptions arise from uncertainty about either (1) the NIH's budget or (2) the quantity and quality of projects that need funding that fiscal year, or both. Some uncertainty over projects is inherent in the review process, as there are three review cycles throughout the fiscal year. However, the federal budgeting process plays a role as well. Figure **??** shows for Fiscal Years 1998 to 2018 the relationship between when the federal budget was passed that fiscal year and when the budgets of NIH grants started on average. When a grant is "Continuing" (e.g. in the third year of a 5-year grant), there is little to no relationship. But for grants that had to be competed for (i.e. "New" or "Renewed" grants), budgets tend to start later in the fiscal year if the federal budget was passed later.

2.4 Why interruptions might or might not affect employee outcomes

The most direct effect of an interruption is that it requires a PI to use the same amount of funding over a longer period of time. Since personnel tend to form the bulk of expenses of a grant this may affect the ability of a PI to continue supporting members of their lab. Many lab members, such as postdocs, are often on short-term contracts that may simply not be renewed if funding is not readily available.

However, interruptions are not a new phenomenon. employees and research institutions (including the NIH) are aware of the possibility of interruptions. Thus, they may have developed ways of mitigating the disruptive effects of interruptions. For example, a PI's home institution may be able to provide bridge funding while a PI waits for an interruption to be resolved. However, the availability of bridge funding is also likely to both across institutions and also across employees within the same institution.

3 Data

The overarching goal of our data construction is to identify personnel who were part of labs that had an expiring R01 *that was eventually renewed* (for brevity, in the rest of the paper we sometimes refer to these R01s as simply "expiring" without qualifying that they were also renewed). We then define whether these personnel were part of an interrupted lab or not based on the length of time between expiry and renewal. Finally, we track the earnings and job placement of these personnel before and after expiration.

The components required to achieve this are data on (1) which R01s were expiring, (2) which personnel were part of labs that depended on those R01s, and (3) the wage and job placement outcomes of those personnel. We obtain these data from (1) ExPORTER – a public database of NIH grants, (2) UMETRICS – administrative data from universities on grant transactions, including payments to personnel, and (3) Census data including the universe of W-2 and 1099 tax records and the universe of unemployment insurance (UI) wage records. Together, these data allow us to identify individuals who work in labs that potentially experienced a funding interruption and track their entire US employment and earnings history.

Figure 1 provides a graphical representation of the basic steps to accomplish this:

1. **ExPORTER** Find all instances of R01s expiring and then being successfully renewed. This returns a set of *focal R01s* that the rest of the data construction builds on.
2. **ExPORTER** Each focal R01 has at least one PI. For each PI, form a focal-R01-by-PI pair.
3. **ExPORTER** For each focal-R01-by-PI, find all grants awarded to the PI in the 12 months prior to expiry of the focal R01
4. **ExPORTER + UMETRICS** Find all employees paid by focal-R01-by-PI through any of the grants in the previous step. These employees constitute the PI's "lab."
5. **UMETRICS + Census** Merge employees with Census earnings and placement data

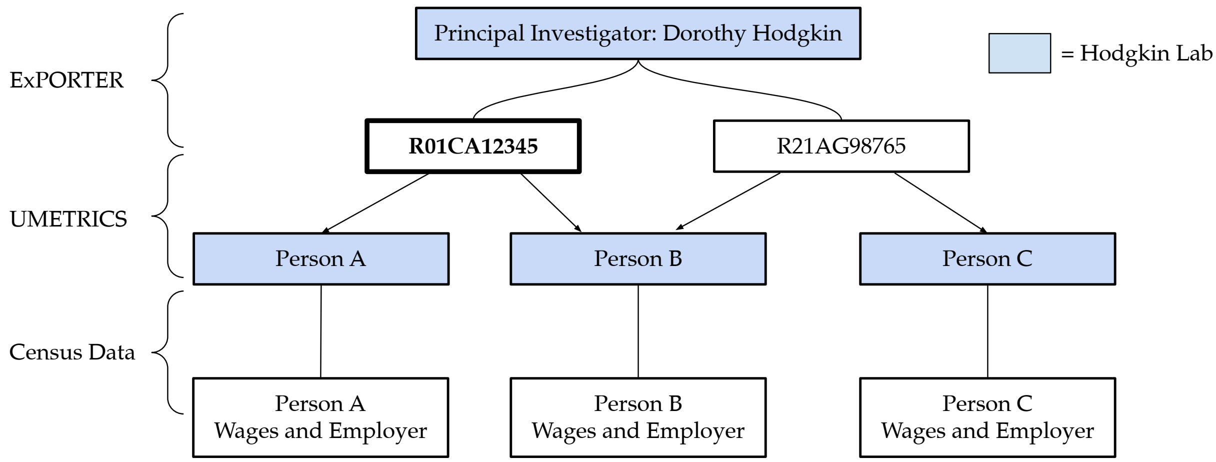

Figure 1: This diagram shows the process of linking R01 grants to employees and their outcomes, starting with NIH ExPORTER data at the top and ending with tax and unemployment insurance records stored at the US Census Bureau.

The remainder of the section goes into more detail about this process and variable construction. Additional detail is also available in the Data Appendix.

3.1 ExPORTER

ExPorter is publicly available data provided by the NIH on NIH grants.[11] We use the information in ExPORTER to generate the following variables (details in Data Appendix).

Length of funding gaps and interruptions. Number of calendar days between the end of a project period and the beginning of the next project period. This is used to define whether an R01 was interrupted or continuously-funded.

PI grant portfolio (Number of R01s). Tham (2021) shows that PIs with multiple R01s decrease spending by a lower magnitude, which suggests that employees working for PIs with multiple R01s may be less affected by funding interruptions. Since ExPORTER has PI identifiers (PI IDs), we can find all NIH grants awarded to a given PI. Specifically, we find all grants that a PI was awarded between the dates of one year before and after the focal R01 expired. We include R01s after focal R01 expiry since given the time it takes from

[11] https://exporter.nih.gov/

applying for a grant to receiving the funds, it is unlikely that a funding interruption would affect the number of R01s a PI has within that time window. We define the size of the PI's grant portfolio based on the number of "R01-equivalent" grants, including the focal R01. For brevity, we refer to this variable as the "Number of R01s" without explicitly defining the other types of grants included.

PI lab grants. We want to identify employees who are part of a PI's lab and thus most likely to be affected by an interruption. For a given PI ID of an expiring R01, we start by identifying all NIH grants awarded to the PI ID in the 12 months prior to the focal R01's expiry. We then define all employees paid by any of those grants in the 12 months prior to R01 expiry as the PI's "lab."

3.2 UMETRICS

UMETRICS is a database of administrative transaction-level data on payments made from university research grants to employees and vendors. It is housed at the Institute for Research on Innovation and Science (IRIS) at the University of Michigan and is derived from university human resources records, sponsored projects, and procurement systems made available by participating universities. We use the 2019 release of UMETRICS, which contains data from 31 universities representing about one-third of US federal research expenditures (IRIS 2019).

For each university, we observe all expenditures from research grants (within the time period for which the data was provided). We are particularly interested in payments from NIH R01 grants to personnel. For each of these payments, we observe the dates of the transaction and the occupation of the employee when the payment occurred, so occupations can change over time (e.g. a post-doc may become a faculty member).[12]

The NIH grants in ExPORTER have been linked to NIH grant transactions in UMETRICS, allowing us to identify employees that were paid a PI whose funding is interrupted. In turn, these UMETRICS employees have been linked to their W-2 and 1099 tax records and

[12]Employee occupations are assigned by the UMETRICS data team using information such as job titles.

their UI wage records, allowing us to observe wages and employer characteristics.

3.3 Census tax and employment data

UMETRICS employees have been linked to Census data using probabilistic matching (Wagner, Layne, et al. 2014), which allows us to track career outcomes both before and after their lab experiences (or does not experience) a funding interruption. We use three sources of confidential administrative data available from the US Census Bureau to derive these career outcomes.

W-2 tax records. Form W-2 is an Internal Revenue Service (IRS) form that US employers have to file listing the wages paid to an employee and taxes withheld them. Each W-2 record is an employee paired with the federal tax identification number (EIN) of an employer, and contains information on yearly wages.

Longitudinal Employer-Household Dynamics (LEHD). LEHD data contain Unemployment Insurance (UI) wage records which track earnings and employment at a quarterly, rather than annual, frequency. Since UI programs are administered at the state level, each record is an employee paired with the state tax identification number (SEIN) of the employer. However, the federal EIN is also available for most employee-SEIN pairs, which enables us to identify which employers are universities. The LEHD is also linked to NUMIDENT data from the Social Security Administration (SSA) which enables us to identify whether an individual was born in the US or not.

1099 tax records. 1099 tax records are available through the US Census Integrated Longitudinal Business Database (ILBD) and contain the population of all non-employer firms in the United States. These 1099 records capture earnings from self-employment.

IPEDS. We link the W-2 and LEHD data to a public-use list of university EINs from the Integrated Postsecondary Education Data System (IPEDS), so that we can determine whether an individual is paid by a university.[13] IPEDS contains EINs for most U.S.-based universities, and all UMETRICS universities are in IPEDS.

[13] The public-use list of university EINs from IPEDS can be found at https://nces.ed.gov/ipeds/datacenter/DataFiles.aspx under the title "Directory Information." We combine the datasets from 2002 to 2018.

With these data, we want to understand how job placement and earnings might be affected by funding interruptions. To do so, we construct the variables below.

Sector indicators. We define three mutually exclusive categories to represent the sectors that an employee can belong to in a given year:

1. US university – received positive earnings from an IPEDS university
2. US non-university sector – *only* received positive earnings from a non-IPEDS US employer
3. Not employed in US – did not receive earnings from an employer in W-2, LEHD, or ILBD data (complement of (1) and (2)).

University placement indicators. To understand employee movement *within* the university sector, we construct two more indicators for university employment. For each, we construct the following indicators:

1. Employee receives positive earnings from *their own university* (i.e. the university employing them at the time of R01 expiry) and
2. Employee receives positive earnings from an IPEDS university other than their own.

Wages. We observe yearly wages for each employee from 2005 to 2018. These are derived from a combination of W-2, LEHD, and 1099 (ILBD) wages. We define an employee's total earnings in a given year to be their wages from self-employment (ILBD) plus the maximum of their W-2 and LEHD wages. That is, $wage_{total} = wage_{ilbd} + max\{wage_{W2}, wage_{lehd}\}$.[14]

3.4 Descriptive Statistics

Figure 2 shows the distribution of funding gaps for R01 renewals across all NIH grants from fiscal years 2005 to 2018. Figure 2A shows that about 20% of R01s were "interrupted" according to our definition of an interruption being a funding gap larger than 30 days. Figure 2B shows the overall distribution, which is right-skewed and has a median of 88 days.

[14] The LEHD receives data from individual states unemployment insurance systems and there are two gaps that are particularly important to this study: a) Massachusetts data is not in the LEHD until 2011, and b) graduate student stipends are not covered by unemployment insurance and thus not reflected in LEHD data.

