

COGNITIVE DISTANCE, SPILLOVERS, AND COMPLEMENTARITY

Evidence from the Deaths of Eminent Life Scientists[★]

Wei Si*

Qi Wang[†]

July 2026

Abstract

This study examines how peer effects, as measured by productivity changes following a shock to human capital, vary with the intellectual structure of scientific partnerships. We measure this structure using a new index of “cognitive distance” between collaborators, derived from their publication and citation patterns. To identify causal effects, we use a difference-in-differences approach that exploits the unexpected deaths of eminent life scientists as a natural experiment. Our analysis reveals two distinct patterns following the loss of a key collaborator. The production of high-quality output declines significantly only for cognitively close partners, consistent with a disruption of knowledge spillovers. In contrast, the decline in research quantity is U-shaped in cognitive distance: significant for both cognitively close and cognitively distant partnerships, and smallest at intermediate distance. The distant-tail decline replicates across alternative bibliometric distance measures. Overall, our findings show that the causal impact of losing a collaborator varies systematically with cognitive distance: losing a cognitively close partner is most detrimental to research quality, while moderate cognitive distance best balances the benefits of knowledge spillovers and complementarities.

Keywords: Cognitive Distance, Peer Effects, Scientific Collaboration, Knowledge Spillovers, Complementarity

JEL Classification: O31, O32, J24, D83

[★] We thank Peter Fredriksson, Andreas Madestam, Anna Sjögren, Arizo Karimi, and the participants in the seminars and conferences held at Stockholm University, Stockholm School of Economics, Tilburg University, ShanghaiTech University, the Econometric Society’s World Congress 2020, and the 2023 Asian Meeting of the Econometric Society for their insightful comments. We also thank Ludo Waltman and the Center for Science and Technology Studies (CWTS) at Leiden University for providing the data used in this study. The financial support provided by the Shanghai Pujiang Talent Program (2020PJC082) and the National Natural Science Foundation of China (No.72303153) is gratefully acknowledged. The usual disclaimer applies.

*Corresponding author: ShanghaiTech University; email: siwei@shanghaitech.edu.cn

[†]KTH-Royal Institute of Technology and Nordea

1 Introduction

The creation of new knowledge is a fundamental driver of economic growth, fueled by the positive externalities and knowledge spillovers generated by human capital (Lucas, 1988; Romer, 1990). While the importance of knowledge is clear, the mechanisms of its production are less well understood. In recent decades, scientific research has shifted from solo work to collaborative teams (Adams et al., 2005; Ductor, 2015; Fortunato et al., 2018; Leahey, 2016; Lee and Bozeman, 2005; Wuchty et al., 2007), a shift partly driven by the increasing “burden of knowledge,” which forces researchers to specialize narrowly and to collaborate to combine complementary expertise (Freeman et al., 2015; Jones, 2009). This raises a central question for both the organization of science and innovation policy: how does the degree of dissimilarity in expertise between collaborators affect the productivity of individual researchers?

This study uses a natural experiment to examine the causal effect of collaborators’ cognitive distance on individual research productivity. We define cognitive distance as horizontal differences in research expertise across fields, distinct from vertical differences in ability or seniority. To measure cognitive distance, we construct a new, field-agnostic index based on the distribution of researchers’ publications and citation links across academic journals. The life sciences provide an ideal setting for this analysis: basic research in this field is widely recognized as a driver of welfare and growth (Akçigit et al., 2021; Pinto and Teixeira, 2020), and its strong tradition of collaboration, combined with comprehensive bibliometric coverage, supports rich coauthorship analysis. Cognitive distance thus provides a tractable construct for studying how knowledge spillovers from similarity and expertise complementarities from diversity shape collaborative outcomes.

Because scientific collaborations are not formed at random, causal analysis faces challenges from selection bias and endogeneity. To address these concerns, we exploit the unexpected and premature deaths of researchers as a natural experiment, following prior studies that use such events to generate exogenous variation (Aizenman and Kletzer, 2011; Azoulay et al., 2010, 2019; Jaravel et al., 2018; Khanna, 2021, 2023; Mohnen, 2022; Oettl, 2012). These deaths provide an exogenous shock to the productivity of surviving coauthors that is plausibly unrelated to the cognitive distance of their partnerships. We identify 63 eminent life scientists who were actively engaged in research but died unexpectedly between 1995 and 2008. We construct a panel dataset of 9,263 coauthor-deceased scientist pairs, using publication and citation records from the Web of Science covering the period 1980 to 2013. This dataset allows us to track changes in coauthors’ annual publication rates and citation impacts before and after the deaths of their collaborators. We estimate effects using difference-in-differences, with heterogeneity-robust diagnostics for staggered treatment.

The effect of collaborator loss depends strongly on cognitive distance. Following the death of a star scientist, the quantity loss is U-shaped in cognitive distance: cognitively close and cognitively distant coauthors both experience significant declines, while coauthors at interme-

mediate cognitive distance are relatively resilient. The effect on research quality, by contrast, is concentrated among cognitively close collaborators and is persistent. Both patterns hold under heterogeneity-robust estimators for staggered treatment, and the distant-tail quantity decline further replicates across alternative bibliometric distance measures. Taken together, the patterns suggest two distinct channels of collaborative production. A spillover channel, rooted in shared expertise and tacit knowledge transfer, maps onto the exploitation side of organizational learning (March, 1991; Nooteboom et al., 2007; Wuyts et al., 2005) and is the dominant input to high-quality output. A complementarity channel, rooted in distinct, hard-to-substitute expertise, maps onto the exploration side and primarily sustains research quantity.

This study makes three primary contributions to the economics of science and the organization of scientific collaboration. First, we document that the productivity loss from a collaborator’s death follows two distinct patterns rather than one: the quantity loss is U-shaped in cognitive distance, while the quality loss is monotonic and concentrated among cognitively close coauthors. Three features of the design make this asymmetry detectable. Our cognitive-distance measure is continuous, field-agnostic, and defined at the researcher-pair level, the unit at which the death shock varies. Four outcomes separate quantity from quality. A continuous-treatment specification and percentile-rank event studies trace the moderation across the entire distance distribution, and the estimated pattern then guides the categorical specification. The pooled post-death decline itself is common ground with the death-shock literature. The moderation structure is new. Prior work has moderated the shock either by coarse proximity groups applied to a single quality-adjusted outcome, within which quantity and quality cannot diverge (Azoulay et al., 2010), or by network position and star characteristics (Azoulay et al., 2019; Khanna, 2021; Mohnen, 2022; Oettl, 2012). We show that this dyadic dimension, rather than network position or star characteristics alone, organizes the magnitude of productivity losses.

Second, we propose a two-channel decomposition of the spillover-complementarity tradeoff at the dyadic level. The tradeoff originates in the organizational-learning literature on exploration and exploitation (March, 1991), its canonical empirical tests use inter-firm alliances (Nooteboom et al., 2007; Wuyts et al., 2005), and subsequent work has operated at the firm, team, or industry level (Bikard et al., 2015; Larivière et al., 2015; Leahey et al., 2017; Mello and Rentsch, 2015; Teodoridis, 2018; Tzabbar and Kehoe, 2014; Wang et al., 2017), in each case aggregating interpersonal spillovers and complementarities rather than identifying them. Our decomposition expresses the tradeoff as a spillover term that decreases with cognitive distance and a complementarity term that increases with it (Belkhouja and Yoon, 2018; Lee et al., 2015; Uzzi et al., 2013), with channel weights that differ across outcomes. It yields two testable predictions: a monotone quality loss under a spillover-dominated regime and a U-shaped quantity loss under a spillover-leaning regime. The unexpected death of a collaborator dissolves the partnership exogenously, allowing both predictions to be tested within a causal design at the unit where interpersonal exchange occurs.

Third, we introduce a continuous, field-agnostic, researcher-pair-level cognitive distance index built from journal-citation flows over the entire Web of Science catalog. Prior measures of intellectual proximity have used coarse field classifications (Borjas and Doran, 2012; Fafchamps et al., 2010; Jaffe, 1986), biomedical-specific MeSH overlap (Azoulay et al., 2010), reference overlap (Ahuja and Katila, 2001), or text-based distance (Blei et al., 2003; Kelly et al., 2021), each constraining either cross-discipline portability or dyadic identification. Our construction takes the other side of that tradeoff: it gives up article-level granularity in exchange for cross-discipline portability and dyadic identification, as Table 1 in Section 3 makes explicit. Three alternative cognitive-distance constructs (a Web of Science–restricted reference cosine, a soft cosine over reference vectors, and a direct citation overlap) reproduce the distant-tail productivity decline, supporting the robustness of the index to measurement choice. The index is portable across disciplines and applicable to any setting where researcher-level publication and citation data are available.

The remainder of this paper is organized as follows. Section 2 develops the conceptual framework. Section 3 introduces our measure of cognitive distance and contrasts it with alternative proximity measures. Section 4 presents the empirical strategy and specifications. Section 5 introduces the data and sample construction. Section 6 reports the empirical results, and Section 7 concludes.

2 Conceptual Framework

2.1 The Cognitive Similarity-Diversity Tradeoff

Cognitive distance is the dissimilarity in research expertise between collaborators. It creates a central tradeoff. Cognitive similarity promotes efficient communication and knowledge transfer: a shared intellectual foundation enhances absorptive capacity and allows collaborators to integrate complex ideas and benefit from knowledge spillovers (Kaplan and Vakili, 2015; Nooteboom et al., 2007; O'Connor et al., 2021; Teodoridis et al., 2019; Wu et al., 2019). This explains why many scientists form partnerships with peers whose research expertise overlaps with their own (Leahey et al., 2017). Cognitive diversity, by contrast, supplies expertise complementarity: partners with distinct skills broaden the problem-solving scope, enable novel recombinations of ideas, and foster radical innovation (Belkhouja and Yoon, 2018; Bikard et al., 2015; Freeman et al., 2015; Lee et al., 2015; Singh and Fleming, 2010).

Diversity also carries costs. Greater distance leads to communication frictions, higher coordination costs, and difficulty integrating contributions (Boschma, 2005; van der Wouden and Youn, 2023; Wu et al., 2019). Prior studies of the tradeoff typically work at the firm or industry level; we study the same tension between individual collaborators, the smallest unit at which spillovers and complementarities operate through direct interpersonal exchange. This dyadic resolution allows us to isolate how cognitive distance shapes peer effects when a

collaboration unexpectedly dissolves.

2.2 A Two-Channel Decomposition

The tradeoff above implies that the loss from a star’s death flows through two channels whose mix differs across outcomes. We formalize this intuition with a stylized two-channel decomposition that adapts the two-force structure of [Wuyts et al. \(2005\)](#) to a setting with quality and quantity outcomes, combining the forces additively rather than multiplicatively. The model is reduced-form: it yields sign and shape predictions for the cognitive-distance gradient of productivity losses but does not pin down the values of its underlying parameters.

Consider a surviving coauthor i and a deceased star j at cognitive distance $d \equiv d_{ij} \in [0, 1]$. The pre-death contribution of j to i ’s output operates through two monotonic channels:

$$S(d) = (1 - d)^\alpha, \quad \alpha > 0, \quad (1)$$

$$C(d) = d^\beta, \quad \beta > 0, \quad (2)$$

where $S(d)$ is the spillover channel through which j transmits tacit knowledge ([Cohen and Levinthal, 1990](#); [Nooteboom et al., 2007](#)) and is decreasing in d , and $C(d)$ is the complementarity channel through which j bridges to distant expertise ([Boschma, 2005](#); [Singh and Fleming, 2010](#)) and is increasing in d . The curvature parameters α and β govern how steeply each channel attenuates away from its peak. We maintain $\alpha > 1$ and $\beta > 1$ throughout. This strict-curvature condition makes each channel convex in d , so that each channel’s contribution is concentrated near its own peak. Together with the loading assumptions introduced below, the condition is sufficient for the comparative statics in [Appendix C](#).¹

When j dies, i loses the contribution from each channel. The total loss for outcome y combines the two channels with outcome-specific weights:

$$L_y(d) = w_y^S S(d) + w_y^C C(d), \quad (3)$$

where w_y^S and w_y^C are the spillover and complementarity loadings on outcome $y \in \{q, Q\}$. For research quality ($y = Q$), the spillover channel strongly dominates, $w_Q^S \gg w_Q^C$ (Assumption A1 in [Appendix C](#)), because high-impact research requires deep tacit understanding of the star’s frontier expertise. For research quantity ($y = q$), the spillover channel modestly dominates, $w_q^S > w_q^C > 0$ (Assumption A2 in [Appendix C](#)): routine co-production draws on shared expertise, while a smaller share requires complementary specialist input from distant expertise.

Under Assumption A1, quality loss reduces to $L_Q(d) \approx w_Q^S S(d)$ and is monotonically

¹The power-function parameterization is chosen for analytical convenience. [Appendix C.6](#) shows that the qualitative predictions hold under alternative functional forms (e.g., exponential), so the framework should be read as a family of two-channel decompositions sharing the spillover-complementarity structure rather than as a single specific model.

decreasing in d : close coauthors had the most spillover-intensive partnerships and bear the largest decline, while distant coauthors had only weak spillovers to begin with and bear the smallest.

Quantity loss combines a spillover term $w_q^S S(d)$ that is largest at $d = 0$ and vanishes at $d = 1$, with a complementarity term $w_q^C C(d)$ that is zero at $d = 0$ and largest at $d = 1$. Under the curvature condition $\alpha, \beta > 1$ and the strictly positive loadings in Assumption A2, L_q is strictly convex on $(0, 1)$ with a unique interior minimum d^* , producing a U-shape: loss is highest at both tails and lowest at d^* .

The theory pins down the location of the minimum only partially. In the symmetric baseline $\alpha = \beta$ and $w_q^S = w_q^C$, it sits at $d^* = 1/2$. Holding the curvatures equal, the spillover-leaning regime $w_q^S > w_q^C$ from Assumption A2 shifts d^* rightward of $1/2$, because balancing the heavier spillover loss requires a larger d . Beyond this rightward shift, the position of d^* relative to the three empirical distance bands depends on all four parameters $(\alpha, \beta, w_q^S, w_q^C)$ and is not determined by the theory. The band-level implication is conditional on that position: the loss is smallest in the band nearest d^* and larger in the bands farther from it. Hypothesis 2 states this pattern with the middle band as the resilient one, a location consistent with the rightward shift implied by Assumption A2, and Section 6 estimates where d^* falls. Appendix C reports the sufficient conditions, derivations, and formal characterization of d^* .

The framework yields two testable predictions.

Hypothesis 1 (Quality, monotonic decline). *Following the death of a star coauthor, the productivity loss in research quality decreases monotonically in cognitive distance; cognitively close coauthors experience the largest decline.*

Hypothesis 2 (Quantity, U-shape). *Following the death of a star coauthor, the productivity loss in research quantity is U-shaped in cognitive distance; cognitively close and cognitively distant coauthors both experience significant declines, while coauthors at intermediate cognitive distance are relatively resilient.*

The additive structure in Equation (3) is not the only way to combine the two channels, and the leading alternative makes the opposite prediction about the middle of the distance distribution. In the multiplicative tradition of Nooteboom et al. (2007) and Wuyts et al. (2005), partnership value is the product of novelty and absorbability and peaks at intermediate cognitive distance; the loss from dissolution would then be largest for intermediate-distance pairs, an inverted U, with close and distant pairs comparatively resilient. The data can therefore discriminate between the two structures: the resilient middle band and the monotone quality decline documented in Section 6 are consistent with the additive-convex decomposition and inconsistent with the multiplicative benchmark in this setting.

The two channels parallel the exploitation-exploration distinction in organizational learning (March, 1991; Nooteboom et al., 2007; Wuyts et al., 2005): the spillover channel embodies exploitation through deep refinement of shared knowledge, and the complementarity chan-

nel embodies exploration through novel recombination across distant expertise. Spillover-dominated quality loss reflects exploitation alone; the U-shape in quantity loss reflects both modes, with the close tail tracing exploitation losses and the distant tail tracing exploration losses. We treat the empirical patterns in Section 6 as consistent with the framework rather than as a direct test of it.²

3 Measuring Cognitive Distance

3.1 Comparison with Alternative Proximity Measures

The typology of Boschma (2005) distinguishes five proximity dimensions: cognitive, organizational, social, institutional, and geographical. We focus on the cognitive dimension, within which four empirical implementations recur: idea-space overlap on shared MeSH terms in biomedicine (Azoulay et al., 2010), reference-overlap measures over cited works (Ahuja and Katila, 2001), text-based distance over abstracts or full text (Blei et al., 2003; Kelly et al., 2021), and field-code overlap based on assigned classifications (Borjas and Doran, 2012; Fafchamps et al., 2010; Jaffe, 1986). Table 1 contrasts our journal-citation construction with these four alternatives, together with the geographical and social proximity dimensions that are common in studies of scientific collaboration. The organizational and institutional dimensions do not enter the comparison set: researcher pairs are not embedded in formal organizational hierarchies, and any time-invariant institutional differences within a pair are absorbed by the dyad fixed effects in our identification strategy.³

[Table 1 about here]

Our index operationalizes the conceptual cognitive distance of Section 2. Table 1 shows that no empirical implementation dominates the others: each occupies a different point on the tradeoff between article-level granularity on one side and cross-discipline portability with dyadic identification on the other. Our construction takes the side that the death-shock design requires. We build the index from journal citation flows because journal positioning encodes institutionalized knowledge structures that are stable across years and portable across fields.

²The framework is falsifiable within a defined domain. Its predictions concern how the losses at the three distance bands compare with one another, the object that the pooled categorical specification (Equation (5)) estimates as contrasts relative to the middle band. Two observable patterns would falsify it. First, a middle-worst quantity pattern, in which the middle band loses strictly more than both the close and distant bands. No non-negative loading combination of the two channels can generate this pattern under the maintained curvature condition, because strict convexity of L_q places the middle-band loss strictly below a weighted average of the two tail losses. This is also precisely the pattern the multiplicative benchmark above predicts, so the condition doubles as the discriminating test between the two structures. Second, a U-shaped quality loss in cognitive distance, which would contradict (A1). The band-specific estimates in Table A.4 lie outside this domain: they measure each band's absolute within-band trajectory rather than the cross-band comparison, so their positive distant-band coefficient is not a violation on its face. Section 6.3 shows in addition that the positive coefficient does not survive a count-data specification (Table A.10), attributing it to the linear estimator.

³The institutional dimension instead enters the analysis as a binary moderator, through the same-institution split of Section 6.3.

3.2 Index Construction

The construction of the index involves three steps. First, we create a citation matrix $[c_{ab}]$ that records direct citation relationships between any two journals active from 1980 to 2013 in the Web of Science (WoS) database.⁴ For each citing journal (*Journal a*) and cited journal (*Journal b*), c_{ab} denotes the total number of citations of *Journal b* by *Journal a* over the entire period 1980–2013.

In the second step, we convert the journal citation matrix into a journal-by-journal distance matrix $[d_{ab}]$. We calculate the similarity between every pair of journals by applying the cosine similarity measure, a standard metric from computer science used to assess the relatedness of documents (Salton and McGill, 1983), to their citation vectors from the previous step.⁵ This yields a similarity matrix $[s_{ab}]$, from which the distance matrix is derived. The similarity between *Journal a* and *Journal b* is defined as

$$s_{ab} = \frac{\sum_k c_{ak}c_{bk}}{\sqrt{(\sum_k c_{ak}^2)(\sum_k c_{bk}^2)}},$$

where k indexes all journals included in the database. The distance between *Journal a* and *Journal b* is then defined as

$$d_{ab} = 1 - s_{ab}.$$

Each element d_{ab} of the distance matrix $[d_{ab}]$ ranges between 0 and 1, with lower values indicating greater similarity and higher values indicating greater distance.

The third step then computes the cognitive distance between a researcher i and a deceased collaborator j as a weighted average of journal distances, based on their publication distributions prior to the year of j 's death. Let n_i and n_j denote the total number of papers published by i and j ; n_i^a and n_j^b the number of papers they published in journals a and b , respectively; and d_{ab} the journal distance from the second step. The index of cognitive distance between researchers i and j is defined as

$$Distance_{ij} = \frac{\sum_a \sum_b n_i^a n_j^b d_{ab}}{n_i n_j}.$$

This formula computes the average journal distance across all pairwise combinations of publications by i and j , weighted by their respective publication distributions. The index

⁴Our dataset is limited to this period because complete WoS publication and citation records are available only for 1980–2013.

⁵The cosine similarity measure is broadly employed in computer science and related fields and has also been applied in bibliometric and economic research. For instance, some studies have used cosine similarity to quantify the distance between journals (Leydesdorff and Rafols, 2011; Rafols et al., 2012), while others have applied it to develop an index of research overlap among economists using JEL codes (Fafchamps et al., 2010).

ranges from 0 (if every pairwise combination of their journals has perfectly similar citation profiles, as when both researchers publish only in the same journal) to 1 (if they publish in entirely unrelated journals whose citation profiles share no common cited journals). A potential concern is that frequent co-publications could mechanically reduce the measured distance even when collaborators have distinct research backgrounds. To address this concern, our primary specification uses a “leave-out” version of the index, which excludes all joint publications when calculating the distance for each dyad. The same logic appears in [Azoulay et al. \(2010\)](#), whose keyword-overlap measure is computed on the pair’s nonjoint publications. The non-leave-out version is reported in [Table A.5](#) as a robustness check.

[Figure 1](#) shows the distribution of the cognitive distance index for the analysis sample, calculated between each deceased scientist and their coauthors using the procedure described above. The distribution is unimodal, apart from a spike in the lowest bin, and slightly left-skewed, with a mean of approximately 0.62 and a median of 0.64.⁶

[[Figure 1](#) about here]

3.3 Properties

The index has three properties that make it well-suited to dyadic productivity analysis. First, by aggregating journal-level citation patterns rather than article-level content, it captures institutionalized knowledge structures that persist across periods. In fast-moving fields such as the life sciences, papers have relatively short citation half-lives ([Bohannon, 2013](#); [Davis and Cochran, 2015](#)), so researchers active in different periods may appear intellectually distant on article-level measures even when they share a common focus. A molecular biologist publishing in *Cell* and a geneticist publishing in *Nature Genetics* decades later may rarely cite one another directly, yet strong citation ties between their journals signal shared cognitive foundations.⁷ Second, the broad multidisciplinary coverage of the WoS catalog makes the index portable across disciplines, applicable wherever researcher-level publication and citation data are available. Third, the construction reduces to three reproducible steps (citation matrix, cosine-similarity distance matrix, publication-weighted aggregation), making the index straightforward to audit and to apply across studies of scientific collaboration.

⁶The bar corresponding to the lowest range of cognitive distances (0 to 0.2) shows a notably higher count than adjacent bins. This anomaly is driven by data from one deceased scientist, but including or excluding this case does not affect the results.

