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¶

2022-02-23

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 present evidence that interruptions lower wages by 20%.

*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

Economists have long recognized the importance of innovation to economic growth (Romer 1990; Weitzman 1998) as well as its likely underproduction due to its public goods nature (Arrow 1962). These ideas underpin the justification for government intervention in the production of knowledge, such as government funding of R&D. Recent attempts to quantify the returns to R&D spending suggest that they are substantial (Jones and Summers 2020; Azoulay et al. 2015; Myers and Lanahan 2021). However, assessments of the institutions through which the US government funds research are less sanguine.

One major concern is the effects that a "hypercompetitive" environment – arising due to a mismatch in the number of scientists and amount of funding (Alberts et al. 2014) – might have on the stability of a career in science (Milojević, Radicchi, and Walsh 2018). In fields such as the biomedical sciences, this instability is amplified by the growth of "soft money" positions, where faculty are expected to raise part (or sometimes all) of their salary by winning research grants.

In addition to these systemic concerns, Principal Investigators (PIs) face the yearly uncertainty of how much funding Congress will provide and when it will approve the federal budget (not to mention the non-negligible threat of a government shutdown), which make it harder for them to hire personnel (Fikes 2018; DrugMonkey 2009, n.d.).

Research and policy solutions have tended to focus on whether faculty or principal investigators would do more innovative or risky science in less uncertain environments (e.g.

Azoulay, Graff-Zivin, and Manso (2011); Thompson (2021)). However, the importance of team science (Wuchty, Jones, and Uzzi 2007) means that these concerns extend beyond faculty to the careers of personnel supported by research grants, including graduate students, postdocs, research scientists, and research assistants. Even in non-emergency times, the predictability of funding may still matter greatly to the scientific workers who rely on grants.

It is important to understand how the combination of funding uncertainty and scientific workers' dependence on grant funding affects their career trajectories. Much work on the careers of scientists focus on bibliometric data such as publications, patents, and citations as a measure of scientific ability to explain one's progression - or lack thereof - on the academic ladder (Balsmeier and Pellens 2014; Conti and Visentin 2015). However, research on the general labor market has found that "luck" (initial labor market conditions at graduation) matters greatly for workers (Oreopoulos, Von Wachter, and Heisz 2012). More narrowly, there is evidence that scientific careers can be affected meaningfully by events that initially seem small (Azoulay, Greenblatt, and Heggeness 2021; Hill 2019). Focusing on bibliometric data provides a limited understanding of the career choices that scientific workers face. It leaves out personnel who have not published yet or whose primary job function does not involve publishing. It also does not capture other factors that scientific workers value in their job placement such as compensation, research flexibility, and job stability (Janger and Nowotny 2016; Roach and Sauermann 2017; Stern 2004; Cheng 2021).

We make progress on this front by using a difference-in-differences design to study how funding delays ("interruptions") to a scientist's major NIH grant – the "R01" – affects the personnel they hire (including the scientists themselves). We use university administrative data linked to W2 tax records, 1099 tax records, and unemployment insurance (UI) wage records, which enables us to observe earnings and *all* job placement within the US.

Tham (2021) shows that when PIs of an R01 experience an interruption, PI spending across all NIH grants decreases substantially until the grant is renewed. This decrease in spending is concentrated among PIs with only one R01, while spending decreases much less for PIs with multiple R01s. This indicates that interruptions constrain a PI's ability to

spend and thus possibly prevent them from continuing to employ existing personnel or from hiring new personnel. This raises the question of how an employee's outcomes are affected. Consistent with these spending patterns, we find that the impacts are substantial for employees working with single-R01 PIs while the employees of multiple-R01 PIs are unaffected on any of the outcomes we observe.

We first describe the results for employees working with a single-R01 (i.e. more funding constrained) lab. 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.

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.[1] Since we do not observe immigration status, we instead estimate results separately for US-born and foreign-born employees.[2] 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, re-

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

[2]We use the term "foreign-born" instead of "non-US-born" for easier reading.

search 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 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 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 effect. 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.[3]

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.

2 Background[4]

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.[5] 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.

[3] 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.

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

[5] 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.

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 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.[6] 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.[7] 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.

As discussed, the focus of our paper is on the effects of lapses in funding between two project periods – that is "funding interruptions." Thus, we exclusively analyze projects that are successfully renewed at least once and so span multiple project periods.

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

[7] 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.

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 the rest of the paper, we

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.

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).[8]

Crucially, the NIH grants in UMETRICS can be linked to NIH grants in ExPORTER, allowing us to identify employees that have a PI whose funding is interrupted. Moreover, we can link these UMETRICS employees to their W2/1099 tax records and their UI wage

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

records, allowing us to observe various career outcomes, such as wages and employer characteristics.

ExPORTER. ExPORTER is publicly available data provided by the NIH.[9] and contains a wide range of information about NIH grants such as grant amounts and grant dates.

First, we identify all NIH R01 grants that are set to expire in a given year. We divide these expiring grants into two categories: (1) "interrupted" grants that were renewed after more than 30 calendar days and (2) "continuously funded" or "uninterrupted" grants that were renewed in 30 calendar days or less.[10]

Next, we link the PI(s) of each expiring R01 to any other NIH grants on which they were a PI in the 12 months prior to the focal R01's expiry. We define each PI's "lab" as the set of employees that were paid by any one of these grants. We use as an example the Principal Investigator Terrie E. Moffit, who was PI of the following three grants in 2015: R01AG048895, R01AG032282, and R01AG049789. By our definition, the employees paid by these grants comprise the "Moffit Lab' ' in 2015.

