



Essays on the impact of star scientist-mobility universities, peers and stars

Title	Essays on the impact of star scientist-mobility universities, peers and stars
Author(s)	Sasidharan, Akhil
Publication Date	2024-09-25
Publisher	University of Galway

Essays on the Impact of Star Scientist Mobility- Universities, Peers and Stars

A Thesis Submitted in Application for the Degree of Doctor of Philosophy to the J.E. Cairnes School of
Business and Economics at the University of Galway, Ireland

Akhil Sasidharan

Supervised By

Prof. John McHale (University of Galway)

Dr. Jason Harold (University of Galway)

May, 2024



Table of Contents

Table of Contents	2
List of Tables	4
List of Figures	7
Declaration	10
Acknowledgements	11
Dedication	12
Abstract	14
Chapter 1 Introduction	16
1.1 Context and Background.....	16
1.2 Overview of Science policies : Ireland , Denmark and New Zealand	18
1.3 Understanding Star Spillover: Gaps in the Current Literature.....	21
1.4 Overview of the Data	24
1.5 Overview of the Thesis Contributions	28
1.6 Thesis Outputs.....	32
Chapter 2 Stars as Catalysts: An Event-Study Analysis of the Impact of Star-Scientist Recruitment on Local Research Performance in a Small Open Economy	33
2.1 Introduction	33
2.2 Econometric Methodology	38
2.3 Data	43
2.4 Results	48
2.5 Robustness.....	57
2.6 Conclusions	60
Chapter 3 Star Help and Knowledge Transfer: An Event Study Analysis of Star Interactions Observed from Acknowledgement Texts	63
3.1 Introduction	63
3.2 Related Literature and Hypotheses	67
3.3 Data and Methodology	73
3.4 Results	85
3.5 Robustness.....	95
3.6 Summary and Policy Implications	107

Chapter 4 Estimating the Impact of a Star Scientist Move on their own Research Productivity	111
4.1 Introduction	111
4.2 Related Literature and Framework.....	113
4.3 Data and Methodology	119
4.4 Results	123
4.5 Robustness.....	127
4.6 Summary and Policy Implications	134
Chapter 5 Star Arrival Effects on Scientists’ Productivity: Does the Nature and Extent of Relatedness to the Star Matter?	136
5.1 Introduction	136
5.2 Model and Econometric Specification	142
5.3 Data	149
5.4 Results	153
5.5 Robustness.....	159
5.6 Concluding Comments	164
Chapter 6 Conclusion	166
6.1 Summary and Main Findings	166
6.2 Drawing the Threads Together: The Importance and Design of Star-Recruitment Policies	170
6.3 Limitations and Future Research.....	172
Appendix A. Supplementary material for chapter 2	174
Appendix B. Supplementary material for chapter 3	191
Appendix C. Supplementary material for chapter 4	200
Appendix D. Supplementary material for chapter 5	205
Bibliography	213

List of Tables

Table 1.1 Counts of star scientists based on their origin countries from which the stars are moving from into the SOEs.....	26
Table 2.1 Descriptive statistics.....	45
Table 2.2 Summary statistics: control and treated subfields at baseline (k-to-k matched)	47
Table 2.3 Dynamic star arrival effects for FNTC (department level)	49
Table 2.4 Dynamic star arrival effects for FNTC (individual level).....	53
Table 3.1 Acknowledgement types that define the types of helpful interactions.....	77
Table 3.2 Summary statistics: control and treated group (k-to-k matched)	79
Table 3.3 Dynamic star help effects for FNTC (homogenous and heterogeneous effects).....	81
Table 3.4 Dynamic star help effects for FNTC (multiple and one-time)	85
Table 3.5 Dynamic star help effects for normalised raw publications (overall, multiple and one-time)	88
Table 3.6 Dynamic star help effects for FNTC (five types of helpful interactions).....	91
Table 3.7 Summary statistics: control and treated group (k-to-k matched) (a) balanced panel, (b) alternate matching criteria using the availability of publication match at t=0, (c) alternate matching criteria using acknowledgement text availability in publication at t=0.	94
Table 3.8 Robustness test of dynamic star help effects for FNTC (balanced panel and alternative matching criteria)	96
Table 3.9 Robustness test of dynamic star help effects for FNTC (country level)	99
Table 3.10 Robustness test of dynamic star help effects for FNTC (excluding the authors with previous indirect relations with a star (column1) and excluding the co-authored publications with a star after the initial acknowledgement event (column 2 & 3)).....	101
Table 3.11 Robustness test of dynamic star help effects for FNTC (comparing the initial productivity effects of authors who acknowledge stars- single and multiple times).....	103
Table 3.12 Summary statistics: by quartile of the cumulative FNTC distribution for the treated group at year t-1.	105
Table 3.13 Robustness test of dynamic star help effects for FNTC (for quartile levels)	105
Table 4.1 Mean value comparison of treated, control group, and broad subject field specifications	120
Table 4.2 Dynamic star move effects on star's own productivity for FNTC (overall, early and late career moves).....	122

Table 4.3 Dynamic star move effects on star’s own productivity for FNTC (broad subject fields) ..	125
Table 4.4 Robustness test of dynamic star move effects on star’s own productivity for FNTC (heterogeneous effects, first and subsequent star moves)	128
Table 4.5 Robustness test of dynamic star move effects on star’s own productivity for FNTC (alternative control group)	131
Table 4.6 Robustness test of dynamic star move effects on star’s own productivity for FNTC (country level)	133
Table 5.1 Balancing test between treatment and control groups	150
Table 5.2 Distributional characteristics for the relatedness measures	152
Table 5.3 Scientific relatedness to star: journal overlap	154
Table 5.4 Correlation matrix for different scientific relatedness measures	158
Table 5.5 Correlation matrix between relatedness across alternative knowledge spaces	160
Table 5.6 Robustness of alternative knowledge spaces	161
Table A.1 Country-level composition of department-year units 1996-2017	180
Table A.2 Characteristics of the treatment group (1996-2017)	181
Table A.3 Dynamic star arrival effects (restricted specification). Event window: 1996-2017. Observation window: 2001-2013	186
Table B.1 Yearly distribution of publications with acknowledgement texts and percentage of publications in which star names are acknowledged	191
Table B.2 Identification of help words used to describe the type of helpful interaction from a star scientist	193
Table B.3 Summary statistics: control and treated group (k-to-k matched) for an unbalanced panel for the types of helpful interactions	194
Table B.4 Dynamic star help effects for normalised raw publications (types of helpful interactions)	196
Table B.5 Robustness test of dynamic star help effects for FNTC (excluding the authors with co-authored publications with a star after initial contact)	198
Table C.1 Summary statistics: control and treated group (k-to-k matched) for the alternate control group	201
Table C.2 Robustness test on dynamic star move effects on star's own productivity for normalised raw publications	203
Table D.1 Summary of missing data for keywords by department	206

Table D.2 Correlation matrix between relatedness across alternative knowledge spaces.....	206
Table D.3 Robustness: quantity-adjusted symmetric (cosine) scientific relatedness to the star	207
Table D.4 Robustness alternative productivity measures: journal overlap	209
Table D.5 Robustness of alternative sample	211

List of Figures

Figure 1.1 Star catalyst framework	23
Figure 2.1 Event study model with the homogeneous (top) vs. heterogeneous (bottom) arrival effects across cohorts (department level)	50
Figure 2.2 Size of estimated star arrival effect four years post star arrival based on alternative thresholds for the field.	52
Figure 2.3 Event study model for matched incumbents (unbalanced panel)	54
Figure 2.4 Event study model for matched incumbents with star co-authoring effects (binary)	56
Figure 2.5 Related vs. unrelated.....	57
Figure 2.6 Event study model with homogenous arrival effects across cohorts for alternative outcome measures.....	58
Figure 2.7 Event study model with homogenous arrival effects across cohorts for alternative star definitions	59
Figure 3.1 Star help and observed scientific productivity.....	71
Figure 3.2 Event study model with homogenous (left) vs. heterogeneous (right) star help effects at an individual level.....	83
Figure 3.3 Event study model with multiple (left) vs. one-time (right) star help effects at an individual level.....	86
Figure 3.4 Event study model with homogenous star help effects at an individual level: overall star help (row 1); multiple star help (row 2, left); one-time star help (row 2, right)	89
Figure 3.5 Event study model with channels of help identified from the acknowledgement texts.....	93
Figure 3.6 Robustness test on event study model with homogenous star help effects at an individual level for a balanced panel	97
Figure 3.7 Robustness test on event study model with homogenous star help effects at the individual level using alternate matching criteria; the matched control author should have a publication in year $t=0$ (left), and the control author should have a publication with acknowledgement text in year $t=0$ (right), where year $t=0$ is the year when the treatment starts.	98
Figure 3.8 Event study model with homogenous star help effects at an individual level by country	100
Figure 3.9 Robustness test on event study model with homogenous star help effects at an individual level: excluding the authors with indirect star relationship before the acknowledgement event.....	102

Figure 3.10 Robustness test on event study model with homogenous star help effects at an individual level: excluding the future star co-authored publications for authors who acknowledge star help for: one-time (left); multiple times in the years after the first point of contact (right).	102
Figure 3.11 Robustness test on event study model with homogenous star help effects at an individual level comparing low-high productivity clusters by equal four quartiles of cumulative FNTC received for the treated authors one year before the treatment.....	107
Figure 4.1 Example of productivity effects under the institution-building mechanism.....	117
Figure 4.2 Event study plot with the dynamic effects on star movers' own productivity	123
Figure 4.3 Event study plot with the dynamic effects on star movers' own productivity for early career star movers (left) vs late career star movers(right).....	124
Figure 4.4 Event study plot with the dynamic effects on star movers' own productivity for broad field of research.....	126
Figure 4.5 Event study plot with the dynamic effects on star movers' own productivity for heterogeneous star moves	129
Figure 4.6 Event study plot with the dynamic effects on star movers' own productivity for first star (left) vs subsequent star movers(right).....	129
Figure 4.7 Event study plot with the dynamic effects on star movers' own productivity for alternate control group (left) and matched sample of treated and control units(right).	132
Figure 4.8 Event study plot with the dynamic effects on star movers' own productivity for Ireland Denmark and New Zealand	134
Figure 5.1. Event-study model for star arrival effects on incumbents' output.....	153
Figure 5.2 Non-linear relationship between star arrival and scientific relatedness effects – journal overlap space.....	156
Figure 5.3 Non-linear relationship between star arrival and scientific relatedness effects – alternative knowledge spaces.....	163
Figure 5.4 Nonparametric relationship between star arrival and scientific relatedness effects based on the journal co-occurrence measure	164
Figure A.1. Arrival and leaver rates as a function of talent stock.....	179
Figure A.2. Evolution of talent stock following arrival of the star	179
Figure A.3: Number of departments with first star arrival by country (1996-2017).....	182
Figure A.4: Event study model based on split between earlier and later career age star arrivals	182
Figure A.5: Event study model with raw output measure.....	183
Figure A.6 Event study model with homogenous arrival effects across cohorts by country (Ireland, Denmark, and New Zealand)	184

Figure A.7: Event study model for matched incumbents by country.....	185
Figure A.8: Estimates of the star arrival effect from the restricted model (see Table A.3)	188
Figure A.9: Comparison of autoregressive-distributed lag (ADL) with the baseline specification. ...	188
Figure B.1 Event study model with channels of help identified from the acknowledgement texts. ...	197
Figure B.2 Event study model with homogenous star help effects at an individual level :excluding the authors with future star co-authored publications that happen after the star help (t=0): overall star help effect (row1, left) , one-time star help effect (row1, right), multiple star help effect (row 2, left)...	199
Figure C.1 Event study model with star's own productivity effects for the alternative dependent variable.....	204
Figure D.1 Non-linear relationship between star arrival and scientific relatedness effects	208
Figure D.2 Non-linear relationship between star arrival and scientific relatedness effects-alternative productivity measures.	210
Figure D.3 Parametrically testing for non-linearity in effects for all three measures using journal overlap space.....	212

Declaration

I declare that this thesis, submitted to the University of Galway, Ireland for the degree of Doctor of Philosophy (Ph.D.), has not been submitted as an exercise for a degree at this or any other university. All research contained herewith is entirely my own and the use of all material from other sources has been properly and fully acknowledged.

Acknowledgements

I want to express my sincere gratitude to my supervisors, Professor John McHale and Dr. Jason Harold. Professor McHale's deep expertise in the field and extensive experience and networking skills proved invaluable. His meticulous attention to detail during editing further strengthened my thesis. I am also incredibly grateful to Dr. Harold for his vast knowledge of econometric modelling and his unwavering academic support throughout my PhD journey. In particular, his guidance through complex concepts was instrumental to my success.

My thanks also extends to my colleagues, Jen, Jefferson, Eoin, and Olha. Their willingness to share ideas and offer help during challenging situations was a constant source of encouragement. I am particularly thankful to Anil, my friend who accompanied me throughout this PhD journey. His unwavering mental support and extensive knowledge were crucial in helping me achieve my goal. His contributions were essential to the completion of this thesis.

I thank the University of Galway and the Discipline of Economics for providing a supportive research environment. The facilities and opportunities to interact with other PhD students fostered a stimulating academic atmosphere. My gratitude goes to my GRC committee, Cian Twomey, Ashley Piggins, and Aidan Kane. Their insightful comments and unwavering support throughout each GRC review process motivated me to strive for excellence and complete this thesis. Additionally, I am thankful to Claire Noone, Imelda Howley, and Charmain Byrne for facilitating access to research grants and travel opportunities, which significantly aided my research endeavours.

Finally, on a more personal note, I express my deepest gratitude to my parents and brother for their constant encouragement throughout my academic journey. Their unwavering belief in me provided invaluable strength. I am also incredibly grateful to my partner, Arathi. Her unwavering support, especially during difficult times, was crucial in helping me finish this thesis. To all my friends in Ireland, thank you for making me feel at home during my PhD experience. Your friendship made this journey even more meaningful.

Dedication

The thesis is dedicated to my parents Sasidharan, Bindhu, my brother Nikhil and all my friends.

Abbreviations

ATT	Average Treatment Effect on the Treated
CATT	Cohort Average Treatment Effect on the Treated
CEM	Coarsened Exact Matching
DiD	Difference-in-Differences
DNRF	Danish National Research Foundation
ESRI	Economic and Social Research Institute
FNTC	Field Normalised Total Citations
FWCI	Field Weighted Citation Impact
IW	Interaction Weighted
HRB	Health Research Board
NLP	Natural Language Processing
SFI	Science Foundation Ireland
SOE	Small Open Economy
R&D	Research and Development
TWFE	Two-Way Fixed Effect estimator

Abstract

In the ever-evolving global innovation landscape, small open economies (SOEs) face a unique challenge: fostering robust research ecosystems despite limited resources. A critical strategy for SOEs is maximising the impact of "star scientists" - who can catalyse research productivity and knowledge creation. This thesis investigates the influence of star scientists within SOEs, examining their influence on departmental output, incumbent productivity, and their own research trajectory following relocation. This research sheds light on the complex dynamics at play by employing novel methodologies and analysing data from three SOEs - Denmark, Ireland, and New Zealand. The findings offer valuable insights for policymakers and academic institutions, informing strategies for star scientist recruitment, integration, and fostering vibrant research clusters that propel knowledge-driven economic growth within SOEs.

The first essay investigates the impact of star arrival on departments in the SOEs. The arrival of a star scientist is expected to help increase departmental productivity. An event-study model estimates the dynamic effects of a star arrival on quality-adjusted research output at both the department and matched individual incumbent levels. The analysis considers the broader influence of star scientists, encompassing potential channels such as fostering new research norms, facilitating access to superior training and knowledge, and stimulating collaboration opportunities within the department.

Building on the prior investigation into the impact of star arrivals on departmental research output, the following essay explores the specific mechanisms through which star scientists influence their peer's productivity. While co-authorship offers a clear path for knowledge exchange, this chapter sheds light on the more nuanced effects of "star help" - assistance provided by a star scientist that falls outside the traditional co-authorship framework. I leverage natural language processing (NLP) techniques to analyse acknowledgement sections within a vast corpus of published papers. Employing an event-study framework with matched data, I estimate the causal effect of star help on author productivity. Furthermore, the chapter explores the influence of sustained star support, examining how ongoing engagement with a star scientist shapes an author's

long-term research output. These findings provide crucial insights for policymakers and academic institutions, informing strategies for star scientist recruitment that go beyond co-authorship.

While the previous essays explore the influence of star scientists on departmental productivity and incumbent output, the following essay addresses this gap in the literature by examining the research productivity of star scientists following their relocation to SOEs. I use an event-study model to isolate the causal effects of mobility on the star's productivity, comparing their performance to non-mobile star scientists within the same SOEs. This analysis explores the potential cost of mobility, investigating whether star scientists prioritise institution-building activities upon arrival, potentially at the expense of their immediate research output. Furthermore, the essay explores how career stage and scientific field may influence the impact of mobility on a star scientist's performance. By examining these dynamics, this research provides policymakers and academic institutions with evidence to support the design of star-scientist integration strategies that balance institution-building with continued research excellence from the stars themselves.

The final essay investigates how the knowledge space relatedness between a star and their peers intermediates the impact on those colleagues' productivity. In particular, this chapter investigates factors that affect the "intensity" of the treatment effect. I posit a non-linear relationship between relatedness to the star and the observed productivity effect on co-located scientists. While a stronger relationship can enhance the incumbent's absorptive capacity, it might also lead to knowledge redundancy. I employ a difference-in-difference model to examine this, estimating the treatment effects for varying degrees of relatedness (proxied by my relatedness measure). Furthermore, an event-study model with coarsened exact matching is utilised to help establish causal evidence of star arrivals on incumbent scientist productivity. By disentangling these dynamics, this research offers insights for academic institutions in fostering an environment that maximises the positive spillover effects of star scientists on their colleagues.

Chapter 1 Introduction

1.1 Context and Background

One of the ways science policy plays a role in supporting a growing economy is by facilitating the emergence of dynamic research clusters that underpin knowledge accumulation. Policies that support technological advance are central to modern theories of economic growth. The role of technology in the growth process has been debated since the Solow-Swan model (Solow, 1956; Swan, 1956). While hugely influential in shifting attention to the importance of technological change, the model had the limitation that it treated technology as an exogenous variable. In the 1980s, this approach was challenged by the emergence of new theories that endogenised the process of technological change (Romer, 1990; Lucas, 1988). The emergence of new growth theory had the effect of shifting attention to policies – including science policies – that could affect the pace of technological change and thereby the rate of economic growth over the long term. Rather than treating technology as a residual, the new theories treated such change as the result of deliberate Research and Development (R&D) efforts in the context of the given state of scientific knowledge. Endogenous growth theory thus reoriented attention from the traditional emphasis on capital accumulation to a new emphasis on the factors that support the generation of new knowledge and its diffusion among firms, sectors and countries. This, in turn, brought renewed attention to the role of national science and technology policies in terms of the direct production of knowledge in public institutions and the incentive structures for private knowledge-production efforts.

Science policy can be seen as one component of a broader national knowledge infrastructure. In part, such policies aim to create new knowledge by developing local research clusters. However, most of the literature focuses on the effects of science policies in larger economies. These studies focus on public funding in research, investment in the supply of Science, Technology, Engineering and Mathematics (STEM) graduates, direct R&D grants, R&D tax credits, facilitation of skilled migration and mission-oriented policies (Bloom et al., 2019). Relatively few studies focus specifically on the importance of such policies in the distinctive context of a small economy. In this thesis, I use data from small open economies (SOEs) to analyse the impacts of particular science policy

measures – policies to recruit star scientists – as a contribution to bridging this gap in the literature.

There is no single definition of what constitutes an SOE. Most definitions highlight the size of a country's population size, with a population threshold of 5 million being a widely used threshold (Roolaht, 2011). Whatever the precise definition, given the fixed-cost and public-good nature of investments in science, small size will inevitably limit the benefits of economies of scale that accrue to local knowledge production (Agrawal and Cockburn, 2003; Grossman and Helpman, 1991). Complicating the task of governments seeking to support the local production of science is that these investments compete with areas that tend to have greater short-term political salience, including health, primary/secondary education, and infrastructure.

An additional challenge is that the benefits of investments in science – especially at the more basic end of the basic-applied continuum – tend to spill over to other countries. The fraction of the benefits retained by the investing country will tend to be lower for a smaller economy than for a larger economy. Policymakers in SOEs may, therefore, look to investments with benefits that exhibit greater "local stickiness" than would their counterparts in larger economies. These policies could include a focus on more applied science that could more easily flow to local knowledge industries, integration into international knowledge networks (e.g., participation in European funding networks for smaller countries within the European Union), or more concentrated investments in targeted sectors. This thesis focuses on another complementary element of an SOE's science policy mix – star scientist-recruitment policies designed to catalyse local research clusters in targeted areas. By developing strong local clusters, such policies can, in principle, help overcome both the disadvantages of small size in terms of local scale and the disproportionate spillover of benefits beyond the country's borders.

This thesis provides an empirical investigation of the effects of star scientist arrivals on scientific productivity, looking at the level of science departments, incumbent scientists in those departments and the post-move productivity of the stars themselves. It also explores the conditions under which star recruitments are successful (e.g., the extent to which the expertise of the star is related to the expertise of scientists in the

receiving departments) and the mechanisms through which stars can influence incumbent productivity (notably through direct indications of star help as recorded in the acknowledgement texts of publications). While the main focus is on how star arrivals impact scientists at the receiving department, I also examine the effects of star-scientist moves on the productivity of the stars. Ultimately, the effects of star recruitments will depend on both the direct effects of the outputs produced by the star and the myriad indirect effects they have on the development of the local research cluster. While I do not attempt an overall cost-benefit assessment of star recruitment policies, by exploring a number of channels through which stars impact their receiving institution, we can form a better understanding of the strengths and weaknesses of star recruitment as an element of the science-policy mix of SOEs.

1.2 Overview of Science policies : Ireland , Denmark and New Zealand

Ireland has demonstrated notable advancements in its science policy landscape, transitioning from historically under-resourced conditions by the concerted efforts by of its academic institutions. These institutions have introduced courses in science policy and analysis, aligning with efforts by Irish research funding agencies such as the SFI Policy Research Programme and the Health Research Board's (HRB) policy investigations aimed at evaluating and enhancing science policy. Contributions from the Economic and Social Research Institute (ESRI) in building research capacity in the social sciences further reflect Ireland's commitment to evidence-based policymaking, addressing earlier criticisms regarding the lack of such integration (Ruane and Whelan, 2011; MacCarthaigh, 2013). Key strategies in Ireland's science policy include the strategic recruitment and retention of leading scientific talent through programs such as the SFI Research Professorships Programme and the President of Ireland Future Research Leaders Programme, significantly boosting research capabilities since 2003.

Recent initiatives in Ireland reflect a comprehensive approach to addressing both structural and societal challenges in science policy. The creation of 40 STEM professorships in 2019 to increase women's representation in STEM fields highlights the use of science policy to address gender disparities in higher education. Significant

investments in frontier research, such as the €29.6 million allocated through the Irish Research Council's Laureate Awards in 2018, demonstrates a commitment to filling gaps in the research and innovation landscape. Emphasising measurable outcomes like star publications, patents, and spin-outs, as outlined in SFI's Agenda 2020 strategy document, underscores the focus on tangible outputs driving economic and societal benefits.

Similarly, Denmark has strategically developed its science policies to foster dynamic and innovative research environments. Targeted recruitment of star scientists enhances research quality and impact (Schmidt et al., 2003). The Danish National Research Foundation's initiatives, such as the Niels Bohr Professorships, have been instrumental in attracting internationally renowned researchers to Danish institutions. These professorships provide substantial funding and resources, allowing star scientists to conduct high-impact research and collaborate extensively with local researchers, thereby elevating the overall research output and innovation within the country. In 2006, the Danish National Research Foundation (DNRF) established a five-year professorship extension to the Niels Bohr Visiting Professorships, aimed at bolstering the internationalisation and competitiveness of Danish research. This initiative, spanning from 2007 to 2012, sought to attract elite international scientists, including both foreigners and Danes abroad, to permanent positions at Danish universities. The appointed professors were tasked with developing and advancing research in prioritised areas, stimulating high-level international research cooperation, and strengthening research education programs. Each professorship also included positions for one or two younger scientists to further support research development. The DNRF provided financial support for the first five years of employment, after which the host institutions would cover the costs. The initiative was supported with a total funding of approximately 63.4 million Danish Krone, distributed across three grants, significantly enhancing the Danish research landscape by integrating global expertise.

Moreover, the Novo Nordisk Foundation has played a critical role in supporting the recruitment and integration of top-tier researchers through substantial financial investments in research centres and fellowship programs. In addition, these policies have facilitated important collaborations with global research entities, positioning Denmark as a leader in health and life sciences. Such initiatives underscore the importance of well-crafted science policies in driving both scientific excellence and societal benefits.

Despite these advancements, a critical challenge in the formulation of science policies is the lack of comprehensive evidence. This gap is particularly notable in Ireland, where the depth and consistency of policy analysis are often criticised (Boden and Fitzgibbon, 1995; Kane, 2005, 2014). While measurable outcomes like star publications and patents are highlighted, the integration of research evidence into policymaking remains a work in progress (Georghiou, 1995). In Denmark, although policies are well-crafted and focused on attracting elite researchers, there is a continuous need for robust evidence-based approaches to ensure the sustainability and effectiveness of these policies (European Commission, 2019).

New Zealand, despite its smaller market size and comparatively lower public and business research investment relative to the OECD average, has made strides in enhancing its science and innovation policies. The country's policy landscape is marked by an emphasis on fostering collaborations between research institutions and commercial organisations, aiming to leverage global research effectively. One significant initiative is the government's investment in the National Science Challenges, which focus on addressing major national and global challenges such as climate change, aging populations, and natural hazards. These challenges aim to drive collaborative research and innovation by bringing together researchers from diverse disciplines to tackle complex problems.

However, New Zealand faces ongoing challenges in science policy, particularly in attracting top international researchers. Unlike countries like Denmark, which have established targeted recruitment programs to draw elite scientists, New Zealand has not yet developed similarly focused strategies. This absence of targeted recruitment policies limits the country's ability to capitalise on international expertise and integrate it into its research ecosystem. The lack of such policies is a significant impediment to maximising New Zealand's research potential and achieving substantial economic and societal benefits (Jaffe, 2013).

Addressing these challenges, the 2013 report from the Prime Minister's Chief Science Advisor, Sir Peter Gluckman, highlighted the importance of strategic investments in research to support effective policy formation. The report emphasised the need for a more systematic approach to integrating robust evidence into policymaking. It pointed

out the variability in how scientific evidence is used across different government agencies and stressed the necessity for a consistent, whole-of-government approach to evaluating and implementing science policies (Gluckman, 2013). The report also underscored the need for increased investment in research infrastructure and human capital to enhance the country's innovation capacity. Furthermore, New Zealand's shift towards a more innovation-focused approach has sometimes led to the marginalisation of basic research, as noted in the literature. This shift has resulted in a preference for applied research that directly benefits business, often at the expense of long-term scientific inquiry and basic research (Leitch et al., 2014). This trend is indicative of the broader global shift towards demand-side policies where the National Innovation Systems are harnessed to drive economic growth.

New Zealand's approach to science policy includes initiatives such as the Research, Science, and Innovation (RSI) system, which aims to create a more coordinated and strategic research environment. The government's emphasis on increasing collaboration between research institutions and industries is designed to enhance the impact of research outputs and drive innovation. However, to achieve significant improvements, there is a recognised need for more proactive and targeted recruitment efforts to attract high-calibre international researchers. Addressing this gap is essential for boosting New Zealand's research profile and ensuring that it remains competitive on a global scale.

1.3 Understanding Star Spillover: Gaps in the Current Literature

Knowledge spillover is the process of sharing and transferring novel ideas and discoveries among different individuals and groups. It is essential to advancing scientific knowledge and fostering innovation. In academia, one of the main sources of knowledge spillover is the presence of star scientists, who are researchers with exceptional achievements and influence in their fields. Moreover, these elite scientists are responsible for a large stock of published research (Lotka, 1926). Star scientists have been shown to improve the performance of the organisations they belong to, as well as the quality and impact of their research outputs (Zucker and Darby, 2006; Hess and

Rothaermel, 2007). They also have significant direct and indirect peer effects on their collaborators and other researchers in their fields.

The direct peer effects of star scientists are mainly mediated by co-authorship, which allows the exchange of knowledge and skills between the star and their collaborators (Agrawal et al., 2017). However, co-authorship also entails a dependency on the star, which can have negative consequences for the collaborators in case of the star's departure or death. Several studies have found that the unexpected death of a star scientist leads to a decline in the productivity and quality of their co-authors (Azoulay et al., 2010; Oettl, 2012). This effect can be mitigated by the size and diversity of the collaboration network, which provides alternative sources of knowledge and support (Khanna, 2021). However, the indirect peer effects of star scientists are less explored but potentially more widespread and lasting. One recent study by Betancourt et al. (2023) found that the co-authors of Nobel Prize winners in Physics receive a citation boost for their publications prior to their first collaboration with the winner, indicating that the star's recognition enhances the visibility and credibility of their collaborators. This effect may also extend to other researchers who work on similar topics or use methods similar to those of the star, as they benefit from the increased attention and legitimacy that the star brings to their field (Simcoe and Waguespack, 2011; Merton, 1968). These findings suggest that star scientists play a crucial role in stimulating knowledge spillover in academia, directly through their own outputs and indirectly through various network effects.

One way for an SOE to access a star's expertise is by recruiting them to a local research cluster. Recruiting stars enable local scientists to acquire tacit knowledge that is hard to communicate or codify due to its complexity and specificity. Proximity to the star facilitates a smoother transfer of tacit knowledge, allowing more frequent and informal interactions. Despite the advances in information technology, proximity still matters for scientific productivity (Bourdeau et al., 2017; Catalini, 2018). Recruiting a star to an institution can generate positive spillovers for the institution and the region, as it enhances the opportunities for collaboration, mentoring, and networking with the star (Glaeser, 1999; Agrawal et al., 2014).

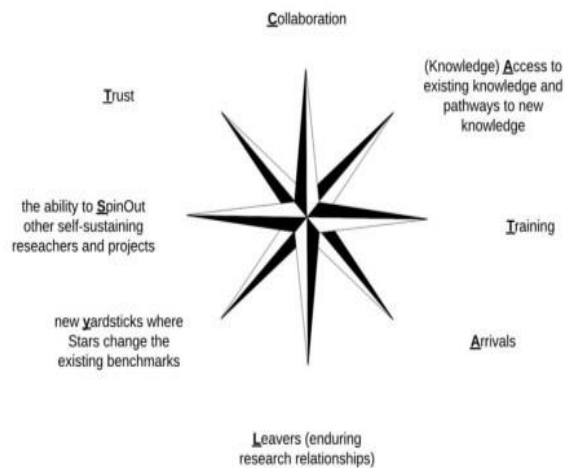


Figure 1.1 Star catalyst framework

Figure 1.1 illustrates the types of network relations that can be enhanced by recruiting a star to an institution in an SOE. As mentioned in the previous section, collaboration is a key mechanism for producing new knowledge. We have empirical evidence to show that co-authorship with a star boosts the productivity of scientists even when joint publications with the star are excluded (Yadav et al., 2023). Thus, a star's arrival at an institution in an SOE is expected to create new connections with local scientists. Furthermore, a star's presence can help develop the local network and attract further talented recruits to improve the institution's performance (Agrawal et al., 2017). Along with attracting talented scientists, stars can help retain the institution's high-performing scientists (i.e., reduce the number of "leavers").

Identifying and recruiting scientific talent is challenging for policymakers and institutions. They rely on the paper trail of publications, citations, accomplishments, grants, etc., left by these elite scientists to signal their work quality (Jaffe and Trajtenberg, 2002; Lehmann et al., 2006). However, it is uncertain whether these scientists can maintain their high-quality work after moving to a new institution. One of the contexts where this question is relevant is in the small open economy (SOE), which invests in star scientists to improve their research performance and attract more funding. The post-move productivity of these stars is not well understood and the stars may face – indeed embrace – additional academic duties, such as training and mentoring young scientists or developing and restructuring the research bodies, which could affect their own research productivity. Recognising this dual role of star scientists, this thesis seeks to improve our

understanding of the effects of star recruitment both on other scientists at the receiving institution and on the productivity of the stars themselves.

The thesis delves into four novel research questions, each forming the basis of one of the four main chapters:

1. What is the effect of the recruitment of a star on the productivity of the receiving department and on the productivity of individual incumbents in that department?

2. What is the role of "star help" in supporting the productivity of peer scientists, where such help is indicated by acknowledgements to the star in the acknowledgment texts of non-co-authoring star scientists?

3. What is the effect of a star move on the arriving star scientist's own productivity?

4. How does the degree of relatedness between the work of the star and the incumbent scientist mediate the impact of the star's arrival on an incumbent's productivity?

In the remainder of this introduction, I first describe the data used to address these four questions and then outline the main contributions of each of the four chapters.

1.4 Overview of the Data

To address the research questions in this thesis, I use publication data collected from Scopus, a large and comprehensive database of abstracts and citations from peer-reviewed journal articles, books, trade journals, patent records, and conference publications. It covers various disciplines in life sciences, social sciences, physical sciences, and health sciences. Given the richness of these data, they can be used for various purposes, such as finding relevant and authoritative research, identifying experts and collaborators, accessing reliable data and metrics, and analysing research trends and patterns.

In particular, given the focus of the thesis on SOEs, I use publication data collected from Scopus for three SOEs: Ireland, Denmark, and New Zealand from 1990 to 2017. Denmark and Ireland are small open economies that have implemented national programmes to recruit star researchers. For example, the Danish National Research Foundation (DNRF) has launched various programmes, including the Niels Bohr Visiting Professorships, the DNRF Professorships, and the latest DNRF Professorships (2021), with the aim of enriching Danish research communities with top-class researchers from abroad. Similarly, Science Foundation Ireland (SFI) initiated the Research Professorship and Future Research Leadership Programmes, which have supported research institutions in attracting outstanding international research talent since 2003. In contrast, New Zealand is a small open economy with no formal star recruitment programme.

For each collected publication, I record the year, authors, affiliations, subject field, citation count to 2019, references, abstracts, acknowledgements, keywords, journal information, and, importantly, two unique identifiers – the EID Scopus publication identifier and the Scopus Author(s) ID. This gives us approximately 1.43 million publications divided over 219,582 unique authors. For each author, I also access their catalogue of publications back to 1990, including publications that took place prior to joining the department in Denmark, Ireland or New Zealand. For each chapter, I use the citation data to define a star scientist as someone who is at or above the 95th percentile of scientists in the cumulative distribution of citations received since 1990 for their subject field in any given year. Further, in each chapter, I use the information tagged to each author ID, such as affiliations, keywords, journal details and acknowledgement texts, to explore the different mechanisms of star scientist spillover. Table 1.1 provides a detailed breakdown of the origin country from which the star scientists are arriving to the focal countries: Ireland, Denmark and New Zealand. Notably, 50% of the star scientists are relocating from United States and United Kingdom, with Denmark receiving the highest number of star arrivals.

Table 1.1 Counts of star scientists based on their origin countries from which the stars are moving from into the SOEs.

Origin Country	Country of Arrival			Total
	Denmark	Ireland	New Zealand	
Australia	1	0	3	4
Canada	2	3	2	7
China	1	1	2	4
Czech Republic	0	1	0	1
Denmark	6	0	0	6
Finland	4	0	1	5
France	1	0	1	2
Germany	8	9	0	17
Hong Kong	2	1	1	4
Ireland	0	1	0	1
Italy	0	1	2	3
Netherlands	2	2	0	4
New Zealand	0	0	6	6
Norway	1	0	0	1
Singapore	1	0	0	1
South Korea	1	0	0	1
Spain	0	1	0	1
Sweden	8	0	0	8
Switzerland	2	0	0	2
United Kingdom	9	14	13	36
United States	33	9	10	52
West Indies	0	0	1	1
Total	82	43	42	167

Note: The data is obtained from the affiliation details tagged to each of the publications of the star scientists. There are a total of 13 internal moves within the focal countries recorded. (6 in Denmark, 1 in Ireland and 6 in New Zealand)

In Chapter 2, I focus on the effect of the first star arrival observed in the data on the productivity of the recruiting department and on the productivity of incumbent scientists in that department. The departments are defined based on the Scopus-defined subject fields rather than the departmental structure at a particular institution. Furthermore, the study identifies a star arriving at the department as a star that has not previously published with an affiliation to the department and is now recording an affiliation to that department for the first time. In total, 81 first-star arrivals are recorded as their year of arrival and publication output. Here, the identification of star arrivals is restricted to those who publish with an affiliation to the department for at least the next four years. The chapter also identifies the individual publication output of incumbent scientists at the receiving department. I use the same departmental identification method as before to determine the department associated with each scientist in the data.

In Chapter 3, the research explores the mechanism of star help using the acknowledgements texts included in the publications. To expand the set of star acknowledgement events, this study is not limited to the stars who arrive at the SOEs; instead, I identify all the stars (981) who passed the threshold for star identification outlined above. Notably, the chapter employs a text analysis of acknowledgement texts to identify the star names and the help words that describe the type of help provided by the star. One limitation of the data here is that acknowledgement texts are unavailable for all the Scopus publications. More details of the available publications with acknowledgements and the identification of star names in each text are provided in the chapter.

Next, in Chapter 4, the focus shifts to the effects of star scientist moves on the star's own productivity. While in Chapter 2, the research explores the impact on incumbent scientists connected with the star who arrives at the department, this chapter examines the differences in the productivity between two groups of star scientists – star movers and star non-movers. Hence, the chapter also requires the identification of non-mover stars in the data in addition to the previously identified star movers/arrivals. For consistency, like in Chapter 2, star arrivals are restricted to those who continue to publish in the arriving institution for at least four years. Using the affiliation information for each identified star's publication, I identify the locations of these stars across time from 1990-2017. Then, I divide the stars between movers and non-movers, where the non-movers constitute the comparison group for the analysis.

Finally, in Chapter 5, the research employs a number of relatedness measures between stars and incumbent scientists. Similar to the data used in Chapter 2, I use the authors' information from the star arrival departments to measure their degree of relatedness to the arriving star. Therefore, the chapter uses each scientist's publications record to build a scientific knowledge space that proxies individuals' accumulated knowledge sets. For that, I use the information from the journals associated with each publication provided by Scopus so that scientists who publish in the same journals are more likely to share a 'related' knowledge set. Furthermore, as robustness tests on the measure of relatedness, the chapter uses the journal category overlap, author-reported keyword overlap, and Scopus-reported keyword overlap.

1.5 Overview of the Thesis Contributions

This thesis aims to provide an evidence-based analysis of star scientists and their impacts through various mechanisms in Ireland, Denmark and New Zealand. In addition to advancing our understanding of this relatively neglected element of the science policy mix, an important objective is to help the policymakers in the SOEs to make better evidence-based design decisions relating to star recruitment policies. With this aim, the remainder of the thesis consists of four main chapters and a concluding chapter and follows an article-based format. Chapter 2 evaluates the peer productivity effects of star scientists' arrival in one of these SOEs. In contrast, Chapter 3 analyses the specific channel of star help using acknowledgement texts as an indicator of the provision of help. In Chapter 4, I examine the impact of star mobility on their own productivity and in Chapter 5, I explore the implications of relatedness to a star and its effect on the incumbent's productivity. Finally, in Chapter 6, I summarise my overall findings. I next provide more detail on the contributions of the individual chapters.

In Chapter 2, I discuss the impacts and policy implications of recruiting star scientists to catalyse the successful development of local clusters. I hypothesise that the productivity of scientists depends on their position in scientific networks, and an essential aspect of these networks is being co-located with star scientists who are arriving at the departments in the three focal countries. The chapter uses a (panel) event-study framework to explore the dynamic impacts of star arrivals on the arriving institution's departmental and individual incumbent performance. Using Scopus data on publications and citations from 1990 to 2017, I define star scientists as scientists above the 95th percentile in the cumulative distribution of quality-adjusted output, where the quality measure is based on the observed citations to the end of the sample period. To help address the potential endogeneity bias involving recruiting star scientists to a department, the chapter also focuses on individual scientists who are present in the department when the star arrives ("incumbents") to make more credible causal inferences for the overall star-arrival effect on performance. The various event-study analyses of first-star arrivals to the departments support the hypothesis of a positive productivity effect of star arrivals at both the departmental and incumbent levels. Star recruitment has positive and sustained productivity effects on the departmental output,

which is between 12% and 25% higher after four years (excluding the star's output). The absence of pre-arrival effects suggests a causal effect of star arrivals at the departmental level; however, we cannot rule out the possibility that stars are attracted to departments with high productivity potential. This concern should be diminished at the individual incumbent scientist level. Using a coarsened-exact-matching (CEM) procedure (to match individuals at departments that either receive a star and those that do not), I find that star arrivals impact the output of incumbents, increasing their quality-adjusted output by 5% with no evidence of pre-trends. Stars arriving relatively early in their career have a larger positive effect. The results are robust to various alternative specifications and estimation methods, such as staggered treatment effects, alternative output measures and alternative star identification strategies. The chapter discusses the various policy implications for star recruitment and their integration into the receiving SOE institutions.

In Chapter 3, I examine how receiving informal help from a star scientist affects the productivity of authors who acknowledge the star in their publications. Informal help refers to any assistance that is not classified as co-authorship, such as providing conceptual guidance, technical support, material resources, grants and funds, or other types of acknowledgements. I use natural language processing (NLP) to extract the names of stars and manually extract the help words from the acknowledgement texts of the publications available in the three countries: Ireland, Denmark, and New Zealand. I then apply an event study methodology to estimate the pre-and post-treatment effects of the star's help on the author's productivity, measured by the quality and quantity of publications. This chapter contributes to the literature on knowledge spillovers and star scientists by focusing on the informal channel of help, which is often overlooked compared to the more direct effects of co-authorship. By providing informal help, star scientists can also share their expertise, resources, and networks with other researchers, which can enhance their productivity. Moreover, informal help can also capture the effects of star scientists who are not co-located with the authors, as they can still provide help remotely or through visits. The results of this chapter support my hypothesis that receiving informal help from a star scientist has a positive and significant effect on the author's productivity in the year of acknowledgement, both in terms of the raw publication output and the quality-adjusted output. The effect decreases in the subsequent years but remains positive and substantial for the authors who continue to

acknowledge the star. I also find that different forms of informal help affect the author's productivity differently. Conceptual help has the largest effect in the year of acknowledgement, while help recognised from the donation of materials has the most sustained effect in the subsequent years. Finally, the results indicate the value of policy measures that increase the star's opportunities for helpful interactions with incumbents.

Next, in Chapter 4, I examine the scientific productivity of star scientists who move to the SOEs and how their move affects their own productivity. Here, I propose three mechanisms through which the move could affect the star's own productivity, either positively or negatively: the matching mechanism, which supposes that moves are associated with better matches between the star and their institution; the disruption mechanism, which supposes there is a short-term negative effect on star's productivity from the disruptive effects of the move; and the institution-building mechanism, which assumes there is a more sustained loss in the productivity of star as they reallocate their time and effort towards institution-building activities and away from activities more directly focused on their own productivity. Using an event study methodology, the chapter uses the Scopus data to identify 161-star movers in the three countries and compares their quality-adjusted productivity with 75 star non-movers. The results show a striking and persistent post-move decrease in the star's productivity. Taken together with the earlier findings on the positive effect of the star's arrival on departmental and incumbent productivity, these results are most consistent with the institution-building mechanism. The results show a contemporaneous decrease in the quality-adjusted productivity of star scientists in the year of the move of approximately 58%, which remains negative for the next four years but becomes less in magnitude. In addition, this initial dip is more pronounced for the established stars who move to the SOEs later in their career. Also, the chapter finds similar patterns across different scientific fields – physical, health, and life sciences. For social sciences, however, the results show only small, move-related productivity changes and are suggestive of the greater relative importance of the matching mechanism. The results are robust to various alternative analyses, including allowing for heterogeneous treatment effects, alternative control samples and different definitions of the dependent variable. The chapter concludes with a discussion of the implications of the results for policy design. I suggest that the policymakers should anticipate a decrease in the productivity of the star after the move.

However, the fall in productivity could be mitigated by ensuring that the necessary supporting infrastructure is in place contemporaneously with the star's arrival to reduce the observed disruption costs. More positively, assuming an institutional-building motivation on the part of the arriving star, the institution should facilitate interactions between the star and incumbents and possibly take advantage of the star's experience to take on leadership roles at their new institutions.

In Chapter 5, I investigate how the productivity of incumbents is affected by their relatedness with the star. The chapter aims to understand how the intensity of the treatment (exposure), as measured by relatedness, helps the incumbents access novel knowledge upon the star's arrival, thereby creating new knowledge through combining different ideas. I use various co-occurrence measures and develop suitable indicators of relatedness. For example, the baseline measure of relatedness for the analysis is the co-occurrences of publications in the same journal by star and the incumbent. In addition, the chapter also measures the asymmetric relatedness measure, which captures how related the knowledge bases are from the perspective of the incumbent scientists and the time-constrained star. Here, I hypothesise that greater scientific relatedness will mean a greater intensity of treatment from a star's arrival, given that non-redundant knowledge of a star is assumed to be a star's knowledge stock less the fraction of that knowledge already possessed by the incumbent scientists. In this chapter, I apply a difference-in-differences framework to examine the possible non-linear interaction between the degree of relatedness to the star and the size of the star's arrival effect. The results show that the greater the extent of relatedness to the star from the star's perspective, the larger the positive productivity effect from the star's arrival. The results are robust to symmetric and asymmetric relatedness measures and additional relatedness measures using keywords and scientific categories. Finally, the chapter discusses the specific policy implications of the findings, including the importance of the presence of a cadre of scientists at the receiving institution with research interests related to those of the star.

Finally, in Chapter 6, I briefly summarise the findings of the four main chapters and provide an overall conclusion for the thesis. In particular, I discuss the policy implications of my findings regarding their consequences for the design of star recruitment policies. Overall, I highlight the thesis's contribution to the small but growing literature on star recruitment policies and, more particularly, their impacts on small open

economies. I close with a discussion of the thesis's limitations and potential avenues for future research.

1.6 Thesis Outputs

Peer-reviewed Journal Articles

Chapter 2: McHale, J., Harold, J., Mei, J. C., Sasidharan, A., & Yadav, A. (2023). Stars as catalysts: an event-study analysis of the impact of star-scientist recruitment on local research performance in a small open economy. *Journal of Economic Geography*, 23(2), 343-369.

Chapter 3: Sasidharan, A., McHale, J., & Harold, J. (2024). Star help and knowledge transfer: an event study analysis of star interactions observed from acknowledgement texts. *The Journal of Technology Transfer*, 1-49.

Journal Article Under Review

Chapter 4: Sasidharan, A., McHale, J., & Harold, J. (2024). Estimating the impact of a star scientist move on their own research productivity. (Under Review)

Chapter 5: McHale, J., Galetti, J. R., Harold, J., Mei, D. J. C., Sasidharan, A., & Yadav, A. (2023) Star Arrival Effects on Scientists' Productivity: Does the Nature and Extent of Relatedness to the Star Matter? (Under Review.)

Conference and Seminar Presentations

Sasidharan, A., McHale, J., & Harold, J. (2024). Star help and knowledge transfer: an event study analysis of star interactions observed from acknowledgement texts. Presented at:

- The Irish Economic Association 36th Annual Conference, May 2023, Economics and Social Research Institute Ireland
- Economics of science workshop, May 2023, University of Bordeaux, France
- GMIT School of Business Research Group (BRING) workshop, May 2021, GMIT

Sasidharan, A., McHale, J., & Harold, J. (2024). Estimating the impact of a star scientist move on their own research productivity

Presented at:

- PhD seminar Series, September 2023, Discipline of Economics, University of Galway
- The Irish Economic Association 37th Annual Conference, May 2024, University of Galway

Chapter 2 Stars as Catalysts: An Event-Study Analysis of the Impact of Star-Scientist Recruitment on Local Research Performance in a Small Open Economy

2.1 Introduction

Policy makers in small open economies face particular challenges in justifying investments in scientific research. Smallness can act as a disincentive to two ways. First, given that investments add to the global stock of knowledge, smallness means that a large share of the benefits flow outside as spillovers to the global stock. And second, given the importance of scale economies for scientific productivity – and notably the importance of local knowledge-sharing and collaboration networks – a lack of scale could put small-country research institutions at a competitive disadvantage. Strategies such as targeting research with substantial local spillovers (especially more applied and use-inspired basic research), developing of centres of excellence to achieve necessary scale, and integrating into international knowledge and collaboration networks¹ are pursued to overcome this size disadvantage. The targeted recruitment of star scientists is a potential additional element in the strategy mix for policy makers and research institutions. One objective behind such policies is that "stars" – disproportionately productive and connected scientists in their field – will catalyse the development of successful local clusters in targeted areas. In addition to their own direct outputs, it is hoped that stars will raise the productivity of their peers through such channels as increased access to knowledge networks, greater collaboration opportunities, mentoring and even changed norms relating to the conduct of scientific research. It is also hoped that they aid in the attraction and retention of other scientific talent, thereby supporting the growth of targeted high-performing clusters.

As examples of nationally supported star recruitment policies, the Danish National Research Foundation (DNRF) has implemented a series of programmes: the Niels Bohr

¹ For example, participation in EU-wide funding consortia for European Research Council funding.

Visiting Professorships, the D NRF Professorships, the Niels Bohr Professorships and the latest D NRF Professorships (2021). Although the design has shifted somewhat over time, the core purpose has been the "enriching of Danish research communities with top-class researchers from abroad." Attracting outstanding international research talent to Ireland is also the objective behind the Science Foundation Ireland (SFI) Research Professorship and Future Research Leadership Programmes. Since the initial launch in 2003, programme design has also evolved, but a consistent goal has been to support research institutions in their recruitment of world-class researchers.²

We hypothesise that the productivity of a scientist depends on their position in scientific networks and that an important component of these networks is local. The relevant networks could include, for example, those that support knowledge access or provide co-authorship opportunities. The importance of location is assumed to reflect the lower cost of forming network links when scientists are co-located. An arrival of a star at a scientist's department is hypothesized to strengthen that scientist's network in ways that support their productivity. Moreover, if other scientists are attracted by the presence of a star, and if already present scientists are less likely to leave with that presence, the arrival of the star could further catalyse the development of the department through improved recruitment and retention.³

We use a (panel) event-study framework to explore the dynamic impacts of star arrivals (excluding the direct publication impact of the star) on receiving university department and individual incumbent performance in three small countries – Denmark, Ireland and New Zealand. We assemble Scopus data on publications and citations to those publications for the period 1990 to 2017. Star scientists are defined as scientists above a particular percentile (the 95th percentile and above in our baseline specification) in the cumulative distribution of quality-adjusted output, where the relevant quality measure is based on observed citations to a publication to the end of the sample period. In turn, star arrivals are identified as long-term arrivals at institutions where they were not previously affiliated as evidenced by the affiliations recorded on their publications. We implement a standard event-study model with the assumption of homogenous arrival

² Conditional on funding availability, SFI's recent strategy for the years to 2025 targets the recruitment of "20 world-class researchers to Ireland annually" (Science Foundation Ireland, 2021, p. 9).

³ We develop a simple model of such catalytic effects of star arrivals in Appendix A

effects across arrival cohorts and also allow for heterogeneous arrival effects across cohorts using the method of Sun and Abraham (2021).

Notwithstanding the value of the event-study design to protect against a number of sources of endogeneity bias, there remains a concern that the decision of a department to hire a star, and/or the decision of a star to join a department, are not independent of foreseeable but uncontrolled for factors that affect departmental performance. A complementary focus on individual incumbent scientists at the time of the star's arrival allows us to make more credible causal inferences for one important component of the overall star-arrival effect on performance. There are a number of reasons. First, while a star's arrival may be influenced by expectations of future departmental performance (including future planned investments), the arrival is less likely to be influenced by expectations related to individual scientists. Second, we can implement a matching procedure at the individual level where each incumbent is paired with a highly similar scientist at a non-star recruiting department at another institution. As detailed in the next section, we implement this matching with a Coarsened Exact Matching (CEM) procedure. And third, we can better explore particular network-related causal mechanisms through which a star arrival affects productivity. Specifically, we explore how co-authorship with the star affects post-arrival incumbent productivity. We therefore present results both for department-year and (matched) individual-incumbent-year units of analyses.

Our study is related to a number of literatures, including those on stars, networks, mobility and agglomeration. The tendency for economic and social outcomes to follow fat-tailed and skewed distributions – with outsized outcomes for stars – has been observed in many settings.⁴ In science, such fat-tailed distributions have long been observed for publications, citations and collaborations (Lotka, 1926; Price, 1963; Newman, 2001; Goyal et al., 2006). These distributions could be explained by a highly uneven underlying distribution of talent or by network effects that allow small initial differences to become magnified over time through a cumulative advantage process (Merton, 1968; Azoulay et al., 2014). In a market setting with economies of scale, Rosen (1981) shows how small differences in ability can lead to outsized differences in rewards

⁴ See, for example, Gabaix (1999) for city sizes, Gyourko et al. (2013) for city house prices, Malmendier and Tate (2009) for CEO pay and Autor et al. (2020) for firm profits.

for "Superstars" through a "winner-take-all" phenomenon. Adler (1985) provides an alternative account of the emergence of stars that depends on network effects and first-mover advantage – indeed he models the emergence of stars even where initial talents are identical.

The literature on economic and social effects of networks has highlighted the potential effects of network structure – captured by metrics such as density, centrality and clustering – on individual behaviour and outcomes (Jackson, 2008; Jackson et al., 2017). Moreover, the economic approach to network formation emphasises how network participants weigh the costs and benefits of forming connections (Goyal, 2007). Of relevance to our setting, where co-location lowers the cost of forming connections, network clustering will arise at the local level suggesting the potential for arriving stars to become embedded in local networks even while retaining more distant connections (Jackson and Rogers, 2005; 2007).⁵

For scientist networks, an important strand of the literature has used exogenous star departures or arrivals to identify the impact of a disruption to a peer network on the productivity of members of that network. Waldinger (2012) uses the dismissal of scientists in Nazi Germany as a source of exogenous variation in the quality and quantity of the peer group of remaining scientists, but does not find evidence of harm. However, Waldinger (2010) does find evidence of harm from these dismissals on PhD students. Azoulay et al. (2010) find that collaborators suffer a lasting reduction in quality-adjusted output upon the unexpected death of a star, with collaborators closer to the star in "idea space" suffering a sharper decline. Also using unexpected star deaths, Oettl (2012) reports that the unexpected death of "helpful star scientists" negatively impacts the quality of but not the quantity of their co-authors output. In a recent paper, using data from pharmacology and pharmacy, Khanna (2021) finds an across-the-board decline in co-author productivity following a star death, but the effect is attenuated by the size of the scientist's collaboration network. In an event-study setting, Agrawal et al. (2017)

⁵ A large literature has investigated the tendency of knowledge flow to be disproportionately localised (see Grilliches, 1992, for an early survey). In an influential paper, Jaffe et al. (1993) focus on the geographic localisation of citations to patents and find that citing and cited patents are disproportionately co-located after controlling for the geographic distribution of activity in the relevant area. Breschi and Lissoni (2009) find that a significant fraction of knowledge flows across firms and locations is due to labour mobility that is geographically bounded.

report evidence that in the field of evolutionary biology star arrivals increase the productivity of incumbents working in related areas and also increase the quality of subsequent hires.⁶

The traditional migration literature has generally treated mobility as a one-time, all-or-nothing event, examining, for example, the impact of an immigration shock on local wages (e.g. Card, 1990; Borjas, 2003) or the impact of the "brain drain" on sending destinations. In contrast, a more recent literature has emphasised the dynamic nature of migration (or "brain circulation") – including, for example, anticipation effects relating to the prospect of future moves and the importance of return migration (Kapur and McHale, 2006; Docquier and Rapoport, 2012; Kerr, 2019). The nature of moves can also be more partial with, for example, members of diaspora networks retaining connections to networks in the country or region of origin even as they form connections in their new locations. Agrawal et al. (2006) find evidence of enduring links that support knowledge flow between departed inventors and their original locations. Using data for Indian inventors, Agrawal et al. (2011) examine the trade-off between increased access to knowledge from connections to the diaspora and the thinning of local knowledge networks that results from inventor out migration.

