

The Impact of Research Grants on Publications and Patents Across Disciplines*

Jérôme Adda [†] Pedro Oliboni[‡] Marco Ottaviani[§]

September 6, 2023

Abstract

This paper investigates the effect of research funding on researcher productivity in terms of both the quantity and quality of publications and patents across disciplines. We employ unique data on grant applications with evaluator grade information, matched with data on the career of researchers, before and after application, including publications and patenting records. We show that research grants substantially increase publication outcomes by exploiting a fuzzy regression discontinuity design. Regarding patenting, we find a significant effect for the field of technology, where research outcomes are typically eligible to patent protection. The analysis also reveals heterogeneous effects on the quantity and quality of publications by field, age, and type of grant program, but not by gender.

*We gratefully acknowledge research support by the NBER-Sloan Foundation Science of Science Funding program, the European Patent Office (EPO) Academic Research Programme, and the European Research Council through Grant 101055295 InfoEcoScience (Information Economics for Science). We thank without implication Christian Bjerke, Kristin Oxley, and Eystein Strømme for sharing their knowledge about the RCN data; James Atkins, Mohammed Badaoui, Francesco Bilotta, Alberto Frison, Saeid Kazemi, Marta Mojoli, Federico Pessina, and Biao Yang for excellent research assistance; as well as Yann Ménière, Neus Palomera, and Bart Van Looy for helpful discussion and comments.

[†]Bocconi University, jerome.adda@unibocconi.it.

[‡]Global Priorities Institute (University of Oxford), pedro.oliboni@economics.ox.ac.uk.

[§]Bocconi University, marco.ottaviani@unibocconi.it.

Contents

1	Introduction	1
2	Data	3
2.1	Data Sources	3
2.2	Data Construction	5
3	Econometric Methodology	7
4	The Effect of Funding on Research Productivity	9
4.1	Graphical Evidence	9
4.2	Regression Analysis	14
4.3	Fuzzy Regression Discontinuity Design	16
4.4	Effects by Field, Program Type, Age, and Gender	19
5	Discussion of Contribution Within the Literature	24
6	Conclusion	27
	Appendices	29
A	OLS Restricted by Grade	29
A.1	Heterogeneity by Field	31
B	Robustness	33
B.1	Mean outcome by pre-treatment covariates	33
B.2	Fixed application cohort: 2014	33
C	Institutional Background	36

1 Introduction

Governments worldwide invest substantially in research grants. Financial support to researchers is widely believed to play a crucial role in advancing science and innovation. But how effective are grants-based funding schemes? This paper leverages unique data from the Research Council of Norway (RCN) to investigate the causal effect of research grants on individual researchers' productivity, focusing on the impact on the quantity and quality of publications and patents across different disciplines.

Estimating the causal effect of research grants poses a significant challenge due to the tendency for the most productive scientists to be the most successful in securing grants (Jaffe, 2002). We overcome this challenge using a Fuzzy Regression Discontinuity (FRD) design. This is possible because our RCN data includes detailed information on peer review evaluations. Our FRD exploits the fact that for different grant programs, there are different cut-offs in terms of evaluation grades such that projects just above the cut-off are much more likely to get funded.

To study the effect of grants on both quantity and quality of publications and patents we combine this RCN data with three other data sets containing information on the careers of researchers, patenting, and citations. Overall, our data has two key advantages. First, since the RCN is a large, multidisciplinary funder, we can investigate how the causal impact of research grants varies across different disciplines. Second, our estimates are plausibly not significantly threatened by the issue of unobserved funding as the RCN is the main funder for all scientific disciplines in Norway—such that researchers in Norway have relatively few alternative sources of funding.

Overall, we find that research grants have substantial effects on both the quantity and quality of publications. Concerning quantity, we find that a grant adds 5.14 articles (control group mean of 9.07). Concerning quality, a grant adds 11.88 publications weighted by journal influence score (control group mean of 17.86), and 303.3 citation-weighted publications (control group mean 179.16). Separating by field, there are substantial effects on both quantity

and quality of publications in medicine and health science, mathematics and natural sciences, and technology. While for humanities, social sciences, and agricultural and fisheries subjects we only see an effect on the quantity of publications.

Concerning patents, we find that funding does not have an impact on patents in general. This is not surprising as most disciplines do not have much patenting activity. However, for the field of technology grants do have a substantial, positive effect on the quality of patents as measured by patent-to-patent citations.

Moreover, the effectiveness of research grants on quantity and quality of scientific output varies substantially by type of grant, seniority, but not by gender. Grant programs at the RCN that are open to all topic areas and primarily focus on excellence of research generally have a higher effect on publications than other types of programs. In particular, we see a large effect on quality, as the effect on journal influence score-weighted publications is more than doubled for these open-ended programs as compared to other programs. Concerning the seniority of researchers, we find that having less than 10 years of scientific experience is associated with a smaller effect on both quantity and quality of publications.

We contribute to a growing literature investigating the causal impact of research grants on scientific and innovative output. Jacob and Lefgren (2011) were the first to employ a Fuzzy Regression Discontinuity design to estimate the causal impact of research grants. They exploit the fact that the realized score-based cutoff for funding at the NIH varies. Their key finding suggests a moderate impact of grants: only about one additional publication over the next five years. However, they highlight that for several reasons, importantly potential unobserved external alternative funding, their findings underestimate the actual effect of funding.

Since Jacob and Lefgren (2011), other authors have employed the regression discontinuity approach to estimate the causal impact of research grant programs in different national contexts. Benavente et al. (2012) study the Chilean National Science and Technology Research Fund and find that funding affects the quantity but not quality of publications. Lanser et al.

(2013) investigate the effectiveness of research grants by a relatively narrow grant program by the Dutch Technology Foundation focused on application-oriented scientific research with a budget of 80 million Euros a year. They find no effects for a grant program that is open to all disciplines, but find that thematic (or directed) calls have a significant effect on publications. Ganguli (2017) applies the RD design to estimate the causal effect of a historic emergency grant program that funded over 28,000 Soviet scientists following the end of the USSR. She finds that grants more than double publications and induce scientists to remain in science. Onishi and Owan (2020) investigate research grants to the field of economics in Japan. In general, they find increases in the count of papers and citations as well as heterogeneous effects such as differences between junior researchers with and without tenure.

We expand on this prior work in three key ways. First, we investigate how the impacts of grants varies across scientific fields. Second, we estimate effects on both the quantity and quality of publications and patents produced by grantees. Third, our estimates are plausibly less threatened by the issue of unobserved funding due to the Research Council of Norway being the main funder for all scientific disciplines in Norway.

The paper proceeds as follows. Section 2 describes the data. Section 3 explains our econometric methodology. Section 4 presents the results. Section 5 frames our contribution within other related literature. Section 6 concludes.

2 Data

2.1 Data Sources

We use three data sources. Our main data source consists of grant applications to the Research Council of Norway (RCN), the agency operating on the behalf of the Norwegian government to fund research and innovation projects across all fields. RCN, with a yearly budget of 11.9 billion Norwegian Krone (NOK) in 2021, was established in 1993 as a result of the merging of five research councils covering different research areas. Following common

practice at research funding organizations, RCN evaluates applications within peer review panels. We have access to data including the evaluation grades given to applications by each individual evaluator inside a panel, the agreed grade for each application, as well as whether the application was funded. As explained in Appendix C in more detail, the agreed grade of an application is the grade assigned to the application after the evaluators in the panel discuss their assessment. The agreed grade is often close to, but need not coincide with, the average of the grades given by each individual evaluator. Final grant decisions are made by portfolio boards, who are not composed by panel peer reviewers, and is largely based on the ranking of applications inside each panel. However, some applications are funded out-of-order.