Third, we identify cohorts of labs that have one or more R01s set to expire in a given year. For each of these labs, we then determine whether *all* expiring grants were continuously funded and the lab belongs to the control group, or whether *at least one* of the expiring grants was interrupted and the lab belongs to the treatment group.[11] To continue the example from above, suppose that two of the grants that support the Moffit Lab (say, R01AG048895 and R01AG032282) are set to expire in 2018. Then the Moffit Lab will be part of the 2018 expiration-year cohort. Moreover, if *neither* of these two grants are interrupted, then the Moffit Lab is part of the control group for the 2018 cohort. If *either* of the grants are interrupted, then the lab is part of the treated group for this cohort.

As suggested, grant expiration years define cohorts of labs. That is, all labs with a grant set

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

[10] 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.

[11] Of course, some grants are never renewed. Though all grants are used to characterize the lab of a PI, the never-renewed grants play no role in determining whether a lab is interrupted or continuously-funded (that is, treated or untreated).

to expire in a given year belong to one cohort, all labs with a grant set to expire the next year belong to a second cohort, and so on. In our data, there are 13 expiration-year cohorts between 2006 and 2018. It is important to emphasize that the *expiration-years* that we use to define treatment cohorts are distinct from the *calendar-years* that we use to measure the employment outcomes of employees. This distinction will be particularly relevant when we discuss our final "stacked" estimation sample in the Estimation section.

Finally, we link UMETRICS personnel to NIH R01 grants in ExPORTER, which allows us to identify the individuals employed in each lab.[12] Since grant expiration years create cohorts of labs, the linkage of grants to UMETRICS employees creates cohorts of individuals whose PI's grants expire in different years. For instance, all UMETRICS employees employed in a lab with a grant set to expire in 2013 are members of the 2013 cohort. Moreover, we can use this data to create a treatment variable indicating whether a given employee is employed in a lab that experiences a funding interruption in a given expiration-year, thus creating a set of treated and control employees for each expiration-year cohort of labs. Again continuing our example from above, suppose that one of the grants (say, R01AG048895) supporting the Moffit Lab and set to expire in 2018 is interrupted. Then the *entire* Moffit Lab is interrupted in 2018, not only the employees paid by the interrupted grant. That is, the employees paid on the focal interrupted grant (R01AG048895) are defined as interrupted *and* the employees paid by other grants supporting the lab in 2018 are also defined as interrupted.[13] To be more precise, we confine our analysis to employees working in a lab (i.e. paid by a grant supporting the lab) at some point during the 12 months prior to expiration. This restriction makes it more likely that employees actually

[12]Since ExPORTER includes the universe of NIH R01 grants and UMETRICS only includes the subset of NIH R01s that pay at least one employee at a UMETRICS university, there will be some grants in ExPORTER that are not in UMETRICS (not all universities are included in UMETRICS). Similarly, since UMETRICS includes the universe of all grants (not just NIH R01 grants) at each UMETRICS university, some employees (i.e. those paid on grants other than NIH R01 grants) will not be linked to a grant in ExPORTER.

[13]It is worth pointing out that the same UMETRICS employee can be employed in labs supported by multiple grants with different expiry years and thus they can belong to multiple expiration-year cohorts. Moreover, they can be in the treatment (interrupted) group in one cohort and in the control (continuously funded) group in another cohort. For instance, an employee may be employed by a lab with a grant set to expire in 2008 and thus belong to the 2008 expiration-year cohort. The same employee may also later be employed by a lab with a grant set expire in 2013 and thus also belong to the 2013 expiration-year cohort. In this case, the employee belongs to both the 2008 and 2013 cohorts and whether they are a treated or control employee within each cohort depends on whether the expiring grant is interrupted or continuously funded.

experience the effects (or non-effects) of an interruption to the lab's funding.

For most of our results, we obtain estimates of 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 be better able than single-R01 labs to cushion the shock of a funding interruption. Tham (2021) shows that funding interruptions cause a substantial reduction in grant expenditures for single-R01 labs but a small or no reduction in spending for multiple-R01 labs.

Census Data Career Outcomes. UMETRICS employees can be linked to Census data, 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.

First, we link UMETRICS employees to the universe of W2 tax records.[14]Each W2 record is an employee paired with the federal tax identification number (EIN) of an employer, and contains information on yearly wages. Second, we link employees to UI wage records through the Longitudinal Employer-Household Dynamics (LEHD) program. The LEHD data augment the W2 data and allow us to 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 (see below). Finally, we link UMETRICS employees to the universe of 1099 tax records, which are available through 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. Combined, the W2, LEHD, and 1099 (ILBD) data allow us to track the entire domestic employment and earnings history for every UMETRICS employee between 2005-2018.

The first set of variables we construct from this data are three mutually exclusive (in a

[14]These available to qualified employees through agreements between the US Census Bureau and the Internal Revenue Service (IRS)

given calendar year) employment outcome variables. First, we create an indicator for whether the employee is absent from all three of our administrative wage sources – in other words, this variable indicates that, for a given year, the employee is *not* in the W2, LEHD, or ILBD data. Second, we use the EINs in the W2s and LEHD, along with a public-use list of university EINs from the Integrated Postsecondary Education Data System (IPEDS), to determine whether the employee receives positive earnings from an IPEDS university.[15] Lastly, we create a third, catch-all, category that is the complement of these two variables and essentially captures self-employment or employment in industry or government.[16] These three variables allow us to exhaustively track where employees go after their lab experiences an interruption.