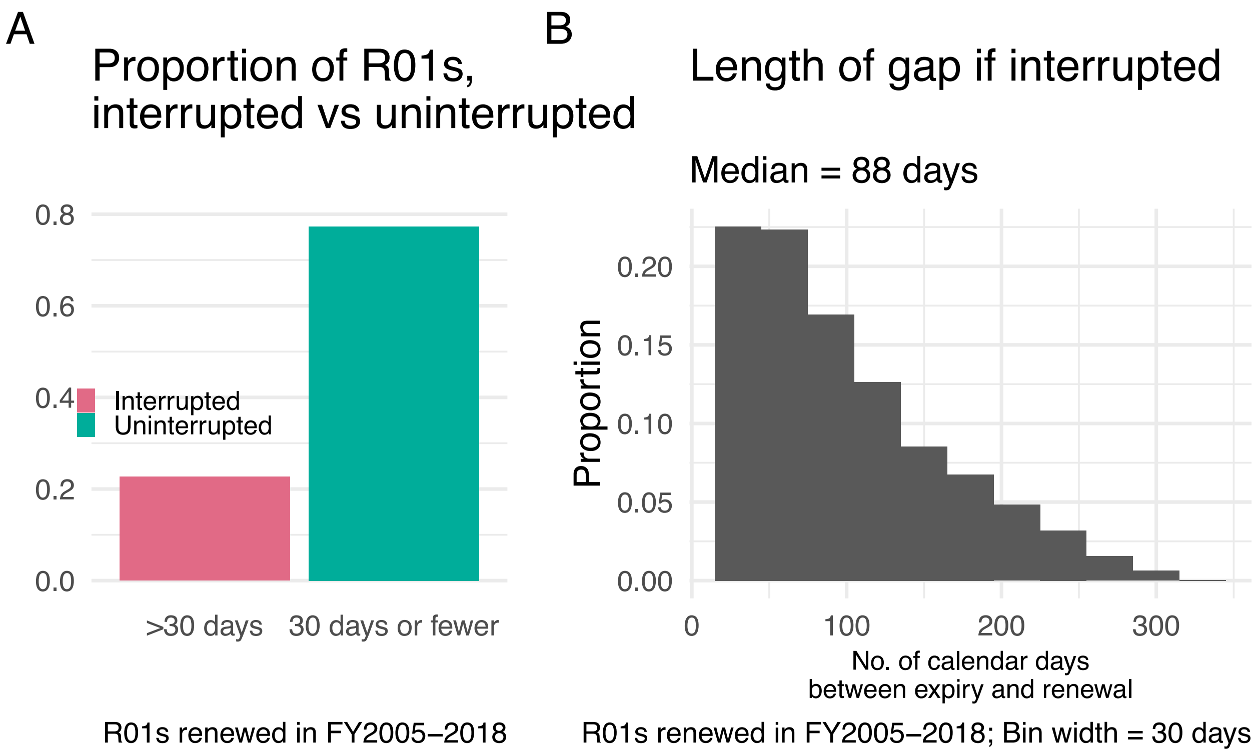

Figure 2: This figure shows the distribution of funding gaps for renewed R01s expiring in Fiscal Years 2005 to 2018. The figure on the left, A, shows the proportion of R01 grants that are interrupted, according to our threshold of 30 days. Figure B on the right shows the overall distribution of funding gaps.

Table 1: This table shows the number of unique individuals by occupation category (at the time of interruption) and birthplace for individuals belonging to single-R01 labs.

Type	N	type
Faculty	900	Occupation
Postdoc/Grad students	1300	Occupation
Others	2000	Occupation
US-born	2700	Birthplace
Foreign-born	1400	Birthplace

Our final sample of individuals used in the analysis consists of about 4,200 unique individuals belonging to single-R01 labs and about 13,500 unique individuals belonging to multiple-R01 labs. Table 1 shows the breakdown of the individuals belonging to single-R01 labs by occupation and then by birthplace.[15]

4 Estimation

4.1 Stacked Difference-in-differences

We use a difference-in-differences (DiD) design to estimate the effect of funding interruptions on employees. Our estimation method involves two steps:

1. "Stacking" the data by event (i.e. R01 expiration year)

2. Using a modified version of the Callaway and Sant'Anna (2020) estimator ("CS estimator" from here on) to compare treatment and control units within the same event.

Step (1) is a data formatting step where we make explicit choices about what control units to use for each cohort of treated units and Step (2) is the application of an estimator to the reformatted data. Each step is an independent choice. That is, one could stack the data and use a different estimator (e.g. OLS with two-way fixed effects (A. C. Baker, Larcker, and Wang 2022; Cengiz et al. 2019)) or directly apply the CS estimator on a standard unit-time

[15]The summation across categories may not always equal the total because of rounding.

panel dataset.

In a staggered DiD, the CS estimator estimates disaggregated "group-time" treatment effects, where groups are defined by time of treatment, which can then be aggregated as desired (e.g. as a static treatment effect or by time relative to treatment for an event study). However, control units are not assigned to groups and are used as a comparison as long as they remain untreated. This is in line with many applications where it is harder to define a counterfactual treatment time (e.g. the year in which a state might have but did not pass a minimum wage increase). In our setting, however, both interrupted and continuously-funded R01s have an expiration date. The expiration year can therefore be used as a well-defined treatment period for both treated and control employees.

We take advantage of this feature to group treatment and control employees with the same expiration year into cohorts.[16] Each cohort can be thought as a separate DiD with only one treatment period. This means that we are more likely to be comparing employees in labs with projects and budgets that are at similar stages in their lifecycle (e.g. employees may be less likely to leave their job at the beginning of an R01 than at the end of an R01, so using them as control units will overstate the effect of an interruption).

These cohorts are then "stacked" to form a unit-by-time panel dataset where units are defined by an employee-cohort pair and time is defined by years relative to treatment year. Once the data is formatted, estimation is typically done with a modified two-way fixed effects estimator (with unit-cohort fixed effects and time-cohort fixed effects)(A. C. Baker, Larcker, and Wang 2022). We instead use the CS estimator because it allows for transparent and flexible aggregation of the group-time treatment estimates and has more modeling options (including a doubly robust estimator).[17]

Within the set of expiring R01 grants that were eventually renewed, we define R01s as either (1) "interrupted" if were renewed after more than 30 calendar days or (2) "continuously funded" or "uninterrupted" if they were renewed in 30 calendar days or less.[18] In turn, an

[16]This can be thought of as exact matching on treatment time.

[17]CS aggregates treatment effects by group size. Two-way fixed effects implicitly uses OLS weights which is more efficient at the cost of bias (A. C. Baker, Larcker, and Wang 2022).

[18]Our choice of 30 calendar days is meant to approximate a month – grants are usually funded on the first

employee is interrupted or treated if they were part of a lab with at least one interrupted R01, and they are in the control group if they were part of a lab with an expiring R01 that was continuously funded, but not part of any labs with an interrupted R01.

One consideration in the construction of treatment cohorts is that control units are "clean" i.e. not treated (or experiencing the effects of treatment) in the time window of interest. For example, suppose an employee is in a lab with two R01s that are expiring in consecutive years, 2001 and 2002. The first R01 is continuously-funded in 2001, but the second R01 is interrupted in 2002. This employee would be a control unit in the 2001 cohort but be treated in 2002. We address this with an additional requirement that to be included in a cohort, control units must not be treated two years before or after (e.g. in the cohort with expiration year 2001, control units must not have been treated in any year from 1999 to 2003.).

For most of our results, we estimate the effects of funding interruptions on career outcomes separately for individuals employed in labs supported by multiple R01 grants and labs supported by a single R01 grant. The idea is that multiple-R01 labs may have more funds available to continue hiring their current employees. Tham (2021) shows that funding interruptions cause a substantial reduction in grant expenditures for single-R01 labs but a much smaller or no reduction in spending for multiple-R01 labs, indicating that this is a possible mechanism.

4.2 Identification

We rely on two main assumptions to identify the average treatment effect on the treated (ATT) of an interruption on the career outcomes of employees: (1) parallel trends and (2) no anticipation.

Parallel Trends. The parallel trends assumption requires that the average outcome among the treated and comparison populations would have followed parallel trends in the absence of treatment. In our context, this means that outcomes such as employment and wages

of the month, thus the arrival of new grant funding can be thought of as occurring on a monthly basis.

for interrupted and continuously-funded employees would have evolved in parallel if the funding interruption had not occurred.

The parallel trends assumption allows treatment to be non-random based on characteristics that affect the *level* of the outcome but requires that the treatment be mean independent of characteristics that affect the *trend* of the outcome. For instance, highly organized PIs may select into the continuously-funded control group because they are more likely to submit their paperwork on time and avoid a funding interruption. Their high level of organization may also affect employee outcomes (e.g. by ensuring that postdocs and graduate students are regularly publishing in a timely manner). However, as long as PI organizational skills affect employee outcomes in the same way both before and after the treatment, it does not violate the parallel trends assumption.

Thus, we are mainly worried about time-varying unobserved confounders. For instance, suppose that a PI selects into an interrupted lab because they are more likely to be offered another job as their grant's expiration date nears, so the probability of a new job offer varies over time. This offer may interrupt funding as the PI sets up their new lab at their new university and may also affect the trends of the post-expiration potential outcomes of employees working in their old lab, *even if the interruption had not taken place*.

Though the parallel trends assumption cannot be tested, we provide a variety of evidence suggesting that it is plausible in our setting. First, we produce raw means and event studies, neither of which show evidence of diverging trends prior to expiration-year. Second, we test the sensitivity of our estimates to violations of the parallel trend assumption using the methods of Rambachan and Roth (2019). The idea of this sensitivity analysis is to formalize violations of the parallel trend assumption that are most likely to occur in our context. These can then be used to partially identify the estimates we would obtain if we allowed for those violations.

No or Limited Anticipation. The no anticipation assumption requires that there is no effect of the treatment prior to the treatment actually taking place, for instance if units are aware the treatment will occur and change their behavior in advance. The limited anticipation assumption allows for anticipation if we are willing to assume that anticipation occurs at a fixed length of time before the treatment.[19] In our setting, this means that funding interruptions cannot impact employment and wage outcomes before the grant expires. For example, if the PIs or employees have advance knowledge that their R01 is likely to be interrupted because they notice an experiment is not working out and that would hurt their renewal application, they might choose to find a new job before waiting for the results of the renewal application to come in. We do not see any differences between treated and control outcomes before treatment, suggesting that anticipation of treatment is not violated.

5 Results

5.1 Employment Outcomes

Figure 3 displays the fraction of employees that, in a given year, fall into one of three mutually exclusive employment categories: (1) they receive positive earnings from an IPEDS university, (2) they do not receive earnings from any of our administrative data sources (W-2, 1099, LEHD PHF), or (3) they do not fall into (1) or (2) and receive positive earnings from some other non-university source (e.g. industry, self-employment, government).

The first row of graphs (Figure 3A) shows that the probability of working at an IPEDS university increases as grant expiration approaches, peaks around expiration, and then declines after expiration. This pattern holds regardless of whether individuals were in a single- or multiple-R01 lab and their lab was interrupted (treated) or continuously-funded (control). More specifically, around 65% of UMETRICS employees receive positive university earnings five years before expiration, increasing to about 90% around expiration, and then declining to around 60% five years after expiration. This inverted-U shape with a peak near expiration is mechanical, determined by the way we define our sample – we are very likely to observe positive university earnings around expiry (and less likely to observe positive university earnings as we move away from expiry) because, by definition, research personnel are receiving payments from a grant at their UMETRICS university at

[19] In this case, one might redefine the treatment period to be at the point when units are aware of treatment.

