Science, Interrupted: Funding Delays Reduce Research Activity but Having More Grants Helps

Wei Yang Tham*

2020-11-02

Abstract

I study how scientists respond to interruptions in the flow of their research funding, focusing on research grants at the National Institutes of Health (NIH), which awards multi-year grants that researchers can renew. However, there can be delays during the renewal process. In a two-year period around these delays, I find that interrupted labs reduce spending on inputs by over 20% in an average month but over 90% at the lowest point. This change in spending is mostly driven by a decrease in payments to employees that is partially mitigated when scientists have other grants to draw on.

*Harvard Business School, wtham@hbs.edu, ORCID. I am grateful to Bruce Weinberg, David Blau, and Kurt Lavetti, for their invaluable guidance when I started this project as a graduate student. I am thankful to Kyle Myers for his feedback and support in the later stages of this project. This work would not have been possible without the support of IRIS at the University of Michigan, with special thanks to Natsuko Nicholls and Beth Uberseder. The paper benefitted from helpful comments by discussants and participants at the: Aug 2018 UMETRICS meeting, Census ARiS Brown Bag, NBER-IFS International Network on the Value of Medical Research meeting, OSU Micro Lunch, SUNY Albany Department of Economics Seminar, MTEI Seminar at EPFL, RISE2 Workshop at the Max Planck Institute for Innovation and Competition (MPI), the Economics of Science & Engineering Seminar at Harvard Business School, and the 14th Workshop on the Organisation, Economics, and Policy of Scientific Research at MPI. Research reported in this paper was supported by the National Institute on Aging of the National Institutes of Health under Award Number R24AG048059 to the National Bureau of Economic Research. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Institutes of Health or the NBER.

1 Introduction

In many fields of science, research is a resource-intensive endeavour. It requires people, capital, and the management of these resources. In addition, scientists have to obtain and manage the funding necessary to acquire these inputs. This includes dealing with the possibility that funding may not arrive in the amount or at the time they want it to. How scientists respond to uncertainty and liquidity constraints in funding is therefore an important part of the research production function.

In general, however, this aspect of a scientist's job is difficult to observe on a large scale. The UMETRICS dataset (Lane et al. 2015), which consists of administrative data from universities on transactions from sponsored projects, helps to bridge this gap. I use UMETRICS to study how scientists respond to funding delays or "interruptions" in the context of research funded by the National Institutes of Health (NIH).

I first document that when funding is guaranteed and available, scientists tend to maintain spending at a steady level (Figure 1). On average, after a "ramping up" period in the first year of a *project period* (the NIH term for a multi-year grant), spending is relatively flat until the final year of the project period, when it steadily decreases. This suggests that in the absence of uncertainty about funding or liquidity constraints (and conditional on how the NIH disburses funds), scientists have a preference for a stable rate of spending.[1]

Next, I study how scientists respond to funding delays or "interruptions", focusing on a particular type of NIH grant, the "R01". R01 grants are generally regarded as being necessary to establish an independent research lab in the biomedical sciences. They have to be renewed periodically (typically every four to five years) at the end of each project period, at which point the following scenarios may occur:

a. The project's new funding stream begins as soon as its previous one ends.

b. The project is *interrupted*: its new funding stream only begins some time after its

[1] By policy, the NIH funds projects in annual increments so this places some restrictions on the spending trajectory. The spending trajectory is also shaped by the fact that employee costs, which entail a longer commitment, tend to be the largest costs in a project.

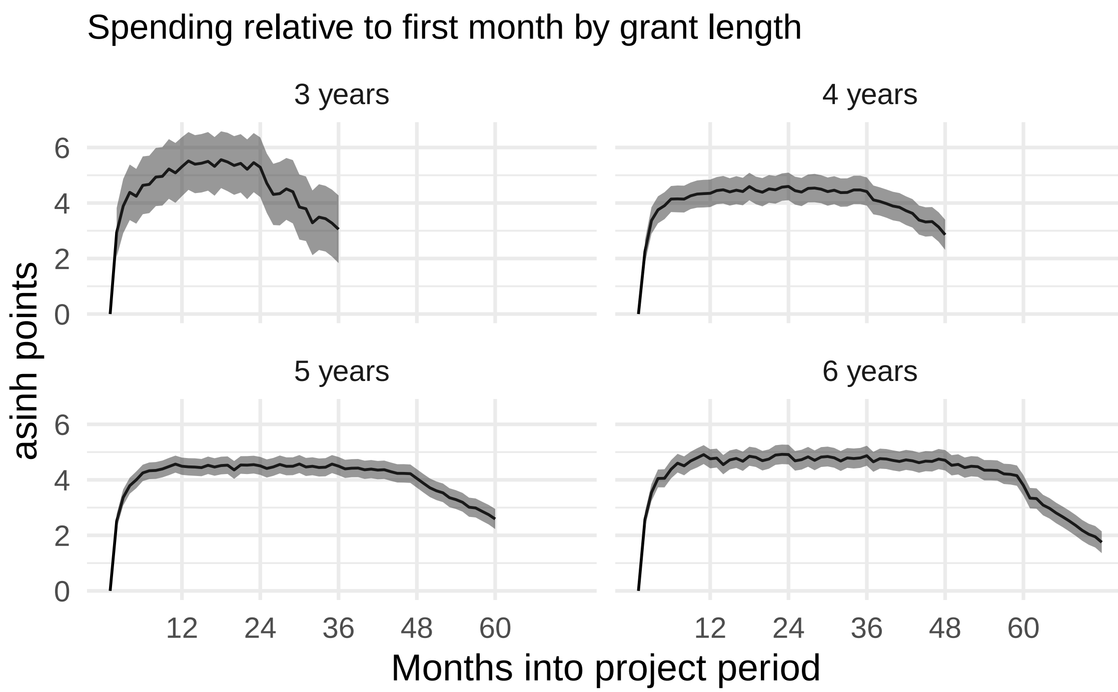

Figure 1: This figure shows spending per month for R01 project periods that last three, four, five, or six years, relative to spending in the first month of the project period. Estimates are from a regression of total expenditure (arcsinh-transformed) on a set of dummies for each month in a project period with project period fixed effects, with the first month of the project period as the excluded category. Separate regressions are run by project period length. Standard errors are clustered at the project-period-level.

previous one ends.[2]

Interruptions can arise for different reasons. A funding agency that is uncertain about its budget may engage in "precautionary saving" and delay spending to the end of the fiscal year (Liebman and Mahoney 2017). There may also be disruptions to the funding allocation process, such as the government shutdown of Fiscal Year 1996 (Mervis and Marshall 1996), which slowed down the processing of paperwork and led to peer review meetings being postponed.