Finally, our study is related to the literature on agglomeration (Krugman 1991; and Black and Henderson, 1999; Duranton and Puga, 2004; Glaeser, 2008). Agglomeration models typically combine scale economies (including through network effects) and mobility (individuals and/or firms are attracted to highly performing regions). In a typical model, the spatial distribution of activity is affected by the balance between centripetal forces – pre-existing scale attracts newcomers leading to further scale – and centrifugal forces resulting from some form of congestion effect at the expanding location. Fujita et al. (2000) examine how the core-periphery pattern can emerge as the equilibrium outcome of such forces. Audretsch and Feldman (1996) report evidence on the importance of knowledge spillovers to the spatial clustering of activity in knowledge-intensive industries. One question is how a positive feedback loop that leads to a successful cluster of activity gets started. Of particular relevance to the importance

⁶ Lacetera et al. (2004) and Hess and Rothaermel (2011) examine the effect of star arrivals on firm-level innovation performance.

of star scientists, Zucker et al. (1998) identify the pre-existing geographical distribution of star scientists as an important determinant of the geographical distribution of activity in the biotechnology industry.⁷

The remainder of the chapter is structured as follows. Section 2.2 sets out the panel event-study framework that we use for the estimation of the dynamic effects of star arrivals under the assumptions of both homogenous and heterogeneous cohort arrival effects. Section 2.3 describes our Scopus-derived data and outlines our star and star arrival identification strategies as well as our department/incumbent research output metrics and incumbent matching procedure. Section 2.4 reports our results for both the combined three-country sample and each of the three countries individually. Section 2.5 conducts a number of robustness tests on our baseline results. Section 2.6 concludes with a discussion of the relevance of our findings for policy and institution-level strategies.

2.2 Econometric Methodology

2.2.1 Baseline distributed lead/lag specification

The gradual strengthening of network effects provides a mechanism through which the arrival of a star could impact departmental performance over time. To measure these dynamic effects, we begin with a general baseline two-way fixed effects (TWFE) specification that assumes the multiplicative infinite distributed lead/lag form:

$$\ln Y_{i,t} = \alpha + \sum_{j'=-\infty}^{\infty} \gamma_{j'} S_{i,t-j'} + X'_{i,t} \lambda + \phi d_i + \varphi d_t + u_{i,t} \quad (2.1)$$

where $Y_{i,t}$ is a measure of quality-adjusted output (e.g. citations-weighted publications) in department i in year t , $S_{i,t-j'}$ is a dummy variable indicating if a (first-arriving) star is present in the department in year $t - j'$, $X'_{i,t}$ is a vector of controls, ϕd_i is a department fixed effect, φd_t is a year fixed effect and $u_{i,t}$ is a zero mean error term. Depending on the

⁷ Also in the knowledge industry context, Agrawal and Henderson (2003) find that the presence of a large R&D intensive firm – what they call an “anchor tenant” – enhances the ability of the regional innovation system to absorb university research and stimulate knowledge-intensive output.

specification, the unit of analysis can be departments in a given year or individual incumbents attached to a particular department in a given year.

The estimated dynamic multipliers, $\gamma_{j'}$, have the convenient interpretation of semi-elasticities. The cumulative dynamic multiplier j periods after a star arrives (where j can be negative if the star has yet to arrive) is⁸:

$$\beta_j = \sum_{j'=-\infty}^j \gamma_{j'}, \quad \text{or} \quad \beta_j = \beta_{j-1} + \gamma_{j'} \quad (2.2)$$

We initially make the strong assumption of strict exogeneity so that the contemporaneous error is independent of past, present and future star arrivals. We assume that if a star arrives in a department that they remain present in the department from that time period forward so that the arrival of a star is an "absorbing" event.⁹ The multiplicative form allows for star arrival and time effects to be proportional to current department-level performance. These (proportional) star arrival effects in our baseline model are initially assumed to be homogeneous across time and departments/incumbents, but we then relax this assumption to allow for the possibility that star arrival effects are heterogeneous across arrival cohorts.

The presence of lead effects in our general specification allows in principle for anticipation effects of star arrivals on productivity performance. However, in our empirical analysis we take a conservative approach and assume such anticipation effects are zero. Instead, we suppose that any observed pre-arrival effects – or, even more concerning, a positive trend in pre-arrival effects – is evidence of a failure of the parallel trends assumption. The absence of an observed pre-trend therefore helps mitigate concerns of a correlation between star recruitment and unobserved factors that independently affect department performance. However, even in the absence of a pre-trend, it is possible that contemporaneous (or subsequent) performance-enhancing developments in the department are associated with star recruitment. This could be because, for example, these developments attract the star (reverse causality) or possibly

⁸ Therefore, j has the interpretation of "event time."

⁹ This assumption can be weakened to allow for star exits recognising that a star's arrival (and temporary presence) can have effects on the department even after they have left.

because the recruitment of the star is part of a package that includes other performance-enhancing changes in the department (omitted variables). Therefore, in the absence of a credible instrument for star recruitment, we will be cautious in inferring causality from any observed star effects on performance at the department level. However, the use of CEM allows us to make stronger causal inferences at the individual incumbent level.

2.2.2 Event-study specification with homogenous arrival effects across arrival cohorts and binning

The cumulative dynamic multipliers of interest (and their standard errors) can be estimated directly by reparametrizing (2.2) as an event specification where the *event*¹⁰ of a star arrival can take place in the past or in the future¹¹:

$$\begin{aligned} \ln Y_{i,t} &= \alpha + \sum_{j=-\infty}^{\infty} \left(\left(\sum_{j'=-\infty}^j \gamma_{j'} \right) \Delta S_{i,t-j} \right) + X'_{i,t} \lambda + \phi d_i + \varphi d_t + u_{i,t} \\ &= \alpha + \sum_{j=-\infty}^{\infty} \beta_j \Delta S_{i,t-j} + X'_{i,t} \lambda + \phi d_i + \varphi d_t + u_{i,t}. \end{aligned} \quad (2.3)$$

To empirically implement this framework we need to make assumptions about the effects of star arrivals when the arrivals take place in both the distant past and the distant future. We assume that the cumulative effect is constant at β_{-j} for all years up to \underline{j} years up to and including the arrival of the star and constant at $\beta_{\bar{j}}$ from \bar{j} years after the arrival. With this "binning" of effects, we can rewrite (2.3) as:

$$\ln Y_{i,t} = \alpha + \beta_{-j} \sum_{j=-\infty}^{-j} \Delta S_{i,t-j} + \sum_{j=-\underline{j}+1}^{\bar{j}-1} \beta_j \Delta S_{i,t-j} + \beta_{\bar{j}} \sum_{j=\bar{j}}^{\infty} \Delta S_{i,t-j} + X'_{i,t} \lambda + \phi d_i + u_{i,t} \quad (2.4)$$

Moreover, using our assumptions that no department is initially treated and all departments are eventually treated and once treated stay treated, we can use the fact that

¹⁰ There is a rapidly growing literature applying and developing panel event studies. Influential early studies include Jacobson (1993), Autor (2003) and Stevenson and Wolfers (2006).

¹¹ See Schmidheiny and Siegloch (2020) for a discussion of relationship between an infinite distributed lead/lag model and an event-study specification.

after cancelling terms in the first and third summations of changes these summations can be written simply as:

$$\sum_{j=-\infty}^{-j} \Delta S_{i,t-j} = S_{i,t+\infty} - S_{i,t+\underline{j}-1} = 1 - S_{i,t+\underline{j}-1}, \quad (2.5)$$

and

$$\sum_{j=\bar{j}}^{\infty} \Delta S_{i,t-j} = S_{i,t-\bar{j}} - S_{i,t-\infty-1} = S_{i,t-\bar{j}}. \quad (2.6)$$

This allows us to simplify (2.4) as:

$$\ln Y_{i,t} = \alpha + \beta_{-\underline{j}}(1 - S_{i,t+\underline{j}-1}) + \sum_{j=-\underline{j}+1}^{\bar{j}-1} \beta_j \Delta S_{i,t-j} + \beta_{\bar{j}} S_{i,t-\bar{j}} + X'_{i,t} \lambda + \phi d_i + \varphi d_t + u_{i,t}, \quad (2.7)$$

For our empirical implementation, we assume that binning occurs at five leads and five lags, so that $\underline{j} = \bar{j} = 5$. Normalising β_{-1} to zero, our estimation equation becomes¹²:

$$\begin{aligned} \ln Y_{i,t} = & \alpha + \beta_{-5}(1 - S_{i,t+4}) + \sum_{j=-4}^{-2} \beta_j \Delta S_{i,t-j} + \beta_0 \Delta S_{i,t} + \sum_{j=1}^4 \beta_j \Delta S_{i,t-j} \\ & + \beta_5 S_{i,t-5} + X'_{i,t} \lambda + \phi d_i + \varphi d_t + u_{i,t} \end{aligned} \quad (2.8)$$

We will present our results with event graphs that show the proportionate effect of the event of a star arrival from four years before the arrival to four years afterwards – i.e. a plot of $\hat{\beta}_{-4}$ to $\hat{\beta}_4$ with associated 95 percent confidence intervals. The binning variables play important roles as controls, but given their level form we are concerned they may be correlated with other excluded level variables and thus may not be strictly comparable to the leads and lags of the arrival dummies in the estimation of relevant dynamic effects. Statistical inference will be based on robust standard errors clustered at the department level to allow for arbitrary forms of serial correlation and heteroscedasticity (Kripfranz, 2016).

¹² Freyaldenhoven et al. (2019) employ a similar specification for various applications but do not derive it formally in the paper.

2.2.3 Heterogeneous star arrival effects by arrival cohort

Our baseline model has assumed star arrival effects that are homogenous across the timing of arrivals. An active recent literature has questioned the causal interpretation of estimates of coefficients such as the β_j coefficients in equation (2.8) in two-way fixed effects models (Athey and Imbens, 2018; Borusyak and Jaravel, 2017; Callaway and Sant'Anna, 2021; de Chaismartin and D'Haultfoeuille, 2020; Goodman-Bacon, 2021; and Sun and Abraham, 2021). A major concern is that in settings with variation in the timing of treatments – star arrivals in our case – and heterogeneous effects across different timings the coefficients on a given lead or lag can be contaminated by effects from other periods (see, e.g., Sun and Abraham, 2021). Using the approach of Sun and Abraham (2021), we therefore allow specifically for staggered treatments and our dynamic star arrival effects are derived as an appropriately weighted average of "cohort average treatment effects on the treated" (CATT). This approach has the advantage of being relatively easy to implement in a regression framework and allows straightforward comparisons to our baseline results.

In allowing for staggered star arrival effects we first define E_i as the year in which an arrival in department i occurs. The appropriate weighted average of cohort effects for a given time relative to the arrival event, j – or Interaction-Weighted (IW) estimator – is then obtained using a three step procedure.

First, in place of equation (2.8), we run a regression that allows for estimated arrival effects to vary based on the year that the arrival event occurs:

$$\begin{aligned} \ln Y_{i,t} = \sum_e \left[\delta_{e,-5} \left((\mathbf{1}\{E_i = e\}(1 - S_{i,t+4})) \right) + \sum_{j=-4}^{-2} \delta_{e,j} (\mathbf{1}\{E_i = e\} \Delta S_{i,t-j}) \right. \\ \left. + \sum_{j=0}^4 \delta_{e,j} (\mathbf{1}\{E_i = e\} \Delta S_{i,t-j}) + \delta_{e,5} (\mathbf{1}\{E_i = e\} S_{i,t+5}) \right] \\ + X'_{i,t} \lambda + \phi d_i + \varphi d_t + u_{i,t} \end{aligned} \quad (2.9)$$

where $\mathbf{1}\{E_i = e\}$ is an indicator variable that takes the value 1 if the arrival to a department i takes place in year e and 0 otherwise, and $\delta_{e,j}$ is the star arrival effect j years after an arrival that occurs in year e . Our observation window for the analysis is

from 1996 to 2017 and there are 19 treated cohorts in our sample after the first cohort (1996) is excluded from estimation since it is always treated across the observation window. Second, we estimate a set of weights $Pr\{E_i = e | E_i \in [-j, T - j]\}$ that are equal to the sample shares of each cohort for the relevant periods j . And third, to obtain the IW estimator, we take a weighted average of the $\hat{\delta}_{e,j}$ (or $CATT_{e,j}$) estimates from the first step with the relevant weights from the second step:

$$\beta_j^* = \sum_e [\hat{\delta}_{e,j} Pr\{E_i = e | E_i \in [-j, T - j]\}] \quad (2.10)$$

To allow comparison between the homogenous and heterogeneous effects models, we also exclude the 1996 arrival cohort from our homogenous effects model. This exclusion has minimal effects on the results.

2.3 Data

Our data consists of all publications (and citations to those publications at collection) recorded in Scopus in 27 subject fields for the period 1990 to 2017 where at least one author's affiliation is recorded as being in Denmark, Ireland or New Zealand. We chose Denmark and Ireland since both these countries are small open economies that have and are engaged in nationally supported formal programmes for the recruitment of star researchers. Then we selected New Zealand which is a small open economy with no record of any formal star recruitment programme. For each collected publication, we record the year, authors, affiliations, subject field¹³, citation count to 2019, references, abstracts and acknowledgements, as well as, importantly, two unique identifiers – the EID Scopus publication identifier and the Scopus Author(s) ID. This procedure gives us approximately 1.43 million publications divided over 219,582 unique authors. For each author, we also access their catalogue of publications back to 1990, including publications that took place prior to joining the department in Denmark, Ireland or New Zealand. We

¹³The 27 subject fields are: Agricultural and Biological Sciences, Arts and Humanities, Biochemistry, Business, Chemical Engineering, Chemistry, Computer Science, Decision Science, Dentistry, Earth and Planetary Sciences, Economics, Energy, Engineering, Environmental Science, Health Professions, Immunology and Microbiology, Multidisciplinary, Materials Science, Mathematics, Medicine, Neuroscience, Nursing, Pharmacology, Physics and Astronomy, Psychology, Social Sciences, and Veterinary.

identify a "department" with Scopus-defined subject field at a given institution so our departments need not coincide with the departmental structure at particular institutions. However, in some institutions there was no activity in the subject area and these "departments" were dropped from our dataset where there was not at least one cited publication in each sample year. Our data is spread across 457 remaining departments, comprising of 153 departments in Denmark, 141 in Ireland and 163 in New Zealand.¹⁴

2.3.1 Identification of stars

The first step is the identification of star scientists. We define a star scientist as a scientist who is at or above the 95th percentile of scientists in the cumulative distribution of citations received since 1990 for their subject field in any given year. To allow reasonable time for citations to publications to accumulate in our data, we only identify stars from 1996 onwards in our analysis. We allocate a star to a subject area (department) based on where they have the largest number of their publications. Where stars have an equal plurality of publications across multiple subject areas, then the star is assigned to multiple departments. It is also important to note that as our data only relates to scientists in Denmark, Ireland and New Zealand, our identification of stardom is relative to scientists in the subject field for these countries only. However, this procedure allows us to identify high performing scientists for this reference group.

A natural concern is that this relative measure may not identify true international stars. As part of our robustness analysis we therefore also explore the effects of augmenting our identification of star arrivals with an additional (variable) filter that requires that the star's Scopus Field-Weighted Citation Impact (FWCI) is above a threshold. The FWCI weights a scientist's citation score by the international average for the field. We compare the size of the star-arrival effect after four years based on different thresholds for this index. Although setting a higher threshold reduces the number of identified star arrivals and thus the precision of the estimated star arrival effects, an examination of how the size of the estimated effects varies with the threshold allows us to see how the effect size varies as we become more demanding in terms of this measure of international stardom.

¹⁴ See Appendix Table A.1 for additional detail.

Table 2.1 Descriptive statistics

	Total Sample 1996-2017			Never-Treated and Ever-Treated 1997-2017			
	Mean (1)	S.D (2)	N (3)	Never- Treated (4)	Ever- Treated (5)	p-value (6)	% Treated (7)
<i>Unit of analysis: department-year</i>							
Publications (log)	3.603	1.346	10,054	3.346	4.794	<0.001	17.7%
Field normalized (log)	3.406	1.582	10,040	3.108	4.787	<0.001	17.7%
Euclidean (log)	3.487	1.764	10,041	3.160	5.000	<0.001	17.7%
Incumbents' Publications (log)	1.258	1.264	6,781	0.962	2.387	<0.001	20.7%
Incumbents' Field normalised (log)	0.879	1.683	6,720	0.522	2.226	<0.001	20.9%
Incumbents' Euclidean (log)	0.957	1.872	6,742	0.581	2.382	<0.001	20.9%
University output (log)	7.079	1.137	10,054	6.968	7.588	<0.001	17.7%

Notes: The percentage of observations for "ever-treated" groups are reported in column (7) (in % of total observations 1996-2017). Column (6) presents the p-value for the null hypothesis that there is no significant difference between the two groups.

2.3.2 Identification of star arrivals

We identify a star arrival at a department as a star that had not previously published with an affiliation to the department now recording an affiliation to that department. We record the year of the first such star arrival (if any) to each department. For our empirical analysis, we restrict identified star arrivals to those who publish with an affiliation to the department for at least the next four years, or for those that publish in 2015-2016, they must also be observed to publish at the same affiliation at the end of our observation window in 2017.¹⁵ Overall, we identify 167 star arrivals over the period 1996 to 2017 in total using this procedure. As our empirical implementation uses only first star arrivals to a department, we record 81 such first star arrivals, excluding the 1996 cohort, across the three countries – 38 in Denmark, 19 in Ireland and 24 in New Zealand.¹⁶

2.3.3 Output measures

Our output measures are at the department-year or individual incumbent level depending on the specification. We consider three output measures: (i) total count of publications in the department/incumbent year; (ii) field-normalised total citations (the sum across department/incumbent publications of citations divided by the average

¹⁵ In terms of the same first star arrival at multiple departments, Columns 5 and 6 in Appendix Table A.2 present such occurrences. Specifically, there are 6 instances of the same individual star being allocated to two separate subject fields in this analysis. Additionally, there are 3 instances where more than one star first arrived at the same department in a given year (see Column 4, Appendix Table A.2); however, it is worth highlighting that our estimation results are robust to both their inclusion and exclusion as treated departments.

¹⁶ Appendix Figure A.1 shows the distribution across time of first star arrivals in aggregate and for each of the three countries. The year with the highest number of departments with a first star arrival is 1998 with 9 stars first arriving across the three countries.

citations to a publication for that field in that year);¹⁷ and (iii) a scaled Perry-Remy Euclidean Index (Perry and Remy Perry, 2016) that places greater weight on highly cited publications than measure (ii).¹⁸

More formally, the three output measures are defined as follows:

Publications:
$$Y_{i,t}^P = P_{i,t};$$

Field Normalized Total Citations:
$$Y_{i,t}^{FNTC} = \sum_{p_{i,t}=1}^{P_{i,t}} \frac{c_{p_{i,t}}}{\bar{c}_{s,t}};$$

Perry – Remy Euclidean Index:
$$Y_{i,t}^E = \left(\sum_{p_{i,t}=1}^{P_{i,t}} \left(\frac{c_{p_{i,t}}}{\bar{c}_{s,t}} \right)^2 \right)^{\frac{1}{2}};$$

where $P_{i,t}$ is the total number of publications in department i or incumbent i in year t , $c_{p_{i,t}}$ is the total subsequent citations (or "forward" citations recorded at 2019) to a publication $p_{i,t}$ that occurs in department i in year t , and $\bar{c}_{s,t}$ is the average citations to a publication in the relevant subject field, s , for publications that occur in year t .

These output measures provide the dependent variables for our regressions. Our estimation period is selected as 1996 to 2017, which also allows for the inclusion of the chosen number of leads and lags ($\underline{j} = \bar{j} = 5$) for all post-1996 star arrivals. Descriptive statistics for the department-level output measures are shown in Table 2.1 (columns 1-3). Table 2.1 columns 4-7 split departments into "never-treated" and "ever-treated" subgroups where the treatment refers to the first arrival of a star. The two groups are found to be different with the treated departments having larger outcome values on all dimensions. In addition to controlling for time-invariant department characteristics using department fixed effects, we also control for department-specific linear time trends and a time-varying university-level control (ex the focal department). Pre-arrival trends

¹⁷ See Radicchi, Fortunato and Castellano (2008).

¹⁸ See Perry and Remy (2016).

are examined to check for violations of the parallel trends assumption in our estimated specifications.

Table 2.2 Summary statistics: control and treated subfields at baseline (k-to-k matched)

Variable	Control	Treated	Diff. in mean	P-value
<i>Panel A: Unbalanced panel on 16,730 matched authors, with 8,365 in each group</i>				
Citation received	34.742	35.221	-0.478	0.6988
Career age (Bins)	2.209	2.209	0.000	1.000
Cumulated citation received (Bins)	3.455	3.455	0.000	1.000
Career age (Num.)	8.907	8.876	0.031	0.7807
Cumulated citation received (Num.)	350.902	345.197	5.704	0.7149
<i>Panel B: Balanced panel on 4,696 matched authors, 2,348 in each group</i>				
Citation received	52.029	53.689	-1.659	0.5997
Career age (Bins)	2.892	2.892	0.000	1.000
Cumulated citation received (Bins)	5.621	5.621	0.000	1.000
Career age (Num.)	12.451	12.408	0.042	0.8537
Cumulated citation received (Num.)	670.287	675.770	-5.482	0.8916

Notes: We report the t-test for mean differences between the treated and control groups at one year before the star arrival for these time-varying covariates. For career age, we create thirteen bins: less than 5 years; between 5 and 15 years; between 15 and 20 years; between 20 and 25 years; between 25 and 30 years; between 30 and 45 years; between 45 and 50 years; between 50 and 55 years; between 55 and 60 years; and over 60 years of career age. Similarly, we coarsen the distribution of cumulated citations at baseline 26 mutually exclusive bins; between one and 100 citations; between 200 and 300 citations; between 300 and 400 citations; between 400 and 500 citations; between 5,000 and 10,000 citations; and larger than 10,000 citations.

2.3.4 Coarsened exact matching procedure (CEM) for incumbent sample

For the matched individual-level analysis, we use the following CEM procedure to identify appropriate matches. First, we identify all incumbents in the department in the year, t , that the star arrives. (Incumbents are required to be present four years before the star's arrival and four years after the arrival.) For each identified incumbent in a star-receiving department, we identify an appropriate match based on covariates measured in $t - 1$. The following covariates are used to identify the match: scientist career age (i.e. years since first publication); citations received by the scientist in $t - 1$; cumulative citations received by the scientist up to $t - 1$; subject field; country; and year. We utilize relatively coarse bins for career age (13 bins) and for cumulative citations (26 bins), but require an exact match for citations received in $t - 1$, subject field, country and year. Our matching procedure successfully matches 16,730 incumbent scientists. The resulting panel is unbalanced as all scientists will not be present in the department over the entire 22-year sample period. As a robustness check we repeat the analysis for a balanced panel of 4,696 incumbents whom were present at a star-receiving department over the entire sample period. Table 2.2 provides summary statistics for the matching process as applied

to both the unbalanced and balanced panels. Overall, the application of CEM is successful in identifying good matches for our sample of incumbents.

2.4 Results

2.4.1 Star arrival effects on department output

Our star arrival results for under both the homogenous and heterogeneous arrival effects by cohort are shown in Table 2.3. The department output variable is field-normalised total citations to the output published in year t . (We compare the results with other output measures as part of robustness tests in Section 2.5.) We present two specifications to allow the reader to compare the homogenous and heterogeneous models with and without the inclusion of control variables. The control variables are a set of department specific time trends and university-level output excluding the focal department.

Column 1 of Table 2.3 shows the results for the homogenous model without controls. The star arrival effects are also shown in the top left panel of Figure 2.1. While the results show some visual evidence of a pre-trend, we see evidence of economically and statistically significant star arrival effects, with a star arrival being associated with a 12.2 log points (12.9 percent) increase in output after four years. Column 2 of Table 2.3 and the top right panel of Figure 2.1 show the results with the controls included. With the inclusion of the controls we do not see evidence of a pre-trend and the estimated arrival effects are larger, reaching 18.9 log points (20.8 percent) after four years. We treat this specification as our baseline in conducting our robustness tests below. These results are also robust to the separate inclusion of the set of department-specific time trends and the university-level controls.

As our parallel trends assumption is critical to the identification of a casual effect of star arrivals, it is important to validate the assumption for our data. A common method to assess the credibility of the parallel trends assumption is to test for significant pre-trends in the event study setting; however, a recent body of literature has raised important concerns about this test for pre-trends. In general, it warns that the test may have low power to detect meaningful violations of parallel trends (Bilinski and Hatfield,

2020; Freyaldenhoven et al., 2019; Kahn-Lang and Lang, 2018; and Roth, 2021). Using an approach proposed by Roth (2021), we present diagnostics to evaluate whether the limitations of pre-trends testing are an important concern for our analysis.

Table 2.3 Dynamic star arrival effects for FNTC (department level)

	Homogenous Arrival Effects		Heterogeneous Arrival Effects	
	No controls	With Controls	No controls	With Controls
	(1)	(2)	(3)	(4)
$1 - S_{i,t+4}$ (Binned lead)	-0.176 (0.118)	-0.128 (0.106)	-0.142 (0.110)	-0.118 (0.108)
$\Delta S_{i,t+4}$	-0.0718 (0.0808)	-0.00404 (0.0881)	-0.047 (0.085)	-0.092 (0.090)
$\Delta S_{i,t+3}$	-0.0739 (0.0761)	-0.0426 (0.0795)	-0.055 (0.072)	-0.082 (0.095)
$\Delta S_{i,t+2}$	-0.0736 (0.0644)	-0.0295 (0.0648)	-0.046 (0.055)	-0.042 (0.069)
$\Delta S_{i,t}$	0.0507 (0.0619)	0.0634 (0.0641)	0.044 (0.054)	0.067 (0.058)
$\Delta S_{i,t-1}$	0.115* (0.0589)	0.138** (0.0637)	0.101** (0.050)	0.163*** (0.062)
$\Delta S_{i,t-2}$	0.111* (0.0638)	0.153** (0.0679)	0.104* (0.054)	0.197*** (0.068)
$\Delta S_{i,t-3}$	0.113 (0.0722)	0.166** (0.0777)	0.104* (0.060)	0.181** (0.071)
$\Delta S_{i,t-4}$	0.122 (0.0743)	0.189** (0.0823)	0.112* (0.067)	0.220*** (0.076)
$S_{i,t-5}$ (Binned lag)	-0.0360 (0.0934)	0.226** (0.0965)	-0.013 (0.089)	0.241** (0.098)
Log University Output (excl. dept.)		0.857*** (0.0701)		0.838*** (0.071)
R-squared	0.4337	0.6561	-	-
Pre-test against hypothesized trend (Roth 2021)				
Power	0.50	0.50	-	-
Hypothesized trend	0.05	0.05	-	-
Bayes factor	0.56	0.56	-	-
Likelihood ratio	0.40	0.09	-	-
Department FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Department time trend	No	Yes	No	Yes
Observations	10,040	10,040	10,040	10,040

Event Window: 1997-2017 (1996 cohort dropped). Estimation Window: 1996-2017.

Dependent variable is Field Normalised Total Citations.

Notes: This table reports the estimates based on the model specification in equation 2.8 in Columns 1 and 2. The dependent variable always excludes the star's own output. Robust standard errors are clustered at department-level and reported in parentheses. Columns 3 and 4 are the estimates based on the Sun and Abraham (2020) method (equation 2.10). *, **, and *** represent significance levels at the 10%, 5% and 1% respectively.

Specifically, we implement a power calculations analysis to construct a hypothesized linear violation of parallel trends and then compare the likelihood of the estimated pre-trend coefficients under both the hypothesized trend and parallel trends. The results from this analysis for a relevant violation of parallel trends are also reported in Table 2.3. For our analysis, the slope of the hypothesized trend is constructed so that a

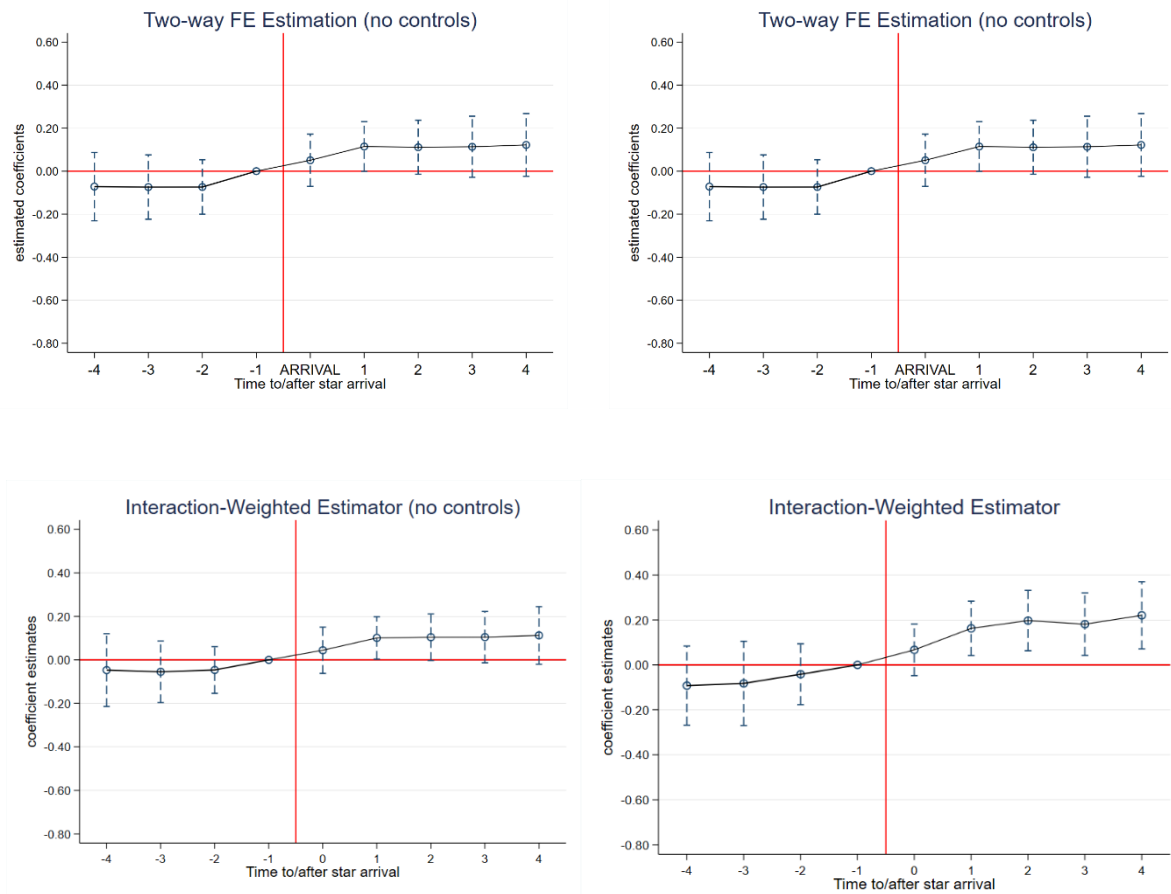


Figure 2.1 Event study model with the homogeneous (top) vs. heterogeneous (bottom) arrival effects across cohorts (department level)

Notes: This figure plots the dynamic effect of a first star arrival with 95% confidence intervals. The event and observation window is from 1996-2017 (the 1996 cohort is omitted). The dependent variable is the field-normalised total citations and excludes the output of the star arrival. First star arrivals are restricted to stars staying in the department for four years or more. The omitted category is one year before the first star arrival. University-level controls (excluding the focal department) and department time-specific trend are included in the model specification on which the right-hand plot is based.

pre-trend test based on the statistical significance of the pre-trend coefficients would detect a significant pre-trend 50% of the time. This slope estimate is 0.05 both with and without the inclusion of controls. The likelihood ratio which reports the ratio of the likelihood of the observed coefficients under the hypothesized trend relative to under parallel trends suggests that, for both specifications, we are much more likely to observe our estimated coefficients under parallel trends rather than the hypothesized trend with 50% power in detecting a pre-trend, which gives us confidence in the parallel trends assumption. It is also worth highlighting that the likelihood ratio for the model that identifies a star arrival based on the baseline criteria decreases from 0.40 to 0.09 when the University level controls and the department time-specific trend are included in the specification, underlining the importance of including these controls in our preferred specification.

Column 3 of Table 2.3 (and the bottom left panel of Figure 2.1) records the estimated star arrival effects using the Sun and Abraham methodology without the controls. We now see little evidence of a pre-trend and the estimated star arrival effects are quantitatively similar to the homogenous effects model results, with an 11.2 log point (11.8 percent) increase in department output four years after the arrival of the star. As noted above, we provide for comparison purposes the results of the heterogeneous effects model with controls in Column 4 (and bottom right panel of Figure 2.1). This specification again shows little evidence of a pre-trend and the largest star arrival effects of the specification shown, with a star-arrival effect after four years of 22.0 log points (24.6 percent). Overall, the results appear reasonably robust to the choice of specification with star arrival effects four years after the arrival of the star of between approximately 11 and 22 log points (or 11.8 and 24.6 percent).

One possible concern with our method of star identification is that it applies a relative standard – i.e. where the scientist fall in the distribution of total (eventual) citations to their published work in their respective fields across the three small open economies. It is reasonable to argue that star recruitment policies aim to recruit true international star, where the relevant distribution would be the full international distribution in that field. As an additional robustness test we therefore add an additional threshold (or filter) to our identification method based the Scopus SciVal reported Field Weighted Citation Impact (FWCI). For each scientist, Scopus measures their citations relative to the international average for the field, so that a scientist with a FWCI of 1 would exactly match the international average.

Figure 2.2 (A) shows the evolution of the star arrival effect measured four years post arrival based on different thresholds for the FWCI. The figure shows that the number of identified stars does indeed fall as we move to higher thresholds, indicating that a number of our relative stars appear less outstanding from the truly international

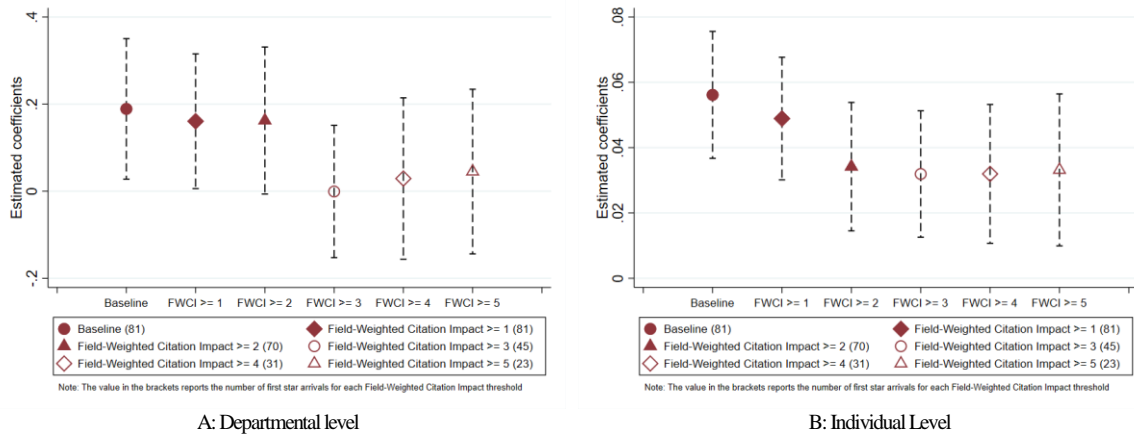


Figure 2.2 Size of estimated star arrival effect four years post star arrival based on alternative thresholds for the field.

Notes: The graph shows the star arrival effects on the departmental level (Panel A) and individual level (panel B) for the FWCI, which compares the entity's number of citation received for a publication compared with the world average. The FWCI is calculated as an average 4 years prior to the star arrival year (1996 onwards, given Scival's limitation) from Scival metrics. While the number of star arrivals are the same between the baseline and the group $FWCI \geq 1$ (i.e., 81 star arrivals), the groups are marginally different since the threshold restriction ($FWCI \geq 1$) omits a small number of first star arrivals identified under the baseline criteria, yet includes a number of star arrivals otherwise not identified under the baseline.

perspective. While care must be taken in interpreting the results because of the falling number of identified stars as we adopt a stricter threshold, it is notable that the estimated star arrival effect is smaller at the higher thresholds, with an apparent step drop in the estimated effect at a FWCI of ≥ 3 . These results suggest that the largest star arrival effects may not occur for the most accomplished international stars, possibly because such stars embed less deeply in local networks. This result would also appear consistent with the result reported below that star arrivals effects are generally larger for earlier career stars (career age ≤ 20), whom are also likely to rank lower in the international distribution. In terms of policy design, this underlines the importance of identifying stars that have a strong incentive to embed in local networks, which may depend, inter alia, on their career horizons and the strength of their (possibly competing) existing international networks.

For our star arrivals, one interesting split is between stars that arrive relatively early in their careers and those that arrive later. We measure career age as the number of years since the scientist recorded their first publication, and split arriving stars into "earlier" career age stars (career age ≤ 20) and "later" career age stars (career age > 20). Due to space limitations the results are shown in Appendix A.3. Appendix Figure A.4 indicates that the star arrival effect on department productivity is large for the earlier career arrivals. As discussed further in Section 2.5, this may reflect a greater willingness for scientists with longer career horizons at their new institution to embed into local

scientist networks leading to more pronounced catalytic effects on department performance.

Table 2.4 Dynamic star arrival effects for FNTC (individual level)

	Homogenous Arrival Effects		Heterogeneous Arrival Effects	
	Balanced Panel (1)	Unbalanced Panel (2)	Balanced Panel (3)	Unbalanced Panel (4)
$1 - S_{i,t+4}$ (Binned lead)	0.003 (0.020)	0.007 (0.009)	-0.020 (0.027)	-0.008 (0.011)
$\Delta S_{i,t+4}$	-0.007 (0.019)	0.008 (0.009)	-0.009 (0.021)	0.006 (0.010)
$\Delta S_{i,t+3}$	-0.006 (0.018)	0.008 (0.008)	-0.006 (0.019)	0.009 (0.009)
$\Delta S_{i,t+2}$	0.021 (0.016)	0.008 (0.007)	-0.014 (0.017)	0.005 (0.008)
$\Delta S_{i,t}$	0.043*** (0.015)	0.014** (0.007)	0.043*** (0.015)	0.015** (0.008)
$\Delta S_{i,t-1}$	0.051*** (0.016)	0.029*** (0.007)	0.053*** (0.017)	0.035*** (0.008)
$\Delta S_{i,t-2}$	0.056*** (0.017)	0.042*** (0.008)	0.065*** (0.018)	0.050*** (0.009)
$\Delta S_{i,t-3}$	0.064*** (0.018)	0.045*** (0.009)	0.073*** (0.018)	0.052*** (0.009)
$\Delta S_{i,t-4}$	0.066*** (0.018)	0.056*** (0.009)	0.075*** (0.019)	0.063*** (0.010)
$S_{i,t-5}$ (Binned lag)	0.108*** (0.018)	0.070*** (0.009)	0.119*** (0.019)	0.080*** (0.011)
R-squared	0.0473	0.0347	-	-
Pre-test against hypothesized trend (Roth 2021)				
Power	0.50	0.50	-	-
Hypothesized trend	0.01	0.005	-	-
Bayes factor	0.58	0.57	-	-
Likelihood ratio	0.57	0.04	-	-
Department FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Department time trend	No	No	No	No
Observations	98,978	260,629	98,978	260,629

Event Window: 1997-2017 (1996 cohort dropped). Estimation Window: 1996-2017.

Dependent Variable is Field-Normalised Total Citations

Notes: This table reports the estimates based on the model specification in equation 2.8 in Columns 1 and 2. The dependent variable always excludes the star's own output. Robust standard errors are clustered at individual-level and reported in parentheses. Columns 3 and 4 are the estimates based on the Sun and Abraham (2021) method (equation 2.10). *, **, and *** represent significance levels at the 10%, 5% and 1% respectively

One concern is that the results for the aggregate of the three countries may hide different dynamic patterns at the country level. The results are shown in Appendix A.4. Appendix Figure A.6 shows the results of our baseline regression with inclusion of the controls for the three countries separately. The broad pattern for Denmark mirrors those found for the aggregate results. However, the patterns for Ireland and New Zealand are notably different with no evidence of star arrival effects for Ireland and a notable positive pre-trend visible and at best limited evidence of post-arrival star effects for New Zealand.

While care must be taken in over interpreting the country-specific results due to the relatively small number of treatments, the differences across countries suggests the importance of country and institution-specific policy design. However, as we report below, at the individual incumbent level there is a more consistent pattern of star arrival effects across the three countries.




Figure 2.3 Event study model for matched incumbents (unbalanced panel)

Notes: This figure plots the dynamic effect of a first star arrival with 95% confidence intervals based on the matched sample, with 16,351 (379 treated author in 1996 cohort removed) authors. It consists of 260,629 observations. The event and observation window is from 1996-2017. The left panel shows the homogenous effects results; the right panel shows the heterogeneous effects results based on Sun and Abraham (2021). The dependent variable is the field-normalised total citations in logarithm (with 1 added across all observations). The omitted category is one year before the first star arrival. Robust standard errors are clustered by author.

2.4.2 Star arrival effects on incumbent output

Overall, our department-level results point to a significant effect of star arrivals on subsequent department productivity performance. Moreover, the event-study results appear robust to a misspecification of the pre-trend and allowance for heterogeneity in the arrival effect across arrival cohorts. However, there remains a concern that observed arrivals are not independent of productivity developments that occur contemporaneous with or subsequent to the star arrival. For the reasons noted earlier, we therefore augment our department-level analysis with an individual-level analysis of the effect of star arrivals on matched incumbents. We require incumbents to be present in the department four years before the arrival of the star and also present in the four years following the arrival. The use of matching allows us to give a firmer causal interpretations to the observed star arrival effects.

Table 2.4 records the results for the individual-level incumbent regressions. Column (1) and (2) show the results for the balanced and unbalanced panels under the

assumption of homogenous effects by arrival cohorts. In discussing the results, we concentrate on the unbalanced panel as the long-stay incumbents in the balanced panel are a selected subgroup that are more likely to be older and individuals who have never moved institutions.

Figure 2.3 shows the event-study graph for a sample of 16,730 matched incumbents in the unbalanced panel, where the dependent variable is an individual's field-normalised total citations to their publications in a given year. Visual inspection and tests for significance of the pre-trend coefficients reveals no evidence of a pre-trend, and star arrival effects are statistically significant after one year and grow over time. The star-arrival effect is 5.6 log points (5.8 percent) after four years.

Paralleling our approach to the department-level analysis, we also apply the Sun and Abraham (2021) methodology to allow for heterogeneous effects by arrival cohort. Column (4) of Table 2.4 shows the results for the unbalanced panel. For ease of visual inspection, Figure 2.3 (on the right) graphs the results for the unbalanced panel under the assumption of heterogeneous effects by arrival cohort. We again find our results are robust to the allowance for such effects.

Table 2.4 also reports the Roth (2021) test for the sensitivity to the violation of the parallel trends assumption. For the unbalanced panel with homogenous effects, the slope of the hypothesised trend required for finding a significant pre-trend 50 percent of the time is 0.005. Not surprisingly given the absence of visual evidence of a pre-trend, the likelihood ratio of the observed pre-trend coefficients under the hypothesised trend relative to under parallel trends is low at 0.04. This gives us confidence that the parallel trends is supported in our individual-level data after matching.

As we did for the department-level results, Figure 2.2 (B) shows the implications of applying an increasingly strict filter to our identification of a star based on the SciVal FWCI measure. The results across different thresholds are similar to what we observe at the departmental level, with evidence of smaller star arrival effects as we apply an increasingly stringent threshold for being identified as a true international star.

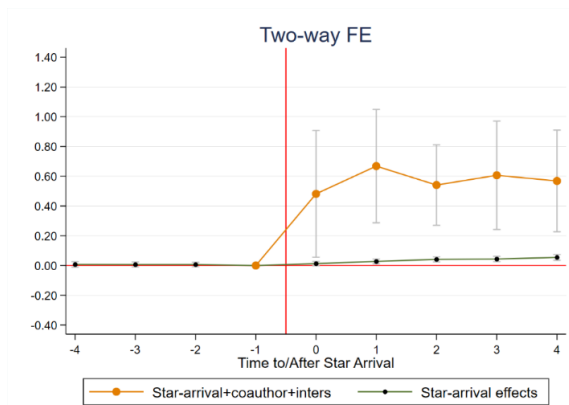


Figure 2.4 Event study model for matched incumbents with star co-authoring effects (binary)

Notes: This figure plots the dynamic effect of a first star arrival with 95% confidence intervals based on the matched sample, with 16,351 (379 treated author in 1996 cohort removed) authors. It consists of 260,629 observations. The event and observation window is from 1996-2017. The dependent variable is the field-normalised total citations in logarithm (with 1 added across all observations). The omitted category is one year before the first star arrival. Robust standard errors are clustered by author

Appendix Figure A.5 also confirms the qualitative picture from the department-level results that shows that the star-arrival effect is larger for stars arriving earlier in the careers, which again is suggestive of early-career arriving stars to embed in local networks when their horizons at the institution are longer.

Turning to the country-specific results, Appendix Figure A.7 shows a more consistent star arrival effect on incumbents than we found for the department level analysis. While standard errors are of course larger than for our three-country sample, we find evidence of a statistically significant star arrival after three years for each country and broadly similar evolutions over time. The size of the effect is also roughly similar across the three countries (≈ 5 percent) and consistent with the combined sample.

We can further increase confidence that we are identifying a causal effect of star arrivals if there is a clear mechanism through which arrivals are affecting individual productivity. One plausible mechanism is through opportunities to co-author with the star. A co-authorship relationship should indicate that the work of the incumbent and star are related, and also collaboration with the star should directly boost incumbent productivity. As expected, Figure 2.4 shows a large additional effect on incumbent productivity if the incumbent develops a co-authorship relationship with the star.




Figure 2.5 Related vs. unrelated

Notes: This figure plots the dynamic effect of a first star arrival with 95% confidence intervals based on the matched sample. The panel consists of 32 authors related to the star arrival (estimated 125,698 observations) and other unrelated authors with 259,916 estimated observations. The dependent variable is the field-normalised total citations in logarithm (with 1 added across all observations). The omitted category is one year before the first star arrival. Robust standard errors are clustered by author.

To strengthen our confidence in a causal interpretation of the star-arrival effect, a further hypothesis is that incumbent scientists who are working in areas related to the star should experience larger beneficial productivity impacts from the star’s arrival. Although there are many possible indicators of relatedness, an obvious possibility is to identify incumbents that have cited the star at some point in the sample. Using this indicator of "relatedness," Figure 2.5 compares the results for unrelated and related incumbents. The hypothesis that related incumbents benefit more from an arriving star is clearly supported, with related incumbents increasing their productivity after four years by approximately 40 percent.

2.5 Robustness

In this section, we conduct several robustness tests on our baseline model. We limit our reporting to the department-level analysis due spaces constraints, although we also find our incumbent-level analysis is robust to these alternative specifications. Specifically, we compare our results to the estimates from specifications with alternative output measures and alternative star identification methods. Further to this, we conduct two robustness tests in the appendix that accompanies this analysis: we explore a restricted specification that imposes a piecewise linear structure on the dynamic

effects¹⁹; and we test for miss-specified dynamics²⁰. Overall, the comparisons outlined below and in the appendix indicate that our results are generally robust across the different measures.

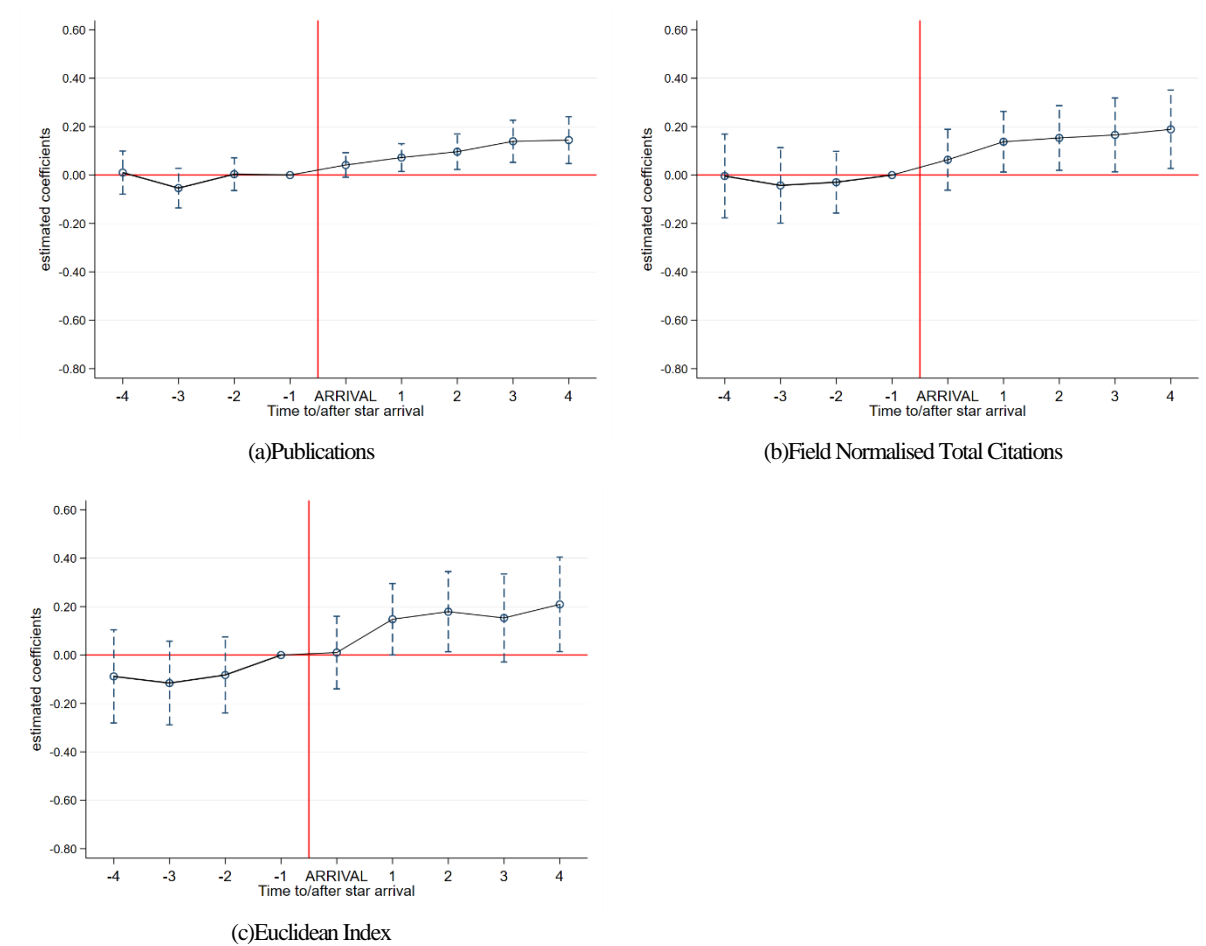


Figure 2.6 Event study model with homogenous arrival effects across cohorts for alternative outcome measures

Notes: This figure plots the dynamic effect of a first star arrival with 95% confidence intervals and assuming homogenous arrival effects across cohorts based on the specification in equation 2.8 for alternative outcome measures. The event and observation window is from 1996-2017 (the 1996 cohort is omitted). Each alternative outcome measure excludes the output of the star arrival. First star arrivals are restricted to stars staying in the department for four years or more. The omitted category is one year before the first star arrival. University-level controls (excluding the focal department) and department time-specific trend are included in the model specification. There are 10,054 observations in Publications measure and 1,782 are treated; 10,040 observations in Field Normalised Total Citations measure and 1,782 are treated; 10,041 observations in Euclidean Index measure and 1,782 are treated. See Tables 3a and 3b for more details.

2.5.1 Robustness to alternative outcome measures

Our baseline regression uses department-level field-normalised total citations as the dependent variable. This measure can be viewed as a quality-adjusted publications

¹⁹ Under specific conditions, the efficiency of our estimates can be improved by imposing appropriate restrictions on equation (2.8) - see Deryugina (2017). Appendix A.5 describes our restricted specification and presents the related results.

²⁰ Appendix A.6 describes our robustness test for miss-specified dynamics in equation (2.8) and reports the results.

measure of department output, where a publication with the average number of citations for the field gets a weight of 1, a publication with twice the average gets a weight of 2, etc. As a first test of the robustness of our results we examine the sensitivity of these results for the baseline specification of the model to alternative output measures: raw publication counts and the scaled Euclidean Index due to Perry and Remy (2016) that puts greater weight on highly cited publications relative to the field-normalised total citations measure. These measures therefore provide reasonable bounds in terms of the relative weighting of quantity versus quality.

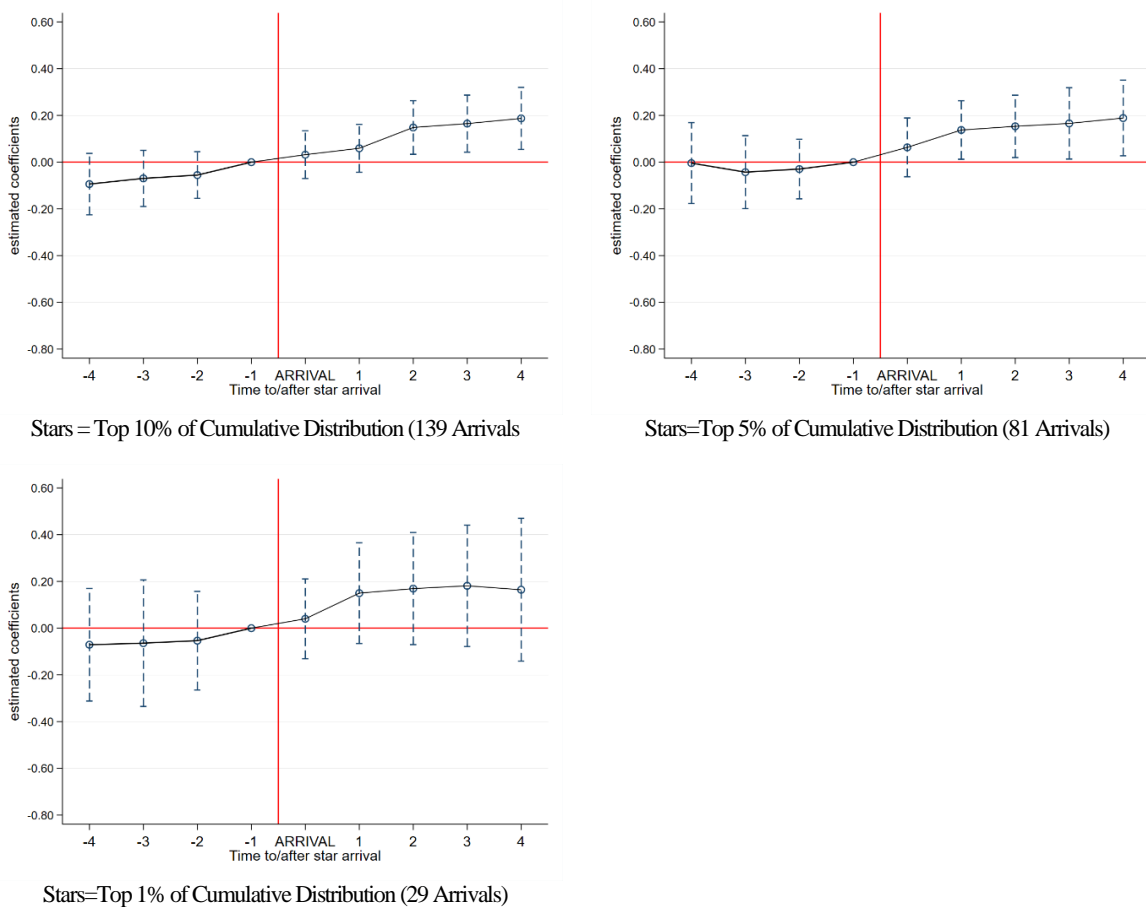


Figure 2.7 Event study model with homogenous arrival effects across cohorts for alternative star definitions

Notes: This figure plots the dynamic effect of a first star arrival with 95% confidence intervals and assuming homogenous arrival effects across cohorts based on the specification in equation 2.8 for alternative star definitions. The event and observation window is from 1996-2017 (the 1996 cohort is omitted). The dependent variable is the field-normalised total citations and excludes the output of the star arrival. First star arrivals are restricted to stars staying in the department for four years or more. The omitted category is one year before the first star arrival. University-level controls (excluding the focal department) and department time- specific trend are included in the model specification.