We use data from applications received after 2011, when RCN moved to a new internal database, until 2018. Unfortunately, we lack complete information about the panel under which each application was evaluated. Yet, as we explain below, our main econometric method (the Fuzzy Regression Discontinuity) relies on identifying the funding cutoff in terms of the grade associated to each application - which is in fact determined via the panel each application is evaluated under at the RCN. Thus we employ an algorithm (described in more detail below) in order to approximate the relevant cutoff for each application.

Our two other data sources allow us to link grant applicants to outcome measures in terms of publications and patents. For publications and researcher characteristics, we use the Cristin database, the Norwegian national repository of research publications. This database covers all researchers working in Norway and includes various types of academic output such as articles, books, dissertations, and academic lectures. To match applicants to publications we rely on a unique person identifier that relates the Principal Investigator of an application to the author of a publication.

For patents, we use the Orbis database, from which we obtain patent counts and the forward citations associated to each patent. The matching process is more difficult in this case as we have to rely on the full name of researchers. We match researchers and patenting outcomes based on a score of similarity, but we only include matches that have an almost

Table 1: Descriptive Statistics: Sample Size

	Total
Number of researchers	6,665
Number of funded researchers	3,051
Number of applications (with grade)	13,474
Number of publications	757,157
of which articles	303,130
Number of patents	2,229
Cite-weighted patents (modified)	5,109
Cite-weighted pubs (modified)	9,952,014

Notes: Data from RCN applications and from Cristin repository.

perfect score. Moreover, we only match individuals who have a unique name in our researchers data. There is no reason to believe that the quality of the matching is different across fields.

Table 1 presents information on the totals for different variables in our analysis sample once all of our datasets have been merged and we have kept only researchers who apply to the RCN at least once (and for whom we have information about evaluation grades).

2.2 Data Construction

Our final sample is constituted by records of researcher-year, alongside information about any applications made that year as well as publishing and patenting outcomes.¹

Innovation can be defined and measured in several different ways and as discussed in the literature section, it is important to consider both direct and indirect measures. Therefore, we use multiple outcomes variables. We use four different measures related to academic publications on a per researcher-year basis: simple publication counts, simple article counts, journal-impact-weighted publication counts, and cite-weighted publication counts.²

¹In the case a researcher applies for several grants in the same year, we define their (effective) application outcome as successful if and only if they were successful in at least one such application. For such researchers, we define our three grade measures as follows: restricting attention to the applications which achieved the researcher’s effective outcome, the overall grade is the maximum out of these applications, and the mean and normalized grade are defined as the averages of these respective grades out of the relevant applications. Moreover, for researchers with 4 or more applications in a year, we exclude that researcher-year from consideration.

²We preform a minor adjustment for the weighted measures: we add one to each patent or publication. That is, if a patent or publication received 2 citations, then we take it to 3, and similarly for journal impact scores for publications. The reason for this is that otherwise patents (and similarly publications) which

To weight the influence of journal articles we use the Article Influence Score (AIS), which counts recursively for the average article published in a given journal the average number of citations obtained per article over the first five years after publication, relative to the citations given (Clarivate, n.d.). We also use two measures related to patenting: simple patent counts and cite-weighted patent counts (in which we use the forward citations received by the focal patent from other patents). Moreover, for each measure, we will define the outcome in terms of an aggregate of a period after a given grant application. In particular, we consider a period of 2 to 4 years after application. For instance, the outcome in terms of article counts for person i who obtains a grant in year $t = 2005$ is the count of articles i publishes in years 2007 – 2009. The simple idea here is that we expect it to take a few years for research output to “show up” after a grant.

We construct fields and broad fields based on the fields assigned to grant applications at the RCN. For fields, we classify each researcher in one of 7 RCN fields according to the mode of the fields assigned to their grant applications. These 7 fields are: humanities, agricultural and fisheries subjects, mathematics and natural sciences, medicine and health sciences, social science, technology, and other. Then, we group these 7 fields into 3 domains which broadly correspond to the domains used by the European Research Council: (1) social sciences and humanities; (2) physical sciences and engineering, including technology, mathematics, and natural sciences; (3) life sciences, including medicine, health sciences, as well as agricultural and fisheries subjects.

In our Fuzzy RD method we use a normalized application grade. This measure is constructed from the mean of the grades given by individual reviewers to each application. Normalization is done by subtracting from the mean grade the relevant cutoff above which an application is significantly more likely to be accepted for funding than below it. Since we do not have complete access to the panel each application is associated to, we use information about programs, which sometimes are broader and other times narrower than panels,

receive no citations would count the same as zero patents (or zero publications).

but which are good enough approximations such that we are still able to identify cutoffs associated to a discrete jump in the probability of funding. Our rule for finding such cutoffs is simple: for each program in each year, we set the cutoff as the 15th percentile mean grade out of the funded applications. When there are no funded applications in a program-year, we set the cutoff as the 85th percentile mean grade out of the unfunded applications. We use this percentile-based rule in order to avoid setting the cutoff based on applications which were funded out-of-order.

3 Econometric Methodology

We are interested in how receiving a grant in year t affects a researcher's innovative output in terms of both patents and academic publications in some period $(t+k, T)$ after the grant is received. With varying methods, we will build on the following econometric model to obtain an estimate of the parameter τ associated with the indicator of treatment assignment (funding) w , and which thus represents the impact of funding on innovative output,

$$y_{i,t^+} = \beta_0 + \tau w_{i,t} + g(x_{i,t}) + X_{i,t}B + \epsilon_{i,t^+} \quad (1)$$

where t^+ denotes an aggregation of periods from $t+k$ to T , $x_{i,t}$ is a grade for researcher i in year t , $g(x_{i,t})$ is a function of $x_{i,t}$ which changes depending on our method, $X_{i,t}$ includes our researcher-year covariates, and ϵ_{i,t^+} is a mean zero residual. By letting $g(x_{i,t}) = \beta_1(x_{i,t} - c) + \delta_1[x_{i,t} \geq c] \cdot (x_{i,t} - c)$ (a more flexible version of $g(x) = \beta x$), which allows for different linear trends for the grade below and above the threshold c , we obtain the following model

$$y_{i,t^+} = \beta_0 + \tau w_{i,t} + \beta_1(x_i - c) + \delta_1[x_i \geq c] \cdot (x_i - c) + X_{i,t}B + \epsilon_{i,t^+} \quad (2)$$

Application grades aim to capture the quality of projects. If we believe that such grades, along with other measures of applicant quality we have access to (such as prior publications),

account for all of the factors that influence both funding assignment and innovative output, then this method would give unbiased estimates of the effect of funding. However, the coefficient of interest obtained in such a way may be biased if the decision to fund an application is based on a broader information set, such as scientific reputation, or if some amount of affirmative action takes place. Our Fuzzy Regression Discontinuity (FRD) design aims to solve the potential issue of selection on unobservables relative to OLS estimation. The method exploits the fact that, for a given panel, funding is typically done on the basis of a ranking of applications such that there is a discrete jump in the probability of funding at a particular cutoff c of a measure of the application grade. While at the RCN such rankings largely reflect the agreed grade, they also take other factors into account such as demographics and more nuanced grades for different criteria of relevance. Moreover, the agreed grades are always integers, which would be an issue for a Fuzzy RD strategy. Fortunately, however, we also have access to the mean grade of each application. This is plausibly a better approximation of the information used in creating the rankings given that these depend on more information than that contained in the agreed grade and the mean grade aggregates the opinion of multiple evaluators in a more fine-grained way since it may take non-integer values. We thus select the mean grade as our grade measure for the FRD method.