An interruption can be thought of as a combination of (a) a *liquidity* shock from the scientist's inability to access funding for some period of time and (b) an *uncertainty* shock about when or even whether they will get the funding in the first place. Although I do not distinguish between these two mechanisms, their combined effect in the form of interruptions is an important policy question. Delays in NIH funding are a real concern among researchers (DrugMonkey 2009). Understanding their role as a potential impediment to science helps policymakers to determine how much attention should be paid to this issue.

I estimate the effect of interruptions with a difference-in-differences design that compares outcomes for interrupted and uninterrupted labs, defining "interrupted" projects as those where funding was renewed after more than 30 days belong to scenario.[3] To allow for the possibility that principal investigators (PIs) can dampen the effects of interruptions with other grants, I run the analysis separately on PIs with one R01 and PIs with multiple R01s.[4]

I find that interrupted labs with one R01 reduce spending significantly. In the average month over a two-year period centered on the expiry of an R01 project period, interrupted labs spend 26% less than uninterrupted labs.[5] However, the change in spending is not uniform over time but V-shaped.[6] At its nadir, spending is 94% lower for interrupted

[2] It is also possible that the project is not renewed. I focus solely on projects that are eventually renewed by the end of the fiscal year.

[3] The 30-day threshold is chosen to approximate a month. Most funding begins on the first of the month, thus the "arrival" of new grant funding can be thought as occurring on a monthly basis.

[4] I include "R01-equivalent" and "P01" grants in this measure. The Data section elaborates.

[5] This figure is the arithmetic mean of event study estimates after each estimate has been converted from log points to percentage changes.

[6] For longer interruptions, this is more like a U-shape. The Online Appendix contains estimates from

labs. This decrease starts before the official expiry date of their grant, indicating that labs learn about the possibility of an interruption before the official expiry date of their grant reallocate their spending in response. After R01 expiry, spending drops sharply before starting to recover, taking about nine months after expiry to catch up with uninterrupted labs.

When a PI has multiple R01s, spending remains stable throughout for interrupted PIs, indicating that there is fungibility across research grants. This is supported by the employee-level analysis. Across occupations, employees who are linked to multiple R01s either experience a lesser or zero decline in the probability of being paid by a grant, whether by their PI or by any grants at all.

I also look at whether PIs adjust different components of spending differently in response to interruptions. For PIs with one R01, both vendor and labor spending decrease substantially (by over 90% at their lowest points), but vendor spending does not recover as quickly. For PIs with multiple R01s, there is some decrease in the number of employees but vendor payments are relatively stable. The decrease in employees for both PI types may be due to labor expenditure constituting a larger share of spending and entailing longer-term commitments that PIs are unwilling to commit to until they know their funding status. However, this does not necessarily mean that employees are being disemployed by their institution or even removed from the research team, as there may be alternative sources of funding for some employees (e.g. teaching positions for graduate students).

In my final set of results, I estimate the impact of a funding interruption on research output as measured by publications and citation-weighted publications. However, these estimates are not precise enough to determine whether interruptions affect output and, if they do, to what extent. This illustrates how traditional measures such as publications and patents may not provide the complete picture because they occur at a lower frequency and only capture one aspect of the research production process.

One limitation of this paper is interruptions are not randomly assigned. While interrup-

specifications that allow the length of an interruption to vary.

tions are driven in part by external events such as government shutdowns, the NIH may prioritize projects or PIs that are perceived to be of higher quality. This is less likely to be a problem for the results where inputs are an outcome, given that the use of inputs is more likely to be driven by budget constraints and the specific needs of the project. This is more likely to bias upward (in magnitude) the results involving publications, although the pre-trends for publications do not indicate that interrupted PIs were on a less productive trajectory leading up to the year they were interrupted.

Another limitation is that I do not observe the full set of funds that are available to a scientist. Spending of internal funds, such as those directly provided to a scientist by their institution, is not observed within UMETRICS. In addition, my measure of PI spending is limited to spending through NIH grants. This ensures a high degree of accuracy in linking NIH PIs to transactions. This naturally raises the question of whether the amount of funding from non-NIH grant has a substantive effect on the results discussed so far. This is unlikely for two reasons.

First, the results on whether interrupted employees continue to be paid on *any grant* are consistent with the overall set of results and thus do not give us a reason to think that there is a substantial pool of non-NIH grants being used to offset the effects of interruptions. Second, other research and government statistics show that if a research group is federally funded, it is also mostly federally funded and the NIH is the largest federal funder of life sciences research (details in Online Appendix). In short, focusing on NIH funding provides substantial coverage of researcher funding.

This paper builds on work using granular data to unpack the role of the "lab" in science, dating back to the anthropological work of Latour and Woolgar (2013) in 1979. On a larger scale, Conti and Liu (2015) use a complete personnel roster of principal investigators in the MIT Department of Biology from 1966 to 2000 to study the role of different types of personnel in research production. Bae, Sattari, and Weinberg (2020) use the UMETRICS dataset as well to estimate the marginal product of scientific funding and show how employee composition changes when funding increases. This paper adds to the body of work by examining how uncertainty and liquidity constraints affect the use of research inputs and showing the usefulness of high frequency data in studying the knowledge production process.

This paper is also part of a literature in innovation economics studying how uncertainty affects the productivity and choices of innovators. Azoulay, Graff-Zivin, and Manso (2011) study the Howard Hughes Medical Institute (HHMI) Investigator Program, which gives grantees more freedom over research direction and effectively gives them longer grant cycles compared to R01s, thus insulating them from the type of disruptions that can arise in the R01 renewal process. They find that HHMI scientists are more likely to produce high-impact papers and explore new research directions. While the insights from Azoulay, Graff-Zivin, and Manso (2011) are important, there are practical difficulties to expanding a resource-intensive program like the HHMI's. Thus, understanding where improvements can be made within the current system is important as well.

These results highlight that when a funding agency decides to delay renewal of a project, the decision may not be costless. Even when the project is eventually funded, there can be disruptions to the use of inputs, team capital (Jaravel, Petkova, and Bell 2018), and the employment or training of personnel. This has two major implications for how we fund projects. Firstly, it suggests that there is value to having the budgets of science funding agencies planned over a longer-term horizon to reduce uncertainty. Secondly, funding agencies delay projects if they expect that higher quality projects may be available later in the fiscal year. Agencies should consider that the cost of disrupting a project could be bigger than the improvement in quality from delaying its decision, especially if their measures of project quality are imperfect.