⁷A potential concern is that broad, highly central journals such as *Science* and *Nature* publish across many disciplines ([Larivière and Gingras, 2010](#)), and that citations to and from such outlets may convey weaker disciplinary signals than citations between specialized journals. We view this as unlikely to substantially distort our measure for two reasons. First, publications in these multidisciplinary outlets represent a small share of a typical author’s output in our sample: the majority of each author’s publications appear in more specialized, field-specific journals. Second, in the cosine-similarity step each generalist outlet enters as one element of a high-dimensional citation vector, so its influence on the resulting distance is bounded by its citation frequency rather than by its disciplinary breadth.

4 Empirical Strategy

4.1 Identification Strategy

Scientific collaboration is not random. Partnerships form and dissolve through deliberate choices, so productivity comparisons across pairs, or before and after a chosen separation, conflate the contribution of the collaboration with the selection that formed and ended it. We address this selection problem by exploiting the unexpected and premature death of an active life scientist as an exogenous shock to collaboration: the timing of such a death is plausibly unrelated to the evolution of the pair’s research. This event allows us to estimate how the resulting productivity changes vary with the surviving coauthor’s cognitive distance from the deceased.

Our identification strategy directly extends [Azoulay et al. \(2010\)](#), who exploit the unexpected deaths of eminent life scientists and estimate a 5 to 8 percent decline in the citation-weighted publication count of surviving collaborators. Two elements of that design are common ground: the shock construction (sudden, premature deaths of actively publishing stars) and the use of publication-count and citation-impact outcomes. The moderation analysis is where the designs part. [Azoulay et al. \(2010\)](#) interact the death indicator with a binary indicator for the top quartile of MeSH keyword overlap and note that a four-quartile version lacked the statistical power to distinguish the quartile coefficients. Our continuous pair-level moderator and the event studies in Section 6 are built to trace the gradient that a binary contrast cannot. Subsequent applications have largely preserved the original death-shock template ([Aizenman and Kletzer, 2011](#); [Azoulay et al., 2019](#); [Jaravel et al., 2018](#); [Khanna, 2021, 2023](#); [Mohnen, 2022](#); [Oettl, 2012](#)). A broader literature documents peer-effect declines following the forced dismissal of scientists in Nazi Germany ([Waldinger, 2010, 2012](#)) and the loss of international knowledge networks during World War I ([Iaria et al., 2018](#)), as well as mixed peer effects from inflows of immigrant scientists ([Borjas and Doran, 2012, 2015](#); [Moser et al., 2014](#)). These mixed findings suggest that peer effects depend on the relationship between the scientists involved, the heterogeneity that our moderation analysis measures along the cognitive dimension.

We implement this approach using a generalized difference-in-differences framework in which the “treatment intensity” for each coauthor is their cognitive distance from the deceased scientist, treated as a continuous variable ([Wooldridge, 2010](#)). In practice, our model compares changes in productivity before and after the death of a focal scientist across coauthors with different levels of cognitive distance.

This design identifies the moderator effect within the treated dyads rather than the level effect of “ever losing a coauthor.” We adopt the within-treated specification because cognitive distance is defined relative to a deceased star coauthor. Adding a never-treated control arm would therefore require an arbitrary modeling choice for the moderator (for instance, distance to the closest treated dyad’s deceased star, or to a hypothetical reference profile) that the data

cannot discipline. The closest precedent for this design choice comes from [Azoulay et al. \(2010\)](#), who compare their estimates with and without a coarsened-exact-matching never-treated control arm and report that “the magnitudes of the estimates presented below are very similar whether or not control scientists are added to the estimation sample” (p. 558). That comparison concerns the average level effect of losing a coauthor rather than a moderation gradient. Identification of our moderator estimand instead rests on within-unit time variation: the imputation diagnostic of [Borusyak et al. \(2024\)](#) in [Table A.1](#) constructs counterfactual trajectories from the not-yet-treated observations of the eventually treated dyads, sidestepping the cross-unit comparison concern.

The validity of our causal interpretation rests on the key identification assumption of parallel trends: absent the collaborator’s death, the productivity of coauthors with varying cognitive distance would have evolved in parallel. We assess the plausibility of this assumption in [Section 6](#) by graphically examining the pre-event trends in our event study analysis. The staggered timing of the deaths across 1995–2008 provides a further safeguard: the year fixed effects absorb calendar-time shocks common to all scientists, and because the post-death window differs across dyads, no such shock aligns systematically with the death event in a way that could be mistaken for a moderated effect.

4.2 Empirical Models

Our baseline empirical model is as follows:

$$Y_{ijt} = \exp[\beta_0 + \beta_1 \text{Death}_{jt} + \beta_2 \text{Distance}_{ij} \times \text{Death}_{jt} + f(\text{Career_Age}_{it}) + \gamma_t + \delta_{ij}] \epsilon_{ijt}, \quad (4)$$

where the dependent variable Y_{ijt} is a measure of research output in year t of coauthor i , who previously collaborated with the now deceased scientist j . The key variables of interest are Death_{jt} , an indicator equal to 1 for all years following the death of scientist j , and Distance_{ij} , our index of cognitive distance. β_1 captures the main effect of a collaborator’s death at the baseline, where cognitive distance is zero, while the interaction coefficient β_2 measures how this effect varies with the level of cognitive distance. The main effect of Distance_{ij} does not enter separately because it is time-invariant and therefore absorbed by the dyad fixed effects. The function $f(\text{Career_Age}_{it})$ flexibly controls for the career age of scientist i , which captures lifecycle factors such as tenure incentives and changes in productivity over the career trajectory.⁸ Year fixed effects, γ_t , absorb time shocks common to all scientists. Dyad fixed effects, δ_{ij} , absorb all time-invariant characteristics of the coauthor, the star, and the unique coauthor–deceased scientist pair, holding constant any vertical differences between partners, such as their respective levels of seniority or innate talent, and thereby isolating the moderating

⁸Since information on the year of PhD completion is not available, we define career age as the difference between the publication year t and the year of scientist i ’s first publication. Career age cohort dummies are constructed in three-year intervals.

role of horizontal cognitive distance on peer effects. ϵ_{ijt} is a multiplicative error term with conditional mean one, so the exponential index specifies the conditional mean of Y_{ijt} .

We measure research output along two dimensions: quantity and quality. Our primary quantity measure is the annual number of journal articles published by scientist i . As a first step toward capturing research quality, we also weight each publication by its journal’s impact factor (JIF), thereby constructing a measure of JIF-weighted publications.⁹

Recognizing that JIF is a crude proxy for article-level quality, we supplement it with more precise citation-based measures of research impact. Absolute citation counts, however, are affected by truncation, since earlier publications have had more time to accumulate citations. To correct for this bias, we normalize citations within each publication year and field. Specifically, we compute the citation distribution for all life sciences papers in our dataset published between 1980 and 2013, grouped by broad WoS subject categories. For each researcher and year, we then count the number of publications that rank above the 50th percentile and above the 85th percentile of these field-year distributions. These normalized indicators capture different aspects of research impact. The number of publications above the 50th percentile captures a researcher’s ability to produce consistently solid, above-average work, serving as a broad measure of impactful quantity. In contrast, the number of publications above the 85th percentile identifies “home runs” and provides a more stringent measure of research quality and breakthrough impact.

The interaction term in [Equation \(4\)](#) restricts the moderation to be linear in cognitive distance within the exponential index. To detect potential nonlinearity, we first inspect the relationship visually in [Figure 2](#). The figure plots post-death changes in four measures of research output against the percentile rank of cognitive distance: the number of publications, JIF-weighted publications, publications above the 50th citation percentile, and publications above the 85th percentile. The patterns reveal clear nonlinearities. Guided by these plots, we classify coauthors into three groups: cognitively close coauthors in the 0–40th percentiles, middle-distance coauthors in the 40–85th percentiles, and cognitively distant coauthors in the 85–100th percentiles. [Figure A.2](#) in the appendix presents the same coefficient estimates at a finer 5-percentile resolution, where the 85th percentile cutoff coincides with a bin boundary; the broad nonlinear pattern is consistent across both resolutions, providing additional support for our chosen categorization.¹⁰

[[Figure 2](#) about here]

Given this evidence of a nonlinear relationship, we also estimate the model in [Equation \(5\)](#)

⁹The impact factor of an academic journal measures the average number of citations received by articles published in that journal over a specific period. We use the two-year JIF for each journal to calculate the number of JIF-weighted publications.

¹⁰The results are robust to alternative cutoff choices. For instance, shifting the boundary for cognitively close coauthors from the 40th percentile to the lower third, 45th, or 50th percentile, or redefining the cognitively distant group as the upper third, 75th, or 90th percentile, does not alter the results.

based on the three cognitive distance groups defined above.

$$Y_{ijt} = \exp[\theta_0 + \theta_1 Death_{jt} + \theta_2 Close_{ij} \times Death_{jt} + \theta_3 Distant_{ij} \times Death_{jt} + f(Career_Age_{it}) + \gamma_t + \delta_{ij}] \epsilon_{ijt}. \quad (5)$$

The indicator variable $Close_{ij}$ equals one if coauthor i is cognitively close to deceased scientist j , and $Distant_{ij}$ equals one if i is a cognitively distant coauthor of scientist j . The middle-distance coauthor group is the omitted reference category. As with $Distance_{ij}$, the group indicators are time-invariant and therefore enter only through their interactions with $Death_{jt}$. The other variables in Equation (5) retain the same definitions as in Equation (4). The coefficient θ_1 measures the post-death productivity change for middle-distance coauthors, while θ_2 and θ_3 capture the differential effects for close and distant coauthors, respectively, relative to their middle-distance peers. This specification allows for a direct comparison of the shock's impact across the three cognitive distance groups.

Consistent with standard practice in studies of R&D and scientific productivity, and considering the nonnegative and highly skewed nature of our dependent variables, we estimate the models using the fixed effects Poisson quasi-maximum likelihood estimator (QMLE) developed by Hausman et al. (1984). The estimator requires only that the conditional mean in Equation (4) be correctly specified: it delivers consistent parameter estimates even when the dependent variables are not true counts (Silva and Tenreyro, 2006), and inference based on robust standard errors remains valid under departures from the Poisson distribution and under arbitrary forms of serial correlation (Wooldridge, 1999). We cluster the standard errors at the level of the deceased scientist.

Recent advances in the staggered difference-in-differences literature show that two-way fixed effects event-study coefficients can be biased weighted averages of cohort-specific treatment effects, with negative weights and forbidden comparisons of later-treated against earlier-treated units when treatment effects are heterogeneous (Callaway et al., 2024; Roth et al., 2023). Because the deaths in our sample occur at different dates and every dyad is eventually treated, these concerns apply directly to our design. We address them with three heterogeneity-robust event studies on the pooled sample and a within-band check.

The imputation estimator of Borusyak et al. (2024) (henceforth BJS) delivers a heterogeneity-robust event study for the binary-event version of the design (Table A.1). The continuous-treatment estimator of de Chaisemartin and D'Haultfœuille (2026) (henceforth dCDH) adapts this logic to the continuous-intensity specification in Equation (4), using $Distance_{ij} \times Death_{jt}$ as the treatment intensity (Table A.2). The interaction-weighted estimator of Sun and Abraham (2021) is reported alongside a Poisson event-study sensitivity (Correia et al., 2020) (Table A.3); the Poisson variant confirms that the post-event dynamics are not artifacts of the linear functional form. BJS event studies estimated separately on the three cognitive-distance bands (Table A.4) avoid the strong-parallel-trends assumption that any continuous-treatment speci-

fication imposes across distance values (Callaway et al., 2024). The results of these checks are reported in Section 6.2, and the band-specific pattern is taken up in Section 6.3.

5 Data and Sample Construction

5.1 Data

We collected the data for our empirical analysis from the Web of Science database, which is one of the most comprehensive multidisciplinary research platforms. WoS provides access to major citation databases and bibliometric information on scientific publications worldwide. Our analysis focuses on the life sciences, broadly defined to encompass 30 WoS subject categories such as biology, biomedical sciences, microbiology, and medicine, and we restrict the sample to publications in these categories.¹¹ Our dataset covers the period from 1980 through 2013, which reflects the window of complete WoS records available to us. To focus on original, peer-reviewed research, we include only articles published in academic journals; we accordingly exclude non-article document types such as letters, meeting abstracts, and editorial materials.

5.2 Sample of Deceased Scientists

A fundamental step in our analysis is identifying a sample of active life scientists who passed away unexpectedly and prematurely. We limit the year of death to the period between 1995 and 2008 to observe changes in the productivity of their coauthors within a sufficient time window before and after these events. Moreover, our goal is to identify star scientists who had established distinguished research records. This approach is consistent with prior work, which suggests that the impact on coauthors of losing a less prominent researcher may be less detectable than that of losing a star (Azoulay et al., 2010, 2019; Khanna, 2021; Mohnen, 2022; Oettl, 2012).

Our search for these scientists therefore began by examining trends in annual publication counts. Our guiding assumption is that after the sudden death of a preeminent researcher, their annual publication count will decrease sharply, eventually reaching zero. However, this decline is not always instantaneous, as articles published posthumously likely represent unfinished studies that were already in the pipeline at the time of death. To identify potential star scientists fitting this pattern, we first selected researchers who had produced more than 50 total publications and then identified those whose annual publication trends exhibited a sharp and permanent decline over a two-year period. We defined this decline as an annual count of more than 20 publications at the start of the two-year window and fewer than three thereafter. Figure B.1 illustrates an example of such a publication trend.

¹¹See Table B.1 in Appendix B for the full list of subject categories included in the life sciences sample.

Using these criteria, we identified an initial sample of 102 researchers in the life sciences. However, we recognized that other factors could also lead to a similar publication decline, such as retirement or departure from academia to practice medicine, an event that is not uncommon among researchers with an M.D. degree. Therefore, to distinguish sudden, premature deaths from other career exits, we manually verified each case in this initial sample. Our verification process involved searching for obituaries, personnel information from their departments, personal websites, and “In Memoriam” articles from journals in the WoS database. We then excluded researchers who retired, left academia, or passed away at an advanced age or in a manner that was not sudden, as well as those for whom definitive information was unavailable. Our final sample consists of 63 active and preeminent life scientists who died unexpectedly and prematurely, with an exact year of death between 1995 and 2008.

[Table 2](#) reports summary statistics for this sample of deceased scientists: year of death, years of first and last publication, career age (measured as the difference between the year of death and the year of first publication), total publications, cumulative citations, citations per career year, and the number of distinct coauthors.

[[Table 2](#) about here]

5.3 Sample of Coauthors

Our primary focus is on the coauthors of the deceased scientists. We define these coauthors by the following criteria. First, a coauthor must be part of the one-degree coauthorship network of a deceased star, meaning that they published at least one paper with one of the 63 scientists from 1980 onward. Second, we include only researchers who also have publications independent of the deceased scientist, either as a solo author or with other collaborators. This criterion is intended to exclude individuals who may have been graduate students of the deceased, which helps distinguish faculty peer effects from mentorship impacts, given that authorship norms in the life sciences frequently include students on papers. Applying these criteria yields a sample of 9,006 coauthors.

The fundamental unit of observation in our study is the coauthor-star dyad, which represents a unique collaborative relationship. Since some of the 9,006 coauthors in our sample collaborated with more than one star scientist, we identify 9,263 unique coauthor-star dyads. Our panel analysis tracks these dyads from 1980 to 2013, making the coauthor-star-year triad the effective unit of analysis. The resulting estimation sample is an unbalanced panel with 155,294 coauthor-star-year observations. The Poisson estimation samples for the two citation-percentile outcomes are slightly smaller, at 154,895 and 150,256, because the conditional fixed-effects likelihood drops dyads that record no positive outcome in any year. Dyads enter the panel in the coauthor’s first publication year and continue through the 2013 data cutoff irrespective of the star’s death year; we do not impose a fixed post-death window such as ten years. Within these bounds, the panel is unbalanced because coauthors have heterogeneous

entry years and because years in which a coauthor records no Web of Science publication do not yield observations. The panel is therefore publication-driven: a year with no recorded publication contributes no row rather than a zero-count row. To verify that this construction does not drive the results, [Table A.9](#) re-estimates the headline specifications on a rectangular panel in which every interior zero-output year between a coauthor’s first and last observed publication year is filled in explicitly. [Section 6.2](#) reports the results.

[Table 3](#) presents summary statistics for the full coauthor sample and for subsamples based on cognitive distance. Unless otherwise noted, all covariates are measured in the year of the star scientist’s death. The average cognitive distance is 0.62 for the full sample, with mean values of 0.46, 0.69, and 0.86 for the close, middle, and distant coauthor groups, respectively.

[[Table 3](#) about here]

Our measures of research outcomes include four indicators: raw and JIF-weighted publication counts, and the number of publications ranking above the 50th and 85th citation percentiles for their publication year. We also report cumulative measures for each outcome.

We construct several variables to characterize the coauthors and their collaborative relationships. A coauthor’s career age is defined as the difference between the year of the star’s death and the coauthor’s first year of publication. Based on this measure, we define a senior coauthor as a researcher whose career age is greater than that of their deceased partner. To capture a coauthor’s scientific standing, we calculate their cumulative citations up to the year of the star’s death and define an elite coauthor as one whose cumulative citations are in the top decile of the full coauthor sample. We also measure the intensity and recency of each partnership. We count the number of collaborations for each dyad from 1980 to 2013; as a histogram of these counts reveals a highly skewed distribution ([Figure A.1](#)), we create a categorical variable to capture high-intensity ties. We define a regular coauthor as one who collaborated more than five times, and a recent coauthor as one who published with the star within the three years before the star’s death. Finally, we create indicators for the organizational and structural context, identifying whether partners were affiliated with the same institution and whether the deceased star was consistently the first or last author on their joint publications. As shown in [Table 3](#), these covariates are broadly comparable across the close, middle, and distant coauthor groups, suggesting the groups are well balanced and providing support for our research design.

A comparison of the summary statistics in [Table 2](#) (deceased scientists) and [Table 3](#) (coauthors) confirms the preeminence of the former group, revealing substantial differences in average total publications, cumulative citations, and citations per career year.

6 Results

6.1 Main Results

This section presents the results from our main empirical models, specified in [Equation \(4\)](#) and [Equation \(5\)](#). [Table 4](#) presents the estimated effects of a star scientist’s death on coauthor productivity across four outcome variables: (i) raw publication counts, (ii) JIF-weighted publications, (iii) publications above the 50th citation percentile, and (iv) publications above the 85th citation percentile. For each outcome, we estimate models using both a continuous measure of cognitive distance and categorical groups (close, middle, distant) to test for linear and nonlinear effects.

[[Table 4](#) about here]

Our findings reveal two distinct patterns. First, for outcomes that emphasize research quality, the negative impact of a star’s death is concentrated among cognitively close coauthors and is attenuated by distance. This is evident for JIF-weighted publications and publications above the 85th citation percentile. The continuous interaction models for these outcomes (columns (3) and (7)) show a negative main effect of death that is mitigated by increasing cognitive distance, with both coefficients significant at the five percent level or better for the JIF-weighted outcome and at the ten percent level for the top-85%ile outcome. The categorical models corroborate this finding. For JIF-weighted publications (column (4)), only the close-distance group experiences a statistically significant productivity decline of approximately 18.3%¹² relative to the middle-distance group. Similarly, for publications above the 85th percentile (column (8)), the close-distance group is the only one to show a significant decrease (a drop of approximately 10.6%).

Second, a different, nonlinear pattern emerges for outcomes related to the overall quantity of research. For these measures, we find a non-monotonic relationship in the pooled sample: both the closest and most distant collaborators are adversely affected, while those at intermediate cognitive distance are not. The continuous models hint at this nonlinearity: when the moderation is restricted to be linear, the interaction is statistically indistinguishable from zero for raw publication counts (column (1)) and only weakly positive for the top-50%ile outcome (column (5)), as opposing losses at the two tails largely cancel in a linear fit. For raw publication counts (column (2)), both the close-distance and distant-distance groups experience significant declines of 11.1% and 13.5%, respectively, relative to the middle-distance group, which shows no statistically significant change. A similar non-monotonic pattern appears for publications above the 50th percentile (column (6)), where both the close and distant groups

¹²As the models are Poisson specifications, the percentage change is calculated as $1 - \exp(\theta)$. For instance, this 18.3% decline is derived from the coefficient in column (4): $1 - \exp(-0.202) \approx 0.183$. This calculation applies to all subsequent percentage changes reported.

are hurt significantly more than their middle-distance peers, with declines of 13.3% and 9.8%, respectively.

Overall, the results show that the consequences of losing a star collaborator vary systematically with cognitive distance, though the specific pattern depends on the type of research output. For highly specialized, high-impact research (JIF-weighted publications and publications above the 85th percentile), the negative effects are concentrated among cognitively close collaborators, consistent with a disruption of knowledge spillovers. For broader measures of research output (raw publications and publications above the 50th percentile), the effect is nonlinear, suggesting that both cognitively close (spillover-dependent) and cognitively distant (complementarity-dependent) partnerships are vulnerable, while those at a medium distance appear most resilient. Cognitive distance is a pre-determined feature of the partnership rather than a randomly assigned treatment. Section 6.3 sets out what that implies for reading the distance profile and applies the same qualification to the other dyadic moderators.

To further explore the timing and persistence of these effects, we conduct an event study analysis, plotting the year-by-year coefficients in Figure 3. This specification interacts the close- and distant-distance group indicators with a full set of time dummies, relative to the middle-distance group and the year of the star's death ($t = 0$).

[Figure 3 about here]

First, the plots support the validity of our research design. In the five years preceding the event, the estimated coefficients for both the close- and distant-distance groups are statistically indistinguishable from zero for all four outcome variables, and the close-distance group shows no significant lead at any horizon. The only exceptions among the eighty plotted lead coefficients are three isolated distant-group leads for the top-85%ile outcome at long horizons (six, seven, and ten years before the event), plus one borderline distant-group lead for the JIF-weighted outcome at seven years ($p = 0.050$); none of these forms a trend running into the event. This absence of differential pre-event trends in the window where anticipation or divergence would surface supports the core assumption of our difference-in-differences model.

Second, after the star's death, the plots reveal the evolution of the treatment effects, which visually corroborates our main findings. For the quality-focused outcomes (JIF-weighted publications and publications above the 85th percentile), a sharp, significant, and persistent decline in productivity emerges only for the close-distance coauthors, while the distant-distance group shows no persistent statistically significant change. For the quantity-focused outcomes (raw publications and publications above the 50th percentile), the close-distance group suffers a lasting and significant drop in output relative to the middle group, and the distant-distance group displays a gap that is negative in every post-event year but reaches conventional significance only intermittently in the year-by-year estimates; the pooled contrasts in Table 4 provide the sharper summary of the distant-band loss.