Furthermore, a well known result by Hahn et al. (2001) establishes that the fuzzy regression discontinuity estimator is identical to a local IV estimator of τ in equation 2, where the instrument for $w_{i,t}$ is $z_{i,t} \equiv 1[x_{i,t} \geq c]$ and we restrict data such that $x_{i,t}$ is in a small neighborhood of c . The underlying assumption is that in a neighborhood of c , conditional on observables X_{it} , funding is random (Wooldridge, 2010).

Finally, we also employ a different method made possible by the fact that RCN ranks applications using, among other factors, the agreed grade, which is always an integer number (see Appendix A). This model aims to solve the issue of selection on unobservables by using the agreed grade as opposed to the the normalized mean grade and making a different assumption. We assume here that, save for selection on observables, once we restrict attention

to a particular agreed grade (e.g. 6), applications are selected into funding in a process that is as good as random.

4 The Effect of Funding on Research Productivity

4.1 Graphical Evidence

We begin with preliminary descriptive information about our analysis sample. As presented in Table 1, our sample includes 6665 different researchers, all of whom apply at least once for RCN funding (and for whom we have at least one grade), and 3051 of whom obtain RCN funding at least once.

In Table 2, we display the table of means for our outcome variables and covariates depending on the status of each application. Applications that get funded tend to have higher agreed grades as well as higher normalized grades than those who do not. We also see evidence that those who get funded tend to be less experienced in research (scientific age) and tend to have published more (cumulative publications). Moreover, the means of all outcome measures are greater for applications which are funded as compared to those which are not. However, the difference in means is only significant for cite-weighted publications, AIS-weighted publications, count of articles, and count of general publications.

Figure 1 displays the frequencies with which applications have each possible cutoff based on the mean grade given by individual reviewers (recall that cutoffs are computed at the level of program-year groups). We see that a cutoff of 6 is the most common and that the vast majority of cutoffs is greater than or equal to 5.

Graphical analysis In line with recommendations in Imbens and Lemieux (2008) and Lee and Lemieux (2010), we begin with a series graphs to assess the validity of our main identification strategy based on the Fuzzy Regression Discontinuity design. We focus on three kinds of graphs: (1) probability of assignment to treatment by forcing variable, (2)

Table 2: Descriptive Statistics by Funding Status

	Not funded	Funded	Diff	p-value
<i>Applicant/application characteristics</i>				
Normalized grade	-0.616	0.484	1.100	0.0000
Agreed grade	4.694	6.117	1.423	0.0000
Gender (female=1)	0.381	0.389	0.008	0.4558
Scientific age	13.568	13.169	-0.398	0.0282
Cumulative publications	88.440	100.715	12.275	0.0000
Cumulative patents	0.344	0.320	-0.025	0.6641
<i>Outcome measures</i>				
Patents years 2-4	0.042	0.044	0.002	0.8222
Cite-weighted pats years 2-4	0.068	0.114	0.046	0.1961
Cite-weighted pubs years 2-4	179.161	296.701	117.540	0.0000
AIS-weighted pubs years 2-4	17.859	26.693	8.834	0.0000
Articles years 2-4	9.069	12.116	3.047	0.0000
Pubs years 2-4	20.227	29.374	9.147	0.0000

Notes: Data from RCN applications and from Crislin repository.

mean outcome by the forcing variable, and (3) baseline covariates by forcing variable (see Appendix B). The point of (1) is to verify whether there is indeed a jump in the probability of treatment at the cutoff c . The point of (2) is that in general, in order to obtain a robust and credible estimate with a statistically significant magnitude of the effect, a jump in the conditional mean of the outcome variable given the forcing variable should be visible (Imbens and Lemieux, 2008). Finally, the point of (3) is that a credible analysis requires that there be no discontinuities at the cutoff in variables that are determined prior to selection into funding (Lee and Lemieux, 2010).

In Figure 2, we can see a clear jump in the probability of funding at the cutoff, for the normalized mean grade. This provides evidence in favor of the crucial assumption behind our Fuzzy RD design: at the cutoff there is a discontinuity in the probability of assignment to treatment.




Figure 1: Frequencies of cutoffs by application.




Figure 2: Probability of funding by normalized mean grade

In Figure 3 we present plots of our outcome measures in terms of the normalized agreed grade. Each of our outcome measures here aggregates post-application outcomes from year $t+2$ to year $t+4$, where t is the application year. There does not seem to be a clear pattern in terms of patenting outcomes. While for cite-weighted patents we do see a jump at the cutoff when comparing the two regression lines, this seems largely driven by a few observations. For our publication measures, an interesting pattern emerges: the slope of mean outcomes in terms of normalized grade is essentially flat up until the normalized grade of 0, beyond which it generally starts to increase. Moreover, we can see that for each of our publication measures there is an identifiable jump in the mean outcome functions at the cutoff. Indeed these plots of patenting and publishing outcomes foreshadow that we generally find positive effects of funding for publishing outcomes while we do not generally find an effect in patenting. It is worth noting, however, that as our sample includes mostly researchers focused on academic research, we have relatively few patents. It is thus not surprising that we do not find effects on patenting in general. For this reason, we later also investigate heterogeneous effects in terms of scientific fields - as we might expect researchers working on technology or medicine, for example, to produce patents.




Figure 3: Plots of mean outcome variables against normalized mean grade

Notes: The fits takes into account number of observations within each grade. These plots do not include all of our data points as we are explicitly focusing on on a window of ± 2 around the cutoff.

We now report our estimates for each of the three general methods described in our methods. Throughout, standard errors are clustered by researcher, and our field fixed effects

are based on RCN broad fields, unless otherwise specified. Recall that our outcome variables are all counts of the respective measure (e.g. cite-weighted patents) for the years 2-4 after the application year. That is, if i receives a grant in 2005, the outcome measure counts the number of patents i received between 2007 – 2009.

4.2 Regression Analysis

Table 3 presents our OLS estimates. Each row reports the coefficient of funding for each outcome variable. Our baseline specification includes only the indicator of funding. We then progressively include more controls.

We do not find an effect for patenting, both for simple counts and cite-weighted patents, at any level of controls. Moreover, we find a significant association between every measure of publications and funding. For publication counts, we see that, out of our levels of control, the minimum estimate has that funding is associated with an increase in 5.5 publications (27% of the control group mean, significant at the 1% level). For AIS-weighted publications, our minimum estimate is of 5.02 additional AIS-weighted publications, as compared to a control group mean of 17.86 (28%, significant at the 1% level). For articles, our minimum estimate is of 1.47 (significant at the 1% level), which is 16% of the control group mean. Finally, for cite-weighted publications we see considerable variation in the level of significance across specifications, with the coefficients ranging from 72.37 to 117.5 (for a control group mean of 179.16).

These OLS estimates suggest that grants by the RCN increase various publication outcomes. However, we observe considerable variation in our coefficients as we include or not application cohort fixed effects - an issue that we will investigate further within our FRD method. Moreover, although we are able to account for various measures which are correlated with output, such as prior publications and the application mean grade, there might still be concerns about selection on unobservables (Jacob and Lefgren, 2011). Our Fuzzy Regression Discontinuity design, which we turn to next, addresses this issue.