2 Background and conceptual framework

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 impact 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.[7] 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 choose to delay funding for new or renewed projects in response to uncertainty about the federal budget since this implies uncertainty about the size of the NIH's budget for the fiscal year.

2.2 Scientist perspective

The R01 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.[8] The same *project* can last for multiple *project periods*.

[7] 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.

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

Ideally, a researcher wants to maintain R01 funding for as long as possible. Toward the end of each project period, the principal investigator (PI) has to apply to renew their project for another project period of 4-5 years. PIs typically apply for renewal 1-2 years before a project period ends in order to receive funding continuously. In addition to the time taken to prepare the application itself, PIs have to take into account other factors such as potentially having to resubmit an application that is rejected the first time.

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.

3 Data and variable construction

3.1 UMETRICS

I use the 2019 release of the UMETRICS data, which is housed at the Institute for Research on Innovation and Science (IRIS). UMETRICS core files are administrative records of transactions from sponsored projects from 31 member universities. The time span covered by each university's records varies, spanning 2001 to 2018 overall. Payments from a project can go to one of three categories: vendors, subawards, or personnel.[9]

[9] Summary documentation for 2019 UMETRICS data is also available here.

3.1.1 Lab/PI total direct expenditure

A key outcome variable is direct expenditure from grants, which excludes the overhead costs that are paid to universities as a percentage of a grant award. Although I define the timing and length of funding delays around the R01 grant, I sum up outcomes to the level of the PI/lab. Specifically, for each R01, I find its associated PIs at the point of renewal. I then sum up spending for each PI across all NIH grants that they are associated with at a given point in time.[10]

3.1.2 Lab/PI vendor and labor expenditure

I repeat the procedure above for payments to vendors and payments to labor. Payments to vendors includes purchases of equipment or services. UMETRICS does not include salaries, so payments to labor are backed out as the remainder after subtracting vendor and subaward payments from total expenditure (i.e. $Labor = Total - Vendor - Subaward$).

3.1.3 Lab/PI employee counts

Count of the number of employees paid by a PI through the PI's NIH grants.

3.1.4 Employee-level outcomes

The next part of my analysis is at the individual employee level. The UMETRICS data contains unique employee identifiers so that I can follow an employee's employment status over time. For each PI-R01-renewal combination, I identify employees paid by the PI *every month* in the 10-12 months before R01 expiry (i.e. the first three months of the panel). This is a heuristic to identify personnel who are more likely to be long-term members of the PI's lab or who were not already scheduled to end their tenure with the focal PI. I then create a monthly panel following their employment status from 9 months before expiry to 12 months after expiry.

[10] Some transactions in UMETRICS are negative amounts. These can appear for a number of reasons including returns, discounts, reversing a purchase that was wrongly assigned, or money that was unused and refunded. I discuss this more in the Online Appendix.

I construct two outcome variables. The first is whether the employee is paid by the focal PI through any of the PI's grants in a given month. This can be thought of as a proxy for whether employee-PI matches are disrupted. The unit of interest is an employee-PI combination, and the data structure is an employee-PI-R01-renewal monthly panel.

The second outcome measure is whether the employee is paid by *any grants* from any PI in any given month. Even if an employee is separated from their usual PI they may be shifted to a different project, so this captures the overall "employment status" of the employee. The unit of interest here is an employee, so the data structure is an employee-R01-renewal monthly panel.

Both of these outcomes are non-absorbing. That is, the indicator can be on-then-off or off-then-on in consecutive months and do not necessarily indicate that an employee as "exited" employment, which we do not observe.

3.2 ExPorter

ExPorter is publicly available data provided by the NIH.[11] It contains data on NIH-funded projects, including identifiers that IRIS has used to link projects to their transcations in UMETRICS. It also provides links to publications, patents, and clinical studies that cite support from the NIH.

3.2.1 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

3.2.2 Grant portfolio ("Number of R01s")

The effect of an interruption on a PI may vary by the size of their grant portfolio. I count the NIH grants that a PI was on between the dates of one year before and after the focal R01 expired. I define the size of the PI's grant portfolio based on the number of "R01-equivalent"

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

or "P01" grants that the PI had, including the focal R01. I follow the NIH definition of R01-equivalent grants.[12] P01 grants provide funding for multiple research projects with a common theme.[13] For brevity, I will refer to this variable as the "Number of R01s" without explicitly defining the other types of grants included.

Since employees are not necessarily in charge of their own grants, I need to define funding support for employees. When the unit of interest is an employee-PI combination, this is the grant portfolio of the focal PI, as defined above. When the unit of interest is an employee, I identify all PIs that the employee was paid by at any time during the 10-12 months before R01 expiry. I then count the number of R01s (including R01-equivalents and P01s) that those PIs were in charge of during the 24-month period used in the analysis.

3.2.3 Lab/PI publications

I use publication counts as a proxy for research output. ExPorter provides a crosswalk between NIH projects and publications in PubMed, a database for publications in biomedical research and the life sciences. These can then be aggregated up to the PI or lab-level. I also weigh publications by 2-year forward citations i.e. citations up to two years from publication year, including publication year.

3.3 Author-ity & Web of Science

Author-ity is a dataset of disambiguated author names based on a snapshot of MEDLINE, a bibliographic database with a focus on "biomedicine and health," in 2009.[14] Each observation in Author-ity is a cluster of names and articles that are predicted to belong to the same author. Clarivate Analytics Web of Science (WoS) is a citation indexing database that is not publicly available. I use a version of WoS that indexes articles and citations up

[12] "R01-equivalent grants are defined as activity codes DP1, DP2, DP5, R01, R23, R29, R37, R56, RF1, RL1, U01 and R35 from select NIGMS and NHGRI program announcements (PAs) Not all of these activities may be in use by NIH every year." (Link)

[13] A P01 can roughly be thought of as a combination of R01 projects.

[14] Author names cannot be used directly to identify individuals because multiple versions of the same author's name may appear in the literature (e.g. "Adam Smith" and "A. Smith") and multiple authors may have the same name.

to and including 2013.