We also construct two secondary employment outcomes that are related to the indicator for receiving positive university earnings. For each year, we identify whether the employee receives positive earnings from *their own UMETRICS university* and whether the employee receives positive earnings from *some other university in IPEDS*. These two variables help us understand whether the impact of a grant funding interruption on attachment to academia is driven by its impact on own-university employment or also has effects on academic attachment more broadly.

In addition to employment outcomes, we also measure yearly wages for each employee between 2005 and 2018. These are derived from a combination of W2, LEHD, and 1099 (ILBD) wages.[17] By tracking wages, we can measure not only whether interruptions alter employment patterns, but also whether interrupted employees are propelled toward careers with higher or lower earnings than they would have received in a counterfactual world in which they were employed by a continuously-funded lab. If an interruption pushes employees out of academia and into industry, then wages may rise. However, if an interruption derails an employee's academic career, wages could decline as they settle for

[15] For clarity, note that IPEDS contains EINs for most U.S.-based universities, not just UMETRICS universities (though all UMETRICS universities are included in IPEDS).

[16] As noted, these three outcomes are mutually exclusive in a given year, so for a given employee-year, they always add to 1.

[17] Specifically, we define a employee's total earnings in a given year to be their wages from self-employment (ILBD) plus the maximum of their W2 and LEHD wages. That is $wage_{total} = wage_{ilbd} + max\{wage_{w2}, wage_{lehd}\}$

less remunerative academic jobs (e.g. they settle for a non-tenure track faculty position).

Stacked Panel Estimation Sample. Our final estimation sample is constructed by linking the employee-expiration-year data to the employee-calendar-year data. This combination yields a "stacked" panel structure, where observations are uniquely defined by an employee, calendar-year, and grant expiration-year. For a given grant expiration-year, the sample includes all employees employed in labs with an expiring grant in that year. The employment and earnings outcomes for each of these employees are then tracked over calendar years between 2005 and 2018. Thus, within an expiration-year, we have a standard balanced panel structure where observations are uniquely defined by an employee and calendar-year.[18] This stacked data structure is becoming more common in difference-in-differences (DiD) settings with multiple time periods and staggered treatments that affect different groups at different times.

As noted, our stacked data implies that the same employee/calendar-year can belong to multiple expiration-year cohorts and thus appear multiple times in our data, which can happen when the employee switches labs or is employed in a lab supported by multiple grants with different expiration years. This means that it is possible for a employee to belong to an interrupted (treated) lab in one expiration-year and then belong to a continuously-funded (control) lab in a later expiration-year (and vice versa), making it possible for an already-treated employee to serve as a control for a currently-treated employee. To mitigate this, we impose a 2-year "clean window" around each expiration-year. That is, we do not include an employee as a control within an expiration-year cohort if they themselves are treated two years before or after the focal expiration-year. For instance, an employee in a continuously-funded lab in the 2010 expiration-year cohort is not used as a control if they were employed by an interrupted lab in any of the 2008-2012 expiration-year cohorts.

[18] In other words, each expiration-year cohort is treated as its own balanced panel data set, and then each of these data sets are stacked on top of each other.

4 Estimation

4.1 Estimation with Stacked Data

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. expiration year), which was discussed in the data section; (2) using a modified version of the estimator developed by Callaway and Sant'Anna (2020) for DiD with multiple treatment periods which compares treatment and control units within the same event.[19] (1) is a data formatting step and (2) uses the data format to avoid the issues laid out in Goodman-Bacon (2018).

This two-step approach takes advantage of the fact that, in our setting, both treated and untreated employees have well defined "treatment" periods (i.e. expiration-years). In effect, stacking the data by expiration-year lets us view each cohort as a separate quasi-experiment with employees being either treated or untreated in a specific expiration-year. Since the observations in this setting are uniquely defined by a employee/calendar-year/expiration year, but the the standard Callaway and Sant'Anna (2020) estimator directly applies to a more traditional un-stacked panel data structure (where observations are uniquely defined by an individual/year) we slightly modify the estimator to account for our stacked structure.[20] Specifically, we relabel the events (i.e. expiration years) and the calendar years in a way that ensures no overlap in the calendar years.

To ensure no overlap in calendar years – and thus ensure that treated units within an expiration-year cohort are only compared to control units in the same cohort – we first relabel the twelve expiration-year cohorts between 2006 and 2018. Let g index expiration years between 2006 and 2018. Then for each g, the newly labeled "expiration-year'' is $g^{'} = g + (g - 2006) * (2018 - 2005 + 1) = g + (g - 2006) * 14$. The $g - 2006$ is the first expiration-year subtracted from each subsequent expiration-year. The $2018 - 2005 + 1 = 14$

[19] Stacking can be thought of as exact matching of treatment and control units on (counterfactual) treatment timing.

[20] We could have reformatted our data into a standard employee/year panel if we were willing to regard a employee's *first* interruption as the treatment of interest. However, this approach ignores valuable information on employees who experience multiple interruptions at different points in their careers.

is the number of years spanned by the calendar years. This process results in the 2006 expiration-year cohort remaining labeled as 2006, the 2007 cohort being relabeled as 2021, the 2008 cohort being relabeled as 2036, and so on. This ensures that there is always a distance of 15 periods between each cohort, which will accommodate the 14 calendar years with no overlap.