Table 3: OLS estimates of the effect of funding

	(1)	(2)	(3)	(4)	(5)	(6)
	Patents	Cite-wei pats	Pubs	Cite-wei pubs	AIS-wei pubs	Articles
(0): no controls	0.00199 (0.0112)	0.0463 (0.0429)	9.147*** (0.860)	117.5*** (34.09)	8.834*** (1.650)	3.047*** (0.488)
(1): grade controls	-0.00546 (0.00997)	0.0333 (0.0440)	7.226*** (0.893)	83.13** (38.50)	5.387*** (1.704)	1.823*** (0.485)
(2): (1) and gender and sci age	-0.00387 (0.00989)	0.0346 (0.0434)	7.453*** (0.888)	85.58** (38.40)	5.596*** (1.719)	1.925*** (0.487)
(3): (2) and RCN broadfield FE	-0.000748 (0.00995)	0.0394 (0.0439)	7.572*** (0.886)	100.1*** (37.99)	6.665*** (1.716)	2.266*** (0.482)
(4): (3) and prior pub and pats	0.0000524 (0.00892)	0.0404 (0.0431)	6.853*** (0.791)	91.43** (37.84)	5.919*** (1.617)	1.925*** (0.442)
(5): (4) and app cohort FE	-0.00158 (0.00904)	0.0339 (0.0406)	5.500*** (0.726)	72.37* (38.06)	5.019*** (1.594)	1.468*** (0.428)
Control group mean	0.04	0.07	20.23	179.16	17.86	9.07
Control group SD	0.35	1.21	25.53	854.59	30.04	13.18
Observations	11573	11573	11573	11573	11573	11573

Standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; Num obs is 11576 for (0) and (1)

4.3 Fuzzy Regression Discontinuity Design

Baseline Table 4 presents our Fuzzy RD estimates. As in the baseline OLS case, our baseline FRD specification includes only the indicator of funding. We then progressively include more controls.³ In all specifications, column (1) reports the first stage results. This column shows that our instrument of receiving a normalized mean grade of 0 or more is associated with an increase in 22% the probability that an application receives funding (for specification (4), for example). This estimate is significant at the 1% level and for all specifications the F-statistic on the excluded instrument is of at least 360, which suggests no weak instrument bias.

In the remaining columns, as a general pattern our second stage estimates suggest that funding might have a substantial impact in publication outcomes, while we find no statistically significant association with patenting outcomes. In our specification (0), with no controls, we find a positive association between funding and publishing which is significant at the 1% level for all publishing outcome measures. However, as we include more controls up until specification (2), we lose significance for cite-weighted publications, AIS-weighted publications, and simple article counts. Interestingly, as we further add more controls in specification (3) and then (4)—our preferred specification—our point estimates increase and we again find statistically significant relation for all publication measures (at at least the 10% level). However, once we include application cohort fixed effects, in specification (5), all of our point estimates for publication outcomes decrease and we see no statistically significant association between funding and any patenting or publishing outcome.

Overall, save for our specification that includes application cohort fixed effects, our results suggest a positive and substantial effect of funding on publishing, although not patenting, outcomes. In particular, for specification (4) our estimates of the impact of funding on publication outcomes in terms of percentages relative to the control group means are the following: 76% increase in publications, 169% increase in cite-weighted publications, 67%

³In all of our Fuzzy RD estimates, we employ an estimation band of 1.5 on both sides of the cutoff.

increase in AIS-weighted publications, and 57% increase in articles.

These findings are importantly not in line with what has been previously found in this literature. In particular, in the work closest to our setting, the results of Jacob and Lefgren (2011) suggest that NIH research grants at most moderately affect publishing outcomes of marginal applicants. For instance, the authors' OLS estimate suggests the impact of a grant on the number of publications is of about 8% of the control group mean (with their RD setting suggesting an even smaller effect). Yet, even when we include application cohort fixed effects and thus obtain smaller coefficients, our estimated effects are larger than what has been previously found. For instance, our OLS estimates including all controls suggest a 16% increase in article counts.

As discussed in the work of Jacob and Lefgren (2011), there could be various reasons for why even under a true substantial effect, estimates would be small or not significant. One such reason concerns the fact that researchers often have outside options to funding from the NIH in the form of other funding agencies, their own academic institutions, or even coauthors Jacob and Lefgren (2011). Researchers who apply to funding by the RCN may also have outside options. And as discussed in our methods and in Jacob and Lefgren (2011), this should lead us to interpret well-identified Fuzzy RD estimates as lower bounds on the effect of funding. Thus, so far our results suggest a relatively higher lower bound for the effect of funding than what has been previously found. Here it is important to highlight that the RCN plausibly plays a larger role in research funding in Norway as compared to the role the NIH plays in the US. This could explain why our estimated effects are in general larger.⁴

⁴As we mentioned, we find no significant effects with our FRD design when we control for application cohort fixed effects. In Appendix B.2 we further explore the issue regarding application cohort fixed effects.

Table 4: Fuzzy Regression Discontinuity estimates of the effect of funding

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Funded	Patents	Cite-weigh pats	Pubs	Cite-weighted pubs	AIS-weigh pubs	Articles
(0): no controls	0.376*** (0.00722)	0.0112 (0.0206)	0.162 (0.123)	13.53*** (1.504)	213.5*** (60.35)	16.19*** (2.513)	5.441*** (0.870)
(1): grade controls	0.227*** (0.0104)	-0.0412 (0.0495)	0.277 (0.302)	13.76*** (4.074)	234.3 (176.8)	6.490 (5.516)	3.227 (2.218)
(2): (1) and gender and sci age	0.228*** (0.0104)	-0.0413 (0.0494)	0.279 (0.303)	13.63*** (4.063)	234.1 (177.0)	6.496 (5.492)	3.215 (2.212)
(3): (2) and RCN broadfield FE	0.226*** (0.0103)	-0.0286 (0.0500)	0.304 (0.310)	14.36*** (4.076)	290.8 (177.2)	10.85** (5.521)	4.668** (2.211)
(4): (3) and prior pub and pats	0.226*** (0.0103)	-0.0263 (0.0467)	0.308 (0.309)	15.35*** (3.702)	303.3* (174.1)	11.88** (5.210)	5.135** (2.015)
(5): (4) and app cohort FE	0.224*** (0.0104)	-0.0408 (0.0473)	0.252 (0.291)	3.827 (3.415)	132.9 (181.5)	4.109 (5.001)	1.234 (1.929)
Control group mean	-	0.04	0.07	20.23	179.16	17.86	9.07
Control group SD	-	0.35	1.21	25.53	854.59	30.04	13.18
Observations	12076	10274	10274	10274	10274	10274	10274

Standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Num observations for first stage of (0) and (1) is 12080, for second stage of (0) and (1) it is 10277

4.4 Effects by Field, Program Type, Age, and Gender

Within our Fuzzy RD design, we examine heterogeneity in the effect of funding alongside four dimensions: scientific field, program type, scientific age, and gender. We adopt specification (4) and interact the indicator of funding, $w_{i,t}$, with the relevant categorical or indicator variable (e.g. gender indicator). We uncover several interesting patterns which warrant further study.

Scientific field In Table 5 we report estimates of the heterogeneity of effects by RCN field. For this estimation, we modify specification (4) by having finer grained fixed effects at the RCN field level as opposed to RCN broad-field level and we add a control for FRI vs. not FRI as the prevalence of different fields varies in the two broad program types. Recall that we have seven RCN fields, which are (in the order that they appear in Table 5): other, humanities, agricultural and fisheries subjects, mathematics and natural sciences, medicine and health sciences, social science, and technology. First, we note that for the “other” field, there are only 24 observations.

Concerning patents, while in general we do not find an effect, here we find an effect of funding for the field of technology: an increase in 0.5 cite-weighted patents (significant at the 5% level), which is substantial given that the mean for unfunded applications in the field of technology is 0.18.

Concerning publications, we find substantial effects on both quantity and quality in the fields of medicine and health science, mathematics and natural sciences, and technology. While for humanities, social sciences, and agricultural and fisheries subjects we only see an effect on the quantity of publications.