The results for the alternative measures are shown in Figure 2.6. The top right panel (b) repeats our baseline results (with controls) from Figure 2.1 for ease of comparison. The top left panel (a) shows the results with raw publication counts as the dependent variable. The results are broadly similar to the baseline but with a steadier

rise post arrival and a smaller positive impact on department output in the arrival year and at each lag. The bottom panel (c) shows the results for the Euclidean Index measure. Again, the results are broadly similar to the baseline with a similar dynamic pattern and somewhat larger post-arrival effects. At the fourth lag, the star arrival effect is 14.4 log points (15.4 percent) for raw publications, 18.9 log points (20.8 percent) for field-normalised total citations and 20.9 log points (23.2 percent) for the Euclidean Index. Thus, our results are broadly robust to the choice of output measure but with evidence of greater star-arrival impacts the more weight that is put on quality (as measured by subsequent citations to a publication) compared to publication quantity.

2.5.2 Robustness to alternative star identification methods

For our baseline results, we identified a star arrival as the arrival of a scientist in the top 5 percent of the relevant cumulative distribution of total citations to their work measured at the time of arrival. Figure 2.7 provides a comparison to our baseline results (middle panel) to a more permissive definition of a star arrival (top 10 percent) and a stricter definition of a star arrival (top 1 percent). Of course, the number of star arrivals will fall as we adopt a progressively stricter definition, which impacts the standard errors.

The results are again broadly similar across the alternatives, although the estimated star arrival effects are somewhat smaller for both the top 10 percent definition and the top 1 percent definition compared to our baseline. For example, four years after arrival, the impact on department output is 18.7 log points (10 percent definition), 18.8 log points (5 percent definition) and 16.4 log points (1 percent definition). The wider confidence intervals reflecting the smaller number of arrivals is also apparent in the bottom panel.

2.6 Conclusions

Our various event analyses – where the event is the arrival of a star – are supportive of the hypothesis that star recruitments have positive and sustained effect on productivity. Although formal tests support the parallel-trends assumption at the department level, we have been cautious in inferring causality at that unit of analysis. However, the employment of matching at the individual incumbent unit of analysis,

combined with strong support for the parallel trends assumption, gives us confidence in offering a causal interpretation of the observed star arrival effects. This interpretation is further strengthened by the finding of the expected larger effects when certain intermediating effects – such as the formation of local collaborations with the star – are present. Moreover, both sets of results are robust to explicitly allowing for heterogeneous effects by arrival cohort.

These findings have potentially important implications for the adoption and design of star recruitment policies. In general, we hypothesize that the beneficial effects of star arrivals depend on stars embedding in local networks. The results indicate that it is not necessarily the most accomplished stars that have the largest local impacts. In particular, we find that stars arriving relatively early in their careers have a larger positive effect on both departments and individual incumbents, possibly reflecting the effects of longer horizons and less developed existing networks on the willingness to make the investments necessary to embed in local networks.

While relatively small numbers of arrivals at the individual-country level means that country-specific results need to be interpreted with care, we did find notable differences in the size of the star arrival effect across our three countries. At the departmental level a clear star-arrival effect is only present for Denmark, although there is more consistent evidence of positive effects for individual incumbents. The differences across countries do point to the importance of the design details of recruitment strategies and post-arrival supports. The importance of these design questions is heightened by the large national-level investments made to support star recruitments in some countries. Although our study does not specifically focus on the relatively small number of stars recruited through these national programs, in ongoing related work we have interviewed stars recruited through these programs in Denmark and Ireland as well as relevant institutional stakeholders. The early results from this work point to stars successfully embedding in local networks. However, the findings also point to various challenges. Several categories of challenges were evident in both programmes. As examples, a number of recruits in Ireland identified challenges such as poorly planned post-arrival supports and overtly bureaucratic institutional support processes; whereas some recruits in Denmark identified cultural challenges in internationalising research focus and challenges posed by local institutional recruitment strategy and structure. With

cluster-forming star-recruitment policies likely to become an increasingly important component of the science-policy mix, it is essential to better understand the policy and institutional strategy designs that maximize the effectiveness of these investments.

Chapter 3 Star Help and Knowledge Transfer: An Event Study Analysis of Star Interactions Observed from Acknowledgement Texts

3.1 Introduction

Star scientists have long been recognised as catalysts for knowledge creation within academia. There is also a growing policy interest in smaller economies and regions in increasing their connections to academic star scientists as a means to catalyse dynamic local research clusters. Although explicit policy programmes have been implemented, there is limited understanding of the mechanisms through which star scientists have an effect on local scientific productivity. One candidate mechanism is the direct effect of help from star scientists to incumbent local scientists, even where an explicit co-authorship relationship does not exist. Existing studies have generally not focused on the impacts of star scientists in smaller economies and have tended to overlook modes of knowledge transfer beyond co-authorship. Moreover, the importance of network building, both within and across institutions, emerges as a strategy for facilitating access to star help and unlocking its potential for productivity enhancement. This research contributes insights into knowledge and technology transfer by providing new empirical evidence on the knowledge transfer mechanisms facilitated by star scientists, and offering actionable recommendations for improving academic and research management practices in small open economies. This chapter addresses these gaps with new empirical evidence on knowledge transfer that occurs through a star-help mechanism in small economies, where star help is indicated by acknowledgements to various forms of star help in the acknowledgement sections of published journal articles. To broaden the evidence base and recognise that developing connections to international star scientists has been a particular policy focus for smaller open economies (SOEs), our data are drawn from academic departments in Ireland, Denmark and New Zealand. As we use the term in this chapter, 'star scientist' refers to those individuals who are highly productive in terms of their publications and citations and are located chiefly at a university (Schiller and Diez, 2010; Zucker and Darby, 2006).

Scientific collaboration through formal (e.g., co-authorships) or informal means (e.g., discussions and helpful comments) connects a star with other scientists. These interactions could play an essential role in generating knowledge spillovers. This chapter focuses on one particular channel of star impact and aims to investigate the effects of the star's help on their colleagues' productivity. Helpful interaction – either with co-located or non-co-located scientists – is a relevant and usually neglected channel for learning and knowledge exchange in socially interactive environments. Interactions with stars, even if they are informal, influence behaviour, research topics, collaborations, and the productivity of scientists (Bramoullé et al., 2020). More frequent and long-lasting interactions impact the extent of knowledge overlap (Nooteboom, 2000), facilitate learning processes, and enhance knowledge creation through the recombination of existing ideas between stars and scientists (Nelson and Winter, 1982; Weitzman, 1998). The empirical challenge is that evidence of help does not leave a paper trail; however, explicit acknowledgements to help from a star provide a (noisy) indicator of star interaction (Bozeman and Corley, 2004; Oettl, 2012). We broadly interpret such acknowledgements as evidence of helpful interaction. From acknowledgements in the publications, we use natural language processing (NLP) techniques to identify the star scientist names that are being acknowledged. We then draw on recent advancements in econometric techniques to estimate the productivity impacts of star scientists' help on the scientific productivity of the acknowledging authors. We also explore the effects of different kinds of star help by categorising acknowledgements into five types: Conceptual, Technical, Material, Funds & Support, and Other Types.

The basic idea behind the empirical approach we take in the chapter is that help (in various forms) from star scientists can raise the productivity of the assisted scientists and that an indicator of the presence and form of such help can be gleaned from the acknowledgement texts of publications. Our core measure of productivity is (forward) citation-weighted publications in a given year. However, a key empirical challenge is that an observed correlation between acknowledgements and productivity might not reflect a causal productivity effect. An alternative possibility is that an acknowledgement leads to a signalling effect, whereby a publication with a star acknowledgement is more likely to be cited due to a perceived endorsement from a star. To distinguish between the productivity and signalling hypotheses, we repeat the analysis using a measure of raw

publications as our dependent variable. Assuming blind refereeing, where referees do not have sight of the acknowledgement texts, raw publications should not be contaminated by a signalling channel. We can, therefore, attain some confidence that we are observing a true productivity effect if the results are robust to switching to a citation-insensitive dependent variable. We find similar qualitative dynamic patterns across the two dependent variables.

Our data consists of all Scopus-recorded academic publications by authors affiliated with a university in the three SOEs across 27 subject fields for the period 1990-2017. Scopus provides acknowledgement data under the 'Funding Text' heading. Using this pool of acknowledgement data, we identify star names and their type of helpful interaction for each publication. We classify the type of acknowledgements based on the keywords often used to acknowledge the star's contribution to the publication. This help by the star can take many forms: an author might acknowledge the star for helpful discussions throughout research; there may be an acknowledgement to comments received during a seminar or a conference; there may be acknowledgements to various forms of technical assistance, especially in the fields of biochemistry, pharmaceuticals, and medicine; or there could be theoretical or mathematical help received during the analysis stages of research. It is important to note that these types of help are not formally captured as a recognised co-authorship. The co-authorship channel has been a major focus of the literature exploring the effects of star connections on productivity (see, e.g., Oettl, 2012). In this chapter, we explore the effects of star help where a formal co-authorship relationship does not exist. Our study thus extends the literature by exploring the productivity effects of star help even where the star does not become a co-author. We investigate both the immediate effects of observations of acknowledgements of star help and any enduring effects over time using a dynamic event-study approach.

Using this event-study methodology, where the event is the acknowledgement by the author to a star in the publication, we test a number of hypotheses relating to the effect of star help on the acknowledging scientist's productivity. We employ a coarsened exact matching (CEM) of star-acknowledging and non-star-acknowledging scientists to create a similar treatment and control group. The findings support the hypothesis of a positive productivity effect and find an interesting dynamic pattern of productivity effects following an acknowledgement to a star. First, the results show a positive star-

help effect on the authors who acknowledge a star, with the impact on productivity evident in the year of publication but largely disappearing in the subsequent years. Second, the results indicate that an author who continues to acknowledge a star in the years after the initial interaction has a higher and more persistent productivity effect. The results are robust to using raw publications as the dependent variable. In considering the types of acknowledgements, we find a significant positive but transitory effect on author productivity in the year of acknowledgement across all types, except materials where the effect is shown to be more persistent. We speculate that materials-related acknowledgements may reflect more enduring relationships that are not repeatedly acknowledged in publications over time. Finally, we find that the estimated productivity impact of star help is greater for assisted authors that are in the lower quartiles of the relevant field-specific productivity distribution in the year before the (initial) acknowledgement occurs, suggesting that less productive authors get greater benefit from a star's help.

This chapter contributes to the literature in several ways. First, the chapter investigates the impacts of star help on peer productivity. Prior research mainly focuses on the star scientist's proximity with limited consideration of the precise mechanisms involved (Agrawal et al., 2017; Azoulay et al., 2005). Second, in contrast to the existing literature, the chapter focuses on the effects of star help in small open economies. This fills a significant gap in the literature, given the importance that policies to improve connections to star scientists (including but not limited to star recruitment policies) play in the policy mix of smaller national and regional economies. Third, the chapter identifies non-co-authorship-based star connections using acknowledgements in the publication, where we hypothesise that such non-co-authorship-based interactions have an impact on a scientist's productivity. Moreover, the chapter differentiates these informal interactions into Conceptual, Technical, Material, Funds & Support, and Other Types. Finally, the chapter employs the growing literature on heterogeneous treatment effects accounting for staggered treatment by using Sun and Abraham's (2021) estimator. This allows us to consistently estimate the effects of interest in the potential presence of heterogeneity across both cohorts and periods.

The remainder of the chapter is structured as follows. Section 3.2 reviews the related literature and describes the framework and hypotheses. Section 3.3 describes the

data and the use of NLP in extracting acknowledgement data from the acknowledgement texts and outlines the event-study-based empirical framework used to identify the causal effects of star help. Section 3.4 reports the results and Section 3.5 discusses various robustness tests for the analysis. Finally, Section 3.6 concludes with a review of the main findings and their implications for science policy.

3.2 Related Literature and Hypotheses

3.2.1 Related literature

A key determinant of the growth of an economy is the accumulation and transmission of knowledge (Romer, 1986, 1990). Romer identifies two processes, human capital spillovers and R&D investments, where knowledge accumulation and spillovers happen. Human capital spillovers benefit others in the economy through increased productivity and innovation. These innovations are often a result of interaction between individuals with a considerable stock of previously accumulated knowledge. As potential key figures in this knowledge spillover process, star scientists may play a central role in the workflow structure due to their unique expertise and social status (Paruchuri, 2010). These spillover agents are often highly skilled pioneers in their fields who induce knowledge flows across organisational and regional boundaries and affect ongoing research projects inside academia and firms (Bergman and Schubert, 2005; Maier et al., 2007).

Waldinger (2012) and Azoulay et al. (2010) discuss the adverse impacts of losing ties with a star on collaborators. They find a decline in productivity and quality-adjusted publications for those affected by a star's externalities. Zucker and Darby (2006) find a positive impact of star scientists in start-up biotech firms. In these studies, researchers focus more on the impact on the productivity distribution rather than the mechanism enabling the knowledge flow. Further studies have provided evidence that the influence of a star might depend upon the collaborative strength and breadth of the star's expertise (Kehoe and Tzabbar, 2015) and have documented significant long-term effects on the performance of junior researchers (Li et al., 2019).

Although the presence of stars could catalyse the development of peers' scientific output, helpful interactions are essential for some knowledge transfer mechanisms. First, human capital externalities generate knowledge flows that must be recognised as valuable, then assimilated and recombined with previous knowledge to generate new scientific knowledge. In other words, to take advantage of stars' helpfulness, scientists have to increase their absorptive capacity and their capacity to apply knowledge to their research projects (Cohen and Levinthal, 1990). Although a certain level of previous related knowledge is required for learning to take place (Cohen and Levinthal, 1990; Torre and Rallet, 2005), recurrent helpful interactions and the resulting knowledge exchange adjust the optimal level of knowledge overlap between scientists over time (Nooteboom, 2000), which is fundamental to facilitate the transfer of knowledge (Boschma, 2005).

Second, these interactions also take advantage of spatial proximity to transmit formal and tacit knowledge forms (Bathelt et al., 2004). Scientists working in the same department go through face-to-face (F2F) communications, allowing efficient transmission of complex and tacit knowledge (Gertler, 2003; Storper and Venables, 2004). However, organised proximity, i.e., belonging to the same community of practice or sharing the same system of representations, could alleviate the requirements of being co-located for knowledge transfer and learning to occur. Information and communication technologies and temporary co-locations such as meetings and academic congresses could solve communication problems (Torre and Rallet, 2005). In the case of star help, both channels (F2F and organised proximity) are helpful for knowledge transfer that affects the productivity of incumbent scientists.

Finally, information sharing is the key to scientific progress in any spillover mechanism. Merton (1973) defines the norms of unconditional knowledge-sharing as one of the critical processes in academic life. At the same time, there is tension among academic individuals to share their valuable information due to incentives associated with the research (Blumenthal et al., 1996; Walsh and Cohen, 2008; Murray, 2010). Haeussler et al. (2013) consider two possibilities in which such information sharing happens: specific sharing, where researchers share their private work on a specific request, which in turn will get recognised for future accomplishments, and general

sharing, where such information is shared publically, such as sharing unpublished data, materials, etc.

This study discusses how star help affects scientists' productivity. Through this particular mechanism, the absorptive capacity of scientists is enhanced, the optimal level of knowledge overlap is adjusted through repeated social interactions, and knowledge sharing is facilitated. However, while networks can positively affect productivity, their specific nature can also play a crucial role in determining their impact. For example, a study by Horta et al. (2010) found that academic inbreeding, or hiring PhD graduates as faculty members at the same university where they received their degree, can harm scholarly output. Of course, although understudied, informal interaction is just one example of a network-intermediated productivity effect. More broadly, networks help recombine specialised ideas and knowledge transformation (Groysberg et al., 2008; Rothaermal and Hess, 2007). Another widely studied network effect is co-authorship. Grigoriou and Rothaermal (2014), Agrawal et al. (2017), and Yadav et al. (2023) find increased productivity gains and show very high positive spillovers of co-authorship with a star. Also, the network position of the co-author is crucial for the productivity that facilitates access to non-redundant knowledge (Mohnen, 2022). However, the impact of stars on productivity through non-co-authorship-based network channels remains comparatively understudied.

One important exception to the neglect of the helpful interaction channel is Oettl (2012). He defines a taxonomy of stars based on these social interactions: maven (highly helpful and average productivity), all-star (highly helpful and highly productive), lone wolf (highly productive and less helpful), and non-star. Following Oettl's approach, we use acknowledgement texts in publications to provide a paper trail on helpfulness activities. Other studies have also utilised the information available in acknowledgement texts. Mackintosh (1972) classifies acknowledgements based on the facilities, access to data, and help of individuals. McCain (1991) classifies acknowledgements from experimental papers in genetics into five typologies: access to research-related information, unpublished results or data, peer interactive communication, technical assistance, and manuscript preparation. Cronin (1991) (later modified in Cronin, McKenzie and Rubio (1993)) introduces six typologies for acknowledgements: access, peer interactive communication, moral support, technical support, clerical support, and

financial support. These studies focus on the typology of acknowledgements, while our focus is on how different types of acknowledgements are associated with the observed productivity effect on the acknowledging author. Moreover, we build on this literature by using observed acknowledgements as not just thank-you notes but indicators of important productivity-affecting activities that are distinct from co-authorships and citations (Cronin, 1991; Paul-Hus et al., 2019).

Nevertheless, the information in the acknowledgements of academic publications should be applied with caution and with sensitivity to context. Using acknowledgements as a sign of helpful interaction with a star scientist could be either a sign of intellectual debt or a signalling effect. For example, along with co-authorship, acknowledgements could also reveal other forms of collaboration (Laudel, 2002). The content of such acknowledgements is intended to repay debt towards formal and informal collaboration. In addition, acknowledgements could reflect personal relationships and individual preferences. Hellqvist (2010) conducted a qualitative analysis of acknowledgements in sociology journals and suggested that personal style, editorial guidelines, cultural norms, and ethical principles influence the pattern of acknowledgements.

The literature also considers the various motivations for acknowledging prior work in a paper. Berg and Faria (2008), for example, argue that the scientist names chosen in acknowledgements are based on the effect they may have on readers. We, therefore, recognise that acknowledgement of star help is a noisy indicator of actual star help but assume there is a sufficient signal in the acknowledgement data to make them useful in the estimation of star-help effects on research productivity. Furthermore, the type of helpful interaction also matters to those who receive such help from star scientists. As argued by Oettl (2012), the type of helpfulness being acknowledged is also relevant to the productivity effects of these acknowledgements, and we also distinguish between different types of star help in our empirical analysis.

3.2.2 Framework and hypotheses

We conceptualise a star scientist as a scientist who is usually productive and connected within their network. From the point of view of a non-star scientist, we assume, following the literature cited above, that forming a connection with a star scientist changes their network position and potentially their productivity. As an example, a scientist will obtain access to knowledge through their network, with their eigenvalue centrality within their network being one commonly used measure of how easily they can access knowledge (Newman, 2018). Forming a connection with a star will change their centrality and thus their access to knowledge. As another example, forming a collaborative connection with a star can directly increase productivity and may have further indirect effects through access to the star's broader collaborative network.

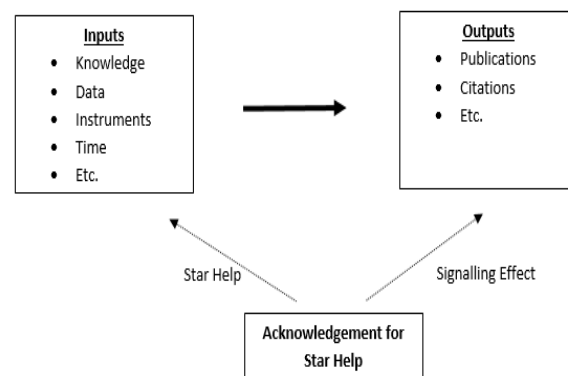


Figure 3.1 Star help and observed scientific productivity

This chapter focuses on connections to a star that deliver potentially productivity-improving star help, separate from any direct effect through co-authorship. We think of such star help as augmenting the scientist's inputs into their production of scientific outputs (as measured by citation-weighted publications). Figure 3.1 captures schematically the relationship between inputs and outputs, where star help is viewed as augmenting the scientist's inputs. We also allow for the possibility that a positive effect of observed star acknowledgements could reflect a signalling effect, where the positive effect on subsequent output – particularly citations to a publication – reflects a signalling (or reputational spillover) effect from the star.

Our empirical approach is to look for evidence of the presence and type of star help in the acknowledgement text of publications. By carefully matching each scientist that is "treated" with star help to an observationally similar scientist who is not observed to receive such help, we use an event-study framework to estimate the dynamic effects of treatment.

Our first hypothesis is that an observation of acknowledgement for star help is associated with a contemporaneous increase in author's productivity, with productivity measured by citation-weighted publications.

Hypothesis 1a: The observation of an acknowledgement for star help is associated with an increase in citation-weighted publications in the year of acknowledgement.

In addition, an important advantage of our event-study framework is that we can observe the effects of an acknowledgement for star help in a given year on the evolution of productivity over time. This leads to our second hypothesis:

Hypothesis 2a: Repeated observations of acknowledgements for star help are associated with a sustained increase over time in productivity as measured by citation-weighted publications.

A significant challenge is disentangling the input-enhancing role implied by observed star acknowledgements from signalling effects. On the assumption that acknowledgement texts are not observable to referees, raw publication data (i.e., output not weighted by subsequent citations) should not be contaminated by signalling effects. This leads to the following hypotheses based on a simple raw publications measure of output:

Hypothesis 1b: The observation of an acknowledgement for star help is associated with an increase in raw publications in the year of acknowledgement; and

Hypothesis 2b: Repeated observations of acknowledgements for star help are associated with a sustained increase over time in productivity as measured by raw publications.

Finally, acknowledgements to a star scientist can be classified by type based on the keywords in each publication's acknowledgement text. These keywords assist with identifying the nature of the help between the author and star, leading to our third hypothesis that the effect of star help on an author's productivity will vary across the types of helpful interaction with the author.

Hypothesis 3: The observation of an acknowledgement for star help is associated with an increase in citation-weighted publications in the year of acknowledgement, with the magnitude of the increase potentially varying across the different types of helpful interaction.

3.3 Data and Methodology

Our dataset consists of all the publications and their citations across the 27 subject fields identified by Scopus, where each publication contains at least one author affiliation recorded in Ireland, Denmark, or New Zealand. As discussed previously, Ireland and Denmark are chosen because they are small open economies with nationally funded star recruitment programmes, while New Zealand is chosen because it is also a small open economy with no formal star recruitment programmes. We collect variables at the publication level for the year of publication, authors, affiliations, subject field, citation count to 2019, references, abstracts, keywords, and acknowledgements. We identify authors and publications using unique Scopus identifiers. The dataset contains approximately 1.43 million publications divided over 219,582 unique authors. The unit of observation is at the author level rather than the publication level, and the dependent variable tracks the author's productivity before and after the star help revealed from the acknowledgement texts.

To help capture an author's performance over time, we restrict our analysis to the authors present in one of the three countries for at least four years. For each author who satisfies this condition, we access their catalogue of publications dating back to 1990, including those outside their affiliations in the three focal countries. The final dataset contains an unbalanced panel of 889,479 unique publications divided over 59,122 authors across the years. Using the affiliation data of the author, we determine the

department (defined based on the Scopus definition of subject fields²¹). If an author publishes in more than one subject field, we assign them to the department in which they have the most publications.

3.3.1 Identification of a star scientist

Given that this study focuses on the impact of star scientists on their peers through channels of knowledge spillover identified by the acknowledgement text in an author's publication, the first step is to identify the star scientists in our dataset. The dataset contains the citation distribution of each author from 1990 for each subject field in any given year. A star is defined as one who is at or above the 95th percentile of scientists in the cumulative distribution of citations for the relevant subject field received for all publications up to that year. In measuring the citations to any given publication, we measure all citations up to the end of our observation period in 2019, so our publication quality measure depends on the subsequent citations to that publication.

We only identify star scientists from 1996 onward to allow enough time for the accumulation of citations, given that our publication data began in 1990. Similar to the department allocation for each author, we determine the department for each star as the department in which they had the most publications. The star scientist identified is relative to scientists in Denmark, Ireland, and New Zealand. Hence, these stars are highly productive scientists in the context of the overall distribution of scientist citations for these three countries. Through this method, we identify 981 stars from the dataset, 205 in Ireland, 548 in Denmark, and 228 in New Zealand.

3.3.2 Identification of star names in the acknowledgements

We identify star names and their author interaction type from the pool of available acknowledgement text. Since the study focuses on the impacts of star scientists on their peers through informal collaborations, we think of the influence of these interactions as occurring through intense discussions, critical reviews, data sharing, collaborations, and supervision. Broad forms of funding texts are available in Scopus, which is translated into

²¹The 27 subject fields are Agricultural and Biological Sciences, Arts and Humanities, Biochemistry, Business, Chemical Engineering, Chemistry, Computer Science, Decision Science, Dentistry, Earth and Planetary Sciences, Economics, Energy, Engineering, Environmental Science, Health Professions, Immunology and Microbiology, Multidisciplinary, Materials Science, Mathematics, Medicine, Neuroscience, Nursing, Pharmacology, Physics and Astronomy, Psychology, Social Sciences, and Veterinary.

acknowledgement texts. A limitation of the Scopus database is that the acknowledgement text is not available for every publication. Table B.1 in the Appendix.B.1 outlines the yearly distribution of publications and available acknowledgement texts in Scopus.

Using a Natural Language Processing (NLP) tool, we extract the names²² of scientists from our acknowledgement text data. We use the open-source library, Spacy, an advanced NLP tool in Python designed to process a large volume of text data. More specifically, we utilise Spacy's Named Entity Recognition (NER) feature to extract named entities from each acknowledgement text. Named entities are phrases that contain the names of persons, organisations, and locations. Spacy's NER feature has reported an accuracy of 89.8 (OntoNotes²³) and 91.6 (CoNLL-2003 corpora²⁴) percent and compares well to other packages available for NLP. The tool helps to token the words in an acknowledgement text, and we group them using a unique publication identifier. Using the NER feature, we tag the author entities in each publication's acknowledgement text, and then we match these names to those of the star scientists identified in Section 3.3.1. Overall we identify 331 stars acknowledged over 971 publications by 1815 authors.

3.3.3 Identification of acknowledgement types

While prior studies primarily focus on the typology of the overall acknowledgement text (Paul-Hus et al., 2016, 2017; Desrochers et al., 2018; Cronin, 1991), we focus on the types of interaction between a star and an author. Using a similar approach as Oettl (2012) we identify the type of star help based on their helpfulness. Moreover, these acknowledgement types are classified in our study based on the available keywords that determine the star help in our data. Here we define five types of acknowledgements to a star who is not a co-author: Conceptual, Technical, Materials,

²²We provide the procedure for identifying star names and the type of star help from the acknowledgement texts in the Appendix B.2.

²³OntoNotes project is a collaborative effort between BBN Technologies, the University of Colorado, the University of Pennsylvania, and the University of Southern California's Information Sciences Institute. The goal of the project was to annotate a large corpus comprising various genres of text (news, conversational telephone speech, weblogs, USENET newsgroups, broadcast, and talk shows) in three languages (English, Chinese, and Arabic) with structural information (syntax and predicate argument structure) and shallow semantics (word sense linked to an ontology and coreference).

²⁴CoNLL-2003 shared task: concerns language-independent named entity recognition. They concentrate on four types of named entities: persons, locations, organizations, and names of miscellaneous entities that do not belong to the previous three groups (Tjong Kim Sang, Fien De Meulder, 2003)

Funds & Support, and Other Types. Each acknowledgement type indicates the interaction between an author and a star during the research project.

Conceptual acknowledgements recognise star scientists for their intellectual contribution to the research project. For example, "We thank P. Di Vecchia, G. Grignani, C. Kristjansen, and N. Obers for useful discussions" (Harmark T and Orselli M, 2006) and "I would like to thank Prof. P. Sigmund for stimulating and enlightening discussions of the topic and useful comments on the manuscript" (Glazov L.G, 1998). Using keywords, we identify this type of interaction between a star and the author. Here, keywords such as "discussions," "comments," "feedback," "critique," "advice," "suggestions," are used. Keywords that categorise help in the technical form are "technical assistance," "statistical assistance," "excellent assistance," and "expert assistance." For example, "We thank P. Rasmussen, B. Jensen, and D. Bardenfleth for expert technical assistance" (Sorensen M.R. et al., 2010). In this example, P. Rasmussen is the star, and the author acknowledges his help through their technical assistance.

Material acknowledgement shows the debt of gratitude for materials and data shared by the star for the research purpose. Keywords used to identify interactions based on this specification are "data sample," "antibodies," and "cells." For example, "Professor Klaus Bendtzen is thanked for providing antibodies for the cytokine measurements" (Theander T.G. et al., 1997). Another acknowledgement type captures the author's "funding and supports" through the star interaction. Keywords such as "grants," "funds," "support," and "financial." are used to identify the interaction with the stars. For example, "This research was sponsored through the contestable research fund of the Waikato Management School of the University of Waikato and Professor Chris Ryan, Waikato Management School, for his continued support" (Lockyer T, 2005). Finally, acknowledgements using the keywords such as; "committee," "contributions," "permission," "facilities," "director," "supervisor," "founder," "help," "dedication," are classified as "Other Types." For example, "The study was initiated by Torben Jørgensen, DMSc (PI); Knut Borch-Johnsen, DMSc (Co-PI); Troels Thomsen, Ph.D., and Hans Ibsen, DMSc" (Baumann S. et al., 2015). A similar approach is adopted by (Paul-Hus et al., 2019) to classify acknowledgement types as 'vague.'

Table 3.1 Acknowledgement types that define the types of helpful interactions

Subject Field	Conceptual	Technical	Materials	Funds & Support	Other Types	Total
Agricultural and Biological Sciences	177	105	53	29	58	422
Arts and Humanities	2	0	1	0	0	3
Biochemistry	234	113	104	37	37	525
Business	22	0	1	2	2	27
Chemical Engineering	8	0	0	0	3	11
Chemistry	51	40	14	3	24	132
Computer Science	10	0	0	4	5	19
Dentistry	5	6	6	0	7	24
Earth and Planetary Sciences	83	16	34	6	18	157
Economics	25	3	2	1	1	32
Energy	2	2	0	0	3	7
Engineering	73	11	1	3	22	110
Environmental Science	17	20	2	5	9	53
Health Professions	1	0	0	0	0	1
Immunology and Microbiology	18	28	17	4	3	70
Materials Science	13	14	12	0	13	52
Mathematics	13	2	0	1	3	19
Medicine	157	133	142	96	509	1,037
Multidisciplinary	1	4	2	0	0	7
Neuroscience	44	23	19	5	11	102
Pharmacology	30	9	1	2	3	45
Physics and Astronomy	157	26	32	11	32	258
Psychology	14	1	0	4	8	27
Social Sciences	15	5	5	5	7	37
Veterinary	1	14	3	2	1	21
Total	1,173	575	451	220	779	3,198

Note: Distribution of acknowledgement types over the subject fields based on the interaction with a star. The table includes multiple acknowledgements from 1815 unique authors over 971 publications. The authors acknowledge a total of 331 stars from 1990-2017.

Table 3.1 presents the variation across the different types of acknowledgements by 1815 authors to 331 stars over the period 1990 to 2017. Star help classified as the conceptual type constitutes 37% of the overall acknowledgements, where the subject areas: Agricultural and Biological Sciences; Biochemistry; Medicine; Physics; and, Astronomy account for 62% of conceptual acknowledgements to a star scientist. Star help based on technical acknowledgements occur primarily in the subject areas: Agricultural and Biological Sciences, Biochemistry, and Medicine, where the research in these departments can include considerable laboratory work as well as the use of complex instruments. Material acknowledgements account for 14% of the general acknowledgements and are mainly observed in the fields of Medicine and Biochemistry.

Only 7% of the acknowledgements to star scientists in our sample show an acknowledgement to a star for the funding the author received.

3.3.4 Output Measure

We use two output measures to test our hypotheses discussed in Section 3.2.2: (i) Field Normalised Total Citations (FNTC) - the sum of an individual's publication citations divided by the average citations to a publication for that subject field in that year for all countries combined; and (ii) total count of normalised raw publications of the individual author in that year. We calculate the dependent variable as follows:

Field Normalized Total Citations:
$$Y_{it}^{FNTC} = \sum_{p_{i,t}=1}^{P_{i,t}} \frac{c_{p_{i,t}}}{\bar{c}_{s,t}};$$

Publications:
$$Y_{i,t}^P = \frac{P_{i,t}}{\bar{P}_{s,t}};$$

where $P_{i,t}$ is the total number of publications by individual i published in year t , $c_{p_{i,t}}$ are the subsequent total citations (or "forward" citations recorded in 2019) to a publication $p_{i,t}$ that occur for individual i in year t , $\bar{c}_{s,t}$ is the average citations to a publication in the relevant subject field, s , for publications that occur in year t , and $\bar{P}_{s,t}$ is the average number of publications in subject field s .

These output measures provide the dependent variables for our regressions. Using both dependent variables, we use the author's publication data from 1996-2017 (our estimation window) to calculate the overall impact of the star's helpful interaction. In our analysis, we split the individuals into two subgroups - Treated: Authors who acknowledge the star for the helpful interaction for their publications and who did not co-author with the star before the acknowledgement; and Never-Treated: Authors who received neither helpful interaction nor a star co-authorship interaction. Section 3.3.5 below discusses the matching procedure to identify matched pairs of treated and control authors before the treatment event happens.

Table 3.2 Summary statistics: control and treated group (k-to-k matched)

Variable	Control	Treated	Diff in mean	P-value
<i>An unbalanced panel of 1258 matched authors, with 629 in each group</i>				
Year	2007.122	2007.099	0.024	0.944
Subject	10.906	10.906	0.065	0.672
Country	1.988	1.986	0.000	1.000
Total Career Age	20.541	20.655	-0.114	0.841
Total Career Age Bins	5.730	5.730	0.000	1.000
Cumulative Publication experience	9.068	9.049	0.019	0.959
Cumulative Publication experience Bins	2.366	2.366	0.000	1.000
Cumulative citations received per cumulative publications	44.226	44.769	-0.543	0.954
Cumulative citations received per cumulative publications Bins	2.035	2.035	0.000	1.000

Notes: Reports the t-test for the mean difference between control and treated groups one year before forming the star interaction

3.3.5 Coarsened Exact Matching Procedure (CEM)

Matching treated units to control units in observational data helps to mitigate the confounding influence of pre-treatment control variables with the primary goal of improving the balance between the treated and control groups. We, therefore, employ matching to help control for the endogeneity issue that might arise in acknowledging a star. For example, acknowledging a star can be considered random in the cases such as tips and help received from a seminar series or international conferences. At the same time, prior connections with the author also involve the star being a part of the research as an informal contributor. To address this, we create a panel data which has a control group comprising authors who never received any star exposure in terms of star's helpful interaction or co-authorship and a treated group - a similar set of authors in terms of characteristics who received star help that is identified from the acknowledgement texts of the publications

In our study, we ensure that the treated authors had no star interaction in terms of star co-authors before or during the event of the star's helpful interaction. To minimize the effect of confounding in our observational causal inference, we employ Coarsened Exact Matching (CEM), which is a monotonic imbalance-reducing matching method (Blackwell et al., 2009; M. Iacus et al., 2011; M. Iacus et al., 2012). CEM is a design strategy that involves matching on a set of covariates that have been "coarsened," meaning that they have fewer possible values for matching, which increases the number of matches (G.

King et al., 2011; G. King et al., 2019). This technique has been shown to improve the balance of covariates between exposure groups and to guarantee balance for each covariate, only limited by the coarseness of the grouping (Fini et al., 2018). In contrast, other matching techniques, such as Propensity Score Matching (PSM), do not guarantee each variable such a balance guarantee. They may require repeated iterations to achieve balance (Fini et al., 2023). Moreover, CEM ensures balance for higher-order terms, such as interactions of covariates, while such a guarantee does not exist in a propensity score approach.

We find an author in the control group at $t - 1$ year who matches similar characteristics (based on the observables) with an author from the treated group at $t - 1$ who receives star help at year t . The CEM procedure allows us to define the covariates to match a categorical variable rather than a continuous variable. We identify a match based on five covariates: subject field, country, total career age, cumulative publication experience, and cumulative citations received on prior publications. Each covariate uses a categorical variable with course bins²⁵(M. Iacus et al., 2012). Table 3.2 shows the difference in the mean value and p-value of each matching variable for both control and treated groups one year before the star interaction occurs. The unbalanced matched panel data set contains 1258 authors, with 629 authors in the treated group who acknowledge a star for the first time. Furthermore, 75 of these authors acknowledge the stars in the years after the first acknowledgement. Similarly, for the five cohorts based on acknowledgement types: Conceptual, Technical, Material, Funds & Support, and Other Types, we use the CEM matching procedure to identify matching pairs of control and treated author. Table B.3 in the Appendix B.3 reports the matching statistics of matched pairs for each cohort.

²⁵ We create twenty-seven bins for the subject fields and three for the countries. Eleven bins for the total career age of the scientist ranging from 4 years to 50 years; less than 5 years, between 5 and 10 years, between 10 and 15 years, between 15 and 20 years, between 20 and 25 years, between 25 and 30 years, between 30 and 35 years, between 35 and 40 years, between 40 and 45 years, between 45 and 50 years, 50 years and above. Cumulative publications experience captures cumulative years since the first year of publication (here in this data from 1990 onwards). Seven bins range from 0 to 1 year, 1 to 5 years, 5 to 10 years, 10 to 15 years, 15 to 20 years, 20 to 25 years, and 25 and above. Finally, cumulative citations received per cumulative publication are categorised based on the distribution of resulted average value, which shows the citations received by the author per year based on the number of publications. Sixteen bins are created for values less than 1, between 1 to 25, 25 to 50, 50 to 75, 75 to 100, 100 to 150, 150 to 200, 200 to 250, 250 to 300, 300 to 350, 350 to 400, 450 to 500, 500 to 750, 750 to 1000, and 1000 and above.

Table 3.3 Dynamic star help effects for FNTC (homogenous and heterogeneous effects)

	Homogenous Star Help Effects (1)	Heterogeneous Star Help effects (2)
<i>Staracknw</i> _{<i>i,t-3</i>}	-0.00829 (0.0255)	-0.0105 (0.0256)
<i>Staracknw</i> _{<i>i,t-2</i>}	-0.0226 (0.0240)	-0.0205 (0.0242)
<i>Staracknw</i> _{<i>i,t</i>}	0.247*** (0.0243)	0.2463*** (0.0244)
<i>Staracknw</i> _{<i>i,t+1</i>}	0.0432* (0.0255)	0.0505** (0.0256)
<i>Staracknw</i> _{<i>i,t+2</i>}	0.0450* (0.0263)	0.0581* (0.0267)
<i>Staracknw</i> _{<i>i,t+3</i>}	0.0344 (0.0278)	0.0432 (0.0285)
Constant	0.402*** (0.0237)	0.460*** (0.00916)
R-squared	0.029	0.521
Pretest against the hypothesized trend (Roth 2022)		
Power	0.50	0.50
Hypothesized trend	0.02	0.02
Bayes factor	0.55	0.55
Likelihood ratio	0.30	0.38
Observations	20,451	20,451
Number of authors	1,238	1,238
Author FE	YES	YES
Year FE	YES	YES

Event Window: 1997-2017 (1996 cohort dropped for comparison only). Estimation Window: 1996-2017.

The dependent variable is Field-Normalised Total Citations.

Note: The table reports the estimates based on the model specification in the equation 3.1 and equation 3.3. Column 1 reports the homogenous star-help effects at the individual level. Column 2 reports the estimates using Sun and Abraham's (2021) method. Also, the 1996 cohort is dropped to compare the two methods, although it does not affect the final estimated coefficients since only ten authors are treated from 1996. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively.

3.3.6 Econometric methodology

The empirical goal of our econometric analysis is to estimate the effects of helpful interactions with star scientists on the productivity of authors over time, where we capture the interactions from the acknowledgement texts of the author's research publications. As we discussed, a star might impact the productivity of their peers through various channels, and our focus is to identify whether an author's informal non-co-authorship interactions with a star could affect their productivity measured in terms of

(forward) citations to their publications (hypothesis 1) & raw publications output (hypothesis 2).

To measure these dynamic effects, we utilise an 'event-study' specification where the event is an acknowledgement to a star scientist that can occur in the past or future. Our event-study design is a staggered adoption design where units (authors) are treated at different times, and some units have never been treated. We estimate the dynamic treatment effects of the helpful interaction with a star on their peers from three years before the published acknowledgement to three years afterwards in the following event study specification:

$$\begin{aligned} \ln Y_{it} = & \alpha + \beta_{\leq -4} \text{staracknw}_{i,-4} + \sum_{j=-3}^{-2} \beta_j \text{staracknw}_{ij} + \sum_{j=0}^3 \beta_j \text{staracknw}_{ij} \\ & + \beta_{\geq 4} \text{staracknw}_{i,4} + \delta_t + \mu_i + \epsilon_{it}, \end{aligned} \quad (3.1)$$

where the dependent variable $\ln Y_{it}$ is a measure of the citation-weighted/normalised raw publication output of author i at year t , staracknw_{ij} is a binary variable equal to 1, if a star is acknowledged by the author i as of year t , j years ago, δ_t is a year fixed effect, μ_i is an author fixed effect, and ϵ_{it} is a zero mean error term. The coefficients of interest, β_j , show the proportionate effect of the helpful interaction with the star on productivity from three years before the acknowledgement to three years afterwards. We normalise the effect of the star interaction to zero for the year before the star acknowledgement, and we assume that the cumulative effect is constant at $\beta_{\leq -4}$ and $\beta_{\geq 4}$ by binning at four leads and four lags. The binning variables may not be comparable to the leads and lags of the acknowledgement binary variables in estimating the dynamic effects since they could be correlated with other excluded level variables; however, they act as essential controls in our specification (Kurt Schmidheiny et al., 2019). Furthermore, standard errors are clustered at the author level and are robust to arbitrary forms of serial correlation and heteroscedasticity. McHale et al. (2021) adopt a similar econometric approach that estimates the effects of star arrival on the departments' productivity in SOEs.

An important assumption of our event study model is the generalised form of the parallel trends assumption, whereby without a star acknowledgement, the quality-




Figure 3.2 Event study model with homogenous (left) vs. heterogeneous (right) star help effects at an individual level.

Note: The figure plots the dynamic effect of the star's helpful interactions with an author at 95% confidence intervals. The event and observation window is from 1996-2017 (the 1996 cohort dropped is omitted). The dependent variable is field normalised total citations.

adjusted output in the treatment group would have changed in the same way as it did in the non-treatment group (the authors that did not acknowledge a star scientist). The estimated lead coefficients in our specification allow us to examine indirect evidence to support the parallel trends assumption, with any observed pre-acknowledgement effects considered evidence of a failure of this assumption. On the other hand, an observed pre-trend could indicate anticipation effects of star acknowledgement on productivity; for example, if an author had prior knowledge of possible helpful interaction with a star in the future, they might change their productivity behaviour. However, in our case, this is considered to be unlikely since the author acknowledges a star for his helpfulness which could not be anticipated before the interaction, and therefore, we assume that anticipation effects are zero.

Another critical assumption in this event-study setting is that the star acknowledgement effect on author productivity is homogeneous across the timing of the acknowledgements. However, recent literature has shown that the coefficients of given leads and lags can be contaminated by the effects from other periods in the presence of heterogeneous effects across different treatment timings (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021). In our econometric analysis, star acknowledgement has a staggered treatment timing, and heterogeneity in the effects could arise if different cohorts experience different treatment paths. Therefore, we also adopt the approach of Sun and Abraham (2021) to estimate the heterogeneous treatment effects. This approach derives the dynamic effects of star acknowledgement in a three-

step estimation that is robust to treatment effect heterogeneity and calculates a weighted average of 'cohort average treatment effects on the treated' (CATT). First, we define the year in which an author i acknowledges a star scientist as e_i . Second, we estimate the weighted average of cohort effects for a given time relative to the acknowledgement event. To allow the estimated star acknowledgement effects to vary by cohort based on the year that the acknowledgement event occurs, we estimate the following equation:

$$\ln Y_{i,t} = \sum_e \left[\delta_{e,-4} (\mathbf{1}\{E_i = e\} \text{staracknw}_{i,-4}) + \sum_{j=-3}^{-2} \delta_{e,j} (\mathbf{1}\{E_i = e\} \text{staracknw}_{ij}) + \sum_{j=0}^3 \delta_{e,j} (\mathbf{1}\{E_i = e\} \text{staracknw}_{ij}) + \delta_{e,4} (\mathbf{1}\{E_i = e\} \text{staracknw}_{i,4}) \right] + \delta_t + \mu_i \quad (3.2)$$

where $\mathbf{1}\{E_i = e\}$ is an indicator variable that takes the value 1, if the individual i receives star help in the year e and 0 otherwise. $\delta_{e,j}$ is the star help effect on productivity j year after author acknowledges a star in year e . The 1996 treated cohort is dropped from the analysis since it is always treated across the observation window. A further set of weights are estimated $Pr\{E_i = e | E_i \in [-j, T - j]\}$ that are equal to sample shares of each cohort for the relevant periods j . Finally, to obtain the IW estimator, we take a weighted average of the $\hat{\delta}_{e,j}$ (or $CATT_{e,j}$) and estimate Equation 3.2 with relevant weights calculated.

$$\beta_j^* = \sum_e [\hat{\delta}_{e,j} Pr\{E_i = e | E_i \in [-j, T - j]\}] \quad (3.3)$$

Table 3.4 Dynamic star help effects for FNTC (multiple and one-time)

	Multiple acknowledgements to stars (1)	One-Time acknowledgement to stars (2)
<i>Staracknw_{i,t-3}</i>	-0.0891 (0.0757)	-0.00244 (0.0270)
<i>Staracknw_{i,t-2}</i>	-0.000964 (0.0762)	-0.0262 (0.0252)
<i>Staracknw_{i,t}</i>	0.348*** (0.0634)	0.232*** (0.0260)
<i>Staracknw_{i,t+1}</i>	0.201*** (0.0686)	0.0269 (0.0274)
<i>Staracknw_{i,t+2}</i>	0.237*** (0.0736)	0.0145 (0.0276)
<i>Staracknw_{i,t+3}</i>	0.239*** (0.0838)	0.00363 (0.0289)
Constant	0.541*** (0.0705)	0.387*** (0.0243)
R-squared	0.055	0.028
Pretest against the hypothesized trend (Roth 2022)		
Power	0.50	0.50
Hypothesized trend	0.07	0.03
Bayes factor	0.55	0.55
Likelihood ratio	1.68	0.19
Observations	2,681	18,160
Number of authors	150	1,108
Author FE	YES	YES
Year FE	YES	YES

Estimation Window: 1996-2017.

The dependent variable is Field-Normalised Total Citations.

Note: The table reports the estimates based on the model specification in equation 3.1. Column 1 reports the subsequent star-help effects at the individual level after the first year of star acknowledgement. Column 2 reports the first-time star acknowledgement without any subsequent interaction. Also, the 1996 cohort (always treated) is not dropped in this case. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively.

3.4 Results

In this section, we present the results for our different hypotheses outlined in Section 3.2.2. We use FNTC as the dependent variable for testing hypotheses 1a and 2a and the number of raw publications for testing hypotheses 1b and 2b. Furthermore, we only use FNTC as the dependent variable to test hypothesis 3, where we explore the effect of star help across the types of helpful interaction identified.




Figure 3.3 Event study model with multiple (left) vs. one-time (right) star help effects at an individual level

Note: The figure plots the dynamic effect of the star's helpful interactions with an author at 95% confidence intervals. The event and observation window is from 1996-2017. The dependent variable is field normalised total citations.

3.4.1 Star help effect on individual output (with FNTC as dependent variable)

We examine hypotheses 1a and 2a in this section. We present the estimated results from the event study specification for both the homogenous and heterogeneous star acknowledgement effects in Table 3.3. The dependent variable is FNTC to the output published in year t . Column 1 of Table 3.3 shows the results for the homogeneous model, and the star acknowledgement effects are also shown on the graph in the left panel of Figure 3.2. We find evidence that a helpful interaction with a star, as captured by a star acknowledgement, significantly affects an author's productivity supporting hypothesis 1a. A star acknowledgement is associated with an economically and statistically significant contemporaneous increase in the quality-adjusted output of the author of 24.70 log points, which translates to a 28.02% increase in output. In the years after the star acknowledgement, the estimated coefficients remain positive but decrease substantially to between 0.04 and 0.05 log points and are only statistically significant at the 10% level.

Regarding the heterogeneous model using the approach of Sun and Abraham (2021), the event study results are reported in Column 2 of Table 3.3 and displayed in the right panel of Figure 3.2. Under the heterogeneous model, we find a similar contemporaneous effect as in the homogenous model, with a 28% increase in the quality-adjusted output of the author in the year of acknowledgement. Moreover, the results lend further support to hypothesis 1a when the differential timing of treatment is taken into account in our estimation. Furthermore, the estimated coefficients for the years following

a star acknowledgement are again similar to the homogenous model (0.05 to 0.06 log points), and these coefficients are statistically significant at the 5% and 10% levels, respectively. In considering the staggered timing of treatment on authors that have a helpful interaction with a star, the homogenous results are robust to the heterogeneous specification in equation 3.3. Overall, the results suggest that the quality-adjusted output of an author increases contemporaneously when a star is acknowledged, but the effect tends to fall in subsequent years. We also dropped the 1996 (always treated) cohort from the homogenous estimation for comparison purposes, which has minimal effects on the final results. Overall, we find supporting evidence for hypothesis (1a) that there is a contemporaneous increase in the citation-weighted publications from star help.

One advantage of an event study setting is that it allows visual evidence to support or contradict the parallel trends assumption. Our results in Figure 3.2 depicts that both the homogeneous and heterogeneous models show very little evidence of a pre-trend; therefore, visual inspections support the parallel trends assumption. However, Roth (2022) raises some important concerns about relying on insignificant pre-trends in the event study setting to assess the credibility of parallel trends. In particular, the test may have low power to detect meaningful violations of parallel trends. Roth (2022) suggests a diagnostic approach to determine whether such concerns could be warranted, and we implement these diagnostics in our analysis.

Roth's (2022) technique involves constructing a hypothesised linear violation of parallel trends and then comparing the likelihood of the observed coefficients under the hypothesised trend relative to under-parallel trends. We construct a hypothesised trend with a power of 50%, which would detect a significant pre-trend 50 percent of the time. With a power of 50%, the slope estimate is found to be 0.02 for both the homogenous and heterogeneous effects models. The likelihood ratio reporting the ratio of the likelihood of the observed coefficients under the hypothesised trend relative to under-parallel trends is 0.30 and 0.38, respectively, which supports that the estimated coefficients are more likely observed under a parallel trend. The results of this diagnostic test are also shown in Table 3.3.

Table 3.5 Dynamic star help effects for normalised raw publications (overall, multiple and one-time)

	Overall Star-Help Effect on individual productivity (1)	Multiple Star-Help Effect (2)	One-Time Star-Help Effect (3)
<i>Staracknw_{i,t-3}</i>	-0.00799 (0.0192)	-0.00407 (0.0622)	-0.0113 (0.0202)
<i>Staracknw_{i,t-2}</i>	-0.0227 (0.0181)	-0.0280 (0.0566)	-0.0225 (0.0191)
<i>Staracknw_{i,t}</i>	0.202*** (0.0159)	0.251*** (0.0469)	0.194*** (0.0168)
<i>Staracknw_{i,t+1}</i>	0.0340* (0.0192)	0.173*** (0.0496)	0.0118 (0.0207)
<i>Staracknw_{i,t+2}</i>	0.0218 (0.0207)	0.134** (0.0617)	0.00301 (0.0218)
<i>Staracknw_{i,t+3}</i>	0.0220 (0.0213)	0.110* (0.0581)	0.00604 (0.0229)
Constant	0.447*** (0.0186)	0.589*** (0.0503)	0.424*** (0.0199)
R-squared	0.048	0.064	0.049
Pretest against the hypothesized trend (Roth 2022)			
Power	0.50	0.50	0.50
Hypothesized trend	0.02	0.02	0.02
Bayes factor	0.55	0.55	0.55
Likelihood ratio	0.38	0.26	0.52
Observations	20,841	2,681	18,160
Number of authors	1,258	150	1,108
Author FE	YES	YES	YES
Year FE	YES	YES	YES

Estimation Window: 1996-2017.

The dependent variable is normalised raw publications.

Note: The table reports the estimates based on the model specification in equation 3.1. Also, the 1996 cohort (always treated) is not dropped in this case. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively

Next, we examine hypothesis 2a, which proposes that the effect of star help has a sustained increase in the quality-adjusted output of the author. Table 3.4, Column 1, reports the results based on the sample of the author's productivity that continues to acknowledge a star after the first point acknowledgement. After their initial interaction, the authors who maintain this relationship with a star show a substantial increase in their quality-adjusted publications. Their output is found to increase by 41.62% in the year of acknowledgement and is statistically significant at 1%. These results show the importance of maintaining the relationship with a star. The estimated coefficients

decrease in the years after the first acknowledgement (20.12, 23.74, 23.94 log points) but are significant and persistent.

Additionally, the event study plot in Figure 3.3 left panel shows no evidence of a pre-trend as well as a strong help effect on productivity up to 3 years after the first acknowledgement in the case where multiple acknowledgements take place. This is in contrast to an increase in productivity that is broadly limited to the year of acknowledgement in the right hand panel, which limits the sample to cases of a single acknowledgement. Overall, the evidence suggests that once-off help is associated with a time-limited increase in productivity (supporting hypothesis 1a), and that sustained help is associated with a sustained increase in productivity (supporting hypothesis 2a).