4 Empirical Strategy

I compare how outcomes change for labs/PIs that were interrupted versus those that were not. I begin by finding instances where an R01 was successfully renewed within the fiscal year it expires. I then stack all combinations of renewed R01s and PIs of those R01s to create a balanced PI-R01-renewal monthly panel spanning 24 months—one year before and one year after the expiry. I define an R01 as (a) "interrupted" if took more than 30 calendar days to be renewed or (b) "uninterrupted" or "continuous" if it took fewer than 30 calendar days to be renewed. I then estimate event study specifications that allow us to see how labs respond to interruptions month-to-month.

4.1 Event-study: Research inputs

The main specification I estimate is an event study centered around the expiry month of a project period.

$$y_{LRt} = \sum_{e=-10}^{12} \beta_m * \mathbf{1}(e = t - t_{expiry})\mathbf{1}(Interrupted) + \delta_{LR} + \gamma_e + \epsilon_{LRt}$$

I index PIs as L, R01s as R, and the year-month as t. t_{expiry} is the year-month that the R01 grant R expires. e is the number of months before expiry i.e. $e = 0$ when R expires amd $e < 0$ before the grant expires. I restrict the sample to the one year before and after the R01 R expires, i.e. e starts at month -11 and ends at 12. $e = -11$ is excluded from the specification.

y_{LRt} is a variable at the PI-level, such as total spending across all of the PI's grants), δ_{iR} are PI-R01-renewal fixed effects, and γ_e are fixed effects for months relative to expiry. The coefficients of interest are β_t, $t = -10, -9, \ldots, 11, 12$.

4.1.1 Employee-level outcomes

Continuing with the same notation, I index employees as i. For looking at whether the employee-PI relationship was affected, I estimate the following regression specification with employee-PI-R01-renewal fixed effects and relative time fixed effects:

$$\mathbf{1}(PI - paid - employee)_{iLRt} = \beta_m * \mathbf{1}(e = t - t_{expiry})\mathbf{1}(Interrupted) + \delta_{iLR} + \gamma_e + \epsilon_{iLRt}$$

For looking at the effects on the employee, I estimate:

$$\mathbf{1}(Any - grant - paid - employee)_{iRt} = \beta_m * \mathbf{1}(e = t - t_{expiry})\mathbf{1}(Interrupted) + \delta_{iR} + \gamma_e + \epsilon_{iRt}$$

4.2 Event-study: Research outputs

I draw on the universe of NIH-sponsored scientists (through ExPorter) to estimate the effect of funding restrictions on research output. I use a similar "stacking" procedure as described above to construct a PI-R01-renewal by year panel that starts 4 years before an interruption and ends 5 years after, restricted to all interruptions that took place from 1989 to 2006, maximizing the time span covered without truncating any outcome variables. I then apply the Coarsened Exact Matching (CEM) procedure (Iacus, King, and Porro 2012) and estimate the following event-study specification and its "static" counterpart with the matching weights:

$$y_{iRt} = \sum_{k=-5}^{5} \beta_k * \mathbf{1}(k = t - t_{expiry})\mathbf{1}(Interrupted) + \delta_{iR} + \gamma_{tc} + \epsilon_{iRt}$$

$$y_{iRt} = \beta_{static} * \mathbf{1}(t - t_{expiry} \geq 1)\mathbf{1}(Interrupted) + \delta_{iR} + \gamma_{tc} + \epsilon_{iRt})$$

$$t - t_{expiry} = -4, -3, \ldots, 0, \ldots, 5$$

I index PIs with i, renewed R01s with R, and calendar year with t. c indexes treatment cohorts or calendar year R expired. Treatment cohorts are defined as the set of PIs that were renewed in the same calendar year, some of whom were interrupted and some not. y_{iRt} is a measure of research production (e.g. publications), δ_{iR} are PI-R01-renewal fixed effects, and γ_{tc} are treatment cohort by calendar year fixed effects. $t - t_{expiry}$ is years since interruption, 0 being the interruption year.

This is a variant of the two-way fixed effects specification but with time-by-treatment-cohort fixed effects instead of time fixed effects to address issues arising from staggered treatment timing in difference-in-differences (Goodman-Bacon 2018; Callaway and Sant'Anna 2018; Abraham and Sun 2018). The Online Appendix discusses these in more detail.

4.3 Inverse hyperbolic sine

Unless otherwise stated or the outcome is a binary variable, I apply an inverse hyperbolic sine (also *asinh* or *arcsinh*) transformation to the outcome variable for all regressions, which approximates a natural logarithm and is defined at zero.[15]

5 Descriptive statistics

The analysis within UMETRICS uses a sample of 362 PI-R01-renewals with one R01 (288 uninterrupted, 74 interrupted) and 687 PI-R01-renewals with multiple R01s (524 uninterrupted, 163 interrupted). The timing of renewals ranges from 2003 to 2017.

Figure 2 compares characteristics of the expiring (and eventually renewed) R01 project periods from the UMETRICS sample used in the upcoming analysis. Figures 2A and 2B show the distribution of funding gaps. Overall, a little over 20% of R01s were "interrupted" or renewed more than 30 days after their expiry date. The funding gap experienced by interrupted R01s has a wide range with a maximum of over 300 calendar days.

[15]For "large" outcomes i.e. spending amounts, I convert estimates to percentage changes using the standard $exp(\hat{\beta}) - 1$ for log transformations. More details in Online Appendix.

Figures 2C and 2D show that interrupted and continuous R01s are similarly funded, whether in terms of total funding over the entire project period or funding per year, with interrupted projects being slightly bigger. Interrupted projects are more likely to have had a six-year project period (Figure 2E).

R01s are awarded for a maximum of five years, thus six-year projects periods are likely to have been six-year awards that exercised a one-year no-cost extension and more likely to be interrupted because they no longer have that option. This might affect the results if spending trends differ by project length. I repeat the analysis on lab spending with exact matching on project period lengths and obtain similar results (available in Online Appendix).

In the first month of the panel, the median team for PIs with one R01 had 4 employees in total, 1 faculty, and 1 research employee, and 0 for the remaining occupations.[16] The median team for PIs with multiple R01s had 8 employees, 2 faculty, 1 research staff, 1 postgraduate researcher, and 0 for the remaining occupations. The most common occupations on the average team are faculty, research staff, graduate students, postgraduate researchers, and research facilitators, after which other occupations are much less likely to be on a team.[17] Median expenditure at the beginning of the panel for PIs with one R01 was $17,100 for total direct expenditure, $14,200 for labor payments, $900 for vendor payments, and 0 for subaward payments. For PIs with multiple R01s, the same statistics were $40,500 for total direct expenditure, $31,200 for labor payments, $4,100 for vendor payments, and $0 for subaward payments.