Together, these results are in line with the predictions of the two-channel decomposition

introduced in Section 2.2. The quality results, summarized in columns (4) and (8) of Table 4 and the corresponding panels of Figure 3, are consistent with Hypothesis 1: only cognitively close coauthors experience a statistically significant decline in JIF-weighted publications and high-impact publication counts, while distant coauthors are largely unaffected. The continuous estimates in columns (3) and (7) corroborate this pattern in its most direct form: the negative death effect attenuates significantly and approximately linearly with cognitive distance, and the quadratic terms for these outcomes are statistically indistinguishable from zero.¹³ This monotonic, distance-decreasing pattern is consistent with quality output being spillover-dominated, sourced primarily through deep tacit knowledge transfer that is irreplaceable when the partnership ends (Azoulay et al., 2010; Cohen and Levinthal, 1990; Nooteboom et al., 2007; Wu et al., 2019).

The quantity results, in columns (2) and (6) of Table 4 and the corresponding panels of Figure 3, are consistent with Hypothesis 2: the loss is U-shaped in cognitive distance, with significant declines at both close and distant pairs and resilience at intermediate distance. This is the empirical counterpart of the U-shape in L_q that Proposition 2 in Appendix C describes. The middle-band null locates the interior minimum d^* near $d_{\text{middle}} \approx 0.69$, with the close-band loss sitting on the falling arm of the U and the distant-band loss on the rising arm. The null is precisely estimated rather than merely underpowered: the middle-band post-death changes in columns (2) and (6) are 0.037 and 0.030 (standard errors of 0.027), so the 95 percent confidence intervals exclude middle-band declines larger than roughly two percent.

A quadratic specification of the continuous interaction provides a descriptive check on this location, independent of the categorical cutoffs: the estimated turning point for unweighted publications is $\hat{d}^* = 0.70$ (delta-method 95 percent confidence interval [0.48, 0.92]), interior to the observed distance support of [0.12, 0.99] and coinciding with the middle band. The individual quadratic terms are imprecisely estimated, however, and the formal intersection-union test of an inverse-U effect profile (Sasabuchi, 1980; Lind and Mehlum, 2010) does not reject monotonicity at conventional levels ($p = 0.16$). The non-rejection reflects the imprecision of the quadratic terms rather than evidence against the U-shape.¹⁴ The categorical estimates of Table 4 therefore remain the primary evidence on the shape, with the turning-point estimate as corroborating description.

The band-specific BJS event studies in Table A.4 offer a complementary cut: estimated separately within each band, the absolute post-death effects order monotonically in distance, negative for the close band, null for the middle band on three of the four outcomes, and positive for the distant band. Section 6.3 shows that the positive distant-band coefficient does not survive a count-data specification (Table A.10) and attributes it to the linear estimator. Both arms of the pattern also survive the further scrutiny of Section 6.2: the close and distant

¹³The estimates from the quadratic specification are available upon request.

¹⁴The test requires the effect profile to slope significantly upward at the lower end of the distance support and significantly downward at the upper end. This joint requirement is demanding in our design, with standard errors clustered on 63 deceased scientists and the distant arm containing fifteen percent of the dyads.

quantity contrasts retain significance under wild-cluster bootstrap inference, and the distant-band decline in unweighted publications, the empirically novel arm of the U, is directionally preserved and statistically significant at the ten percent level or better under all three alternative cognitive-distance measures (Table A.6).

The monotonic quality decline in columns (4) and (8) and the U-shaped quantity loss in columns (2) and (6) together form the central asymmetry that the two-channel decomposition organizes. In the language of organizational learning, the U-shape is the dyadic counterpart of the firm-level exploitation-exploration tension (March, 1991; Nooteboom et al., 2007): cognitively close partnerships are exploitation-intensive and are vulnerable to the loss of shared tacit knowledge, cognitively distant partnerships are exploration-intensive and are vulnerable to the loss of irreplaceable complementary expertise, and intermediate-distance partnerships balance both modes.

The strength of these channel-specific losses, however, varies systematically across subgroups of coauthors. Section 6.3 examines how channel loadings shift with the surviving coauthor’s absorptive capacity, the geography of the partnership, the intensity of past collaboration, and the deceased star’s authorship role; Appendix C.5 provides the channel-weight reading of these patterns.

6.2 Robustness Checks

This section reports three groups of checks. We first re-estimate the headline specifications on four subsamples, then subject the design to three heterogeneity-robust event-study estimators, and finally report six further checks on variable measurement, functional form, pre-event productivity dynamics, panel construction, and inference.

To check the robustness of our main findings, we re-estimate our primary specifications on four different subsamples, with the results presented in Table 5 through Table 8. First, to isolate the shock to a coauthor’s individual productivity from the mechanical loss of the star’s contribution, we exclude all joint papers with the deceased scientist from the outcome variables and consider only the coauthor’s publications with other researchers or solo work (Table 5). Second, to ensure our results are not driven by individuals who experienced the loss of more than one prominent collaborator, we restrict the sample to coauthors linked to only a single star in our dataset (Table 6). Third, to ensure that the effects are not driven solely by the most elite superstars, we exclude coauthors of stars in the top decile of cumulative citations (Table 7). Finally, to ensure the results are not driven by network hub effects, we exclude coauthors of stars in the top decile of total coauthor count (Table 8). In each subsample the cognitive-distance bands are re-cut on that subsample’s own percentile distribution, so dyads near a band boundary can be classified differently than in Table 4. This affects at most 2.4 percent of dyads in any of the four subsamples, and no dyad moves between the close and the distant band: every reassignment is with the middle reference band.

[Table 5, Table 6, Table 7, and Table 8 about here]

Across all four subsamples, the results are highly consistent with our main findings, and both patterns identified above persist. For outcomes related to research quantity, the U-shaped pattern recurs in each subsample: the close contrast is significant in every table, the distant contrast for raw publication counts is significant in all four (at the ten percent level in Table 8), and the distant contrast for the top-50%ile outcome is significant at the ten percent level in Tables 5 and 7 while attenuating below conventional significance in Tables 6 and 8, with the point estimates remaining negative throughout. For outcomes related to research quality, the negative effect remains concentrated among cognitively close coauthors. These checks provide strong evidence that our results are not artifacts of specific sample compositions or the mechanical loss of a star’s contribution to joint papers.

We next report the results of the three heterogeneity-robust event studies introduced in Section 4.2. Figure 4 overlays the three pooled estimators across the four outcomes. The three estimators play different roles: the BJS imputation estimator is the primary heterogeneity-robust diagnostic for the binary-event version of the design, the dCDH estimator probes the continuous-treatment specification of Equation (4), and the Sun-Abraham estimator, reported alongside a Poisson variant, isolates the role of the linear functional form. In each case the validity check is the profile of pre-event placebo coefficients, the estimated leads for the years before the death, which should be statistically indistinguishable from zero if surviving coauthors were not already on differential trajectories before the event. All three estimators reproduce the pooled-sample pattern in sign and rough magnitude. The primary diagnostic passes without qualification, and the pre-event irregularities that arise in the two remaining estimators trace to their linear specifications, as the following paragraphs document.

Imputation. The BJS estimator delivers significantly negative ten-year coefficients across all four outcomes ($\tau_{10} = -1.49^{***}$ for unweighted publications, -5.10^{***} for JIF-weighted publications, -1.21^{***} for publications above the 50th citation percentile, and -0.30^{***} above the 85th), with all twenty pretrend coefficients statistically insignificant (Table A.1). Because the estimator imputes each dyad’s counterfactual from its own not-yet-treated observations, this check addresses the heterogeneity and forbidden-comparison concerns of Section 4.2 directly, and it passes in full: significant long-run declines on every outcome with no detectable pre-event drift.

Continuous treatment. The dCDH estimator returns average total effects that are negative and statistically significant on all four outcomes, with ten-year horizons that are similar in magnitude to the BJS values for the two quantity-focused outcomes and larger for the two quality-focused outcomes (Table A.2). Its placebo diagnostics are mixed: the placebo at $t = -1$ is statistically significant for the unweighted and top-50%ile outcomes, with two further leads marginal for the unweighted count, although the placebos are jointly insignificant in every column. If these leads reflected genuinely different pre-death trajectories across the distance bands, a specification that allows each band to follow its own nonparametric trend should

absorb them. The sensitivity reported in the note to [Table A.2](#) implements this specification, and the leads persist: the average effects are essentially unchanged, the significant $t = -1$ lead remains, and the $t = -3$ lead strengthens to significance at the five percent level for the same two outcomes. The leads are therefore structural to the linear continuous-treatment specification rather than evidence of differential pre-death trends across bands, and we read the dCDH results as corroborating the sign and rough magnitude of the BJS estimates rather than as an independent heterogeneity-robust diagnostic.

Functional-form sensitivity. The Sun-Abraham linear specification displayed in [Figure 4](#) returns positive pre-event coefficients that are statistically significant at $t = -2$ for all four outcomes and at additional leads for the unweighted and top-50%ile counts. Two explanations are possible: the linear functional form applied to count outcomes, or the estimator’s reliance on the latest-treated cohort as the comparison group. Two diagnostics point to the former. First, the Poisson re-estimation in column (5) of [Table A.3](#), available for the unweighted count, returns statistically insignificant pretrends with the same event-time structure. The band-specific estimates repeat the reversal for the JIF-weighted outcome: the positive, individually significant leads that the linear estimator returns for the close band in [Table A.4](#) turn negative and become jointly insignificant in the Poisson event studies of [Table A.10](#), with two leads still individually significant under the reversed sign ([Section 6.3](#)). A genuine pre-death trend is a feature of the data and cannot vanish when the link function changes; an artifact of the functional form can. Second, if the leads reflected the comparison against the latest-treated cohort, dropping that cohort should move the coefficients. A sensitivity that excludes the year-of-death 2008 cohort and re-estimates the specification with the new latest-treated cohort (year of death 2006, $n = 3,978$ comparison-cohort observations) instead leaves the $t = -2$ point estimates essentially unchanged, and three of the four leads remain statistically significant. The one exception, the unweighted count, loses significance through a wider standard error rather than through coefficient attenuation.¹⁵ The lead pattern is therefore an artifact of fitting a linear event study to skewed counts, not a feature of the design. Because the headline specifications of [Table 4](#) are Poisson rather than linear, this artifact cannot affect the main estimates.

[[Figure 4](#) about here]

We now turn to the six remaining checks.

Non-leave-out cognitive distance. We use an alternative “non-leave-out” measure of cognitive distance, which does not remove publications co-authored with the deceased star. The results in [Table A.5](#) remain consistent with our main findings, with three attenuations: the distant top-50%ile contrast falls below conventional significance, and the two close contrasts on the unweighted and top-85%ile outcomes weaken but remain statistically significant.

Alternative bibliometric measures. We replace the leave-out cognitive distance with three methodologically distinct bibliometric similarity measures built from our available data: a Web

¹⁵The estimates excluding the 2008 cohort are available upon request.

of Science—restricted reference cosine similarity, a soft cosine similarity over reference vectors, and a direct citation overlap measure (Table A.6). The reference cosine treats two researchers as similar when their publications draw on the same cited literature. The soft cosine additionally credits citations to related rather than identical sources. The direct citation measure, the coarsest of the three, records whether one researcher’s publications cite the other’s. The two cosine-family measures share the vector-similarity logic of our preferred index, with dyad-level Pearson correlations of 0.66 and 0.78 with the leave-out distance, whereas the direct-citation measure is the least correlated, at 0.12. For each alternative, we reassign the close, middle, and distant indicators using the same percentile cutoffs as the production specification and re-estimate Equation (5) with the production estimator and controls.

These alternatives serve as robustness checks rather than as candidate replacements for the production index, and the direction of their measurement error makes the check a conservative one. All three are built from the pair’s own citing behavior, so they inherit the two limitations that Table 1 attributes to reference-overlap measures: the choice of references is endogenous to publication strategy, and overlap between distant fields is mechanically near zero, which limits resolution precisely in the distant tail where our design needs it. The direct citation measure carries both limitations to their extreme, since whether one researcher cites the other is a relational choice that reflects social as well as cognitive proximity. Both limitations bias the reassigned contrasts toward zero, so patterns that survive under the alternatives are informative. Our index instead positions each researcher in the aggregate journal citation network (Section 3.1), which no single pair’s citing decisions can move, and it alone supports the leave-out construction of Section 3.2 that prevents the collaboration itself from compressing measured distance.

The distant-pair contrast is the robust feature of this check. The *Distant* $\times$ *Death* coefficient remains negative in every panel and outcome. It is statistically significant at the ten percent level or better for the unweighted publication count under all three measures and for the top-50%ile outcome under the two cosine-family measures. The JIF-weighted and top-85%ile contrasts are null throughout, as they are in the production specification. The distant-tail decline in research quantity, the empirically novel arm of the U-shaped pattern, therefore appears under every alternative measure, including the one least correlated with our preferred index.

The *Close* $\times$ *Death* coefficient attenuates under the alternatives, and a misclassification calibration shows that the attenuation is quantitatively what band reclassification implies. Bands assigned under an imperfectly correlated measure mix dyads from different production bands, pulling the estimated contrasts toward zero. Using only the Panel A coefficients as the data-generating process and the observed dyad-level overlap between production and alternative bands, the calibration predicts the close contrast that each alternative should deliver. All twelve observed close contrasts lie within 1.3 standard errors of these predictions, and nine lie within half a standard error. The distant contrasts, by comparison, exceed their predicted magnitudes in every panel. The observed attenuation is therefore the pattern the production

estimates themselves imply. The $Close \times Death$ coefficient is statistically null under the WoS-restricted reference cosine and direct-citation measures and marginally significant for one outcome under the soft cosine, and even the small positive close estimates under the direct-citation measure fit the calibration, since its near-uniform band overlap predicts contrasts of approximately zero. Panels B–D therefore support rather than contradict the production estimates: the close attenuation is what reclassification implies, and the distant decline is larger than reclassification alone can explain.

OLS sensitivity. We re-estimate the models using ordinary least squares instead of the Poisson specification. The OLS estimates in [Table A.7](#) reproduce the categorical significance pattern of [Table 4](#): the same six contrasts are significant at the ten percent level or better, and the two distant contrasts on the quality outcomes remain null. Our core conclusions do not depend on the choice of estimation method.

Mean reversion. The close-band attenuation in publication outcomes could in principle reflect regression to the mean rather than a substantive treatment effect. We re-estimate the headline specification on two restricted samples: an Active-Pair sample, comprising dyads with strictly positive pre-event total output for the relevant outcome, and a Top-Quartile sample, comprising dyads in the top quartile of pre-event mean output for the relevant outcome, computed among the active dyads. If mean reversion were driving the close-band effect, restricting to high-pre-period-output pairs would attenuate the $Close \times Death$ coefficient toward zero. Instead, the coefficient is stable or strengthens under both restrictions ([Table A.8](#)), alleviating the concern. [Figure A.3](#) plots the corresponding event-study dynamics for the Active-Pair sample.

Panel construction. As described in [Section 5](#), the estimation panel is publication-driven: coauthor-years with no recorded publication contribute no row, so the estimates condition on publication-active years. [Table A.9](#) re-estimates the headline specifications on a rectangular panel in which every interior zero-output year between a coauthor’s first and last observed publication year is filled in explicitly. All eight categorical contrasts of [Table 4](#) retain their signs with slightly larger magnitudes, the six contrasts significant in [Table 4](#) remain significant at least at the five percent level, and the two distant contrasts on the quality outcomes remain statistically null. A full rectangle through 2013, which also fills the years after each coauthor’s last observed publication, strengthens the estimates further.¹⁶ Precision falls only in the continuous specification. The distance interactions for the two citation-percentile outcomes, and the death main effect in the top-85%ile column, lose their ten percent significance in the rectangular panel, with point estimates little changed and standard errors roughly twenty percent wider. The share of filled zero years is balanced across the three distance bands (23, 20, and 25 percent for the close, middle, and distant bands) and slightly lower after the death year in all three, consistent with the stability of the estimates.

Few-cluster inference and influence. Standard errors throughout are clustered on the 63

¹⁶The full-rectangle estimates are available upon request.

deceased scientists, a cluster count at which conventional cluster-robust inference can over-reject (Cameron et al., 2008). Two corrections address this concern. First, wild-cluster bootstrap p -values, computed on the linear analogue of Equation (5), reproduce the significance pattern of Table 4 exactly. The close contrast remains significant on all four outcomes ($p = 0.015, 0.015, 0.001$, and 0.004 for the unweighted, JIF-weighted, top-50%ile, and top-85%ile outcomes), and the distant contrast remains significant on the two quantity outcomes ($p = 0.011$ and 0.044) and null on the two quality outcomes. Second, cluster-jackknife (CV3) standard errors (MacKinnon et al., 2023), computed from 63 re-estimates of the Poisson specification that each exclude one deceased scientist, confirm the same pattern with one exception: the distant top-50%ile contrast is significant at the ten rather than the five percent level ($p = 0.081$). The same leave-one-out estimates double as an influence check. All 63 re-estimates keep every focal contrast within a narrow band around the full-sample value, and no single exclusion overturns the significance of any close contrast (Figure A.5).

6.3 Heterogeneous Effects and the Underlying Mechanisms

This section proceeds in two parts. We first complete the band-level evidence, adjudicating the one result that appears to conflict with the pooled U-shape, and then examine how the losses vary across types of coauthors and collaborative relationships.

We begin with the band-specific event studies introduced in Section 4.2 and reported in full in Table A.4. Applied separately to the dyads in each cognitive-distance band, the BJS imputation estimator recovers ten-year coefficients that are strongly negative for cognitively close coauthors on all four outcomes and essentially null for middle-distance coauthors on three of the four. Both results are in line with the pooled evidence: the close-band decline is consistent with the spillover-dependent interpretation of Hypothesis 1 and the joint spillover-complementarity interpretation of Hypothesis 2, and the middle-band null is the direct counterpart of the resilient middle-distance group in Table 4 and Figure 3. The one discordant result is the distant band, where the linear estimator returns significantly positive long-run coefficients. This positivity is an artifact of the linear functional form rather than a feature of the data. Table A.10 re-estimates the band-specific event studies with a fixed-effects Poisson specification, the same count-data sensitivity applied to the pooled Sun-Abraham estimates in Table A.3. Under the count-data link, the distant-band ten-year coefficients are small, of mixed sign, and uniformly insignificant (-0.19 for unweighted publications, $+0.09$ for JIF-weighted publications, $+0.09$ above the 85th citation percentile, and -0.07 above the 50th), while the close-band decline persists, most clearly for the JIF-weighted outcome (-0.40 at ten years, significant at the five percent level). The Poisson specification also resolves the close-band pretrend anomaly. For the close band's JIF-weighted outcome, the linear estimator returns pre-event coefficients that are positive, rising toward the event, and individually significant at every lead. Under the count-data link they turn negative or vanish, and the joint pretrend test does not reject ($p = 0.19$), although two of the leads remain individually significant with the

sign now reversed. Anticipation cannot generate such a reversal, but functional form can.

The band-specific event studies and the pooled categorical specification estimate different objects. [Table 4](#) estimates each band’s average *proportional* moderation contrast relative to the resilient middle group, which is the profile the two-channel framework speaks to, while the event studies estimate each band’s *absolute* within-unit trajectory. Once the distant-band positivity is attributed to the linear estimator, the band-specific evidence no longer conflicts with the pooled U-shape: close coauthors decline, middle-distance coauthors are resilient, and distant coauthors show no long-run gain. [Appendix C.5](#) accordingly retains the within-field competition-relief mechanism, the one channel a genuine distant-band gain would have supported, only as a theoretical possibility that the evidence does not require.

With the role of cognitive distance itself established, we estimate our nonlinear model from [Equation \(5\)](#) on subsamples split by professional standing, institutional proximity, collaboration intensity, and the deceased’s authorship position, and plot the key interaction coefficients for the cognitively close and distant groups in [Figure 5](#). Each plotted coefficient is a contrast against middle-distance coauthors within the same subsample. A subgroup in which no coefficient is significant is therefore one in which cognitive distance does not moderate the response, not one in which the response is absent.¹⁷ In what follows, we call a subgroup estimate statistically significant when its 95 percent confidence interval, plotted in [Figure 5](#), excludes zero. The death provides plausibly exogenous variation in collaborator loss, but the dyadic moderators, cognitive distance included, are not randomly assigned. Subgroup differences in the estimated death effect therefore describe how the average response covaries with observable dyadic characteristics rather than identify causal moderating effects in the strict econometric sense reviewed by [Roth et al. \(2023\)](#). We treat the patterns below as descriptive evidence about the mechanisms suggested by the channel-weight framework of [Section 2.2](#).

[[Figure 5](#) about here]

Professional standing. We define “elite” coauthors as those in the top decile of cumulative citations at the time of the star’s death and “senior” coauthors as those with a greater career age than the deceased scientist. The two splits show the same pattern in [Figure 5](#). Every statistically significant contrast occurs among non-elite or junior coauthors, and not one of the sixteen estimates for their elite and senior counterparts is statistically significant on any outcome. Cognitive distance therefore moderates the response among less established coauthors and does not moderate it among their established counterparts.¹⁸ Two complementary mechanisms

¹⁷Prior work motivates each split. Superstar-extinction studies find that the loss from a star’s death varies with professional standing, collaboration intensity, and authorship position ([Azoulay et al., 2010, 2019; Oettl, 2012](#)), and the spatial-knowledge-flow literature shows that proximity shapes knowledge transmission between institutions ([Fernandez et al., 2016; Singh and Marx, 2013](#)). We build on this evidence by examining each dimension jointly with cognitive distance, testing the subgroup contrasts formally ([Tables A.11 to A.13](#)), and interpreting the same-institution close-pair pattern through resource competition.

¹⁸[Table A.13](#) reports a negative *Elite* $\times$ *Death* interaction, ranging from -0.118 to -0.189 in the fully joint specification and significant at the one percent level on all four outcomes. That coefficient is a level shift in the

can account for this asymmetry. First, knowledge spillovers and complementarities flow predominantly from more established scientists to less established colleagues, with weaker flows in the reverse direction (Azoulay et al., 2010; Borjas and Doran, 2015), so the two channels that make the loss sensitive to cognitive distance carried comparatively little of an elite or senior coauthor’s output. Second, established researchers possess extensive experience and diverse professional networks (Khanna, 2021; Mohnen, 2022), so their capacity to replace a lost source of expertise depends less on how cognitively close or distant that source was. In channel-weight terms (Section 2.2), this pattern is consistent with elevated weights w^S and w^C among junior and non-elite coauthors: lower own absorptive capacity (Cohen and Levinthal, 1990) and narrower external networks imply that a larger share of pre-event output flowed through the spillover and complementarity channels, which sharpens the distance profile that Figure 5 traces.