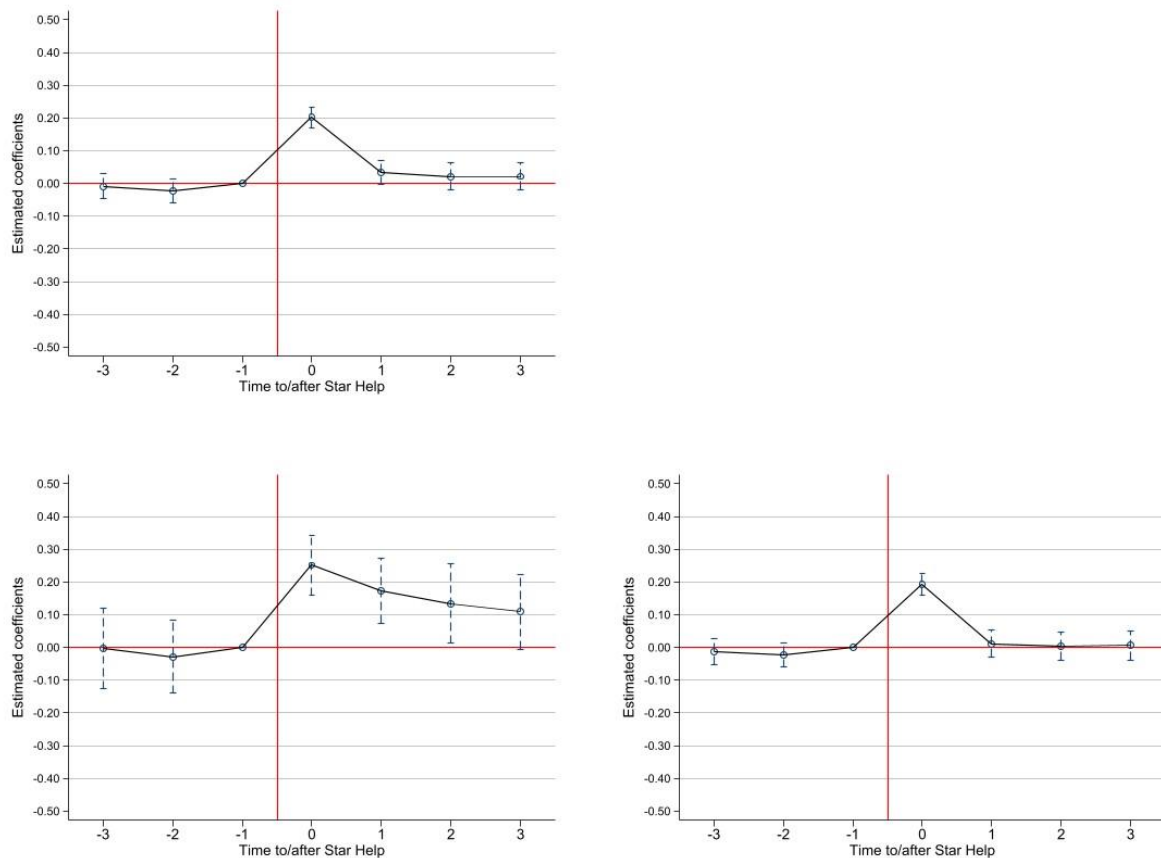


Figure 3.4 Event study model with homogenous star help effects at an individual level: overall star help (row 1); multiple star help (row 2, left); one-time star help (row 2, right)

Note: The figure plots the dynamic effect of the star's helpful interactions with an author at 95% confidence intervals. The event and observation window is from 1996-2017. The dependent variable is raw publications.

3.4.2 Star help effect on individual output (with raw publications as dependent variable)

In our estimation above, we examine the effect of star help on the quality-adjusted productivity of authors, where we use the field-weighted citations as our dependent variable to study this effect. However, there remains a concern that the positive effect of observed star acknowledgement on citation-weighted output found above could reflect a signalling effect, whereby acknowledgement to a star scientist is used as a signal or indicator to convey information about the potential quality of the publication. For the reasons outlined in Section 3.2.2, we augment our analysis using FNTC with an analysis using field-normalised total raw publications to help disentangle the input-enhancing role implied by observed star acknowledgements from a signalling effect. Based on our assumption that the blinded manuscript submitted to a journal for peer review contains no acknowledgement texts and, thus, the decisions by the reviewers cannot be impacted by the acknowledgement to a prominent star scientist, we now present the results from testing hypothesis 1b and 2b using raw publications count.

Column 1 in Table 3.5 reports the estimated coefficients for the baseline regression using authors' field normalised raw publications as the dependent variable. We find that in the year of the acknowledgement event, the raw publication output of authors increased by 22.38% compared to the authors who never acknowledged any help from a star. We also find a decrease in the magnitude of the star-help effect on the raw publications output of the author. These results support hypothesis 1b, where we propose that there will be an increase in the raw publications output in the year of star help. However, similar to our findings for FNTC, we do not find any evidence of a significant effect on raw publications output in the years after the event.

For hypothesis 2b, the results based on the treated sample of authors that acknowledge a star on one occasion and those that acknowledge a star on multiple occasions are presented in Table 3.5 Columns 2 and 3, respectively. Regarding the authors that acknowledge the star for help in the years after the initial event, the results suggest a sustained effect on raw publications output broadly similar to our analysis with field-weighted citation as the dependent variable. An acknowledgement of star help is associated with a 28.53% increase in raw publications in the event year, and the coefficients are found to be statistically significant in the years after. This provides

evidence to support our hypothesis that there is a sustained increase in the raw publications output for authors who maintain an informal collaboration with a star.

Table 3.6 Dynamic star help effects for FNTC (five types of helpful interactions)

	Conceptual (1)	Technical (2)	Materials (3)	Funds & Support (4)	Other Types (5)
<i>Staracknwtype_{i,t-3}</i>	-0.0488 (0.0394)	0.0121 (0.0448)	0.0704 (0.0756)	-0.00899 (0.202)	-0.0211 (0.0551)
<i>Staracknwtype_{i,t-2}</i>	-0.0529 (0.0412)	0.0106 (0.0486)	-0.00552 (0.0589)	0.0200 (0.127)	-0.0299 (0.0496)
<i>Staracknwtype_{i,t}</i>	0.316*** (0.0432)	0.172*** (0.0404)	0.183*** (0.0626)	0.306** (0.124)	0.253*** (0.0569)
<i>Staracknwtype_{i,t+1}</i>	0.0589 (0.0453)	0.0528 (0.0436)	0.0937 (0.0703)	0.0822 (0.121)	-0.0172 (0.0581)
<i>Staracknwtype_{i,t+2}</i>	-0.00485 (0.0426)	0.0224 (0.0434)	0.191** (0.0746)	0.138 (0.137)	0.0428 (0.0661)
<i>Staracknwtype_{i,t+3}</i>	0.0627 (0.0502)	0.0333 (0.0495)	0.0947 (0.0704)	0.0505 (0.113)	-0.0559 (0.0616)
Constant	0.365*** (0.0345)	0.379*** (0.0388)	0.315*** (0.0545)	0.310*** (0.107)	0.357*** (0.0578)
R-squared	0.046	0.021	0.033	0.047	0.048
Pretest against the hypothesized trend (Roth 2022)					
Power	0.50	0.50	0.50	0.50	0.50
Hypothesized trend	0.04	0.04	0.07	0.18	0.05
Bayes factor	0.55	0.55	0.55	0.55	0.55
Likelihood ratio	1.40	0.08	0.03	0.19	0.43
Observations	7,665	4,942	3,001	947	4,152
Number of authors	468	300	176	60	248
Author FE	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES

Estimation Window: 1996-2017.

The dependent variable is Field-Normalised Total Citations.

Note: The table reports the estimates based on the model specification in equation 3.4. Also, the 1996 cohort (always treated) is not dropped in this case. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively.

In contrast, in considering the effect of star help on the raw publications output for the authors who acknowledge a star for help in a single event, the results are aligned with our prior analysis using field-weighted citations as the dependent variable in that we do not find any evidence for a statistically significant effect in the years after the event. Figure 3.4 illustrates the event study plots for the two cases. Overall, by separating the input-enhancing role implied by observed star acknowledgements from a signalling effect using the raw publications measure, the results indicate that the positive effect of

the observed star acknowledgement on output reflects more than just a signalling effect in both the cases of once-off acknowledgement and multiple acknowledgements.

3.4.3 Effects of star help through different channels

In this Section, we consider hypothesis 3, which states that the effect of star help is different in magnitude across the different types of star help identified based on the keywords contained in the acknowledgement text for the publication. We identify five types of acknowledgement to understand the effects of the helpful interaction between the author and the star: Conceptual, Technical, Material, Funds and Support, and Other Types. Our dataset is split into five cohorts, each containing treated and control authors identified from the CEM procedure. This analysis estimates the effect of a star help that happens through five different channels.

For each channel, we modify equation 3.1

$$\begin{aligned} \ln Y_{it} = & \alpha + \beta_{\leq -4} staracknwtype_{i,-4} + \sum_{j=-3}^{-2} \beta_j staracknwtype_{ij} \\ & + \sum_{j=0}^3 \beta_j staracknwtype_{ij} + \beta_{\geq 4} staracknwtype_{i,4} + \delta_t + \mu_i + \epsilon_{it}, \end{aligned} \quad (3.4)$$

where $staracknwtype_{ij}$ is a binary variable that indicates if a star is acknowledged for the specific type of help in year j by an author i .

The results from the estimates are presented in Table 3.6. They show that a conceptual acknowledgement to a star is associated with an increase in output of 37.16%, and this is the most significant increase in output compared to all other types examined in this analysis. The help from the star through sharing their knowledge positively impacts the author's productivity. In contrast, a technical acknowledgement type is associated with the smallest effects on output, with an increase of 18.77%. This acknowledgement type is observed primarily in subject fields such as medicine and biochemistry, where the research methods involve practical examination and laboratory experiments, and stars can provide their direct expertise. Material acknowledgement to a star is associated with an increase in output of 20.08% in the year of acknowledgement

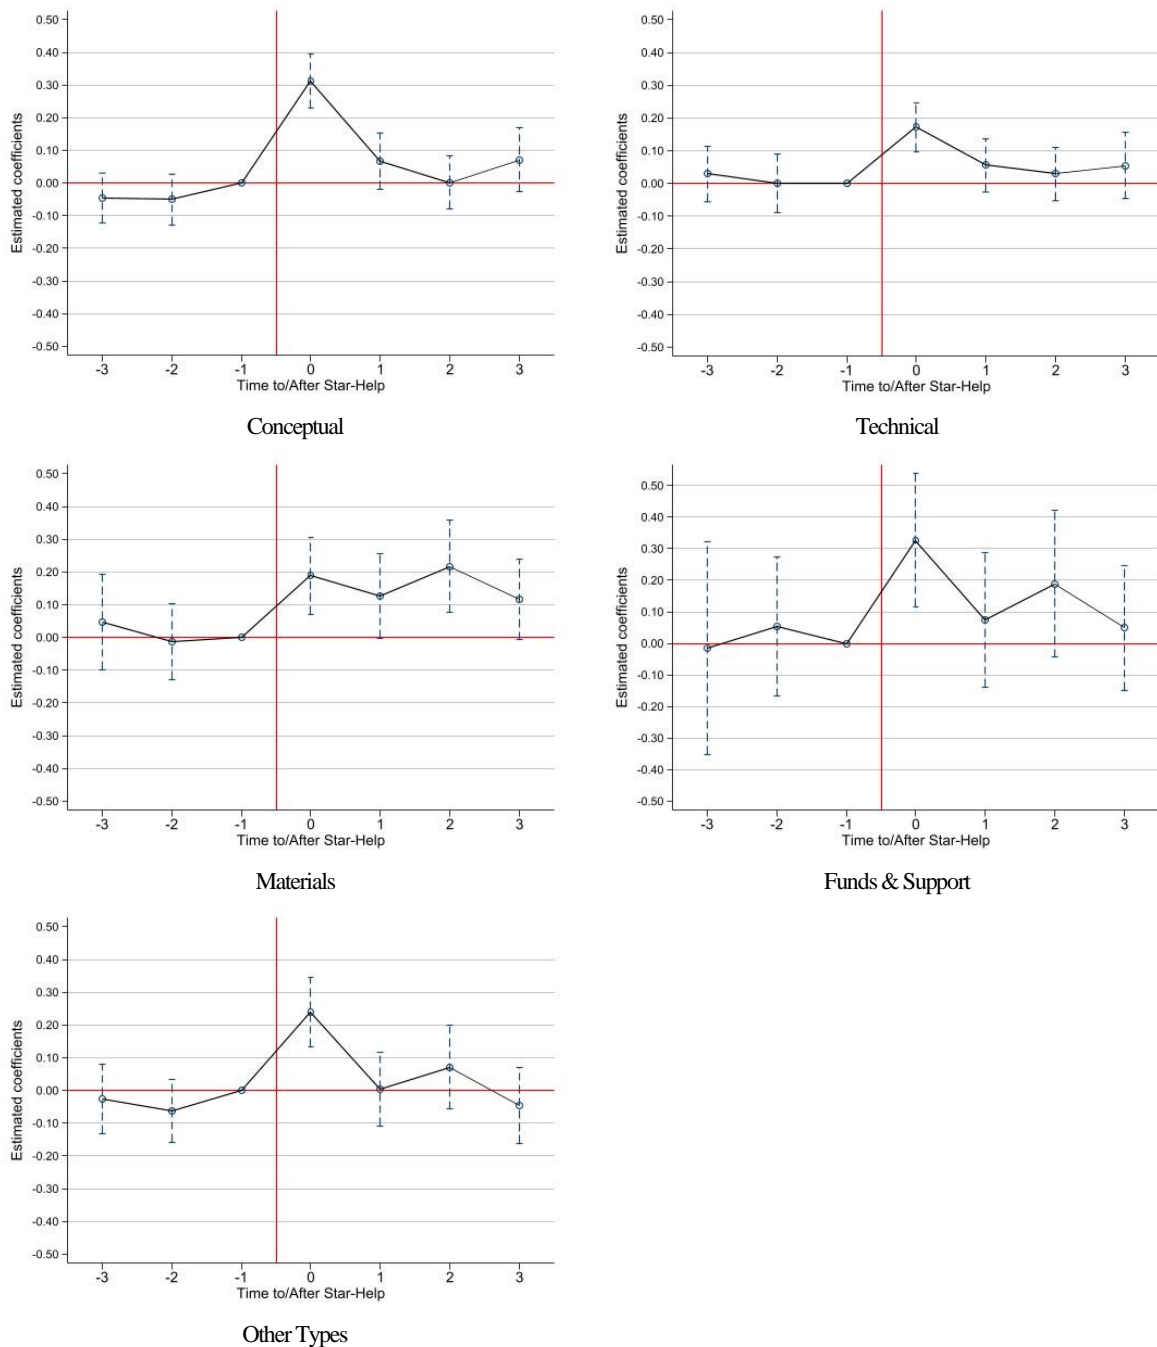


Figure 3.5 Event study model with channels of help identified from the acknowledgement texts.

Note: The figure plots the dynamic effect of the star's helpful interactions with an author at 95% confidence intervals. The event and observation window is from 1996-2017. The dependent variable is field normalised total citations.

and unlike the other types of help, there was a 21.05% increase in output in the year afterwards. Also, the event study plot for material acknowledgement in Figure 3.5 shows a persistent productivity rise in the years after star acknowledgement, which contrasts with the absence of a sustained effect for the other types of acknowledgement. Material acknowledgements to a star are observed mainly in the fields of medicine and

Table 3.7 Summary statistics: control and treated group (k-to-k matched) (a) balanced panel, (b) alternate matching criteria using the availability of publication match at t=0, (c) alternate matching criteria using acknowledgement text availability in publication at t=0.

Variable	Control	Treated	Diff in mean	P-value
<i>(a) A balanced panel of 520 matched authors, with 260 in each group</i>				
Year	2005.080	2005.060	0.027	0.961
Subject	10.219	10.219	0.000	1.000
Country	2.054	2.054	0.000	1.000
Total Career Age	30.450	30.727	-0.277	0.632
Total Career Age Bins	7.738	7.738	0.000	1.000
Cumulative Publication experience	13.515	13.554	-0.038	0.945
Cumulative Publication experience Bins	3.293	3.300	-0.007	0.953
Cumulative citations received per cumulative publications	43.732	44.364	-0.632	0.814
Cumulative citations received per cumulative publications Bins	2.246	2.246	0.000	1.000
<i>(b) An unbalanced panel of 1208 matched authors, with 604 in each group</i>				
Year	2007.331	2007.296	0.035	0.920
Subject	10.942	10.942	0.000	1.000
Country	2.013	2.013	0.000	1.000
Total Career Age	20.608	20.755	-0.147	0.801
Total Career Age Bins	5.750	5.750	0.000	1.000
Cumulative Publication experience	9.353	9.262	0.091	0.808
Cumulative Publication experience Bins	2.414	2.414	0.000	1.000
Cumulative citations received per cumulative publications	43.560	44.257	-0.697	0.943
Cumulative citations received per cumulative publications Bins	2.007	2.007	0.000	1.000
<i>(c) An unbalanced panel of 990 matched authors, with 495 in each group</i>				
Year	2007.911	2007.889	0.022	0.953
Subject	10.820	10.820	0.000	1.000
Country	1.974	1.974	0.000	1.000
Total Career Age	20.198	20.418	-0.220	0.729
Total Career Age Bins	5.677	5.677	0.000	1.000
Cumulative Publication experience	9.479	9.543	-0.065	0.877
Cumulative Publication experience Bins	2.471	2.471	0.000	1.000
Cumulative citations received per cumulative publications	34.323	34.262	0.062	0.966
Cumulative citations received per cumulative publications Bins	1.881	1.881	0.000	1.000

Notes: Reports the t-test for the mean difference between control and treated groups one year before forming the star-help relation.

biochemistry, where stars can conveniently share unpublished data and loan specimens. This could suggest that the star and the author develop a more persistent connection through the materials-sharing channel than through the other acknowledgement channels.

A project funded by the star's support also impacts positively on an author's productivity, with a funding acknowledgement to a star found to be associated with a 35.80% increase in output in the year of acknowledgement. These findings suggest that a star who supports the research through financial means also has an effect on the author's productivity. The number of these acknowledgements are fewer when compared to other types of interaction channels in our data. Finally, all the other types of acknowledgement that could not be accurately classified under conceptual, technical, materials, or funds and support are categorised as 'Other Types' for the purpose of this study. In the year of these type of acknowledgement events, the author's quality-adjusted output increases by 28.79%, similar to our overall star acknowledgement effect in Section 3.4.1. These results suggest that the 'Other' category proxies for the other more concrete forms of help the stars provide. Alternatively, the relative robustness of the results across acknowledgement types might indicate that the precise form of help is less important than the close interaction with the star. Overall, these results provide evidence to support hypothesis 3 that the effect of star help on the author's quality-adjusted productivity²⁶ is present, albeit variable across the types of help. Furthermore, we note that the estimated coefficients are statistically significant in the acknowledgement event year for each type, which can be taken as further support for hypothesis 1a.

3.5 Robustness

3.5.1 Robustness to a balanced panel

Our baseline estimation uses an unbalanced panel of 1258 authors to estimate the effects on the dependent variable. In a balanced panel, we observe units (in this case,

²⁶ We also check the effect of star help on the author's productivity on raw publication output (Hypothesis 2). We find a similar trend in the results for each helpful interaction with less magnitude- See Appendix B.4

Table 3.8 Robustness test of dynamic star help effects for FNTC (balanced panel and alternative matching criteria)

	Balanced panel (1)	Alternative matching using the availability of publication match in control at t=0 (2)	Alternative matching using availability of acknowledgement text match in control at t=0 (3)
<i>Staracknw_{i,t-3}</i>	-0.0188 (0.0421)	-0.0115 (0.0258)	-0.00760 (0.0253)
<i>Staracknw_{i,t-2}</i>	-0.0304 (0.0405)	-0.0279 (0.0243)	-0.0225 (0.0237)
<i>Staracknw_{i,t}</i>	0.240*** (0.0414)	0.245*** (0.0247)	0.249*** (0.0247)
<i>Staracknw_{i,t+1}</i>	0.0946** (0.0417)	0.0527** (0.0263)	0.0648** (0.0274)
<i>Staracknw_{i,t+2}</i>	0.0743* (0.0437)	0.0493* (0.0269)	0.0579** (0.0280)
<i>Staracknw_{i,t+3}</i>	0.105** (0.0451)	0.0437 (0.0288)	0.0462 (0.0300)
Constant	0.502*** (0.0245)	0.405*** (0.0224)	0.389*** (0.0245)
R-squared	0.024	0.029	0.034
Pretest against the hypothesized trend (Roth 2022)			
Power	0.50	0.50	0.50
Hypothesized trend	0.04	0.02	0.02
Bayes factor	0.55	0.55	0.55
Likelihood ratio	0.40	0.36	0.34
Observations	11,440	20,023	16,356
Number of authors	520	1,208	990
Author FE	YES	YES	YES
Year FE	YES	YES	YES

Estimation Window: 1996-2017.

The dependent variable is Field Normalised Total Citations.

Note: The table reports the estimates based on the model specification in equation 3.1. Also, the 1996 cohort (always treated) is not dropped in this case. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively

authors) every time period, reducing the noise introduced by unit heterogeneity. To analyse whether this variation in the dataset affects our results, we use a panel of 520 authors present in the data throughout our estimation period (1996-2017). The matching statistics of these 520 authors are reported in Table 3.7 (a). We present the estimated results of the event study specification based on the balanced panel in Table 3.8 Column 1. Similar to the star acknowledgement associated with the unbalanced panel data, we again observe a significant contemporaneous effect (27.12%). Though the estimated coefficients in the further years fall (9.46, 7.43, 10.5 log points), they are

statistically significant up to three years after the year of acknowledgement. The event study plot in Figure 3.6 again shows no evidence of a pre-trend. In addition, Roth's pretest diagnostic analysis reports a likelihood ratio of 0.40, suggesting that the estimated coefficients are likely to follow a parallel trend before the treatment.




Figure 3.6 Robustness test on event study model with homogenous star help effects at an individual level for a balanced panel

Note: The figure plots the dynamic effect of star interactions with an author at 95% confidence intervals. The event and observation window is from 1996-2017. The dependent variable is field normalised total citations.

3.5.2 Robustness to alternative matching criteria

As a further robustness check, we examine the sensitivity of the results to alternative matching criteria. One concern is that authors are unlikely to publish their research work annually throughout their careers. It takes time to publish a paper following the review process and editing. As an acknowledgement can only occur in a year a paper is published, acknowledgements may be simply partly picking up the fact that a paper was published in a particular year, thus biasing our estimate of the productivity effect of an acknowledge relative to the control authors. To this end, we identify the authors in the control group that have also published at least one paper in the event year ($t = 0$) when the authors in the treated group have a publication that acknowledges a star. Table 3.7 (b) reports the matching statistics of the new set of control and treated authors. Here we find 1208 matching authors that satisfy the matching criteria. In Table 3.8, Column 2, we present the estimated results from the event study specification. In considering the additional matching criteria, we still find a similar contemporaneous effect from star acknowledgement on output. An increase of 27.76% in the quality-adjusted output is estimated once the author gets exposure from the star, which is similar to our baseline results.




Figure 3.7 Robustness test on event study model with homogenous star help effects at the individual level using alternate matching criteria; the matched control author should have a publication in year $t=0$ (left), and the control author should have a publication with acknowledgement text in year $t=0$ (right), where year $t=0$ is the year when the treatment starts.

Note: The figure plots the dynamic effect of the star's helpful interactions with an author at 95% confidence intervals. The event and observation window is from 1996-2017. The dependent variable is field normalised total citations.

As discussed previously in Section 3.3.2, an additional concern is that the availability of acknowledgement text from the Scopus database is limited. The credibility of the star help effects can be questioned due to the comparison between a treated author who gets a match from the control group and has no acknowledgement text available in that particular year. To check the possible implications for our results, we compare the authors in the control group that have a publication with an available acknowledgement text in the event year ($t = 0$) when the treated author acknowledges a star. Table 3.7 (c) reports the matching statistics of 990 authors at year $t - 1$ with these additional criteria. The event study estimation reported in Table 3.8, Column 3 again shows similar results to our analysis in Section 3.4.1. We find a contemporaneous increase in the quality-adjusted output of the author by 28.27% associated with star help. Also, the plots for the estimation under both alternative matching criteria are presented in Figure 3.7, and they indicate no evidence of a pre-trend. Furthermore, the likelihood ratios of observing any pre-trend are 0.36 and 0.34 from the Roth test analysis, which supports the appropriateness of the parallel trend assumption.

Table 3.9 Robustness test of dynamic star help effects for FNTC (country level)

	Ireland (1)	Denmark (2)	New Zealand (3)
<i>Staracknw</i> _{<i>i,t-3</i>}	-0.0881 (0.0620)	0.0184 (0.0306)	-0.0254 (0.0650)
<i>Staracknw</i> _{<i>i,t-2</i>}	-0.00941 (0.0562)	-0.0322 (0.0295)	-0.0107 (0.0575)
<i>Staracknw</i> _{<i>i,t</i>}	0.213*** (0.0636)	0.271*** (0.0288)	0.198*** (0.0574)
<i>Staracknw</i> _{<i>i,t+1</i>}	-0.0749 (0.0754)	0.0782*** (0.0287)	0.0701 (0.0599)
<i>Staracknw</i> _{<i>i,t+2</i>}	0.0288 (0.0825)	0.0623** (0.0293)	0.0117 (0.0611)
<i>Staracknw</i> _{<i>i,t+3</i>}	-0.0825 (0.0712)	0.101*** (0.0339)	-0.0433 (0.0599)
Constant	0.330*** (0.0581)	0.418*** (0.0283)	0.443*** (0.0545)
R-squared	0.042	0.026	0.044
Pretest against the hypothesized trend (Roth 2022)			
Power	0.50	0.50	0.50
Hypothesized trend	0.06	0.03	0.06
Bayes factor	0.55	0.55	0.55
Likelihood ratio	2.30	0.05	0.34
Observations	3,851	12,873	4,117
Number of authors	226	784	248
Author FE	YES	YES	YES
Year FE	YES	YES	YES

Estimation Window: 1996-2017.

The dependent variable is Field Normalised Total Citations.

Note: The table reports the estimates based on the model specification in equation 3.1. Also, the 1996 cohort (always treated) is not dropped in this case. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively

3.5.3 Robustness to country-specific cohorts

The publication data from Scopus comprises those authors that have published in Ireland, Denmark, and New Zealand, and the stars we identified are those individuals who are highly productive relative to their peers in these three countries. Therefore we should be able to see a similar effect of star help on the author's productivity in the year of star acknowledgement. The results reported in Table 3.9 and presented graphically in Figure 3.8 show some differences across countries – with the largest acknowledgement effect observed for Denmark and the smallest for New Zealand – but overall, we find evidence of a significant productivity effect in the year of acknowledgement.

3.5.4 Robustness to co-authorship relations

While our baseline estimation excludes authors that have co-authored with a star prior to the star acknowledgement event from the treatment group and the control group comprises of matched authors that have never co-authored or acknowledged a star scientist at any time in the sample, there could still be the potential for endogeneity from

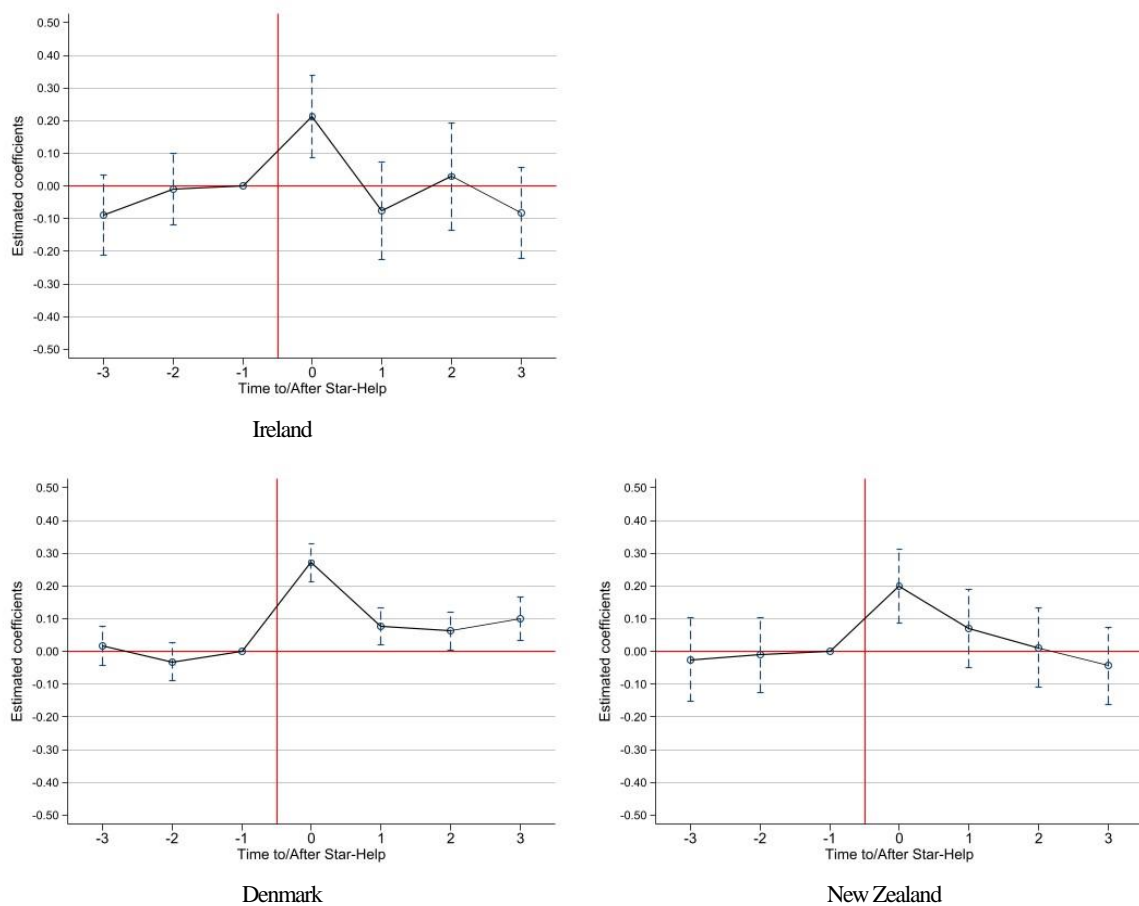


Figure 3.8 Event study model with homogenous star help effects at an individual level by country

Note: The figure plots the dynamic effect of the star's helpful interactions with an author at 95% confidence intervals. The event and observation window is from 1996-2017. The dependent variable is raw publications.

other types of social connection between the author and the star prior to the acknowledgement event and separate to direct co-authorship. For example, an author might have had an indirect social relationship with the star by co-authoring with a non-star and thus it would be important to control for social distance in the estimation more generally.

Table 3.10 Robustness test of dynamic star help effects for FNTC (excluding the authors with previous indirect relations with a star (column1) and excluding the co-authored publications with a star after the initial acknowledgement event (column 2 & 3))

	Indirect Relationship with a star before the event (1)	One-Time Star Help Effect: Excluding the publications (2)	Multiple Star Help Effect: Excluding the publications (3)
<i>Staracknw</i> _{<i>i,t-3</i>}	-0.0122 (0.0251)	0.00172 (0.0269)	-0.110 (0.0752)
<i>Staracknw</i> _{<i>i,t-2</i>}	-0.0245 (0.0244)	-0.0236 (0.0253)	-0.0217 (0.0743)
<i>Staracknw</i> _{<i>i,t</i>}	0.250*** (0.0250)	0.235*** (0.0260)	0.330*** (0.0624)
<i>Staracknw</i> _{<i>i,t+1</i>}	0.0589** (0.0266)	0.00767 (0.0273)	0.129** (0.0634)
<i>Staracknw</i> _{<i>i,t+2</i>}	0.0556** (0.0274)	-0.00126 (0.0276)	0.196*** (0.0707)
<i>Staracknw</i> _{<i>i,t+3</i>}	0.0463 (0.0292)	-0.00401 (0.0284)	0.203** (0.0827)
Constant	0.402*** (0.0238)	0.386*** (0.0244)	0.540*** (0.0692)
R-squared	0.029	0.028	0.056
Pretest against the hypothesized trend (Roth 2022)			
Power	0.50	0.50	0.50
Hypothesized trend	0.02	0.07	0.02
Bayes factor	0.55	0.55	0.55
Likelihood ratio	0.38	2.90	0.14
Observations	19,021	18,160	2,681
Number of authors	1,146	1,108	150
Author FE	YES	YES	YES
Year FE	YES	YES	YES

Estimation Window: 1996-2017.

Note: The table reports the estimates based on the model specification in equation 3.1. Also, the 1996 cohort (always treated) is not dropped in this case. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively

To explore this possibility, we examine the network of co-authors for the treated authors before the acknowledgement of star help (before $t = 0$) to identify the individuals not directly connected to the star through co-authorship but sharing a mutual co-authorship connection with someone else in the second degree of separation. We identify a total of 56 out of the 629 treated authors reported in Table 3.2 that share a mutual co-authorship connection with the star. Then in a robustness check, we exclude these 56 treated authors and their matched control pairs from the analysis and re-estimate the baseline regression. Column 1 in Table 3.10 presents the estimated coefficients after their exclusion. Consistent with the baseline results in Table 3.3, we find

that acknowledging star help is associated with a 28.40% (0.25 log points) increase in the quality-adjusted output of the author in the year of acknowledgement. In addition, Figure 3.9 shows the estimated event study plot after the exclusion of the 56 treated authors and their matched pairs, and it also indicates that the baseline results are robust to their exclusion.




Figure 3.9 Robustness test on event study model with homogenous star help effects at an individual level: excluding the authors with indirect star relationship before the acknowledgement event.




Figure 3.10 Robustness test on event study model with homogenous star help effects at an individual level: excluding the future star co-authored publications for authors who acknowledge star help for: one-time (left); multiple times in the years after the first point of contact (right).

Note: The figure plots the dynamic effect of star interactions with an author at 95% confidence intervals. The event and observation window is from 1996-2017. The dependent variable is field normalised total citations.

Another concern for our baseline results is the potential emergence of a co-authorship relationship between the author and the star post the acknowledgement event. This could have important implications for our results where the estimated increase in author output associated with the acknowledgement of star help could instead be partially attributed to the later emergence of a direct co-authorship relationship between author and star. Authors who acknowledge a star for their help

could also produce more co-authored publications with the same star or indeed any star in the years after the event ($t = 0$). In a further test of the robustness of our results, we examine the effect of once-off and multiple-star help on author's output, excluding all co-authorship publications that occur between a star and author after the acknowledgement event from our estimation.²⁷

Table 3.11 Robustness test of dynamic star help effects for FNTC (comparing the initial productivity effects of authors who acknowledge stars- single and multiple times).

	Overall Star Help Effects (1)		Star Help Effects ;Interacted with Dummy Variable for Multiple Star Help (2)
<i>Staracknw_{i,t-3}</i>	0.000487 (0.0269)	<i>M.Staracknw_{i,t-3}</i>	-0.0962 (0.0812)
<i>Staracknw_{i,t-2}</i>	-0.0243 (0.0253)	<i>M.Staracknw_{i,t-2}</i>	0.00500 (0.0795)
<i>Staracknw_{i,t}</i>	0.237*** (0.0260)	<i>M.Staracknw_{i,t}</i>	0.0874 (0.0688)
<i>Staracknw_{i,t+1}</i>	0.0312 (0.0274)	<i>M.Staracknw_{i,t+1}</i>	0.132* (0.0718)
<i>Staracknw_{i,t+2}</i>	0.0196 (0.0277)	<i>M.Staracknw_{i,t+2}</i>	0.190** (0.0772)
<i>Staracknw_{i,t+3}</i>	0.0129 (0.0288)	<i>M.Staracknw_{i,t+3}</i>	0.172** (0.0860)
Constant	0.409*** (0.0238)		
R-squared	0.030		
Observations	20,841		
Number of authors	1,258		
Author FE	YES		
Year FE	YES		

Estimation Window: 1996-2017.

The dependent variable is Field Normalised Total Citations.

Note: The table reports the estimates based on the model specification in equation 3.5. Also, the 1996 cohort (always treated) is not dropped in this case. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively

Column 2 in Table 3.10 reports the results for one-time star help after excluding these publications. Again, the results are similar to the baseline and also have a similar dynamic pattern (see Table 3.4, Column 2 and Figure 3.3 (right) for comparison). Overall,

²⁷ In an additional robustness check, we also examine whether our results are robust when we exclude the 179 matched pairs of treated and control authors from the analysis that co-author with a star after the acknowledgement event rather than just excluding the publications, and we find similar supporting evidence for the baseline results – See appendix.

they show a star help effect of 26.5% for the quality-adjusted productivity of the author in the year $t = 0$ as well as no evidence of a sustained productivity effect in the following years. Furthermore, Column 3 in Table 3.10 reports the estimated coefficients for multiple star help after excluding the post acknowledgement event star co-authorship publications from the analysis. Here, the star help effect is estimated to be a 39.1% increase in the quality-adjusted output of the author. Additionally, in the years following the initial acknowledgement event, the results show statistically significant coefficients broadly comparable with the baseline findings but with a smaller positive impact at each lag. Figure 3.10 also shows the event study plots of the star help effect for both once-off and multiple star help after removing these co-authored publications, and these are comparable to the baseline plots in Figure 3.3. Therefore, our initial results are largely robust to the exclusion of co-authored publications between an author and star in the years after the acknowledgement event.

3.5.5 Robustness to single and multiple acknowledging authors

To test whether the effects of a star acknowledgement are different where the star is acknowledged multiple times, we define a dummy variable M that takes a value of 0 if the star is acknowledged just once and 1 where if the star is acknowledged more than once. This leads to a revised estimating equation:

$$\begin{aligned}
\ln Y_{it} = & \alpha + \beta_{\leq -4} \text{staracknw}_{i,-4} + \sum_{j=-3}^{-2} \beta_j \text{staracknw}_{ij} + \sum_{j=0}^3 \beta_j \text{staracknw}_{ij} \\
& + \beta_{\geq 4} \text{staracknw}_{i,4} + \gamma_{\leq -4} (\text{staracknw}_{i,-4} \times M) \\
& + \sum_{j=-3}^{-2} \gamma_j (\text{staracknw}_{ij} \times M) + \sum_{j=0}^3 \gamma_j (\text{staracknw}_{ij} \times M) \\
& + \gamma_{\geq 4} (\text{staracknw}_{i,4} \times M) \delta_t + \mu_i + \epsilon_{it},
\end{aligned} \tag{3.5}$$

At any given lead or lag, a simple test of difference in effect for the cases of single or multiple acknowledgements is then a test of the statistical significance of the relevant γ coefficient. A positive and statistically significant coefficient in period 0 (the period in which the acknowledgement occurs) would indicate that even in the period in which the acknowledgement occurs, a scientist that makes multiple acknowledgements receives a greater productivity boost than a scientist who only makes a single acknowledgement. This could reflect unobserved heterogeneity between single and multiple acknowledging

scientists (that is not picked up by our controls) or indicate that the quality of initial help received is greater for multiple acknowledging scientists. However, a finding of no statistically different initial productivity effects of acknowledgements between the two groups helps allay concerns of unobserved heterogeneity between the single and multiple citing scientists.

Table 3.12 Summary statistics: by quartile of the cumulative FNTC distribution for the treated group at year t-1.

Quartile	N	Mean	SD	Median	Min	Max	IQR
1	2943	.544	0.312	.589	0	1.017	.58
2	2938	1.452	0.224	1.474	1.021	1.828	.364
3	2926	2.252	0.255	2.23	1.833	2.738	.445
4	2933	3.556	0.816	3.298	2.744	6.705	.959

Table 3.13 Robustness test of dynamic star help effects for FNTC (for quartile levels)

	Quartile-1 (1)	Quartile-2 (2)	Quartile-3 (3)	Quartile-4 (4)
<i>Staracknw_{i,t-3}</i>	0.0246 (0.0224)	-0.0641 (0.0415)	-0.119** (0.0515)	0.0981 (0.0891)
<i>Staracknw_{i,t-2}</i>	-0.0107 (0.0189)	-0.0224 (0.0408)	-0.0893* (0.0510)	0.0125 (0.0888)
<i>Staracknw_{i,t}</i>	0.323*** (0.0270)	0.297*** (0.0463)	0.184*** (0.0541)	0.127 (0.0774)
<i>Staracknw_{i,t+1}</i>	0.0770*** (0.0252)	0.0837* (0.0431)	0.0301 (0.0588)	-0.0212 (0.0901)
<i>Staracknw_{i,t+2}</i>	0.159*** (0.0333)	0.0482 (0.0455)	-0.0246 (0.0574)	-0.0691 (0.0852)
<i>Staracknw_{i,t+3}</i>	0.110*** (0.0304)	0.0395 (0.0500)	-0.0317 (0.0651)	-0.0144 (0.0922)
Constant	0.136*** (0.0252)	0.280*** (0.0300)	0.520*** (0.0424)	0.772*** (0.0622)
R-squared	0.059	0.056	0.038	0.022
Observations	5,507	5,348	5,069	4,917
Number of authors	420	328	270	240
Author FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES

Estimation Window: 1996-2017.

The dependent variable is Field Normalised Total Citations.

Note: The table reports the estimates based on the model specification in equation 3.1. Also, the 1996 cohort (always treated) is not dropped in this case. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively

We report the event study estimates of the model in Table 3.11. In Column 1, we show the estimated beta coefficients from Equation 3.5, while in Column 2, we show the estimated gamma coefficients on the interaction between the indicator variable for star help and the multiple acknowledgement dummy. Most importantly, in the year of star help ($j = 0$), we don't find any evidence of a significantly different initial productivity effect for the authors who acknowledge a star multiple times. We take this as evidence that the observation of multiple acknowledgements does not imply that the scientist is more productive and/or can make better use of star help, suggesting that the sustained productivity effect associated with multiple acknowledgements reflects a causal effect of the help rather than being due to any selection effect.

3.5.6 Robustness to the productivity of acknowledging authors

In a final robustness test, we examine for potential heterogeneous treatment effects across the distribution of the outcome variable. The effect on productivity from acknowledging star help could differ depending on the productivity of the acknowledging authors. To investigate this, we look at the cumulative FNTC of the treated authors one year before treatment, the acknowledgment to a star. This productivity indicator provides a measure of where an author is in the initial productivity distribution before treatment. Table 3.12 reports summary statistics for each of the four quartiles of treated authors from the distribution of cumulative FNTC at year $t - 1$. We divide our overall sample of treated authors into four sub-samples based on these quartiles and then, using Equation 3.1, we estimate the results for each sub-sample of treated authors and their matched pairs from the control group.

The event study estimates for each quartile are reported in Table 3.13 and their related event-study plots are presented in Figure 3.11. The results show that the contemporaneous effect of star help on author productivity for authors in the first and second quartiles (column 1 and 2) is larger compared to the effect for authors in the third and fourth quartiles (columns 3 and 4) of the cumulative FNTC distribution. Star help is associated with a statistically significant increase of 38.13% for authors in quartile 1 and a statistically insignificant increase of 13.54% for authors in quartile 4. The event study plots also indicate that the effect is sustained in the years after the acknowledgement event for quartile 1 authors on average. In contrast, it is not sustained for the authors in

later quartiles. In comparing the low- and high-productivity author clusters, we thus find that star help has a greater effect on authors that have low productivity before the star help is received relative to the authors that have high initial productivity. We interpret this result as indicating that lower productivity authors gain most from star help. We hope to further investigate this finding in future work, as it may have important implications for policies that can be used to support lower-productivity colleagues.

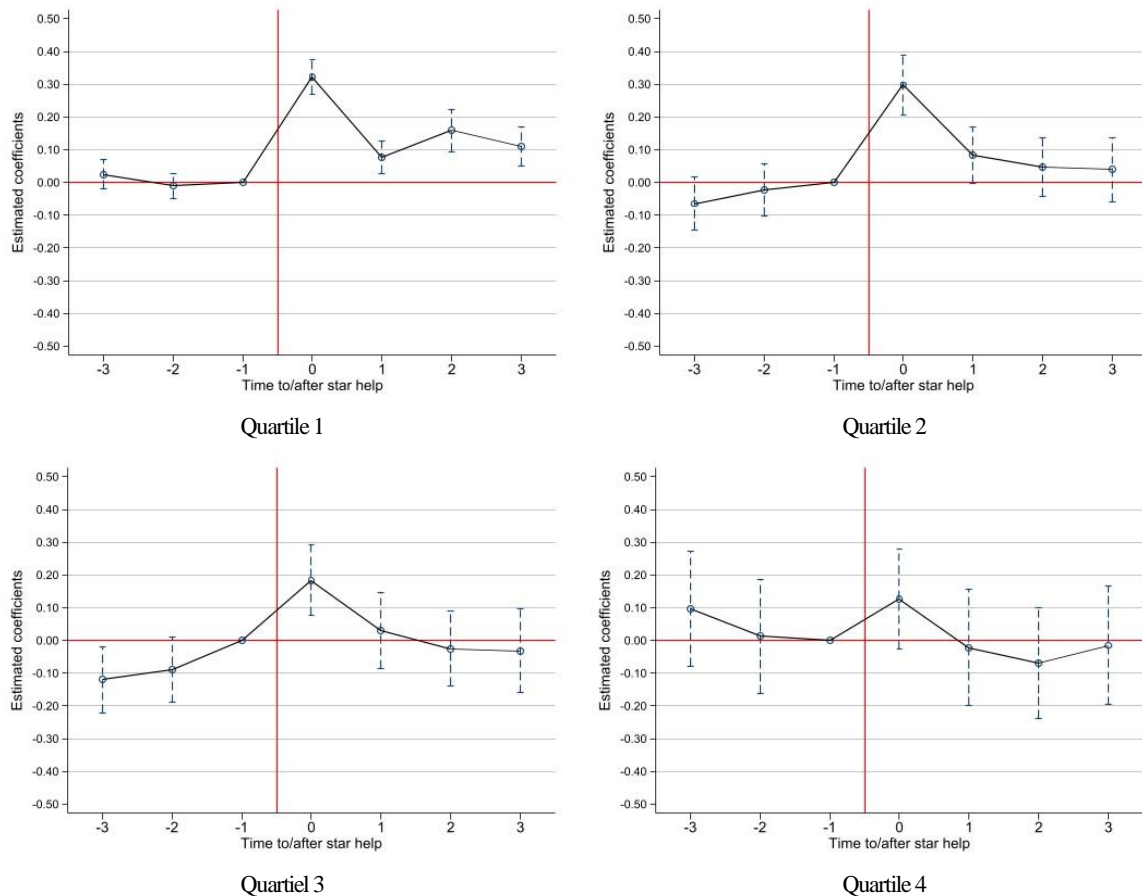


Figure 3.11 Robustness test on event study model with homogenous star help effects at an individual level comparing low-high productivity clusters by equal four quartiles of cumulative FNTC received for the treated authors one year before the treatment

Note: The figure plots the dynamic effect of the star's helpful interactions with an author at 95% confidence intervals. The event and observation window is from 1996-2017. The dependent variable is field normalised total citations.

3.6 Summary and Policy Implications

Although there is a growing body of literature that estimates the effects of connections to stars on the productivity of peers, a star's role in helping their peers through non-co-author-based relationships has been largely neglected. While focusing on

the performance gains of co-authoring with a helpful star, Oettl (2012) discusses the possibility of extending the research to colleagues and students who are not co-authoring with the star. Following this suggestion, this chapter investigates the effects of star help on the productivity of scientists receiving the help.

To implement this empirically, we identify interactions between a star and an author identified from the acknowledgement texts in a publication. We treat these acknowledgements as an indicator of helpful interaction and examine their impacts on the author's quality-adjusted scientific output.

We analyse the output of authors who publish in three countries: Ireland, Denmark, and New Zealand. Using Natural Language Processing (NLP) techniques, we extract the names of individuals who get acknowledged for their contributions to the research. We match these individuals to previously identified stars' names relative to the peers who published in the three countries. We estimate the effect of star help by comparing scientists who acknowledge the star with a carefully matched control group of scientists who do not.

The overall effect of star help is estimated in an event study setting, where the event is the acknowledgement by an author of an identified star. We observe an increase in the quality-adjusted output of authors in the event year. Although we find evidence of a sharp fall in the productivity effect in subsequent years, the initial effect of star interaction is significant. In addition, analysis of authors who continue to acknowledge a star in the years after the initial acknowledgement shows a higher and sustained helpfulness effect on their output.

In addition, we examine the effects of different types of star help. We classify these interactions from a star based on the keywords in the acknowledgement text. We identify five interaction channels through which the knowledge transfer occurs. A disaggregated analysis of these interactions suggests that all types of help positively affect the author's output. However, authors who acknowledge a star for the conceptual help show the highest increase in the output in the year of acknowledgement. Unlike other forms of star help, authors who acknowledge a star for materials received show consistent productivity effects even without evidence of sustained acknowledgements.

A second contribution of our analysis is to examine star-help effects within an explicitly dynamic framework. Our staggered event-study findings show significant productivity effects in the year that star help is acknowledged. However, the observed effects drop sharply in the years following the acknowledgement unless that help is sustained. Moreover, these findings are robust to recent techniques that explicitly account for heterogeneity across the years in which the star help is observed to occur.

Previous work has established that the occurrence and nature of star help can be an important mediating factor in determining the effects of co-authorship relationships with a star (e.g., Oettl, 2012). Our study focuses on the productivity effects of star help in the absence of a co-author relationship. We find a robust relationship across forms of star help between the acknowledgement of help and the productivity of non-co-authoring researchers. Moreover, the effects on productivity are more significant where there is evidence of the help being sustained over time and where the researcher receiving the help is positioned in the lower quartiles of the relevant field-specific productivity distribution.

These results have important implications for the recruitment and organisational strategies of academic departments. To the extent that stars are co-located with non-star researchers, they should have more opportunities to provide help to the benefit of incumbents at the receiving department, with the beneficial effects being very impactful for less productive researchers. This points to one source of potential value from star recruitment policies in addition to the more widely studied co-authorship and recruitment-quality channels, at least to the extent that co-located stars are better positioned to provide help to their non-star colleagues (Agrawal et al., 2017). Furthermore, as the results are not specific to co-located researchers, organisational policies that help embed researchers in networks that increase the probability of interaction with stars should have productivity benefits. Where the star and non-star are co-located, formal mentoring programmes and more informal means such as regular departmental seminars, workshops and social gatherings should help initiate and develop the relationships that support the provision of star help. Star help can also be provided even where the star and non-star are located in different institutions. In this regard, it is likely beneficial for the non-stars to embed in broader networks that allow for interactions with stars, thus supporting the development of the relationships that

facilitate productivity-enhancing star help. This suggests the importance of organisational support for network building, such as providing funding for conference travel and encouraging membership of cross-department societies or international funding consortia. Our results suggest that such policies could be especially beneficial for scientists in the lower parts of their field-specific productivity distributions, especially where this help can be sustained over time.

Chapter 4 Estimating the Impact of a Star Scientist Move on their own Research Productivity

4.1 Introduction

In recent years, star-recruitment policies have received increasing attention as an element of the science-policy mix. Such policies have received particular attention in smaller countries and regions, where local research clusters can struggle due to a lack of scale when compared to competing jurisdictions. Two prominent examples are Denmark's Niels Bohr research professorships and Ireland's Science Foundation Ireland research professorships. Such policies aim to catalyse local research clusters by deepening local knowledge networks and also connecting to broader international networks of scientists. The specific network-related channels include enhanced local knowledge and co-author networks, improved recruitment and retention, and greater access to external knowledge and non-knowledge resources.

Star scientists are outstanding researchers with the potential to generate and transfer knowledge across regions and sectors. They are involved in various forms of scientific and industrial collaborations as well as in training and mentoring activities (Zucker and Darby, 2007). The impact of star scientists on regional development depends on the availability and quality of local partners, the type and scope of their interactions, and the characteristics of their research fields. Not all-star scientists are equally engaged in regional knowledge spillovers, and not all regions are equally attractive to star scientists (Schiller and Diez, 2008). The existing literature shows evidence of the positive impacts that star scientists have on incumbents and institutions and how they transform the research cluster (Zucker and Darby, 1996, 2001; Agrawal et al., 2014, 2017; Oettl, 2012; Trippel and Maier, 2010).

In the small-country context, recent evidence has shown evidence of positive effects of star recruitments on departmental and individual incumbent scientist productivity (McHale et al., 2022). Evidence is also emerging on the particular channels

underlying such overall effects, including the positive impact of co-authorship with stars (Yadav et al., 2023) and the positive impact of star help, as indicated by acknowledgements in publication texts (Sasidharan et al., 2024). However, there is limited evidence on the effects of mobility on a star's own productivity. In an overall assessment of star recruitments, policymakers and recruiting institutions are interested in the direct effects of the star's own productivity in addition to these broader catalytic effects, notwithstanding the possible trade-offs between these two types of outcomes given the inevitable limitations on the time of the star.

In this chapter, we estimate the effect of a star's move on their own productivity using the same data source that has been used in recent studies to estimate the indirect effects of star arrivals. We use an event-study methodology, where the event is defined as the star scientist's move to an institution in one of three small-open economies: Denmark, Ireland, and New Zealand. In our baseline specification, we identify stars that move into these three countries for the first time as our treated group and non-mover stars that are located in these countries as our control group. Furthermore, we examine the robustness of the results to alternative matching methods for identifying the control group. We also examine the robustness of our baseline two-way fixed-effects estimates to methods that are robust to temporally staggered moves by potentially heterogeneous cohorts of stars.

The results are striking and surprising given our priors. As measured by Field Normalised Total Citations (FN_{TC}), the productivity of stars falls dramatically in the year of the move and regains only a relatively small fraction of the decrease in productivity after four years. In our baseline specification, star productivity is 58.23% lower in the year of the move than the year immediately prior to the move and 42.42% lower after four years. We find that star scientists who make a move in the later stages of their careers experience a larger negative impact on their productivity compared with stars who move at an earlier stage in their careers. Additionally, we observe a similar effect on stars' productivity across most scientific fields, though stars in the social sciences appear to show no reduction in their productivity. The findings are robust across various specifications and control strategies.

To establish our priors, we first develop a simple framework that highlights three mechanisms for pre- and post-move time paths. The *matching mechanism* postulates that a move results in better matches, leading to the expectation of an increase in the star's own productivity following a move. The *disruption mechanism* postulates that a move initially disrupts the star's research activities, leading to an initial drop in productivity followed by a gradual return to levels at or above the pre-move level. Lastly, the *institution-building mechanism* postulates that a star experiences a permanent drop in own-productivity as they reallocate their time towards institution-building activities (e.g. helping younger incumbent stars or recruiting new rising stars to the institution) and away from activities focussed on their own research output.

The remainder of the chapter is structured as follows. Section 4.2 discusses the related literature and develops a simple framework that captures the three postulated mechanisms that affect the star's relative pre- and post-move productivities. In Section 4.3, we describe our data and methodology and then, report our results in Section 4.4. Section 4.5 conducts several robustness tests on our baseline results. Finally, Section 4.6 concludes with a discussion of the policy implications of our findings.

4.2 Related Literature and Framework

4.2.1 Scientific mobility and productivity

Prior literature has come to diverse findings on the effects of scientist mobility on own-productivity, suggesting that the effects are sensitive to the context in which a move occurs. A number of studies have found evidence that scientist mobility is associated with higher own-productivity. Using patent measures, Trajtenberg et al. (2006) find that inventor moves are associated with an increase in productivity. Their study examines the causal link between mobility and productivity using data on over 1.5 million inventors listed on U.S. patent documents. Their findings show that the patents of inventors who moved received more citations compared to the patents of non-movers. Similarly, Hoisl (2007) reports that inventors who move are more productive than non-movers, although she also finds that more productive inventors are less likely to move, possibly because they are more likely to have found a better match. In contrast, Azoulay et al. (2017) find that more productive scientists are more likely to move (see also Zucker et al., 2002;

Coupé et al., 2006; Lenzi, 2009; Ganguli, 2015), though their study does not examine the effects of mobility on productivity. In a study of 63,976 inter-firm moves of engineers and scientists, Chang (2023) finds evidence of positive effects on post-move performance for some movers, with the effects depending on the team design at the receiving institution. Likewise, in a study of a Spanish university, De Filippo and Casado (2009) find that researchers who had moved were more productive and received more citations than their non-mobile counterparts.

While there is evidence that scientists' moves are associated with increased productivity in some settings, a number of other studies have found that such moves have no effect and in some instances a negative effect on productivity. For example, Zubieta et al. (2015) do not find any evidence of enhanced post-move productivity in a study of 171 UK researchers. Their findings show that moving to a low-ranked university has an adverse effect on productivity in terms of both publication quantity and quality. Moreover, Allison and Long (1990) find that relocating to a less prestigious institution adversely affects the productivity of highly productive scientists. In addition, Hunter et al. (2009) show that the highly-cited UK physicists who moved to the US show no significant differences in productivity compared to domestic US physicists. Similar to this, Aksnes et al. (2013) find that mobility has only a marginal effect on the research performance of 11,000 Norwegian university researchers. Other studies point to an initial dip in performance due to transaction costs and a new working environment (Halevi et al., 2016). However, the extent of the impact is found to vary with the area of expertise, the countries they have worked, and the length and characteristics of their international experiences (Netz et al., 2020). One possibility is that following a move, a scientist takes time to adapt to and learn new tasks and administrative processes at their new institution, taking time and focus from their own research activities (Groysberg et al., 2008).