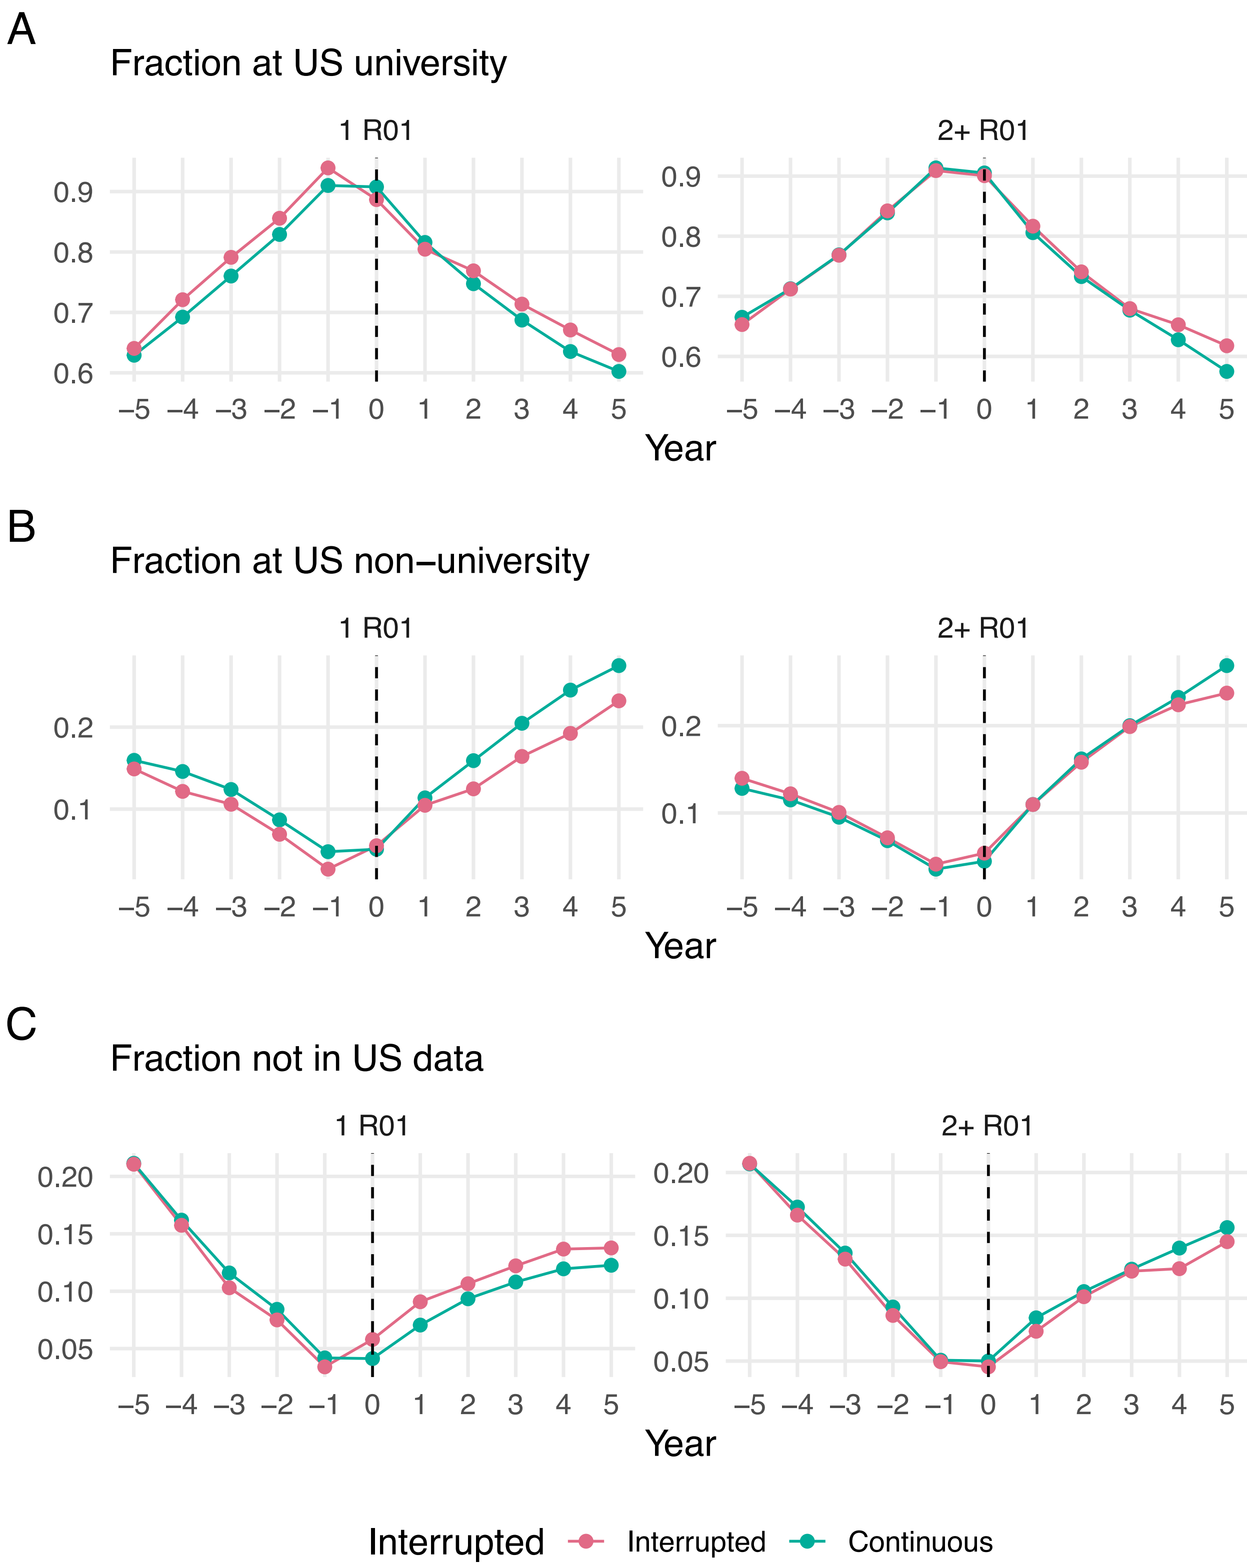

Figure 3: This figure shows the average probability that an individual is in one of three categories (paid by a US university, paid by a US non-university, or not paid in the US) from five years before interruption to five years after interruption. Figures in the left column are for individuals in a single-R01 lab and the right column is for individuals in a multiple-R01 lab.

some point during the 12 months prior to expiry.

The second row of graphs (Figure 3B) shows that the probability of receiving positive earnings from a non-university source has a U-shape with a trough around grant expiry. Specifically, about 15% of UMETRICS employees receive positive non-university earnings five years before expiration, which declines to about 5% around expiration, and then increases to about 25% after expiration. The third row of graphs (Figure 3C) shows a similar pattern for the probability of an employee not appearing in any of our administrative data sources. In this case, the probability starts at 20% five years before expiration, declines to about 5% around expiration, and then increases to near 15% five years after expiration. Thus, the probability of receiving positive earnings from a non-university and the probability of not appearing in our administrative data have U-shaped patterns – mirror images of the probability of receiving university earnings – which, again, is driven by the way we define our sample.

Though the general patterns are similar for both single-R01 and multiple-R01 labs, there are some differences. Specifically, research personnel in multiple-R01 labs experience employment trends (and levels) that are very similar before, at, and after grant expiration – crucially, these patterns are similar whether their lab is continuously-funded or interrupted, suggesting that interruptions have no employment impacts on research personnel in multiple-R01 labs.

In contrast, for all three outcomes there are sharp changes for research personnel of interrupted (treated) single-R01 labs relative to personnel of continuously-funded single-R01 labs. Indeed, personnel in a single-R01 lab are slightly more likely to receive university earnings prior to expiration. However, in the year of expiration, there is a sudden drop in this probability, followed by a recovery after two years. We also see that, prior to expiry, research personnel in an interrupted single-R01 lab are less likely to receive positive non-university earnings than are employees in a continuously funded single-R01 lab. In the year of expiration, the gap closes for about two years before opening up again. Finally, we see that single-R01 research personnel are less likely to be absent from our administrative data prior to expiry and then, in the year of expiry, become permanently more likely to be absent. Thus, the raw means suggest that funding interruptions to single-R01 labs significantly alter employment patterns for the personnel of those labs. These individuals become temporarily less likely to work at a university, temporarily more likely to work outside a university, and permanently more likely to disappear from our administrative data, reflecting detachment from the US labor market, possibly due to emigration from the US.

Moving on from the visual evidence in Figure 3, we estimate the effects of funding interruptions with the stacked difference-in-differences method described in the Estimation section. The event studies in Figure 4 show how the effects evolve over time. These results are consistent with the patterns we observe in the raw means. We see from the rightside column of graphs that, for multiple-R01 labs, interruptions have a precisely estimated zero effect on all three employment outcomes. In contrast, the left column of graphs show that the research personnel in single-R01 labs experience an immediate, though temporary, 5 percentage point drop in the probability of receiving positive university earnings – that is, they temporarily leave academia. Where do these detached personnel go? We see that there is an immediate, though temporary, 2.5 percentage point increase in the probability of receiving earnings from a non-university source. More strikingly, there is a permanent 2.5 percentage point increase in the probability of being absent from our administrative data. Thus, it appears that, immediately after an interruption, about half the research personnel who leave academia go to work in industry and half disappear from our data. The half that go into industry filter back into academia over time, but the half that exit the data appear to permanently detach from the US labor market.

To give a sense of magnitudes, recall that, in the year prior to expiry, about 90% of research personnel receive positive university earnings. Thus, a five percentage point decline is about 5.6% of the mean. In contrast, only about 5% of personnel receive positive non-university earnings and 5% were not found in any administrative data. Thus, 2.5 percentage point increases in these outcomes are about 50% of the pre-interruption means.

Table 5 presents the aggregated post-treatment effects up to 5 years after treatment for our three mutually exclusive employment outcomes. We again observe that personnel in