Institutional proximity. The institutional split shows the sharpest asymmetry in Figure 5. Among same-institution pairs, the distant-band decline is larger than anywhere else in the figure, ranging from -0.33 to -0.46 across the four outcomes and statistically significant on three of them. The close-band estimates move in the opposite direction. All four shift toward zero relative to the different-institution subsample, the estimate for publications above the 85th percentile turns positive, and none is statistically distinguishable from zero. Among different-institution pairs, the pattern of Section 6.1 broadly reappears, with the close band declining significantly on all four outcomes and the distant band on the unweighted count. We interpret the asymmetry through substitution-set logic. Same-institution close-cognitive coauthors share two scarce inputs, intellectual capital and institutional resources such as laboratory space, students, and internal grants. When one partner dies, the institutional resources are released to the survivor, and this relief can offset the spillover loss (Azoulay et al., 2010, 2019; Borjas and Doran, 2012, 2015; Teodoridis et al., 2019). Same-institution distant-cognitive coauthors face the opposite margin. Such matches are rare, and the search frictions that make distant complementarity hard to assemble also make replacement within the same campus difficult (Fernandez et al., 2016; Singh and Marx, 2013; Sosa and Maoret, 2023; van der Wouden and Youn, 2023), so the complementarity loss compounds. Different-institution pairs sit on neither margin. Their substitution pool is wider in both the cognitive and the geographic dimension, so neither force operates and the baseline contrast survives. In framework terms, the two same-institution tails move in opposite directions relative to the different-institution baseline, matching the resource-competition-relief mechanism formalized in Appendix C.5. The relief channel activates for cognitively close pairs, where the released resources were previously contested between the partners, while match rarity inflates the

average post-death response of elite relative to non-elite coauthors with the distance bands held fixed, a different object from the within-subgroup band contrasts plotted in Figure 5. Its inclusion leaves the close and distant contrasts essentially unchanged (columns 3 and 4 of each outcome block). Because the elite flag is defined by a realized performance stock (cumulative citations at the time of the star’s death), the level shift is also consistent with regression to the mean among coauthors selected on a high realization. The null *Senior* $\times$ *Death* interaction is in line with this reading: the senior flag selects on career age, not on realized performance. The sample restrictions of Table A.8 and Figure A.3 address such reversion for the band contrasts themselves.

complementarity weight w^C for cognitively distant pairs.

Collaboration intensity. This dimension has two margins: how recently and how often the pair collaborated. On the recency margin, Figure 5 shows that the losses are largest among “recent” collaborators, those who published with the star within the three years before the death. Within this group the pattern of Table 4 reappears in amplified form: the close-band decline is significant on all four outcomes, and the distant-band decline is confined to the quantity outcomes. “Former” collaborators are affected as well, though less severely. Their close-band declines are smaller but remain significant on three of the four outcomes, and their distant contrasts do not reach the five percent level on any outcome. The recency pattern is not a mechanical artifact of losing joint work in progress: excluding all joint publications with the deceased from the outcomes, as in Table 5, leaves the recent-pair estimates essentially unchanged.¹⁹ Active collaboration likely sustains the strongest dependency on the star’s input, a dependency that attenuates but does not vanish once the collaboration lapses. On the frequency margin, the split separates the losses on some margins but not others. “Regular” collaborators, those with more than five joint publications, show no significant distant-band contrast on any outcome and no significant close-band contrast on either citation-percentile measure, consistent with established routines absorbing the shock. Their close-band declines on the unweighted and JIF-weighted outcomes are the exception: -0.19 against -0.10 and -0.24 against -0.19 , measured against “infrequent” collaborators, with the latter the largest close-band estimate in the split. Routinization therefore buffers the margins for which established procedure can substitute, while the loss persists, and if anything deepens, where joint output was most tightly coupled to shared tacit knowledge. In channel-weight terms, active dependency sustains high contemporaneous values of both w^S and w^C for recent dyads, whereas routinization and skill redundancy attenuate the effective weights for regular dyads on the substitutable margins (Appendix C.5).

Authorship role. In the life sciences, authorship order typically signals a researcher’s contribution: the first author is often the main executor of the project, while the last author is the principal investigator who directs the laboratory and secures funding. We create indicators for dyads in which the star was consistently the first or the last author on the joint publications. In Figure 5, the significant declines concentrate in the complementary subsamples, those in which the star did not consistently occupy the role. Dyads with the star consistently first author show no significant contrast on any outcome, and dyads with the star consistently last author are largely resilient, with significant contrasts confined to two of the eight coefficient cells (the distant contrast for unweighted publications and the close contrast for JIF-weighted publications). Consistent first or last authorship likely marks a partnership embedded in a larger research program, in which organizational routines and skill redundancy buffer the surviving coauthor against the loss. A middle-author role, by contrast, marks a peer-to-peer relationship in which the coauthor depends on the star’s specific intellectual contributions. The

¹⁹The estimates excluding joint publications are available upon request.

resilience of coauthors whose partners were consistently the last author is consistent with the shock operating primarily through the loss of intellectual inputs rather than tangible resources such as funding or laboratory access. However, consistent last authorship is also correlated with the star’s seniority, and our categorical specification cannot separate leadership of a larger research program from greater average seniority. Both readings are in line with the muted response in this subsample, and we treat the distinction as a limit of the data rather than a settled interpretation. In channel-weight terms ([Appendix C.5](#)), the surviving coauthor’s pre-event dependency on the star’s specific intellectual input is plausibly smaller in first-author and last-author dyads, attenuating both w^S and w^C relative to peer-to-peer dyads.

To corroborate the visual contrasts in [Figure 5](#) with formal inference, we apply two sets of Wald tests to the subgroup estimates. The first set tests the equality $H_0 : \beta_{\text{close}} = \beta_{\text{distant}}$ within each subgroup; [Table A.11](#) reports the results, and [Figure A.4](#) plots the corresponding 95 percent confidence intervals. The test rejects the equality in two of the seven subgroups. Among same-institution pairs, the close-versus-distant gap is statistically significant at the one percent level for unweighted publications and at the five percent level for the two citation-percentile outcomes, consistent with the asymmetric institutional pattern discussed above. Among recent collaborators, the gap is significant at the five percent level on the JIF-weighted outcome ($p = 0.024$), where a significant close-band decline coexists with a distant-band null. The equality is not rejected even at the ten percent level in any other cell, although the point estimates retain the interpretive direction outlined above. Within most subgroups, then, the close and distant responses are statistically indistinguishable from each other. The heterogeneity documented above lies in how each band’s contrast with the middle group varies across subsamples, and institutional proximity is the one partition in which the two bands diverge across multiple outcomes.

The second set of Wald tests approaches the same question from the opposite direction. Rather than comparing the two bands within a subgroup, these tests ask whether a single band’s coefficient differs between the two sides of a partition. Each is estimated from a fully interacted specification in which every regressor is allowed to differ between the flagged subgroup and its complement, with standard errors still clustered on the deceased scientist. [Table A.12](#) reports all fifty-six resulting gaps, seven partitions by four outcomes for each of the two bands.

The institutional split is the only partition that rejects at the five percent level, and it rejects with the sign pattern that the substitution-set reading predicts. The distant coefficient is more negative among same-institution pairs on all four outcomes, the direction match rarity predicts, significantly so at the five percent level for the unweighted and JIF-weighted outcomes ($p = 0.014$ and 0.043) and at the ten percent level for the two citation-percentile outcomes ($p = 0.073$ and 0.058). The close coefficient is less negative on the two citation-percentile outcomes ($p = 0.052$ and 0.032), the direction the relief channel predicts. Of the forty-eight gaps outside the institutional split, one falls below the ten percent threshold: the

close coefficient in the deceased-last-author split on the unweighted outcome ($p = 0.076$).

On their own, three rejections at the five percent level out of fifty-six tests would be unremarkable. What separates this pattern from noise is that the three do not scatter across the seven partitions. All fall in the institutional split, and each carries the sign its mechanism predicts. The differences in statistical significance across subsamples that are visible in [Figure 5](#) for the standing and authorship-role splits, by contrast, do not correspond to coefficients that are statistically distinguishable at the five percent level.

Finally, we address the concern that cognitive distance might proxy for the dyadic correlates examined above rather than operate as an independent moderator. We re-estimate Equation (5) with the seven moderator-by-Death interactions added in three nested layers: tie strength (regular and recent collaboration), professional standing (elite and senior coauthors), and roles (same-institution, deceased first author, deceased last author). [Table A.13](#) reports the results. Across all four outcomes, both contrast coefficients are stable between the baseline (column 1) and the fully joint specification (column 4), the largest movement being 0.009. The close-by-Death coefficient retains significance at the five percent level or better in every cell. The distant-by-Death coefficient retains its negative sign throughout and preserves the [Table 4](#) profile in every specification: significant at the one percent level in every column for the unweighted count, significant at the five or ten percent level in every column for the top-50-percentile outcome, and statistically null throughout for the JIF-weighted and top-85-percentile outcomes. The moderating role of cognitive distance is therefore not absorbed by the seven correlated dyadic dimensions. It operates as an independent moderator of the death-shock response, in line with the channel-weight framework of [Section 2.2](#).

Taken together, the heterogeneity patterns fall under the two interpretive mappings formalized in [Appendix C.5](#). The absorptive-capacity, collaboration-intensity, and authorship-role patterns map onto shifts in the channel weights w^S and w^C : professional standing appears to determine the direction of knowledge flows, collaboration intensity reflects the need for active engagement to sustain both channels, and authorship role tracks how much the coauthor depended on the star's specific intellectual input. The institutional split engages both mappings. The disappearance of the same-institution close-pair decline, which turns into a positive point estimate for the most stringent quality measure, is captured by the competition-relief activation mode, while match rarity inflates the complementarity weight w^C for same-institution distant pairs. The framework is one parsimonious interpretation, not the only one. The same patterns could also reflect variation in the deceased scientists' attributes across subgroups, or selection of coauthors into subgroups on dimensions correlated with absorptive capacity. The value of the decomposition is to organize the heterogeneity dimensions under a unified framework rather than to identify the underlying channel weights.

7 Discussion and Conclusion

Star scientists' unexpected deaths produce two distinct productivity patterns in surviving coauthors, each tracing a different channel of collaborative production. First, the loss of a cognitively close partner severely damages the production of high-impact research, a finding that is both statistically and economically significant. This suggests that knowledge spillovers, which depend on a shared intellectual foundation, are a central input for high-quality research. Second, for broader measures of research quantity, we find a non-monotonic relationship in the pooled sample in which both cognitively close and distant coauthors are adversely affected while coauthors at intermediate cognitive distance are resilient. The distant-tail decline replicates across alternative bibliometric distance measures. This pattern points to the importance of complementary expertise from distant partners in broadening the scope of research and enabling novel combinations of ideas.

The U-shaped quantity loss in cognitive distance is the dyadic counterpart of the firm-level exploitation-exploration tension in organizational learning (March, 1991; Nooteboom et al., 2007). Extreme cognitive closeness, which facilitates exploitation of existing knowledge, creates a dependency on spillovers that is vulnerable to disruption. Extreme cognitive distance, which facilitates exploration, creates a reliance on a single source of complementary knowledge that is difficult to substitute. Collaborations at a moderate cognitive distance appear to strike a balance, benefiting from both mechanisms while mitigating the vulnerabilities of either extreme. This balance allows for effective knowledge exchange without excessive redundancy and supports innovation without prohibitive communication barriers.

Our heterogeneity analyses further indicate how the professional and organizational context of a partnership moderates these mechanisms. The sensitivity of the loss to cognitive distance is confined to less established coauthors, consistent with knowledge spillovers and complementarities flowing primarily from elite and senior researchers to their less established colleagues. Physical proximity cuts in opposite directions at the two tails: the release of contested institutional resources can offset the close pairs' spillover loss, while the rarity of cognitively distant matches within the same institution deepens the complementarity loss. The recency of collaboration suggests that the benefits of both spillovers and complementarities require active engagement to be sustained, while established routines buffer regular collaborators only on the margins where procedure can substitute for the lost partner's input. Finally, the muted response in dyads where the star consistently held the first or the last author role points to partnerships embedded in larger research programs, in which the surviving coauthor depends less on the star's specific intellectual input.

These findings have important implications for science policy and the management of research organizations. They caution against a "one-size-fits-all" approach to partnership formation. Assembling groups of cognitively similar researchers may be effective for producing high-impact work within an established paradigm, but doing so creates fragile dependencies.

Conversely, fostering highly diverse partnerships may increase the breadth of research but leave these partnerships vulnerable to the loss of unique, complementary members. Our results suggest that a portfolio approach to collaboration may be most effective, with a particular emphasis on moderately distant partnerships as a source of organizational resilience.

Our study has several limitations that also suggest avenues for future research. First, its focus is exclusively on the life sciences. While this field provides a rich context for studying collaboration, its norms may not be representative of all research domains, limiting the external validity of our findings. Future research should therefore examine whether these patterns generalize to other contexts, such as less collaborative fields (for instance, the humanities) or nonacademic settings (for instance, corporate R&D). Second, the unexpected death of a collaborator is a particularly severe shock, and future work could explore responses to less severe shocks, such as a collaborator moving to another institution. Finally, we deploy the cognitive distance index only in the setting of team dissolution. The index is broadly applicable, however, and could be used to study not only the dissolution of teams but also their formation, the diffusion of knowledge across different fields, and the structure of collaboration within and between industry and academia.

By introducing a new, generalizable measure of cognitive distance into a natural experiment framework, this paper shows that the causal impact of losing a collaborator is not uniform: it varies systematically with the intellectual architecture of the partnership. Accounting for this intellectual dimension is therefore essential for researchers, organizations, and policymakers who seek to build collaborations that remain productive after the loss of key members.

References

- Adams, J. D., Black, G. C., Clemmons, J. R., and Stephan, P. E. (2005). Scientific teams and institutional collaborations: Evidence from US universities, 1981–1999. *Research Policy*, 34(3):259–285.
- Agrawal, A., Kapur, D., and McHale, J. (2008). How do spatial and social proximity influence knowledge flows? Evidence from patent data. *Journal of Urban Economics*, 64(2):258–269.
- Ahuja, G. and Katila, R. (2001). Technological acquisitions and the innovation performance of acquiring firms: A longitudinal study. *Strategic Management Journal*, 22(3):197–220.
- Aizenman, J. and Kletzer, K. (2011). The life cycle of scholars and papers in economics — the ‘citation death tax’. *Applied Economics*, 43(27):4135–4148.
- Akcigit, U., Hanley, D., and Serrano-Velarde, N. (2021). Back to basics: Basic research spillovers, innovation policy, and growth. *The Review of Economic Studies*, 88(1):1–43.
- Azoulay, P., Fons-Rosen, C., and Graff Zivin, J. S. (2019). Does science advance one funeral at a time? *American Economic Review*, 109(8):2889–2920.
- Azoulay, P., Zivin, J. S. G., and Wang, J. (2010). Superstar extinction. *The Quarterly Journal of Economics*, 125(2):549–589.
- Belkhouja, M. and Yoon, H. D. (2018). How does openness influence the impact of a scholar’s research? An analysis of business scholars’ citations over their careers. *Research Policy*, 47(10):2037–2047.
- Bikard, M., Murray, F., and Gans, J. S. (2015). Exploring trade-offs in the organization of scientific work: Collaboration and scientific reward. *Management Science*, 61(7):1473–1495.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.
- Bohannon, J. (2013). The Secret Half-Lives of Scientific Papers. *Science*, ScienceInsider, <https://www.science.org/content/article/secret-half-lives-scientific-papers>.
- Borjas, G. J. and Doran, K. B. (2012). The collapse of the Soviet Union and the productivity of American mathematicians. *The Quarterly Journal of Economics*, 127(3):1143–1203.
- Borjas, G. J. and Doran, K. B. (2015). Which peers matter? The relative impacts of collaborators, colleagues, and competitors. *Review of Economics and Statistics*, 97(5):1104–1117.
- Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *The Review of Economic Studies*, 91(6):3253–3285.
- Boschma, R. (2005). Proximity and innovation: A critical assessment. *Regional Studies*, 39(1):61–74.
- Callaway, B., Goodman-Bacon, A., and Sant’Anna, P. H. C. (2024). Difference-in-Differences with a Continuous Treatment. NBER Working Paper 32117, National Bureau of Economic Research.
- Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3):414–427.
- Cohen, W. M. and Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1):128–152.
- Correia, S., Guimarães, P., and Zylkin, T. (2020). Fast Poisson estimation with high-dimensional fixed effects. *The Stata Journal*, 20(1):95–115.
- Davis, P. M. and Cochran, A. (2015). Cited half-life of the journal literature. *arXiv preprint arXiv:1504.07479*.
- de Chaisemartin, C. and D’Haultfœuille, X. (2026). Difference-in-Differences Estimators of Intertemporal Treatment Effects. *The Review of Economics and Statistics*, pages 1–18.
- Ductor, L. (2015). Does co-authorship lead to higher academic productivity? *Oxford Bulletin of Economics and Statistics*, 77(3):385–407.
- Fafchamps, M., Leij, M. J., and Goyal, S. (2010). Matching and network effects. *Journal of the European Economic Association*, 8(1):203–231.

- Fernandez, A., Ferrándiz, E., and León, M. D. (2016). Proximity dimensions and scientific collaboration among academic institutions in europe: The closer, the better? *Scientometrics*, 106:1073–1092.
- Fortunato, S., Bergstrom, C. T., Börner, K., Evans, J. A., Helbing, D., Milojević, S., Petersen, A. M., Radicchi, F., Sinatra, R., Uzzi, B., Vespignani, A., Waltman, L., Wang, D., and Barabási, A.-L. (2018). Science of science. *Science*, 359(6379):eaao0185.
- Freeman, R. B., Ganguli, I., and Murciano-Goroff, R. (2015). Why and Wherefore of Increased Scientific Collaboration. In Jaffe, A. B. and Jones, B. F., editors, *The Changing Frontier: Rethinking Science and Innovation Policy*, pages 17–48. University of Chicago Press eBooks.
- Hausman, J., Hall, B. H., and Griliches, Z. (1984). Econometric models for count data with an application to the Patents-R&D relationship. *Econometrica*, 52(4):909–38.
- Iaria, A., Schwarz, C., and Waldinger, F. (2018). Frontier knowledge and scientific production: Evidence from the collapse of international science. *The Quarterly Journal of Economics*, 133(2):927–991.
- Jaffe, A. B. (1986). Technological opportunity and spillovers of R&D: Evidence from firms’ patents, profits, and market value. *The American Economic Review*, 76(5):984–1001.
- Jaravel, X., Petkova, N., and Bell, A. (2018). Team-specific capital and innovation. *American Economic Review*, 108(4-5):1034–1073.
- 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):283–317.
- Kaplan, S. and Vakili, K. (2015). The double-edged sword of recombination in breakthrough innovation. *Strategic Management Journal*, 36(10):1435–1457.
- Kelly, B., Papanikolaou, D., Seru, A., and Taddy, M. (2021). Measuring technological innovation over the long run. *American Economic Review: Insights*, 3(3):303–320.
- Khanna, R. (2021). Aftermath of a tragedy: A star’s death and coauthors’ subsequent productivity. *Research Policy*, 50(2):104159.
- Khanna, R. (2023). Passing the torch of knowledge: Star death, collaborative ties, and knowledge creation. *Research Policy*, 52(1):104649.
- Larivière, V. and Gingras, Y. (2010). On the relationship between interdisciplinarity and scientific impact. *Journal of the American Society for Information Science and Technology*, 61(1):126–131.
- Larivière, V., Haustein, S., and Börner, K. (2015). Long-distance interdisciplinarity leads to higher scientific impact. *Plos One*, 10(3):e0122565.
- Leahey, E. (2016). From sole investigator to team scientist: Trends in the practice and study of research collaboration. *Annual Review of Sociology*, 42:81–100.
- Leahey, E., Beckman, C. M., and Stanko, T. L. (2017). Prominent but less productive: The impact of interdisciplinarity on scientists’ research. *Administrative Science Quarterly*, 62(1):105–139.
- Lee, S. and Bozeman, B. (2005). The impact of research collaboration on scientific productivity. *Social Studies of Science*, 35(5):673–702.
- Lee, Y.-N., Walsh, J. P., and Wang, J. (2015). Creativity in scientific teams: Unpacking novelty and impact. *Research Policy*, 44(3):684–697.
- Leydesdorff, L. and Rafols, I. (2011). Indicators of the interdisciplinarity of journals: Diversity, centrality, and citations. *Journal of Informetrics*, 5(1):87–100.
- Lind, J. T. and Mehlum, H. (2010). With or without U? The appropriate test for a U-shaped relationship. *Oxford Bulletin of Economics and Statistics*, 72(1):109–118.
- Lucas, R. E. (1988). On the mechanics of economic development. *Journal of Monetary Economics*, 22(1):3–42.
- MacKinnon, J. G., Nielsen, M. Ø., and Webb, M. D. (2023). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics*, 232(2):272–299.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1):71–87.

- Mello, A. L. and Rentsch, J. R. (2015). Cognitive diversity in teams: A multidisciplinary review. *Small Group Research*, 46(6):623–658.
- Mohnen, M. (2022). Stars and brokers: Knowledge spillovers among medical scientists. *Management Science*, 68(4):2513–2532.
- Moser, P., Voena, A., and Waldinger, F. (2014). German Jewish émigrés and US invention. *American Economic Review*, 104(10):3222–3255.
- Nooteboom, B., Van Haverbeke, W., Duysters, G., Gilsing, V., and Van den Oord, A. (2007). Optimal cognitive distance and absorptive capacity. *Research Policy*, 36(7):1016–1034.
- O'Connor, M., Doran, J., and McCarthy, N. (2021). Cognitive proximity and innovation performance: Are collaborators equal? *European Journal of Innovation Management*, 24(3):637–654.
- Oettl, A. (2012). Reconceptualizing stars: Scientist helpfulness and peer performance. *Management Science*, 58(6):1122–1140.
- Pinto, T. and Teixeira, A. A. (2020). The Impact of Research Output on Economic Growth by Fields of Science: a Dynamic Panel Data Analysis, 1980–2016. *Scientometrics*, 123(2):945–978.
- Rafols, I., Leydesdorff, L., O'Hare, A., Nightingale, P., and Stirling, A. (2012). How journal rankings can suppress interdisciplinary research: A comparison between innovation studies and business & management. *Research Policy*, 41(7):1262–1282.
- Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5):S71–S102.
- Roth, J., Sant'Anna, P. H. C., Bilinski, A., and Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2):2218–2244.
- Salton, G. and McGill, M. (1983). *Introduction to Modern Information Retrieval*. McGraw-Hill computer science series. McGraw-Hill.
- Sasabuchi, S. (1980). A test of a multivariate normal mean with composite hypotheses determined by linear inequalities. *Biometrika*, 67(2):429–439.
- Silva, J. S. and Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics*, 88(4):641–658.
- Singh, J. (2005). Collaborative networks as determinants of knowledge diffusion patterns. *Management Science*, 51(5):756–770.
- Singh, J. and Fleming, L. (2010). Lone inventors as sources of breakthroughs: Myth or reality? *Management Science*, 56(1):41–56.
- Singh, J. and Marx, M. (2013). Geographic constraints on knowledge spillovers: Political borders vs. spatial proximity. *Management Science*, 59(9):2056–2078.
- Sosa, M. and Maoret, M. (2023). Close to me: The impact of the interplay of physical and social proximity on dyadic collaboration effectiveness. *INSEAD Working Paper*.
- Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.
- Teodoridis, F. (2018). Understanding team knowledge production: The interrelated roles of technology and expertise. *Management Science*, 64(8):3625–3648.
- Teodoridis, F., Bikard, M., and Vakili, K. (2019). Creativity at the knowledge frontier: The impact of specialization in fast-and slow-paced domains. *Administrative Science Quarterly*, 64(4):894–927.
- Tzabbar, D. and Kehoe, R. R. (2014). Can opportunity emerge from disarray? An examination of exploration and exploitation following star scientist turnover. *Journal of Management*, 40(2):449–482.
- Uzzi, B., Mukherjee, S., Stringer, M., and Jones, B. (2013). Atypical combinations and scientific impact. *Science*, 342(6157):468–472.
- van der Wouden, F. and Youn, H. (2023). The impact of geographical distance on learning through collaboration. *Research Policy*, 52(2):104698.