[16]The full set of UMETRICS occupation codes is: Faculty, Research, Graduate Student, Postgraduate, Research Facilitation, Undergraduate, Technical Support, Clinical, Instructional, Other, and Other Staff.

[17]Online Appendix contains more summary statistics about lab size.

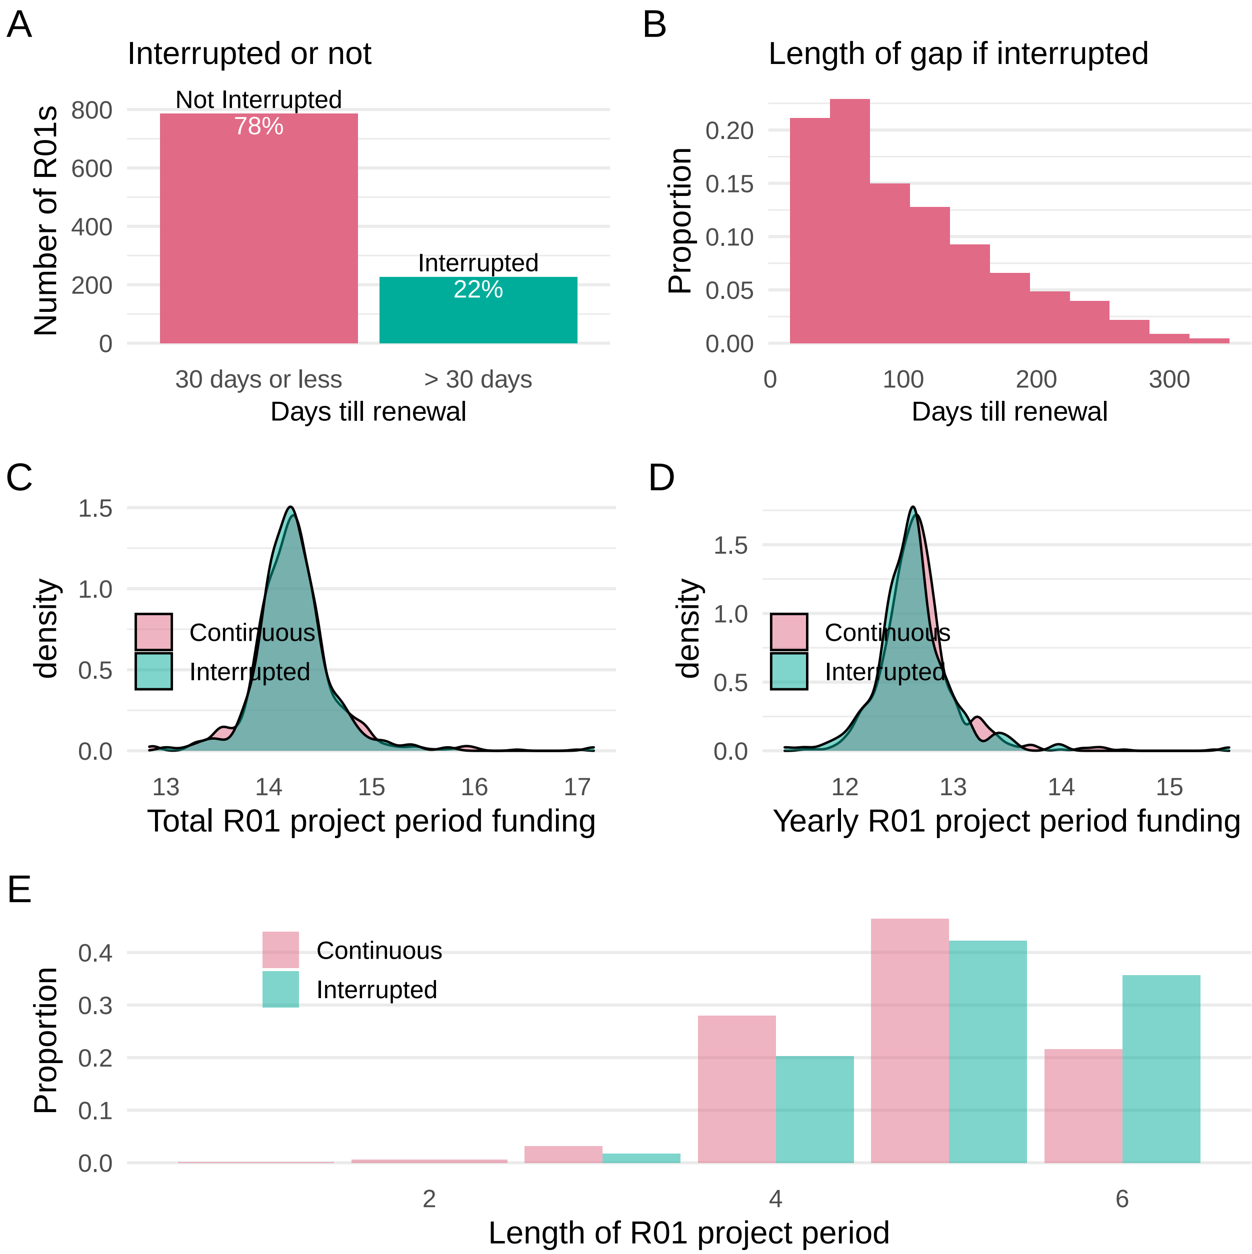

Figure 2: This figure compares the characteristics of expiring (and eventually renewed) R01 project periods that are used in the UMETRICS analysis. Each unit of observation is an R01 project period. Figure 1A shows the number of R01s that were renewed within 30 calendar days of their expiry. Figure 1B is a histogram (30-day bins) for the number of days till renewal for R01s not renewed within 30 days. Figure 1C shows the smoothed density for total funding in the expiring project period for interrupted and uninterrupted R01s. Figure 1D shows the same figure for funding per year (total funding / length of expiring R01 project period). Figure 1E shows the proportion of projects by the expiring project period's length.

6 Results

6.1 Spending

Figure 3A shows the event study estimates for total spending by PIs, with separate estimates by whether the PI had one R01 (green) or multiple R01s (brown) around the time of expiry. The "1 R01" graph (green) shows that for PIs with one R01, total spending starts decreasing about three months before the official expiry date. At the lowest point, spending is 94% lower for PIs with interrupted R01s.

A priori, we might expect not to observe any differences between the spending of interrupted and uninterrupted PIs if the funds for a project are restricted to being spent within the originally designated time period and there is therefore no reason for PIs to adjust the path of their spending.