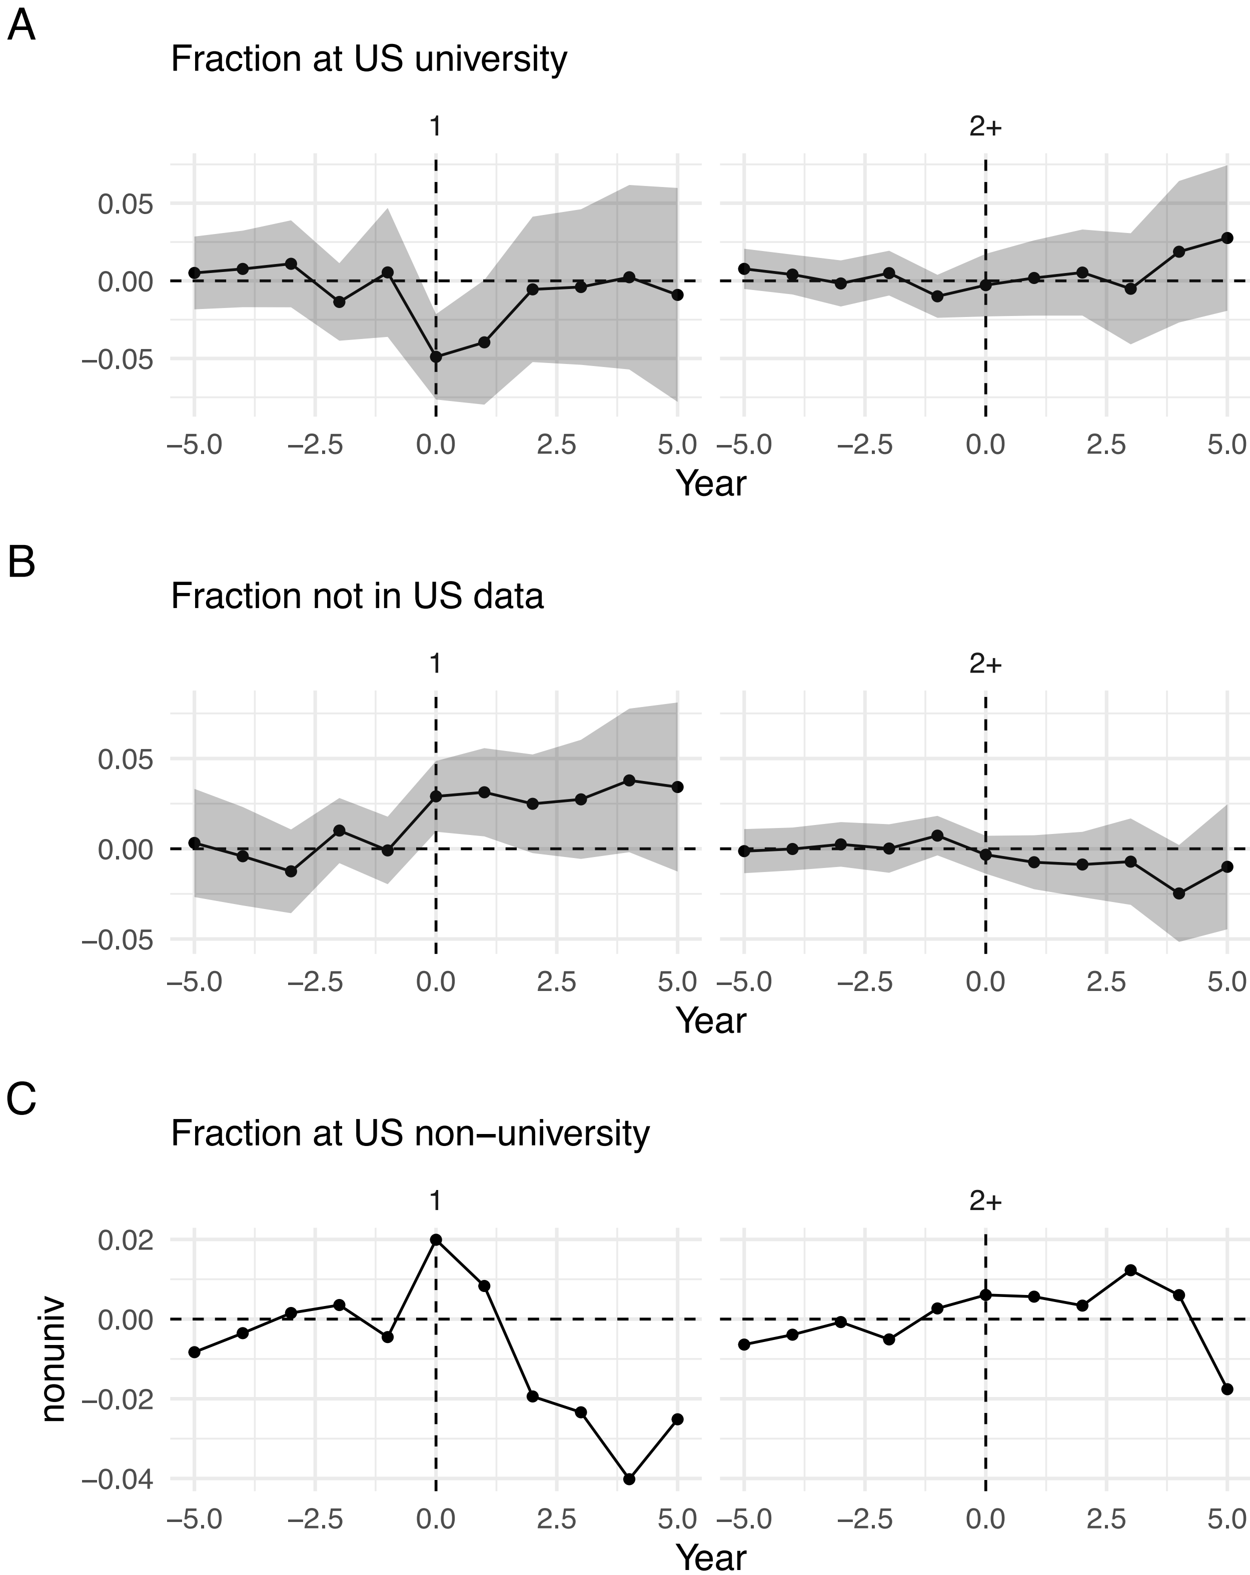

Figure 4: This figure shows event studies of the effects of interruptions on employees estimated with our modified Callaway-Sant'Anna (2020) estimator. The outcomes are the probability that an individual in one of three categories: paid by a US university (A), not paid in the US (B), or paid by a US non-university entity (C). The expiration of a lab's grant takes place at year 0. The left column is for individuals in a single-R01 lab and the right column is for individuals in a multiple-R01 lab. Standard errors are bootstrapped and clustered at the expiring-R01-level. Standard errors not availble for Figure C for disclosure reasons.

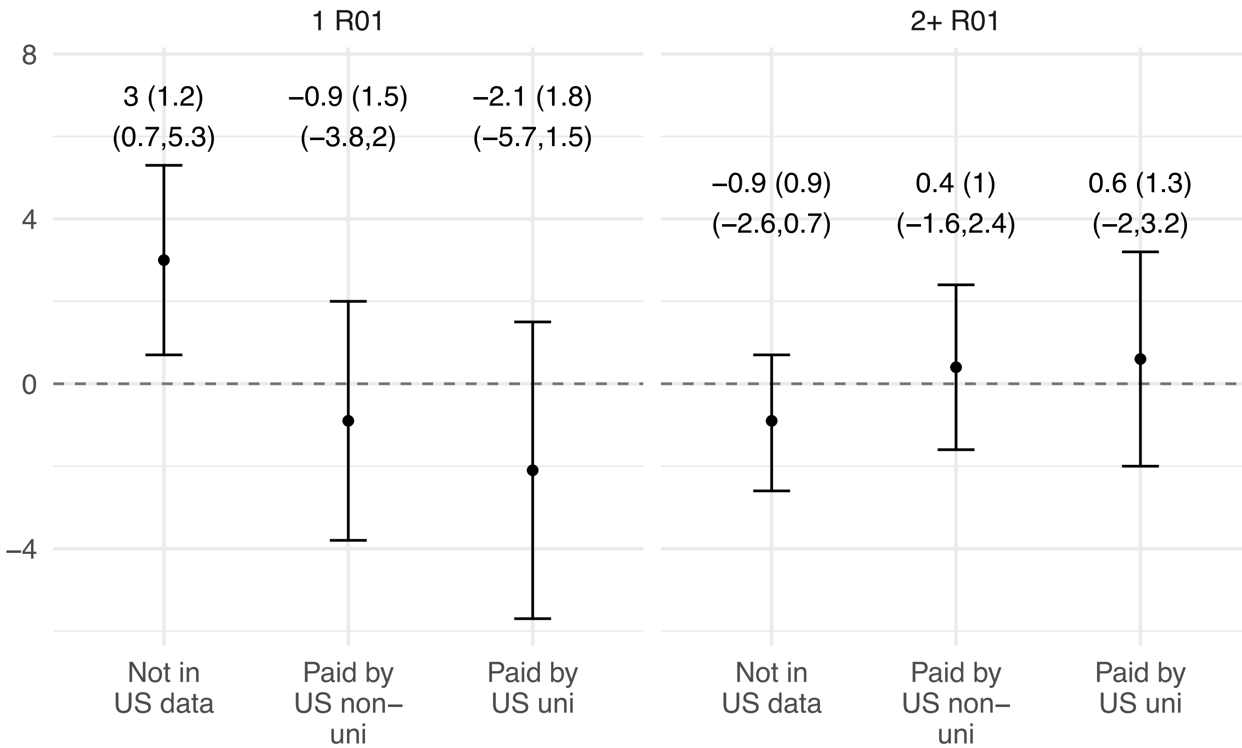

Figure 5: This figure shows aggregated estimates of the average treatment effects (ATTs) of an interruption on three mutually exclusive employment outcomes: 1) paid by a US university, 2) paid by a US non-university, or 3) not paid in the US. The effects are aggregated over the 5 years after the expiration of a lab's grant. The left graph is for individuals in a single-R01 lab and the right graph is for individuals in a multiple-R01 lab. Standard errors are bootstrapped and clustered at the R01-renewal level.

a single-R01 lab that experiences a funding interruption are about 3.0 percentage points more likely to be absent from any administrative data source, confirming the seemingly permanent exit of these individuals from the US labor market. In contrast, there are precisely estimated null effects for research personnel in labs supported by multiple R01s. Note that the recovery visible in the event studies of Figure 4 leads to a statistically insignificant drop in the probability of an individual working at a university for the aggregated treatment effects, despite the sharp and statistically significant initial (but temporary) drop in the event study.

The striking increase, after a funding interruption to a single-R01 lab, in the probability of research personnel being absent from our data raises the question of where these individuals go. Given the comprehensiveness of our administrative data, which includes the universe of W-2 and 1099 tax records along with the universe of UI wage records, it is clear that these individuals become permanently detached from the US labor market. One possibility is that they simply no longer work. Another possibility is that they leave the US to find employment in another country and thus exit our US-based administrative data. To shed light on this possibility, we break out our results by whether the research personnel are US-born or foreign-born. The idea is that, due to both preferences and constraints, the US-born are likely more attached to the US labor market than the foreign-born and are thus less likely to leave the US and to go missing from our administrative data after a funding interruption to their lab.

Figure 6 displays the aggregate five-year post-treatment effects for single-R01 labs, by whether the individual is US- or foreign-born. We see that the impacts of an interruption are driven by the US-born. Rows 1-2 show that, while US-born single-R01 research personnel are 3.5 percentage points more likely to be absent from our administrative data, the impact on foreign-born personnel is much smaller (statistically insignificant 1.8 percentage point increase). Rows 3-4 show that the foreign-born personnel of a single-R01 lab are 2.9 percentage points less likely to receive positive non-university earnings after an interruption. Rows 5-6 show that US-born personnel in a single-R01 lab are 3.5 percentage points less likely to work at a university after an interruption whereas the foreign-born are