- 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):787–831.
- 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.
- Wang, P., Van De Vrande, V., and Jansen, J. J. (2017). Balancing exploration and exploitation in inventions: Quality of inventions and team composition. *Research Policy*, 46(10):1836–1850.
- Wooldridge, J. M. (1999). Distribution-free estimation of some nonlinear panel data models. *Journal of Econometrics*, 90(1):77–97.
- Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. MIT Press, 2 edition.
- Wu, L., Wang, D., and Evans, J. A. (2019). Large teams develop and small teams disrupt science and technology. *Nature*, 566(7744):378–382.
- Wuchty, S., Jones, B. F., and Uzzi, B. (2007). The increasing dominance of teams in production of knowledge. *Science*, 316(5827):1036–1039.
- Wuyts, S., Colombo, M. G., Dutta, S., and Nooteboom, B. (2005). Empirical tests of optimal cognitive distance. *Journal of Economic Behavior & Organization*, 58(2):277–302.




Figure 1: Dyadic Cognitive Distance between a Coauthor and a Deceased Scientist

Note: Figure 1 presents a histogram of the cognitive distance calculated for each coauthor-deceased scientist dyad in the sample.




Figure 2: Impact of the Unexpected Deaths of Scientists on the Research Output of Coauthors

Note: Figure 2 plots the research output of coauthors after the unexpected death of a prominent scientist against the percentile rankings of the cognitive distance between coauthors and deceased scientists in the sample. The blue dots correspond to the coefficient estimates obtained through a conditional fixed effects QML Poisson model controlling for year fixed effects, career age fixed effects, and interaction terms between a dummy variable for the “death” of the collaborator and indicators of the different percentile ranks of the cognitive distance between a coauthor and their deceased collaborator. The 95% confidence interval (which corresponds to robust standard errors, clustered at the level of the deceased scientist) is illustrated by the dark blue vertical bars and their caps.




Figure 3: Dynamics Underlying the Impact of the Unexpected Deaths of Scientists on the Research Output of Coauthors

Note: Figure 3 illustrates the dynamics underlying the impact of the unexpected death of scientists on the research output of their coauthors who are at the close or distant cognitive distances in each year before and after the year of death, as compared to coauthors at the middle cognitive distance. The horizontal axis measures the number of years since the scientist's unexpected death. The plots connected by each solid line correspond to coefficient estimates drawn from conditional fixed-effects quasi-maximum likelihood Poisson regressions, in which context each dependent variable representing a coauthor's research output is regressed onto year fixed effects, career age effects, and interaction terms between indicators for the year relative to the year of the scientist's unexpected death and each indicator of a coauthor being sorted into the close-distance group (red-circle line) or the distant-distance group (purple-square line), respectively. The specification includes ten leads, ten lags, and two endpoint indicators binning the years beyond ten in each direction; the twenty non-binned coefficients are plotted. Coauthors at a middle level of cognitive distance from the deceased scientist serve as the reference group. The interaction terms for the coauthor group indicators and the year of death are omitted. The 95% confidence intervals (which correspond to robust standard errors clustered at the level of the deceased scientist) are illustrated by the vertical bars and their caps (in red for the close-distance group and in purple for the distant-distance group).




Figure 4: Heterogeneity-Robust Event-Study Estimates: BJS, dCDH, and Sun-Abraham

Note: Figure 4 overlays year-by-year coefficient estimates from three heterogeneity-robust difference-in-differences estimators: the imputation estimator of [Borusyak et al. \(2024\)](#) (navy circles), the continuous-treatment estimator of [de Chaisemartin and D'Haultfoeuille \(2026\)](#) (teal diamonds), and the interaction-weighted estimator of [Sun and Abraham \(2021\)](#) (burnt-orange triangles). Each panel reports estimates for one of the four research output measures. The horizontal axis measures event time relative to the year of the deceased scientist's death; the Sun-Abraham estimator omits the reference period at $t = -1$ by construction. Vertical bars depict 95% confidence intervals from standard errors clustered at the level of the deceased scientist. Numerical estimates are reported in Tables A.1, A.2, and A.3.




Figure 5: Heterogeneous Effects of Collaborator Death on Coauthor Productivity, by Subgroup

Note: Figure 5 plots the coefficient estimates and 95% confidence intervals for the interaction terms $Close \times Death$ (red circles) and $Distant \times Death$ (purple squares) from a series of conditional quasi-maximum likelihood Poisson regressions. Each panel represents a different dependent variable measuring the annual productivity of a coauthor of a deceased star scientist. The subgroups are defined as follows: "Elite" coauthors are in the top decile of cumulative citations at the time of the star's death. "Senior" coauthors have a greater career age than the star. "Same Inst." indicates both authors were at the same institution. "Recent" coauthors published with the star within three years prior to the death. "Regular" coauthors had more than five joint publications. "Deceased First/Last" indicates the star was always the first or last author on collaborations with that specific coauthor. All specifications include coauthor-dyad fixed effects, publication year fixed effects, and flexible controls for the coauthor's career age interacted with their cognitive distance group. The reference group for each estimate is the set of coauthors at a middle-level cognitive distance from the deceased star. Robust standard errors are clustered at the level of the deceased scientist.

Table 1: Cognitive Distance vs. Alternative Proximity Measures

Measure	Built from	Strengths	Limitations
Cognitive distance (this paper)	Journal-to-journal citation flows over the Web of Science (WoS) catalog, aggregated to the researcher pair	Field-agnostic across the WoS catalog; researcher-pair construct; stable over time because journal positioning shifts slowly; broad dyadic coverage	Coarsens within-journal heterogeneity; generalist outlets such as <i>Nature</i> and <i>Science</i> are absorbed into the journal-similarity matrix; not article-level
Idea-space overlap (MeSH) (Azoulay et al., 2010)	Shared MeSH terms on PubMed papers	Article-level granularity; controlled vocabulary supports retrospective term stability	Biomedical-only; the controlled vocabulary is itself curated and contested; not portable to non-Medline disciplines
Reference overlap (Ahuja and Katila, 2001)	Overlap in cited references between publication or patent portfolios	Article-level inputs; theoretically grounded in the recombinant search tradition	Highly endogenous to publication strategy; overlap between distant fields is mechanically near zero, limiting resolution at high distance
Text-based distance (Blei et al., 2003 ; Kelly et al., 2021)	Document-term matrices over abstracts or full text, via topic models or weighted term-overlap similarity	Continuous and data-driven; no a priori category list	Requires text access; modeling choices such as topic count and term weighting are researcher degrees of freedom; measures drift with corpus updates
Field-code overlap (Borjas and Doran, 2012 ; Fafchamps et al., 2010 ; Jaffe, 1986)	Assigned classification codes (JEL codes, patent classes, WoS subject categories)	Inexpensive; precomputed; widely understood	Coarse; codes evolve; inter-field comparability is weak
Geographical proximity (Agrawal et al., 2008 ; Singh and Marx, 2013)	Affiliation distance (kilometers or co-location)	Straightforward to measure; well-supported in the localized-spillovers literature	Orthogonal to cognitive content; co-location is an increasingly imperfect proxy for interaction as remote collaboration expands
Social proximity (Singh, 2005 ; Singh and Fleming, 2010)	Coauthorship history and shared collaborators	Behavioral, observable, granular	Endogenous to outcomes, since collaborators are similar by selection; difficult to separate from cognitive proximity (Boschma, 2005)

Notes: Rows 2 through 5 are alternative empirical implementations of cognitive proximity. Rows 6 and 7 cover the geographical and social proximity dimensions of the [Boschma \(2005\)](#) typology that feature most often in the scientific-collaboration literature.

Table 2: Summary Statistics for Deceased Scientists

	Mean	Median	Std. Dev.	Min.	Max.	Observations
Year of death	2000	2000	2.84	1995	2008	63
Year of first publication	1981	1980	2.85	1971	1990	63
Year of last publication	2004	2004	4.06	1995	2013	63
Career age	19	19	3.63	10	31	63
Cum. no. of publications	122	99	91.96	5	570	63
No. of cum. citations at the time of death	5497	3759	5490.02	59	25692	63
No. of citations by career age	280	197	269.24	5	1352	63
No. of coauthors	147	115	119.67	5	460	63

Note: The sample of deceased scientists contains 63 life scientists who passed away unexpectedly and prematurely between 1995 and 2008. The first and last year of publication refer to the publication years of a scientist's first and last academic paper. Career age is defined as the difference between a scientist's year of death and the year of his/her first publication in years. Cumulative number of publications refers to the journal articles produced by a deceased scientist from 1980 onward that are indexed in our analysis extract; other Web of Science document types (editorials, letters, biographical items) are excluded. The screening thresholds of Section 5.2 were applied to candidates' complete Web of Science publication records at the identification stage, before these restrictions, so individual values reported here can fall below those thresholds. Number of cumulative citations represents the total number of citations of the papers published by the deceased scientist as of the year of death. Number of citations by career age indicates the number of cumulative citations of the work of a deceased scientist at the time of death divided by his/her career age. Number of coauthors is defined as the total number of researchers with whom a deceased scientist has coauthored published academic papers.

Table 3: Summary Statistics for Coauthors in the Year of Deceased Scientists' Death

	All Coauthors		Close-distance Coauthors		Middle-distance Coauthors		Distant-distance Coauthors	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Cognitive distance	0.62	0.16	0.46	0.11	0.69	0.06	0.86	0.05
No. of pubs in each pub year	3.27	3.24	3.00	3.13	3.55	3.37	3.16	3.03
Cum. no. of pubs	54.69	120.55	45.77	116.38	62.87	125.01	54.23	116.09
No. of JIF-weighted pubs in each pub year	10.99	15.32	12.23	17.52	11.01	14.30	7.49	10.54
Cum. no. of JIF-weighted pubs	184.15	544.93	186.26	605.69	198.69	534.33	133.15	366.46
No. of pubs abv 50th citation %ile in each pub year	2.31	2.51	2.20	2.40	2.52	2.69	1.97	2.19
Cum. no. of pubs above 50th citation %ile	39.07	91.73	33.39	87.55	45.46	97.96	35.09	81.36
No. of pubs abv 85th citation %ile in each pub year	1.05	1.39	1.06	1.37	1.13	1.48	0.78	1.10
Cum. no. of pubs above 85th citation %ile	17.79	46.92	15.88	44.58	20.77	51.31	13.87	37.51
Career age	13.99	6.43	12.78	6.41	14.74	6.16	15.06	6.77
Senior coauthor	0.08	0.28	0.07	0.26	0.09	0.29	0.09	0.28
No. of cumulative citations	1882.40	4046.60	1865.64	4269.22	2198.86	4077.36	1636.58	3185.37
No. of citations by career age	118.86	216.74	117.38	238.43	127.26	208.87	96.92	170.74
Elite coauthor	0.10	0.30	0.09	0.29	0.12	0.32	0.07	0.26
No. of collaborations	3.16	5.95	3.21	5.31	3.29	6.96	2.60	3.76
Regular coauthor	0.13	0.33	0.14	0.34	0.13	0.34	0.10	0.30
Recent coauthor	0.45	0.50	0.45	0.50	0.45	0.50	0.43	0.49
Same institute	0.10	0.31	0.09	0.29	0.11	0.31	0.14	0.34
Deceased first author	0.03	0.17	0.02	0.15	0.03	0.16	0.04	0.21
Deceased last author	0.21	0.41	0.19	0.39	0.22	0.41	0.26	0.44
Observations	9263		3751		4169		1343	

Note: The sample of coauthors includes scientists who served as coauthors with the 63 deceased life scientists included in this research on academic papers published between 1980 and 2013. All variables are measured as of the deceased scientists' death. Percentiles for citations are calculated based on the total citations of articles published in a given year in the life sciences. Career age is defined as the difference between the deceased scientist's year of death and the author's year of first publication, in years. Senior coauthors include individuals whose career age is greater than that of the deceased scientist at the time of death. The number of cumulative citations is calculated as the total number of citations of the scientist's papers by the year of their deceased coauthor's death. The number of citations by career age refers to the total number of cumulative citations divided by the scientist's career age. Elite coauthors are defined as those whose number of cumulative citations ranks above the top 10th percentile of the full sample of coauthors. The number of collaborations indicates the number of coauthored articles published by a coauthor-deceased scientist dyad between 1980 and 2013. Regular coauthors are defined as those who collaborated with their deceased coauthor more than five times. Recent coauthors refer to those who published at least one collaboration within the three years before their coauthor's death. Same institution indicates that a coauthor and a deceased scientist worked at the same institution. Deceased first author indicates that the deceased scientist was consistently listed as the first author on all joint publications with a given coauthor. Deceased last author indicates that the deceased scientist was consistently listed as the last author on all joint publications with a given coauthor.

Table 4: Impact of the Unexpected Deaths of Scientists on Coauthors' Research Productivity

<i>Dependent Variable:</i>	No. of Publications				Distribution of Citations			
	Unweighted		JIF-weighted		50th %ile of Cits		85th %ile of Cits	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
main								
Death	-0.122 (0.115)	0.037 (0.027)	-0.476*** (0.181)	0.030 (0.043)	-0.238** (0.116)	0.030 (0.027)	-0.247* (0.137)	-0.005 (0.031)
Distance × Death	0.152 (0.159)		0.669** (0.261)		0.319* (0.163)		0.308* (0.185)	
Close × Death		-0.118*** (0.044)		-0.202*** (0.068)		-0.143*** (0.041)		-0.112*** (0.040)
Distant × Death		-0.145*** (0.047)		-0.036 (0.061)		-0.103** (0.051)		-0.040 (0.057)
Log pseudo-likelihood	-333342.1	-333153.9	-861489	-861036.7	-296243.1	-296136.3	-204420.5	-204395.6
No. of Observations	155294	155294	155294	155294	154895	154895	150256	150256

Note: Robust standard errors (shown in parentheses) are clustered at the level of the deceased scientist. The estimates are drawn from conditional quasi-maximum likelihood Poisson specifications. Observations are made at the coauthor_{*i*}-deceased scientist_{*j*}-year_{*t*} level. The dependent variable of the models in columns (1) and (2) is the raw number of publications authored by a coauthor (of a deceased scientist) in the year of observation. The dependent variable of the models in columns (3) and (4) is the number of JIF-weighted publications authored by a coauthor (of a deceased scientist) in the year of observation. The dependent variable of the models in columns (5) and (6) is the number of papers authored by a coauthor (of a deceased scientist) in a given year that rank above the 50th percentile of papers published that year in terms of total citations. The dependent variable of the models in columns (7) and (8) is the number of papers authored by a coauthor (of a deceased scientist) in a given year that rank above the 85th percentile of papers published that year in terms of total citations. All columns control for publication year dummy variables, indicator variables for career age (year of publication minus year of first publication; career age zero is the omitted category), and the interaction terms between the career age indicators and each covariate of interest. In columns (2), (4), (6), and (8), the reference group contains coauthors at the middle level of cognitive distance from the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table 5: Robustness Check: No Publications of Deceased Scientists

<i>Dependent Variable:</i>	No. of Publications				Distribution of Citations			
	Unweighted		JIF-weighted		50th %ile of Cits		85th %ile of Cits	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Death	-0.081 (0.110)	0.076** (0.030)	-0.397** (0.174)	0.077 (0.047)	-0.194* (0.110)	0.080** (0.032)	-0.248** (0.117)	0.058 (0.041)
Distance $\times$ Death	0.144 (0.156)		0.625** (0.258)		0.322** (0.162)		0.403** (0.173)	
Close $\times$ Death		-0.121*** (0.044)		-0.192*** (0.065)		-0.147*** (0.042)		-0.123*** (0.045)
Distant $\times$ Death		-0.131*** (0.043)		-0.011 (0.051)		-0.089* (0.046)		-0.012 (0.043)
Log pseudo-likelihood	-317718.7	-317512.7	-816839.8	-815951.7	-281488.3	-281353.7	-192649	-192614.6
No. of Observations	146808	146808	146808	146808	146413	146413	141677	141677

Note: Robust standard errors (shown in parentheses) are clustered at the level of the deceased scientist. The estimates are drawn from conditional quasi-maximum likelihood Poisson specifications. Observations are made at the coauthor_{*i*}-deceased scientist_{*j*}-year_{*t*} level. The dependent variable in columns (1) and (2) is the raw number of publications authored by a coauthor (of a deceased scientist) in the year of observation. The dependent variable in columns (3) and (4) is the number of JIF-weighted publications authored by a coauthor (of a deceased scientist) in the year of observation. The dependent variable in columns (5) and (6) is the number of papers authored by a coauthor (of a deceased scientist) in a given year that rank above the 50th percentile of papers published that year in terms of total citations. The dependent variable in columns (7) and (8) is the number of papers authored by a coauthor (of a deceased scientist) in a given year that rank above the 85th percentile of papers published that year in terms of total citations. All columns control for publication year dummy variables, indicator variables for career age (year of publication minus year of first publication; career age zero is the omitted category), and the interaction terms between the career age indicators and each covariate of interest. In columns (2), (4), (6), and (8), the reference group contains coauthors at the middle level of cognitive distance from the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table 6: Robustness Check: No Multiple Losses

<i>Dependent Variable:</i>	No. of Publications				Distribution of Citations			
	Unweighted		JIF-weighted		50th %ile of Cits		85th %ile of Cits	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Death	-0.078 (0.096)	0.033 (0.030)	-0.392*** (0.147)	0.024 (0.047)	-0.200** (0.095)	0.027 (0.030)	-0.211* (0.112)	0.001 (0.039)
Distance $\times$ Death	0.105 (0.132)		0.569*** (0.212)		0.281** (0.133)		0.277* (0.153)	
Close $\times$ Death		-0.078* (0.040)		-0.148** (0.059)		-0.106*** (0.038)		-0.086** (0.043)
Distant $\times$ Death		-0.105** (0.041)		-0.006 (0.044)		-0.064 (0.046)		-0.029 (0.062)
Log pseudo-likelihood	-289371.1	-289286.8	-716194.1	-715416.9	-255751.5	-255694.2	-173324.4	-173277.7
No. of Observations	143584	143584	143584	143584	143185	143185	138592	138592

Note: Robust standard errors (shown in parentheses) are clustered at the level of the deceased scientist. The estimates are drawn from conditional quasi-maximum likelihood Poisson specifications. Observations are made at the coauthor_{*i*}-deceased scientist_{*j*}-year_{*t*} level. The dependent variable in columns (1) and (2) is the raw number of publications authored by a coauthor (of a deceased scientist) in the year of observation. The dependent variable in columns (3) and (4) is the number of JIF-weighted publications authored by a coauthor (of a deceased scientist) in the year of observation. The dependent variable in columns (5) and (6) is the number of papers authored by a coauthor (of a deceased scientist) in a given year that rank above the 50th percentile of papers published that year in terms of total citations. The dependent variable in columns (7) and (8) is the number of papers authored by a coauthor (of a deceased scientist) in a given year that rank above the 85th percentile of papers published that year in terms of total citations. All columns control for publication year dummy variables, indicator variables for career age (year of publication minus year of first publication; career age zero is the omitted category), and the interaction terms between the career age indicators and each covariate of interest. In columns (2), (4), (6), and (8), the reference group contains coauthors at the middle level of cognitive distance from the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table 7: Robustness Check: Exclusion of Highly Cited Deceased Scientists

<i>Dependent Variable:</i>	No. of Publications				Distribution of Citations			
	Unweighted		JIF-weighted		50th %ile of Cits		85th %ile of Cits	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Death	-0.144 (0.142)	0.033 (0.029)	-0.506** (0.217)	0.038 (0.048)	-0.266* (0.145)	0.032 (0.029)	-0.272 (0.175)	0.005 (0.034)
Distance $\times$ Death	0.165 (0.197)		0.705** (0.310)		0.345* (0.203)		0.338 (0.236)	
Close $\times$ Death		-0.137*** (0.049)		-0.222*** (0.071)		-0.167*** (0.045)		-0.143*** (0.044)
Distant $\times$ Death		-0.141*** (0.050)		-0.034 (0.066)		-0.101* (0.057)		-0.041 (0.062)
Log pseudo-likelihood	-274144.7	-273979.2	-708180.6	-707615.5	-242744.4	-242643.1	-166069.8	-166046.3
No. of Observations	126867	126867	126867	126867	126497	126497	122418	122418

Note: Robust standard errors (shown in parentheses) are clustered at the level of the deceased scientist. The estimates are drawn from conditional quasi-maximum likelihood Poisson specifications. Observations are made at the coauthor_{*i*}-deceased scientist_{*j*}-year_{*t*} level. The dependent variable in columns (1) and (2) is the raw number of publications authored by a coauthor of a deceased scientist in the year of observation. The dependent variable in columns (3) and (4) is the number of JIF-weighted publications authored by a coauthor of a deceased scientist in the year of observation. The dependent variable in columns (5) and (6) is the number of papers authored by a coauthor of a deceased scientist in a given year that rank above the 50th percentile of papers published that year in terms of total citations. The dependent variable in columns (7) and (8) is the number of papers authored by a coauthor of a deceased scientist in a given year that rank above the 85th percentile of papers published that year in terms of total citations. All columns control for publication year dummy variables, indicator variables for career age (year of publication minus year of first publication; career age zero is the omitted category), and the interaction terms between the career age indicators and each covariate of interest. In columns (2), (4), (6), and (8), the reference group contains coauthors at the middle level of cognitive distance from the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.