Next, for each of the expiration-year cohorts, we relabel every calendar-year so that each cohort gets its own uniquely labeled calendar years. Let t index calendar years between 2005 and 2018. Then for every expiration-year g and every calendar-year t, the newly labeled calendar years are generated in the following way: $t'_{gt} = t + (g - 2006) * (2018 - 2005 + 1) = t + (g - 2006) * 14$. This ensures that each expiration-year cohort gets its own set of calendar years. For instance, the calendar years for the 2006 cohort remain 2005-2018, the calendar years for the 2007 cohort are relabeled as 2019-2032, the calendar years for the 2008 cohort are relabeled as 2033-2046, and so on. A full listing of newly labeled expiration years and their accompanying newly labeled calendar years are available in the appendix.

The stacking procedure gives us a unit-by-time panel dataset where each (employee, expiration-year) combination is a unit and the newly-labeled calendar years are the time units. The newly-labeled expiration years define groups/treatment cohorts. This approach incorporates the advantages of the Callaway and Sant'Anna (2020) estimator with a well-defined event year for control employees.

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 for interrupted and continuously-funded employees would have evolved in parallel if the

funding interruption had not occurred. Given the stacked structure of our data, we require parallel trends to hold for *all* expiration-year cohorts between 2006 and 2018.

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.[21] 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 1 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 (W2, 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 1A) 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