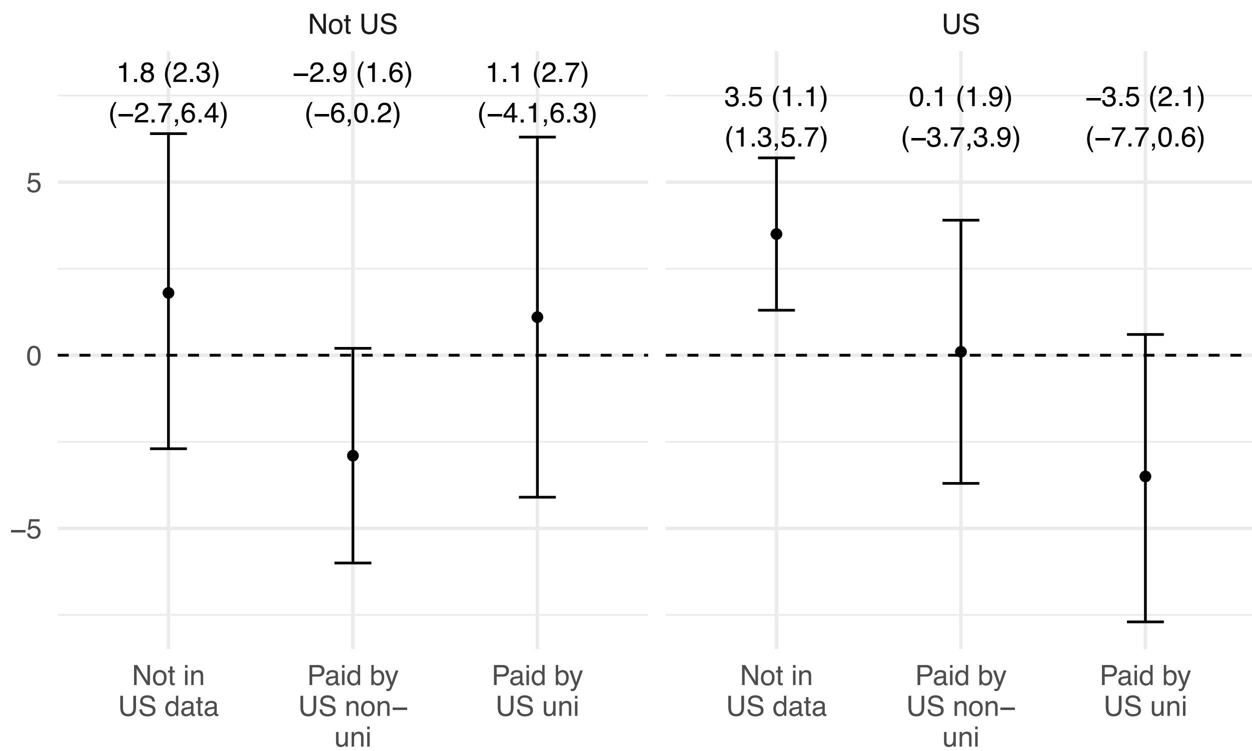

Figure 6: This figure shows, separately for US-born and foreign-born research personnel, aggregated estimates of the average treatment effects (ATTs) of an interruption on three mutually exclusive employment outcomes: 1) paid by a US university, 2) paid by a US non-university, or 3) not paid in the US. The effects are aggregated over the 5 years after the expiration of a lab's grant. All estimates are for individuals in a single-R01 lab. The estimates are obtained using a modified Calloway-Sant'Anna (2020) estimator (see Estimation section). Standard errors are bootstrapped and clustered at the interrupted R01 level.

only 1.1 percentage points (statistically insignificant) less likely.

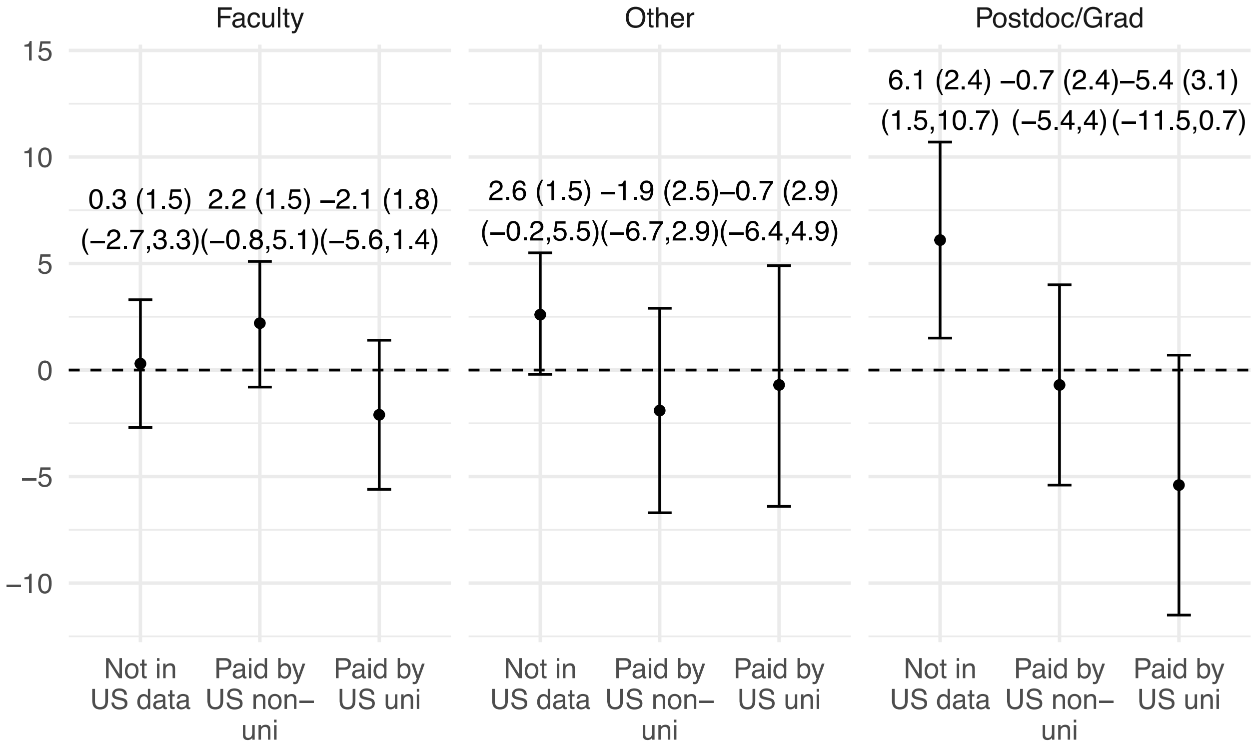

Figure 7: This figure shows, separately by occupation category of research personnel, aggregated estimates of the average treatment effects (ATTs) of an interruption on three mutually exclusive employment outcomes: 1) paid by a US university, 2) paid by a US non-university, or 3) not paid in the US. The effects are aggregated over the 5 years after the expiration of a lab's grant. All estimates are for individuals in a single-R01 lab. The estimates are obtained using a modified Calloway-Sant'Anna (2020) estimator (see Estimation section). Standard errors are bootstrapped and clustered at the R01-renewal level.

Next, we examine whether particular occupations drive our results. In the UMETRICS data, an employee is classified into one of six occupation categories: faculty, post-doc, grad student, undergrad, other student, and staff. Figure 7 presents aggregate post-treatment effects, for single-R01 labs, broken out by three broad occupation categories: faculty, postdoc/grad student, and other.

The effect of an interruption on the research personnel of single-R01 labs is mainly driven by post-docs and grad students. They are 6.1 percentage points more likely to be absent from our administrative data and are 5.4 percentage points less likely to receive positive university earnings. "Other" research personnel are 2.6 percentage points more likely to be absent from US administrative data but their probability of being paid by a US university is unaffected. Finally, the faculty of single-R01 labs are unaffected by a funding interruption, with precisely estimated null effects for all three main employment outcomes.

5.2 Wage Outcomes

In addition to employment outcomes, we are interested in whether funding interruptions impact the subsequent earnings of research personnel. Ex ante, the effects are ambiguous. If the interruption of a lab's funding knocks off course an individual's career progression through academia, then earnings may decrease relative to the counterfactual of belonging to a continuously-funded lab. For instance, if a post-doc's lab is interrupted, and they must scramble to find another job, their progression to a tenure-track faculty position may be delayed or derailed entirely. On the other hand, if the interruption spurs employees to get jobs in the higher-paying private sector, then the interruption could actually cause wages to rise.

Figure 8A displays event studies for research personnel in both multiple- and single-R01 labs. As with employment outcomes, we see that interruptions to multiple-R01 labs have no effect on the arcsine total wages of the personnel in those labs. Once again, the results are starkly different for research personnel of a single-R01 lab – interrupted personnel experience a very sharp decline in earnings that reaches about 50% after 1-2 years and about 75% after 5 years.

We stress that these estimates almost surely overstate the impacts of interruptions on wages. The sample includes individuals with zero earnings in any of our data sources but most of these individuals probably earn positive earnings outside of the US, which we do not observe.

We try to get a better estimate of the impact on wages as follows. We identify subsets of research personnel defined by their attachment to the U.S. labor market before and after the expiry of their grants. The first subset is composed of personnel that are "partially attached" to the US labor market. Specifically, we require them to receive positive earnings