Table 8: Robustness Check: Exclusion of the Most Popular Deceased Scientists

<i>Dependent Variable:</i>	No. of Publications				Distribution of Citations			
	Unweighted		JIF-weighted		50th %ile of Cits		85th %ile of Cits	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Death	-0.165*	0.029	-0.418***	0.023	-0.273***	0.022	-0.320***	-0.023
	(0.100)	(0.031)	(0.126)	(0.054)	(0.082)	(0.032)	(0.084)	(0.036)
Distance × Death	0.224*		0.615***		0.381***		0.431***	
	(0.132)		(0.188)		(0.111)		(0.116)	
Close × Death		-0.099**		-0.146**		-0.120***		-0.085*
		(0.048)		(0.063)		(0.045)		(0.047)
Distant × Death		-0.093*		0.012		-0.049		0.057
		(0.052)		(0.074)		(0.055)		(0.058)
Log pseudo-likelihood	-235385.6	-235227.6	-620438.7	-620388.4	-208854.7	-208746	-143649.9	-143604.9
No. of Observations	109028	109028	109028	109028	108703	108703	105033	105033

Note: Robust standard errors (shown in parentheses) are clustered at the level of the deceased scientist. The estimates are drawn from conditional quasi-maximum likelihood Poisson specifications. Observations are made at the coauthor_{*i*}-deceased scientist_{*j*}-year_{*t*} level. The dependent variable in columns (1) and (2) is the raw number of publications authored by a coauthor of a deceased scientist in the year of observation. The dependent variable in columns (3) and (4) is the number of JIF-weighted publications authored by a coauthor of a deceased scientist in the year of observation. The dependent variable in columns (5) and (6) is the number of papers authored by a coauthor of a deceased scientist in a given year that rank above the 50th percentile of papers published that year in terms of total citations. The dependent variable in columns (7) and (8) is the number of papers authored by a coauthor of a deceased scientist in a given year that rank above the 85th percentile of papers published that year in terms of total citations. All columns control for publication year dummy variables, indicator variables for career age (year of publication minus year of first publication; career age zero is the omitted category), and the interaction terms between the career age indicators and each covariate of interest. In columns (2), (4), (6), and (8), the reference group contains coauthors at the middle level of cognitive distance from the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.


A Supplementary Figures and Tables

Figure A.1 Distribution of Coauthorship Intensity




Note: Histogram of the number of collaborations for the 9,263 coauthor-deceased scientist dyads in the sample.

Figure A.2 Impact of the Unexpected Deaths of Scientists on the Research Output of Coauthors (5-Percentile Resolution)




Note: This figure replicates [Figure 2](#) at 5-percentile bin resolution (20 bins instead of 10). Vertical dashed lines mark the analytical cutoffs at the 40th and 85th percentiles used to define the cognitively close (0–40), middle-distance (40–85), and distant (85–100) groups; both cutoffs sit at bin boundaries. The empirical specification, sample, and 95% confidence-interval construction are identical to those in [Figure 2](#).

Figure A.3 Mean-Reversion Robustness: Dynamic Event-Study Estimates under the Active-Pair and Top-Quartile Restrictions




Note: Dynamic event-study coefficients on $Close \times Death$ and $Distant \times Death$ across the four research output measures defined in Table A.1, estimated on the Active-Pair subsample (dyads with strictly positive pre-event total output for the relevant outcome). The static estimates for all three sample definitions, including the headline (Base) sample and the Top-Quartile restriction, are reported in Table A.8. Vertical bars denote 95 percent confidence intervals; standard errors are clustered at the level of the deceased scientist. Section 6.2 discusses the results.

Figure A.4 Subgroup Forest Plot: *Close* $\times$ *Death* and *Distant* $\times$ *Death* Coefficients with 95% Confidence Intervals



Note: Forest plot of the seven subgroup-specific *Close* $\times$ *Death* (navy circles) and *Distant* $\times$ *Death* (maroon squares) coefficients across the four research output measures defined in Table A.1, shown with their 95 percent confidence intervals. Within each row, asterisks next to the subgroup label denote rejection of the within-row Wald test of $H_0 : \beta_{\text{close}} = \beta_{\text{distant}}$, not the significance of the individual coefficients. Numerical estimates of the gap and the corresponding Wald-test p-values are reported in Table A.11. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Figure A.5 Leave-One-Out Sensitivity: Excluding Each Deceased Scientist in Turn



Note: Each panel plots the 63 leave-one-out estimates of the *Close × Death* (red circles) and *Distant × Death* (purple diamonds) coefficients from the categorical specification of Equation (5), obtained by excluding each deceased scientist in turn; estimates are sorted by coefficient value within each series. Dashed horizontal lines mark the full-sample estimates. Section 6.2 discusses the influence check and the cluster-jackknife (CV3) standard errors computed from these estimates.

Table A.1: BJS Imputation Event-Study Estimates, Pooled Sample

	(1) Unweighted	(2) JIF-weighted	(3) Top-85%ile	(4) Top-50%ile
$t = 0$	0.0718 (0.0893)	0.3872 (0.3897)	0.0430 (0.0371)	0.1169 (0.0716)
$t = +1$	-0.1148 (0.1504)	-0.7854 (0.6845)	-0.0802 (0.0737)	-0.0980 (0.1152)
$t = +2$	-0.3497 (0.2285)	-1.2314 (0.7992)	-0.1435** (0.0694)	-0.3519** (0.1786)
$t = +3$	-0.2383 (0.2933)	-1.3779 (1.2931)	-0.1418 (0.1276)	-0.2739 (0.2381)
$t = +4$	-0.5162* (0.3081)	-2.0675* (1.0684)	-0.1358* (0.0767)	-0.4525** (0.2294)
$t = +5$	-0.6011** (0.2999)	-2.8679** (1.1796)	-0.1505 (0.1071)	-0.5089** (0.2385)
$t = +6$	-0.8944*** (0.3265)	-3.4478*** (1.2927)	-0.2935*** (0.1016)	-0.8054*** (0.2467)
$t = +7$	-0.6865* (0.3530)	-2.6318** (1.3340)	-0.1470 (0.1128)	-0.5772** (0.2616)
$t = +8$	-0.7205** (0.3038)	-2.8122** (1.1534)	-0.1613 (0.0996)	-0.6468*** (0.2374)
$t = +9$	-0.8217** (0.4131)	-4.0884** (1.9598)	-0.3805* (0.1981)	-0.7939** (0.3369)
$t = +10$	-1.4920*** (0.4124)	-5.1002*** (1.1978)	-0.2968*** (0.1055)	-1.2147*** (0.2913)
$t = -1$	0.1520 (0.3990)	3.6531 (2.5877)	0.1502 (0.1736)	0.2625 (0.3396)
$t = -2$	0.2496 (0.3437)	4.1476 (2.6818)	0.2488 (0.1628)	0.3614 (0.2958)
$t = -3$	0.1954 (0.2775)	2.8940 (1.9393)	0.1217 (0.1176)	0.2486 (0.2425)
$t = -4$	0.1767 (0.1860)	1.9555 (1.3947)	0.1520 (0.0942)	0.2057 (0.1691)
$t = -5$	0.1014 (0.1365)	1.2691 (1.0433)	0.0867 (0.0804)	0.1314 (0.1329)
N	130478	130478	130478	130478

Note: [Borusyak et al. \(2024\)](#) imputation estimator applied to the binary-event specification. Cohort is defined as the deceased scientist's year of death. Coefficients at $t = 0$ through $t = +10$ report dynamic ATTs, and coefficients at $t = -1$ through $t = -5$ report pretrend placebos used to assess parallel trends. The four columns show, respectively, the unweighted publication count, JIF-weighted publications, the count of papers above the 85th citation percentile, and the count above the 50th citation percentile. Standard errors (in parentheses) are clustered at the level of the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table A.2: dCDH Continuous-Treatment Event-Study Estimates

	(1) Unweighted	(2) JIF-weighted	(3) Top-85%ile	(4) Top-50%ile
$t = +1$	-0.2801** (0.1134)	-1.3717** (0.6184)	-0.1626** (0.0711)	-0.3075*** (0.1077)
$t = +2$	-0.3869*** (0.1324)	-1.6393** (0.6821)	-0.2049** (0.0835)	-0.4565*** (0.1290)
$t = +3$	-0.4218* (0.2412)	-2.6534* (1.4292)	-0.2506** (0.1256)	-0.5635** (0.2538)
$t = +4$	-0.6553** (0.2595)	-3.4744*** (1.1826)	-0.2368** (0.0979)	-0.7638*** (0.2087)
$t = +5$	-0.6278* (0.3382)	-3.6960** (1.6886)	-0.3252*** (0.0877)	-0.7425*** (0.2711)
$t = +6$	-1.0888 (0.6895)	-5.3319** (2.7157)	-0.4025 (0.3285)	-1.0040* (0.6021)
$t = +7$	-0.7728** (0.3680)	-5.1266** (2.4109)	-0.3860 (0.3883)	-0.8374* (0.4473)
$t = +8$	-0.8043** (0.3397)	-4.9900*** (1.9003)	-0.4199 (0.3512)	-0.8423*** (0.3097)
$t = +9$	-1.0512** (0.4459)	-6.4854*** (2.4364)	-0.6567** (0.2838)	-0.8951** (0.3664)
$t = +10$	-1.4892*** (0.5121)	-8.6886*** (2.6379)	-0.7098*** (0.2466)	-1.4438*** (0.3922)
$t = -1$	-0.2194** (0.1108)	-0.1396 (0.4602)	-0.1071 (0.0657)	-0.2240** (0.0923)
$t = -2$	-0.2230* (0.1299)	1.2708 (0.8383)	0.0207 (0.0738)	-0.1077 (0.1162)
$t = -3$	-0.2141* (0.1248)	-0.3786 (0.5880)	-0.2078 (0.1403)	-0.3040 (0.1868)
$t = -4$	0.0667 (0.1715)	-0.6283 (0.8440)	0.0075 (0.0792)	-0.1947 (0.1348)
$t = -5$	0.2706 (0.4397)	-0.3508 (1.2372)	0.1055 (0.2400)	0.1585 (0.2916)
Average Total Effect	-1.0104*** (0.3433)	-5.6499*** (1.8113)	-0.4915** (0.1945)	-1.0883*** (0.2945)
Joint nullity of effects (p -value)	0.006	0.057	0.003	<0.001
Joint nullity of placebos (p -value)	0.297	0.595	0.221	0.124
N	155294	155294	155294	155294

Note: de Chaisemartin and D'Haultfoeulle (2026) heterogeneity-robust estimator with continuous treatment intensity defined as $Distance_{ij} \times Death_{jt}$. The table reports 10 post-event horizons ($t = +1$ through $t = +10$), 5 placebo periods ($t = -1$ through $t = -5$), the average total effect across all post-event periods, and the p -values from tests of the joint nullity of the ten dynamic effects and of the five placebos. Because the treatment indicator switches on in the year following death, the estimator's first dynamic effect corresponds to event time $t = +1$ and no death-year ($t = 0$) effect is estimated. The four columns correspond to the four research output measures defined in Table A.1. Standard errors (in parentheses) are clustered at the level of the deceased scientist. The estimator is linear; the headline TWFE specifications in the main text are Poisson, so the reported coefficients serve as directional checks on sign and significance. A sensitivity allowing regime-specific nonparametric trends (`trends_nonparam` over the three cognitive-distance bands of Table A.4) returns average total effects of -0.84^{**} , -5.29^{***} , -0.50^{***} , and -0.99^{***} across columns (1) to (4); the placebo pattern is qualitatively similar, with the $t = -1$ lead remaining significant for the unweighted and top-50%ile outcomes and the $t = -3$ lead strengthening to significance at the five percent level for the same two outcomes. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table A.3: Sun-Abraham Interaction-Weighted Estimates and Poisson Sensitivity

	(1) Unweighted	(2) JIF-weighted	(3) Top-85%ile	(4) Top-50%ile	(5) Unweighted (Poisson)
$t = -5$	0.4422* (0.2445)	-0.2775 (0.6753)	0.0917 (0.1431)	0.2864 (0.2202)	-0.0482 (0.0457)
$t = -4$	0.3570** (0.1580)	0.0342 (0.4981)	0.1383 (0.1033)	0.2363* (0.1431)	-0.0105 (0.0306)
$t = -3$	0.2358** (0.1032)	0.2398 (0.3560)	0.0336 (0.0545)	0.1621* (0.0860)	-0.0022 (0.0188)
$t = -2$	0.1453** (0.0598)	1.0970*** (0.3963)	0.1496*** (0.0465)	0.1547*** (0.0513)	0.0040 (0.0182)
$t = 0$	0.1251 (0.0865)	0.1505 (0.2598)	0.0669** (0.0327)	0.1788*** (0.0592)	0.0210 (0.0146)
$t = +1$	-0.0383 (0.1064)	-0.7685* (0.4461)	-0.0280 (0.0779)	-0.0044 (0.0949)	0.0037 (0.0235)
$t = +2$	-0.1009 (0.1062)	-0.6789 (0.4208)	-0.0458 (0.0673)	-0.0922 (0.1025)	-0.0265 (0.0289)
$t = +3$	-0.0139 (0.1299)	-0.4106 (0.5507)	-0.0646 (0.0736)	-0.0293 (0.1217)	-0.0307 (0.0361)
$t = +4$	-0.1721 (0.1763)	-1.1744 (0.7660)	-0.0914 (0.1004)	-0.1712 (0.1601)	-0.0539 (0.0416)
$t = +5$	-0.1740 (0.2108)	-1.1877 (0.9313)	-0.1149 (0.1030)	-0.1806 (0.1840)	-0.0820* (0.0480)
$t = +6$	-0.1285 (0.2126)	-0.9775 (1.1269)	-0.1686 (0.1148)	-0.1669 (0.2017)	-0.0935* (0.0528)
$t = +7$	-0.1277 (0.2209)	-1.4430 (1.0887)	-0.1736 (0.1098)	-0.2240 (0.1933)	-0.0999 (0.0670)
$t = +8$	-0.2161 (0.2398)	-1.9110 (1.2930)	-0.2577** (0.1245)	-0.3122 (0.2135)	-0.1331* (0.0776)
$t = +9$	-0.1763 (0.2681)	-2.5397* (1.3500)	-0.3561*** (0.1200)	-0.2881 (0.2360)	-0.1553* (0.0858)
$t = +10$	-0.7371*** (0.1335)	-4.4689*** (0.9193)	-0.4269*** (0.0765)	-0.7747*** (0.1213)	-0.2415** (0.1183)
N	155294	155294	155294	155294	155294

Note: Columns (1)–(4) report the [Sun and Abraham \(2021\)](#) interaction-weighted estimator implemented via `eventstudyinteract`, applied to the four research output measures defined in Table A.1. The latest-treated cohort serves as the comparison group because the within-treated design contains no never-treated dyads. Column (5) reports a Poisson event-study sensitivity check estimated via the `ppmlhdfe` routine of [Correia et al. \(2020\)](#) on the unweighted publication count, included as a count-data confirmation that the heterogeneity-robust dynamics survive the Poisson link function used in the headline tables. The reference period is event time -1 , which is omitted from the table by construction. Visual overlays of the [Borusyak et al. \(2024\)](#), [de Chaisemartin and D’Haultfoeuille \(2026\)](#), and [Sun and Abraham \(2021\)](#) estimators are reported in [Figure 4](#). Standard errors (in parentheses) are clustered at the level of the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table A.4: BJS Imputation Event Studies by Distance Subsample

	Close [0,40]				Middle (40,85]				Distant (85,100]			
	(1) Unwt	(2) IPF	(3) T85	(4) T50	(5) Unwt	(6) IPF	(7) T85	(8) T50	(9) Unwt	(10) IPF	(11) T85	(12) T50
$t = 0$	-0.0854 (0.1783)	0.4472 (0.7264)	0.0016 (0.0598)	0.0068 (0.1253)	0.1832 (0.1465)	0.2848 (0.4876)	0.0405 (0.0670)	0.1403 (0.1341)	0.2967 (0.2232)	0.5635 (0.5621)	0.1351* (0.0795)	0.2691* (0.1556)
$t = +1$	-0.1772 (0.2391)	-0.8596 (0.9581)	-0.1094 (0.0820)	-0.2466* (0.1265)	-0.0780 (0.2698)	-1.3304 (1.2036)	-0.1589 (0.1599)	-0.0726 (0.2517)	-0.0160 (0.2603)	0.5978 (0.7238)	0.0710 (0.0736)	0.1071 (0.1762)
$t = +2$	-0.5422* (0.3257)	-1.8169* (0.9768)	-0.2558*** (0.0656)	-0.5979*** (0.1418)	-0.2492 (0.4060)	-1.4831 (1.3978)	-0.1342 (0.1436)	-0.3811 (0.3914)	-0.2785 (0.3580)	0.1161 (0.8871)	-0.0211 (0.0787)	-0.1136 (0.2358)
$t = +3$	-0.4861 (0.3643)	-1.7855** (0.7662)	-0.1810*** (0.0624)	-0.5082*** (0.1204)	0.0504 (0.5425)	-1.6445 (2.4676)	-0.2725 (0.2961)	-0.2169 (0.5077)	0.0479 (0.3632)	1.0448 (1.0156)	0.1103 (0.0910)	0.0837 (0.2460)
$t = +4$	-1.0848*** (0.3942)	-4.1586*** (0.7465)	-0.3466*** (0.0573)	-0.9441*** (0.1305)	0.0235 (0.5565)	-0.8388 (2.1226)	-0.1116 (0.1777)	-0.2060 (0.5066)	0.2832 (0.3645)	1.8516* (0.9663)	0.2199*** (0.0794)	0.2445 (0.2340)
$t = +5$	-1.6500*** (0.3469)	-6.5146*** (0.8851)	-0.4581*** (0.0744)	-1.2843*** (0.1479)	0.3469 (0.5306)	-0.6464 (2.0695)	-0.0969 (0.2297)	-0.0764 (0.4994)	0.5699 (0.3495)	2.9184*** (1.0900)	0.2730*** (0.1023)	0.3498 (0.2382)
$t = +6$	-1.8954*** (0.3296)	-7.1196*** (0.8558)	-0.5357*** (0.0807)	-1.4576*** (0.1649)	-0.0223 (0.6764)	-2.2358 (2.9754)	-0.4322 (0.3076)	-0.4914 (0.6161)	-0.1716 (0.4774)	2.2802 (1.4358)	0.2491** (0.1146)	0.1042 (0.3334)
$t = +7$	-2.1086*** (0.4222)	-7.6635*** (1.4213)	-0.3841*** (0.1330)	-1.4958*** (0.2495)	0.3901 (0.7159)	0.2432 (2.7499)	-0.0663 (0.1934)	-0.1169 (0.6283)	0.7036 (0.4986)	4.1245*** (1.5327)	0.2733* (0.1509)	0.5723 (0.3496)
$t = +8$	-2.2745*** (0.4622)	-9.2851*** (1.2056)	-0.6338*** (0.1578)	-1.7805*** (0.2334)	1.1957** (0.4905)	2.3585 (1.8535)	0.0251 (0.1706)	0.3620 (0.4430)	0.6741 (0.4440)	6.1238*** (1.5947)	0.5091*** (0.1484)	0.6413** (0.3146)
$t = +9$	-1.4117*** (0.4619)	-6.6676*** (1.2731)	-0.6369*** (0.1716)	-1.2712*** (0.2602)	0.8276 (0.6078)	0.2038 (2.6257)	-0.2833 (0.2680)	0.0880 (0.5197)	0.9736*** (0.3474)	4.9579*** (1.1468)	0.1272 (0.1120)	0.7153*** (0.2296)
$t = +10$	-2.3276*** (0.7770)	-10.0344*** (1.2873)	-0.8055*** (0.1587)	-2.0044*** (0.4050)	-0.1539 (0.6109)	-1.2208 (2.2728)	-0.3306** (0.1301)	-0.5560 (0.5337)	0.9082** (0.3803)	8.6057*** (0.8396)	1.1149*** (0.0765)	0.8048*** (0.2363)
$t = -1$	0.3549 (0.6898)	8.9653** (4.4253)	0.3849 (0.2539)	0.5544 (0.5500)	0.1347 (0.4758)	0.4724 (1.7797)	-0.0170 (0.2310)	0.1135 (0.4128)	-0.3857 (0.5512)	-1.5056 (1.2377)	0.0145 (0.1490)	-0.1616 (0.3079)
$t = -2$	0.4211 (0.6482)	9.3796* (4.8449)	0.3888 (0.2452)	0.6091 (0.5268)	0.1713 (0.3520)	0.8159 (1.4086)	0.1598 (0.1912)	0.2224 (0.3167)	-0.0243 (0.3974)	-0.3267 (0.8762)	0.1284 (0.1991)	0.0939 (0.2519)
$t = -3$	0.3825 (0.5039)	7.2369** (3.3921)	0.2675 (0.1731)	0.5093 (0.4039)	0.1450 (0.2900)	0.2085 (1.1809)	-0.0040 (0.1490)	0.0845 (0.2648)	-0.2281 (0.3141)	-1.0621 (0.7847)	0.0039 (0.1044)	-0.1361 (0.2020)
$t = -4$	0.4023 (0.3400)	5.0064** (2.4996)	0.3389** (0.1513)	0.4692 (0.2965)	0.1381 (0.1901)	0.1375 (0.7997)	0.0129 (0.1049)	0.0863 (0.1674)	-0.1767 (0.2548)	-0.6087 (0.6157)	0.0612 (0.1063)	-0.1258 (0.1900)
$t = -5$	0.2813 (0.2545)	3.4289* (1.9268)	0.1720 (0.1346)	0.3033 (0.2342)	0.0358 (0.1391)	-0.3104 (0.6446)	0.0277 (0.0890)	0.0473 (0.1315)	-0.0992 (0.1631)	0.3737 (0.6374)	0.0632 (0.0776)	-0.0565 (0.1517)
N	48106	48106	48106	48106	63847	63847	63847	63847	18525	18525	18525	18525

Note: BJS imputation estimator applied separately to dyads in each cognitive-distance band. Columns (1)–(4) cover the close band (0–40th percentiles), columns (5)–(8) the middle band (40–85th percentiles), and columns (9)–(12) the distant band (85–100th percentiles). Within each band the four columns correspond to the unweighted publication count, JIF-weighted publications, the count of papers above the 85th citation percentile, and the count above the 50th citation percentile. Cohort, horizon, and clustering definitions are as in Table A.1. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table A.5: Robustness Check: Non-Leave-Out Index of Cognitive Distance

<i>Dependent Variable:</i>	No. of Publications				Distribution of Citations			
	Unweighted		JIF-weighted		50th %ile of Cits		85th %ile of Cits	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Death	-0.096 (0.115)	0.027 (0.028)	-0.480*** (0.178)	0.027 (0.048)	-0.227* (0.118)	0.014 (0.029)	-0.264* (0.146)	-0.018 (0.035)
Distance $\times$ Death	0.113 (0.160)		0.680*** (0.260)		0.302* (0.167)		0.335* (0.200)	
Close $\times$ Death		-0.093** (0.043)		-0.197*** (0.071)		-0.109*** (0.042)		-0.089* (0.051)
Distant $\times$ Death		-0.111*** (0.040)		-0.014 (0.058)		-0.061 (0.047)		-0.003 (0.055)
Log pseudo-likelihood	-335749.9	-335684	-865867.6	-866596.2	-298282.3	-298288.3	-205761.5	-205811.5
No. of Observations	157336	157336	157336	157336	156918	156918	152033	152033