Table 5: Fuzzy RD estimates of effect of funding, heterogeneity by field

	(1)	(2)	(3)	(4)	(5)	(6)
	Patents	Cite-weighted pats	Pubs	Cite-weighted pubs	AIS-weighted pubs	Articles
Fund*other	-0.0370 (0.0410)	0.126 (0.183)	-6.386 (20.15)	-58.65 (189.1)	-5.511 (8.593)	6.151 (5.153)
Fund*Hum	-0.0993 (0.0733)	0.0120 (0.276)	15.91** (6.246)	162.9 (191.0)	-1.565 (6.172)	3.026 (2.826)
Fund*Ag&Fish	-0.0320 (0.0507)	0.131 (0.182)	13.99 (9.165)	208.1 (143.8)	3.927 (5.070)	4.493** (2.241)
Fund*Mat&NatSci	-0.0407 (0.0595)	0.217 (0.233)	19.62*** (4.786)	292.8* (172.7)	13.19** (6.222)	6.902** (2.904)
Fund*Med	-0.0209 (0.0763)	0.684 (0.737)	14.18*** (4.890)	672.5** (336.1)	32.31*** (10.80)	8.531*** (3.255)
Fund*Soc	-0.0514 (0.0348)	0.100 (0.166)	10.03*** (3.423)	59.81 (131.9)	-1.274 (3.666)	1.182 (1.619)
Fund*Tech	0.0950 (0.0855)	0.497** (0.242)	18.18*** (4.383)	186.0 (119.8)	7.227* (4.011)	4.290** (2.025)
Observations	10274	10274	10274	10274	10274	10274

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Fuzzy RD estimates of effect of funding, heterogeneity by program type

	(1)	(2)	(3)	(4)	(5)	(6)
	Patents	Cit-we pats	Pubs	Cit-we pubs	AIS-we pubs	Articles
(a) Fund*Not FRI	-0.0158 (0.0434)	0.246 (0.209)	14.37*** (3.566)	258.4 (168.8)	8.783* (5.159)	4.282** (1.887)
(b) Fund*FRI	-0.0418 (0.0624)	0.462 (0.538)	15.81*** (4.738)	394.6* (216.2)	18.62*** (6.682)	6.612** (2.754)
(b) - (a)	-0.0260 (0.0381)	0.217 (0.340)	1.434 (3.041)	136.2 (127.4)	9.835** (4.883)	2.330 (1.857)
Observations	10274	10274	10274	10274	10274	10274

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Program type The success of a research funding program depends on how it is designed. Should grant-makers be open to applications about any topic or problem? Or should they implement targeted programs, aiming to solve particular scientific or practical problems?

We contribute to this issue of program design by estimating how the effect of funding varies in terms of two broad types of programs at the RCN. While the RCN has various programs enacting calls for proposals, we classify the program calls into two large categories: those part of the broader FRIPRO category, and those which are not. The FRIPRO category represents open calls for ground-breaking research. It is thus characterized by a heightened focus on excellence as well as openness to any topic areas, basic and applied. This is contrasted with other programs which target specific scientific areas or problems and which may have other aims in addition to excellence in research.

Table 6 shows our IV estimates for the heterogeneous effects of funding depending on program type (FRIPRO vs. not FRIPRO). We see evidence that funding via FRI programs has a higher effect on publishing outcomes. In particular, all of the point estimates for the FRI group are higher than those for the non FRI group, and for AIS-weighted publications, we find that the difference in effect size between the two groups is of 9.84 (significant at the 5% level).

Scientific age Another important issue for designing scientific grant-making programs concerns the level of experience of researchers. It is plausible that the effects of funding a more experienced versus a less experienced researcher would change—and this could be for various reasons. It may be that funding less experienced researchers is generally more risky, as the grant evaluators have less of a track record to base their decisions on. But also, it might be that younger researchers more often provide novel perspectives, and thus are better placed at making groundbreaking discoveries. Obviously, many other reasons exist for experience to modulate the impact of funding.

In order to get at experience, we determine each researchers *scientific age* as the number of years since their first publication (which may have been, for instance, their doctoral dissertation). We then define a researcher as less experienced (“young”) if their scientific age is strictly less than 10, and those with scientific age greater than or equal to 10 are classified as “old”.

Table 7 presents our IV estimates for how the effects of funding depend on whether a researcher is more or less experienced (old vs. young). The overall pattern that emerges is that the effect of a grant, for publication output, is larger for more experienced researchers. Indeed, consider how a grant tends to add 6.03 articles (significant at the 1% level) to “old” researchers, but 3.52 for “young” researchers (with the difference in effects significant at the 10% level).

This pattern should be interpreted carefully, however. We must consider that the RCN has different types of funding (for instance, FRI vs. not FRI), and these plausibly vary systematically depending on scientific age. For example, funding for training is plausibly more common among younger researchers. Thus, the smaller effect for younger researchers might be explained by an only temporary drop in productivity, perhaps laying the ground for longer-term publications. A more in depth analysis about age would thus benefit from outcome measures of longer duration.

Table 7: Fuzzy RD estimates of effect of funding, heterogeneity by scientific age

	(1)	(2)	(3)	(4)	(5)	(6)
	Patents	Cit-we pats	Pubs	Cit-we pubs	AIS-we pubs	Articles
(a) Fund*Old	-0.0403 (0.0454)	0.152 (0.189)	17.92*** (3.926)	321.6 (199.4)	13.87** (5.763)	6.028*** (2.146)
(b) Fund*Young	-0.00111 (0.0556)	0.591 (0.533)	10.70*** (3.807)	270.1* (145.2)	8.266* (4.840)	3.518* (2.015)
(b) - (a)	0.0391 (0.0330)	0.439 (0.353)	-7.212*** (2.436)	-51.49 (110.2)	-5.608* (3.354)	-2.510** (1.229)
Observations	10274	10274	10274	10274	10274	10274

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Gender We also study heterogeneity of effects of funding in terms of the gender of applicants. Table 8 reports on our findings. We find mixed evidence on how the effect of funding varies by gender. On the one hand, the estimated effects for publications and cite-weighted publications are larger for women. On the other hand, the effects on AIS-weighted publications and articles are larger for men. In any case, however, we do not obtain statistically significant estimates of the difference between the two groups.

Table 8: Fuzzy RD estimates of effect of funding, heterogeneity by gender

	(1)	(2)	(3)	(4)	(5)	(6)
	Patents	Cit-we pats	Pubs	Cit-we pubs	AIS-we pubs	Articles
(a) Fund*Men	-0.0269 (0.0562)	0.369 (0.392)	14.84*** (3.789)	272.3 (176.3)	13.75** (5.847)	5.369** (2.230)
(b) Fund*Woman	-0.0255 (0.0407)	0.223 (0.200)	16.07*** (4.170)	347.0* (198.1)	9.243* (5.279)	4.806** (2.015)
(b) - (a)	0.00139 (0.0382)	-0.146 (0.213)	1.234 (2.803)	74.73 (130.3)	-4.506 (4.268)	-0.563 (1.481)
Observations	10274	10274	10274	10274	10274	10274

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5 Discussion of Contribution Within the Literature

There is a growing literature investigating the connection between R&D investments and innovation. This broadly concerns estimating causal effects of different push and pull incentives, such as grants and patent protection, for investment in research. In this section we consider studies beyond those employing an RD design to estimate the causal effect of academic research grants and highlight how the former relate to our contribution.

Causal impact of industry R&D grants

Two studies have investigated the causal impact of grants for industry research. Howell (2017) uses a Fuzzy Regression Discontinuity design to study the effectiveness of R&D grants from the Small Business Innovation Research (SBIR) program of the US Department of Energy. She finds that early stage grants increase a firm's citation-weighted patents by 30%. In a similar vein, Santoleri et al. (2020) use a Regression Discontinuity design to study the effectiveness of the SME Instrument, an European R&D grant program. They find that early stage grants (for business concept development) have no effect, while later stage grants (for product development) increase citation-weighted patents between 15 and 31%. Similar to these studies, we also investigate outcomes in terms of patents and citation-weighted patents. However, we largely consider academic grants, as opposed to grants to industry, and measure effects on publication outcomes as well.