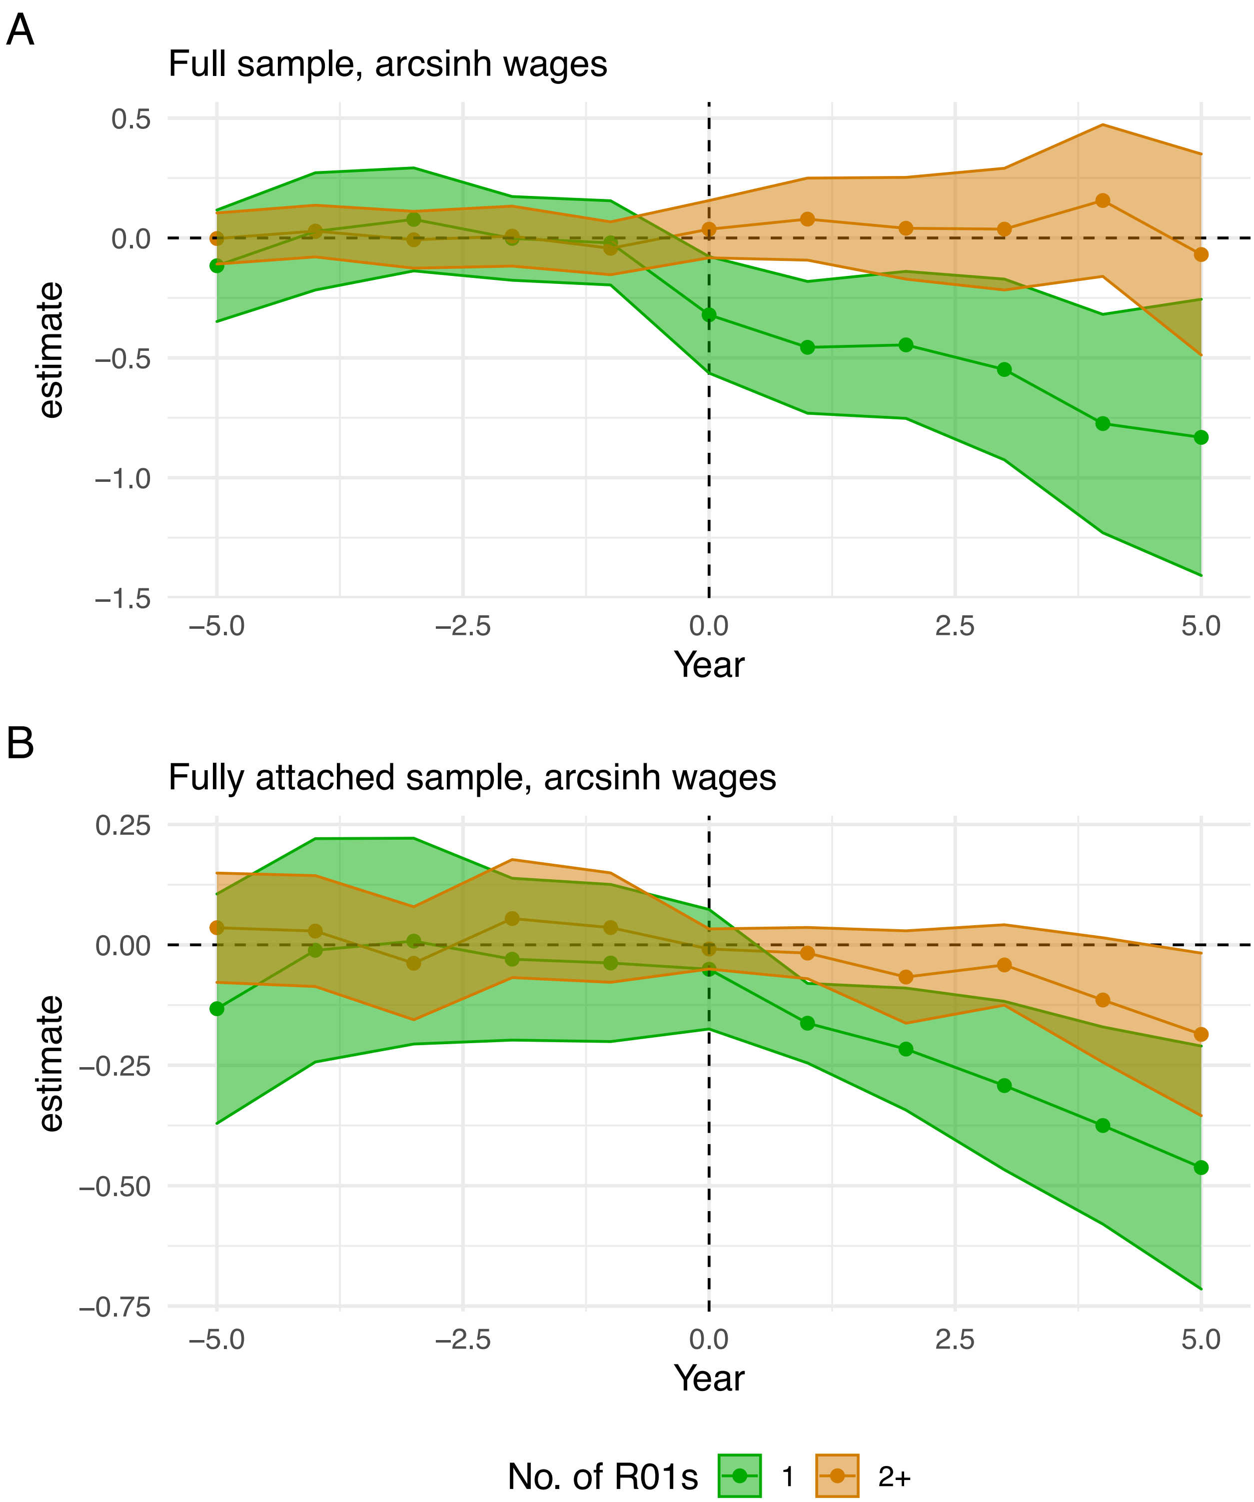

Figure 8: This figure shows time-varying effects (event studies) of an interruption on the arcsine of total earnings. The expiration of a lab's grant takes place at time 0, and the estimated interruption effects range from 5 years before to 5 years after an interruption. The green series is for individuals in a single-R01 lab and the orange series is for individuals in a multiple-R01 lab. The top graph (Panel A) is for the full sample of research personnel and the bottom graph (Panel B) is for the sub-sample of research personnel that have positive earnings in all periods from one period before expiry to five years after expiry (i.e. the fully attached sub-sample). The estimates are obtained by using our modified Calloway-Sant'Anna (2020) estimator, which accounts for well defined event years in the continuously-funded control group. Standard errors are bootstrapped and clustered at the

in the year before expiration (-1) and to receive positive earnings in *at least one* of the six years after expiration (0-5). The second subset is composed of personnel that are "fully attached" to the US labor market. Specifically, we require them to receive positive earnings in *all* years between -1 and 5 around expiry.

This restriction requires additional assumptions to be interpreted. The core question here is whether comparing employees who remained in the US after a funding interruption is equivalent to comparing the full sample of employees. This may not be the case if there is differential selection into which employees remain in the US after an interruption, thus altering the sample composition of who remains in the US workforce. For example, if employees without US citizenship or permanent residency are willing to accept a lower-paying job in order not to leave the US, then our estimates would overstate the effect of interruption on wages (after adjusting for cross-country wage differentials).

Figure 8B displays event studies for the fully attached subsample (for disclosure avoidance reasons at Census, we cannot display the event studies for the partially attached sample). As with the full sample, the fully attached research personnel in multiple-R01 labs are mostly unaffected by an interruption. The point estimates are precisely estimated zeros in periods 0 through 3, but they drift down slightly in periods 4 and 5. The fully attached research personnel in single-R01 labs again experience a sharp decline in earnings after an interruption, but it is much more modest than the decline experienced by the full sample. This suggests that our wage results are not completely driven by the exit of research personnel from the US labor market (extensive margin) and that, even conditional on employment in the US, research personnel from single-R01 labs experience substantial and permanent wage declines after a funding interruption.

Figure 9 confirms these results, displaying the aggregate post-treatment effects, for both multiple-R01 and single-R01 labs, for the full sample, the partially attached subsample, and the fully attached subsample. The wages of research personnel in multiple-R01 labs do not change in response to an interruption for the full sample and the partially attached subsample. The wages of personnel in the fully-attached subsample decline, relatively modestly, by about 5.9%, which is driven by the decline in the later periods of

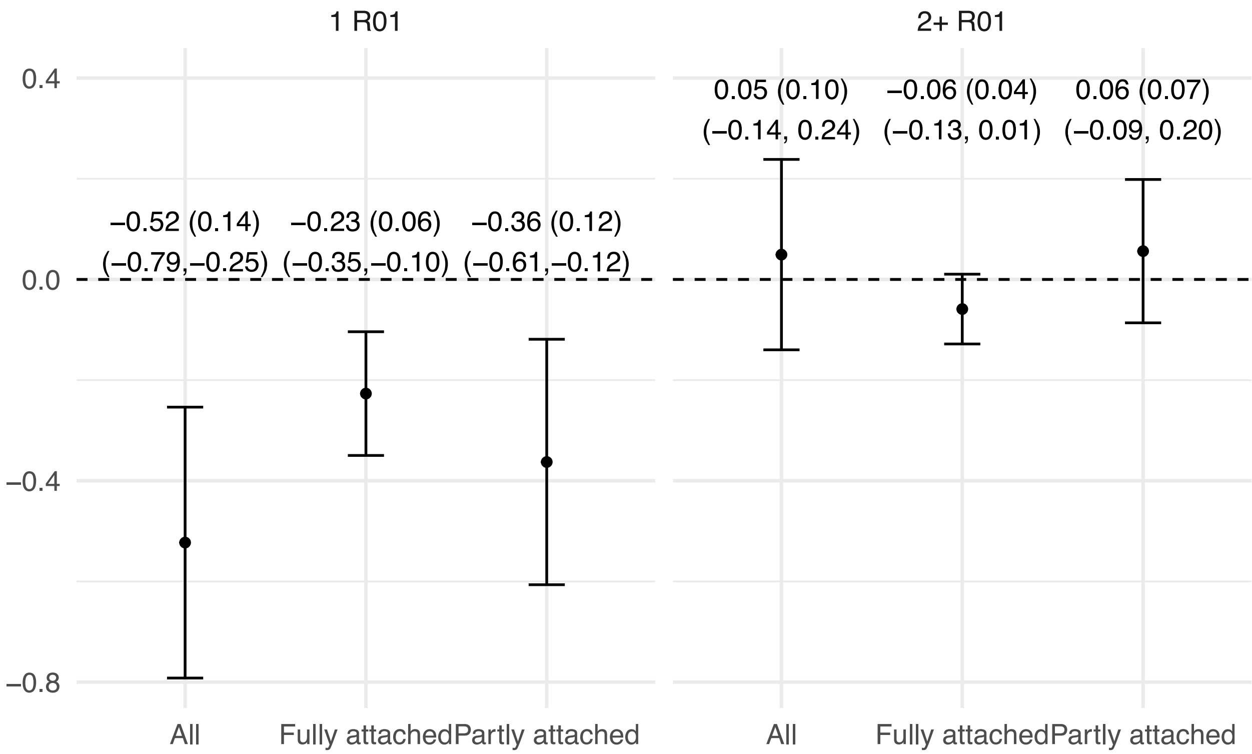

Figure 9: This figure shows aggregated estimates of the average treatment effects (ATTs) of an interruption on the arcsine of total wages. The effects are aggregated over the 5 years after the expiration of a lab's grant. The first three rows are for individuals in a single-R01 lab and the bottom three rows are for individuals in a multiple-R01 lab. The sample labeled 'All' is the full sample of all research personnel, the sample labeled 'Partly attached' is the sub-sample that have positive earnings one year before expiration and in at least one of the five years after expiration, and the sample labeled 'Fully attached' is the sub-sample that have positive earnings in all periods from one period before expiry to five years after expiry. The estimates are obtained using a modified Calloway-Sant'Anna (2020) estimator. Standard errors are bootstrapped and clustered at the interrupted R01 level.

the event study. Consistent with previous results, we see large and statistically significant interruption effects, across all samples, for research personnel in single-R01 labs. The full sample experiences a wage decline of about 52%, the partially-attached sample a decline of 36%, and the fully attached sample a decline of about 23%. Thus, again, even conditional on positive earnings in all post-interruption time periods (i.e. the fully-attached sample), we observe substantial wage declines for research personnel in single-R01 labs that experience an interruption to funding.

5.2.1 Wage outcomes by occupation

We repeat the estimation on the fully attached subsample, further divided into subsamples by our three occupation categories. Figure 10 shows that the wage estimates vary substantially by occupation. Wages decline by 5% for faculty but the estimate is statistically insignificant. Taken together with small employment effects, this suggests that interruptions do not have much of an effect on faculty labor market outcomes. For trainees, the wage decline is substantially larger at 20%, though statistically insignificant. For other occupations, the wage decline is 30% and statistically significant, though the confidence intervals are also wide. The larger estimate for other occupations may be in part driven by undergraduates taking on lower paying jobs if they are released from a research job.

6 Robustness

6.1 Alternative control group

One concern about our identification strategy is whether there are unobserved differences between interrupted and uninterrupted labs that both cause interruptions and affect labor market outcomes of lab employees irrespective of interruptions. As a robustness test, we use as a control group employees in labs that also had interrupted R01s but had multiple R01s. Using interrupted labs instead of uninterrupted labs allows us to adjust for factors that lead to labs having interrupted R01s, assuming those are the same for labs with one R01 and multiple R01s. It also requires employees in multiple R01 labs to be less affected