Note: Robust standard errors (shown in parentheses) are clustered at the level of the deceased scientist. The estimates are drawn from conditional quasi-maximum likelihood Poisson specifications. Observations are made at the coauthor_{*i*}-deceased scientist_{*j*}-year_{*t*} level. The dependent variable of the models in columns (1) and (2) is the raw number of publications authored by a coauthor (of a deceased scientist) in the year of observation. The dependent variable of the models in columns (3) and (4) is the number of JIF-weighted publications authored by a coauthor (of a deceased scientist) in the year of observation. The dependent variable of the models in columns (5) and (6) is the number of papers authored by a coauthor (of a deceased scientist) in a given year that rank above the 50th percentile of papers published that year in terms of total citations. The dependent variable of the models in columns (7) and (8) is the number of papers authored by a coauthor (of a deceased scientist) in a given year that rank above the 85th percentile of papers published that year in terms of total citations. All columns control for publication year dummy variables, indicator variables for career age (year of publication minus year of first publication; career age zero is the omitted category), and the interaction terms between the career age indicators and each covariate of interest. In columns (2), (4), (6), and (8), the reference group contains coauthors at the middle level of cognitive distance from the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table A.6: Robustness Check: Alternative Bibliometric Distance Measures

<i>Dependent Variable:</i>	No. of Publications		Distribution of Citations	
	Unweighted	JIF-weighted	50th %ile of Cits	85th %ile of Cits
<i>Panel A: Production leave-out cognitive distance (baseline specification)</i>				
Death	0.037 (0.027)	0.030 (0.043)	0.030 (0.027)	−0.005 (0.031)
Close × Death	−0.118*** (0.044)	−0.202*** (0.068)	−0.143*** (0.041)	−0.112*** (0.040)
Distant × Death	−0.145*** (0.047)	−0.036 (0.061)	−0.103** (0.051)	−0.040 (0.057)
Log pseudo-likelihood	−333153.9	−861036.7	−296136.3	−204395.6
No. of Observations	155294	155294	154895	150256
<i>Panel B: WoS-restricted reference cosine similarity</i>				
Death	0.002 (0.028)	−0.007 (0.040)	0.004 (0.029)	−0.019 (0.037)
Close × Death	−0.032 (0.038)	−0.092 (0.067)	−0.061 (0.042)	−0.055 (0.043)
Distant × Death	−0.089* (0.046)	−0.058 (0.054)	−0.089* (0.047)	−0.071 (0.050)
Log pseudo-likelihood	−333423.1	−867473.9	−296369.3	−204520.7
No. of Observations	155287	155287	154888	150256
<i>Panel C: Soft cosine similarity</i>				
Death	0.023 (0.028)	0.002 (0.037)	0.015 (0.029)	−0.027 (0.035)
Close × Death	−0.063 (0.042)	−0.107 (0.066)	−0.081* (0.043)	−0.040 (0.047)
Distant × Death	−0.154*** (0.041)	−0.071 (0.064)	−0.114** (0.046)	−0.053 (0.065)
Log pseudo-likelihood	−333367.9	−866186.9	−296332.2	−204502
No. of Observations	155294	155294	154895	150256
<i>Panel D: Direct citation overlap</i>				
Death	−0.028 (0.029)	−0.041 (0.038)	−0.038 (0.026)	−0.037 (0.035)
Close × Death	0.039 (0.033)	−0.028 (0.059)	0.034 (0.034)	−0.035 (0.036)
Distant × Death	−0.070* (0.037)	−0.025 (0.044)	−0.051 (0.040)	−0.024 (0.052)
Log pseudo-likelihood	−315328.5	−828308.8	−280981.4	−195397.8
No. of Observations	141599	141599	141515	139037

Note: Robustness of the headline categorical specification (Equation (5)) under four cognitive-distance measures: the production leave-out distance in Panel A and three alternatives in Panels B–D. Panel A reproduces the production specification using the leave-out journal-citation-flow distance (matches Table 4 columns 2, 4, 6, 8). Panels B–D substitute the production distance with three bibliometric similarity measures constructed at the dyad level: a WoS-restricted reference cosine similarity, a soft cosine similarity over reference vectors, and a direct citation overlap measure (Section 6.2 defines the measures). For each alternative, the dyad-level distance distribution is percentile-ranked separately, and the close, middle, and distant indicators are assigned using the same percentile cutoffs as the production specification (close at or below the 40th percentile, distant at or above the 85th percentile, middle in between). All columns retain the production estimator (conditional quasi-maximum likelihood Poisson with pair fixed effects), the production controls (publication-year dummies, fourteen career-age indicators, and full distance-by-career-age interactions), and standard errors clustered at the level of the deceased scientist. Pairwise dyad-level Pearson correlations between the leave-out distance and the alternative measures are 0.66, 0.78, and 0.12 for Panels B, C, and D, respectively. Section 6.2 discusses the results, including the misclassification calibration that interprets the cross-panel differences. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table A.7: Robustness Check: OLS

<i>Dependent Variable:</i>	No. of Publications				Distribution of Citations			
	Unweighted		JIF-weighted		50th %ile of Cits		85th %ile of Cits	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Death	−0.088 (0.098)	0.001 (0.026)	−0.318* (0.190)	−0.015 (0.035)	−0.166** (0.072)	0.006 (0.019)	−0.110** (0.054)	−0.009 (0.016)
Distance × Death	0.080 (0.142)		0.387 (0.279)		0.211** (0.104)		0.121* (0.071)	
Close × Death		−0.069* (0.037)		−0.144** (0.058)		−0.083*** (0.027)		−0.051*** (0.017)
Distant × Death		−0.095*** (0.033)		−0.017 (0.049)		−0.050* (0.030)		−0.033 (0.024)
R ²	0.09	0.09	0.10	0.10	0.06	0.06	0.03	0.03
No. of Observations	155294	155294	155294	155294	155294	155294	155294	155294

Note: Robust standard errors (shown in parentheses) are clustered at the level of the deceased scientist. The estimates are drawn from coauthor-deceased author pairwise fixed effect specifications. Observations are made at the coauthor_{*i*}-deceased scientist_{*j*}-year_{*t*} level. The dependent variable of the models in columns (1) and (2) is the raw number of publications authored by a coauthor (of a deceased scientist) in the year of observation. The dependent variable of the models in columns (3) and (4) is the number of JIF-weighted publications authored by a coauthor (of a deceased scientist) in the year of observation. The dependent variables in columns (1) to (4) are measured in logarithmic form. The dependent variable of the models in columns (5) and (6) is the number of papers authored by a coauthor (of a deceased scientist) in a given year that rank above the 50th percentile of papers published that year in terms of total citations. The dependent variable of the models in columns (7) and (8) is the number of papers authored by a coauthor (of a deceased scientist) in a given year that rank above the 85th percentile of papers published that year in terms of total citations. The dependent variables in columns (5) to (8) are measured in $\log(x + 1)$ form. All columns control for publication year dummy variables, indicator variables for career age (year of publication minus year of first publication; career age zero is the omitted category), and the interaction terms between the career age indicators and each covariate of interest. In columns (2), (4), (6), and (8), the reference group contains coauthors at the middle level of cognitive distance from the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table A.8: Mean-Reversion Robustness: Active-Pair and Top-Quartile Pre-Period Intensity Restrictions

	Unweighted			JIF-weighted			Top-85%ile			Top-50%ile		
	(1) Base	(2) Act.	(3) TopQ	(4) Base	(5) Act.	(6) TopQ	(7) Base	(8) Act.	(9) TopQ	(10) Base	(11) Act.	(12) TopQ
Death	0.0370 (0.0272)	0.0378 (0.0272)	0.0153 (0.0303)	0.0298 (0.0432)	0.0308 (0.0432)	-0.0045 (0.0472)	-0.0052 (0.0306)	-0.0109 (0.0309)	-0.0529 (0.0366)	0.0302 (0.0268)	0.0313 (0.0266)	-0.0072 (0.0340)
Close $\times$ Death	-0.1184*** (0.0438)	-0.1193*** (0.0438)	-0.1580*** (0.0568)	-0.2022*** (0.0675)	-0.2032*** (0.0676)	-0.2265*** (0.0812)	-0.1115*** (0.0402)	-0.1080*** (0.0413)	-0.1707*** (0.0612)	-0.1434*** (0.0409)	-0.1444*** (0.0410)	-0.1820*** (0.0538)
Distant $\times$ Death	-0.1451*** (0.0474)	-0.1451*** (0.0475)	-0.1575** (0.0637)	-0.0361 (0.0613)	-0.0360 (0.0612)	-0.0500 (0.1185)	-0.0404 (0.0565)	-0.0495 (0.0566)	-0.0442 (0.0741)	-0.1029** (0.0510)	-0.1049** (0.0508)	-0.1022 (0.0665)
<i>N</i>	155294	154936	62055	155294	154936	56211	150256	145056	47690	154895	153936	59954

Note: Mean-reversion robustness check applied to the headline Equation (5) specification across four research output measures: the unweighted publication count, JIF-weighted publications, the count of papers above the 85th citation percentile, and the count above the 50th citation percentile. Within each three-column outcome block, the columns labeled Base, Act., and TopQ correspond, respectively, to the full sample (which reproduces the corresponding columns of Table 4), the Active-Pair restriction (dyads with strictly positive pre-event total output for the relevant outcome), and the Top-Quartile restriction (dyads in the top quartile of pre-event mean output for the relevant outcome, computed among the active dyads). Section 6.2 describes the mean-reversion logic and discusses the results. Estimates are drawn from conditional quasi-maximum likelihood Poisson specifications. Standard errors (in parentheses) are clustered at the level of the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table A.9: Robustness Check: Rectangular Panel with Zero-Output Years Filled

<i>Dependent Variable:</i>	No. of Publications				Distribution of Citations			
	Unweighted		JIF-weighted		50th %ile of Cits		85th %ile of Cits	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Death	-0.134 (0.134)	0.023 (0.032)	-0.494** (0.199)	0.020 (0.046)	-0.243* (0.137)	0.015 (0.032)	-0.242 (0.163)	-0.020 (0.034)
Distance $\times$ Death	0.141 (0.191)		0.671** (0.290)		0.294 (0.199)		0.268 (0.227)	
Close $\times$ Death		-0.127** (0.051)		-0.214*** (0.075)		-0.151*** (0.049)		-0.117** (0.049)
Distant $\times$ Death		-0.165*** (0.057)		-0.055 (0.065)		-0.124** (0.061)		-0.065 (0.067)
Log pseudo-likelihood	-413494.7	-413279	-1104561	-1104043	-351003.5	-350877.7	-228343.3	-228305
No. of Observations	199118	199118	199118	199118	198270	198270	189561	189561

Note: Re-estimation of the Table 4 specifications on a rectangular panel in which every interior zero-output year between a coauthor's first and last observed publication year enters as an explicit zero-count row (43,824 filled zeros, 22 percent of the resulting panel). Years after the last observed publication are excluded because exit from the Web of Science sample cannot be distinguished from exit from research. The share of filled zero years is 23, 20, and 25 percent in the close, middle, and distant bands, respectively, and is slightly lower after the death year in all three bands. Column structure, controls, and clustering follow Table 4: odd columns report the continuous specification of Equation (4) and even columns the categorical specification of Equation (5), with the middle-distance group as the reference category; all columns control for publication year dummies and career-age indicators interacted with the covariates of interest. Robust standard errors (in parentheses) are clustered at the level of the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table A.10: Poisson Event Studies by Distance Subsample

	Close [0,40]				Middle (40,85]				Distant (85,100]			
	(1) Unwt	(2) IPF	(3) T85	(4) T50	(5) Unwt	(6) IPF	(7) T85	(8) T50	(9) Unwt	(10) IPF	(11) T85	(12) T50
$t \leq -6$	-0.1121 (0.0901)	-0.2066** (0.1037)	-0.1536* (0.0889)	-0.1457 (0.0962)	-0.0223 (0.0623)	0.0075 (0.0946)	0.0306 (0.1152)	-0.0271 (0.0771)	0.0350 (0.0894)	0.0265 (0.1210)	-0.0539 (0.1270)	0.0076 (0.0791)
$t = -5$	-0.0581 (0.0614)	-0.1232** (0.0581)	-0.0610 (0.0722)	-0.0685 (0.0670)	-0.0136 (0.0491)	-0.0185 (0.0687)	0.0482 (0.0939)	-0.0107 (0.0620)	0.0395 (0.0728)	0.1087 (0.1240)	0.0236 (0.1148)	0.0114 (0.0764)
$t = -4$	-0.0307 (0.0455)	-0.0950* (0.0513)	0.0095 (0.0816)	-0.0379 (0.0513)	0.0027 (0.0440)	0.0106 (0.0601)	0.0436 (0.0775)	0.0016 (0.0525)	0.0255 (0.0571)	0.0560 (0.0864)	0.0561 (0.0671)	-0.0024 (0.0584)
$t = -3$	-0.0261 (0.0285)	-0.0221 (0.0365)	-0.0502 (0.0489)	-0.0218 (0.0316)	0.0052 (0.0306)	-0.0004 (0.0431)	0.0194 (0.0558)	0.0020 (0.0373)	0.0360 (0.0453)	0.0194 (0.0604)	-0.0298 (0.0777)	0.0059 (0.0470)
$t = -2$	-0.0145 (0.0215)	0.0247 (0.0489)	0.0050 (0.0371)	-0.0070 (0.0249)	0.0002 (0.0272)	0.0311 (0.0351)	0.1257** (0.0637)	0.0286 (0.0324)	0.0523 (0.0438)	0.0925** (0.0447)	0.0849 (0.0801)	0.0722 (0.0489)
$t = 0$	0.0021 (0.0288)	-0.0262 (0.0306)	-0.0290 (0.0340)	-0.0002 (0.0302)	0.0267 (0.0179)	0.0027 (0.0234)	0.0417 (0.0312)	0.0420** (0.0206)	0.0290 (0.0324)	0.0529 (0.0435)	0.0961 (0.0725)	0.0606 (0.0392)
$t = +1$	0.0004 (0.0443)	-0.0718 (0.0508)	-0.0236 (0.0762)	-0.0330 (0.0399)	0.0024 (0.0303)	-0.0416 (0.0373)	0.0067 (0.0390)	0.0206 (0.0325)	-0.0153 (0.0425)	0.0588 (0.0511)	0.0458 (0.0684)	0.0232 (0.0399)
$t = +2$	-0.0740 (0.0606)	-0.1263* (0.0668)	-0.1247* (0.0651)	-0.1260** (0.0537)	0.0121 (0.0352)	-0.0594 (0.0427)	-0.0314 (0.0398)	0.0038 (0.0361)	-0.0753 (0.0668)	0.0226 (0.0845)	-0.0334 (0.0868)	-0.0502 (0.0634)
$t = +3$	-0.0403 (0.0744)	-0.1113 (0.0842)	-0.0742 (0.0755)	-0.0807 (0.0675)	-0.0201 (0.0536)	-0.0981 (0.0606)	-0.0660 (0.0580)	-0.0315 (0.0598)	-0.0925 (0.0808)	0.0233 (0.1067)	0.0140 (0.0966)	-0.0413 (0.0871)
$t = +4$	-0.1035 (0.0888)	-0.1863* (0.1009)	-0.1281 (0.0913)	-0.1311 (0.0863)	-0.0259 (0.0637)	-0.1215* (0.0727)	-0.0943 (0.0733)	-0.0325 (0.0712)	-0.0852 (0.0948)	0.0215 (0.1494)	0.0736 (0.1221)	-0.0128 (0.0880)
$t = +5$	-0.1756* (0.0965)	-0.2648** (0.1100)	-0.2042* (0.1097)	-0.2168** (0.0946)	-0.0440 (0.0820)	-0.1395 (0.0890)	-0.1127 (0.0925)	-0.0578 (0.0889)	-0.0729 (0.1115)	0.0813 (0.1487)	0.0850 (0.1224)	-0.0117 (0.0985)
$t = +6$	-0.1535 (0.1054)	-0.2463** (0.1245)	-0.2254* (0.1311)	-0.1863* (0.1068)	-0.0711 (0.0937)	-0.1796* (0.1019)	-0.1778 (0.1198)	-0.0953 (0.1037)	-0.1128 (0.1365)	0.1051 (0.1930)	0.1120 (0.1661)	-0.0108 (0.1273)
$t = +7$	-0.1600 (0.1181)	-0.2746** (0.1385)	-0.1901 (0.1489)	-0.2034 (0.1241)	-0.0754 (0.1223)	-0.2115 (0.1373)	-0.1871 (0.1349)	-0.1096 (0.1358)	-0.1321 (0.1586)	0.0333 (0.2133)	0.0425 (0.2041)	-0.0362 (0.1477)
$t = +8$	-0.1763 (0.1259)	-0.3256** (0.1510)	-0.2905* (0.1640)	-0.2399* (0.1371)	-0.1148 (0.1416)	-0.2681* (0.1627)	-0.2230 (0.1528)	-0.1483 (0.1539)	-0.1950 (0.1757)	0.0317 (0.2352)	0.0093 (0.2079)	-0.0903 (0.1679)
$t = +9$	-0.1851 (0.1326)	-0.3662** (0.1560)	-0.3188* (0.1802)	-0.2542* (0.1484)	-0.1644 (0.1645)	-0.3655* (0.1973)	-0.3106* (0.1847)	-0.2020 (0.1803)	-0.1827 (0.1830)	0.0290 (0.2495)	-0.0493 (0.2112)	-0.0823 (0.1662)
$t = +10$	-0.2238 (0.1436)	-0.3970** (0.1693)	-0.3137 (0.1939)	-0.2966* (0.1627)	-0.2258 (0.1918)	-0.4537* (0.2437)	-0.3629* (0.2179)	-0.2747 (0.2112)	-0.1851 (0.1971)	0.0885 (0.2825)	0.0924 (0.2626)	-0.0723 (0.1826)
$t \geq +11$	-0.3036 (0.1863)	-0.4947** (0.2206)	-0.4429 (0.2718)	-0.4164* (0.2196)	-0.2705 (0.2360)	-0.5058* (0.3073)	-0.3944 (0.2673)	-0.3231 (0.2615)	-0.2657 (0.2328)	0.0220 (0.3311)	-0.0766 (0.2886)	-0.1753 (0.2211)
Joint pretrend p (leads -5 to -2)	0.888	0.191	0.228	0.872	0.846	0.720	0.036	0.802	0.774	0.104	0.625	0.366
No. of Observations	57639	57639	56076	57531	75798	75798	73679	75622	21857	21857	20501	21742

Note: Fixed-effects Poisson event studies estimated separately on each cognitive-distance band via `ppmlhdf`, with dyad and publication-year fixed effects: the count-data counterpart of the linear BJS estimates in Table A.4 and the band-level analogue of the pooled Poisson sensitivity in Table A.3. Event-time dummies are binned at $t \leq -6$ and $t \geq +11$, with $t = -1$ as the reference period. Columns (1)–(4) cover the close band (0–40th percentiles), columns (5)–(8) the middle band (40–85th percentiles), and columns (9)–(12) the distant band (85–100th percentiles); within each band the four columns correspond to the unweighted publication count, IIF-weighted publications, the count of papers above the 85th citation percentile, and the count above the 50th citation percentile. The joint pretrend row reports the p-value of the Wald test that the four non-binned leads are jointly zero; the middle-band top-85%ile cell rejects at the five percent level ($p = 0.036$), the only rejection among the twelve cells, within the expected false-rejection rate at this significance level. Section 6.3 discusses the results. Standard errors (in parentheses) are clustered at the level of the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table A.11: Within-Subgroup Tests of Equality Between Close and Distant Coefficients

	Unweighted	JIF-weighted	Top-85%ile	Top-50%ile
Elite vs Non-elite	-0.0543 (0.1595) [0.734]	-0.1417 (0.2201) [0.520]	-0.0923 (0.1347) [0.493]	-0.0689 (0.1382) [0.618]
Senior vs Junior	-0.1549 (0.2011) [0.441]	-0.2322 (0.2340) [0.321]	-0.1648 (0.1920) [0.391]	-0.1664 (0.2016) [0.409]
Same vs Different Inst.	0.3786*** (0.1269) [0.003]	0.3329 (0.2299) [0.148]	0.4070** (0.1669) [0.015]	0.3889** (0.1595) [0.015]
Recent vs Former	0.0121 (0.0764) [0.874]	-0.2319** (0.1029) [0.024]	-0.1126 (0.1036) [0.277]	-0.0701 (0.0774) [0.365]
Regular vs Infrequent	-0.0559 (0.1074) [0.603]	-0.2142 (0.1538) [0.164]	0.0536 (0.1276) [0.674]	-0.0217 (0.1037) [0.834]
Deceased First vs Non-first	-0.1239 (0.1648) [0.452]	0.0847 (0.2045) [0.679]	0.1351 (0.2435) [0.579]	-0.1277 (0.1837) [0.487]
Deceased Last vs Non-last	0.1347 (0.0964) [0.162]	-0.0824 (0.1015) [0.417]	-0.0243 (0.1142) [0.831]	-0.0009 (0.0958) [0.992]

Note: Each cell reports the gap $\hat{\beta}_{\text{close}} - \hat{\beta}_{\text{distant}}$ obtained from a single regression on the first-named subgroup of each row label (elite, senior, same-institution, recent, regular, deceased-first-author, and deceased-last-author dyads, respectively), with the standard error of the gap in parentheses and the p-value of the within-row Wald test of $H_0 : \beta_{\text{close}} = \beta_{\text{distant}}$ in brackets. The seven subgroup definitions partition the 9,263 dyads along, respectively, coauthor scientific standing (elite vs. non-elite, defined as top-decile cumulative citations at the time of the star's death), career stage measured at the time of the star's death (senior vs. junior coauthor), institutional proximity (same vs. different institution), recency of the most recent pre-event collaboration (recent vs. former), collaboration intensity (regular vs. infrequent), the deceased scientist's first-author position (first vs. non-first), and the deceased scientist's last-author position (last vs. non-last). The four columns correspond to the four research output measures defined in Table A.1. Estimates are drawn from conditional quasi-maximum likelihood Poisson specifications matching Equation (5); standard errors are clustered at the level of the deceased scientist. **Stars apply to the within-row Wald test of equality, not to the individual $\text{Close} \times \text{Death}$ and $\text{Distant} \times \text{Death}$ coefficients;** see Figure A.4 for the corresponding 95 percent confidence intervals and Section 6.3 for the discussion. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