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

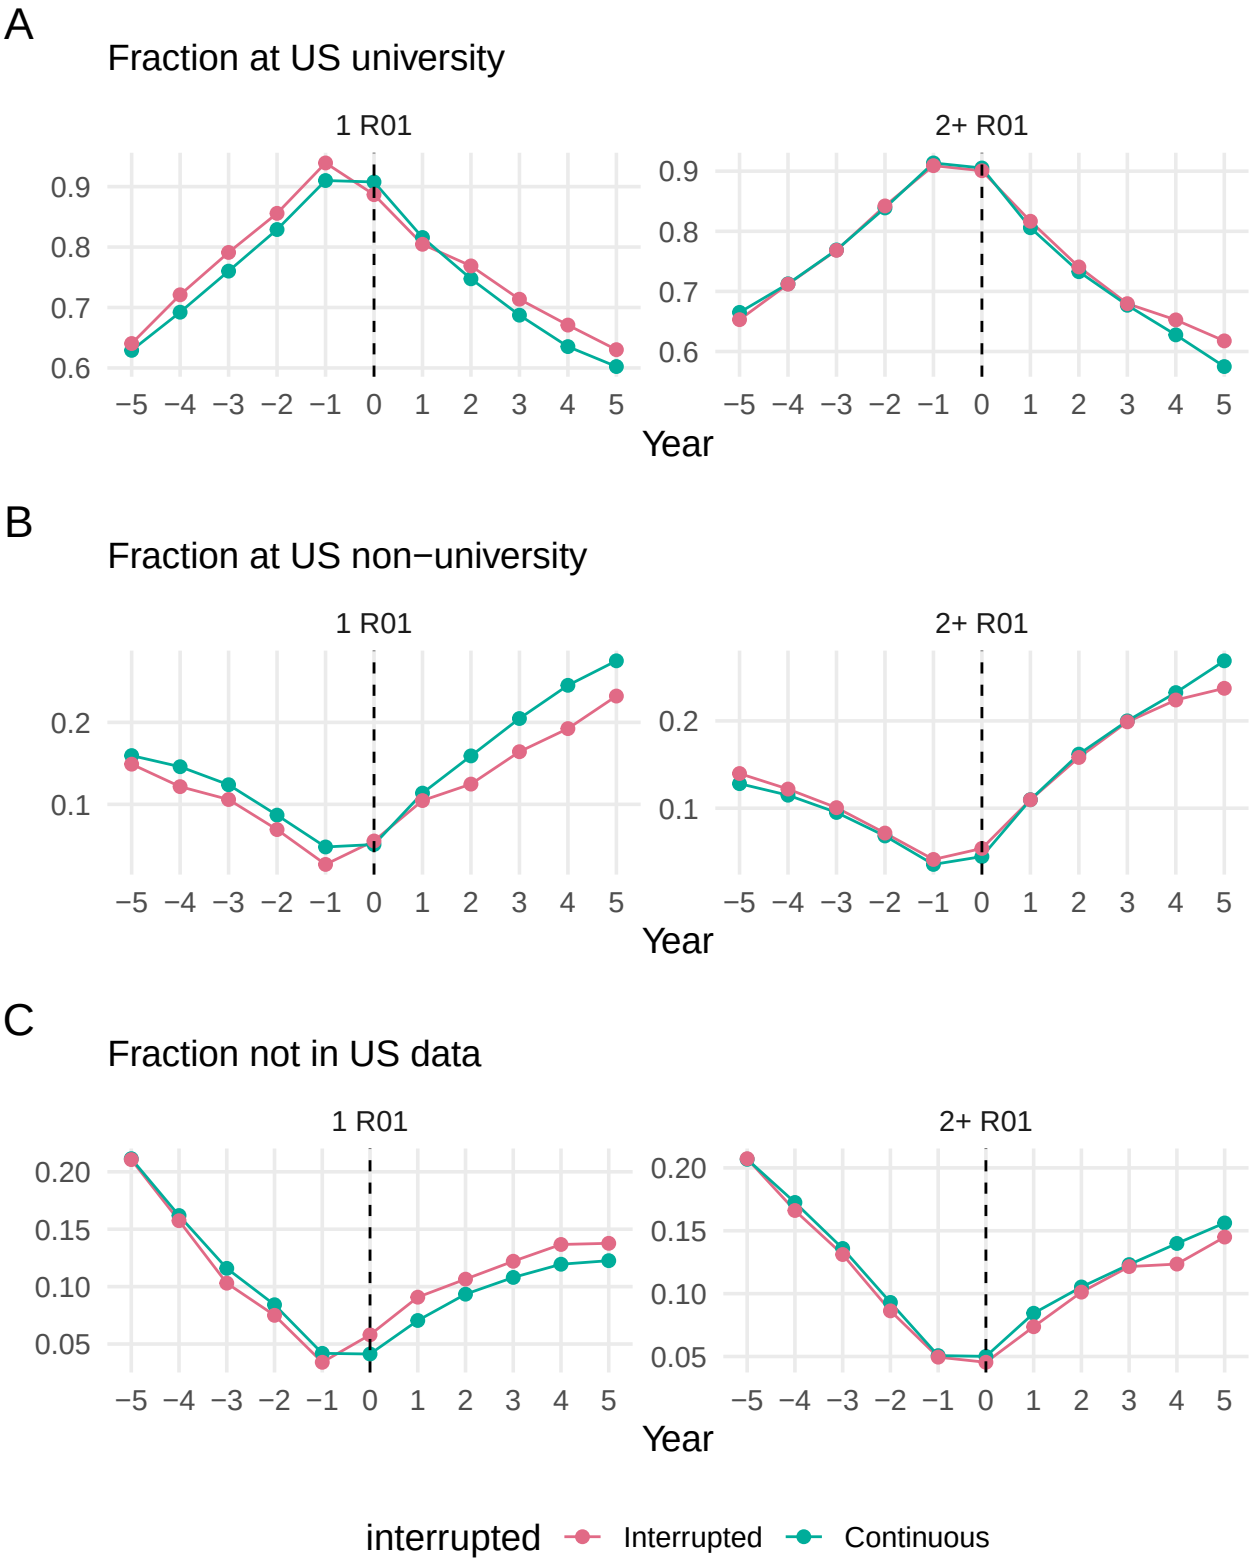


Figure 1: This figure shows the average probability that an individual is in one of three categories (paid by a US university, paid by a US non-university, or not paid in the US) from five years before interruption to five years after interruption. Figures in the left column are for individuals in a 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 1B) 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 1C) 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 crucial 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.

Though the raw means presented in Figure 1 provide compelling visual evidence for the effects of funding interruptions on the employment outcomes research personnel in single-R01 labs, we clarify the magnitudes and precision of these effects in Figure 2, which shows event studies derived from our adaptation of the the Callaway and Sant'Anna (2020) estimator to our stacked data structure. 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.