The importance of context together with human and social capital for the relationship between mobility and research productivity is directly highlighted in some existing literature. For example, research examining mobility patterns and performance among applicants to the Ramón y Cajal programme in Spain revealed important discipline-specific variations. Specifically, applicants in physics and space science, philosophy, and philology were all found to have significantly higher rates of mobility

when compared to those in molecular biology. While mobility correlated positively with participation in international research projects, the study did not find conclusive evidence supporting a direct link between mobility and publication productivity (Cañibano et al., 2008). In another study, Bäker (2015) argues that the negative effects on publication output are short-term and primarily depend on the researcher's context, such as research discipline, department size, and co-authors. The chapter also finds some support for the positive long-term effect of mobility, especially for researchers who have co-authors at the new institution.

The current literature also highlights various factors that affect the post-move productivity of academic scientists. Tartari et al. (2020) argue that a dynamic interplay of positive treatment effects, negative disruptions, and positive selection effects shapes post-mobility performance. These forces are influenced by knowledge spillovers, improved employee matching, disruption of routines, loss of firm- and relationship-specific human capital, and worker autonomy when voluntarily relocating. Furthermore, Kato and Ando (2016) observe that the mobility choices of elite scientists might be affected by the more supportive research environments in the destination country and the possibilities for more international collaboration. Additionally, Geuna (2015) highlights that international mobility can allow for better matches that result in an improved profile and productivity. There is also existing evidence that moving to a new institution can affect a scientist's productivity by disrupting their collaboration network and routines. (Feldman, 2000; Nelson and Winter, 1985). Consequently, experienced researchers may be reluctant to change their working methods to align with their new institutions in ways that impede productivity in the immediate post-move phase (Cattell and Tiner, 1949; Sorensen and Stuart, 2000; Walsh, 1988)

4.2.2 Framework for mobility effects on star scientists

The review of the existing literature suggests that a variety of mechanisms can affect how a move affects a scientist's productivity, with the balance of these mechanisms depending on the context. To motivate our empirical application, we draw on this literature to postulate three distinct mechanisms that could be at work to varying degrees in any observed instance of star mobility:

- A standard productivity-improving matching mechanism whereby, a star moves to an institution that provides a better match for their skills.
- A disruption mechanism, whereby a move initially disrupts the star scientist's activities, leading to a temporary reduction in productivity.
- An institution-building mechanism, whereby a star scientist has preferences over their own productivity and institution-building activities and faces improved opportunities for institution-building activities at their new institution.

The matching mechanism

The standard assumption in the labour economics literature is that voluntary job changes occur due to a worker finding a better match for their skills (Jovanovic, 1979; Liu, 1986; Topel and Ward, 1992; Robst, 1995). While from the worker's viewpoint, some elements of a good match could reflect preferences, part of observed mobility will reflect the worker finding a better employment match for their skills. We thus hypothesise that, all else equal, a star scientist move should be associated with an increase in observed productivity as the star moves to environments that better match their skills, at least over the medium term.

The disruption mechanism

Even where observed mobility results from better scientist-institution matches, job moves can potentially disrupt productivity in the short term. These costs arise due to factors such as professional and personal dislocation (Azoulay et al., 2017), adjustment challenges (Zubieta et al., 2016), and disruptions to established routines and social networks (Azoulay et al., 2017; Groysberg et al., 2008). Such disruption is likely to be more significant when there is a greater need for complementary assets, such as the establishment of a dedicated lab to support the productivity of the arriving star. Where disruption effects are significant, we hypothesise that, all else equal, there will be an initial fall in productivity followed by levels at or exceeding the pre-move level.

The institution-building mechanism

A potential mechanism that is consistent with a sustained negative impact on the star's own productivity is that the star optimally reallocates their time to reflect the greater opportunities for positive institutional impact – e.g., a larger number of young colleagues to be mentored or more significant opportunities to be involved in new hiring at their new institution. The idea is that the star's preferences extend beyond their own productivity to include their potential institution-building impacts and that the relative opportunities for own productivity and institution-different are different at different institutions. This allows for a richer view of matching than the simple own-productivity-based matching considered above. If the enhanced opportunities for institution-building are sufficiently valued by the star, we could observe a fall in the star's own productivity following a move. This could be accompanied by positive effects on the productivity of their new colleagues due to the institution-building activities of the star.

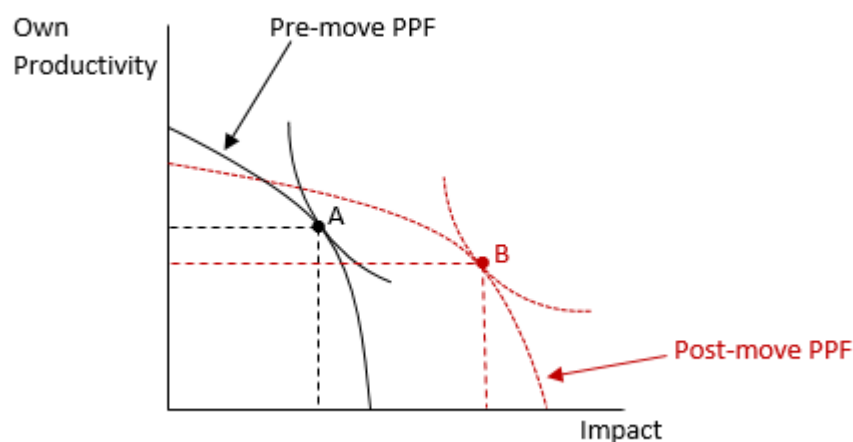


Figure 4.1 Example of productivity effects under the institution-building mechanism.

Figure 4.1 shows an example of the initial productivity effects of a star move in which these different mechanisms are potentially at work. At a given institution, we assume that the star faces a productivity-possibility-frontier (PPF) showing the possibilities for own productivity and institution productivity for a given time endowment. We assume that the PPF has a standard strictly concave shape reflecting the increasing opportunity cost in terms of own productivity as the star increasingly devotes their time to the institution-building activity. For simplicity, we assume preferences are

constant and are represented by a standard indifference map in own-productivity/institution-productivity space with indifference curves that are strictly convex. In the example depicted, the PPF at the pre-move institution is represented by a solid line, and the dashed line represents the PPF at the post-move institution. The relative positions are drawn in the example on the assumption that the move harms the possibilities for own productivity, but the opportunities for institution building are enhanced. One way this can be seen is to consider the productivity possibilities under extreme time allocation cases. If all time is devoted to own-productivity activities, we can see that own-productivity would fall due to a move. On the other hand, if all time is devoted to institution-building activities, the productivity of the star's institution is higher post-move – presumably because they have more opportunities to engage in constructive institution-building roles after the move.

In comparing the optimal pre- and post-move allocation of time. The optimal allocation is given by the point of tangency between the relevant PPF and the highest feasible indifference curve. As shown, utility is higher after the move, leading to a prediction that the move will be made. However, what is most interesting is that when comparing the optimal pre-move allocation (A) with the optimal post-move allocation (B), own-productivity is lower and institution productivity is higher after the move. Moreover, to the extent that the new PPF remains as shown over time, the fall in own productivity will be an enduring result of the move.

This static picture leaves out important dynamic effects that are likely to play out in the time periods following a star's move. In particular, the position of the initial PPF will be affected by any disruption effects that should lessen over time as the star adjusts to their new institution. In addition, it is likely to take time for the star to embed into local networks, which will affect both their own and the new institution's productivity potential. To empirically trace out the dynamic effects, we therefore adopt an event-study design where the focal event is the arrival of the star at their new institution. The event-study approach allows us to examine both the dynamics of the star's pre- and post-move productivity relative to suitably chosen controls. We thus observe the shifting balance of the effects of the different hypothesised mechanisms.

4.3 Data and Methodology

Our analysis utilises publication level data from Scopus collected from 1990 to 2019 for three countries: Ireland, Denmark and New Zealand. We collect all publication data affiliated with these three countries. The data collected includes publication year, authors, affiliations, subject field, and citation count up to 2019. Overall, it comprises approximately 1.43 million publications featuring 219,582 unique authors.

We identify all the author details related to each publication. We then collect each author's back catalogue of citation data since their first available publication in Scopus to include publications affiliated outside of the three focal countries. For this study, we conduct our analysis at the author level rather than the publication level. Finally, we use the affiliation tagged to each author to identify their university location in the three countries. We use the resulting dataset to identify the international stars that move into any of the three countries.

4.3.1 Identification of international stars and their moves

Since our main focus is on the effect that the movement of a star to a new affiliation has on their own productivity in terms of citation-weighted output, we obtain the forward citations to each publication from 1990 for each subject field in any given year up to 2019. In each year, we calculate the cumulative citations to their publications and identify the field-specific distribution of the cumulative citations for publications up to the given year. We then define a *star* as one who is at or above the 95th percentile of scientists in this cumulative distribution. Using this procedure, we identify a total of 981 star scientists in our data.

Next, to identify the star moves from our data, we use the affiliations of each star to determine their institutional affiliation in Ireland, Denmark, or New Zealand. We identify international star moves as instances where star scientists, who previously did not publish with an affiliation to an institution in one of the three focal countries are now publishing with such an affiliation. Moreover, we restrict the identification of star moves to those who continue to publish with an affiliation to the department for at least the next four years. For those stars that moved in 2015-2016, we use the restriction that

Table 4.1 Mean value comparison of treated, control group, and broad subject field specifications

	Variables		
	Total Publications (log)	Total Citations (log)	FNTC (log)
Star-Movers(Treat, n=161)	4.131	1.441	0.915
Non-Movers(Control, n=75)	4.245	1.452	0.870
Diff-in-mean values	0.114	0.011	-0.045
P-Values	0.157	0.741	0.172
Physical Sciences	4.196	1.527	0.886
Health Sciences	4.526	1.601	0.895
Social Sciences	3.334	1.107	0.858
Life Sciences	4.334	1.396	0.929

Note: The table reports the mean average values calculated for the treated and control group for the star authors.

they are also observed to publish at the new affiliation at the end of our observation period 2017 to be included as a star move. As a further restriction, we do not include stars who moved in the first year of our observation period, 1996 (always treated). Through this procedure, we identify 161 unique international star moves. These 161 stars form our treatment group for the analysis. Furthermore, as part of robustness analysis, we separately classify the star moves that represent the first stars received by a department. We identify 83 such stars that move to a department, while the remaining 78 stars are defined as subsequent stars move to a department.

Our biggest challenge in this study is the identification of a control group of non-movers who are star scientists. For the control group, we first identify star scientists in Ireland, Denmark and New Zealand, who have had the same affiliation in these three countries since their first publication in Scopus; and second, we confirm that they continue to publish at the same affiliation. We use the author's affiliation details each year to see whether the star moves between affiliations during their academic careers. It is important to note that we do not consider a star as a non-mover if they have any change in affiliation during their observed career. Overall, we identify 75 stars that meet this definition of a non-mover star, which forms the control group in our baseline analysis. Table 4.1 reports descriptive statistics for both the treated and control groups.

4.3.2 Output Measure

We use the Field Normalised Total Citations (FNTC) as our scientific output measure. The FNTC is the sum of an individual's publication citations divided by the

average citations to a publication for that subject field in that year for all countries combined;

We calculate the dependent variable as follows:

Field Normalized Total Citations:
$$Y_{it}^{FNTC} = \sum_{p_{i,t}=1}^{P_{i,t}} \frac{c_{p_{i,t}}}{\bar{c}_{s,t}};$$

where $P_{i,t}$ is the total number of publications by individual i published in year t , $c_{p_{i,t}}$ are the subsequent total citations (or "forward" citations recorded in 2019) to a publication $p_{i,t}$ that occur for individual i in year t , $\bar{c}_{s,t}$ is the average citations to a publication in the relevant subject field, s , for publications that occur in year t .

4.3.3 Econometric Methodology

The empirical goal of our econometric analysis is to estimate the effects of a star's move on their own productivity. To measure these dynamic effects, we utilise an event-study specification where the event is a star scientist move to any of the three countries. Our event-study design assumes staggered adoption, where units (stars) are treated at different times, and some stars have never been treated. Using methods similar to those used in previous studies (McHale et al.,2022; Yadav et al.,2023), we estimate the dynamic treatment effects on the star's quality-adjusted output four years before and four years after the star move event in the following event study specification:

$$\begin{aligned} \ln Y_{it} = & \alpha + \beta_{\leq -5} \text{starmove}_{i,-5} + \sum_{j=-4}^{-2} \beta_j \text{starmove}_{ij} + \sum_{j=0}^4 \beta_j \text{starmove}_{ij} \\ & + \beta_{\geq 5} \text{starmove}_{i,5} + \delta_t + \mu_i + \epsilon_{it}, \end{aligned} \quad (4.1)$$

where the dependent variable $\ln Y_{it}$ is a measure of the citation-weighted output of star i at year t , starmove_{ij} is a binary variable equal to 1 if a star makes an international move to the SOE as of year t , j years ago, δ_t is a year fixed effect, μ_i is a star fixed effect, and ϵ_{it} is a zero mean error term. The coefficients of interest, β_j , show the proportionate effect of the star's own productivity from four years before the move to four years afterward. We normalise the effect on the star's productivity to zero for the year before the star move, and we assume that the cumulative effect is constant at $\beta_{\leq -5}$ and $\beta_{\geq 5}$ by binning at

Table 4.2 Dynamic star move effects on star's own productivity for FNTC (overall, early and late career moves)

	Overall Star Moves (1)	Early Career Moves (2)	Late Career Moves (3)
<i>Starmove_{i,t-4}</i>	-0.129 (0.103)	-0.178 (0.177)	-0.126 (0.123)
<i>Starmove_{i,t-3}</i>	0.0220 (0.0918)	-0.0252 (0.173)	0.0295 (0.0969)
<i>Starmove_{i,t-2}</i>	-0.0113 (0.0940)	-0.0242 (0.170)	-0.0120 (0.105)
<i>Starmove_{i,t}</i>	-0.873*** (0.113)	-0.602*** (0.205)	-1.063*** (0.122)
<i>Starmove_{i,t+1}</i>	-0.587*** (0.105)	-0.500*** (0.169)	-0.626*** (0.133)
<i>Starmove_{i,t+2}</i>	-0.472*** (0.105)	-0.464*** (0.163)	-0.455*** (0.136)
<i>Starmove_{i,t+3}</i>	-0.667*** (0.105)	-0.695*** (0.165)	-0.597*** (0.136)
<i>Starmove_{i,t+4}</i>	-0.552*** (0.102)	-0.562*** (0.168)	-0.497*** (0.123)
Constant	1.345*** (0.0854)	1.289*** (0.105)	1.302*** (0.0876)
R-squared	0.070	0.055	0.083
Observations	4,635	2,639	3,390
Number of authors	236	143	168
Author FE	YES	YES	YES
Year FE	YES	YES	YES

Event Window: 1997-2017. Estimation Window: 1996-2017.

The dependent variable is Field Normalised Total Citations.

Note: The table reports the estimates based on the model specification in Equation 4.1. Column 1 reports the dynamic effects on the star's own productivity for all 161 international star movers in the three countries. Column 2 reports the effects on the star's own productivity for Early Career Movers, while Column 3 reports on the Late Career Movers. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively.

five leads and five lags. The binning variables may not be comparable to the leads and lags of the binary variables (detecting star moves) in estimating the dynamic effects since they could be correlated with other excluded level variables; however, they act as essential controls in our specification (Kurt Schmidheiny et al., 2019). Furthermore, standard errors are clustered at the author level and are robust to arbitrary forms of serial correlation and heteroscedasticity.

4.4 Results

In this section, we present the results from our main event study specifications estimating the effect of a star move on their own productivity. Additionally, we consider the empirical evidence for each of the three competing mechanisms outlined in Section 4.2.2 above. Using the FNTC²⁸ measure as the dependent variable, we report the estimated homogenous treatment effects of a star move in Column 1 of Table 4.2 and based on the specification in Equation 4.1. We find evidence that a star move is associated with a 58.23% decrease in their citation-weighted output in the year that they move to a new affiliation in one of the three SOEs. We find that the estimated coefficients remain negative in the subsequent years and are statistically significant at the 1% level. When compared to the initial drop in the first year, the decrease in quality-adjusted output diminishes, but remains substantial, with a 43.42% fall on average compared to the pre-move level across the next four years.




Figure 4.2 Event study plot with the dynamic effects on star movers' own productivity

Note: The figure plots the dynamic effect of star's own productivity at 95% confidence intervals. The event window is from 1997-2017. The dependent variable is field normalised total citations.

Figure 4.2 presents the event study plot for these estimates. From this plot, we observe that the treated and control groups exhibit similar output levels in the pre-treatment period, as the estimated lead coefficients are not statistically different to zero. The absence of any pre-trend in the pre-treatment period helps support the assumption of a parallel-trend between treated and control groups. In general, these results provide

²⁸ We also use an alternate dependent variable(normalised raw publications) to test the robustness of the results-see Appendix C.3

support for the institution-building mechanism, with some evidence for an initial disruption effect. The most striking finding is that there is an enduring drop in productivity after a move, which is consistent with the star's involvement in activities that affect the institution's overall performance rather than solely contributing to their own scientific output. This interpretation is also consistent with evidence that the arrival of these stars is associated with increased departmental and individual incumbent productivity at the star-receiving institutions (McHale et al., 2022).




Figure 4.3 Event study plot with the dynamic effects on star movers' own productivity for early career star movers (left) vs late career star movers(right).

Note: The figure plots the dynamic effect of star's own productivity at 95% confidence intervals. The event window is from 1997-2017. The dependent variable is field normalised total citations.

We also investigate differences in the effect of a move based on the career age of the stars at the time of their move. We determine the career age of a star by identifying from their back catalogue the year in which their first publication occurred. Using this career age data, we classify the 161-star moves into two groups - *early-career* and *late-career movers*. *Early-career movers* are defined as stars with a career age of 15 years or less at the time of their move (since the year of their first publication is available in the data), while *late-career movers* are specified as stars with a career age of more than 15 years. We identify 68 moves by early-career stars and 93 moves by late-career stars. For both groups, we estimate the impact of a move using our baseline specification together with our control sample of 75 stars. The results for early and late-career moves are reported in Columns 2 and 3 of Table 4.2 respectively.

Table 4.3 Dynamic star move effects on star's own productivity for FNTC (broad subject fields)

	Physical Sciences (1)	Health Sciences (2)	Social Sciences (3)	Life Sciences (4)
<i>Starmove_{i,t-4}</i>	-0.161 (0.217)	-0.0302 (0.203)	0.159 (0.222)	-0.404** (0.178)
<i>Starmove_{i,t-3}</i>	0.188 (0.159)	0.119 (0.147)	0.00109 (0.236)	-0.250 (0.209)
<i>Starmove_{i,t-2}</i>	-0.0226 (0.192)	0.133 (0.134)	0.226 (0.252)	-0.202 (0.169)
<i>Starmove_{i,t}</i>	-1.108*** (0.183)	-1.110*** (0.161)	-0.221 (0.276)	-0.832*** (0.269)
<i>Starmove_{i,t+1}</i>	-0.780*** (0.213)	-0.631*** (0.171)	-0.174 (0.201)	-0.607*** (0.217)
<i>Starmove_{i,t+2}</i>	-0.635*** (0.184)	-0.545*** (0.188)	-0.0276 (0.266)	-0.524*** (0.231)
<i>Starmove_{i,t+3}</i>	-0.791*** (0.216)	-0.603*** (0.172)	-0.595*** (0.209)	-0.620*** (0.237)
<i>Starmove_{i,t+4}</i>	-0.637*** (0.181)	-0.695*** (0.195)	-0.293 (0.235)	-0.535*** (0.213)
Constant	1.467*** (0.138)	1.236*** (0.160)	0.726*** (0.238)	1.609*** (0.172)
R-squared	0.117	0.130	0.061	0.080
Observations	1,585	1,103	802	1,145
Number of authors	83	55	41	57
Author FE	YES	YES	YES	YES
Year FE	YES	YES	YES	YES

Event Window: 1997-2017. Estimation Window: 1996-2017.

The dependent variable is Field Normalised Total Citations.

Note: The table reports the estimates based on the model specification in Equation 4.1 for broader subject areas associated with the stars. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively.

Interestingly, we find that stars who move at an early-career age experience an output loss of 45.23% at the time of their move, while stars who move at a late-career age are found to have a larger output loss at 65.46% at the year of their move. Similar to the main results, we also observe a significant decrease in the citation-weighted output across the subsequent four years for both groups. Figure 4.3 shows the related event-study plot. The figure highlights the more pronounced contemporaneous decrease in output among late-career movers. Overall, this provides further suggestive evidence to support institution-building mechanism, as it is plausible that stars later in their careers will have a stronger preference for institution-building activities relative to activities that further enhance their own publication records.

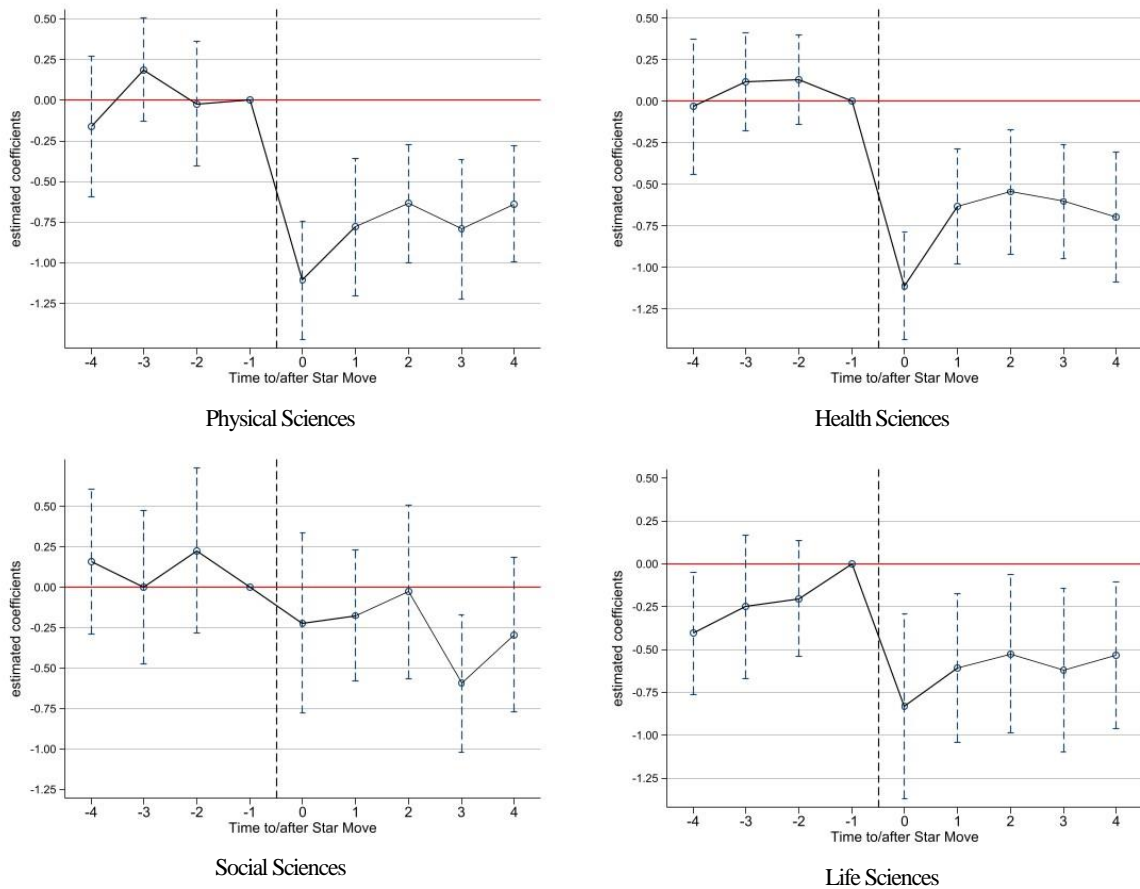


Figure 4.4 Event study plot with the dynamic effects on star movers' own productivity for broad field of research

Note: The figure plots the dynamic effect of star's own productivity at 95% confidence intervals. The event window is from 1997-2017. The dependent variable is field normalised total citations.

Next, we examine the consistency of these productivity findings across scientific fields. Using the star's broad field of research as defined by Scopus²⁹, we examine for differences in the effect of a move for stars located in different fields. We consider four broad field specifications – Physical, Health, Social, and Life Sciences³⁰. Scientific areas

²⁹ We use Scopus-identified subject fields to identify the broad field of star scientists. The 27 subject fields defined by Scopus are Agricultural and Biological Sciences, Arts and Humanities, Biochemistry, Business, Chemical Engineering, Chemistry, Computer Science, Decision Science, Dentistry, Earth and Planetary Sciences, Economics, Energy, Engineering, Environmental Science, Health Professions, Immunology and Microbiology, Multidisciplinary, Materials Science, Mathematics, Medicine, Neuroscience, Nursing, Pharmacology, Physics and Astronomy, Psychology, Social Sciences, and Veterinary.

³⁰ Physical sciences broadly deal with the systematic study of the inorganic world. Publications related to Astronomy, Physics, Chemistry, and Earth Sciences are primarily under Physical sciences. Second, health sciences are related to various disciplines dealing with the study and policy implications of healthcare and medical topics. Often, they conduct original research and help to provide insights into ongoing public health care problems, such as inventions of new medicines and their applicability, and the effects of public health policy. Third, social sciences deal with the study of human society, culture, and the behaviour of the public in relation to various policies. This broad field includes subjects such as sociology, anthropology, economics, and political science. Finally, Life sciences include the study of living organisms. Subject areas related to biochemistry, genetics, molecular biology, and other related fields come under the broad specification of life sciences.

like healthcare, medicine, biochemistry, and neuroscience require the availability of complex lab equipment and facilities. Thus, for the relocation of a star, the new location must have relevant facilities, or it could impact the star's own research productivity since they may need to trade-off between progressing his/her scientific advancement and developing the research infrastructure in the arriving institution. Table 4.3 presents the estimated coefficients for the effects of a star move on its own productivity across the four broad subject areas. For Physical, Health and Life sciences,

We find a significant decrease in the star's citation-weighted output in the year of the move. A star move is associated with a 67% decline in output for both Physical and Health sciences, whereas it is linked to a 56.48% decline in output for Life sciences. The related event study plots in Figure 4.4 further illustrate this effect, showing its sustained nature in subsequent years, albeit with a somewhat reduced magnitude. In contrast, for stars in the Social Sciences, we observe that a move has no effect on their citation-weighted output. While the results for Physical, Health and Life sciences show support for institution-building mechanism, the findings in Social Sciences may align better with the matching mechanism. The results for the Social Sciences suggest that stars in this field do not exhibit the type of disruption costs and enduring adverse own-productivity effects that are observed in the more lab-based sciences, suggesting that both the extent of the disruptions and opportunities for institution-building may be lower in the less equipment- and team-dependent social sciences.

4.5 Robustness

4.5.1 Robustness to heterogeneous treatment effects

The timings of star moves in our data are staggered, and heterogeneity in the effects could occur if different cohorts experience different treatment paths. Furthermore, when the timing of the treatment varies across units (stars), the effect for a specific period can be influenced by the effects from other periods. To reduce this potential 'contamination' from other periods, we use the methodology of Sun and Abraham (2021) to estimate the heterogeneous treatment effects of a star move. This method calculates the dynamic effects of a star move in a three-step estimation that is

robust to treatment effect heterogeneity and computes a weighted average of ‘Cohort Average Treatment effects on the Treated’ (CATT).³¹

Table 4.4 Robustness test of dynamic star move effects on star’s own productivity for FNTC (heterogeneous effects, first and subsequent star moves)

	Heterogeneous Star Moves (1)	First Star Move (2)	Subsequent Star Moves (3)
<i>Starmove_{i,t-4}</i>	-0.085 (0.108)	-0.0313 (0.153)	-0.205 (0.135)
<i>Starmove_{i,t-3}</i>	0.040 (0.095)	0.000570 (0.126)	0.0527 (0.133)
<i>Starmove_{i,t-2}</i>	-0.0055 (0.0963)	-0.0497 (0.140)	0.0293 (0.126)
<i>Starmove_{i,t}</i>	-0.872*** (0.114)	-0.862*** (0.154)	-0.859*** (0.165)
<i>Starmove_{i,t+1}</i>	-0.579*** (0.106)	-0.593*** (0.152)	-0.552*** (0.142)
<i>Starmove_{i,t+2}</i>	-0.461*** (0.108)	-0.492*** (0.154)	-0.394*** (0.137)
<i>Starmove_{i,t+3}</i>	-0.658*** (0.108)	-0.659*** (0.151)	-0.602*** (0.146)
<i>Starmove_{i,t+4}</i>	-0.551*** (0.105)	-0.470*** (0.137)	-0.578*** (0.149)
Constant	1.317*** (0.0508)	1.348*** (0.0934)	1.217*** (0.0982)
R-squared	0.447	0.069	0.064
Observations	4,635	3,089	2,940
Number of authors	236	158	153
Author FE	YES	YES	YES
Year FE	YES	YES	YES

Event Window: 1997-2017. Estimation Window: 1996-2017.

The dependent variable is Field Normalised Total Citations.

Note: Column 1 reports the Sun and Abraham (2021) heterogeneous star-move effects on the star’s own citation-weighted productivity.

Column 2 reports the effects on the star’s own productivity for first star arrivals, while Column 3 reports on the Subsequent stars who arrive at the university. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively.

³¹ We use a similar approach to design the event–study model to calculate the CATT as our previous studies (McHale et al., 2022; Sasidharan et al., 2024). See Appendix C.1.




Figure 4.5 Event study plot with the dynamic effects on star movers' own productivity for heterogeneous star moves




Figure 4.6 Event study plot with the dynamic effects on star movers' own productivity for first star (left) vs subsequent star movers(right)

Note: The figure plots the dynamic effect of star's own productivity at 95% confidence intervals. The event window is from 1997-2017. The dependent variable is field normalised total citations.

In our baseline estimation, we calculate the homogenous treatment effects of a star move on the star's own quality-adjusted output to find that a star's output falls by 58.23% in the year of the move, with the results showing persistent negative effects in subsequent years (see section 4.4). For comparison, we present the heterogeneous treatment effects of a star move using the Sun and Abraham approach in Column 1 of Table 4.4 We find an almost identical 58.19% decrease in the star's output in the event year, with this negative effect again largely persisting in subsequent years. These results are also illustrated in the event study plot in Figure 4.5, which shows a similar pattern to the homogenous treatment effects shown in Figure 4.2.

4.5.2 Robustness to first and subsequent star moves to an affiliation

As the institution or department receives a star scientist for the first time, we expect the productivity to drop due to initial adjustment costs from the star's perspective and the institution's available infrastructure to complement the star. However, we assume that for the subsequent stars arriving at these affiliations, learnings from previous star arrivals could help the star scientist reduce the adjustment time at the new institution. We categorise star moves into two groups: those representing the first stars received by a department and those constituting subsequent star moves to a department. We identify 83 instances of star moves to a department as the first stars, while the remaining 78 stars are defined as subsequent star moves to a department. We estimate the baseline model specification separately for both groups, and the results are presented in Columns 2 and 3 of Table 4.4. Compared to the control sample of non-moving stars, both groups of star moves exhibit a 58% decrease in output in the event year. The event-study plot (Figure 4.6) indicates that both groups have a sustained drop in their research productivity, with this decline again showing only limited moderation over time. On average, both groups also show a decline of between 41% and 42% in their research outputs in the subsequent four years post-move. Overall, the time pattern of productivity changes is almost identical across first- and subsequent-star arrivals. It suggests that the arrival of star scientist is expected to have a drop in their research productivity irrespective of their timing of arrival.

4.5.3 Robustness to alternative control samples

Our primary analysis uses non-mover stars who stay at the same affiliation for the entire period in Ireland, Denmark and New Zealand as our control group. As a further robustness test, we instead define a non-mover star as one who has published at the same affiliation for at least four years after being first identified as a star, starting from the year they reached the threshold for star identification (as explained in section 4.3.1). Using this alternative definition, we identify a total of 144 non-mover stars in our control sample. Column 1 in Table 4.5 reports the estimated effects of a star move using the new control

group. We continue to observe a 58.04% decrease in output for stars after a move, and this reduction in productivity is once again sustained in later years. Therefore, the main results remain robust to this change in the identification of the control group.

Table 4.5 Robustness test of dynamic star move effects on star’s own productivity for FNTC (alternative control group)

	Alternative Control Group (1)	Matched treated and Control Units (2)
<i>Starmove_{i,t-4}</i>	-0.139 (0.103)	-0.217 (0.194)
<i>Starmove_{i,t-3}</i>	0.0132 (0.0910)	0.0325 (0.153)
<i>Starmove_{i,t-2}</i>	-0.0194 (0.0936)	-0.0942 (0.151)
<i>Starmove_{i,t}</i>	-0.868*** (0.112)	-0.888*** (0.213)
<i>Starmove_{i,t+1}</i>	-0.579*** (0.104)	-0.665*** (0.193)
<i>Starmove_{i,t+2}</i>	-0.462*** (0.103)	-0.432** (0.181)
<i>Starmove_{i,t+3}</i>	-0.651*** (0.104)	-0.696*** (0.197)
<i>Starmove_{i,t+4}</i>	-0.532*** (0.1000)	-0.462*** (0.140)
Constant	1.291*** (0.0665)	1.362*** (0.114)
R-squared	0.063	0.074
Observations	6,074	1,971
Number of authors	305	94
Author FE	YES	YES
Year FE	YES	YES

Event Window: 1997-2017. Estimation Window: 1996-2017.

The dependent variable is Field Normalised Total Citations.

The table reports the estimates based on the model specification in Equation 4.1. Column 1 reports the effects on the star's own productivity compared with an alternate control group. Column 2 reports the analysis based on matched treated and control units from the new set of control samples. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively

As an additional robustness measure, we utilise Coarsened Exact Matching (CEM) to enhance the balance of covariates between treated and control groups and bring each covariate into balance across both groups to help minimise any confounding effect in our observational causal inference. CEM is a method that reduces imbalance in a monotonic manner (Blackwell et al., 2009; M. Iacus et al., 2011; M. Iacus et al., 2012). One significant advantage of this procedure is that we can define covariates to match a categorical variable rather than a continuous one. Based on four covariates: the year, the arrival

institution (universities), career age at which the star is making a move, and the cumulative average citations per publication, all measured at $t - 1$, we identify 47 pairs with exact matches for the treatment and control groups one year before the treatment occurs³². We then employ the same event study specification, utilising the newly matched sample to estimate the effect of a move on the star’s output. Column 2 in Table 4.5 reports the estimates obtained from the event study analysis, and Figure 4.7 provides a visual representation of the event study. Once again, the estimated results align with our main findings, increasing our confidence in the robustness of the results to the precise procedure used to identify the control stars.




Figure 4.7 Event study plot with the dynamic effects on star movers' own productivity for alternate control group (left) and matched sample of treated and control units(right).

Note: The figure plots the dynamic effect of star’s own productivity at 95% confidence intervals. The event window is from 1997-2017. The dependent variable is field normalised total citations.

4.5.4 Robustness to the country of star arrival

To assess whether the impacts of a star move vary across the three SOEs, we also estimate our baseline specification independently for each of the three countries: Ireland, Denmark and New Zealand. The results for each country are presented in Table 4.6. The results show that there are some differences observed across countries; for instance, a star move to Ireland is associated with a 67.04% decline in their citation-weighted output in the year of the move, while a star move to Denmark is associated with a smaller decline in output (52.62%). Furthermore, when compared to the stars who move to Ireland (58.37%), the average decrease in the citation-weighted output over the subsequent four

³² See Appendix C.2 for summary statistics for matched treated and control group.

years is also less in Denmark (34.79%). However, in line with our main findings, we find a substantial decrease in a star's output in the year of the move, accompanied by a persistent decline in productivity in subsequent years following the move across all three countries. Figure 4.8 shows these results with separate event-study plots for each country.

Table 4.6 Robustness test of dynamic star move effects on star's own productivity for FNTC (country level)

	Ireland (1)	Denmark (2)	New Zealand (2)
<i>Starmove_{i,t-4}</i>	-0.0836 (0.205)	0.0362 (0.125)	-0.392 (0.263)
<i>Starmove_{i,t-3}</i>	-0.0188 (0.183)	0.184 (0.139)	-0.114 (0.174)
<i>Starmove_{i,t-2}</i>	0.00224 (0.184)	0.185 (0.132)	-0.402** (0.195)
<i>Starmove_{i,t}</i>	-1.110*** (0.255)	-0.747*** (0.144)	-0.887*** (0.241)
<i>Starmove_{i,t+1}</i>	-0.855*** (0.210)	-0.395*** (0.143)	-0.700*** (0.238)
<i>Starmove_{i,t+2}</i>	-0.745*** (0.225)	-0.351** (0.143)	-0.569** (0.236)
<i>Starmove_{i,t+3}</i>	-0.896*** (0.247)	-0.619*** (0.156)	-0.659*** (0.189)
<i>Starmove_{i,t+4}</i>	-1.030*** (0.243)	-0.345** (0.145)	-0.490** (0.190)
Constant	1.305*** (0.245)	1.297*** (0.114)	1.230*** (0.154)
R-squared	0.135	0.075	0.067
Observations	885	2,361	1,389
Number of authors	45	118	73
Author FE	YES	YES	YES
Year FE	YES	YES	YES

Event Window: 1997-2017. Estimation Window: 1996-2017.

The dependent variable is Field Normalised Total Citations.

Note: Table reports the estimated coefficients based on the model specification in Equation 4.1. We report the disintegrated analysis of stars who arrive at Ireland, Denmark and New Zealand in Columns 1, 2 and 3 respectively. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively

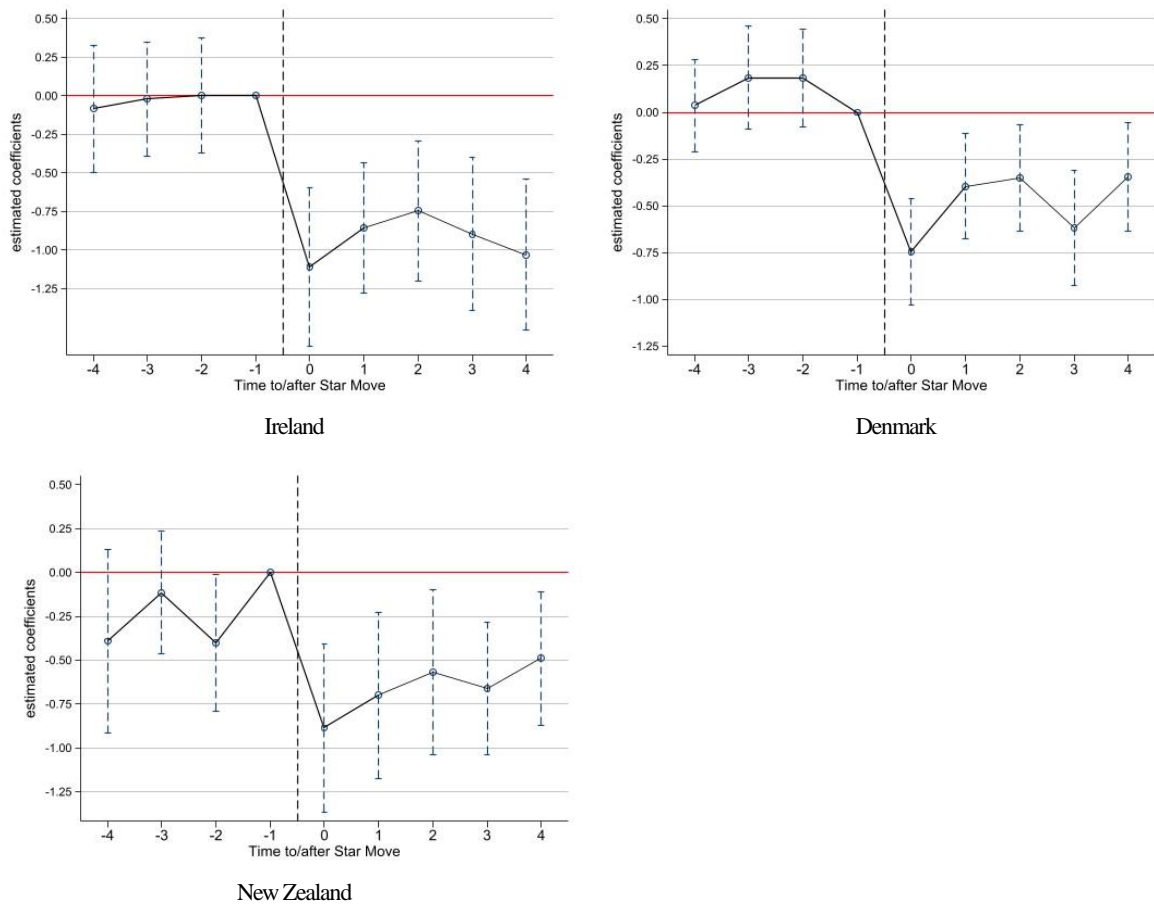


Figure 4.8 Event study plot with the dynamic effects on star movers' own productivity for Ireland Denmark and New Zealand

Note: The figure plots the dynamic effect of star's own productivity at 95% confidence intervals. The event window is from 1997-2017. The dependent variable is field normalised total citations.

4.6 Summary and Policy Implications

This study explored the effects of mobility on star scientists' productivity in three small open economies (SOEs): Ireland, Denmark, and New Zealand. Using data from Scopus, we identified 981 stars and 161 star-movers who join an SOE for the first time and stay for at least four years, as well as 75 star non-movers who remain in the same SOE. We applied an event study approach to compare the field normalised total citations (FNTC) of the two groups before and after moving to an SOE. In general, we find that the movers' FNTC decreased by 58% in the year of the move (event year), with the negative effect broadly enduring in subsequent years. We find that these results are robust to alternative estimation methods, control strategies and productivity measures. We also find that the own-productivity declines are most pronounced for late-career star movers

and are broadly similar across scientific fields with the exception of the social sciences, which is found to exhibit lower declines in star-mover productivity in our analysis.

We also developed a simple framework that identified three mechanisms that could impact the own-productivity of a star who moves to a new institution: matching, disruption, and institution-building. Our results are most consistent with the institution-building mechanism, which postulates that a star's motivation for a move is to avail of greater opportunities for institution-building activities, even where these activities come at a cost in terms of their own productivity. This interpretation is consistent with previous findings that show star arrivals are associated with improvements in departmental and incumbent productivity at the receiving department (McHale et al., 2022). However, we find discipline-specific variations in the post-mobility productivity mechanism. For the broad subject fields, including physical, health and social sciences, the citation-weighted output is reduced significantly and remained negative in the subsequent years, which lends support to the institution-building mechanism. Given the nature of these fields, the mobile stars require time to implement and install their labs and equipment at the new affiliation. More specifically for social sciences, we find no evidence of any adverse productivity effects due to mobility, instead showing evidence of a matching mechanism. These findings suggest mobility and productivity are associated with the subject fields in which the stars work.

Overall, these findings contribute to the literature by estimating the effects of star moves on the star's own productivity in the context of SOEs. While policymakers in these countries are interested in the potential benefits these stars bring to the institution and countries, our results suggest that they should also consider the potential falls in the star's own productivity when they join their new institution. Whether the goal is to support that star's institution-building role or the star's own productivity, the results also highlight the importance of the successful integration of the arriving stars into their new institutions. In addition, the differing results for early- and late-career arriving stars suggests that support policies should be tailored based on the objectives of the recruitment, ensuring that the stars can maximise their contributions, whether that contribution is coming mainly through their own productivity or through their institution-building roles.

Chapter 5 Star Arrival Effects on Scientists' Productivity: Does the Nature and Extent of Relatedness to the Star Matter?

5.1 Introduction

There is increasing interest in star recruitment policies as a means to catalyse the development of clusters in targeted research areas, industries and regions. Lacking broader economies of scale, such targeted policies may be of particular interest in smaller economies and economic regions. In motivating these policies, a core underlying assumption is the importance of local knowledge networks. The presence of stars is expected to strengthen these networks, with possible reinforcing dynamic effects through subsequent improved attraction and retention of talent.

Scholars find that knowledge spillovers tend to occur primarily between activities that draw on "related" knowledge (Autant-Bernard and LeSage, 2011; Frenken et al., 2007). The concept of relatedness is based on a measure of cognitive proximity, in which individuals sharing similar sets of knowledge and expertise increase their ability to interact and learn from each other (Boschma, 2005; Torre and Rallet, 2005). In the present chapter, we extend the focus beyond local economic activities and look at the effects of the scientists' relatedness on scientific productivity within academic departments.

Although spatial proximity could facilitate access to relevant knowledge, co-location is neither necessary nor sufficient for knowledge spillovers (Boschma, 2005; Torre and Rallet, 2005). Substantial spillovers could occur while being geographically distant but related through networks in other ways. On the other hand, scientists could be co-located with a star but receive minimal knowledge spillovers due to their distance from the star in knowledge space³³. It follows that the arrival of a star might have

³³ A knowledge space displays the position of scientists within a network based on the similarity of scientists' knowledge sets. These sets are derived from the co-occurrences of publications in the same dimension of publication records (journals, keywords, subject categories). Scientific relatedness denotes the distance between pairs of scientists within the knowledge space.

substantially different effects on incumbents in the receiving location depending on their degree of scientific relatedness to the star. To motivate our empirical application, we first develop a simple model in which star arrivals affect incumbent scientist productivities through knowledge spillover effects. Such spillovers from the star scientist take place through interactive learning that allows new rounds of ideas combination and knowledge creation (Nelson and Winter, 1982; Weitzman, 1998). The efficiency of the communication and absorption of ideas increases if researchers share references in the knowledge space that are sufficiently 'related' to bridge gaps between them. On the other hand, creating novel knowledge depends on combining ideas that are related but dissimilar to some degree (Cohen and Levinthal, 1990; Nooteboom, 2000). A greater variety of (slightly different) ideas increases the number of ways they might be combined, affecting the amount and quality of knowledge output. We show the possibility of an inverted U-shaped relationship between the degree of relatedness to the star and the size of the star-arrival productivity effect.

Our empirical application builds on previous literature that attempts to measure the average treatment effect of star arrivals to academic departments in the context of a small open economy (McHale et al., 2022). While significant treatment effects have been observed, the literature has not accounted for varying "intensities of treatment" (or "exposures") based on the relatedness of incumbent scientists to the incoming star. Here, modelling the production of scientific output as a process of combining ideas, the arrival of a star can increase the potential access to new ideas, but the ability to actually absorb that knowledge depends on the scientist's degree of relatedness to the star. It follows that incumbent scientists are subject to treatments of varying intensity when a star arrives.

The chapter draws upon and contributes to a number of literatures spanning different disciplines. First, the chapter most directly contributes to the empirical literature on peer effects in science. An influential strand of this literature uses the plausibly exogenous arrivals or departures of accomplished peers to estimate the treatment effect of proximity to a star. Azoulay, Graff Zivin, and Wang (2010) use unexpected deaths of star scientists, and Waldinger (2012) uses the dismissal of scientists in Nazi Germany to estimate the effects on those left behind. The literature has also used a combination of event-study methodologies – where the event is the entry or

exit of a star – and matching strategies to estimate such peer effects (Agrawal et al., 2017; Azoulay et al., 2019).

Second, the chapter relates to the organizational science literature on network structure and performance. A major theme in this literature is the relative values of network "closure" versus "brokerage" in supporting performance (Ahuja, 2000). On the one hand, embeddedness in dense, closed networks can support deeper bonds of trust and may be especially important in transferring tacit knowledge (Kogut and Zander, 1992). Dougherty (1992) points to how divergent "interpretative schemes" can act as a barrier to collaboration. On the other hand, individuals ("brokers" or "boundary spanners") who are able to close structural holes within networks, or span boundaries across networks, may play an important role in bridging gaps in networks and thus in supporting knowledge transfer and collaboration (Burt, 1992, 2004; Obstfeld, 2005; Tushman, 1977).

Recent articles have also stressed how different network structures might differentially support the star's own-productivity compared to their role in catalysing the productivity of others (Sosa, 2011; Tortoriello et al., 2015). Emphasizing this catalytic role of stars, Liu, Mihm, and Sosa (2018) find that the network's social cohesion and expertise similarity are especially important in the transfer of the star's "creative synthesis skills," suggesting a possible magnified importance of relatedness when stars are playing such roles.

The organizational science literature thus points to both the potential costs as well as the potential benefits of greater relatedness of an incoming star to other scientists at the receiving destination (Ahuja, 2000; Uzzi, 1996). The higher the scientific relatedness in the knowledge space, the higher the researchers' ability in communicating with and absorbing ideas from the star, but at the same time, too much proximity in cognitive terms may imply low levels of knowledge novelty, reducing opportunities for ideas combination and leading to a knowledge lock-in (Nooteboom, 2000). It follows that knowledge spillover and creation depend on a balance between an optimal scientific relatedness distance (variety) and proximity (similarity) between researchers that is adjusted over time through continuous interactions. In the context of a simple combinatory model of science production, we attempt to capture this tension by allowing for greater

relatedness to increase an incumbent scientist's ability to absorb knowledge from the star (absorptive capacity) but also to raise the risk of knowledge redundancy (Cohen and Levinthal, 1990). The absorption of knowledge is modelled as the product of absorptive capacity (assumed to be positively associated with relatedness) and non-redundant star knowledge (assumed to be negatively associated with relatedness). The actual nature of the possibly non-linear relationship between relatedness to the star and scientist productivity is thus an empirical question.³⁴

Third, the chapter is related to the burgeoning literature in economic geography that investigates the importance of relatedness on the evolution of local economic outcomes. Although there is no explicit definition of the sense economic actors are to be related, the literature has converged to consider spillovers more likely to occur if actors draw on similar knowledge (Content and Frenken, 2016; Grillitsch et al., 2018). A distinguishing feature of this literature is the attention given to spatial proximity as well as its interactions with related knowledge through the cognitive dimension of proximity (Boschma, 2005; Torre and Rallet, 2005).

While the organizational science literature tends to focus on one-mode networks (i.e., the direct and indirect links between individuals), the economic geography literature has instead tended to emphasize two-mode affiliation networks (where the links arise through various forms of co-membership or co-affiliation). An example of such co-affiliation in our setting would be publishing in the same journals or using the same keywords, indicating joint affiliation to a particular community of practice as well as to the same knowledge space.

A major research question in this literature is how activities sharing related knowledge bases (e.g., skills, technologies) tend to facilitate local spillovers that condition the evolution of the structure of those activities over time (Boschma et al., 2013; Galetti et al., 2021; Neffke et al., 2011). This focus has led to the development of sophisticated co-occurrence metrics to measure various dimensions of relatedness. A variety of co-occurrence measures have been constructed and applied in this context, and we draw on

³⁴ The organizational science literature also highlights how the arrival of a star could affect incumbent scientists through multiple channels, including opportunities for co-authorship, access to information, the development of research know-how through more tacit knowledge transfer, etc. The nature of the absorptive capacity-redundancy tradeoff could differ across these channels.

this literature to develop suitable indicators of relatedness as our measure of the intensity of treatment on incumbent scientists following a star's arrival.

Finally, the chapter draws on the growing methodological literature on the challenges of estimating ATT in difference-in-differences and event-study frameworks in the presence of various forms of heterogeneity (Athey et al., 2018; Borusyak et al., 2022; Callaway and Sant'Anna, 2021; de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021; Sun and Abraham, 2021). These challenges may be especially pronounced when there are varying intensities of treatment. One concern is the possibility of endogenous selection of the intensity of treatment (or "dose"), which complicates the estimation of "dose-response curves" (Callaway et al., 2021). This source of potential bias may be less severe in our context as the degree of scientific relatedness should be largely fixed³⁵ (at least in the short term) at the time of the star arrival.

However, heterogeneity in responses to a given intensity of treatment can also complicate the estimation of ATT (Sun and Shapiro, 2022). Such heterogeneity can violate the linearity assumption typical in two-way fixed effects (TWFE) models leading to misspecification of the estimating equation. As a test of the robustness of our results, we therefore also apply the interaction-weighted estimator of Sun and Abraham (2021), which is designed for the unbiased estimation of the ATT under arrival-cohort heterogeneity.

We present our results in two stages. First, we use an event-study framework to provide *prima facie* evidence of a causal effect on incumbent productivity from star arrivals controlling for a measure of star-incumbent relatedness. In addition to an examination of pre-trends, we use coarsened exact matching to identify a control for each treated incumbent and we also control for university- and department-level output measures to capture factors that might affect incumbent productivity and also be correlated with an investment by the university in a star. We present both two-way-fixed effect (TWFE) estimates and estimates based on Sun and Abraham (2021) that explicitly

³⁵ One reason for relatedness being fixed in the short term could be derived from the fact that knowledge creation expands the social stock and complexity of knowledge. Therefore, researchers should increase efforts to reduce their distance from the frontier of knowledge by investing more time in learning activities or narrowing their specialization and working more regularly in teams. As knowledge is not readily transferable between researchers, this increasing "burden of knowledge" (Jones 2009) requires time for interactive learning to take place.

allow for the staggered nature of star arrivals. We find credible evidence of a positive star arrival effect that is increasing with time since arrival. We then use a difference-in-differences (DiD) analysis to examine the possible non-linear interaction between the degree of relatedness to the star and the size of the star-arrival effect, measuring the relatedness both from the perspective of the incumbent and also from the perspective of the star.

Overall, the difference-in-differences estimates we report support the presence of a non-linear relationship between the size of the star arrival effect and incumbent productivity and the extent of incumbent relatedness to the star. In terms of results not anticipated at the outset, we find that a measure of relatedness taken from the perspective of the star is a better predictor of star-arrival productivity effects than a measure taken from the perspective of the incumbent scientist. We interpret this result as indicating the importance of the presence of related scientists from the star's perspective. Such presence would incentivize the star to engage with scientists more related to the star's knowledge, allocating more time and effort to projects in which the opportunities to combine novel ideas are higher.

A potential concern regarding our baseline measure is that publications in the same journal may only partially represent the scientist's position in a knowledge space. For example, co-occurrences in general periodicals, where knowledge sets might differ from some extent is considered as related to the same degree as co-occurrences in specialized journals, where knowledge sets might be more similar. To address this issue, in addition to our baseline measure, we also explore additional networks with different topologies to ensure the robustness of our estimates. A knowledge space based on keywords expands the network and provides more detailed information on knowledge sets and scientists' network positions. On the other hand, a condensed network based on scientific categories conveys information about relatedness based on aggregated journals' aims and scopes. Deriving our measures from different knowledge spaces and obtaining substantially similar results across them serve to alleviate our initial concerns regarding our baseline indicator.

The rest of the chapter is organized as follows. Section 5.2 sets out our model, estimating framework and measures of relatedness/intensity of treatment. Section 5.3

describes the data and matching strategy. Section 5.4 reports our results, and the Section 5.5 examines their robustness to alternative assumptions and measures. Section 5.6 provides some concluding discussion of the policy relevance of our findings.

5.2 Model and Econometric Specification

5.2.1 Basic model of star arrival effects with variable intensity of treatment

In this section, we develop a simple model of a scientist's output to motivate our estimating equations. We assume that a scientist's output is the result of combining their ideas to produce new ideas. Combinatorial theories of discovery and innovation have a long history in economics and other disciplines (Arthur, 2009; Clancy, 2018; Fleming and Sorenson, 2004; Nelson and Winter, 1982; Schumpeter, 1939; Usher, 1929; Weitzman, 1998).