One way this condition can be violated is through "pre-award spending."[18] This means that a renewed project may be allowed to access funds from its next budget up to 90 days (or approximately three months) before the budget's offical start date. Thus, continuously funded PIs may be spending more when their R01 expires because they use funds from their future budget to smooth spending.

Another non-exclusive reason for the divergence in spending is if interrupted projects are allowed to spend beyond the official end of their budget period. The event study estimates are consistent with a situation where PIs are allowed to spend after the official expiry date and are making use of that to smooth spending or wait for more certainty about funding before they commit to long-term expenditures (e.g. hiring a postdoc).

Figures 3B and 3C show the event study estimates for labor payments and vendor payments respectively. For labor spending, the change in spending patterns look broadly similar to those for overall spending. Vendor payments decrease less than labor payments in *asinh* points but still substantially by percentage, with a 91% decrease at the lowest point. Vendor payments also do not recover as quickly as labor payments.

[18] https://www.niaid.nih.gov/grants-contracts/caveats-consider-preaward-spending

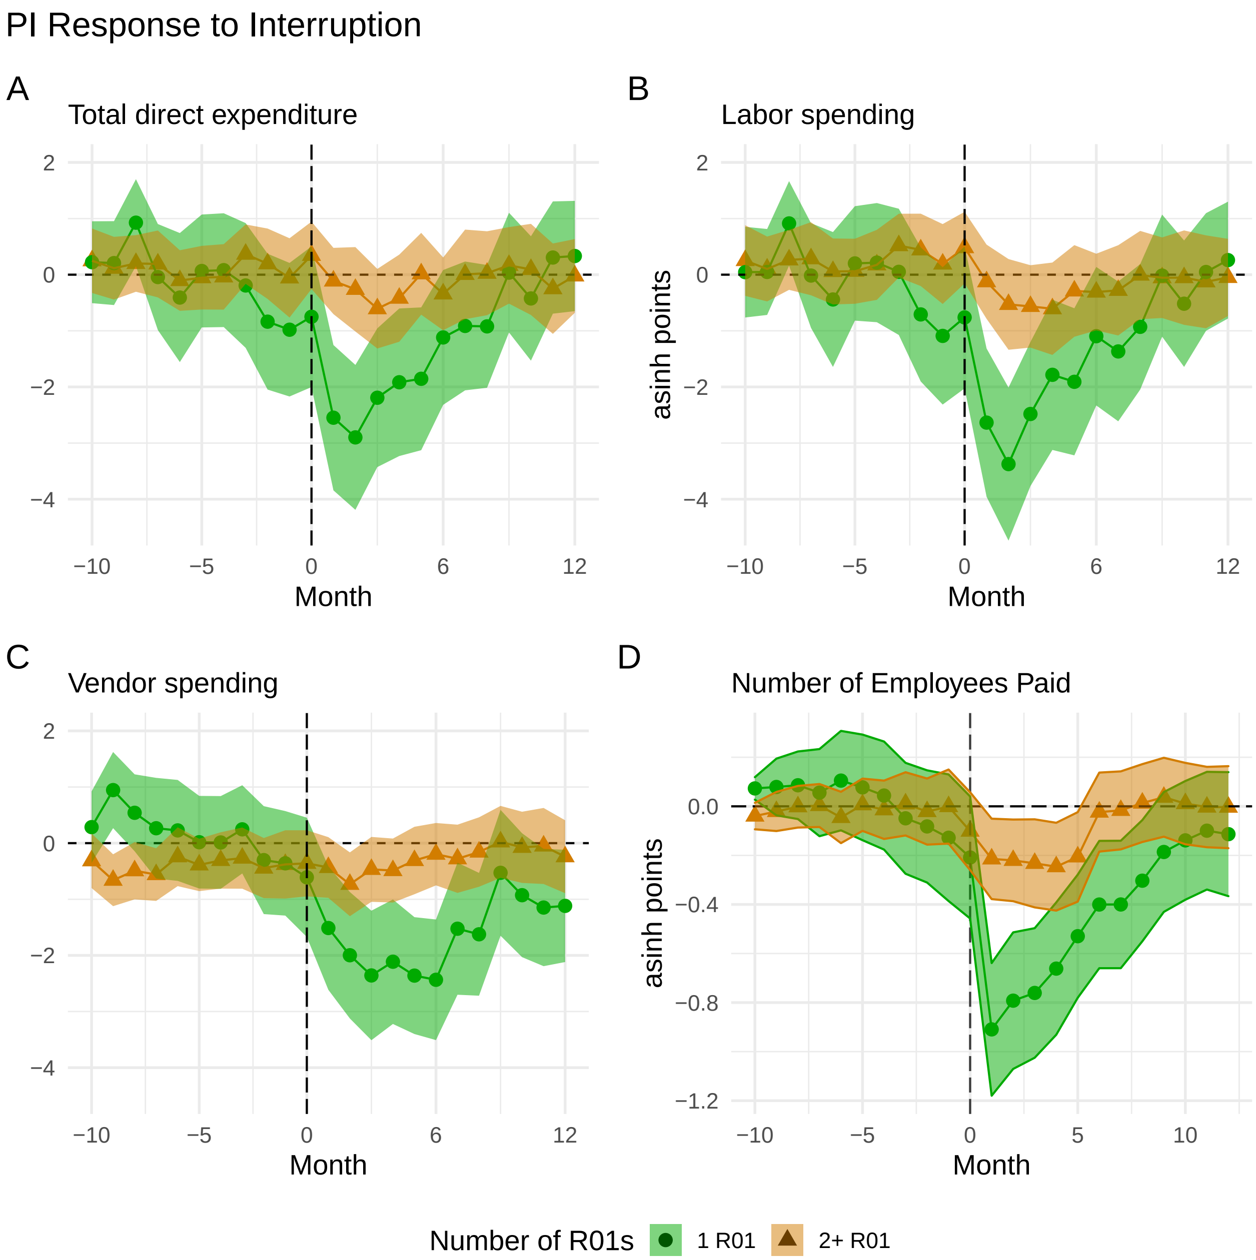

Figure 3: This figure shows event-study estimates (with 95% confidence intervals) of the difference in spending between PIs of interrupted and uninterrupted R01s. Each panel shows the estimates for a different outcome variable: total expenditure by PI (A), total vendor expenditure by PI (B), total labor expenditure by PI (C), and total number of employees paid by PI (D). Month 0 is the month that the focal R01 expires. Month -11 is the excluded category for the regression. Regressions are run separately on subsamples of PIs that have one R01 grant (green) or multiple R01s (brown), including R01-equivalents and P01 grants. Standard errors are clustered at the PI level.