Indirect effects of science funding

Azoulay et al. (2019) aim to broadly quantify the casual impact of publicly funded research and do so by including indirect impacts, which are clearly important given the collective and cumulative nature of innovation. They find that 10 million dollars granted through the NIH lead to 2.3 additional patents in the private sector (Azoulay et al., 2019). They also find strong evidence of serendipitous spillovers: half of patents linked to NIH funding are for disease applications distinct from the one that funded the research. Their findings are

relevant to our own approach as they suggest direct measures such as the ones we use will tend to underestimate the impact of grants. A key difference of their study is that they look at the broad impact of an increase in funding to a particular area, while estimate how grants impact individual researchers.

While we obtain evidence on the direct effectiveness of grants, Li et al. (2017) provides descriptive evidence on the direct and indirect links between public research investments and patenting. In their sample of 365,380 grants from the 1980 – 2007 period, they find that 8.4% of the grants are directly acknowledged by patents (about 17 thousand total patents) but about 31% of the grants produce research that is cited by patents (about 81 thousand total patents). Another study that concerns descriptive evidence on indirect effects of public-sector investments in science funding is Galkina Cleary et al. (2018). They find that each of the 210 new molecular entities approved by the FDA in the 2010-2016 period was associated with research that received NIH funding.

Other incentives and behavior of researchers

While we focus on the causal effect of grants, the study by Adda and Ottaviani (2023) has implications for how the effectiveness of grant programs depends on the rules for allocating funding between fields. They study the strategic behavior of grant applicants when evaluations are noisy. Their key result is that as evaluation becomes noisier, the quantity of applicants increases. If more funds are distributed to topics with more applicants (as in a proportional allocation rule) then topics with more noise in evaluations, and which thus receive more applicants, will end up receiving more funds—and in the end attracting even more applicants.

Another paper that complements our investigation of the effectiveness of research funding is Banal Estañol et al. (2019), which establishes various correlations in the process of funding academic research. They find that researchers who apply more often to grants also tend to publish more. Moreover, they find that researchers with more publications and whose work

is less applied in nature are more likely to obtain a grant.

Like our study, Hvide and Jones (2018) investigate the Norwegian context. But rather than considering research grants, they study another important source of incentives: the rights of researchers to businesses or patents they may generate. They exploit the shift from a model of complete ownership by university researchers over their businesses and intellectual property to one closer to the US model, where 2/3 of such rights go to the university. They find that the reform led to a drop of about 50% in both the rate of start-ups created by university researchers and in patenting output.

Raiteri (2018) considers yet another incentive in research and innovation: public procurement. He finds that public procurement can stimulate innovation complementarities that increase the probability of diffusion of a general purpose technologies among different sectors (thus making them more pervasive, or general).

Contribution to the literature

Our work builds on this emerging literature on the connection between R&D investments and innovation by expanding on the currently available estimates of the causal impact of research grant funding in various ways. As outlined in the introduction, we have three key contributions. First, we consider how the effectiveness of grants varies across scientific fields. Second, our estimates are plausibly less threatened by the issue of unobserved funding due to the RCN being the main funder for all scientific disciplines in Norway. Third, we consider the direct effects of grant funding in terms of both the quantity and quality of publications and patents by individual researchers.

6 Conclusion

In this paper we investigate the causal impact of research grants on the productivity of individual researchers as measured by changes in the quantity and quality of publications and patents across disciplines. In our preferred specification, we exploit non-linear variation in the probability of funding as a function of evaluation averages in a Fuzzy Regression Discontinuity design.

We find evidence of substantial overall effects of grants on the quantity and quality of publications. In terms of quantity, we find that a grant adds 5.14 articles (control group mean of 9.07). In terms of quality, a grant adds 11.88 publications weighted by journal influence score (control group mean of 17.86), and 303.3 citation-weighted publications (control group mean 179.16). This contrasts with the previous findings of Jacob and Lefgren (2011), who do not obtain statistically significant effects in a similar Fuzzy Regression Discontinuity on the impact of research grants by the NIH. This is not so surprising. Because there exist other large alternative funders in the US, such as the National Science Foundation, researchers rejected by the NIH can often still get funded. However, the RCN plays a larger role in Norway compared to the role of the NIH in the US—thus reducing unobserved funding in our data.

The multidisciplinary context of the RCN also allows us to investigate how the effectiveness of grants varies by field. We find substantial effects on both quantity and quality of publications in the fields of medicine and health science, mathematics and natural sciences, and technology. While for humanities, social sciences, and agricultural and fisheries subjects we only see an effect on the quantity of publications. Concerning patents, we find that while funding does not have an impact in general, for the field of technology grants have a substantial, positive effect on quality as measured by patent-to-patent citations.

We also consider how effectiveness varies in researcher and grant characteristics. The effect of grants appears to be smaller for less experienced researchers. We suggest this finding does not necessarily imply inexperienced researchers benefit less from grants. Rather, this

difference calls for a more in depth analysis considering a fine-grained differentiation of kinds of funding as well as the inclusion of longer term outcome measures. The impact of grants on publication outcomes varies widely by type of grant. In particular, the effect is much larger for open-ended grant programs, such as FRIPRO, that focus on excellence in research without targeting specific scientific areas or problems. Finally, we find the impact of grants to not differ by gender.

Finally, this analysis is still an initial step and has uncovered various questions which are important and tractable for future work. First, a more meaningful assessment of the impact of scientific funding on patenting will require indirect measures of patenting (e.g. the count of patent-to-research article citations) as well as more fine-grained sensitivity to differences in grant research programs. These additional, more indirect measures of innovation would also allow us to capture the important spillovers and unexpected discoveries that characterize much of scientific progress (Sampat, 2015; Azoulay et al., 2019). Second, for a better understanding of the impact of research grants on publication outcomes, we need a finer grained categorization of different types of grants (e.g. project-based vs. training). Third, for all types of outcomes, it is important to obtain longer-term measures to capture effects which may naturally be delayed (e.g. moonshot projects or training programs). Fourth, our analysis is insensitive to differences in the size of grants, which is clearly an important factor to be considered in future work.

Appendices

A OLS Restricted by Grade

The specific setting of the RCN allows us to obtain a plausible causal estimate of the effect of funding in addition to the one obtained via the Fuzzy RD design. In particular, rather than focusing on mean grades, which can take non-integer values, we use information from agreed grades, which are always integers. In this way, we are able to exploit the following assumption to obtain plausible causal identification: if two applications have the same agreed grade, and this agreed grade is high enough, then save for some selection on observables, assignment into treatment (funding) is as good as random. Unpacking, we need the grade to be high enough because in general, the RCN funds applications with at least a score of 5 (more commonly the bar is set at 6), out of 7. If we look at the funded applications with a lower score, say 3, it is less plausible that assignment into treatment is effectively random because applications which are selected into treatment while having a lower grade are often done so for special reasons such as affirmative action or explicit preference for certain topics. However, such special selection will bias our estimates as long as these special features are correlated with outcomes. Moreover, we know that there is selection on observables. This is because, for instance, the RCN has an explicit preference for applications submitted by women—given that the applications being compared have received the same grade.