Our central assumption is that the number of additional ideas that an incumbent scientist can access once a star arrives – and thereby the effect on the knowledge base of ideas they combine to produce new output – will depend on how related their knowledge is to that of the arriving star. Greater scientific relatedness will therefore mean greater intensity of treatment from a star arrival. To keep the model simple, we focus on how related the scientist's knowledge is to the star's knowledge viewed from the vantage point of the scientist. However, in reality the star may be more willing to give time and attention to a scientist when that scientist's knowledge is related to their knowledge. As the knowledge bases of scientist and star are likely to be quite different – with the star's knowledge likely to range more widely – the relevant (asymmetric) relatedness measures may be quite different. While we take the scientist's perspective in outlining the model, we investigate both measures (in addition to a geometric average of the two) in the empirical application in Section 5.4.

We assume that a scientist begins (i.e., pre-star arrival) with N_{i0} ideas. These ideas are then combined to produce new scientific output, Y_{i0} . The ideas can be combined one-at-a-time, two-at-a-time up to N_{i0} -at-a-time. The number of ways to combine k ideas is then given by the binomial coefficient,

$$\binom{N_{i0}}{k} = \frac{N_{i0}!}{k!(N_{i0} - k)!} \quad (5.1)$$

We can use the binomial theorem to identify the total number of ways that N_{i0} ideas can be combined (including the combination with zero ideas) as,

$$\sum_{k=1}^{N_{i0}} \binom{N_{i0}}{k} = 2^N \quad (5.2)$$

We next assume that scientist i 's output is given by the isoelastic function,

$$Y_{i0} = \tilde{\lambda}_i \left(\frac{(2^N)^\phi - 1}{\phi} \right) e^{u_i} \quad (5.3)$$

where $\tilde{\lambda}_i$ is an individual productivity parameter, ϕ is the elasticity of output with respect to the total number of possible combinations and u_i is a normally distributed disturbance term with mean zero and variance σ^2 . Of course, as N_{i0} increases to even reasonably high levels the number of possible combinations becomes enormous, making search increasingly challenging. To counter the effects of explosive growth in the number of possible combinations, we focus on the limit as ϕ goes to zero. Using L'Hopital's rule, we can identify this limit as,

$$\lim_{\phi \rightarrow 0} Y_{i0} = \tilde{\lambda}_i \ln(2) N_{i0} = \lambda_i N_{i0} e^{u_i}, \quad (5.4)$$

where $\lambda_i = \tilde{\lambda}_i \ln(2)$. Output is thus a (random) fraction of the scientist's stock of ideas, but we allow for individual specific abilities in combining ideas into output.

We now allow for the arrival of a star in period 1. We define an indicator variable, S_i , that takes a value of 1 if a star has arrived in scientist i 's department, and 0 otherwise. Central to the impact of this star arrival on the productivity of an incumbent scientist is how related the ideas of the star are to those of that incumbent. If a scientist already has a related idea, we assume there is some probability that they will be able to access the star's idea and absorb it into their idea set. Specifically, we assume for each idea that the scientist knows, the probability that they will be able to absorb an idea possessed by the star is $\tilde{\alpha}S_i + \tilde{\beta}S_iR_{is}$, where R_{is} is a measure of scientific relatedness. The overall

absorptive capacity of an incumbent scientist with N_{i0} ideas is then $(\tilde{\alpha}S_i + \tilde{\beta}S_iR_{is})N_{i0}$. Letting N_s represent the knowledge (or idea) stock of the star, we can then write the change in the knowledge stock of the incumbent scientist as,

$$\begin{aligned}\Delta N_i &= \text{Absorptive Capacity} \times \text{Star's Knowledge} \\ &= (\tilde{\alpha}S_i + \tilde{\beta}S_iR_{is})N_{i0}N_s = (\alpha S_i + \beta S_iR_{is})N_{i0},\end{aligned}\quad (5.5)$$

where $\alpha = \tilde{\alpha}N_s$ and $\beta = \tilde{\beta}N_s$.

We can thus write the period 1 output function as,

$$Y_{i0} = \lambda_i N_{i0} \left(1 + \frac{\Delta N_i}{N_{i0}}\right) e^{u_i} = \lambda_i N_{i0} (1 + \alpha S_i + \beta S_iR_{is}) e^{u_i} \quad (5.6)$$

Taking logs and assuming $\ln(1 + \alpha S_i + \beta S_iR_{is}) \approx \alpha S_i + \beta S_iR_{is}$,

$$\ln Y_{i1} \approx \ln \lambda_i + \ln N_{i0} + \alpha S_i + \beta S_iR_{is} + u_i, \quad (5.7)$$

where we henceforth ignore the approximation. The *expected* difference in log output before and after the arrival of the star is then simply equal to $\alpha S_i + \beta S_iR_{is}$.

5.2.2 Non-linear response to an increasing intensity of treatment due to a star-knowledge redundancy effect

As hypothesized in the literature, there may be a non-linear relationship between the ability to transfer ideas between agents and the cognitive similarity (or relatedness) of those agents (Ahuja, 2000; Boschma and Frenken, 2010). From the vantage point of a scientist seeking to access the knowledge of the star, too little relatedness can undermine the capacity to absorb the star's ideas; on the other hand, too much relatedness might mean redundancy, as the scientist's and star's knowledge bases substantially overlap.

We capture this potential non-linearity by assuming that the new ideas transferred from the star to the scientist are the *product* of the scientist's absorptive capacity and non-redundant knowledge. We assume as before the total absorptive capacity is given by $(\tilde{\alpha}S_i + \tilde{\beta}S_iR_{is})N_{i0}$. The difference now is that the relevant absorption

of ideas relates only to the non-redundant knowledge possessed by the star – i.e., the knowledge that the star has but the incumbent scientist does not. The non-redundant knowledge of the star is assumed to be $N_s - \eta R_{is} N_s$; that is, the star's knowledge stock less some fraction of that knowledge already possessed by the scientist, which again depends positively on the relatedness of the scientist to the star. Appropriately amending 5.5, we can then write the actual knowledge absorbed by the scientist as,

$$\begin{aligned}
\Delta N_i &= \text{Absorptive Capacity} \times \text{Star's Non Redundant Knowledge} \\
&= (\tilde{\alpha} S_i + \tilde{\beta} S_i R_{is}) N_{i0} (N_s - \eta R_{is} N_s) \\
&= (\tilde{\alpha} N_s S_i + (\tilde{\beta} - \tilde{\alpha} \eta) N_s S_i R_{is} - \tilde{\beta} \eta N_s S_i R_{is}^2) N_{i0} \\
&= (\alpha S_i + \beta' S_i R_{is} + \gamma S_i R_{is}^2) N_{i0}
\end{aligned} \tag{5.8}$$

where $\alpha = \tilde{\alpha} N_s$, $\beta' = (\tilde{\beta} - \tilde{\alpha} \eta) N_s$ and $\gamma = -\tilde{\beta} \eta N_s$.

Now replacing 5.5 by 5.9 in 5.6 and using the approximation $\ln(1 + \alpha S_i + \beta' S_i R_{is} + \gamma S_i R_{is}^2) \approx \alpha S_i + \beta' S_i R_{is} + \gamma S_i R_{is}^2$, the period 1 log output equation for the scientist becomes,

$$\ln Y_{i1} \approx \ln \lambda_i + \ln N_{i0} + \alpha S_i + \beta' S_i R_{is} + \gamma S_i R_{is}^2 + u_{i1}, \tag{5.9}$$

where we again henceforth ignore the approximation. Taking the partial derivative of log output with respect to the measure of scientific relatedness and setting it equal to zero, we can find the value of relatedness at which the star arrival effect is maximized as $-\beta'/2\gamma$, which will be positive when β' is positive and γ is negative (with the latter following given a redundancy effect). This optimal value of relatedness can be seen as the result of balancing the trade-off between the benefit of increased relatedness due to

improved absorptive capacity and the cost of increased redundancy as the scientist and star become more related.³⁶

5.2.3 Measuring scientific relatedness using co-occurrence metrics

Our theoretical measure of scientific relatedness captures the ease in which the scientist can access the knowledge of the star and is assumed to range from 0 (no relatedness) to 1 (full relatedness). In our empirical analysis, we need to proxy the scientific relatedness measure based on observed indicators of overlap between the scientist's and the star's knowledge. Although we explore alternative indicators, one obvious approach is to use the journals that the scientist and star have published in and to look for co-occurrence (or overlap) of publications in the same journal.

Given C journals ($c = 1, \dots, C$), we define $F_{ic} = 1$ if i has published in journal c and $F_{sc} = 1$ if the star has published in journal c ; each publication indicator takes a value of zero if no publication has occurred. There is a co-occurrence of publication in journal c if $F_{ic}F_{sc} = 1$. We set the horizon for examining occurrences and co-occurrences of publications in the previous \bar{t} years (which we take to be 4 years in our empirical application). The total number of journals that the scientist and the star have published in is then $\sum_{c=1}^C F_{ic}$ and $\sum_{c=1}^C F_{sc}$ respectively. And the total number of co-occurrences is $\sum_{c=1}^C F_{ic}F_{sc}$.

We adopt the classic cosine similarity index as our baseline measure of relatedness.

$$\bar{R}_{is} = \frac{\sum_{c=1}^C F_{ic}F_{sc}}{\sqrt{(\sum_{c=1}^C F_{ic})(\sum_{c=1}^C F_{sc})}} \quad (5.10)$$

³⁶ The model has been motivated assuming that what matters is the relatedness of the incumbent scientist to the star measured from the perspective of the incumbent scientist. An alternative motivation is to assume the relatedness that matters is the measure of relatedness from the perspective of the star. If the latter is what matters, we can reinterpret the absorptive capacity term as the engagement potential with the incumbent from the perspective of the star, with greater engagement potential rising with the cognitive similarity of the knowledge stocks. The redundancy term can be reinterpreted as the network redundancy from the perspective of the star, where we are assuming that the more related the incumbent is to the star the more pronounced is the network redundancy effect. We take both perspectives – as well as a geometric average of the two – in our empirical application.

Similarity is here identified with overlap between the scientist's and star's scientific activity (e.g., they have both published in the same journal). This index has a number of desirable features: (i) it ranges from 0 (no relatedness) to 1 (complete relatedness); (ii) it can usefully be represented as the geometric average of two asymmetric similarity measures, one taken from the perspective of the scientist and one taken from the perspective of the star; and (iii) it can be readily calculated from the publication/keyword information available in Scopus.

As previously noted, two different facets of relatedness are likely to matter – one taken from the perspective of the scientist and one taken from the perspective of the star. From the scientist's perspective, we assume that what is relevant is the overlap between their knowledge base and the knowledge base of the star. We measure this as the fraction of their knowledge base that overlaps with the star, which suggests the asymmetric relatedness measure:

$$R_{is} = \frac{\sum_{c=1}^C F_{ic} F_{sc}}{\sum_{c=1}^C F_{ic}} \quad (5.11)$$

However, from the star's perspective, they have to decide which scientists to devote time and attention in transferring knowledge. While from the scientist's perspective, there might be a large overlap in knowledge bases; from the star's perspective there might be relatively little overlap with their total knowledge base given their research is likely to range more widely across research topics. The asymmetric relatedness measure taken from the perspective of the star then replaces the denominator in 5.11 with the total occurrences of the star (e.g., the total number of journals they have published in) and takes the form:

$$R_{si} = \frac{\sum_{c=1}^C F_{ic} F_{sc}}{\sum_{c=1}^C F_{sc}} \quad (5.12)$$

In reality, both of these asymmetric relatedness measures are likely to matter – how related the knowledge bases are from the perspective of the incumbent scientist *and* how related those knowledge bases are from the perspective of the time-constrained star. To combine these two forms of relatedness in a simple way, we define an effective-

relatedness index (\bar{R}_{is}) given by the symmetric constant-returns-to-scale Cobb-Douglas form,

$$\bar{R}_{is} = R_{is}^{0.5} R_{si}^{0.5} \quad (5.13)$$

The effective-relatedness index is then the geometric mean of the two asymmetric measures. Substituting 5.11 and 5.12 into 5.13 yields 5.10, our baseline cosine similarity measure above.

5.2.4 Difference-in-differences specification

In the case where the heterogeneity is appropriately modelled using a quadratic specification, our baseline estimating equation is,

$$\ln Y_{i,t} = \alpha S_i + \beta \bar{R}_{is} S_{i,t} + \gamma \bar{R}_{is}^2 S_{i,t} + Z'_{i,t} \lambda + \lambda_i + \lambda_t + u_{i,t} \quad (5.14)$$

where $Y_{i,t}$ is a measure of quality-adjusted output (e.g., citations-weighted publications) for researcher i in year t , $S_{i,t}$ is a dummy variable indicating that a star is present in time period t , \bar{R}_{is} , is the (symmetric) measure of relatedness, $Z'_{i,t}$ is a vector of controls, λ_i is an individual scientist fixed effect and λ_t is a time fixed effect. We include (non-interacted) relatedness and relatedness squared terms as controls in the regression as well as a set of university and department level output variables, where the latter are important to include given the potential endogeneity of star arrivals. Statistical inference based on robust standard errors clustered at the individual level to allow for arbitrary forms of serial correlation and heteroscedasticity (Kripfganz, 2016). We apply this specification to the symmetric relatedness measure and also the asymmetric measures from the perspective of the incumbent scientist and the perspective of the star respectively. In our robustness analysis, we also explore alternative estimation methods and also a number of alternative co-occurrence based approaches to produce the relatedness measure.

5.3 Data

Our dataset consists of information on researchers' publications, citations recorded, references, abstracts, and acknowledgements collected from Scopus across 27 subject fields for the period 1990 to 2017. Given the interest in star recruitment policies to catalyse development in smaller regions, we focus on researchers in three small open economies – Denmark, Ireland, and New Zealand. Denmark and Ireland have implemented nationally supported formal programmes for the recruitment of star researchers. New Zealand is included as a small open economy that has not implemented a formal star-recruitment programme. After data cleaning, we obtained approximately 1.43 million publications divided over 219,582 unique authors, spreading across 457 departments, comprising 153 departments in Denmark, 141 in Ireland, and 163 in New Zealand.

5.3.1 Identification of stars and star arrivals

We first define a star scientist as a scientist who is at or above the 95th percentile of scientists in the cumulative distribution of citations received since 1990 for their subject field in any given year, where the citation measure includes all forward citations to a publication up to 2019.³⁷ In order to allow a reasonable time for citations to accumulate, the stars are identified from 1996 onwards. We then identify a star arrival at a department as a star that had not previously published with an affiliation to the department now recording such an affiliation. We restrict identified star arrivals to those who publish with an affiliation to the department for at least the next four years. For those that published in 2015-2016, they must also be observed to publish at the same affiliation at the end of our observation window in 2017. Overall, we identify 167 star arrivals over the period 1996 to 2017 in total using this procedure. As our empirical implementation applies only the first star arrivals to a department, we focus on 91 such first star arrivals, excluding the 1996 cohort³⁸, across the three countries – 45 in Denmark, 20 in Ireland,

³⁷ Stars scientists were allocated to departments (subject fields) where they have the largest number of publications. They were assigned to multiple departments in cases where they have an equal distribution of publications across multiple subject fields. In terms of the same first star arrival at multiple departments.

³⁸ We exclude the 1996 arrival cohort to allow for appropriate comparison between the homogenous and heterogeneous cohort arrival effects models. Using the approach of Sun and Abraham (2021) to account for heterogeneous cohort arrival effects requires the exclusion of cohorts that are always treated across the observation window (1996 cohort here). This exclusion has minimal effects on the baseline results.

and 26 in New Zealand. We then identify those focal incumbents with and without the star arrival in their departments as treatment and (potential) control groups.

Table 5.1 Balancing test between treatment and control groups

Variable	Control set (N = 6,992)	Treated set (N= 6,992)	Diff. in mean	P-value
	Mean	Mean		
Citation received	36.173	37.722	-1.049	0.2649
Career age (<i>Bins</i>)	2.285	2.285	0.000	1.000
Cumulated citation received (<i>Bins</i>)	3.760	3.792	-0.032	0.715
Career age (<i>Num.</i>)	9.343	9.314	0.029	0.8198
Cumulated citation received (<i>Num.</i>)	393.760	398.872	-5.112	0.769

Source: Authors' calculations

5.3.2 Coarsened exact matching procedure (CEM) for incumbent sample

To help alleviate an endogeneity of treatment concern in our empirical analysis, we use a k-to-k Coarsened Exact Matching (henceforth, CEM) approach to identify appropriate matches for incumbent scientists in star-treated departments (Azoulay, Fons-Rosen, and Graff Zivin 2019; McHale et al. 2022). We first identify all incumbents (i.e., the treated authors) in the star-receiving department who are present four years before the star's arrival and four years after the arrival. We then combine these treated incumbents with the potential control authors (i.e., those who are not in a star-receiving department) with selected covariates in the year before the star arrival. The covariates used for matching are as follows: scientist career age (i.e., years since first publication); citations received by the scientist in the year before the star arrival; cumulative citations received by the scientist up to the year before arrival; subject field; country; and year. We utilize relatively coarse bins for career age (13 bins) and for cumulative citations (26 bins) but require an exact match for subject field and year. Importantly, we provide a balancing test for the matched pairs (Table 5.1) and confirm the appropriate CEM matching procedure that identifies good matches for our sample of incumbents. In so doing, we successfully identify 13,984 matched pairs, resulting in a total of 222,287 individual-year observations.

5.3.3 Scientific relatedness measure

To obtain relatedness between stars and incumbent scientists, we therefore employ each researcher's publication record to build a scientific knowledge space that

proxies individuals' accumulated knowledge sets. The baseline method for computing relatedness between a star and an incumbent scientist is the extent of co-occurrence of publications in given journals. Since the journals' aims and scope encompass a variety of related themes, scientists that publish in the same journal are more likely to share a 'related' knowledge set. The relatedness measure between a star-incumbent pair in the Journal Overlap knowledge space therefore increases with the number of co-occurrences of publications in a journal in a given period. We measure relatedness using a four-year moving window from $t - 4$ years to $t - 1$ year for any year t in our observation window.

To deal with concerns of potential mismeasurement of related knowledge sets between scientists by relying only on the journal overlap measure, we introduce three additional knowledge spaces in the robustness section. These spaces are based on Journal Category Overlap, Author-Reported Keyword Overlap, and Scopus-Reported Keyword Overlap. These alternative sources of relatedness differ in their network topology and variety of knowledge sets compared to the journals. Table 5.2 provides some distributional characteristics across the symmetric and asymmetric measures for each of the additional knowledge spaces. Author-Reported³⁹ and Scopus-Reported Keyword overlaps both expand knowledge space by adding more detailed information about related themes discussed in the publication. Conversely, relatedness based on Journal Category Overlap narrows the knowledge space by aggregating different journals with similar aims and scope into the same scientific category.

To construct the additional relatedness measures, we employ information on 15,406 journals that stars and incumbents have published, grouped into 304 Scopus subject categories. Every category encompasses, on average, 120 journals, and the number of journals across each category ranges from one journal in "Nursing – Review and Exam Preparation" (category 2923) to 1,987 journals in "Medicine – Miscellaneous" (category 2701). In addition, each publication provides, on average, 5.10 author keywords and 24.95 index keywords.

³⁹ The data on author provided keywords raises a concern on selection bias due to missing data on keywords - about 47% of publications' keywords are missing. To mitigate this concern, we employ Scopus provided (index) keywords, which reduce missing data issue from 47% to 23%. Table D.1 in the section D.1 of the Appendix provides information regarding missing keywords between the two sources by department.

Table 5.2 Distributional characteristics for the relatedness measures

	Mean	Percentiles					
		25 th	50 th	75 th	90 th	95 th	99 th
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Journal Overlap							
Asymmetric (Scientist perspective)	0.031	0	0	0	0.071	0.222	0.667
Asymmetric (Star perspective)	0.014	0	0	0	0.031	0.083	0.275
Symmetric (Cosine)	0.017	0	0	0	0.057	0.136	0.326
Quantity-Adjusted (QA) Symmetric (Cosine)	0.018	0	0	0	0.021	0.112	0.456
Journal Category Overlap							
Asymmetric (Scientist perspective)	0.242	0	0.133	0.4	0.667	1	1
Asymmetric (Star perspective)	0.124	0	0.055	0.185	0.347	0.500	0.909
Symmetric (Cosine)	0.152	0	0.112	0.267	0.405	0.490	0.667
QA Symmetric (Cosine)	0.157	0	0.041	0.257	0.492	0.636	0.881
Author-Reported Keyword Overlap							
Asymmetric (Scientist perspective)	0.008	0	0	0	0	0.054	0.2
Asymmetric (Star perspective)	0.003	0	0	0	0	0.020	0.076
Symmetric (Cosine)	0.004	0	0	0	0	0.032	0.084
QA Symmetric (Cosine)	0.004	0	0	0	0	0.023	0.088
Scopus-Reported Keyword Overlap							
Asymmetric (Scientist perspective)	0.004	0	0	0	0.010	0.023	0.066
Asymmetric (Star perspective)	0.007	0	0	0	0.023	0.043	0.111
Symmetric (Cosine)	0.004	0	0	0	0.017	0.028	0.058
QA Symmetric (Cosine)	0.004	0	0	0	0.012	0.028	0.069

Source: Authors' calculations

5.3.4 Productivity of researchers

Our dependent variable is a quality-adjusted measure of scientist productivity in a given year. The quality adjustment is based on the citations received by a publication to the end of the sample period. In recognizing that different scientific fields display quite different citation patterns, the productivity measure is normalized based on the average citations to a publication in the relevant field, yielding a measure of the field-normalized total citations (Perry and Reny, 2016; Radicchi et al., 2008). The measure is defined as:

Field Normalized Total Citations:

$$Y_{it}^{FNTC} = \sum_{p_{i,t}=1}^{P_{i,t}} \frac{c_{p_{i,t}}}{\bar{c}_{s,t}};$$

where $P_{i,t}$ is the total number of publications by scientist i in year t , $c_{p_{i,t}}$ is the total subsequent citations (or "forward" citations recorded at 2019) to a publication $p_{i,t}$ by scientist i in year t , and $\bar{c}_{s,t}$ is the average citations to a publication in the relevant subject

field, s , for publications that occur in year t . We explore the robustness of our results to alternative output measures based on raw publication counts and an alternative measure of quality-adjusted publication based on the Euclidean index.

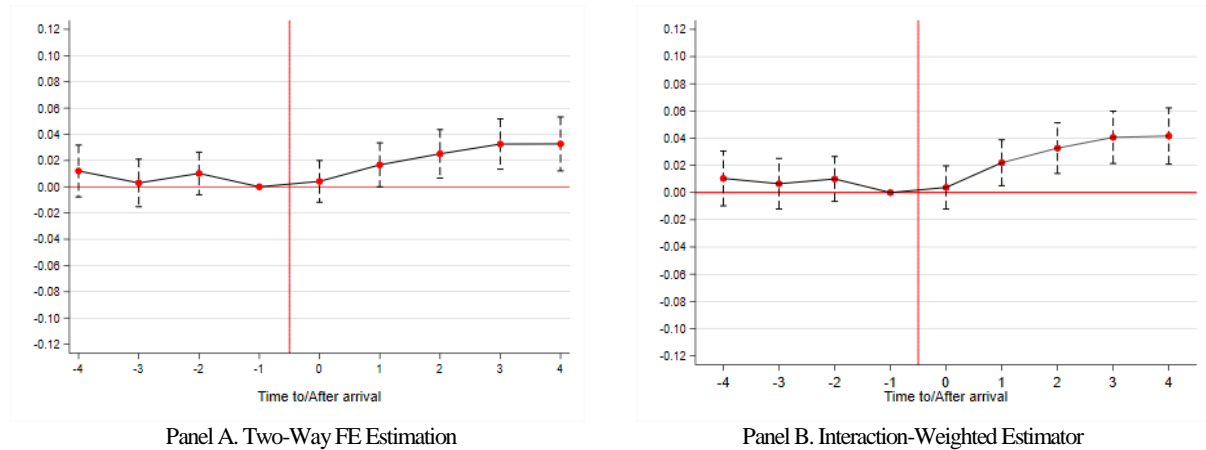


Figure 5.1. Event-study model for star arrival effects on incumbents' output

Note: The figure plots the dynamic effect of the star arrival with 95% confidence intervals based on the matched sample. Panel A shows the homogeneous effects results and Panel B shows heterogeneous effects results based on Sun and Abraham (2021). Both specifications include relatedness, a quadratic term for relatedness, and university and department outputs as control. The results provide prima facie evidence for positive causal effect of star arrivals on incumbent productivity. Moreover, the size of the star effect is generally increasing with time since arrival, potentially reflecting the star's embedding in their new department.

5.4 Results

5.4.1 Star arrival effects on incumbents' output

As a preliminary step to examining how relatedness to the star intermediates the size of the star arrival effect, we first use a staggered event-study framework to examine the evidence for a causal star arrival effect controlling for the level of relatedness. Although we lack a good instrument for star arrivals, we take a number of steps to establish if there is credible evidence of a causal effect. First, we select controls for each of our incumbent scientists using the coarsened exact matching procedure. Second, using both a TWFE estimator and the Sun and Abraham (2021) estimator to explicitly allow for the staggered nature of star arrivals, we carefully examine the evidence for a no-pre-trends assumption, including visual inspection and application of the Roth (2022) test. Finally, we include a department- and university-level outputs in all our regressions to control for university-wide factors that could potentially affect the choice to recruit stars as well the productivity of incumbent scientists independent of any productivity effect through star arrivals. The results of this event-study analysis is shown in Figure 5.1,

Panels A and B. With this evidence for a causal star-arrival effect, we next apply a DiD analysis to examine the possibility of a non-linear relationship between various measures of the relatedness of incumbents to the star and the size of the star-arrival effect on their productivity.

Table 5.3 Scientific relatedness to star: journal overlap

	Baseline	Sun & Abraham estimator	(3)	(4)
<i>Panel A: Symmetric (Cosine)</i>				
$\alpha S_{i,t}$	0.022*** (0.007)	0.027*** (0.008)	0.015** (0.006)	0.013** (0.006)
$\beta \bar{R}_{is,t^*-1} S_{i,t}$			0.231*** (0.087)	0.410*** (0.140)
$\gamma \bar{R}_{is,t^*-1}^2 S_{i,t}$				-0.476* (0.272)
R ²	0.030	0.034	0.031	0.031
F-stat			6.55	4.76
p-value			0.0014	0.0026
<i>Panel B: Asymmetric (Scientist perspective)</i>				
$\alpha S_{i,t}$			0.0176** (0.00700)	0.0155** (0.00694)
$\beta R_{is,t^*-1} S_{i,t}$			0.117*** (0.0386)	0.121 (0.0909)
$\gamma R_{is,t^*-1}^2 S_{i,t}$				-0.0444 (0.0979)
R ²			0.030	0.031
F-stat			8.37	4.15
p-value			0.0002	0.006
<i>Panel C: Asymmetric (Star perspective)</i>				
$\alpha S_{i,t}$			0.0196*** (0.00696)	0.0149** (0.00684)
$\beta R_{si,t} S_{i,t}$			0.0592 (0.0797)	0.469*** (0.166)
$\gamma R_{si,t}^2 S_{i,t}$				-0.529*** (0.200)
R ²			0.031	0.032
F-stat			4.31	4.43
p-value			0.0135	0.0040
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Uni./Dep.	Yes	Yes	Yes	Yes
Relatedness controls	No	No	Yes	Yes
Observations	222,287	222,287	222,287	222,287

Notes: We control for university and department outputs through all specifications (Uni. /Dep.).


5.4.2 Relatedness to the star and the size of the star arrival effects

Our baseline results for the journal co-occurrence relatedness measure are shown in Table 5.3. Panels A, B, and C show the results for the cosine relatedness measure, the asymmetric relatedness measure from the perspective of the scientist, and the

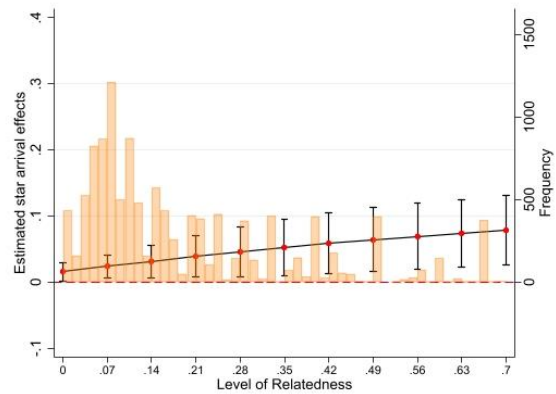
asymmetric relatedness measure from the perspective of the star respectively. Specifications in columns (3) and (4) include the relatedness-star interactions to model the treatment effect heterogeneity across the degrees of relatedness between the incumbent scientist and the arriving star. In addition, both these specifications control for the relatedness variables separately to the interaction terms. Furthermore, across all specifications we control for department and university aggregate output and include a full set of individual scientist and time fixed effects.

Column 1 of Panel A in Table 5.3 records the basic star arrival effect without star-arrival/relatedness interactions. The coefficient estimate of 0.022 indicates that a star arrival is associated with a 2.22 log point (2.22 percent) increase in incumbent output on average. One possible concern is that our basic star arrival effect is assumed to be homogeneous across the timing of star arrivals, whereas a recent literature has indicated that in settings with variation in the timing of treatment and heterogeneous effects across different timings, the coefficient on the ATT could be biased (Athey et al., 2018; Borusyak et al., 2022; Callaway and Sant'Anna, 2021; de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021; Sun and Abraham, 2021). In a robustness test of the homogeneity assumption for our basic star arrival effect, we implement the interaction-weighted estimator of Sun and Abraham (2021) to allow for staggered treatments and any potential treatment heterogeneity across star arrival-cohorts. This approach estimates the aggregate star arrival effect as an appropriately weighted average of "cohort average treatment effects on the treated" (CATT). More specifically, it provides a single overall effect estimate by first calculating the average CATT for each treatment cohort across all post-treatment periods and then computing a weighted average CATT across cohorts, where the weights are based on each cohort's total sample share. Column 2 of Panel A in Table 5.3 reports the unbiased estimate of the ATT under arrival-cohort heterogeneity. The result suggests that the estimated star arrival effect is very similar to the homogeneous basic star arrival effect, with a 2.27 log point (2.27%) increase in incumbent output on average after the arrival of the star. Overall, the star arrival effect is found to be robust under the assumptions of both homogeneous and heterogeneous cohort arrival effects.


In terms of modelling and testing for the heterogeneity in responses to treatment intensity based on the degree of relatedness to the arriving star, Column 3 adds a linear star-arrival/relatedness interaction, with relatedness measured using the different



Panel A. Symmetric (Cosine) scientific relatedness effects through relatedness levels (Max R_{is} =0.43)



Panel B. Asymmetric scientific relatedness effects through relatedness levels (scientist's perspective, Max R_{is} =1.375)



Panel C. Asymmetric scientific relatedness effects through relatedness levels (star's perspective Max R_{si} =0.318)

Figure 5.2 Non-linear relationship between star arrival and scientific relatedness effects – journal overlap space

Notes: We estimate the conditional marginal effects of star relatedness through different levels of relatedness given by the star arrival. Standard errors for the marginal effects are calculated using the delta method

measures described above. Based on the cosine measure, the results show a positive and statistically significant association between the level of relatedness and the estimated star arrival effect. Further to this, to allow for a possible non-linearity between the star-arrival effect on incumbents and the measure of relatedness, Column (4) includes a quadratic term for the star-arrival/relatedness interaction. The coefficient on the quadratic term is negative and statistically significant for the cosine measure, suggesting that there is an increasing relevance of the redundancy effect as relatedness increases.

The interrelationship between the size of the star arrival effect and the measure of relatedness is conveniently illustrated using Figure 5.2. Panel A of Figure 5.2 is based on the estimates using the cosine measure shown in Column (4) of Panel A in Table 5.3. In addition to showing how the estimated star arrival effect evolves with the measure of relatedness, the figure also shows the distribution of the relatedness measure across our sample, indicating the mean and median levels of relatedness are quite low. The estimated star arrival effect is positive at all relatedness levels, although the standard errors around the star arrival effect are large at high relatedness levels where the in-sample observations for relatedness are very small in numbers. These standard errors are calculated using the delta method. The peak star arrival effect occurs at a relatively high level of relatedness at 0.43, around the 99th percentile. Overall, although there is evidence of a redundancy effect, for the large majority of incumbents the overall star arrival effect is increasing in the measure of relatedness to the star.

As discussed in Section 5.3, the cosine measure of relatedness can be usefully represented as a geometric weighted average of two asymmetric relatedness measures – one taken from the perspective of the incumbent scientist and one taken from the perspective of the arriving star. Indeed, the cosine measure is highly correlated with both the asymmetric measures (see Table 5.4). In considering these facets of relatedness, Columns (3) and (4) in Panels B and C of Table 5.3 present the results based on the asymmetric relatedness measures from both the perspectives of the incumbent scientist and the star. Additionally, Figure 5.2 illustrates the star arrival effect across the values of both relatedness measures in Panel B and C respectively. From the estimates in Column (4) of Panel B in Table 5.3 that use the asymmetric measure from the incumbent scientist’s perspective we see a similar pattern to the one observed from using the cosine measure in Panel A, with a positive sign on the coefficient for the star-arrival/relatedness interaction term and a negative sign on the quadratic term. However, we find no evidence that either coefficient is statistically significant.

In contrast, the results reported in Column (4) of Panel C in Table 5.3 – which uses the asymmetric relatedness measure from the perspective of the arriving star – show a very similar pattern to the one observed for the cosine relatedness measure. From Panel C in Figure 5.2, the estimated star-arrival effect peaks at relatedness of 0.318, around the

99th percentile, and indicates that, comparable to the results using the cosine measure, there is evidence of a redundancy effect.

The results suggest asymmetric relatedness measures matter for local interactive learning and may differ because of the differences in incumbents' and stars' knowledge

Table 5.4 Correlation matrix for different scientific relatedness measures

Types of scientific relatedness	Asymmetric (Scientist perspective) (1)	Asymmetric (Star perspective) (2)	Symmetric (Cosine) (3)
Asymmetric (Scientist Perspective)	1		
Asymmetric (Star Perspective)	0.513***	1	
Symmetric (Cosine)	0.843***	0.846***	1

Notes: *, **, *** denote significance at 10%, 5% and 1% levels, respectively.

sets. Less prominent incumbent scientists are expected to be more related to stars than the other way around (Tversky, 1977). On the one hand, high-related incumbents, based on their own perspective, are likely to efficiently absorb some fraction of what the star knows and are less constrained by knowledge redundancy as relatedness increases. On the other hand, to maximize their payoffs, the star needs to allocate more time and attention only to scientists positioned at a relevant range of relatedness from the star's perspective. In this case, less-related incumbents, constrained by low levels of absorptive capacity, and more-related ones, with high levels of redundant knowledge, would receive less attention from the star scientist. This line of reasoning can explain the differences between a more linear and an inverse U-shaped relationship between relatedness and productivity observed from incumbents' and stars' perspectives in Figure 5.2. However, as it is reasonable to consider that incumbents and stars would choose to behave actively to create novel ideas, the cosine measure that combines both perspectives tends to prevail in an interactive learning environment.

In general, the results suggest that both the cosine measure and the asymmetric measure capturing how related the incumbent scientist is to the star *from the perspective of the star* have more power to explain the observed star arrival treatment effect heterogeneity based on the degree of relatedness. This in turn could suggest that a star's decisions in terms of allocating time to sharing knowledge with incumbents – with most time being given to incumbents with knowledge that overlaps most with their own – is an important factor for the impact of the star arrival on the productivity of incumbents.

5.5 Robustness

In this section, we conduct two robustness tests on our baseline model. The first test consists on comparing our results to estimations based on relatedness measures taken from different knowledge spaces. In the second test, we adopt a non-parametrical approach to assess whether the imposition of a quadratic functional form to associate relatedness and the size of the star arrival effect is restrictive in our baseline specification. Due to space constraints, we further include four additional robustness tests in the section D.3 of the Appendix. We first test a quantity-adjusted symmetric relatedness measure based on the number of scientist's publication. In the second set of tests, we examine the sensitivity of the previous results to alternative scientist productivity measures: the raw publication counts for incumbent i in year t , and a scaled Euclidian Index proposed by Perry and Reny (2016). Third, we exclude from the sample all incumbent scientists that are also stars in our definition and repeat our estimates. Finally, as a further test of appropriateness of a quadratic functional form of the relatedness relationship, we separately estimate the size of the star arrival effect for different ranges of the relatedness measure. Overall, the results appear generally robust to the choice of alternative relatedness and productivity measures, and we see evidence consistent with a quadratic pattern in the unrestricted estimates.

5.5.1 Robustness to alternative scientific knowledge spaces

One important concern is that by solely using the journal overlap measure of relatedness we may only be partially representing the similarity between arrival star's and focal incumbents' knowledge sets and disregarding the fact that they may still share related knowledge sets without necessarily publishing in the same journals (Waltman and van Eck, 2012). This could imply that we are underestimating (or overestimating) the formation of scientific knowledge space that impacts on the scientific-relatedness effect. Therefore, we examine the sensitivity of the star arrival/relatedness relationship

to three alternative relatedness measures based on: broader journal categories; authors' keywords; and the keywords reported by Scopus.⁴⁰

It is important to note that these additional dimensions used to capture alternative aspects of similarity between stars and incumbents across different knowledge spaces do not guarantee that related researchers in one specific space are related in another (Tversky, 1977). An incumbent and a star are related if they have published in a broader journal category, e.g., Medicine-Miscellaneous, but they are not related if published in different journals. In the same example, they may or may not be related depending on the keywords they reported, regardless of the journal being the same. Because of this variety of sources used to compute alternative knowledge spaces, the measures may differ from each other and from our baseline Journal Overlap space.

Table 5.5 Correlation matrix between relatedness across alternative knowledge spaces

Knowledge Spaces	Journal Overlap	Journal Category Overlap	Author-Reported Keyword Overlap	Scopus-Reported Keyword Overlap
	(1)	(2)	(3)	(4)
Journal Overlap	1			
Journal Category Overlap	0.483***	1		
Author-Reported Keyword Overlap	0.319***	0.244***	1	
Scopus-Reported Keyword Overlap	0.245***	0.288***	0.453***	1

Notes: *, **, *** denote significance at 10%, 5% and 1% levels, respectively

Though all relatedness measures derive from different knowledge spaces, they are positively (but imperfectly) correlated (see Table 5.5), indicating that scientists in our sample consistently share related knowledge sets across distinct dimensions of publication records, but the relatedness measure will be sensitive to the precise knowledge space that is chosen.

The results for the alternative measures are reported in Columns (2) to (4) in Table 5.6, while Column (1) repeats the baseline results from the journal overlap measure for comparison purposes. In addition, Figure 5.3, Panel A to Panel D, illustrates the results from the separate measures to allow for convenient visual comparisons to the baseline results. The results across the alternative measures are broadly comparable to

⁴⁰ We use exact matching for keywords same in stars and incumbent knowledge set, and Jaro-winkler distance measure to measure the similarity between two keywords. The Jaro-Winkler provides a score between 0 and 1, where 0 equates to exact match and 1 is no similarity. We choose a maximum string distance cut-off of 0.15, as higher values resulted in dissimilar keywords matching.

the baseline pattern of star arrival effects with some notable exceptions. First, in adopting the Journal Category Overlap measure to allow for relatedness between broader knowledge sets, we find that the coefficient on the star arrival binary variable is not statistically significant (see Column (2)). Second, in considering the results for Scopus-Reported Keyword Overlap in Column (4), which is calculated based on the controlled vocabulary terms allocated to the publications by Scopus, we see that the coefficients on the star-arrival/relatedness interaction term and the quadratic term are larger in magnitude compared to the baseline results but the quadratic term is not statistically significant. Notwithstanding the differences in knowledge spaces, the general pattern with positive and statistically significant linear star-arrival/relatedness interactions, and negative coefficients for the quadratic term, although not always statistically significant, suggest that our results are generally robust to the choice of alternative scientific relatedness measures.

Table 5.6 Robustness of alternative knowledge spaces


	Journal Overlap (1)	Journal Category Overlap (2)	Author-Reported Keyword Overlap (3)	Scopus-Reported Keyword Overlap (4)
$\alpha S_{i,t}$	0.0135** (0.00689)	-0.00814 (0.00743)	0.0180*** (0.00696)	0.0133* (0.00692)
$\beta \bar{R}_{is,t^{*-1}} S_{i,t}$	0.410*** (0.140)	0.162*** (0.0537)	0.737** (0.295)	1.632*** (0.509)
$\gamma \bar{R}_{is,t^{*-1}}^2 S_{i,t}$	-0.476* (0.272)	-0.0395 (0.102)	-1.731*** (0.601)	-5.102 (4.523)
Individual FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Relatedness Controls	Yes	Yes	Yes	Yes
Observations	222,287	222,287	222,287	222,287
R ²	0.031	0.032	0.031	0.031
F-stat	4.76	10.20	5.14	6.63
p-value	0.0026	0.0000	0.0015	0.0002

Notes: See Notes in Table 5.2


5.5.2 Robustness to no quadratic functional form restrictions.

While an extensive literature points to an inverted U-shaped relationship between relatedness and the size of the star arrival effect, in our empirical application above we treat the form of the relationship – including the nature of any non-linearity – as an empirical question. We impose a quadratic functional form restriction on the relationship and find that the estimated quadratic term is consistently negative in all specifications reported in Sections 5.4 and 5.5.


In order to test how restrictive the imposed quadratic functional form could potentially be, we also nonparametrically examine the form of any non-linearity between relatedness and the size of the star arrival effect using the semiparametric fixed-effects estimator proposed by Baltagi and Li (2002). We use the same specification as outlined in 5.14 and based on the journal co-occurrence measure; however, instead of imposing the quadratic functional form on the interaction term, $\bar{R}_{is}S_{i,t}$, this term now enters the model nonparametrically. The estimator works by first partialling out the fixed effects and the parametrically independent variables from the error component and then, estimating the net nonparametric relationship between the natural logarithm of quality adjusted output, $\ln Y_{i,t}$ and the interaction term, $\bar{R}_{is}S_{i,t}$, based on a local polynomial fit of the partialled out residuals on the interaction term. Like our baseline specification, we apply this estimation to the symmetric relatedness measure as well as both asymmetric relatedness measures; from the perspective of the incumbent scientist and the perspective of the star.



Panel A. Symmetric (Cosine) scientific relatedness effects through relatedness levels (Max $\bar{R}_{is}=2.05$): Journal Category Overlap



Panel B. Symmetric (Cosine) scientific relatedness effects through relatedness levels (Max $\bar{R}_{is}=0.21$): Author-Reported Keyword Overlap



Panel C. Symmetric (Cosine) scientific relatedness effects through relatedness levels (Max $\bar{R}_{is} = 0.16$): Scopus-Reported Keyword Overlap

Figure 5.3 Non-linear relationship between star arrival and scientific relatedness effects – alternative knowledge spaces

Notes: See Notes in Figure 5.2

The results from the semiparametric fixed effects estimator examining the nonparametric fit for the interaction term, $\bar{R}_{is}S_{i,t}$ are illustrated in Figure 5.3. Similar to the main findings, the relationship between relatedness and the size of the star arrival effect, based on both the cosine relatedness measure and the asymmetric relatedness measure from the perspective of the arriving star, resemble an inverted U-shaped pattern. These results suggest the presence of a redundancy effect as relatedness increases, without imposing any quadratic functional form restriction. Furthermore and not inconsistent with the main results, this pattern is less prominent in the nonparametric fit based on the asymmetric relatedness measure from the perspective of the incumbent scientist. To further validate the presence of the inverted U-shaped relationship, we conduct an additional robustness test by dividing the overall sample into sub-samples based on a range of different levels of relatedness and then, estimating

separate DID models for each subsample to examine for any non-linearity in effects across the levels of relatedness. Similarly, we obtain comparable results that support our earlier findings and these are reported in full in the section D.3.4 of the Appendix. Overall, these results combined indicate that imposing a quadratic functional form in the main specification does not pose a restrictive assumption on the relationship being examined.




Figure 5.4 Nonparametric relationship between star arrival and scientific relatedness effects based on the journal co-occurrence measure

5.6 Concluding Comments

The results of our difference-in-differences estimations show robust evidence that peer effects catalysed by the arrival of the star scientist vary according to the intensity of incumbent relatedness to the star. While, overall, the size of this productivity effect is increasing with relatedness for almost all incumbents, the results reveal that it is relatedness from the perspective of the star that is most important in intermediating the size of the productivity impact, with clear evidence of a peak in the incumbent-productivity effect at less than full relatedness.

These results suggest two important considerations for policymakers and receiving institutions in the design of star-recruitment policies and strategies. First, in selecting stars for recruitment, our findings suggest the importance of targeting individuals with a relevant range of knowledge related to those of scientists already present in the recruiting institution. Higher productivity derived from interactions between high-related scientists seems to be a strong incentive for stars to engage with

incumbents. Recognizing that stars are likely to have expertise across a reasonably broad range of specialized research areas, many incumbents will be likely to find it advantageous to interact with the star. However, a time-constrained star may require substantial level of relatedness from their perspective to be incentivized to make the necessary time investments to engage with an incumbent scientist. The likely level of engagement of the star depends on the level of relatedness to their peers and is therefore an important consideration in the star selection process.

Second, in maximizing the benefit from a star's arrival, the receiving institution will need to increase the level of scientific relatedness as well as lower the cost of engagement. This suggests, for example, the importance of co-location and opportunities for learning interactions in the receiving institution as well as measures to integrate the star into local networks. Other recent evidence supports the importance of star co-authorship (Yadav et al., 2023) and star mentorship (Sasidharan et al., 2022) as mechanisms through which incumbent scientists gain from the presence of a star. But each of these mechanisms requires active engagement from the star, which is facilitated by the level of relatedness. A growing star-arrival benefit over time will also be supported by subsequent recruitments (and other investments) in areas related to the star, so that star recruitments are likely to be most beneficial as part of a broader investment strategy to expand activity in the star's research areas.

In sum, the evidence shows that star recruitments can yield substantial benefits for the receiving institution, potentially helping to catalyse capacity development in targeted research areas. But to be effective, such recruitments need to take into consideration the cognitive proximity from the knowledge sets already existing in the receiving department and be complemented by efforts that ensure active star engagement, including ensuring good matches with the institution in the initial selection process, attention to integration and ongoing investments that complement the recruited star.

Chapter 6 Conclusion

Science policy catalyses a country's knowledge base and the growth of local research communities. However, research on these policies has primarily examined their effects on larger economies. Hence, the thesis focuses on the direct and indirect mechanisms of star scientists' impacts in small open economies. In this section, I summarise my findings from each chapter and discuss their potential policy implications.

6.1 Summary and Main Findings

In Chapter 2, I explored how recruiting star scientists affects the quality-adjusted productivity of departments and individual researchers in the universities of Ireland, Denmark and New Zealand. Utilising a dataset of over 1.4 million publications and citations from Scopus, I employed an event-study model which compared the changes in quality-adjusted output before and after the arrival of a star scientist to identify the causal effect of star recruitments. Additionally, I employed the CEM procedure at an individual incumbent unit of analysis to aid the causal interpretation of the observed star arrival effects combined with the strong evidence to support the parallel trend assumption. The analysis controlled for various other factors that could affect research performance, such as the characteristics of the departments, the individual researchers and the star scientists themselves. However, due to some remaining endogeneity concerns in recruiting star scientists, I exercised caution in inferring causality from the findings, particularly at the department level.

I found that star scientist hires have a positive and sustained effect on the research performance of both the receiving department and the incumbent scientists in that department. Interestingly, star scientists who were relatively early in their careers were found to have larger impacts when compared to more established stars. This could be because the early-career stars have longer horizons and less developed existing networks, making them more likely to invest in local networks. The findings of this study have potential implications for theory, practice, and policy and contribute to the literature on star effects, network formation, and knowledge spillovers, which are key areas in the economics of innovation. Furthermore, the findings provide insights for

academic institutions and policymakers who design and implement star recruitment policies, an increasingly prominent strategy in the science-policy mix of small open economies.

In Chapter 3, I investigated the indirect effects of exposure to a star. The chapter explored the impacts of star help identified from the acknowledgements of publications. Specifically, I compared two groups of authors from Ireland, Denmark, and New Zealand: those who acknowledge stars, and those who do not. Using an event-study methodology where the event is the acknowledgement to a star, I evaluated the differences in the quality-adjusted output for authors before and after the star exposure. Importantly, these authors were selected to ensure they had no previous co-author relationship with any identified stars. I further analysed the effects of different help words within acknowledgements, differentiating between conceptual, technical, material, funding & support, and other types of acknowledgements to a star.

My findings revealed that an author's quality-adjusted output increases on average in the year that star help is acknowledged; however, these effects diminish significantly in subsequent years unless the help is sustained. The chapter also included an analysis of the impacts of the different types of help received from a star, and this further supported the hypothesis that star help had a significant impact on the quality-adjusted productivity of the author, with conceptual help from a star showing the most impact in terms of the magnitude of the effect. Interestingly, I also found that the star help, which recognised the donation of material, showed a more sustained effect on an author's productivity. Again, these results were robust to alternative dependent variables. Overall, I argued that the evidence is most consistent with the effect of star help having a genuine impact on the productivity of the acknowledging author rather than simply the result of signalling effects.

Based on the results in this chapter, I proposed some possible strategies for star recruitment programmes: Institutions should facilitate interactions and collaborations between star and non-star scientists. Examples include co-authorship opportunities, mentoring programs, seminars, workshops, and social events. These institutions should support non-star researchers in building networks with external stars by funding conference travel and encouraging participation in cross-departmental or international

projects. Institutions should target their policies towards non-star scientists who need the most help, such as those at the lower end of their field-specific productivity distribution. Importantly, departments should strive to sustain the help provided by stars over time, as this can have a lasting positive effect on the productivity of non-star researchers.

In the preceding chapters, I focused on the impact of star scientists on the incumbent researchers. Chapter 4, however, explored how star scientists' productivity changes when they move into an institution in one of the three small open economies (SOEs) – Ireland, Denmark, and New Zealand. Using the same data from Scopus, the chapter identified 981-star scientists in the top 5% of the most cited researchers in their fields. Among them, I distinguished between two groups: 161 star-movers, who moved to a department in one of the SOEs for the first time and stayed for at least four years, and 75 star non-movers, who remained at the same department in an SOE. However, in this chapter, due to data limitations, I have chosen the control group of non-moving stars based in the three countries. In an ideal scenario, an optimal control group could be the cohort of star scientists identified from the origin country of the international stars. Consistent with the previous chapters, I measured the productivity of the two groups by their field normalised total citations (FNTC), which is the number of citations they received relative to the average citations in their fields. Similar to previous chapters, I employed an event study methodology to compare the FNTC of the two groups before and after a star move. Additionally, I developed a theoretical framework with three mechanisms to explain the possible productivity trajectories of the star movers: matching, disruption, and institution-building mechanisms.

My findings from Chapter 4 suggested that the institution-building mechanism is likely the main explanation behind the observed output trajectory of star movers. This mechanism proposed that star movers prioritise improving institutional quality before focusing on their own productivity. I found that a star movers' FNTC decreased by approximately 58% in the year of the move and remained negative and significant in the following years. This finding aligns with Chapter 2, which showed a significant positive effect of a star's arrival on the output of the department and the incumbent scientists, further supporting the institution-building mechanism. Furthermore, I found that the productivity effects of the star movers vary depending on their career age and subject

area. Star-movers in the early stage of their careers were less affected than those in the late stage. However, both groups experienced a persistent decline in citation-weighted output. When considering the subject area, the productivity pattern of star-movers in Physical, Health and Life Sciences aligned with the institution-building mechanism. Conversely, the pattern for star-movers in the Social Sciences better aligned with the matching mechanism, indicating that job matching could be more important in these fields. The results remained robust to various sensitivity checks and specifications. This chapter contributed to a significant gap in the literature by examining the effects of star movement on star productivity in SOEs. It also provided policy implications in terms of selecting, attracting and retaining star scientists.

The findings from Chapter 2 motivated the analysis conducted in Chapter 5. This chapter explored how the relatedness to star scientists affects the productivity of scientists in the department. Depending on the relatedness, I hypothesised that the arriving star could have substantially different effects on the incumbents at local departments. Employing a difference-in-differences method, Chapter 5 examined the possibility of a non-linear relationship between the level of relatedness of the incumbent to the star and the size of the star's arrival effect on productivity. Additionally, the chapter utilised both symmetric and asymmetric relatedness measures. It allowed for differences in the measure of relatedness depending on whether it was viewed from the perspective of the incumbent or the star. I also explored the robustness of the findings in relation to different measures of the knowledge spaces, such as keyword similarity and the similarity of publication patterns across journal categories.

My findings revealed that the peer effects of star recruitment depend on the level of relatedness between the star and the incumbents. Moreover, this effect was larger when the star and incumbents shared some common knowledge but not to the point of knowledge redundancy. This suggests that the optimal level of relatedness allows for both learning and complementarity between the researchers. This finding has significant implications for designing and implementing star recruitment policies and strategies. It suggests that policymakers and institutions target stars with relevant knowledge related to the existing researchers in the receiving department. This approach would increase the likelihood of active engagement and collaboration between the star and the incumbents, the key mechanisms for knowledge spillovers. However, the chapter also

identified some limitations of the analysis. These include the variation of the effects by the arrival cohort, the difference of the effects across the three countries in the sample, and the omission of other dimensions of proximity, such as geographical, institutional, and social proximity, which may also affect the peer effects of star recruitments. Therefore, the chapter calls for further research to explore these issues and to examine the long-term outcomes of star recruitments, such as the retention, mobility, and career trajectories of the star and the incumbents, and the impact of star recruitments on the innovation performance and competitiveness of the receiving institutions. Finally, Chapter 5 provided evidence-based and context-specific recommendations for policymakers and institutions seeking to leverage star recruitment to enhance scientific excellence and innovation.

6.2 Drawing the Threads Together: The Importance and Design of Star-Recruitment Policies

This thesis provides evidence for research institutions and policymakers to better understand the impacts of optimal designs for star recruitment policies. Overall, the various chapters provide evidence of the effectiveness of star recruitments in improving productivity at the receiving institutions. However, they also highlight various nuances to this broad finding and potential policy lessons for the successful design of star recruitment policies as a means to catalyse local research clusters.

The findings from each of the four chapters suggest that collaboration and adaptability are the key factors in the success of star recruitments. A recurring theme is that policies should promote co-location and collaboration among star and non-star researchers. These include creating physical and virtual spaces facilitating interaction and communication amongst researchers, such as shared offices, labs, meeting rooms, online platforms, etc. The findings from Chapter 2 suggest that co-location and collaboration can increase the likelihood of star help being provided and received and foster a culture of learning and innovation in the department. Also, the indirect peer effects of star scientists observed from the acknowledgement texts suggest that successful collaboration is not limited to co-authorship relationships with a star; implementing formal and informal mentoring programmes by assigning star researchers

as mentors to non-star researchers ensures collaboration even where a formal co-authorship does not exist. Mentoring programmes offer non-star researchers structured and consistent feedback, guidance, and assistance. They also present opportunities for collaboration, paper co-authorship, grant applications, and participation in conferences alongside established researchers. Formal and informal mentoring programmes can also help build trust and rapport among researchers, enhancing the quality and impact of their work.

Another aspect of collaboration is network building and participation in external activities. This can be done by providing funding and incentives for non-star researchers to travel to conferences, workshops, seminars, and other events to meet and interact with star researchers from other institutions. Network building and participation in external activities can expose non-star researchers to new ideas, methods, and trends in their fields and create potential collaborations and partnerships with star researchers. This can also increase the visibility and reputation of the department and its researchers in the academic community.