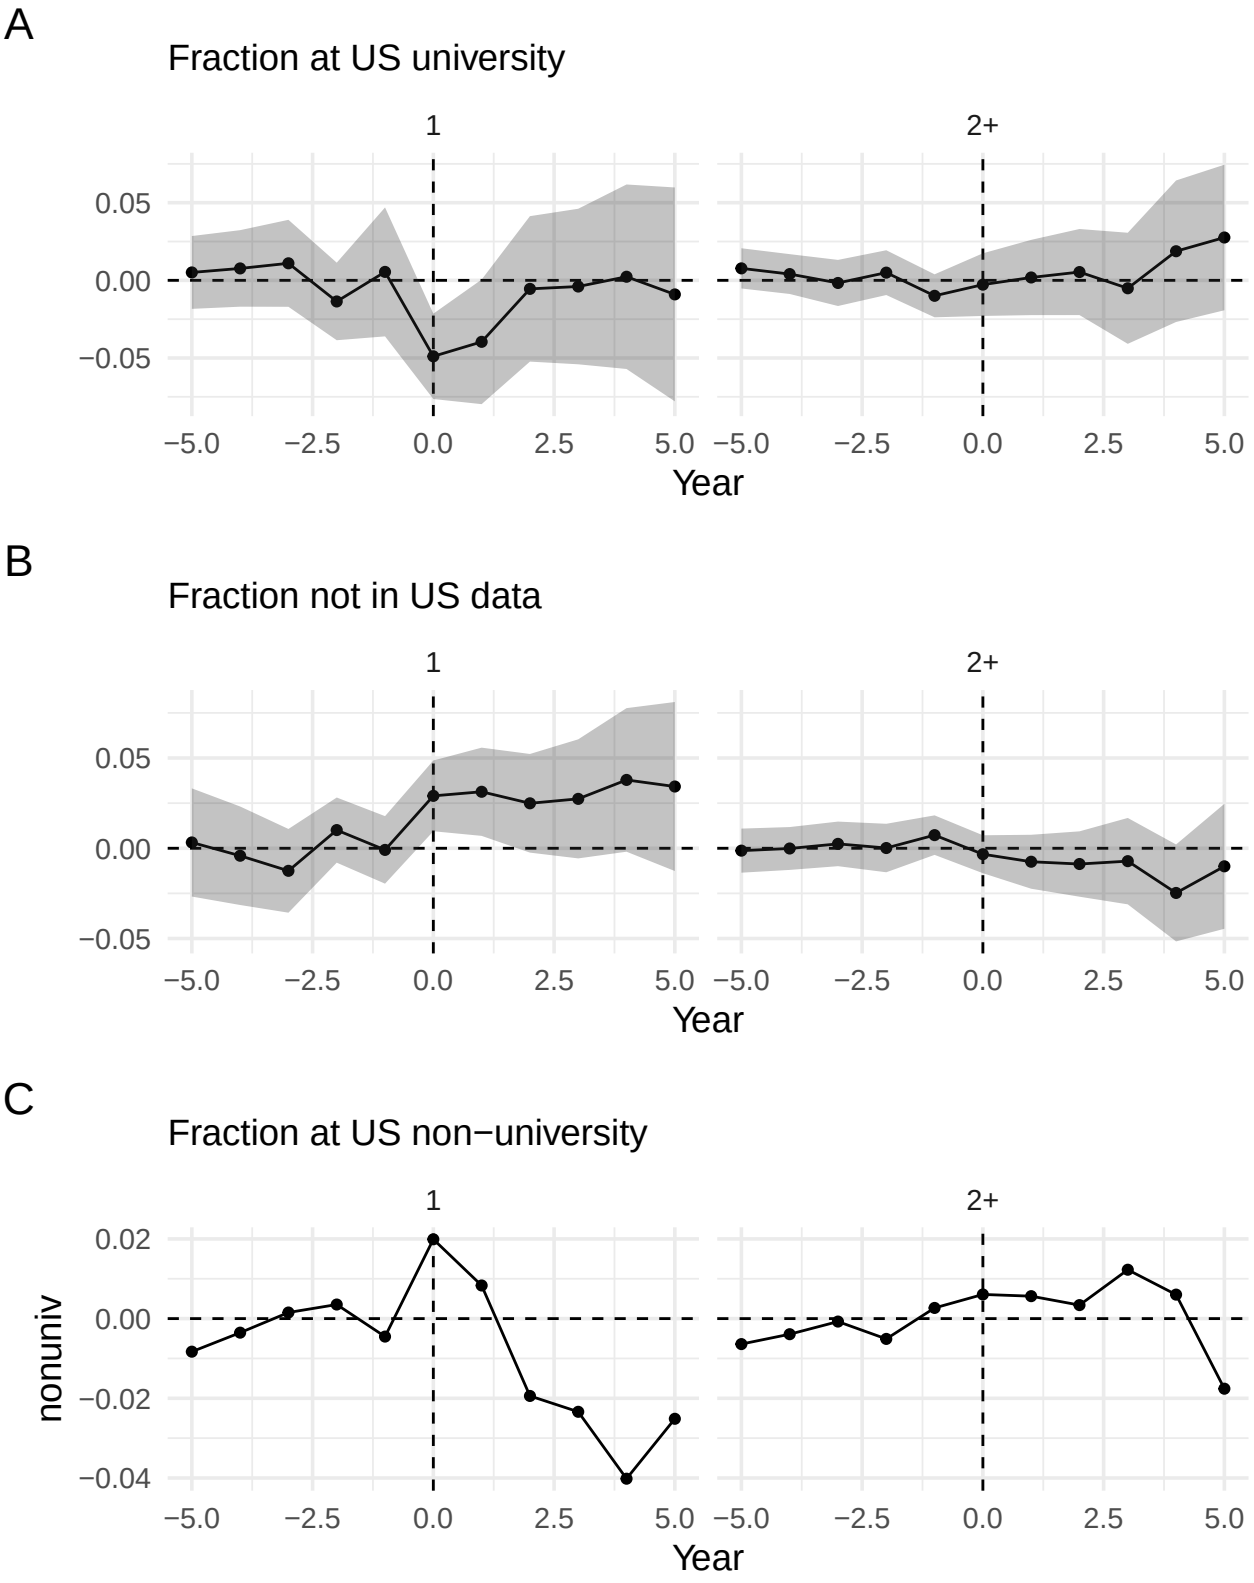


Figure 2: This figure shows event studies of the effects of interruptions on employees estimated with our modified Callaway-Sant'Anna (2020) estimator. The outcomes are the probability that an individual in one of three categories: paid by a US university (A), not paid in the US (B), or paid by a US 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.

Table 1: This table 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 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 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 interrupted R01 level.

Outcome	No. of R01s	Estimate	Std Err	95% CI Lower	95% CI Upper
Paid by US uni	1	-2.1	1.8	-5.7	1.5
Paid by US non-uni	1	-0.9	1.5	-3.8	2.0
Not in US data	1	3.0	1.2	0.7	5.3
Paid by US uni	2+	0.6	1.3	-2.0	3.2
Paid by US non-uni	2+	0.4	1.0	-1.6	2.4
Not in US data	2+	-0.9	0.9	-2.6	0.7

Table 1 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 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 2 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 W2 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

Table 2: This table 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 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 interrupted R01 level.

Row	Outcome	Born	Estimate	Std Err	95% CI Lower	95% CI Upper
1	Not in US data	Not US	1.8	2.3	-2.7	6.4
2	Not in US data	US	3.5	1.1	1.3	5.7
3	Paid by US non-uni	Not US	-2.9	1.6	-6.0	0.2
4	Paid by US non-uni	US	0.1	1.9	-3.7	3.9
5	Paid by US uni	Not US	1.1	2.7	-4.1	6.3
6	Paid by US uni	US	-3.5	2.1	-7.7	0.6

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.

Table 2 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 only 1.1 percentage points (statistically insignificant) less likely.

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. Table 3 presents aggregate post-treatment

Table 3: This table 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 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 interrupted R01 level.