Finally, Figure 3D shows the event study estimates with counts of employees as the outcome variable. These results are consistent with what we see for labor payments. In the month after the expiry of the project period, interrupted PIs with one R01 pay about 64% fewer personnel. In the same month, interrupted PIs with multiple R01s pay about 19% fewer employees.[19]

6.2 Employees

In addition to their effects on research production, interruptions may have disruptive effects on employees. One concern is that it may force employee turnover in a lab. For instance, a staff scientist may have to switch between PIs in order to maintain their salary or employment, or a postdoctoral researcher may be forced to leave if renewal funding does not become available quickly enough to fund their position. In addition to the personal disruption to employees, there may be a loss of team-specific capital (Jaravel, Petkova, and Bell 2018).

To get at these issues, I focus on the following outcome variables: whether an employee was (a) paid by the same PI on an NIH grant or (b) paid on any grants at all. I subset the data by employee occupation and number of NIH grants an employee is associated with (as detailed in the *Employee-level outcomes* section).[20]

Figure 4 shows the event-study estimates by occupation subsample. Across all occupations, interrupted employees associated with one R01 are less likely to be paid by the same PI or by any grant, while for those associated with multiple R01s, the probability of being paid decreases less or not at all. Over time, interrupted and uninterrupted employees converge in their probability of being paid. However, for the "postgraduate", "research facilitation", and "research" occupations, employees with one R01 remain 13, 10, and 7 percentage points less likely to be paid on any grant a year after R01 expiry. This raises the question of whether employees in those occupations are paid less or even leave their

[19]Event study estimates for different interruption lengths are available in the Online Appendix.

[20]Note that the "Number of R01s" variable is defined differently for each outcome variable, thus the subsamples used in the analysis are not identical.

Employee status when interrupted

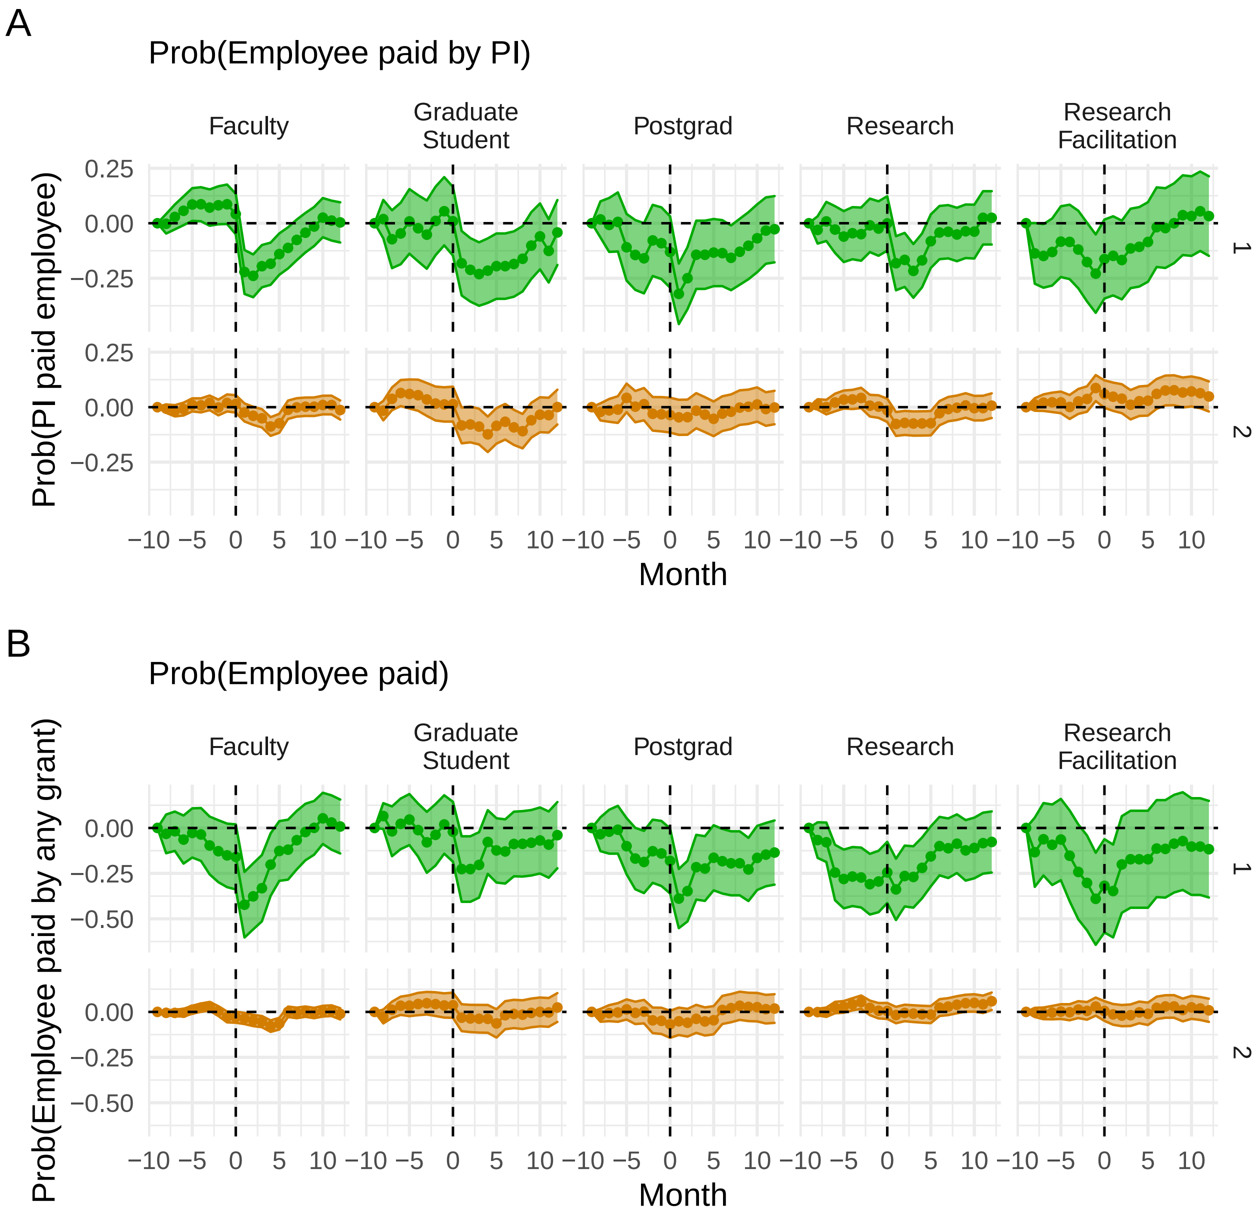

Figure 4: This figure shows event study coefficients with 95% confidence intervals (clustered at employee-level) of the difference in probability of being paid for employees on an interrupted project relative to those on an uninterrupted project. The same event study is estimated on subsamples by occupation and number of R01s. Panel A is for the outcome variable of whether an employee is paid by the PI of the renewed R01. Panel B is for the outcome variable of whether an employee is paid on any grants.