Importantly, the adaptability of the star scientist is key to the return on investment in small open economies. The negative impacts on star's productivity since the move to the institutions in the small open economies are outlined in Chapter 4. It indicates that to ensure suitable matches between stars and institutions in the initial selection process, policymakers and research institutions could involve both the candidates and the existing scientists in the recruitment process, for example, by conducting interviews, site visits, seminars, workshops and feedback sessions to assess the fit and compatibility of the stars with the institutional culture, vision and goals. Policymakers should establish clear and transparent criteria for identifying and selecting star scientists for mobility programmes, considering their research fields, achievements, and potential impact. Orientation programmes, mentoring schemes, networking opportunities, and access to research facilities and equipment support a more seamless integration of the arriving star. More specifically, to attract and retain stars who arrive early in their careers, policymakers and research institutions could offer them long-term contracts, competitive salaries, research grants, career development opportunities and recognition. To evaluate the effectiveness and impact of star recruitment policies, policymakers and research institutions could collect and analyse data on the outputs,

outcomes and impacts of the stars and their departments and colleagues, such as publications, citations, patents, grants, awards, collaborations, innovations and policy influence.

6.3 Limitations and Future Research

A primary limitation of this study is the availability of the data from Scopus. For example, collected publication data contains the affiliation information of the authors. However, it lacks other professional and personal data such as positions they held at the universities, biological (as opposed to career) age, gender and country of origin. Across all the empirical chapters, the analyses lack detailed information on each star's career trajectory. This information would help verify the star scientist's location details and provide other potential insights for their productivity. In Chapter 3, I use star names in the acknowledgement texts and important keywords to identify acknowledgements to a star and the type of help provided. However, Scopus does not have the acknowledgement texts of all publications. While I employ various robustness checks that reduce the potential bias from this data limitation, the incomplete availability of the acknowledgement information for all publications means that the findings should be interpreted with caution due to potential sample selection biases. Also, the collected data does not include the publication details of the star scientist's peers present in their prior universities from which they are arriving from into the three focal countries. Hence, in Chapter 4, due to this data limitation, I use the cohort of stars who originated from the SOEs as the control group.

As is typically the case with observational studies, endogeneity bias remains a concern across the chapters despite the event-study and matching methodologies employed. For example, in chapters 2 and 5, the decision of the star scientist to move to the affiliation in the small open economy could be affected by knowledge of future productivity potential that is available to the star but not to the econometrician. Likewise, in chapter 3, the star scientists voluntarily decide to help or mentor the scientists, which could be influenced by the star's knowledge of the scientist's potential to productively avail of the help. While the thesis attempts to mitigate endogeneity biases through techniques such as the two-way fixed effects event study model and CEM approach, any

causal interpretations of the results are advanced with caution. For future research, causal interpretation would be further advanced if valid instruments for star arrivals can be utilised. Moreover, additional insight and understanding of star recruitment could be gained by complementing the quantitative approach taken in this thesis with qualitative research, such as case studies examining star arrivals.

Appendix A. Supplementary material for chapter 2

A.1. Stars as Catalysts: A Dynamic Model of Cluster Formation Following the Arrival of a Star

To motivate our dynamic empirical specification we develop in this section a simple matching model of how the arrival of a star could catalyse the development of a cluster of scientific activity in their area. We identify a core mechanism based on the attraction and retention of additional knowledge capital – or what we simply call talent. The main idea in the model is that the arrival of a star with talent above a given threshold sets off a self-reinforcing process through which the arrival of talent attracts further talent to apply to open positions in the department, but the department eventually runs into congestion-type constraints determined by the availability of vacancies.⁴¹

A.1.1. Matching and new talent arrivals

We denote the talent in the department as $N(t)$, which is measured in units of effective productivity that we take to be the flow of quality-adjusted publications. At time 0, the department consists of incumbent talent $N(0)$ (which we normalise to zero without loss of generality) and may or may not recruit a star. The recruitment of a star – who is assumed to have talent N^* – is treated as an exogenous event. We assume there is a threshold level of total department talent $\bar{N} > 0$ beyond which new potential recruits are attracted to apply for open positions. Without the recruitment of a star, the department is unable to solve the coordination problem of attracting new applicants. This coordination problem is solved by the recruitment of the star assuming $N^* \geq \bar{N}$.

We model the talent arrivals at a point in time, $A(t)$, with the following assumptions. The department has a budget B for the recruitment of talent including the

⁴¹ By analogy with similar processes in chemistry, it is natural to view this as an autocatalytic process. The defining feature of an autocatalytic reaction is that one of the outputs of a catalysed reaction is further amounts of the catalyst itself. In our setting the catalyst is talent. Autocatalytic reactions are known to follow a logistic growth process and we use the logistic assumption to model the development of the cluster after the arrival of the star.

star.⁴² The wage rate per unit of talent is w and is assumed to be competitively determined so there is no local bargaining over the wage. Therefore the talent capacity of the department is $C = B/w$. However, frictions in the recruitment process mean that the department cannot immediately move to its capacity following the recruitment of the star. Instead, arrivals are the result of a matching process, with the actual number of arrivals (i.e. successful matches) depending positively on both the size of the pool of applicants, $P(t)$, that apply for an open position in department and also the number of vacancies (i.e. unfilled capacity, $C - N(t)$). The pool of applicants can come from new arrivals or scientists separating from other departments. For simplicity, we assume the total number of scientists in the market is fixed, so that the number of new entrants is equal to the number of retirements.

We assume that our focal department is small in relation to the overall talent market and that the rest of the market is in steady state with arrivals balancing leavers. Our smallness assumption is interpreted as meaning that our department has a negligible impact on the market allowing us to ignore any general equilibrium effects. We characterise the determination of the market wage in the steady state in Section A.1.4 below.

Potential recruits are attracted to apply to the department based on the talent already present. Thus a preferential attraction mechanism – the source of increasing returns – operates whereby a department is more attractive the more talent that has already been recruited (including the talent of the star). We capture this mechanism with the simple assumption:

$$P(t) = \phi N(t), \tag{A.1}$$

where ϕ captures the attractive power of the already present talent. The matching function is assumed to take the Cobb-Douglas form:

⁴² We assume that this budget is fixed although this assumption could be relaxed to allow for an exogenously changing budget over time or for an endogenously changing budget that depends on the talent already recruited.

$$\begin{aligned}
A(t) &= \mu P(t)^\alpha (C - N(t))^\nu = \mu \phi^\alpha N(t)^\alpha (C - N(t))^\nu \\
&= a N(t)^\alpha (C - N(t))^\nu.
\end{aligned}
\tag{A.2}$$

The portmanteau attraction parameter, a , thus depends on the efficiency of the matching process, μ , the attractive pull of the star and their local co-located network in generating an applicant pool, ϕ , and the effect of the size of the applicant pool on the actual recruitment of talent, α . The parameter ν measures the effect of vacancies on the actual recruitment. Ignoring separations for the moment, so that $A(t) = dN(t)/dt$, the differential equation given by (A.2) does not have an analytical solution for general choices of the parameters. However, a solution does exist for certain combinations of the α and ν parameters (Blumberg et al., 1959). One particularly tractable case is where $\alpha = \nu = 1$. In this case, conditional on the arrival of a star, the arrival function (A.2) takes the logistic form:

$$A(t) = aN(t)(C - N(t)) \tag{A.3}$$

We assume in what follows that these conditions on the matching function hold.

A.1.2. Leavers

We also allow for talent to leave (or separate from) the department over time. Denoting the number who leave at a point in time as $L(t)$, the total talent stock evolves as:

$$\frac{dN(t)}{dt} = A(t) - L(t) \tag{A.4}$$

We term the *fraction* of talent that leaves at an instant in time the separation rate and denote it as $l(t)$. For a given $l(t)$, the number of leavers at an instant in time is then:

$$L(t) = l(t)N(t). \tag{A.5}$$

However, we also allow for the stock of talent in a cluster (including the initially arriving star) to help retain that talent, which we capture by assuming the separation rate depends negatively on the talent stock according to the linear relationship:

$$l(t) = l_0 - l_1 N(t). \quad (\text{A.6})$$

Substituting (A.6) into (A.5) and rearranging, we can verify that the number of leavers follows:

$$L(t) = l_1 N(t) \left(\frac{l_0}{l_1} - N(t) \right). \quad (\text{A.7})$$

A.1.3. Dynamic evolution of the department following the arrival of a star

Now substituting (A.3) and (A.7) into (A.4), the post-star arrival talent stock evolves according to the logistic differential equation:

$$\frac{dN(t)}{dt} = A(t) - L(t) = (a - l_1)N(t) \left(\frac{aC - l_0}{a - l_1} - N(t) \right), \quad (\text{A.8})$$

where we assume $aC > l_0$ and $a > l_1$. Figure A.1 shows the relationship between both arrivals and the stock of talent and leavers and the stock of talent for $N(t) \geq N^*$. The cluster is in steady state when:

$$\frac{dN(t)}{dt} = (a - l_1)N(t) \left(\frac{aC - l_0}{a - l_1} - N(t) \right) = 0 \quad (\text{A.9})$$

Solving the quadratic equation given by the second equality, the non-zero root gives the steady-state talent stock:

$$N_{ss} = \frac{aC - l_0}{a - l_1} \quad (\text{A.10})$$

When $C \leq l_0/l_1$, the steady state will be at less than or equal to capacity. We assume this condition holds. An interesting special case occurs where $C = l_0/l_1$. In this case, the talent stock will converge towards its capacity level, C .

Solving the differential equation (A.8), and imposing the initial condition, $N(0) = N^*$, yields the logistic functional form for the time path of the talent stock:

$$N(t) = \frac{N_{ss}}{1 + \left(\frac{N_{ss} - N^*}{N^*}\right) e^{-(ac-l_0)t}}. \quad (\text{A.11})$$

The shape of the time path is illustrated in Figure A.2. Following the arrival of the star the catalytic effects are initiated leading the talent stock at the department to increase initially at an increasing rate. Eventually, however, the inhibiting influence of available vacancies comes to dominate, causing the rate of increase in the talent stock to slow down and eventually converge towards the steady state. In terms of individual productivity, we simply assume that individual output is an increasing function of the local talent stock through various local network effects (e.g. access to knowledge or co-authorship networks).

A.1.4. General equilibrium in the talent market

We finally consider the general equilibrium in the market comprising of K departments with positive talent indexed by i . All departments are assumed to be at their steady-state talent levels given by equation (A.10). We assume that the total quantity of talent is fixed at N^T . In the equilibrium, there are no unemployed scientists but there will be vacancies if $C_i = B_i/w < l_0/l_1$. The competitive wage adjusts to balance the total quantity of talent demanded by departments with the total quantity supplied:

$$\sum_{i=1}^K \frac{a \frac{B_i}{w} - l_0}{a - l_1} = N^T \quad (\text{A.12})$$

$$\Rightarrow w = \frac{a}{Kl_0 + N^T(a - l_1)} \sum_{i=1}^K B_{i..}$$

The equilibrium wage is positive (recalling the assumption that $a - l_1 > 0$) and depends positively on the aggregate budget for talent recruitment across departments and negatively on the total talent stock.

Although stark in its features, the model shows how a star arrival could catalyse the development of a department. Notwithstanding a budget for expansion, our focal department is initially unable to expand due to a coordination problem in attracting talent. The arrival of the star can solve the coordination problem and lead to a period of department growth until it reaches its steady state. The numbers of arrivals and leavers

in effective talent units follow inverted U-shaped paths to the steady state. Although we ignored incumbents in the department at the time of the star arrival for simplicity, another important effect is likely to be the impact on incumbent productivity of the expansion in the local talent stock. We next turn to test the predictions of the model for overall department output, incumbent productivity, arrivals and leavers using an event-study framework where event is the arrival of a star.




Figure A.1. Arrival and leaver rates as a function of talent stock




Figure A.2. Evolution of talent stock following arrival of the star

A.2. Descriptive Statistics for Star Arrival

Table A.1 Country-level composition of department-year units 1996-2017

	Total			Treated (1997-2017)		
	<i>N</i> (1)	Departments (2)	Universities (3)	<i>N</i> (4)	Departments (5)	Universities (6)
IRELAND	3,102	141	8	418	19 (13%)	5
DENMARK	3,366	153	12	836	38 (25%)	6
NEW ZEALAND	3,586	163	8	528	24 (14%)	8
TOTAL	10,054	457	28	1,882	81 (17%)	19

Notes: Columns 1-3 reports the total department-year units, total departments and total Universities in the analysis. Columns 4-6 reports the total treated department-year units, total treated departments and total treated Universities in the analysis where treatment is a first star arrival. All first star arrivals in 1996 are omitted since they are 'always treated' throughout the observation window.

Table A.2 Characteristics of the treatment group (1996-2017)

Departments (Deps)	No. of Deps with first star arrivals (Treated)	No. of subsequent star arrivals	Total outputs of first star arrivals (publication count)	No. of first stars arriving at the same time to the same Department	Same first star arrival in multiple Deps	If multiple deps, what are the other Deps?
	(1)	(2)	(3)	(4)	(5)	(6)
AGRI	8 (38%)	2	98	0	-	-
ARTS	1 (6%)	0	2	0	-	-
BIOC	7 (30%)	16	116	2	1	MEDI
BUSI	7 (35%)	5	62	0	-	-
CENG	1 (6%)	0	7	0	1	CHEM
CHEM	4 (17%)	2	76	0	1	CENG
COMP	7 (33%)	3	95	0	3	ENGI, SOCI
DECI	0 (0%)	0	0	0	-	-
DENT	0 (0%)	0	0	0	-	-
EART	6 (38%)	1	196	0	1	PHYS
ECON	4 (27%)	4	41	0	-	-
ENER	0 (0%)	0	0	0	-	-
ENGI	5 (27%)	0	48	0	2	COMP
ENVI	3 (14%)	0	81	0	-	-
HEAL	0 (0%)	0	0	0	-	-
IMMU	1 (6%)	0	3	0	-	-
MATE	2 (10%)	2	12	0	-	-
MATH	2 (10%)	0	4	0	-	-
MEDI	9 (40%)	30	218	2	1	BIOC
MULT	1 (11%)	0	2	0	-	-
NEUR	3 (20%)	1	35	0	-	-
NURS	0 (0%)	0	0	0	-	-
PHAR	0 (0%)	0	0	0	-	-
PHYS	6 (30%)	8	31	2	1	EART
PSYC	1 (8%)	0	2	0	-	-
SOCI	9 (43%)	2	94	0	1	COMP
VETE	1 (11%)	0	3	0	-	-
TOTAL	88 (19%)	76	1,226	6 (across 3 departments)	12	-

Notes: All numbers here include the 1996 cohort. There are 7 departments in the 1996 first star arrival cohort, 1 in Ireland, 3 in Denmark, and 3 in New Zealand.

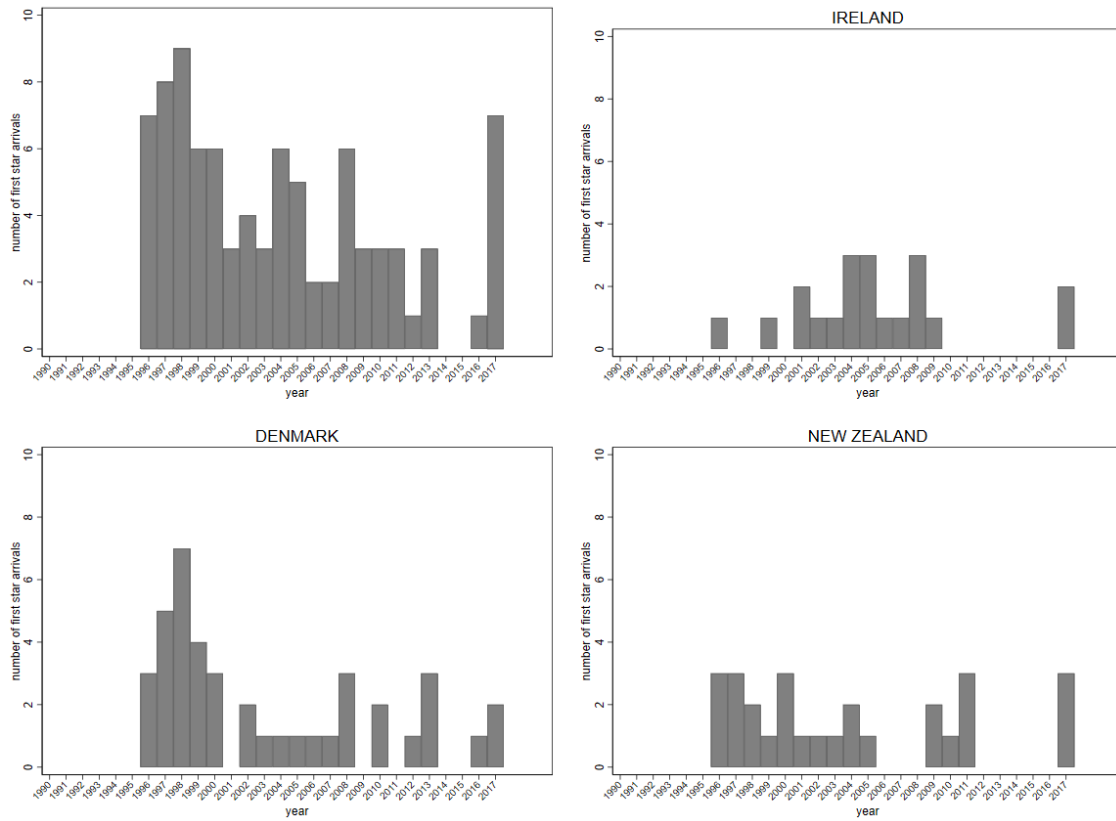


Figure A.3: Number of departments with first star arrival by country (1996-2017)

Notes: There are 88 departments with first star arrival; 20 departments in Ireland, 41 departments in Denmark, and 27 departments in New Zealand.

A.3. Robustness: Career Age




Figure A.4: Event study model based on split between earlier and later career age star arrivals

Notes: See notes to Figure 2.1. Estimation is by TWFE with controls for department specific time trends and university-level output.




Figure A.5: Event study model with raw output measure

Notes: This figure plots the dynamic effect of a first star arrival with 95% confidence intervals based on the matched sample, with 1,062 (23,364 observations for career age ≥ 20) and 3,437 matched authors (75,614 observation for career age < 20) in Panel A, with 1,891 (36,795 observations for career age ≥ 20) and 14,460 (223,834 observations for career age < 20) matched authors in Panel B. Career age is identified since the first publication. The threshold of career age 20 is set at $t-1$, i.e., one year before the star arrival. The event and observation window is from 1996-2017. The dependent variable is the field-normalised total citations in logarithm (with 1 added across all observations). The omitted category is one year before the first star arrival. Robust standard errors are clustered by author.

A.4. Robustness: Country Specific Results

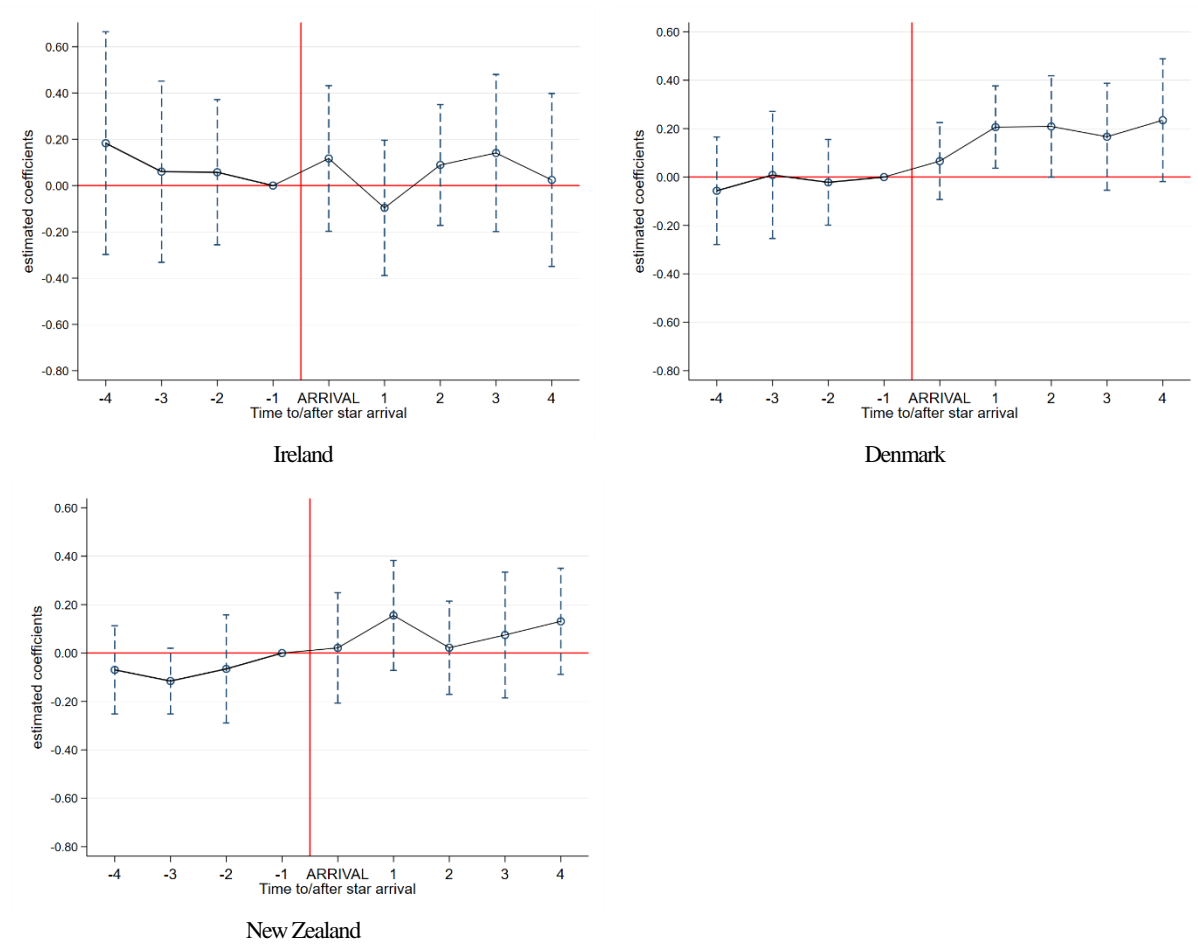


Figure A.6 Event study model with homogenous arrival effects across cohorts by country (Ireland, Denmark, and New Zealand)

Notes: This figure plots the dynamic effect of a first star arrival with 95% confidence intervals and assuming homogenous arrival effects across cohorts based on the specification in equation 2.8 for each country separately. The event and observation window is from 1996-2017 (the 1996 cohort is omitted). The dependent variable is the field-normalised total citations and excludes the output of the star arrival. First star arrivals are restricted to stars staying in the department for four years or more. The omitted category is one year before the first star arrival. University-level controls (excluding the focal department) and department time-specific trend are included in the model specification. There are 3,100 observations in Ireland, where 418 (i.e., 19 departments with star arrivals) are treated; 3,366 observations in Denmark, where 836 (i.e., 38 departments with star arrivals) are treated; 3,586 observations in New Zealand, where 528 (i.e., 24 departments with star arrivals) are treated. See Table A.1 for more details.

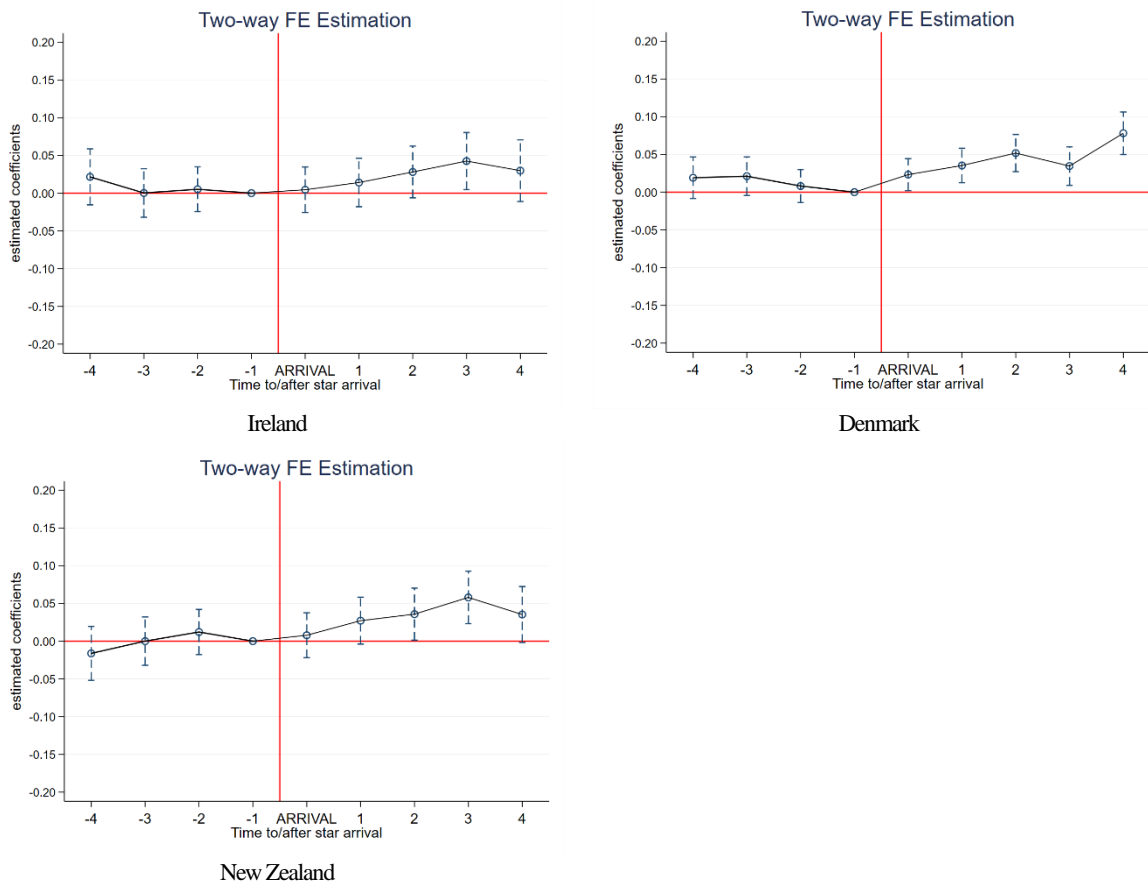


Figure A.7: Event study model for matched incumbents by country.

Notes: This figure plots the dynamic effect of a first star arrival with 95% confidence intervals based on the matched sample. The dependent variable is the field-normalised total citations in logarithm (with 1 added across all observations). The omitted category is one year before the first star arrival. Robust standard errors are clustered by author. There are 61,108 (3,766 authors), 126,127 (8,105 authors), and 73,394 (1,442 authors) observations based on the unbalanced panel for Ireland, Denmark, and New Zealand, respectively. The 1996 treated cohort is removed. Robust standard errors are clustered by author

Table A.3 Dynamic star arrival effects (restricted specification). Event window: 1996-2017. Observation window: 2001-2013

	Field normalised (1)	Field normalised (2)	Field normalised (3)	Field normalised (4)
$S_{i,-4t_0-1} \times (j + 1)$ [Pre-arrival trend]	0.0233 (0.0262)	0.0105 (0.0276)	0.00158 (0.0254)	0.00387 (0.0281)
$S_{i,0t_04}$	0.0928** (0.0399)	0.0951** (0.0421)	0.0894** (0.0392)	0.0974** (0.0423)
$S_{i,0t_04} \times j$ [Post-arrival trend]	0.0146 (0.0135)	0.0260* (0.0152)	0.0131 (0.0124)	0.0282* (0.0144)
Log University Output (excl. dept.)			0.978*** (0.0620)	0.857*** (0.0701)
R-squared	0.434	0.635	0.528	0.656
Department FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Department time trend	No	Yes	No	Yes
Observations	10,040	10,040	10,040	10,040

Notes: This table reports the estimates based on the restricted model specification in equation A.13. The dependent variable always excludes the star's own output. Robust standard errors are clustered at department-level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively.

A.5. A Restricted Event-Study Specification

Our baseline event-study specification places no restrictions on the time pattern of lead and lag effects other than through the binning assumptions. However, if the lead and lag effects are roughly constant or follow linear time trends, then the efficiency of our estimates can be improved by imposing appropriate restrictions on equation (2.8) (see, e.g., Deryugina, 2017). We therefore also explore a restricted specification for the homogenous that imposes a piecewise linear structure on the dynamic effects. More specifically, we assume a linear pre-trend in effects, allow a discontinuous jump in the star arrival effect in the year of arrival and also assume that the post-trend in effects is linear. In addition to the increased efficiency (assuming the restrictions are consistent with the data), this approach offers the possibility of explicitly modelling the pre-trend and extrapolating it to the post-arrival period (see Figure A.8).⁴³ We thus present our results for the restricted model in a way that allows for the post-arrival dynamics to be visualised both with and without an extrapolated pre-trend. To allow comparison with our baseline results, we again normalise the coefficient on the first lead ($j = -1$) to zero. Our usual binning variables are included so that the specification can be viewed as restricting the pattern of leads and lags in the 9-year window around the star arrival. The

⁴³ See Dobkin et al. (2018) for an alternative approach to imposing and extrapolating a linear pre-trend.

restricted specification (where the star arrival effect is again normalised to zero for the first lead) is then:

$$\begin{aligned} \ln Y_{i,t} = & \alpha^* + \beta_{-5}^*(1 - S_{i,t+4}) + \theta_1(S_{i,-4t0-1} \times (j + 1)) + \theta_2 S_{i,0t04} \\ & + \theta_3(S_{i,0t04} \times j) \\ & + \beta_5^* S_{i,t-5} + X'_{i,t} \lambda + \phi d_i + \varphi d_t + u_{i,t}. \end{aligned} \quad (\text{A.13})$$

where $S_{i,-4t0-1}$ and $S_{i,0t04}$ are dummy variables indicating a star arrival in the time span of four years before and an arrival and indicating a star arrival in the time span of the current year to four years after respectively. As this specification allows for a discontinuous jump in the linear trend in the year of arrival and a change in the slope of the trend post arrival, the estimated star arrival effect j years post arrival is then: $\theta_2 + (\theta_3 \times j)$. As noted, the specification also allows for the extrapolation of any estimated pre-trend forward into the post-arrival period. The effect (after accounting for this extrapolation) of a star arrival j years post arrival is: $\theta_2 - (\theta_1 \times (j + 1)) + (\theta_3 \times j)$. We report results of this specification with and without controls.

Table A.3 records the results of our restricted regression where we again present the results with and without controls. We focus on Column 4 – the field-normalised total citations measure of output with department-specific time trends and the university-level control included in the regression. We again do not find evidence that a pre-trend is a significant concern. The overall impact effect of a star arrival is 9.74 log points (which is statistically significant at the 5 percent level). However, we do find evidence of a statistically significant *post*trend at the 5 percent level, with a positive coefficient of 2.82 log points. Overall, the estimated impact of a star arrival after four years is $0.0974 + (0.0282 \times 4) = 0.2102$. As shown in Figure A.8, extrapolating the pre-trend into the post-trend period leads to an effect after four years of $0.0974 - (0.00387 \times (4 + 1)) + (0.0282 \times 4) = 0.19085$.




Figure A.8: Estimates of the star arrival effect from the restricted model (see Table A.3)




Figure A.9: Comparison of autoregressive-distributed lag (ADL) with the baseline specification.

Notes: This figure plots the dynamic effect of a first star arrival with 95% confidence intervals and assuming homogenous arrival effects across cohorts based on the specification in Eq. (C.1) for the left panel and Eq. (8) for the right panel. The event and observation window is from 1996-2017 (the 1996 cohort is omitted). The dependent variable is the field-normalised total citations and excludes the output of the star arrival. First star arrivals are restricted to stars staying in the department for four years or more. The omitted category is one year before the first star arrival. University-level controls (excluding the focal department) and department time-specific trend are included in both model specifications.

A.6. Robustness to Miss-Specified Dynamics

Our primary interest is the estimation of the dynamic effects of star arrivals. As is well known, serial correlation in the error term can indicate misspecification of the dynamics. Consider, for example, the case where random positive shocks to a department's output are correlated with subsequent star arrivals. Serial correlation in these shocks means that the contemporaneous error term will be correlated with the star arrival dummy in the arrival year, and indeed correlated with lags of star arrival dummy as the effects of these shocks persist, leading to biased estimates of the dynamic effects. As a robustness check on our dynamic estimates, we thus relax the assumption of no serial correlation and instead assume that the error term follows the AR(1) process: $u_{i,t} = \rho u_{i,t-1} + \varepsilon_{i,t}$, where ε_t is assumed to be serially uncorrelated and homoscedastic and $|\rho| < 1$.

Lagging (2.8) by one period and multiplying by ρ , and then subtracting the result from (2.8), we obtain, after some rearrangement of the resulting quasi-differenced equation, an amended event-study specification that takes the autoregressive-distributed lead/lag (ADL) form:

$$\begin{aligned}
 \ln Y_{i,t} = & (1 - \rho)\alpha + \rho\varphi + \rho \ln Y_{i,t-1} + \beta_{-5}(1 - S_{i,t+4}) \\
 & - \rho\beta_{-5}(1 - S_{i,t+3}) + \beta_{-4}\Delta S_{i,t+4} \\
 & + \sum_{j=-3}^{-2} (\beta_j - \rho\beta_{j-1})\Delta S_{i,t-j} - \rho\beta_{-2}\Delta S_{i,t+1} + \beta_0\Delta S_{i,t} \\
 & + \sum_{j=1}^4 (\beta_j - \rho\beta_{j-1})\Delta S_{i,t-j} + \beta_5 S_{i,t-5} - \rho\beta_5 S_{t-6} + \lambda X'_{i,t} \\
 & - \rho\lambda X'_{i,t-1} + ((1 - \rho)\phi)d_i + (1 - \rho)\varphi d_t + \varepsilon_{i,t}.
 \end{aligned} \tag{A.14}$$

The combined presence of the lagged dependent variable and the department fixed effects will cause the lagged dependent variable to be correlated with the disturbance term (or "Nickell bias", Nickell (1981)). As shown by Arellano and Bond (1991), consistent estimates of the parameters in (A.14) can be found by applying a General Method of Moments estimator to the equation in first differences under the

assumption of no serial correlation in the transformed regression.⁴⁴ In the first-difference equation used to estimate the ADL model, the instruments for the lagged dependent variable include the lags of dependent variable in levels dated $t-2$ and earlier and the contemporaneous first differences in the levels equation, and all right-hand side variables also serve as the instruments in the transformed equation and untransformed for the levels equation (Bond, 2002). The Hansen j statistic (p -value 0.155) also indicates that the instruments are valid. The estimated value on the (instrumented) lagged dependent variable, $\hat{\rho}$, is -0.038 (s.e. = 0.023), which is significant at the 10 percent level.⁴⁵

We present the results from the application of a GMM estimator graphically in Figure A.9. The relevant dynamic (i.e. beta) coefficients shown in the figure are obtained recursively from the GMM results using equation (A.14) and standard errors are calculated using the delta method. As our robustness test, we compare the implied beta coefficients from the ADL model (left panel) to the baseline model (right panel). The ADL model nests our baseline model where the latter imposes that restriction that $\rho = 0$. Although the pattern of dynamic effects is again broadly similar, the estimated effects are somewhat lower in the ADL model. For example, four years post arrival, the estimated effect of a star arrival is 12.5 [$\hat{\beta}_4 - \hat{\rho}\hat{\beta}_3 = 0.130 - (0.038 \times 0.116)$] log points in the ADL model compared to 18.0 log points in the baseline.

⁴⁴ As noted, first differencing to remove the fixed effects introduces correlation between the lagged dependent variables and the disturbance leading to inconsistent estimates. The GMM uses earlier lagged values of the dependent variable as instruments for the endogenous lagged first difference of the dependent variable, allowing for consistent estimates under the maintained assumptions. The absence of serial correlation can be tested post-estimation to increase confidence that sufficient lags of the dependent variable have been included to ensure a complete dynamic specification.

⁴⁵ The estimator is implemented in Stata using the `xtabond2` procedure (see Roodman, 2009).

Appendix B. Supplementary material for chapter 3

B.1. Yearly Distribution of Acknowledgment Texts

Table B.1 Yearly distribution of publications with acknowledgement texts and percentage of publications in which star names are acknowledged

year	Percentage of publications with acknowledgement texts	Percentage of stars acknowledged publications	Total
1990	5.58%	1.57%	10,284
1991	6.22%	2.58%	11,205
1992	5.75%	1.15%	12,051
1993	5.72%	2.20%	13,480
1994	5.43%	1.47%	15,044
1995	6.47%	2.34%	15,838
1996	6.29%	1.23%	18,053
1997	8.24%	1.85%	19,640
1998	9.31%	1.74%	20,353
1999	9.06%	2.10%	21,076
2000	8.75%	1.02%	22,512
2001	8.63%	1.60%	23,893
2002	8.75%	1.92%	25,018
2003	9.37%	1.03%	28,116
2004	10.29%	1.23%	30,702
2005	9.50%	1.27%	34,805
2006	9.92%	1.01%	36,981
2007	9.84%	0.95%	39,552
2008	10.28%	1.09%	40,251
2009	10.32%	1.27%	42,851
2010	10.69%	0.90%	45,589
2011	12.41%	0.71%	47,952
2012	13.42%	0.71%	49,451
2013	13.22%	0.50%	51,343
2014	13.47%	0.56%	52,660
2015	11.87%	0.59%	53,070
2016	18.81%	0.65%	54,015
2017	24.09%	0.60%	53,694
Total		971	889,479

Source: Scopus Database

B.2. Procedure for Identifying Star Names and the Type of Star Help from the Acknowledgement Texts.

Step 1: Data Collection First, we collect the acknowledgement texts from the Scopus database for all the publications in Ireland, Denmark and New Zealand from 1990 to 2017. The unique identifier for these acknowledgements is the publication ID (EID in Scopus).

Step 2: Data Organisation Once we have collected the data, we organise it in a structured format. This involves creating a dataset containing the publication ID, the acknowledgement text related to that publication, and the authors of these publications.

Step 3: Name Identification In this step, we employ Natural Language Processing (NLP) techniques using the Spacy module available in the Python programming language.

1. To start, the acknowledgement text is broken down into individual units called tokens (Tokenisation) using Spacy; these are typically words or phrases.

2. Then, using Named Entity Recognition (NER), we identify, classify and store the named entities such as PERSON, ORGANISATION, PLACES, etc.

3. Finally, we filter the entities tagged PERSON. Then, using the publication ID as the reference to these names, we check for any wrong identification of entities due to structural deformities in the collected acknowledgement text from the Scopus (such as commas, spaces, hyphens, and non-English names).

Step 4: Star Name Matching After extracting person names, we use fuzzy similarity to identify the exact names of scientists. Fuzzy similarity is a technique used in computing and is based on fuzzy logic. It involves finding strings that are approximately equal to a given pattern. We match the names identified from the acknowledgement texts with those of the scientists we have previously identified as stars.

Step 5: Verification This step involves manually verifying that the star names identified are the same as those in the acknowledgements. This is important to ensure the accuracy of the results.

Step 6: Identifying help keywords. In this step, we manually check the help keywords associated with verified star names and store them. These keywords are help words or indications of star-help to the author. Furthermore, we use these help words to identify the type of helpful interaction from a star (Conceptual, Technical, Material, Funds & support and Other types), as outlined in section 3.3. Table B.2 shows the help words used to identify each type of acknowledgement types that defines the type of helpful interaction.

Step 7: Final Data Compilation Finally, we compile the data, which includes the star name, publication ID, author details and help keywords associated with the star scientists. This final dataset is then merged with the panel data of all authors, and dummy variables are used to tag those authors and publications that acknowledge star help.

Table B.2 Identification of help words used to describe the type of helpful interaction from a star scientist

Classification	Definition	Help words
Conceptual	Primarily thank star scientists for intellectual feedback, critique, and encouragement.	Discussion, Suggestions, Reviews, Comments, Critiques.
Technical	Acknowledging star scientists for their help in the technical side of the study.	Excellent technical assistance, Laboratory assistance, Lab assistance, Mathematical help, Statistical advice.
Materials	Mentioning the star scientists for the help with access to the materials which is required to carry out meaningful research.	Supply and use of antibodies, protein cells, antibodies, Access to unpublished data.
Funds & Support	To acknowledge the star scientists for the funds and support they receive throughout the time of their research.	Funds, Grants, Financial help, Support.
Other Types	The category to cover all the types of acknowledgments where we cannot infer or identify the proper type of helpful interaction from the star.	Contributions, Dedications, Founder of the study.

B.3. Coarsened Exact Matching (CEM) Stats for each Type. of Star Interactions

Table B.3 Summary statistics: control and treated group (k-to-k matched) for an unbalanced panel for the types of helpful interactions

Variable	Control	Treated	Diff in mean	P-value
<i>(a) Conceptual: Unbalanced panel of 468 matched authors, with 234 in each group</i>				
Year	2006.282	2006.278	0.004	0.994
Subject	9.145	9.145	0.000	1.000
Country	2.068	2.068	0.000	1.000
Total Career Age	20.389	20.389	0.043	0.963
Total Career Age Bins	5.697	5.697	0.000	1.000
Cumulative Publication experience	8.175	8.124	0.051	0.932
Cumulative Publication experience Bins	2.171	2.171	0.000	1.000
Cumulative citations received per cumulative publications	56.448	57.712	-1.264	0.960
Cumulative citations received per cumulative publications Bins	2.154	2.154	0.000	1.000
<i>(b) Technical: Unbalanced panel of 300 matched authors, with 150 in each group</i>				
Year	2006.907	2006.880	0.027	0.967
Subject	10.880	10.880	0.000	1.000
Country	1.987	1.987	0.000	1.000
Total Career Age	20.193	20.547	-0.353	0.763
Total Career Age Bins	5.687	5.687	0.000	1.000
Cumulative Publication experience	9.193	9.113	0.080	0.912
Cumulative Publication experience Bins	2.420	2.420	0.000	1.000
Cumulative citations received per cumulative publications	32.358	32.660	-0.302	0.902
Cumulative citations received per cumulative publications Bins	1.813	1.813	0.000	1.000
<i>(c) Materials: Unbalanced panel of 176 matched authors, with 88 in each group</i>				
Year	2006.818	2006.795	0.023	0.980
Subject	11.148	11.148	0.000	1.000
Country	1.955	1.955	0.000	1.000
Total Career Age	21.989	22.068	-0.080	0.961
Total Career Age Bins	6.023	6.023	0.000	1.000
Cumulative Publication experience	9.966	9.727	0.239	0.808
Cumulative Publication experience Bins	2.523	2.523	0.000	1.000
Cumulative citations received per cumulative publications	43.074	44.368	-1.294	0.809
Cumulative citations received per cumulative publications Bins	2.250	2.250	0.000	1.000

Notes: Reports the t-test for the mean difference between control and treated groups one year before forming the star-help relation.

Variable	Control	Treated	Diff in mean	P-value
<i>(d) Funds & Support: Unbalanced panel of 60 matched authors, with 30 in each group</i>				
Year	2010.300	2010.233	0.067	0.964
Subject	11.967	11.967	0.000	1.000
Country	2.000	2.000	0.000	1.000
Total Career Age	18.633	19.267	-0.633	0.798
Total Career Age Bins	5.467	5.467	0.000	1.000
Cumulative Publication experience	10.367	10.367	-0.167	0.928
Cumulative Publication experience Bins	2.633	2.633	0.000	1.000
Cumulative citations received per cumulative publications	40.707	42.375	-1.668	0.864
Cumulative citations received per cumulative publications Bins	2.100	2.100	0.000	1.000
<i>(e) Other Types: Unbalanced panel of 248 matched authors, with 124 in each group</i>				
Year	2008.234	2008.226	0.008	0.991
Subject	13.694	13.694	0.000	1.000
Country	2.008	2.008	0.000	1.000
Total Career Age	20.548	20.871	-0.323	0.811
Total Career Age Bins	7.738	7.738	0.000	1.000
Cumulative Publication experience	9.790	9.895	-0.105	0.903
Cumulative Publication experience Bins	2.500	2.500	0.000	1.000
Cumulative citations received per cumulative publications	37.064	36.567	0.497	0.895
Cumulative citations received per cumulative publications Bins	1.935	1.935	0.000	1.000

Notes: Reports the t-test for the mean difference between control and treated groups one year before forming the star-help relation.

B.4. Effects of Star Help through Different Channels (Raw publications-Hypothesis 2)

Table B.4 Dynamic star help effects for normalised raw publications (types of helpful interactions)

	Conceptual (1)	Technical (2)	Materials (3)	Funds & Support (4)	Other Types (5)
<i>Staracknwtype_{i,t-3}</i>	-0.00830 (0.0297)	-0.00461 (0.0388)	0.0634 (0.0521)	-0.128 (0.0946)	-0.0412 (0.0465)
<i>Staracknwtype_{i,t-2}</i>	-0.0408 (0.0295)	-0.0000 (0.0386)	0.0305 (0.0504)	-0.0843 (0.110)	-0.0543 (0.0344)
<i>Staracknwtype_{i,t}</i>	0.219*** (0.0264)	0.188*** (0.0301)	0.220*** (0.0437)	0.193** (0.0829)	0.173*** (0.0352)
<i>Staracknwtype_{i,t+1}</i>	0.0366 (0.0295)	0.0290 (0.0360)	0.0818 (0.0539)	-0.00607 (0.113)	-0.0110 (0.0484)
<i>Staracknwtype_{i,t+2}</i>	0.0103 (0.0325)	-0.0153 (0.0384)	0.120** (0.0570)	-0.00258 (0.0965)	-0.00547 (0.0531)
<i>Staracknwtype_{i,t+3}</i>	0.0127 (0.0328)	0.0165 (0.0432)	0.133** (0.0602)	-0.0621 (0.103)	-0.0534 (0.0496)
Constant	0.404*** (0.0311)	0.473*** (0.0333)	0.378*** (0.0472)	0.428*** (0.0738)	0.420*** (0.0459)
R-squared	0.051	0.051	0.071	0.095	0.057
Observations	7,665	4,942	3,001	947	4,152
Number of authors	468	300	176	60	248
Author FE	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES

Estimation Window: 1996-2017.

The dependent variable is Raw publications.

Note: The table reports the estimates based on the model specification in equation 3.4. Also, the 1996 cohort (always treated) is not dropped in this case. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively.

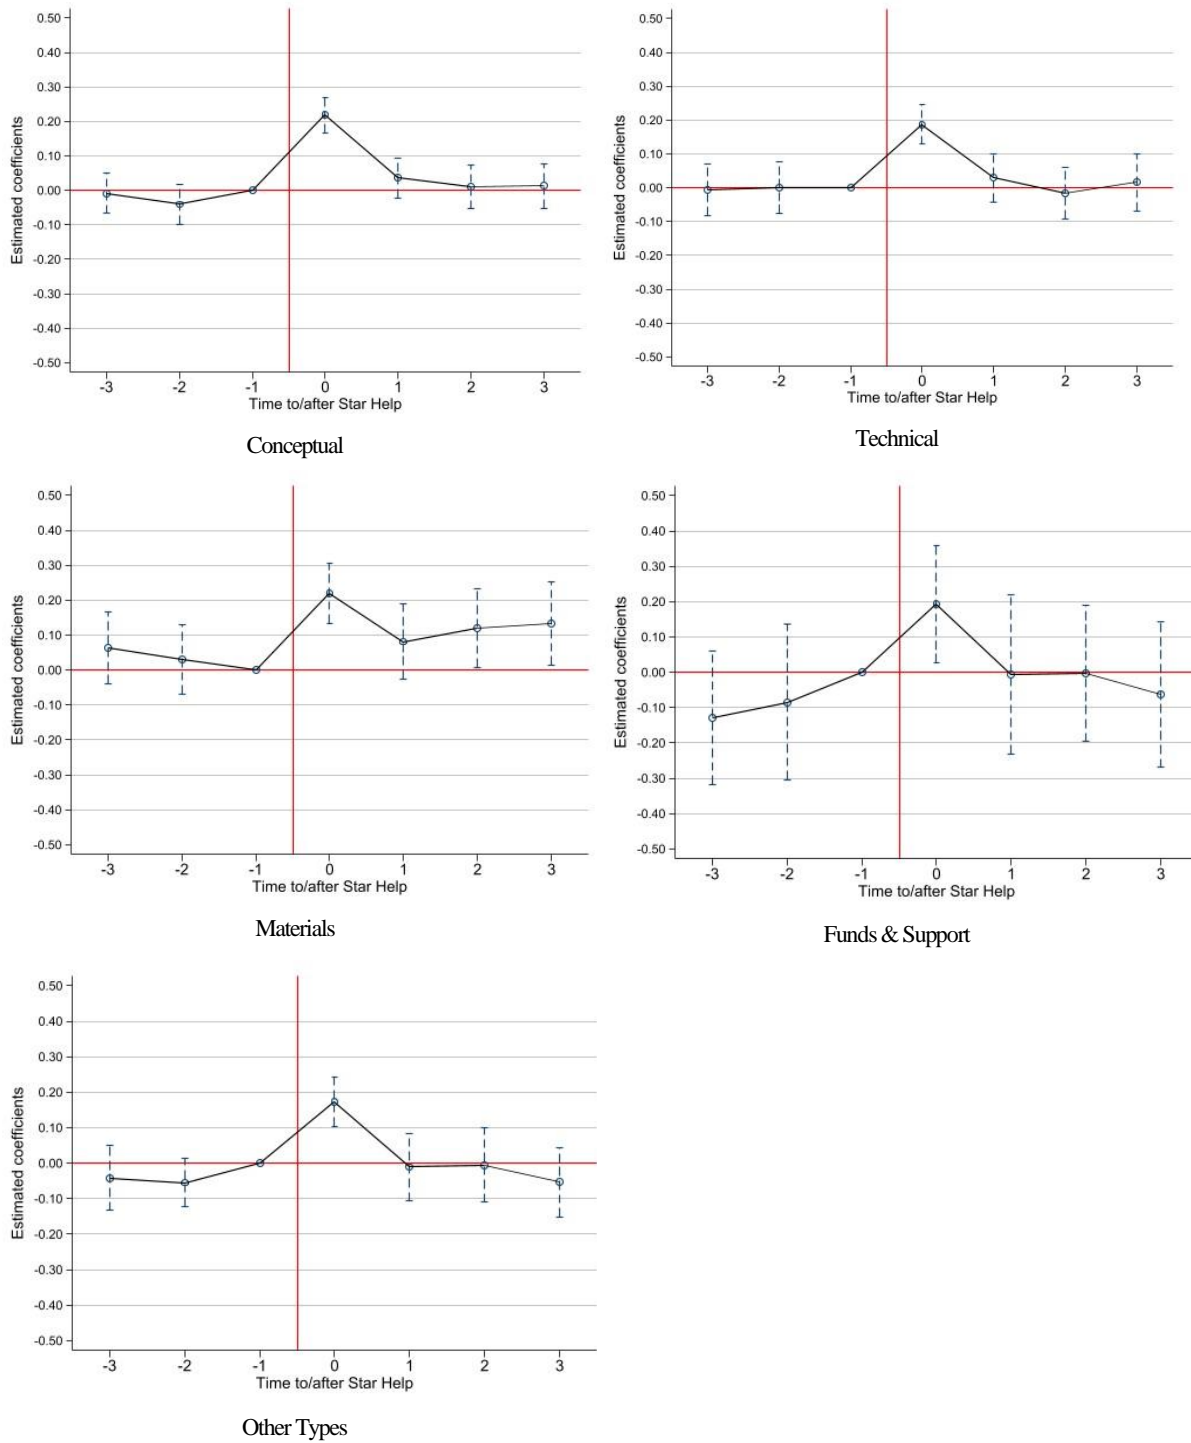


Figure B.1 Event study model with channels of help identified from the acknowledgement texts.

Note: The figure plots the dynamic effect of the star's helpful interactions with an author at 95% confidence intervals. The event and observation window is from 1996-2017. The dependent variable is raw publications.

B.5. Robustness to Co-authorship Relations

Table B.5 Robustness test of dynamic star help effects for FNTC (excluding the authors with co-authored publications with a star after initial contact)

	Indirect Relationship with a star before the event (1)	One-Time Star Help Effect: Excluding the publications (2)	Multiple Star Help Effect: Excluding the publications (3)
<i>Staracknw</i> _{<i>i,t-3</i>}	0.0267 (0.0303)	0.0303 (0.0313)	-0.0550 (0.105)
<i>Staracknw</i> _{<i>i,t-2</i>}	-0.0144 (0.0288)	-0.0128 (0.0297)	-0.0383 (0.110)
<i>Staracknw</i> _{<i>i,t</i>}	0.245*** (0.0285)	0.233*** (0.0295)	0.384*** (0.101)
<i>Staracknw</i> _{<i>i,t+1</i>}	0.00949 (0.0306)	-0.00568 (0.0321)	0.170 (0.103)
<i>Staracknw</i> _{<i>i,t+2</i>}	0.0257 (0.0328)	-0.00409 (0.0335)	0.279** (0.115)
<i>Staracknw</i> _{<i>i,t+3</i>}	-0.00123 (0.0337)	-0.0327 (0.0339)	0.263** (0.124)
Constant	0.424*** (0.0292)	0.400*** (0.0297)	0.643*** (0.105)
R-squared	0.022	0.024	0.061
Observations	14,408	12,946	1,462
Number of authors	900	816	84
Author FE	YES	YES	YES
Year FE	YES	YES	YES

Estimation Window: 1996-2017.

The dependent variable is Field Normalised Total Citations.

Note: The table reports the estimates based on the model specification in equation 3.1. Also, the 1996 cohort (always treated) is not dropped in this case. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively






Figure B.2 Event study model with homogenous star help effects at an individual level :excluding the authors with future star co-authored publications that happen after the star help ($t=0$): overall star help effect (row1, left) , one-time star help effect (row1, right), multiple star help effect (row 2, left)

Note: The figure plots the dynamic effect of the star's helpful interactions with an author at 95% confidence intervals. The event and observation window is from 1996-2017 .The dependent variable is raw publications.

Appendix C. Supplementary material for chapter 4

C.1. Econometric Modelling for Estimating the Heterogeneous Effects on Moving Star's Productivity.

Firstly, we define the year in which a star i moves as e_i . Secondly, we estimate the weighted average of cohort effects for a given time relative to the treatment event. To allow the estimated star's productivity effects to vary by cohort based on the year that the move occurs, we estimate the following equation:

$$\ln Y_{i,t} = \sum_e \left[\delta_{e,-5}(\mathbf{1}\{E_i = e\}starmove_{i,-5}) + \sum_{j=-4}^{-2} \delta_{e,j}(\mathbf{1}\{E_i = e\}starmove_{ij}) + \sum_{j=0}^4 \delta_{e,j}(\mathbf{1}\{E_i = e\}starmove_{ij}) + \delta_{e,5}(\mathbf{1}\{E_i = e\}starmove_{i,5}) \right] + \delta_t + \mu_i \quad (\text{C.1})$$

where $\mathbf{1}\{E_i = e\}$ is an indicator variable that takes the value 1, if the star i is an international star move in the year e and 0 otherwise. $\delta_{e,j}$ is the star move effect on star' productivity j year after star move in year e . The 1996 treated cohort is dropped from the analysis since it is always treated across the observation window. A further set of weights are estimated $Pr\{E_i = e | E_i \in [-j, T - j]\}$ that are equal to sample shares of each cohort for the relevant periods j . Finally, to obtain the IW estimator, we take a weighted average of the $\hat{\delta}_{e,j}$ (or $CATT_{e,j}$) and estimate Equation C.1 with relevant weights calculated.

$$\beta_j^* = \sum_e [\hat{\delta}_{e,j} Pr\{E_i = e | E_i \in [-j, T - j]\}] \quad (\text{C.2})$$

C.2. Coarsened Exact Matching (CEM) Stats for Alternate Control Group

Table C.1 Summary statistics: control and treated group (k-to-k matched) for the alternate control group

Variable	Control	Treated	Diff in mean	P-value
<i>An unbalanced panel of 94 matched stars, with 47 in each group</i>				
Year	2005.170	2005.170	0.000	1.000
University Code	27.405	27.405	0.000	1.000
Career age (at the time of move)	14.532	14.319	0.213	0.879
Career age bins (at the time of move)	4.426	4.426	0.000	1.000
Cumulative citations received per cumulative publications	169.429	175.738	-6.309	0.879
Cumulative citations received per cumulative publications Bins	4.362	4.362	0.000	1.000

Notes: Reports the t-test for the mean difference between control and treated groups one year before the star move event.