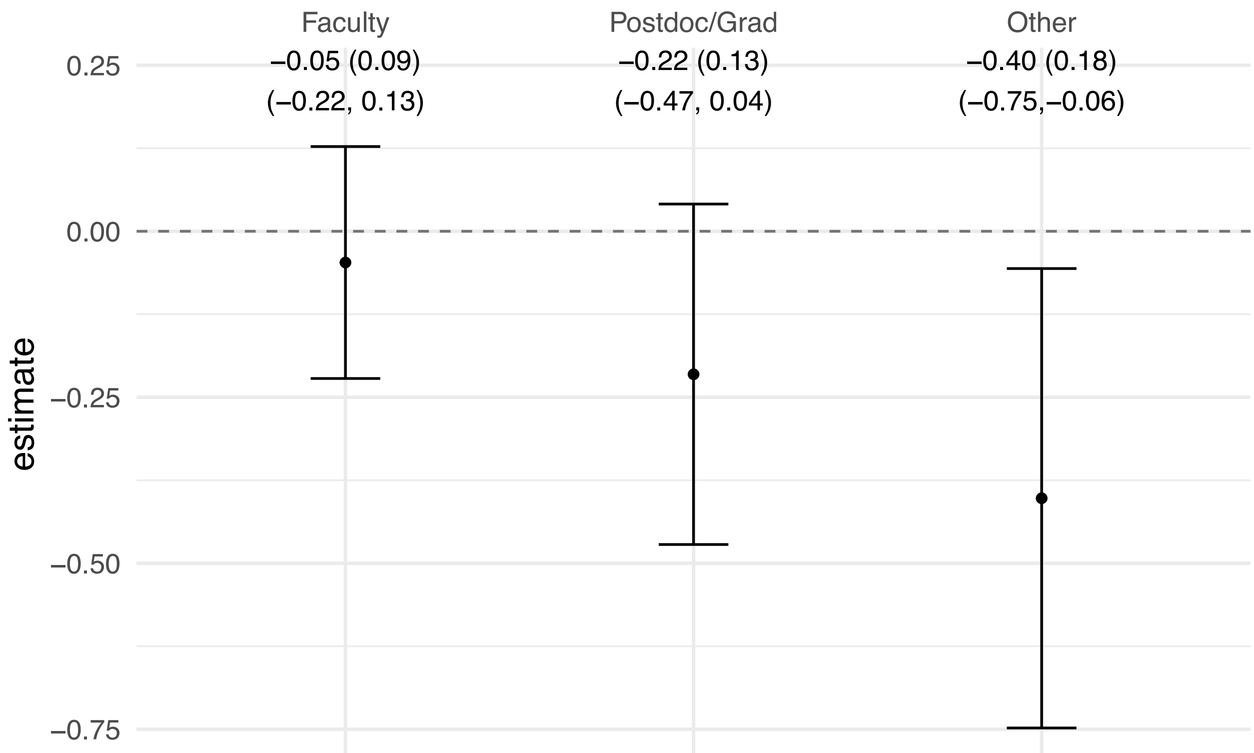

Figure 10: This figure shows aggregated estimates of the average treatment effects (ATTs) of an interruption on the arcsine of total wages, by estimates from splitting the data by occupation. The effects are aggregated over the 5 years after the expiration of a lab's grant. The overall sample is individuals in a single-R01 lab and 'Fully attached' to the US labor force i.e. the subsample that has positive earnings in all periods from one period before expiry to five years after expiry.

(or unaffected) by funding interruptions. Figure 11 in the Appendix shows the event studies for this robustness test. The treatment dynamics for this alternative control group are similar to those in our main results, which suggests that our results are not driven primarily by differences between labs interrupted and uninterrupted R01s.

7 Conclusion

In this paper, we study how grant funding delays or "interruptions" affect the careers of research personnel. Using a combination of public-use grant data, university administrative data, and tax and employment data, we find that when the renewal of a Principal Investigator's major grant – the NIH's R01 – is interrupted, the PI's employees are less likely to work in the US in the long-term, while those remaining in the US are less likely to work in a university initially. Those who remain in the US also earn substantially less.

One lesson from our results is the costs of instability in a system even when that instability is expected. The yearly delays in passing a federal budget and the possibility of funding interruptions are issues that scientists and institutions are keenly aware of and presumably try to prepare for. Yet even within our sample of research-intensive universities, the impact of funding interruptions on employees is non-trivial. While the returns to science are yielded over years or decades (or longer), the costs to scientific workers occur on a more immediate timescale. The benefits of policies and new funding models that reduce the occurrence of similar disruptions may be understated as a result.

References

Azoulay, Pierre, Joshua S Graff-Zivin, and Gustavo Manso. 2011. "Incentives and Creativity: Evidence from the Academic Life Sciences." *The RAND Journal of Economics* 42 (3): 527–54.

Azoulay, Pierre, Wesley H Greenblatt, and Misty L Heggeness. 2021. "Long-Term Effects from Early Exposure to Research: Evidence from the NIH 'Yellow Berets'." *Research Policy* 50 (9): 104332.

Baker, Andrew C, David F Larcker, and Charles CY Wang. 2022. "How Much Should We Trust Staggered Difference-in-Differences Estimates?" *Journal of Financial Economics* 144 (2): 370–95.

Baker, Andrew, David F Larcker, and Charles CY Wang. 2021. "How Much Should We Trust Staggered Difference-in-Differences Estimates?" *Available at SSRN 3794018*.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Did: Difference in Differences." https://bcallaway11.github.io/did/.

Callaway, Brantly, and Pedro HC Sant'Anna. 2020. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics*.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *The Quarterly Journal of Economics* 134 (3): 1405–54.

DrugMonkey. 2009. "Never, Ever, Ever, Nuh-Uh, No Way, Ever Trust a Dec 1 Start Date!" https://web.archive.org/web/20190531000609/http://drugmonkey.scientopia.org/2009/12/15/never-ever-ever-nuh-uh-no-way-ever-trust-a-dec-1-start-date/.

———. n.d. "Never Ever Trust a Dec 1 NIH Grant Start Date: The Sickening." https://web.archive.org/web/20210809210741/http://drugmonkey.scientopia.org/2016/03/09/never-ever-trust-a-dec-1-nih-grant-start-date-the-sickening/.

Fikes, Bradley J. 2018. "Uncertain Federal Budget Curtails Research Hiring in San Diego." *San Diego Union Tribune*, January. https://web.archive.org/web/20180331091854/https://www.sandiegouniontribune.com/business/biotech/sd-me-shutdown-research-20180122-story.html.

Foundation, National Science. 2018. "Science & Engineering Indicators 2018."

Freeman, Richard, and John Van Reenen. 2009. "What If Congress Doubled r&d Spending on the Physical Sciences?" *Innovation Policy and the Economy* 9 (1): 1–38.

Guzey, A. 2019. "How Life Sciences Actually Work: Findings of a Year-Long Investigation." https://guzey.com/how-life-sciences-actually-work/.

Hill, Ryan. 2019. "Searching for Superstars: Research Risk and Talent Discovery in Astronomy." Working Paper. Cambridge, MA: Massachusetts Institute of Technology ….

IRIS. 2019. "IRIS Year in Review, Fall 2019." Institute for Research on Innovation; Science.

Medicine, The Ohio State University College of. 2020. "Appointments, Promotion, and Tenure Criteria and Procedures for the Ohio State University College of Medicine." https://oaa.osu.edu/sites/default/files/uploads/governance-documents/college-of-medicine/Medicine-APT_2020-09-25.pdf.

Mervis, Jeffrey, and Eliot Marshall. 1996. "Science Budget: When Federal Science Stopped." *Science* 271 (5246): 136–36.

Rambachan, Ashesh, and Jonathan Roth. 2019. "An Honest Approach to Parallel Trends." *Unpublished Manuscript, Harvard University.[99]*.

Stephan, Paula. 2013. "The Endless Frontier: Reaping What Bush Sowed?" National Bureau of Economic Research.

Tham, Wei Yang. 2021. "Science, Interrupted: Funding Delays Reduce Research Activity but Having More Grants Helps."

Wagner, Deborah, Mary Layne, et al. 2014. "The Person Identification Validation System (PVS): Applying the Center for Administrative Records Research and Applications'(CARRA) Record Linkage Software." Center for Economic Studies, US Census Bureau.

"What We Learned Doing Fast Grants." n.d. Collison, Patrick and Cowen, Tyler and Hsu, Patrick. https://future.a16z.com/what-we-learned-doing-fast-grants/.

Wuchty, Stefan, Benjamin F Jones, and Brian Uzzi. 2007. "The Increasing Dominance of Teams in Production of Knowledge." *Science* 316 (5827): 1036–39.

Data Appendix

Exporter

ExPORTER is publicly available data provided by the NIH.[20] Multiple categories of data are available on ExPORTER: Projects, project abstracts, publications, link tables from projects to publications, patents, and clinical studies. We use the Projects data to calculate how long it takes for projects to be renewed, which in turn we use to define they were interrupted or not.

Defining project periods

NIH projects are assigned a **core project number** that is used over multiple **project periods**. A project period is what we would conventionally call a "grant": some amount of funding guaranteed over a few years. At the end of each project period, the Principal Investigator (PI) of the project can apply to renew funding for that project. If the renewal application is successful, that begins a new project period. This interval between when a project period ends and when a new project period begins (after successful renewal) is the focus of this project. However, ExPorter does not provide explicit identifiers for project periods so we have to use the other information provided in ExPORTER to determine when a project period started or ended.

The funds for a project period are allocated from the NIH to the project over multiple **budget periods**.[21] Each budget period is recorded as a row in the ExPorter *Projects* data.

For example, project number *R01GM049850*, led by PI Jeffrey A. Simon, was funded from FY 1996 to FY 2017, except for FY 2013. Table 2 below shows the records from its first eight years of funding. Each year the project was funded appears as a new row in the data. In the first year, the project was funded as a new project (application type 1), but then for each of the next three years was funded as a "continuation" (application type 5). The project is then funded as a "renewal" (application type 2) in FY 2000, then again as a "continuation"

[20] https://exporter.nih.gov/
[21] This is laid out in more detail in Section 5.3 of the NIH Grants Policy Statement.

Table 2: Example of NIH ExPorter data before aggregation into project periods

PI Name	Core Project Num	Fiscal Year	Application Type	Comment
Simon, Jeffrey A	R01GM049850	1996	1	New
Simon, Jeffrey A	R01GM049850	1997	5	Continuation
Simon, Jeffrey A	R01GM049850	1998	5	Continuation
Simon, Jeffrey A	R01GM049850	1999	5	Continuation
Simon, Jeffrey A	R01GM049850	2000	2	Renewal
Simon, Jeffrey A	R01GM049850	2001	5	Continuation
Simon, Jeffrey A	R01GM049850	2002	5	Continuation
Simon, Jeffrey A	R01GM049850	2003	5	Continuation

the next three years. Thus, we can infer that FY 1996 to FY 1999 constituted one project period. After that, the project had to be renewed, resulting in a new project periods from FY 2000 to 2003.

The exact steps we use to determine project periods are:

1. Indicate first budget period of a new project period if application type is 1, 2, or 9. Define the start date of the project period as the start date of the budget period
2. Arrange budget periods by budget start date. Assign budgets that start after the first budget of a project period (as indicated by application type) to that project period, until the first budget of a new project period is reached
3. Assign the project period end date as the latest budget end date of all budget periods assigned to the project period

Calculating time to renewal

Our treatment variable is the time between consecutive project periods for a given R01. For each pair of consecutive project periods, we calculate this as the number of calendar days between the expiration date of the earlier project period and the renewal date of the later project period. In cases where the expiration date is after the renewal date, we redefine the expiration date to be one calendar day before the renewal date. We also use this adjusted expiration date as the reference date for defining the periods of time we use to link employees to PIs or to link PIs to their other grants (described below).

Linking PI IDs to project periods

PI IDs in ExPORTER are assigned at the row/budget period level. We assign a PI ID to a project period if a PI was assigned to any of the budget periods that constitute the project period.

Sample Construction

An ideal dataset would allow us identify employees who were part of a PI's lab/research group that went through an R01 renewal. However, UMETRICS allows us to infer those relationships based on which employees a PI was paying around the time of renewal. The overall steps to construct our sample involve decisions at each of the following levels of data:

1. R01
2. PI-R01
3. PI-R01-employee

First, we find all pairs of expiring-renewed R01 project period pairs that were also successfully linked to UMETRICS. We keep all expiring-renewed pairs that were renewed within the same fiscal year to ensure that any observed delays in renewal were not due to unusual circumstances or data errors. This is also consistent with the NIH-level counterfactual we have in mind where the NIH funds the same projects within the same fiscal year without delay.

Next, we link the project periods in each expiring-renewed project period pair to their PI IDs. We retain all units where the PI ID appeared in both the expiring and renewed project periods. This leaves us with a set of (PI, expiring R01 project period, renewed R01 project period) triples. For simplicity, we refer to these as PI-expiring-R01 units.