institution, because it may be harder to find internal sources of funding for them.[21]

6.3 Research output

In my final set of results, I estimate the effect of interruptions on publications at the yearly level. Table 1 shows the estimates from the "static" specification. None of the estimates are statistically different from zero. This is consistent with the event study figures (in Online Appendix), which do not show obvious differential trends or levels in publications before and after an interruption.

One reason for the imprecise estimates may be that interruptions are simply not particularly disruptive in practice. For instance, labs may be able to mitigate the temporary halt in spending by devoting more time to aspects of research that do not immediately need money (e.g. thinking of new ideas, writing).

Another reason is there may be mechanisms in place to cushion the effects of a funding interruption. For example, universities may provide "bridge funding" to researchers while they are waiting for their grant to be renewed. Expenditures using university-provided funds are not captured in the UMETRICS data so we cannot test this directly. In this case, the "true" effect of interruptions on research output is negative and even if bridge funding mitigates their effect, resources have to be diverted from elsewhere to do so.

Finally, standard measures of productivity such as publications may be too coarse relative to the true effects of funding interruptions. In addition, the time lag between the conception of an idea and publication varies. Even if funding interruptions have a meaningful effect on research activity, the effect of an interruption may be "smeared" across several years and therefore hard to detect.

[21] Although this cannot be answered definitively using UMETRICS alone, one suggestive piece of evidence is that for interrupted faculty, the unconditional probability (i.e. not relative to uninterrupted employees) of being paid by a grant increases almost monotonically with time from expiry, whereas interrupted employees in all other occupations hit a plateau even after an initial increase. This may reflect that even though the outcomes pf interrupted and uninterrupted employees converge, for faculty (many of whom would be a PI on the R01) this is due to a "recovery" while for other occupations it is due to a "catching up" of uninterrupted employees to their interrupted counterparts.

Table 1: Effect of interruptions on research production

	No. of Pubs		Cite-weighted Pubs	
	1 R01 (1a)	2+ R01 (1b)	1 R01 (2a)	2+ R01 (2b)
Interrupted-by-Post	0.002	-0.016	0.014	-0.026
	(0.016)	(0.015)	(0.028)	(0.025)
Num.Obs.	111960	112640	111960	112640
R2 Adj.	0.768	0.799	0.674	0.706

This table shows the 'static' difference-in-difference estimates and 95% confidence intervals of the difference in publication output if a PI had an interrupted R01. The regression includes treatment-cohort-by-year and PI-R01-renewal fixed effects and uses weights from coarsened exact matching on age, gender, and pre-interruption publications (raw counts and citation-weighted). Dependent variables are raw publication counts and citation-weighted (2-year forward citations) publications, both arcsinh-transformed. Standard errors clustered at PI level. Event study plots are available in the Online Appendix.

7 Conclusion

I study how NIH-funded researchers respond to funding interruptions. Using transaction-level data, I am able to examine these effects at a level of granularity that was previously unavailable. I find evidence that interruptions are disruptive to research. PIs spend less either because of the uncertainty about whether or when they will be funded again, or because they are not able to draw on funds from their next budget, or both. This may in turn be disruptive to the work and training of employees, who become less likely to be paid on grants. In ongoing work, we investigate whether this affects a wider range of employment outcomes such as earnings or even having to leave their institution.

These results point to two important policy implications. First, policies to reduce uncertainty can help us to avoid the costs of disruptive events such as funding interruptions. An example would be a multi-year appropriation for the budgets of research agencies. Second, given that some amount of uncertainty is unavoidable, how organizations choose to react to uncertainty is an important policy lever that is also more realistically adjusted. This paper shows that the risk aversion of organizations comes with costs that should factor into decision-making.

Bibliography

Abraham, Sarah, and Liyang Sun. 2018. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Available at SSRN 3158747*.

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). Wiley Online Library: 527–54.

Bae, Jung, Reza Sattari, and Bruce A. Weinberg. 2020. "The Marginal Scientific Product of Investments in Science." The Ohio State University.

Callaway, Brantly, and Pedro HC Sant'Anna. 2018. "Difference-in-Differences with Multiple Time Periods and an Application on the Minimum Wage and Employment." *arXiv Preprint arXiv:1803.09015*.

Conti, Annamaria, and Christopher C Liu. 2015. "Bringing the Lab Back in: Personnel Composition and Scientific Output at the Mit Department of Biology." *Research Policy* 44 (9). Elsevier: 1633–44.

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

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

Iacus, Stefano M, Gary King, and Giuseppe Porro. 2012. "Causal Inference Without Balance Checking: Coarsened Exact Matching." *Political Analysis*. JSTOR, 1–24.

Jaravel, Xavier, Neviana Petkova, and Alex Bell. 2018. "Team-Specific Capital and Innovation." *American Economic Review* 108 (4-5): 1034–73.

Lane, Julia I, Jason Owen-Smith, Rebecca F Rosen, and Bruce A Weinberg. 2015. "New Linked Data on Research Investments: Scientific Workforce, Productivity, and Public Value." *Research Policy* 44 (9): 1659–71.

Latour, Bruno, and Steve Woolgar. 2013. *Laboratory Life: The Construction of Scientific Facts*. Princeton University Press.

Liebman, Jeffrey B, and Neale Mahoney. 2017. "Do Expiring Budgets Lead to Wasteful Year-End Spending? Evidence from Federal Procurement." *American Economic Review* 107 (11): 3510–49.

Mervis, Jeffrey, and Eliot Marshall. 1996. "Science Budget: When Federal Science Stopped." *Science* 271 (5246). American Association for the Advancement of Science: 136–36.