In this specification, we will restrict attention to the grade of 6. This is the second highest possible grade and also the grade with the most balanced number of funded vs. unfunded applications. Table 9 reports the means of our key covariates and outcome variables separately by funding status for the sample restricted to applications which obtain an agreed grade of 6. Although we hold constant the agreed grade, we see that there are significant differences in means for the normalized mean grade. This is not surprising given how we determine normalized mean grades. Moreover, we also see a significant difference in the mean

Table 9: Table of means, agreed grade = 6

	Not funded	Funded	Diff	p-value
<i>Applicant/application characteristics</i>				
Normalized grade	0.105	0.423	0.318	0.0000
Overall grade	6.000	6.000	0.000	.
Gender (female=1)	0.342	0.396	0.054	0.0012
Sci age	13.099	13.163	0.064	0.8170
Cumulative publications	97.702	99.418	1.715	0.6599
Cumulative patents	0.449	0.307	-0.142	0.2544
<i>Outcome measures</i>				
Patents years 2-4	0.051	0.048	-0.003	0.8179
Cite-weighted pats years 2-4	0.127	0.168	0.041	0.6877
Pubs years 2-4	22.210	28.971	6.760	0.0000
Cite-weighted pubs years 2-4	228.772	287.074	58.302	0.1853
AIS-weighted pubs years 2-4	22.005	25.336	3.331	0.0478
Articles years 2-4	10.407	11.927	1.520	0.0179

of gender in line with RCN's reported preference for women conditional on the applicants being compared having received the same agreed grade.

Regarding, outcome measures, we see that the difference in means for our patent measures are not significant. Moreover, the means for all the publication outcome measures are greater in the funded group. However, only the difference in means of publication counts, AIS-weighted publication counts, and article counts are statistically significant.

In Table 10 we report our baseline OLS estimates of the effect of funding for the particular agreed grade of 6. This table includes estimates using the levels of controls of specifications (4) and (5) (from our FRD analysis). In line with our other methods, we do not find a significant effect of funding on patenting outcomes. However, we again find effects of funding on publications. In particular, in specification 5, we find that for those who receive an agreed grade of 6, funding adds 4.7 publications, 3.02 AIS-weighted publications, and 1.2 articles (at different levels of significance). Our point estimate for cite-weighted publications from specification 5 is of 46.20. However, this is not statistically significant. These results are more moderate than those suggested by specification 4 in our FRD setting (or the FRD restricted to the cohort of 2014). Nevertheless, they are still larger, especially relative to control group

Table 10: OLS estimates of the effect of funding, agreed grade = 6

	(1)	(2)	(3)	(4)	(5)	(6)
	Patents	Cit-we pats	Pubs	Cit-we pubs	AIS-we pubs	Articles
(4): grade, gender, sci age, field FE, prior pubs and pats	0.00896 (0.0108)	0.0712 (0.0471)	7.095*** (1.008)	76.50 (47.60)	4.881*** (1.699)	2.031*** (0.559)
(5): (4) and app cohort FE	0.00560 (0.0111)	0.0586 (0.0425)	4.740*** (0.927)	46.20 (50.63)	3.019* (1.646)	1.184** (0.542)
Observations	2977	2977	2977	2977	2977	2977

Levels of controls (4) and (5) are the same as those used in Table 4

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

means, than the previous findings of Jacob and Lefgren (2011). Results for specification (4) in this setting are also of higher magnitude and significant at generally lower levels.

A.1 Heterogeneity by Field

Finally, Table 11 reports estimates for heterogeneity of the effects of funding by RCN field. In this case, we omit results for the “other” field as we only have 3 observations. Relative to our field heterogeneity results from the FRD, we see a more robust effect for the humanities, for which we now obtain statistically significant associations for publications, AIS-weighted publications, and articles. This table also provides further evidence on there being substantial effects of funding for mathematics and natural sciences as we obtain large and relatively precise estimates, which are all significant at the 1% level, for each publication outcome in this RCN field. Further research would be important to understand for instance whether this is because this field benefits from less noise in the evaluation of projects, or whether there are more funding constraints.

Table 11: OLS estimates of effect of funding, agreed grade = 6 & heterogeneity by field

	(1)	(2)	(3)	(4)	(5)	(6)
	Patents	Cite-weigh pats	Pubs	Cite-weighted pubs	AIS-weigh pubs	Articles
Fund*Hum	0.00108 (0.00218)	0.00174 (0.0166)	15.56** (6.707)	278.6 (201.1)	4.323* (2.462)	3.297** (1.483)
Fund*Ag&Fish	-0.0411 (0.0307)	-0.0660 (0.0442)	10.99** (4.861)	44.94 (55.50)	1.442 (4.678)	2.355 (1.522)
Fund*Mat&NatSci	0.00845 (0.0159)	0.0385 (0.0397)	9.825*** (1.705)	123.3*** (31.98)	8.862*** (2.432)	3.603*** (0.972)
Fund*Med	0.00389 (0.0330)	0.273 (0.261)	3.984* (2.171)	75.67 (205.8)	7.166 (6.761)	1.485 (1.723)
Fund*Soc	-0.00318 (0.00251)	-0.00177 (0.0110)	2.151 (1.864)	19.63 (24.43)	1.980 (1.396)	0.891 (0.625)
Fund*Tech	0.0611 (0.0574)	0.162 (0.103)	2.067 (3.597)	34.95 (45.96)	-0.828 (3.147)	-1.100 (2.180)
Observations	2977	2977	2977	2977	2977	2977

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B Robustness

B.1 Mean outcome by pre-treatment covariates

Figure 4 shows plots of our baseline covariates. We fit a separate regression line to each side of the cutoff and, in line with what is needed for a valid analysis, observe that for all of our covariates, there is no apparent jump at the cutoff.




Figure 4: Plots of mean covariates against normalized mean grade. The fits takes into account number of observations within each grade.

B.2 Fixed application cohort: 2014

In light of the fact that including application cohort fixed effects in our FRD estimates resulted in no significant coefficients on the effect of funding, we here provide a robustness check by restricting our analysis to a specific cohort. Table 12 presents our FRD estimates

while restricting our data to applications made in the application cohort of 2014. As a result of this restriction we effectively lose 93% of our observations. It is then natural that the variance in our estimates increases. However, for the cohort of 2014 we nevertheless still obtain a significant effect of funding for some of our publication outcomes. For these results, we adopt specification (4) and we report estimates under seven different choices of bandwidth, where the value of h denotes how far the bandwidth of estimation extends to either side.

We see that when using a small bandwidth of 0.8 extending to each side of the cutoff, all of our estimates are significant at the 10% level, and attain magnitudes much higher than those observed when we include all application cohorts. For many selections of bandwidth we obtain no statistically significant association. However, we obtain estimates significant at the 5% level for cite-weighted publications which are of substantial magnitude. For example, for a bandwidth of 2 our coefficient is of 482.4 additional cite-weighted publications. These estimates restricted to the application cohort of 2014 thus provide further evidence suggesting a positive and substantial effect of funding on publication outcomes.

Table 12: Fuzzy Regression Discontinuity estimates of the effect of funding, application cohort = 2014

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Funded	Patents	Cite-weigh pats	Pubs	Cite-weighted pubs	AIS-weigh pubs	Articles
h = 0.8	0.281*** (0.0349)	-0.294 (0.373)	-0.256 (0.377)	42.65** (17.47)	638.3** (281.1)	24.03* (14.41)	12.16* (6.510)
h = 1	0.237*** (0.0337)	-0.129 (0.223)	-0.00846 (0.237)	12.75 (12.26)	505.8 (311.3)	6.077 (11.45)	4.139 (4.962)
h = 1.2	0.238*** (0.0334)	-0.0743 (0.221)	0.0584 (0.236)	15.99 (12.20)	471.9 (306.5)	4.222 (11.18)	3.981 (4.846)
h = 1.5	0.253*** (0.0307)	-0.0279 (0.153)	0.0979 (0.176)	15.94 (10.08)	765.5* (392.9)	7.655 (9.276)	5.091 (4.007)
h = 1.8	0.254*** (0.0303)	-0.0365 (0.148)	0.0744 (0.168)	21.46** (10.87)	700.3** (349.2)	9.554 (9.097)	5.775 (3.964)
h = 2	0.264*** (0.0290)	-0.0223 (0.121)	0.0532 (0.135)	18.14** (9.174)	482.4** (236.8)	10.90 (8.051)	5.227 (3.432)
h = 2.3	0.265*** (0.0289)	-0.0221 (0.119)	0.0508 (0.133)	17.85** (9.087)	469.5** (231.9)	10.40 (7.955)	5.053 (3.397)
Observations	721	721	721	721	721	721	721

Standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C Institutional Background

The RCN has several funding programs with separate calls for applications. Programs vary in terms of scope. While some programs focus on specific themes (such as food safety or environmentally friendly energy), other programs are designed to support excellent research without thematic restrictions (e.g. the program funding Ground Breaking Research, known as FRIPRO). Each year, applicants submit proposals by the relevant deadline for the call they are applying to. In particular, the FRIPRO program has a single yearly submission deadline; while some programs do not have calls every year; and in a few programs there are no submission deadlines (one example is support for research stays abroad).