C.3. Robustness Test with Raw publications Output as Dependent Variable

Our main model uses the star's field normalised total citations as the dependent variable. As a final robustness test, we examine the sensitivity of the results from the main model specification to an alternative output measure: a star's publication-weighted output. Thus, we calculate the dependent variable using the following expression:

Publications:
$$Y_{i,t}^P = \frac{P_{i,t}}{\bar{P}_{s,t}},$$

where $P_{i,t}$ is the total number of publications by individual i published in year t , and $\bar{P}_{s,t}$ is the average number of publications in subject field s .

We present the estimated results using the alternative output measure in Table C.2, along with the corresponding event study plot in Figure C.1. In the event year, a star move is found to decrease their publication-weighted output by 52.72%. Moreover, the results indicate an average decline of 26.71% in their output for the years following a move. Therefore, our findings show that a star move exerts a more significant impact on their quality-adjusted output than on their raw number of publications. To the extent that the fall in productivity reflects a reallocation of time towards institution-building activities, these results suggest the own-productivity effects of this reallocation show up more on the intensive margin (quality of publications) rather than the extensive margin (number of publications).

Table C.2 Robustness test on dynamic star move effects on star's own productivity for normalised raw publications

	Overall Star Move Effect
<i>Starmove_{i,t-4}</i>	-0.0223 (0.0619)
<i>Starmove_{i,t-3}</i>	0.0755 (0.0519)
<i>Starmove_{i,t-2}</i>	0.0328 (0.0445)
<i>Starmove_{i,t}</i>	-0.749*** (0.0756)
<i>Starmove_{i,t+1}</i>	-0.329*** (0.0662)
<i>Starmove_{i,t+2}</i>	-0.223*** (0.0687)
<i>Starmove_{i,t+3}</i>	-0.423*** (0.0757)
<i>Starmove_{i,t+4}</i>	-0.268*** (0.0726)
Constant	1.218*** (0.0566)
R-squared	0.092
Observations	4,635
Number of Stars	236
Star FE	YES
Year FE	YES

Event Window: 1997-2017. Estimation Window: 1996-2017

The dependent variable is raw publication output

Note: The table reports the estimates based on the model specification in Equation 4.1 using raw publications output as the dependent variable. Robust standard errors are clustered at an individual level and reported in parentheses. *, **, and *** represent significance levels at the 10%, 5% and 1% respectively.




Figure C.1 Event study model with star's own productivity effects for the alternative dependent variable

Note: The figure plots the dynamic effect of the star's helpful interactions with an author at 95% confidence intervals. The event window is from 1997-2017. The dependent variable is raw publications.

Appendix D. Supplementary material for chapter 5

In this Appendix, we present the details on the dataset in Section D.1, information on missing data for keywords in Section D.2, and finally, in Section D.3, we present four robustness tests in which we first tested a quantity-adjusted symmetric relatedness measure based on the number of scientist's publication (D.3.1). In the second set of tests, we examine the sensitivity of the previous results to alternative scientist productivity measures: the raw publication counts for incumbent i in year t , and a scaled Euclidian Index proposed by Perry and Reny (2016) (D.3.2). Third, we exclude from the sample all incumbent scientists that are also stars in our definition and repeat our estimates (D.3.3). Finally, as a further test of appropriateness of a quadratic functional form of the relatedness relationship, we separately estimate the size of the star arrival effect for different ranges of the relatedness measure (D.3.4).

D.1. Data

Our dataset consists of information on researchers' publications, citations recorded, references, abstracts, and acknowledgements collected from Scopus across 27 subject fields for the period 1990 to 2017 in Denmark, Ireland and New Zealand. For each collected publication, we record the year, authors, affiliations, subject field, references, abstracts and acknowledgements, as well as, importantly, three unique identifiers – the EID Scopus publication identifier, the journals' Serial Identifier (e.g., ISSN) and the Scopus Author(s) ID. The citation data is count to 2019, the year we collected the data. For each author, we also access their catalogue of publications back to 1990, including publications that took place prior to joining the department in the three countries. Note that McHale et al. (2022) identify a department with Scopus-defined subject field at a given institution and remove some institutions where there was no activity in the subject area and we follow their procedure in this chapter. Stars scientists were allocated to departments (subject fields) where they have the largest number of publications. They were assigned to multiple departments in cases where they have an equal distribution of publications across multiple subject fields. In terms of the same first star arrival at multiple departments.

D.2. Missing Data for Keywords by Department.

Table D.1 Summary of missing data for keywords by department

Department	% Of author reported keywords missing (1)	% Of Scopus reported keywords missing (2)
Agricultural and Biological Sciences	28.34	18.39
Arts and Humanities	60.14	81.85
Biochemistry	45.82	7.88
Business	35.77	78.93
Chemical Engineering	37.18	22.33
Chemistry	55.79	26.71
Computer Science	46.85	15.79
Earth and Planetary Sciences	44.05	21.94
Economics	47.26	66.21
Engineering	45.06	10.05
Environmental Science	28.84	13.78
Immunology and Microbiology	49.01	5.51
Materials Science	49.57	15.99
Mathematics	35.80	59.42
Medicine	46.99	7.57
Neuroscience	27.67	6.02
Physics and Astronomy	64.65	37.61
Psychology	35.35	47.47
Social Sciences	52.80	69.15
Veterinary	47.27	11.08

Sources: Scopus and Author's calculations

Table D.2 Correlation matrix between relatedness across alternative knowledge spaces

Knowledge Spaces	Journal Overlap (1)	Journal Overlap (Quant.-Adj.) (2)	Journal Category Overlap (2)	Author-Reported Keyword Overlap (3)	Scopus-Reported Keyword Overlap (4)
Journal Overlap	1				
Journal Overlap (Q-Adjusted)	0.836***	1			
Journal Category Overlap	0.483***	0.394***	1		
Author-Reported Keyword Overlap	0.319***	0.278***	0.244***	1	
Scopus-Reported Keyword Overlap	0.245***	0.208***	0.288***	0.453***	1

Notes: *, **, *** denote significance at 10%, 5% and 1% levels, respectively.

Table D.3 Robustness: quantity-adjusted symmetric (cosine) scientific relatedness to the star

	<i>Journal Overlap</i>		
	(1)	(2)	(3)
$\alpha S_{i,t}$	0.022*** (0.007)	0.016** (0.007)	0.015** (0.006)
$\beta \bar{R}_{is,t^{*-1}} S_{i,t}$		0.208*** (0.065)	0.260* (0.133)
$\gamma \bar{R}_{is,t^{*-1}}^2 S_{i,t}$			-0.076 (0.201)
Individual FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Relatedness controls	No	Yes	Yes
Uni./Dep.	Yes	Yes	Yes
Observations	222,287	222,287	222,287
R ²	0.030	0.030	0.031
F-stat		8.58	5.77
p-value		0.0002	0.0006

Notes: We control for university and department outputs through all specifications (Uni./Dep.). Relatedness controls are the non-interacted relatedness variables included in the regression. Robust standard errors clustered by individual are provided in parentheses. *, **, *** denote significance at 10%, 5% and 1% levels, respectively. The joint significance test of α , β , and γ is 5.77 (p-value 0.0006).

D.3. Robustness

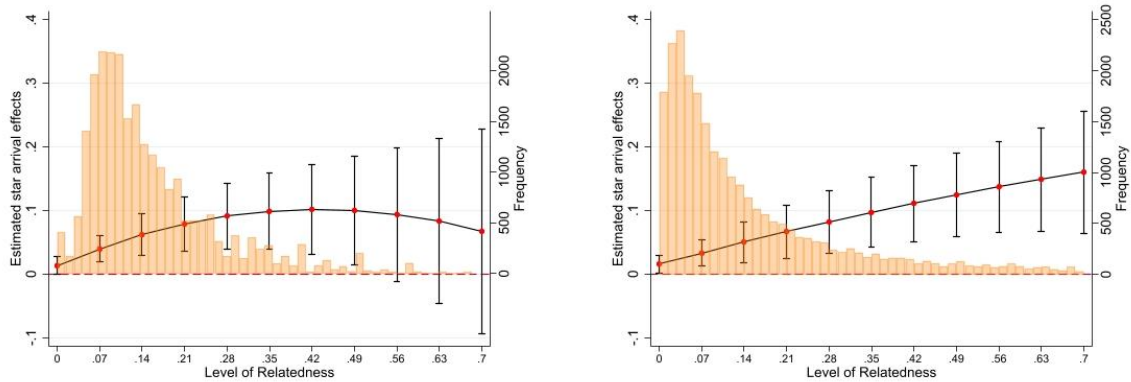
D.3.1. Robustness to alternative relatedness measures

In a first test of robustness, we examine a specification based on a relatedness measure that considers the number of journal occurrences and co-occurrences in a quantity-adjusted measure, where we adjust the cosine measure based on the number of scientist's publications observed in the journal rather than the simple binary indicator of publishing in that journal. Given C journals ($c = 1, \dots, C$), we define G_{ic} as the number of times i has published in journal c and G_{sc} as the number of times the star has published in journal c . The quantity adjusted cosine relatedness measure is then:

$$\tilde{R}_{is} = \frac{\sum_{c=1}^C G_{ic} G_{sc}}{\sqrt{(\sum_{c=1}^C G_{ic})^2 (\sum_{c=1}^C G_{sc})^2}} \quad (\text{D.1})$$

Table D.2 reports the correlation among this new measure, our baseline cosine measure and other three measures used in robustness tests (see below). Table D.3 presents the estimates based on this quantity-adjusted measure and Panel B of Figure D.1 illustrates the relationship between the size of the star arrival effect and the measure of relatedness. The results are broadly similar to the baseline results, with a consistent pattern for the star arrival effect across the values of relatedness. However, the coefficient on the quadratic term is not statistically significant at the conventional levels, which

results in a near linear relationship between the estimated star arrival effects and the level of relatedness.



Panel A. Symmetric (Cosine) scientific relatedness effects through relatedness levels (Max $\bar{R}_{iS} = 0.43$): Journal Overlap

Panel B. Quantity-Adjusted Symmetric (Cosine) scientific relatedness effects through relatedness levels (Max $\bar{R}_{iS} = 1.71$): Journal Overlap

Figure D.1 Non-linear relationship between star arrival and scientific relatedness effects

Notes: We estimate the conditional marginal effects of star relatedness through different levels of relatedness given by the star arrival. Panels A and B show the marginal effects based on the symmetric (cosine) and the Quantity-Adjusted symmetric (cosine) measures in Journal Overlap space. Standard errors for the marginal effects are calculated using the delta method

D.3.2. Robustness to alternative scientist productivity measures

While the baseline estimation uses a quality-adjusted publications measure of scientist productivity as the dependent variable to examine how the estimated star arrival effect evolves with the degree of relatedness between the star and the scientist, we further examine the sensitivity of these results to alternative scientist productivity measures in a further robustness test. We consider two measures to provide reasonable bounds in terms of relative weighting on the quantity versus quality output: (1) raw publication counts for incumbent i in year t (defined as $Y_{i,t}^P$); and (2) a scaled Euclidean Index proposed by Perry and Reny (2016) that places greater weight on highly cited publications relative to the field-normalized total citations measure, defined as

$$\left(\sum_{p_{i,t}=1}^P \left(\frac{c_{p_{i,t}}}{\bar{c}_{s,t}} \right)^2 \right)^{\frac{1}{2}}.$$

The results for the alternative productivity measures are provided in Table D.4, where the scientific relatedness is based on the symmetric (cosine) relatedness measured in the Journal Overlap space. Again, for comparison purposes, we repeat Column (1) with our previous results based on the field-normalized total citations productivity as the dependent variable. Columns (2) and (3) show the results for the raw

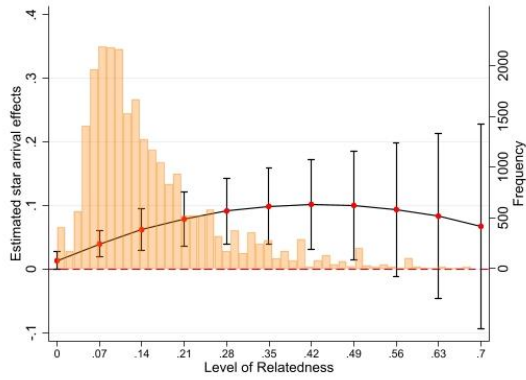
Table D.4 Robustness alternative productivity measures: journal overlap

	FNTC (Baseline) (1)	Raw Publication (2)	Euclidean Index (3)
$\alpha S_{i,t}$	0.0135** (0.0069)	-0.0015 (0.0066)	0.0112** (0.0056)
$\beta \bar{R}_{is,t^{*-1}} S_{i,t}$	0.410*** (0.140)	0.375*** (0.123)	0.298*** (0.110)
$\gamma \bar{R}_{is,t^{*-1}}^2 S_{i,t}$	-0.476* (0.272)	-0.340 (0.229)	-0.377* (0.217)
Individual FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Relatedness Controls	Yes	Yes	Yes
Observations	222,287	222,287	222,287
R ²	0.031	0.059	0.021
F-stat	4.76	4.08	4.36
p-value	0.0026	0.0066	0.0045


Notes: We control for university and department outputs through all specifications (Uni./Dep.). Relatedness controls are the non-interacted relatedness variables included in the regression. Robust standard errors clustered by individual are provided in the parentheses. *, **, *** denote significance at 10%, 5% and 1% levels, respectively.

publications and Euclidean measures, respectively. The two alternative measures show that the estimated coefficients of the (linear) interaction between the star-arrival and relatedness remain positive (0.375 for the raw publication and 0.298 for the Euclidean Index) and statistically significant at the 1% significance level, although we only find a positive and statistically significant coefficient on the star arrival binary variable for the Euclidean output measure (1.12 log points). In addition to the linear star-arrival/relatedness effect, we find that the coefficients on the quadratic term in Columns (2) and (3) are again consistent with the baseline results. While we note that the non-linear relationship between the relatedness measure and the star arrival effect is only identified as statistically significant in the Euclidean output measure, we find the overall patterns generally similar to the baseline.


We plot the interrelationship between the size of the star arrival effect and the measure of relatedness under the different output measures in Figure D.2. Here the results show that the estimated star-arrival effect peaks at two different relatedness levels (0.551 for the raw publication versus 0.395 for the Euclidean Index in Panel B and C, respectively). Nevertheless, they demonstrate a broadly similar pattern to the one observed for our baseline measure which is also presented in Panel A of Figure D.2. Overall, our results appear to be generally robust to the choice of scientist productivity measure.



Panel A. Symmetric (Cosine) scientific relatedness effects through relatedness levels (Max $\bar{R}_{iS} = 0.43$): Field Normalized Total Citations (Baseline)



Panel B. Symmetric (Cosine) scientific relatedness effects through relatedness levels (Max $\bar{R}_{iS} = 0.55$): Raw Publications



Panel C. Symmetric (Cosine) scientific relatedness effects through relatedness levels (Max $\bar{R}_{iS} = 0.39$): Euclidean Index

Figure D.2 Non-linear relationship between star arrival and scientific relatedness effects-alternative productivity measures.

Notes: We estimate the conditional marginal effects of star relatedness through different levels of relatedness given by the star arrival based on the symmetric (cosine) measure for alternative productivity output in Journal Overlap space. Panel A shows the marginal effects on Field Normalized Total Citations. Panels B and C show the marginal effects on Raw Publications and Euclidean Index, respectively. Standard errors for the marginal effects are calculated using the delta method

D.3.3. Robustness to alternative sample of incumbent scientists

As an additional test of robustness, we changed the sample to exclude incumbent scientists who were also stars⁴⁶ and were in the receiving department when a star arrival occurred. There were 42 incumbent scientists also identified as stars. They were removed from the sample, and a new matching process through the k-to-k Coarsened Exact Match (CEM) approach was executed. Table D.5 presents the estimates for this alternative

⁴⁶ Star scientists are those positioned at or above the 95th percentile of scientists in the cumulative distributions of citations received since 1990 for their subject field in any given year, where the citation measure includes all forward citations to a publication up to 2019.

sample of incumbent scientists. The results are similar to our baseline results, which are repeated below for comparison purposes.

Table D.5 Robustness of alternative sample

	Asymmetry (<i>Scientist perspective</i>)		Asymmetry (<i>Star perspective</i>)		Symmetric (<i>Cosine</i>)	
	Full sample	After removing stars	Full sample	After removing stars	Full sample	After removing stars
$\alpha S_{i,t}$	0.0155** (0.00694)	0.0143** (0.00690)	0.0149** (0.00684)	0.0130* (0.00678)	0.013** (0.006)	0.0121* (0.00684)
$\beta \bar{R}_{is,t^*-1} S_{i,t}$	0.121 (0.0909)	0.0974 (0.0856)	0.469*** (0.166)	0.423*** (0.153)	0.410*** (0.140)	0.363*** (0.132)
$\gamma \bar{R}_{is,t^*-1}^2 S_{i,t}$	-0.0444 (0.0979)	-0.0464 (0.0931)	-0.529*** (0.200)	-0.452** (0.191)	-0.476* (0.272)	-0.457* (0.263)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Relatedness Control	Yes	Yes	Yes	Yes	Yes	Yes
Observations	222,287	220,598	222,287	220,598	222,287	220,598

D.3.4. Robustness to no quadratic functional form restrictions

Our baseline specification, ex-ante, places a quadratic restriction on the effect of relatedness. To improve the justification for the inverse U-shaped relationship, we parametrically estimate effects separately by bins based on different levels of relatedness. We run the following DID specification for separate subsamples of the data to test for non-linearity in effects:

$$\ln Y_{i,t} = \alpha S_i + \beta \bar{R}_{is} S_{i,t} + Z'_{i,t} \lambda + \lambda_i + \lambda_t + u_{i,t} \quad (\text{D.2})$$

The results for the Journal Overlap measure are shown in Figure D.3. The point estimates are the linear combination of the star arrival effect and relatedness interaction ($\alpha + \beta$). Based on the cosine (symmetric) measure, Panel A shows the effects follow an inverted U-shaped relationship with the level of relatedness. This quadratic relationship is supported by the asymmetric relatedness from the star perspective (Panel C), but not from asymmetric relatedness from the incumbent perspective. However, as highlighted in the chapter, due to a wider knowledge base, a measure of relatedness taken from the perspective of the star is a better predictor of star-arrival productivity effects than a measure taken from the perspective of the incumbent scientist. Eventually, as both

viewpoints matter for the creation process of new knowledge in an interactive learning environment, the cosine measure combining both perspectives also shows an inverted U-shaped curve.

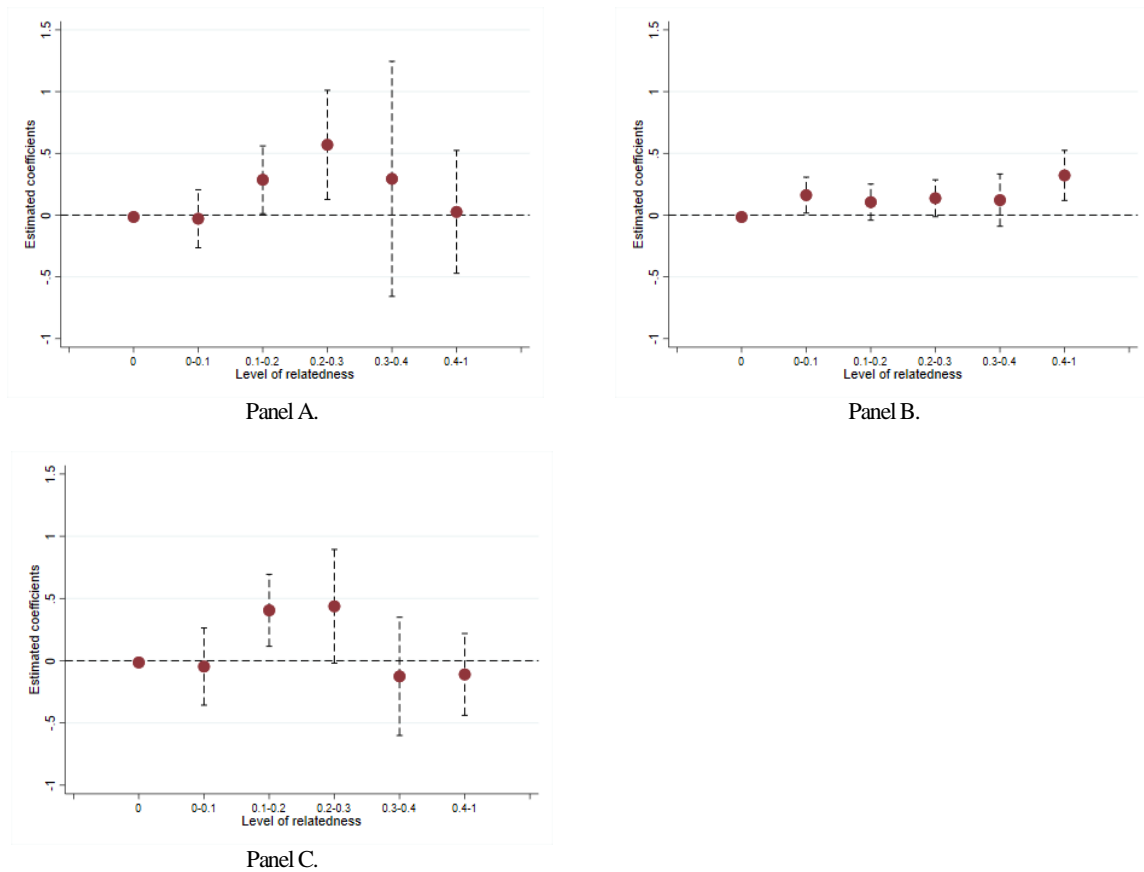


Figure D.3 Parametrically testing for non-linearity in effects for all three measures using journal overlap space

Notes: Panel A: Symmetric (cosine) measure. Panel B: Asymmetric scientific relatedness from the incumbent's perspective. Panel C: Asymmetric scientific relatedness from the star's perspective. The x-axis shows the bins for the level of relatedness. We use average relatedness to assign an incumbent in a specific bin. The incumbents with relatedness greater than 0.4 to 1 are assigned to a single bin, as very few have that high relatedness. Due to many incumbents with 0 relatedness, the confidence intervals for 0 relatedness are very small relative to other bins. Therefore the confidence intervals are not visible on the graph. The figure plots 95 % confidence intervals.

Bibliography

- Adler, M., 1985. Stardom and talent. *The American economic review*, 75(1), pp. 208-212.
- Agrawal, A. and Cockburn, I., 2003. The anchor tenant hypothesis: exploring the role of large, local, R&D-intensive firms in regional innovation systems. *International journal of industrial organization*, 21(9), pp. 1227-1253.
- Agrawal, A., Cockburn, I. and McHale, J., 2006. Gone but not forgotten: knowledge flows, labor mobility, and enduring social relationships. *Journal of Economic Geography*, 6(5), pp. 571-591.
- Agrawal, A. K., McHale, J., & Oettl, A. (2014). *Why stars matter* (No. w20012). National Bureau of Economic Research.
- Agrawal, A., McHale, J. and Oettl, A., 2017. How stars matter: Recruiting and peer effects in evolutionary biology. *Research Policy*, 46(4), pp.853-867.
- Ahuja, G. (2000). Collaboration networks, structural holes, and innovation: A longitudinal study. *Administrative Science Quarterly*, 45(3), 425–455.
- Aksnes, D. W., Rørstad, K., Piro, F. N., & Sivertsen, G. (2013). Are mobile researchers more productive and cited than non-mobile researchers? A large-scale study of Norwegian scientists. *Research Evaluation*, 22(4), 215-223.
- Allison, P. D., & Long, J. S. (1990). Departmental Effects on Scientific Productivity. *American Sociological Review*, 55(4), 469-478.
- Arellano, M. and Bond, S., 1991. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of economic studies*, 58(2), pp. 277-297.
- Arthur, W. B. (2009). *The nature of technology: What it is and how it evolves*. Simon and Schuster.
- Athey, S., Eckles, D., & Imbens, G. W. (2018). Exact p -Values for Network Interference. *Journal of the American Statistical Association*, 113(521), 230–240.
- Audretsch, D.B. and Feldman, M.P., 1996. R&D spillovers and the geography of innovation and production. *American economic review*, 86(3), pp. 630-640.
- Autant-Bernard, C., & LeSage, J. P. (2011). Quantifying knowledge spillovers using spatial econometric models. *Journal of Regional Science*, 51(3), 471–496.
- Autor, D., Dorn, D., Katz, L.F., Patterson, C. and Van Reenen, J., 2020. The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics*, 135(2), pp. 645-709.
- Autor, D.H., 2003. Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of labor economics*, 21(1), pp. 1-42.
- Azoulay, P., & Zivin, J. G. (2005). *Peer effects in the workplace: evidence from professional transitions for the superstars of medicine* (No. y: 2006: x: 1).
- Azoulay, P., Fons-Rosen, C., & Graff Zivin, J. S. (2019). Does Science Advance One Funeral at a Time? *American Economic Review*, 109(8), 2889–2920.

- Azoulay, P., Ganguli, I., & Zivin, J. G. (2017). The mobility of elite life scientists: Professional and personal determinants. *Research Policy*, 46(3), 573-590.
- Azoulay, P., Graff Zivin, J. S., & Wang, J. (2010). Superstar extinction. *The Quarterly Journal of Economics*, 125(2), 549-589.
- Bathelt, H., Malmberg, A., & Maskell, P. (2004). Clusters and knowledge: Local buzz, global pipelines and the process of knowledge creation. *Progress in Human Geography*, 28(1), 31–56.
- Azoulay, P., Stuart, T. and Wang, Y. (2014). Matthew: Effect or fable? *Management Science*, 60(1), pp. 92-109.
- Bäker, A. (2015). Non-tenured post-doctoral researchers' job mobility and research output: An analysis of the role of research discipline, department size, and coauthors. *Research Policy*, 44(3), 634-650
- Baltagi, B. H., & Li, D. (2002). Series estimation of partially linear panel data models with fixed effects. *Annals of Economics and Finance*, 3(6), 103–116.
- Betancourt, N., Jochem, T., & Otner, S. M. (2023). Standing on the shoulders of giants: How star scientists influence their coauthors. *Research Policy*, 52(1), 104624.
- Berg, N., & Faria, J. (2008). Negatively correlated author seniority and the number of acknowledged people: Name-recognition as a signal of scientific merit?. *The Journal of Socio-Economics*, 37(3), 1234-1247.
- Bergman, E. M., & Schubert, U. (2005). Spillovers and innovation, environment and space: policy uncertainties and research opportunities. In *Spillovers and Innovations* (pp. 157-177). Springer, Vienna.
- Bilinski, A. and Hatfield, L. A., 2020. Nothing to see here? Non-inferiority approaches to parallel trends and other model assumptions. arXiv:1805.03273 [stat].
- Black, D. and Henderson, V., 1999. A theory of urban growth. *Journal of Political Economy*, 107(2), pp. 252-284.
- Blackwell, M., Iacus, S., King, G., & Porro, G. (2009). cem: Coarsened exact matching in Stata. *The Stata Journal*, 9(4), 524-546.
- Bloom, N., Van Reenen, J., & Williams, H. (2019). A toolkit of policies to promote innovation. *Journal of economic perspectives*, 33(3), 163-184.
- Blumberg, A. A., 1968. Logistic Growth Rate Functions. *Journal of Theoretical Biology*, 21(1), pp. 42-44.
- Blumenthal, D., Causino, N., Campbell, E., & Louis, K. S. (1996). Relationships between academic institutions and industry in the life sciences—an industry survey. *New England Journal of Medicine*, 334(6), 368-374.
- Boden, M., & Fitzgibbon, M. (1995). Evaluation of science and technology policy in Ireland. *Research Evaluation*, 5(1), 55-62.
- Bond, S.R., 2002. Dynamic panel data models: a guide to micro data methods and practice. *Portuguese economic journal*, 1(2), pp.141-162.

- Borjas, G.J., 2003. The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market. *Quarterly Journal of Economics*, 118(4), pp. 1335-1374.
- Borusyak, K. and Jaravel, X., 2017. Revisiting event study designs. Available at SSRN 2826228.
- Borusyak, K., Jaravel, X., & Spiess, J. (2022). *Revisiting Event Study Designs: Robust and Efficient Estimation* (SSRN Scholarly Paper 2826228)
- Boschma, R. (2005). Proximity and innovation: A critical assessment. *Regional Studies*, 39(1), 61–74.
- Boschma, R., & Frenken, K. (2010). The spatial evolution of innovation networks: A proximity perspective. In R. Boschma & R. Martin (Eds.), *The handbook of evolutionary economic geography* (pp. 120–135). Edward Elgar Publishing.
- Boschma, R., Minondo, A., & Navarro, M. (2013). The emergence of new industries at the regional level in Spain: A proximity approach based on product relatedness. *Economic Geography*, 89(1), Article 1.
- Boudreau, K. J., Brady, T., Ganguli, I., Gaule, P., Guinan, E., Hollenberg, A., & Lakhani, K. R. (2017). A field experiment on search costs and the formation of scientific collaborations. *Review of Economics and Statistics*, 99(4), 565-576.
- Bozeman, B., & Corley, E. (2004). Scientists' collaboration strategies: implications for scientific and technical human capital. *Research policy*, 33(4), 599-616.
- Bramoullé, Y., Djebbari, H., & Fortin, B. (2020). Peer effects in networks: A survey. *Annual Review of Economics*, 12(1).
- Breschi, S. and Lissoni, F., 2009. Mobility of skilled workers and co-invention networks: an anatomy of localized knowledge flows. *Journal of economic geography*, 9(4), pp. 439-468.
- Burt, R. S. (1992). *Structural holes: The social structure of competition*. Harvard University Press.
- Burt, R. S. (2004). Structural holes and good ideas. *American Journal of Sociology*, 110(2), 349–399.
- Callaway, B. and Sant'Anna, P. H., 2021. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), pp.200-230.
- Cañibano, C., Otamendi, J., & Andújar, I. (2008). Measuring and assessing researcher mobility from CV analysis: the case of the Ramón y Cajal programme in Spain. *Research Evaluation*, 17(1), 17-31
- Card, D., 1990. Unexpected inflation, real wages, and employment determination in union contracts. *American Economic Review*, pp. 669-688.
- Catalini, C. (2018). Microgeography and the direction of inventive activity. *Management Science*, 64(9), 4348-4364.
- Cattell, R. B., & Tiner, L. G. (1949). The varieties of structural rigidity. *Journal of Personality*, 17(3).
- Chang, M. H. (2023). Cascading innovation: R&D team design and performance implications of mobility. *Strategic Management Journal*, 44(5), 1218-1253.

- Clancy, M. S. (2018). Inventing by combining pre-existing technologies: Patent evidence on learning and fishing out. *Research Policy*, 47(1), 252–265.
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128–152.
- Content, J., & Frenken, K. (2016). Related variety and economic development: A literature review. *European Planning Studies*, 24(12), 2097–2112.
- Coupé, T., Smeets, V., & Warzynski, F. (2006). Incentives, Sorting and Productivity along the Career: Evidence from a Sample of Top Economists. *Journal of Law, Economics, & Organization*, 137-167.
- Cronin, B. (1991). Let the credits roll: a preliminary examination of the role played by mentors and trusted assessors in disciplinary formation. *Journal of documentation*.
- Cronin, B., McKenzie, G., & Rubio, L. (1993). The norms of acknowledgement in four humanities and social sciences disciplines. *Journal of documentation*.
- De Chaisemartin, C. and d'Haultfoeuille, X., 2020. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), pp. 2964-96.
- De Filippo, D., Casado, E. S., & Gomez. (2009). Quantitative and qualitative approaches to the study of mobility and scientific performance: a case study of a Spanish university, 18(3), 191–200.
- Deryugina, T., 2017. The fiscal cost of hurricanes: Disaster aid versus social insurance. *American Economic Journal: Economic Policy*, 9(3), pp. 168-98.
- Desrochers, N., Paul-Hus, A., Haustein, S., Costas, R., Mongeon, P., Quan-Haase, A., & Larivière, V. (2018). Authorship, citations, acknowledgments and visibility in social media: Symbolic capital in the multifaceted reward system of science. *Social science information*, 57(2), 223-248.
- Dobkin, C., Finkelstein, A., Kluender, R. and Notowidigdo, M.J., 2018. The economic consequences of hospital admissions. *American Economic Review*, 108(2), pp. 308-52.
- Docquier, F. and Rapoport, H., 2012. Globalization, brain drain, and development. *Journal of Economic Literature*, 50(3), pp. 681-730.
- Dougherty, D. (1992). Interpretive barriers to successful product innovation in large firms. *Organization Science*, 3(2), 179–202.
- Duranton, G. and Puga, D., 2004. Micro-foundations of urban agglomeration economies. In *Handbook of regional and urban economics* (Vol. 4, pp. 2063-2117). Elsevier.
- European Commission. (2019). *Peer Review of the Danish R&I System: Ten steps, and a leap forward: taking Danish innovation to the next level*. Directorate-General for Research and Innovation, Directorate G — Research and Innovation Outreach, Unit G1 — ERA and Country Intelligence. ISBN 978-92-76-11921-0. doi:10.2777/603976.
- Feldman, M. S. (2000). Organizational routines as a source of continuous change. *Organization science*, 11(6), 611-629
- Fernández-Zubieta, A., Geuna, A., & Lawson, C. (2015). Mobility and productivity of research scientists. In *Global mobility of research scientists* (pp. 105-131). Academic Press.

- Fernández-Zubieta, A., Geuna, A., & Lawson, C. (2016). Productivity pay-offs from academic mobility: should I stay or should I go?. *Industrial and Corporate Change*, 25(1), 91-114.
- Fini, R., Jourdan, J., & Perkmann, M. (2018). Social valuation across multiple audiences: The interplay of ability and identity judgments. *Academy of Management Journal*, 61(6), 2230-2264.
- Fini, R., Perkmann, M., Kenney, M., & Maki, K. M. (2023). Are public subsidies effective for university spinoffs? Evidence from SBIR awards in the University of California system. *Research Policy*, 52(1), 104662.
- Fleming, L., & Sorenson, O. (2004). Science as a map in technological search. *Strategic Management Journal*, 25(8–9), 909–928.
- Frenken, K., Oort, F. V., & Verburg, T. (2007). Related variety, unrelated variety and regional economic growth. *Regional Studies*, 41(5), Article 5.
- Freyaldenhoven, S., Hansen, C. and Shapiro, J.M., 2019. Pre-event trends in the panel event-study design. *American Economic Review*, 109(9), pp. 3307-38.
- Fujita, M. and Thisse, J.F., 2003. Does geographical agglomeration foster economic growth? And who gains and loses from it? *Japanese Economic Review*, 54(2), pp. 121-145.
- Gabaix, X., 1999. Zipf's law for cities: an explanation. *Quarterly Journal of Economics*, 114(3), pp. 739-767.
- Galetti, J. R. B., Tessarin, M. S., & Morceiro, P. C. (2021). Skill relatedness, structural change and heterogeneous regions: Evidence from a developing country. *Papers in Regional Science*, 100(6), 1355–1376.
- Ganguli, I. (2015). Who Leaves and Who Stays? Evidence on Immigrant Selection from the Collapse of Soviet Science. *Global Mobility of Research Scientists*, 133-154.
- Georghiou, L. (1995). Research evaluation in European national science and technology systems. *Research evaluation*, 5(1), 3-10.
- Gertler, M. S. (2003). Tacit knowledge and the economic geography of context, or The undefinable tacitness of being (there). *Journal of Economic Geography*, 3(1), 75–99.
- Geuna, A. (Ed.). (2015). *Global mobility of research scientists: The economics of who goes where and why*. Academic Press.
- Glaeser, E.L., 2008. *Cities, agglomeration, and spatial equilibrium*. Oxford University Press.
- Glaeser, E. L. (1999). Learning in cities. *Journal of urban Economics*, 46(2), 254-277.
- Gluckman, P. (2013). *The role of evidence in policy formation and implementation: A report from the Prime Minister's Chief Science Advisor*. Office of the Prime Minister's Science Advisory Committee (New Zealand).
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254-277.
- Goyal, S., 2007. *Introduction to connections: An introduction to the economics of networks*. Princeton University Press, Princeton NJ.
- Goyal, S., M.J. van der Leij, J.L. Moraga-González. 2006. Economics: An Emerging Small World. *Journal of Political Economy*, 114(2), pp. 403-412.

- Grigoriou, K., & Rothaermel, F. T. (2014). Structural microfoundations of innovation: The role of relational stars. *Journal of Management*, 40(2), 586-615.
- Griliches, Z., 1992. The Search for R&D Spillovers. *The Scandinavian Journal of Economics*, 94, pp. S29-S47.
- Grillitsch, M., Asheim, B., & Trippl, M. (2018). Unrelated knowledge combinations: The unexplored potential for regional industrial path development. *Cambridge Journal of Regions, Economy and Society*, 11(2), Article 2.
- Grossman, G. M., & Helpman, E. (1991). Trade, knowledge spillovers, and growth. *European economic review*, 35(2-3), 517-526.
- Groysberg, B., Lee, L. E., & Nanda, A. (2008). Can they take it with them? The portability of star knowledge workers' performance. *Management Science*, 54(7), 1213-1230.
- Gyourko, J., Mayer, C. and Sinai, T., 2013. Superstar cities. *American Economic Journal: Economic Policy*, 5(4), pp. 167-99.
- Haeussler, C., Jiang, L., Thursby, J., & Thursby, M. (2014). Specific and general information sharing among competing academic researchers. *Research policy*, 43(3), 465-475.
- Halevi, G., Moed, H. F., & Bar-Ilan, J. (2016). Does Research Mobility Have an Effect on Productivity and Impact?. *International Higher Education*, (86), 5–6.
- Hellqvist, B. (2010). Referencing in the humanities and its implications for citation analysis. *Journal of the American Society for Information Science and Technology*, 61(2), 310-318.
- Hess, A.M. and Rothaermel, F.T., 2011. When are assets complementary? Star scientists, strategic alliances, and innovation in the pharmaceutical industry. *Strategic Management Journal*, 32(8), pp. 895-909.
- Hoisl, K. (2007). Tracing mobile inventors—the causality between inventor mobility and inventor productivity. *Research Policy*, 36(5), 619-636.
- Horta, H., Veloso, F. M., & Grediaga, R. (2010). Navel gazing: Academic inbreeding and scientific productivity. *Management Science*, 56(3), 414-429.
- Hunter, R. S., Oswald, A. J., & Charlton, B. G. (2009). The elite brain drain. *The Economic Journal*, 119(538), F231-F251.
- Iacus, S. M., King, G., & Porro, G. (2011). Multivariate matching methods that are monotonic imbalance bounding. *Journal of the American Statistical Association*, 106(493), 345-361.
- Iacus, S. M., King, G., & Porro, G. (2012). Causal inference without balance checking: Coarsened exact matching. *Political analysis*, 20(1), 1-24.
- Jackson, M.O. and Rogers, B.W., 2005. The economics of small worlds. *Journal of the European Economic Association*, 3(2-3), pp. 617-627.
- Jackson, M.O. and Rogers, B.W., 2007. Meeting strangers and friends of friends: How random are social networks? *American Economic Review*, 97(3), pp. 890-915.
- Jackson, M.O., 2008, December. Average distance, diameter, and clustering in social networks with homophily. In *International Workshop on Internet and Network Economics* (pp. 4-11). Springer, Berlin, Heidelberg.

- Jackson, M.O., Rogers, B.W. and Zenou, Y., 2017. The economic consequences of social-network structure. *Journal of Economic Literature*, 55(1), pp. 49-95.
- Jacobson, L.S., LaLonde, R.J. and Sullivan, D.G., 1993. Earnings losses of displaced workers. *American Economic Review*, pp. 685-709.
- Jaffe, A.B., Trajtenberg, M. and Henderson, R., 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics*, 108(3), pp. 577-598.
- Jaffe, A. B., & Trajtenberg, M. (2002). *Patents, citations, and innovations: A window on the knowledge economy*. MIT press.
- Jaffe, A. (2013). Science and innovation policy for New Zealand. *New Zealand Science Review*, 70(3), 55-61.
- Jones, B. F. (2009). The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? *The Review of Economic Studies*, 76(1), Article 1.
- Jovanovic, B. (1979). Job matching and the theory of turnover. *Journal of political economy*, 87(5, Part 1), 972-990.
- Kahn-Lang, A. and Lang, K., 2018. The Promise and Pitfalls of Differences-in-Differences: Reflections on '16 and Pregnant' and Other Applications. Working Paper 24857, National Bureau of Economic Research.
- Kalpazidou Schmidt, E., Graversen, E. K., & Langberg, K. (2003). Innovation and dynamics in public research environments in Denmark: a research-policy perspective. *Science and Public Policy*, 30(2), 107-116.
- Kane, A. (2005). Indicators and evaluation for science, technology and innovation. *ICSTI Task Force on Metrics and Impact*. Downloaded on, 15.
- Kane, A. (2014). Rationales for science, technology and innovation policy. (accessed on July 2024)
- Kapur, D. and McHale, J., 2006. Should a cosmopolitan worry about the "Brain Drain"? *Ethics & International Affairs*, 20(3), pp. 305-320.
- Kato, M., & Ando, A. (2017). National ties of international scientific collaboration and researcher mobility found in Nature and Science. *Scientometrics*, 110, 673-694.
- Kehoe, R. R., & Tzabbar, D. (2015). Lighting the way or stealing the shine? An examination of the duality in star scientists' effects on firm innovative performance. *Strategic Management Journal*, 36(5), 709-727.
- Kerr, William, 2019. *The Gift of Global Talent: How Migration Shapes Business, Economy & Society*. Stanford: Stanford Business Book.
- Khanna, R., 2021. Aftermath of a tragedy: A star's death and coauthors' subsequent productivity. *Research Policy*, 50(2), p.104159.
- King, G., & Nielsen, R. (2019). Why propensity scores should not be used for matching. *Political analysis*, 27(4), 435-454.

- King, G., Nielsen, R., Coberley, C., Pope, J. E., & Wells, A. (2011). Comparative effectiveness of matching methods for causal inference. *Unpublished manuscript, Institute for Quantitative Social Science, Harvard University, Cambridge, MA.*
- Kogut, B., & Zander, U. (1992). Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science*, 3(3), 383–397.
- Kripfganz, S. (2016). Quasi–maximum Likelihood Estimation of Linear Dynamic Short-T panel-data Models. *The Stata Journal*, 16(4), 1013–1038.
- Krugman, P., 1991. Increasing returns and economic geography. *Journal of Political Economy*, 99(3), pp. 483-499.
- Lacetera, N., Cockburn, I.M. and Henderson, R., 2004. Do firms change capabilities by hiring new people? A study of the adoption of science-based drug discovery. *Business Strategy over the Industry Lifecycle*. Emerald Group Publishing Limited.
- Laudel, G. (2002). What do we measure by co-authorships?. *Research evaluation*, 11(1), 3-15.
- Lehmann, S., Jackson, A. D., & Lautrup, B. E. (2006). Measures for measures. *Nature*, 444(7122), 1003-1004.
- Leitch, S., Motion, J., Merlot, E., & Davenport, S. (2014). The fall of research and rise of innovation: Changes in New Zealand science policy discourse. *Science and Public Policy*, 41(1), 119-130.
- Lenzi, C. (2009). Patterns and determinants of skilled workers' mobility: Evidence from a survey of Italian inventors. *ECONOMICS OF INNOVATION AND NEW TECHNOLOGY*, 18, 161-179.
- Li, W., Aste, T., Caccioli, F., & Livan, G. (2019). Early co-authorship with top scientists predicts success in academic careers. *Nature communications*, 10(1), 1-9.
- Liu, H., Mihm, J., & Sosa, M. E. (2018). Where Do Stars Come From? The Role of Star vs. Nonstar Collaborators in Creative Settings. *Organization Science*, 29(6), 1149–1169.
- Liu, P. W. (1986). Human capital, job matching and earnings growth between jobs: an empirical analysis. *Applied Economics*, 18(10), 1135-1147.
- Lotka, A.J., 1926. The frequency distribution of scientific productivity. *Journal of the Washington Academy of Sciences*, 16(12), pp. 317-323.
- MacCarthaigh, M. (2013). Reform of public policy-making in Ireland. *Journal of the Statistical and Social Inquiry Society of Ireland (2013)*.
- Mackintosh, K. H. (1972). *Acknowledgement patterns in sociology*. University of Oregon.
- Maier, G., Kurka, B., & Tripl, M. (2007). Knowledge spillover agents and regional development: spatial distribution and mobility of star scientists. *DYNREG (dynamic regions in a knowledge-driven global economy)*, 17, 35.
- Malmendier, U. and Tate, G., 2009. Superstar CEOs. *Quarterly Journal of Economics*, 124(4), pp. 1593-1638.
- McCain, K. W. (1991). Communication, competition, and secrecy: The production and dissemination of research-related information in genetics. *Science, Technology, & Human Values*, 16(4), 491-516.

- McHale, J., Harold, J., Mei, J. C., Sasidharan, A., & Yadav, A. (2022). Stars as catalysts: an event-study analysis of the impact of star-scientist recruitment on local research performance in a small open economy. *Journal of Economic Geography*, Volume 23, Issue 2, March 2023, Pages 343–369
- Merton, R. K. (1973). *The sociology of science: Theoretical and empirical investigations*. University of Chicago press.
- Merton, R.K., 1968. The Matthew effect in science: The reward and communication systems of science are considered. *Science*, 159(3810), pp. 56-63.
- Mohnen, M. (2022). Stars and brokers: Knowledge spillovers among medical scientists. *Management Science*, 68(4), 2513-2532.
- Murray, F. (2010). The oncomouse that roared: Hybrid exchange strategies as a source of distinction at the boundary of overlapping institutions. *American Journal of sociology*, 116(2), 341-388.
- Neffke, F., Henning, M., & Boschma, R. (2011). How do regions diversify over time? Industry relatedness and the development of new growth paths in regions. *Economic Geography*, 87(3), Article 3.
- Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change*. The Belknap Press of Harvard University Press.
- Nelson, R. R., & Winter, S. G. (1985). *An Evolutionary Theory of Economic Change*. Harvard University Press.
- Netz, N., Hampel, S., & Aman, V. (2020). What effects does international mobility have on scientists' careers? A systematic review. *Research evaluation*, 29(3), 327-351.
- Newman, M. (2018). *Networks (Second Edition)*. Oxford University Press.
- Newman, M.E., 2011. The structure of scientific collaboration networks. In *The Structure and Dynamics of Networks* (pp. 221-226). Princeton University Press.
- Nickell, S., 1981. Biases in dynamic models with fixed effects. *Econometrica: Journal of the Econometric Society*, pp. 1417-1426.
- Nooteboom, B. (2000). Learning by interaction: Absorptive capacity, cognitive distance and governance. *Journal of Management and Governance*, 4(1), 69–92.
- Obstfeld, D. (2005). Social networks, the tertius iungens orientation, and involvement in innovation. *Administrative Science Quarterly*, 50(1), 100–130.
- Oettl, A. (2012). Reconceptualizing stars: Scientist helpfulness and peer performance. *Management Science*, 58(6), 1122-1140.
- Paruchuri, S. (2010). Intraorganizational networks, interorganizational networks, and the impact of central inventors: A longitudinal study of pharmaceutical firms. *Organization Science*, 21(1), 63-80.
- Paul-Hus, A., & Desrochers, N. (2019). Acknowledgements are not just thank you notes: A qualitative analysis of acknowledgements content in scientific articles and reviews published in 2015. *Plos one*, 14(12), e0226727.

- Paul-Hus, A., Desrochers, N., & Costas, R. (2016). Characterization, description, and considerations for the use of funding acknowledgement data in Web of Science. *Scientometrics*, 108(1), 167-182.
- Paul-Hus, A., Díaz-Faes, A. A., Sainte-Marie, M., Desrochers, N., Costas, R., & Larivière, V. (2017). Beyond funding: Acknowledgement patterns in biomedical, natural and social sciences. *PloS one*, 12(10), e0185578.
- Perry, M. and Reny, P.J., 2016. How to count citations if you must. *American Economic Review*, 106(9), pp. 2722-41.
- Price, D.J.D.S., 1963. *Little science, big science*. Columbia University Press.
- Radicchi, F., Fortunato, S. and Castellano, C., 2008. Universality of citation distributions: Toward an objective measure of scientific impact. *Proceedings of the National Academy of Sciences*, 105(45), pp. 17268-17272.
- Radicchi, F., Fortunato, S., & Castellano, C. (2008). Universality of citation distributions: Toward an objective measure of scientific impact. *Proceedings of the National Academy of Sciences*, 105(45), 17268–17272.
- Robst, J. (1995). Career Mobility, Job Match, and Overeducation. *Eastern Economic Journal*, 21(4), 539–550.
- Roolah, T. (2011). The characteristics of small country national innovation systems. In *Innovation Systems in Small Catching-Up Economies: New Perspectives on Practice and Policy* (pp. 21-37). New York, NY: Springer New York.
- Romer, P. M. (1986). Increasing returns and long-run growth. *Journal of political economy*, 94(5), 1002-1037.
- Romer, P. M. (1990). Endogenous technological change. *Journal of political Economy*, 98(5, Part 2), S71-S102.
- Roodman, D., 2009. How to do xtabond2: An introduction to difference and system GMM in Stata. *The Stata Journal*, 9(1), pp. 86-136.
- Rosen, S., 1981. The economics of superstars. *American Economic Review*, 71(5), pp. 845-858.
- Roth, J. (2022). Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends. *American Economic Review: Insights*, 4(3), 305-22.
- Ruane, F., & Whelan, B. (2011). Building Research Capacity in the Social Sciences-Alternatives Approaches. *Journal of the Statistical & Social Inquiry Society of Ireland*, 40.
- Sasidharan, A., McHale, J., & Harold, J. (2024). Star help and knowledge transfer: an event study analysis of star interactions observed from acknowledgement texts. *The Journal of Technology Transfer*, 1-49.
- Schiller, D., & Diez, J. R. (2008). Mobile star scientists as regional knowledge spillover agents. *IAREG Working Pap*, (11), 3-33.
- Schiller, D., & Diez, J. R. (2010). Local embeddedness of knowledge spillover agents: empirical evidence from German star scientists. *Papers in Regional Science*, 89(2), 275-294.
- Schmidheiny, K., & Siegloch, S. (2023). On event studies and distributed-lags in two-way fixed effects models

- Schumpeter, J. (1939). *Business cycles: A theoretical, historical and statistical analysis of the capitalist process*. McGraw-hill.
- Science Foundation Ireland. 2021. Shaping our Future: Delivering for Today, Preparing for Tomorrow, Science Foundation Ireland Strategy 2025, Dublin.
- Simcoe, T. S., & Waguespack, D. M. (2011). Status, quality, and attention: What's in a (missing) name?. *Management Science*, 57(2), 274-290.
- Solow, R. M. (1956). A contribution to the theory of economic growth. *The quarterly journal of economics*, 70(1), 65-94.
- Sørensen, J. B., & Stuart, T. E. (2000). Aging, obsolescence, and organizational innovation. *Administrative science quarterly*, 45(1), 81-112.
- Sosa, M. E. (2011). Where Do Creative Interactions Come From? The Role of Tie Content and Social Networks. *Organization Science*, 22(1), 1–21.
- Storper, M., & Venables, A. J. (2004). Buzz: Face-to-face contact and the urban economy. *Journal of Economic Geography*, 4(4), 351–370.
- Sun, L. and Abraham, S., 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), pp.175-199.
- Sun, L., & Shapiro, J. M. (2022). A linear panel model with heterogeneous coefficients and variation in exposure. *Journal of Economic Perspectives*, 36(4), 193-204.
- Swan, T. W. (1956). Economic growth and capital accumulation. *Economic record*, 32(2), 334-361
- Tartari, V., Di Lorenzo, F., & Campbell, B. A. (2020). "Another roof, another proof": The impact of mobility on individual productivity in science. *The Journal of Technology Transfer*, 45(1), 276-303.
- Topel, R. H., & Ward, M. P. (1992). Job mobility and the careers of young men. *The Quarterly Journal of Economics*, 107(2), 439-479.
- Torre, A., & Rallet, A. (2005). Proximity and localization. *Regional Studies*, 39(1), 47–59.
- Tortoriello, M., McEvily, B., & Krackhardt, D. (2015). Being a Catalyst of Innovation: The Role of Knowledge Diversity and Network Closure. *Organization Science*, 26(2), 423–438.
- Trajtenberg, M., Shiff, G., & Melamed, R. (2006). The "names game": Harnessing inventors' patent data for economic research.
- Trippl, M. (2013). Islands of innovation as magnetic centres of star scientists? Empirical evidence on spatial concentration and mobility patterns. *Regional Studies*, 47(2), 229-244.
- Trippl, M., & Maier, G. (2010). Star scientists as drivers of the development of regions. In *Innovation, growth and competitiveness: Dynamic regions in the knowledge-based world economy* (pp. 113-134). Berlin, Heidelberg: Springer Berlin Heidelberg
- Tushman, M. L. (1977). Special boundary roles in the innovation process. *Administrative Science Quarterly*, 22(4), 587–605.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4), 327–352.
- Usher, A. P. (1929). *A history of mechanical inventions*.

- Uzzi, B. (1996). The Sources and Consequences of Embeddedness for the Economic Performance of Organizations: The Network Effect. *American Sociological Review*, 61(4), 674–698.
- Waldinger, F. (2012). Peer effects in science: Evidence from the dismissal of scientists in Nazi Germany. *The review of economic studies*, 79(2), 838-861.
- Waldinger, F., 2010. Quality matters: The expulsion of professors and the consequences for PhD student outcomes in Nazi Germany. *Journal of political economy*, 118(4), pp. 787-831.
- Walsh, J. P. (1988). Selectivity and selective perception: An investigation of managers' belief structures and information processing. *Academy of Management journal*, 31(4), 873-896.
- Walsh, J., & Cohen, W. (2008). Real Impediments to Biomedical Research. *Innovation Policy and the Economy*, 8, 1-30.
- Waltman, L., & van Eck, N. J. (2012). A new methodology for constructing a publication-level classification system of science. *Journal of the American Society for Information Science and Technology*, 63(12), 2378–2392.
- Weitzman, M. L. (1998). Recombinant growth. *The Quarterly Journal of Economics*, 113(2), 331–360.
- Wolfers, J., 2006. Did unilateral divorce laws raise divorce rates? A reconciliation and new results. *American Economic Review*, 96(5), pp. 1802-1820.
- Yadav, A., McHale, J., & O'Neill, S. (2023). How does co-authoring with a star affect scientists' productivity? Evidence from small open economies. *Research Policy*, 52(1), 104660.
- Zucker, L. G., & Darby, M. R. (1996). Star scientists and institutional transformation: Patterns of invention and innovation in the formation of the biotechnology industry. *Proceedings of the National Academy of Sciences*, 93(23), 12709-12716.
- Zucker, L. G., & Darby, M. R. (2001). Capturing technological opportunity via Japan's star scientists: Evidence from Japanese firms' biotech patents and products. *The journal of Technology transfer*, 26(1-2), 37-58.
- Zucker, L. G., & Darby, M. R. (2006). Movement of star scientists and engineers and high-tech firm entry.
- Zucker, L. G., Darby, M. R., & Armstrong, J. S. (2002). Commercializing knowledge: University science, knowledge capture, and firm performance in biotechnology. *Management science*, 48(1), 138-153.
- Zucker, L. G., Darby, M. R., Furner, J., Liu, R. C., & Ma, H. (2007). Minerva unbound: Knowledge stocks, knowledge flows and new knowledge production. *Research policy*, 36(6), 850-863.
- Zucker, L.G., Darby, M.R. and Armstrong, J., 1998. Geographically localized knowledge: spillovers or markets? *Economic Inquiry*, 36(1), pp. 65-86.