Outcome	Occupation	Estimate	Std Err	95% CI Lower	95% CI Upper
Paid by US uni	Faculty	-2.1	1.8	-5.6	1.4
Paid by US non-uni	Faculty	2.2	1.5	-0.8	5.1
Not in US data	Faculty	0.3	1.5	-2.7	3.3
Paid by US uni	Postdoc/Grad	-5.4	3.1	-11.5	0.7
Paid by US non-uni	Postdoc/Grad	-0.7	2.4	-5.4	4.0
Not in US data	Postdoc/Grad	6.1	2.4	1.5	10.7
Paid by US uni	Other	-0.7	2.9	-6.4	4.9
Paid by US non-uni	Other	-1.9	2.5	-6.7	2.9
Not in US data	Other	2.6	1.5	-0.2	5.5

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 3A 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 total wages (arcsine-transformed) 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 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.

Figure 3B 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

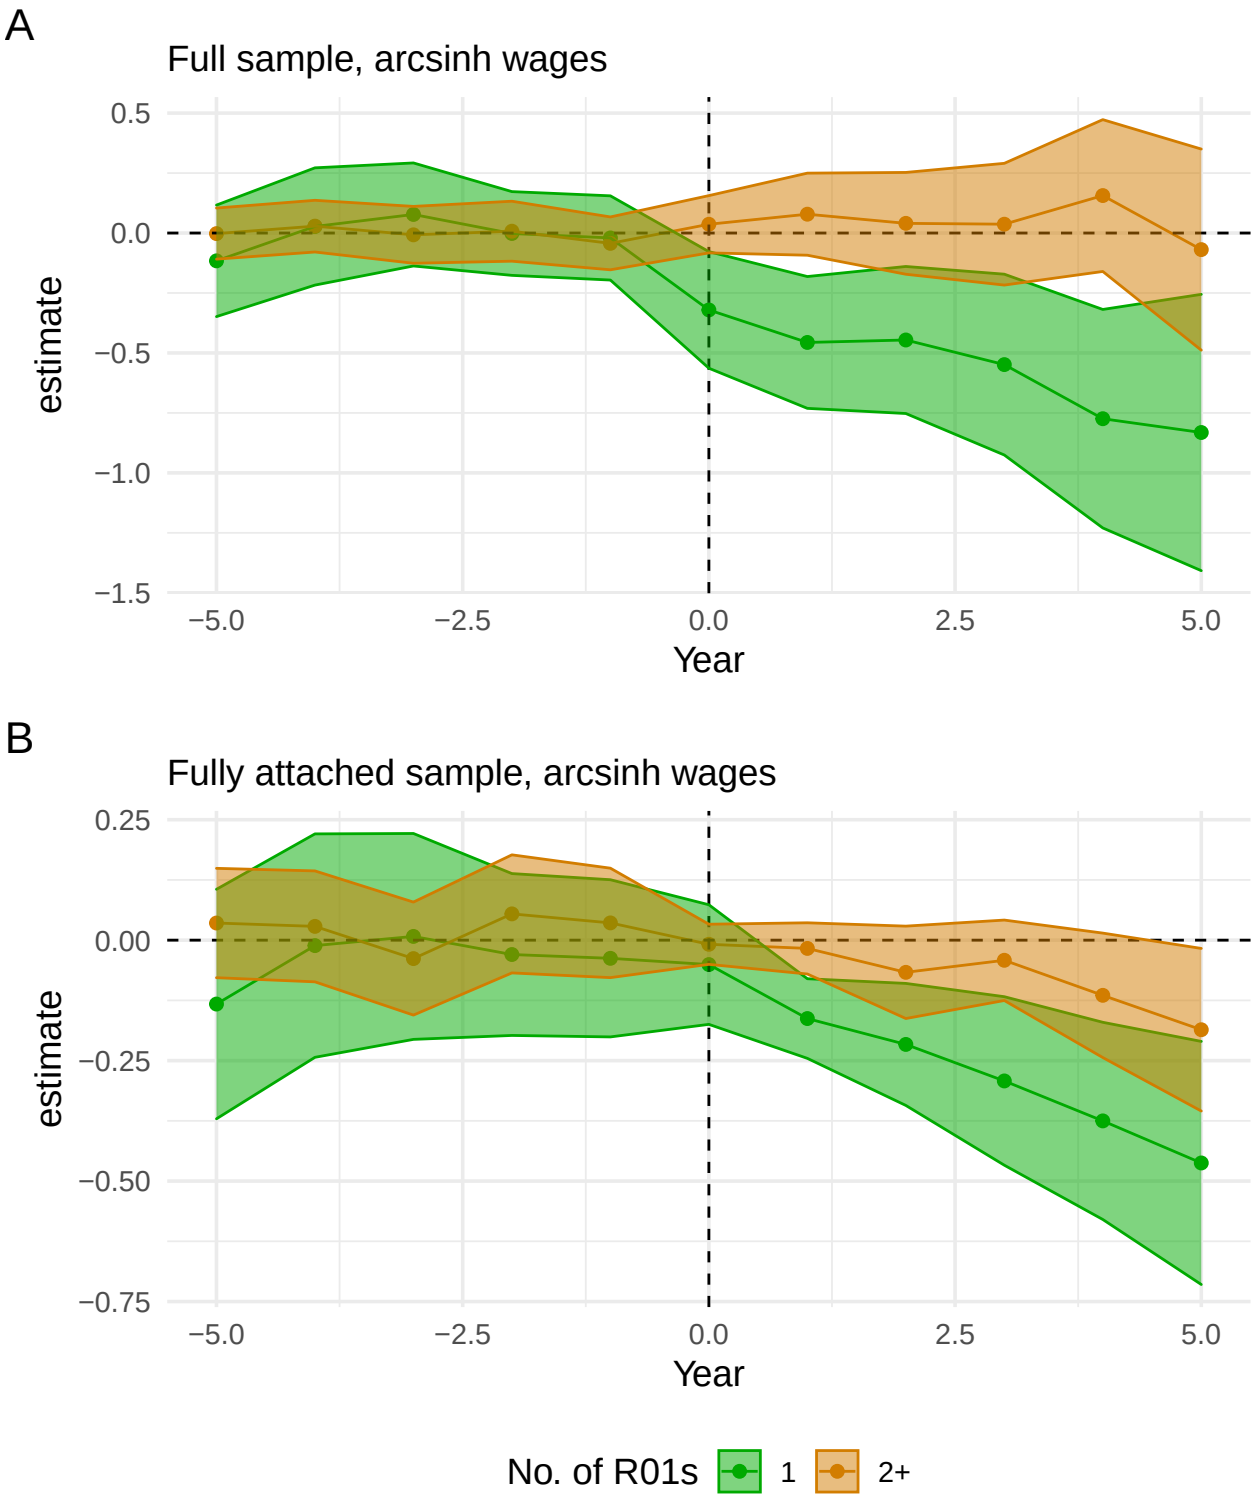


Figure 3: This figure shows time-varying effects (event studies) of an interruption on the arcsine of total earnings. The expiration of a lab's grant takes place at time 0, and the estimated interruption effects range from 5 years before to 5 years after an interruption. The green series is for individuals 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

Table 4: This table 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 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 interrupted R01 level.

# R01s	Sample	Estimate	Std Err	95% CI Lower	95% CI Upper
1	All	-0.52	0.14	-0.79	-0.25
1	Partly attached	-0.36	0.12	-0.61	-0.12
1	Fully attached	-0.23	0.06	-0.35	-0.10
2+	All	0.05	0.10	-0.14	0.24
2+	Partly attached	0.06	0.07	-0.09	0.20
2+	Fully attached	-0.06	0.04	-0.13	0.01

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.

Table 4 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 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.

6 Conclusion

A major theme in the classic *Science, the Endless Frontier*, was how to set up a system that provided the stability for doing basic research (Bush 2020; Stephan 2013). Over time, this stability seems to have been compromised in important ways (Stephan 2013) but our understanding of its effects are still incomplete. We make progress on this issue by studying how grant funding delays (aka "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 tend to be on a longer timescale, the costs to scientific workers occur on a shorter timescale. The benefits of policies and new funding models that reduce the occurrence of similar disruptions may be understated as a result.

References

Alberts, Bruce, Marc W Kirschner, Shirley Tilghman, and Harold Varmus. 2014. "Rescuing US Biomedical Research from Its Systemic Flaws." *Proceedings of the National Academy of Sciences* 111 (16): 5773–77.

Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for Inven-

tion." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 609–26. Princeton University Press.

Azoulay, Pierre, Joshua S Graff Zivin, Danielle Li, and Bhaven N Sampat. 2015. "Public r&d Investments and Private-Sector Patenting: Evidence from NIH Funding Rules." *The Review of Economic Studies*.

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.