For the purposes of reviewing applications, the RCN has a number of panel groups, separated in terms of research areas, each composed of varying numbers of panels. For example, in 2022 there are 29 panel groups and roughly 100 panels. While the panel groups mostly remain constant over time, the panels can change in response to the applications received.

The RCN first uses a machine learning algorithm to pre-sort proposals to panel groups according to the proposals' descriptions. Then, program officers make final assignments of applications to panel groups according to research topics. Each panel group is divided into multiple panels of 20 to 30 applications. Usually panels consist of four to seven panel members. To avoid potential conflicts of interest, panel members are normally not based in Norway. The RCN has a database of international reviewers and invites them to serve on panels based on the expertise needed for the applications received. Usually panel members do not serve for more than three consecutive years.

Panels evaluate applications in two rounds. In a first round, before meeting, an employee of the RCN, the case officer, assigns each application to two panel members (primary and secondary evaluator) who must submit written assessments for it. The selection of primary and secondary evaluator is made with the intent of assigning applications to the evaluators with most relevant expertise in the area of the application. Expertise is determined via a

process whereby panel members state their level of expertise in the topics of the applications being considered. All members of a panel are also expected to submit grades, but only primary or secondary reviewers are expected to submit the written assessments. Panel members cannot view the assessment of other panel members at this stage. Additional reviews may be sought from experts external to the panel if necessary to complement the expertise of the panel for a given application. These external evaluations may be accessed by panel members before they submit their individual evaluations.

In the second round, panel members meet and discuss each application while making use of the written assessments that have been submitted regarding it. During this discussion, the panel members determine an agreed grade and write an overall assessment for each application. The primary reviewer is generally responsible for leading the discussion about the applications she is responsible for.

In the next step, the RCN staff conducts an assessment of relevance to the topic of the call (not included in the case of FRIPRO) and filters out applications that do not meet the requirement of relevance to the priorities of the topic. The filtered out applications can nevertheless be considered for FRIPRO grants. Then, the RCN administration conducts a portfolio assessment, in which they prepare ranked lists of projects for later use by the portfolio boards. These ranked lists consider the agreed grades as well as other factors (e.g. topic priorities, priority to women in cases where grades coincide, and grades for specific criteria) and each ranked list differs by how much weight is given to each factor.

Then, each of a total of 16 portfolio boards receives the ranked lists for the panels under its respective domain. Each portfolio board then uses the rankings according to their own prioritization of different factors in order to make funding decisions. The final funding decision of the portfolio boards may not coincide with the rankings. At the end of the process, only the agreed grade given by panel members to an application is disclosed with the applicant, and not the individual pre-meeting grades. Applicants receive information about the portfolio board which made a decision on their case and the criteria used for this, which

panel they were assessed under, as well as the referees who were part of this panel.

References

- Adda, Jérôme and Marco Ottaviani**, “Grantmaking, Grading on a Curve, and the Paradox of Relative Evaluation in Nonmarkets,” 2023.
- Azoulay, Pierre, Joshua S Graff Zivin, Danielle Li, and Bhaven N Sampat**, “Public R&D Investments and Private-sector Patenting: Evidence From NIH Funding Rules,” *Review of Economic Studies*, 2019, *86* (1), 117–152.
- Benavente, José Miguel, Gustavo Crespi, Lucas Figal Garone, and Alessandro Maffioli**, “The impact of national research funds: A regression discontinuity approach to the Chilean FONDECYT,” *Research Policy*, 2012, *41* (8), 1461–1475.
- Clarivate**, “Article Influence Score,” <https://jcr.help.clarivate.com/Content/glossary-article-influence-score.htm>. Accessed: 2022-11-06.
- Cleary, Ekaterina Galkina, Jennifer M Beierlein, Navleen Surjit Khanuja, Laura M McNamee, and Fred D Ledley**, “Contribution of NIH Funding to New Drug Approvals 2010–2016,” *Proceedings of the National Academy of Sciences*, 2018, *115* (10), 2329–2334.
- Estañol, Albert Banal, Inés Macho-Stadler, and David Pérez-Castillo**, “Funding Academic Research: Grant Application, Partnership, Award, and Output,” 2019.
- Ganguli, Ina**, “Saving Soviet science: The impact of grants when government R&D funding disappears,” *American Economic Journal: Applied Economics*, 2017, *9* (2), 165–201.
- Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw**, “Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design,” *Econometrica*, 2001, *69* (1), 201–209.
- Howell, Sabrina T**, “Financing Innovation: Evidence from R&D Grants,” *American Economic Review*, 2017, *107* (4), 1136–64.

- Hvide, Hans K and Benjamin F Jones**, “University Innovation and The Professor’s Privilege,” *American Economic Review*, 2018, *108* (7), 1860–98.
- Imbens, Guido W and Thomas Lemieux**, “Regression Discontinuity Designs: A Guide to Practice,” *Journal of Econometrics*, 2008, *142* (2), 615–635.
- Jacob, Brian A and Lars Lefgren**, “The Impact of Research Grant Funding on Scientific Productivity,” *Journal of Public Economics*, 2011, *95* (9-10), 1168–1177.
- Jaffe, Adam B**, “Building programme evaluation into the design of public research-support programmes,” *Oxford Review of Economic Policy*, 2002, *18* (1), 22–34.
- Lanser, Debby, Ryanne Van Dalen et al.**, “The effects of research grants on scientific productivity and utilisation,” Technical Report, CPB Netherlands Bureau for Economic Policy Analysis 2013.
- Lee, David S and Thomas Lemieux**, “Regression Discontinuity Designs in Economics,” *Journal of economic literature*, 2010, *48* (2), 281–355.
- Li, Danielle, Pierre Azoulay, and Bhaven N Sampat**, “The Applied Value of Public Investments in Biomedical Research,” *Science*, 2017, *356* (6333), 78–81.
- Onishi, Koichiro and Hideo Owan**, “Heterogenous Impacts of National Research Grants on Academic Productivity,” Technical Report, RIETI Discussion Paper Series, 20-E-052 2020.
- Raiteri, Emilio**, “A Time to Nourish? Evaluating The Impact of Public Procurement on Technological Generality Through Patent Data,” *Research Policy*, 2018, *47* (5), 936–952.
- Sampat, Bhaven N**, “Serendipity,” *Available at SSRN 2545515*, 2015.
- Santoleri, Pietro, Andrea Mina, Alberto Di Minin, and Irene Martelli**, “The Causal Effects of R&D Grants: Evidence from a Regression Discontinuity,” *Available at SSRN 3637867*, 2020.

Wooldridge, Jeffrey M, *Econometric Analysis of Cross Section and Panel Data*, MIT press, 2010.