Our next step is to link PI-expiring-R01 units to employees. For each PI-expiring-R01, we first fix a 12-month window that ends in the month the R01 was expiring. I.e. if the expiring month is Dec 2021, the window is from Jan 2021 to Dec 2021. We then link each PI to all their NIH grants at in that time window based on the overlap between the 12-month window and the start and end dates of any project periods associated with the PI. This gives us a PI's portfolio of NIH grants in the 12-month period prior to expiry.

The next step is to find employees who were part of a PI's lab by finding employees who were paid any of the grants in this portfolio during the 12-month window. We first link the PI's grant portfolio to a crosswalk between NIH core project numbers and UMETRICS award numbers, an identifier in UMETRICS that accompanies each transaction (**XXX** describe crosswalk more?). Through the award numbers, we then link to the UMETRICS employee dataset to obtain all employee numbers paid through the awards in the 12-month window.

Counting R01-equivalents

To take into account that PIs with more grants may have a buffer, we measure the size of a PI's grant portfolio based on the number of R01s they had around the time of R01 expiry. The process of constructing this measure is the same as the one for linking PIs to employees described in the previous section, except that we find all grants within a 24-month window that begins 11 months before and ends 12 months after the expiry date of the focal R01.

We include grants after expiry to allow for the possibility that PIs anticipating receiving more grants may be able or more willing to find ways to continue funding affected employees. This also assumes that the number of R01-equivalents is not affected by interruptions (i.e. not a post-treatment variable), which we think is reasonable in this context given the time lag between applying for and receving an R01.[22]

Defining R01-equivalents

Given the outsized importance of the R01, we use the number of R01s and R01-equivalents as our measure of a PI's grant portfolio. R01-equivalents are defined at the time of writing (2021) as "activity codes DP1, DP2, DP5, R01, R37, R56, RF1, RL1, U01 and R35 from

[22] E.g. A guide by NIAID suggests it can take 8 to 20 months upon applying https://www.niaid.nih.gov/grants-contracts/timelines-illustrated

select NIGMS and NHGRI program announcements."[23] However, the definition of R01-equivalent definitions can change slightly over time. We use the Internet Wayback Machine to find R01-equivalent definitions going as far back as possible (late 2017) and include all activity codes ever defined as an R01. We also include all R35 grants rather than only those from NIGMS or NHGRI, as specified in the definition, as the R35 seems to be used similarly across the NIH (to provide long-term support for outstanding investigators e.g. see https://web.archive.org/web/20211027025938/https://grants.nih.gov/grants/funding/ac_search_results.htm)

Estimating stacked data with CS estimator (Callaway and Sant'Anna 2020)

Stacked data is usually estimated by OLS with cohort-specific unit and time fixed effects.[24] We use the estimator developed by Callaway and Sant'Anna (2020) ("CS estimator") because it has some desirable features such as more transparent weighting in the aggregation of group-time treatment effects, simultaneous confidence intervals, and a doubly robust modeling option. However, the CS estimator as currently implemented via the `did` package (Callaway and Sant'Anna 2021) does not straightforwardly accommodate the case where controls have a well-defined "treatment" date. The rest of this section describes how we implement it.[25]

The main issue at hand is that the `did` package treats all control units as being control units for all treatment cohorts, whereas in our case, control units belong to specific treatment cohorts. This means that if implemented without any modifications, for any given treatment cohort, the `did` package will use observations of control units from another treatment cohort in estimation. Consider, for instance, two treatment cohorts with treatments in 2000 and 2001 and units in each cohort spanning one year before and one year after (so the units

[23] https://web.archive.org/web/20211221215217/https://grants.nih.gov/grants/glossary.htm

[24] See the appendices of Cengiz et al. (2019) for an implementation and A. Baker, Larcker, and Wang (2021) for a discussion

[25] This was originally described in a tweet: https://twitter.com/wytham88/status/1368963248335257600?s=20&t=GfGbNt6iI3xODc5DfZfciQ

in cohort 2000 span 1999 to 2001). The `did` package will compare treatment and control units belonging to the cohort 2001. However, control units from the year 2000 will also have observations in the year 2001 that `did` will use in estimation.

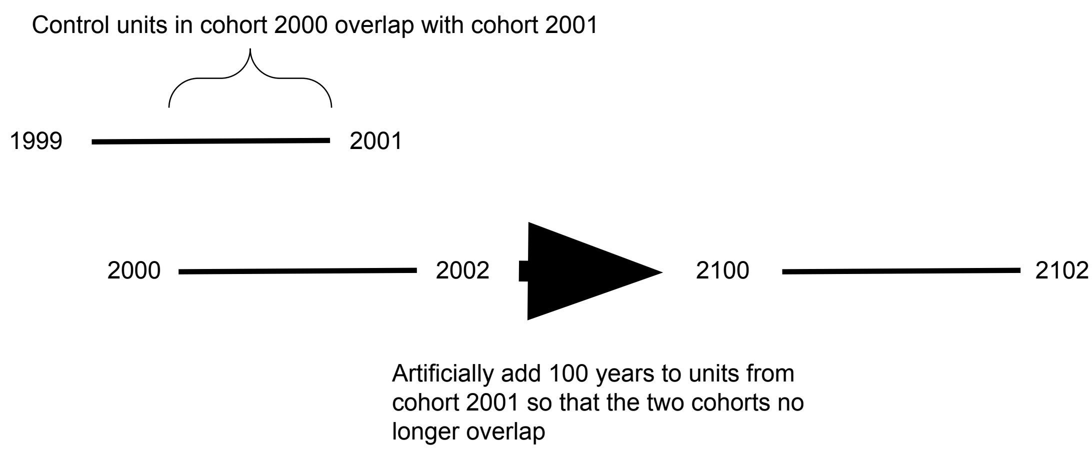

The goal is ensure that treated units within cohort are only compared to control units in the same cohort. This can be done by adding a large number to the calendar years for units in each cohort, thus artificially "separating" them from other cohorts. For instance, consider the example of two cohorts, 2000 and 2001, with the 2000 cohort spanning years 1999 to 2001 and the 2001 cohort spanning years 2000 to 2002. If we articially add 100 years to the 2001 cohort, then it spans the years 2100 to 2102, which no longer overlaps with the 2000 cohort. Thus, the `did` package will no longer use observations from the 2000 cohort in 2001 for estimating effects in the 2001 cohort.

One drawback of this workaround is that it is no longer possible (or at least not straightforward) to aggregate effects by calendar years, since they have been relabelled with fake calendar years.

A Robustness

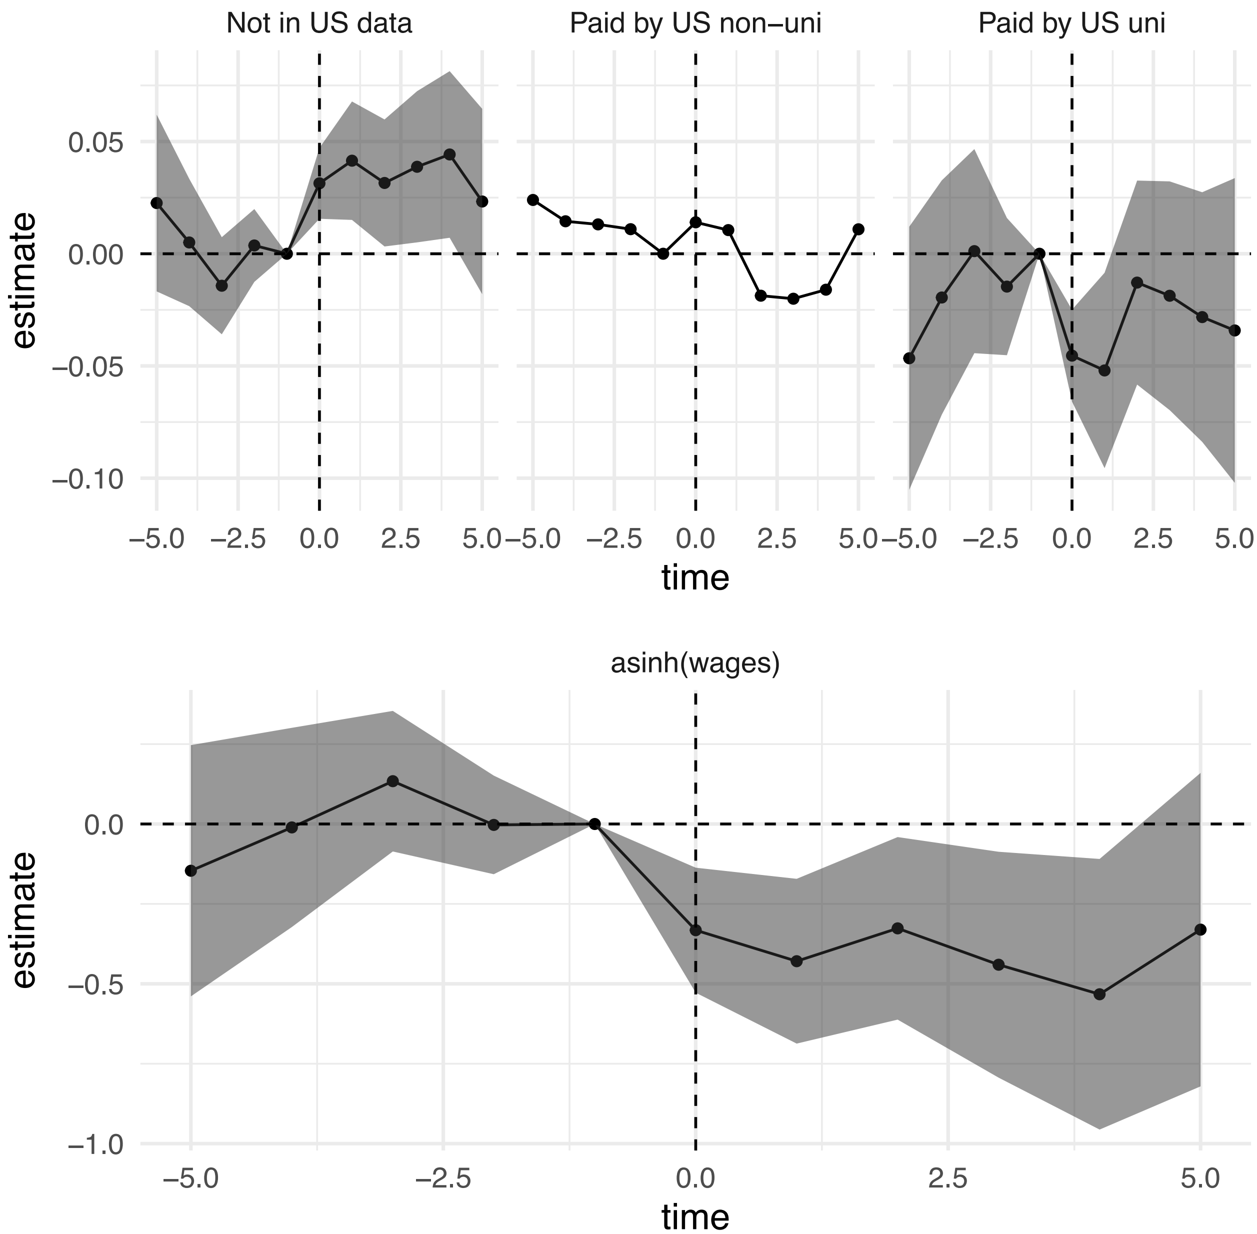

Figure 11: This figure shows the event studies using an alternative control group, employees in labs with interrupted R01s and multiple R01s. The first row shows event studies for the three mutually exclusive placement outcomes: absent from US data, paid by a US university, and paid by a US non-university. The second row shows the event study for arcsinh-transformed wages.