Balsmeier, Benjamin, and Maikel Pellens. 2014. "Who Makes, Who Breaks: Which Scientists Stay in Academe?" *Economics Letters* 122 (2): 229–32.

Bush, Vannevar. 2020. *Science, the Endless Frontier*. Princeton University Press.

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

Cheng, Stephanie. 2021. "What's Another Year? The Lengthening Training and Career Paths of Scientists." Harvard University Working Paper.

Conti, Annamaria, and Fabiana Visentin. 2015. "Science and Engineering Ph. D. Students' Career Outcomes, by Gender." *PloS One* 10 (8): e0133177.

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."

Goodman-Bacon, Andrew. 2018. "Difference-in-Differences with Variation in Treatment Timing." National Bureau of Economic Research.

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.

Janger, Jürgen, and Klaus Nowotny. 2016. "Job Choice in Academia." *Research Policy* 45 (8): 1672–83.

Jones, Benjamin F, and Lawrence H Summers. 2020. "A Calculation of the Social Returns to Innovation." National Bureau of Economic Research.

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.

Milojević, Staša, Filippo Radicchi, and John P Walsh. 2018. "Changing Demographics of Scientific Careers: The Rise of the Temporary Workforce." *Proceedings of the National Academy of Sciences* 115 (50): 12616–23.

Myers, Kyle, and Lauren Lanahan. 2021. "Estimating Spillovers from Publicly Funded r&d: Evidence from the US Department of Energy." *SSRN* 3550479.

Oreopoulos, Philip, Till Von Wachter, and Andrew Heisz. 2012. "The Short-and Long-Term Career Effects of Graduating in a Recession." *American Economic Journal: Applied Economics* 4 (1): 1–29.

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

Roach, Michael, and Henry Sauermann. 2017. "The Declining Interest in an Academic Career." *PloS One* 12 (9): e0184130.

Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy* 98 (5, Part 2): S71–102.

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

Stern, Scott. 2004. "Do Scientists Pay to Be Scientists?" *Management Science* 50 (6): 835–53.

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

Thompson, Derek. 2021. "America Needs a New Scientific Revolution." *The Atlantic*. https://www.theatlantic.com/ideas/archive/2021/11/grants-american-scientific-revolution/620609/.

Weitzman, Martin L. 1998. "Recombinant Growth." *The Quarterly Journal of Economics* 113 (2): 331–60.

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.[22] 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**.[23] 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 5 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"

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

[23] This is laid out in more detail in Section 5.3 of the NIH Grants Policy Statement.

Table 5: 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.[24]

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

[24] 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."[25] 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)

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