Table A.12: Cross-Subgroup Tests of Equality Within Each Cognitive-Distance Band

	Unweighted	JIF-weighted	Top-85%ile	Top-50%ile
<i>Panel A: Close × Death, cross-subgroup gap (flagged minus complement)</i>				
Elite vs Non-elite	-0.0676 (0.0834) [0.418]	-0.0078 (0.0931) [0.933]	-0.0144 (0.0818) [0.860]	-0.0246 (0.0646) [0.703]
Senior vs Junior	-0.0540 (0.1175) [0.646]	0.0314 (0.1296) [0.809]	-0.0333 (0.1521) [0.827]	-0.0120 (0.1340) [0.928]
Same vs Different Inst.	0.0444 (0.0785) [0.572]	0.0847 (0.1017) [0.405]	0.2115* (0.1089) [0.052]	0.1563** (0.0731) [0.032]
Recent vs Former	-0.0679 (0.0779) [0.384]	-0.0501 (0.0836) [0.549]	-0.0541 (0.0938) [0.565]	-0.1143 (0.0836) [0.172]
Regular vs Infrequent	-0.0858 (0.0865) [0.321]	-0.0502 (0.0880) [0.568]	0.0933 (0.1079) [0.387]	0.0135 (0.0852) [0.874]
Deceased First vs Non-first	-0.1252 (0.1620) [0.440]	0.1624 (0.1636) [0.321]	0.2182 (0.2773) [0.431]	-0.1258 (0.1929) [0.514]
Deceased Last vs Non-last	0.1659* (0.0934) [0.076]	0.0445 (0.0860) [0.605]	0.0965 (0.0891) [0.279]	0.1028 (0.0625) [0.100]
<i>Panel B: Distant × Death, cross-subgroup gap (flagged minus complement)</i>				
Elite vs Non-elite	0.0782 (0.1373) [0.569]	0.0198 (0.1926) [0.918]	0.0783 (0.1433) [0.585]	0.0519 (0.1297) [0.689]
Senior vs Junior	0.1910 (0.1892) [0.313]	0.1673 (0.2628) [0.524]	0.1130 (0.2266) [0.618]	0.1722 (0.2046) [0.400]
Same vs Different Inst.	-0.3532** (0.1435) [0.014]	-0.4884** (0.2407) [0.043]	-0.3313* (0.1851) [0.073]	-0.3289* (0.1736) [0.058]
Recent vs Former	-0.0330 (0.0750) [0.660]	0.0913 (0.0979) [0.351]	0.0560 (0.1224) [0.647]	-0.0374 (0.0851) [0.661]
Regular vs Infrequent	0.0152 (0.0714) [0.832]	0.0061 (0.1098) [0.956]	-0.0453 (0.1070) [0.672]	-0.0073 (0.0720) [0.920]
Deceased First vs Non-first	0.0433 (0.1792) [0.809]	-0.0819 (0.2302) [0.722]	0.0246 (0.2141) [0.909]	-0.0247 (0.2076) [0.905]
Deceased Last vs Non-last	0.0451 (0.0824) [0.585]	-0.0435 (0.1242) [0.726]	0.0436 (0.1404) [0.756]	0.0639 (0.1071) [0.551]

Note: Each cell reports the difference in one band's interaction coefficient between the two sides of a partition, $\hat{\beta}(\text{flagged}) - \hat{\beta}(\text{complement})$, with the standard error of the difference in parentheses and the p-value of the Wald test that the difference is zero in brackets. Panel A reports the difference in the *Close* × *Death* coefficient and Panel B the difference in the *Distant* × *Death* coefficient; within a given row and column, the two panels come from the same regression. Each regression is a fully interacted specification in which every regressor is allowed to differ between the flagged subgroup and its complement. The seven partitions are those of Table A.11. The four columns run unweighted publications, JIF-weighted publications, the count of papers above the 85th citation percentile, and the count above the 50th citation percentile, following the ordering of Table A.1 rather than that of Table 4. Estimates are drawn from conditional quasi-maximum likelihood Poisson specifications; standard errors are clustered at the level of the deceased scientist. **Stars apply to the cross-subgroup test of equality, not to the underlying band coefficients**, which are plotted in Figure 5. Section 6.3 discusses the results. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.


Table A.13: Joint Moderation of the Cognitive-Distance Effect by Tie Strength, Standing, and Roles

	Unweighted				JIF-weighted				Top-85%ile				Top-50%ile			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Death	0.0370 (0.0272)	-0.0367 (0.0339)	0.0150 (0.0331)	0.0336 (0.0329)	0.0298 (0.0432)	-0.0433 (0.0499)	0.0507 (0.0462)	0.0729 (0.0468)	-0.0052 (0.0306)	-0.1104** (0.0459)	-0.0282 (0.0441)	0.0029 (0.0433)	0.0302 (0.0268)	-0.0534 (0.0355)	0.0050 (0.0339)	0.0270 (0.0322)
Close × Death	-0.1184*** (0.0438)	-0.1177*** (0.0445)	-0.1190*** (0.0432)	-0.1203*** (0.0437)	-0.2022*** (0.0675)	-0.1966*** (0.0689)	-0.1982*** (0.0657)	-0.1972*** (0.0664)	-0.1115*** (0.0402)	-0.1051** (0.0437)	-0.1097** (0.0433)	-0.1087** (0.0431)	-0.1434*** (0.0409)	-0.1409*** (0.0427)	-0.1438*** (0.0417)	-0.1443*** (0.0421)
Distant × Death	-0.1451*** (0.0474)	-0.1377*** (0.0475)	-0.1464*** (0.0482)	-0.1415*** (0.0485)	-0.0361 (0.0613)	-0.0263 (0.0603)	-0.0372 (0.0617)	-0.0310 (0.0602)	-0.0404 (0.0565)	-0.0302 (0.0570)	-0.0383 (0.0591)	-0.0318 (0.0597)	-0.1029** (0.0510)	-0.0945* (0.0511)	-0.1030** (0.0522)	-0.0979* (0.0532)
Regular × Death		0.0859** (0.0367)	0.0939** (0.0368)	0.1058*** (0.0367)		0.1172** (0.0579)	0.1265** (0.0573)	0.1324** (0.0588)		0.0747 (0.0557)	0.0831 (0.0566)	0.0817 (0.0580)		0.0896** (0.0418)	0.0984** (0.0423)	0.1057** (0.0428)
Recent × Death		0.1072*** (0.0349)	0.1024*** (0.0360)	0.0976*** (0.0344)		0.0925* (0.0511)	0.0871* (0.0520)	0.0811 (0.0502)		0.1628*** (0.0530)	0.1555*** (0.0538)	0.1463*** (0.0503)		0.1225*** (0.0408)	0.1169*** (0.0418)	0.1104*** (0.0393)
Elite × Death			-0.1166*** (0.0353)	-0.1178*** (0.0345)			-0.1900*** (0.0373)	-0.1892*** (0.0358)			-0.1683*** (0.0384)	-0.1695*** (0.0371)			-0.1331*** (0.0358)	-0.1342*** (0.0349)
Senior × Death			-0.0524 (0.0584)	-0.0532 (0.0573)			-0.0712 (0.0714)	-0.0725 (0.0705)			-0.0289 (0.0721)	-0.0273 (0.0698)			-0.0265 (0.0614)	-0.0266 (0.0598)
Same Inst. × Death				-0.1279** (0.0531)				-0.1124* (0.0640)				-0.1163* (0.0615)				-0.1184** (0.0586)
First Author × Death				-0.0135 (0.0843)				0.0006 (0.1001)				-0.0524 (0.1325)				-0.0329 (0.0935)
Last Author × Death				-0.0282 (0.0305)				-0.0550 (0.0490)				-0.0873* (0.0483)				-0.0438 (0.0371)
N	155294	155294	155294	155294	155294	155294	155294	155294	150256	150256	150256	150256	154895	154895	154895	154895

Note: Joint-moderation robustness check applied to the headline Equation (5) specification across four research output measures: the unweighted publication count, JIF-weighted publications, the count of papers above the 85th citation percentile, and the count above the 50th citation percentile. Within each four-column block, columns (1) through (4) progressively add seven moderator-by-Death interactions in three nested layers. Column (1) reproduces Equation (5) and matches the corresponding column of Table 4. Column (2) adds tie-strength interactions (*Regular × Death* and *Recent × Death*). Column (3) further adds standing interactions (*Elite × Death* and *Senior × Death*). Column (4) further adds role interactions (*Same Inst. × Death*, *First Author × Death*, *Last Author × Death*) to yield the fully joint specification. The build-up tests whether the focal *Close × Death* and *Distant × Death* coefficients are absorbed by the seven correlated dyadic moderators discussed in Section 6.3. Estimates are drawn from conditional quasi-maximum likelihood Poisson specifications. Standard errors (in parentheses) are clustered at the level of the deceased scientist. *Significant at the 10 percent level, **significant at the 5 percent level, ***significant at the 1 percent level.

B Supplementary Details Regarding Sample Collection

Figure B.1 Example Trend in the Annual Number of Publications of a Deceased Scientist



Note: Figure B.1 displays the trend in the annual number of publications for one example deceased scientist in the sample, illustrating the sharp and permanent decline used to identify candidate scientists.

Table B.1: Multidisciplinary Scope of the Life Sciences

BIODIVERSITY CONSERVATION
PSYCHOLOGY, BIOLOGICAL
BIOCHEMICAL RESEARCH METHODS
BIOCHEMISTRY & MOLECULAR BIOLOGY
BIOLOGY
BIOPHYSICS
BIOTECHNOLOGY & APPLIED MICROBIOLOGY
CELL BIOLOGY
CHEMISTRY, MEDICINAL
EMERGENCY MEDICINE
DENTISTRY/ORAL SURGERY & MEDICINE
EVOLUTIONARY BIOLOGY
DEVELOPMENTAL BIOLOGY
ENGINEERING, BIOMEDICAL
MEDICINE, LEGAL
MARINE & FRESHWATER BIOLOGY
MEDICAL INFORMATICS
MEDICAL LABORATORY TECHNOLOGY
MEDICINE, GENERAL & INTERNAL
MEDICINE, RESEARCH & EXPERIMENTAL
MATERIALS SCIENCE, BIOMATERIALS
MICROBIOLOGY
RADIOLOGY, NUCLEAR MEDICINE & MEDICAL IMAGING
REPRODUCTIVE BIOLOGY
SOCIAL SCIENCES, BIOMEDICAL
TROPICAL MEDICINE
CRITICAL CARE MEDICINE
MATHEMATICAL & COMPUTATIONAL BIOLOGY
INTEGRATIVE & COMPLEMENTARY MEDICINE
MEDICAL ETHICS

Note: Table B.1 lists the 30 WoS subject categories that define the life sciences sample, broadly defined.

C Derivation of the Two-Channel Decomposition

This appendix formalizes the framework of Section 2.2. It states the sufficient conditions on the curvature parameters and channel weights, derives the comparative statics underlying Hypotheses 1 and 2, develops the channel-weight reading and an auxiliary competition-relief channel that jointly organize the subgroup patterns in Section 6.3, and verifies that the qualitative predictions hold under alternative functional forms.

C.1 Setup

Consider a surviving coauthor i and a deceased star j at cognitive distance $d \equiv d_{ij} \in [0, 1]$. Let

$$S(d) = (1 - d)^\alpha, \quad \alpha > 0, \quad (6)$$

$$C(d) = d^\beta, \quad \beta > 0, \quad (7)$$

denote the spillover and complementarity channels through which j contributed to i 's productivity prior to j 's death. S is strictly decreasing and C is strictly increasing on $(0, 1)$. The curvature parameters α and β govern how steeply each channel attenuates away from its peak ($d = 0$ for S , $d = 1$ for C); larger values of α or β correspond to faster decay.

For outcome $y \in \{q, Q\}$ (quantity and quality, respectively), the productivity loss following j 's death is the linear combination

$$L_y(d) = w_y^S S(d) + w_y^C C(d), \quad (8)$$

where $w_y^S, w_y^C \geq 0$ are outcome-specific channel loadings. Two assumptions govern the loadings:

(A1) *Quality is spillover-dominated*: $w_Q^S > 0$ and $w_Q^S \gg w_Q^C \geq 0$.

(A2) *Quantity is spillover-leaning*: $w_q^S > w_q^C > 0$.

Quality requires deep tacit understanding, sourced primarily through close partnerships. Quantity benefits from both routine co-production (spillover) and rare specialist input (complementarity). The non-monotone relationship between cognitive distance and innovation outcomes documented by Nooteboom et al. (2007) grounds the two-channel structure itself. Assumption (A1) rests on two findings. Absorptive capacity is bounded by prior related knowledge (Cohen and Levinthal, 1990), so high-impact innovation requires depth of tacit transfer between cognitively close partners more than breadth from distant expertise. In our setting, Azoulay et al. (2010) report that quality declines following star-coauthor death are concentrated among collaborators who were close to the deceased along multiple proximity dimensions. The rationale for Assumption (A2) is compositional: the spillover channel mediates the bulk of routine joint output, which constitutes the dominant share of a researcher's pre-event quantity (Boschma, 2005), while a smaller share draws on complementary specialist input, leaving $w_q^C > 0$. We treat (A1) and (A2) as maintained assumptions rather than as derivations from deeper primitives. The empirical patterns in Section 6 are consistent with them but do not separately identify the loadings from the curvature parameters.

C.2 Quality: monotonic decline

Proposition 1 (Quality, spillover-dominated regime). *Suppose $\alpha > 0$, $\beta \geq 1$, and $w_Q^S > 0$. Then for any $\epsilon > 0$, there exists $\bar{r} > 0$ such that $L_Q(d)$ is strictly decreasing on $(0, 1 - \epsilon)$ whenever*

$$w_Q^C/w_Q^S < \bar{r}.$$

Derivation. In the limit $w_Q^C \rightarrow 0$, Equation (8) reduces to $L_Q(d) = w_Q^S (1-d)^\alpha$, which is strictly decreasing on $(0, 1)$ for any $\alpha > 0$. For $w_Q^C > 0$, differentiation yields

$$L'_Q(d) = -\alpha w_Q^S (1-d)^{\alpha-1} + \beta w_Q^C d^{\beta-1}.$$

The first term dominates whenever $w_Q^C/w_Q^S < \alpha (1-d)^{\alpha-1}/[\beta d^{\beta-1}]$. For $\beta \geq 1$, the right-hand side has a strictly positive infimum on $(0, 1-\epsilon)$, so the inequality is satisfied uniformly when the loading ratio is sufficiently small, and L_Q is strictly decreasing on $(0, 1-\epsilon)$. ■

Remark. The conditions $\alpha > 0$ and $\beta \geq 1$ in Proposition 1 are strictly weaker than the working assumption $\alpha > 1, \beta > 1$ maintained in Section 2.2. The quality-monotonicity result therefore holds a fortiori once the strict-curvature condition required by Proposition 2 is imposed. Assumption (A1) is the economic content of the threshold condition: quality loadings are spillover-dominated in the sense that the ratio w_Q^C/w_Q^S lies below $\bar{r}$.

Empirical anchoring. The empirically relevant evaluation points $d_{\text{close}} \approx 0.46$, $d_{\text{middle}} \approx 0.69$, $d_{\text{distant}} \approx 0.86$ from Table 3 all lie in the interior, so the theoretical prediction $L_Q(d_{\text{close}}) > L_Q(d_{\text{middle}}) > L_Q(d_{\text{distant}})$ obtains. Empirically, the close-middle drop is large while the middle-distant drop is small enough to be statistically indistinguishable from zero, matching columns (4) and (8) of Table 4: only the close-distance group experiences a statistically significant quality decline (−18.3% in JIF-weighted publications, −10.6% in publications above the 85th citation percentile).

C.3 Quantity: U-shaped loss with interior minimum

Proposition 2 (Quantity). Suppose $\alpha > 1, \beta > 1$, and $w_q^S, w_q^C > 0$. Then $L_q(d)$ is strictly convex on $(0, 1)$ and has a unique interior minimum $d^* \in (0, 1)$, with L_q strictly decreasing on $(0, d^*)$ and strictly increasing on $(d^*, 1)$.

Derivation. Differentiating Equation (8) once and twice in d yields

$$L'_q(d) = -\alpha w_q^S (1-d)^{\alpha-1} + \beta w_q^C d^{\beta-1}, \quad (9)$$

$$L''_q(d) = \alpha(\alpha-1) w_q^S (1-d)^{\alpha-2} + \beta(\beta-1) w_q^C d^{\beta-2}. \quad (10)$$

At the boundaries: $L'_q(0) = -\alpha w_q^S < 0$ (using $\beta > 1$, which makes $d^{\beta-1}|_{d=0} = 0$); and $L'_q(1) = \beta w_q^C > 0$ (using $\alpha > 1$, which makes $(1-d)^{\alpha-1}|_{d=1} = 0$). Equation (10) is strictly positive on $(0, 1)$ when $\alpha > 1$ and $\beta > 1$, so L_q is strictly convex. By the intermediate value theorem applied to the continuous and strictly increasing L'_q , there exists a unique $d^* \in (0, 1)$ at which $L'_q(d^*) = 0$, and by strict convexity d^* is a strict global minimum on $[0, 1]$.

Setting (9) to zero and rearranging:

$$\frac{(d^*)^{\beta-1}}{(1-d^*)^{\alpha-1}} = \frac{\alpha w_q^S}{\beta w_q^C}. \quad (11)$$

In the symmetric case $\alpha = \beta$ and $w_q^S = w_q^C$, Equation (11) becomes $[d^*/(1-d^*)]^{\beta-1} = 1$. Since $\beta > 1$ by hypothesis, this reduces to $d^*/(1-d^*) = 1$, giving $d^* = 1/2$. More generally, a larger w_q^S moves d^* rightward and a larger w_q^C moves d^* leftward. These comparative statics on the channel loadings hold unconditionally (the direction of $\partial d^*/\partial \alpha$ and $\partial d^*/\partial \beta$ depends on the location of d^* and is not used below). Holding curvatures at the symmetric case $\alpha = \beta$, the spillover-leaning regime $w_q^S > w_q^C$ posited in (A2) therefore shifts d^* rightward of $1/2$ via

Equation (11). For general $\alpha \neq \beta$, the absolute location of d^* depends on the relative magnitudes of all four parameters. ■

Empirical anchoring. The pooled coefficients in columns (2) and (6) of Table 4 are consistent with $d^* \in [0.69, 0.86]$. The close-band loss (−11.1% in raw publications, −13.3% in publications above the 50th percentile) sits well to the left of d^* on the falling arm of the U. The middle-band coefficient is statistically null, locating d^* at or just above the middle regime. The distant-band loss (−13.5%, −9.8%) sits past d^* on the rising arm. The inverted-U pattern in observed productivity changes is the mirror of the U-shape in L_q predicted by Proposition 2. The same coefficients anchor the empirical location of d^* and verify the U-shape prediction. We view this dual role as consistent with the appendix’s interpretive purpose rather than as an independent test. The band-specific BJS estimates in Table A.4 provide a complementary cut: the close-negative, middle-null, distant-positive gradient is monotone in cognitive distance, with the middle-band ten-year coefficient statistically null on three of the four outcomes, but the positive sign in the distant band does not survive the count-data link (Table A.10) and is attributable to the linear estimator rather than to a band-specific treatment effect. The within-field competition-relief mechanism formalized in Subsection C.5 below is therefore retained as a theoretical possibility that the evidence does not require. The location of d^* is not point-identified by the categorical specification (5). Identification would require either a continuous-treatment specification with sufficient curvature or a structural restriction on at least one of the four parameters.

C.4 The boundary case $\alpha = \beta = 1$

When $\alpha = 1$ and $\beta = 1$, the loss function reduces to

$$L_q(d) = w_q^S + (w_q^C - w_q^S) d,$$

which is linear and monotonic (decreasing if $w_q^S > w_q^C$, increasing if $w_q^S < w_q^C$, constant if $w_q^S = w_q^C$). The strict inequalities $\alpha > 1$ and $\beta > 1$ in Proposition 2 are therefore necessary as a pair: when both channels are linear ($\alpha = \beta = 1$), L_q collapses to a linear function and the U-shape vanishes. The boundary case also sharpens the intuition behind the resilience prediction of Section 2.2: intermediate-distance coauthors are most resilient not merely because neither channel peaks at intermediate d , which is equally true in the linear case, but because strict curvature concentrates each channel’s loss near its own peak, pushing the combined loss at intermediate d strictly below the corresponding weighted average of the losses at the two endpoints.

C.5 Heterogeneity in channel weights and competition-relief activation

We organize the subgroup heterogeneity reported in Section 6.3 through two interpretive mappings onto the channel decomposition. The first treats absorptive-capacity, collaboration-intensity, and authorship-role variation as shifts in the channel weights w_y^S, w_y^C , preserving the sign of L_y . The second introduces an auxiliary competition-relief channel that activates at one tail under specific subgroup conditions and can flip the sign of L_y . The framework does not pre-specify the direction of either mapping; we adopt the directions consistent with the empirical patterns and with mechanism-level evidence in the literature.

Channel-weight reading.

- *Junior and non-elite coauthors.* Lower own absorptive capacity (Cohen and Levinthal, 1990) combined with narrower external collaboration portfolios implies that a larger fraction of pre-event output flowed through both channels, so both w_y^S and w_y^C are larger and the distance profile of the loss is correspondingly sharper on both the quality and the quantity margin. Elite and senior coauthors maintain external networks that buffer both channels, so their response varies little across the distance bands.
- *Recent collaboration.* Active dependency on j at the moment of death scales both w_y^S and w_y^C upward; recent coauthors (collaboration within the three years before the death) face larger losses than former coauthors, whose pre-event dependency had already attenuated.
- *Collaboration frequency.* Regular collaboration generates established routines and skill redundancy that absorb the loss of any single channel, attenuating the effective w_y^S and w_y^C for the routinized portion of joint output; infrequent dyads do not benefit from this routinization. Section 6.3 reports an exception to this prediction: cognitively close regular coauthors still experience significant losses, suggesting that the routinization mechanism has limits when joint output is tightly coupled to shared tacit knowledge.
- *Deceased first or last author.* Dyads in which the deceased was consistently the first or last author plausibly belong to larger research programs in which the surviving coauthor's pre-event dependency on the star's specific intellectual input is smaller; the star's contribution is concentrated in roles that substitute more readily across team members, with managerial inputs (funding, laboratory infrastructure) for last-author dyads and hands-on expertise within established workflows for first-author dyads. Both w_y^S and w_y^C are attenuated relative to peer-to-peer dyads, consistent with the muted death effects in Figure 5 for these subgroups.

Competition-relief activation. Once active, an auxiliary competition-relief channel modifies Equation (8) to

$$\tilde{L}_y(d) = w_y^S S(d) + w_y^C C(d) - w_y^R R(d), \quad (12)$$

where $R(d) \geq 0$ measures the relief magnitude at distance d . The channel operates under two modes, each specified at the three cognitive-distance bands $\{d_{\text{close}}, d_{\text{middle}}, d_{\text{distant}}\}$.

- *Within-field competition relief (not required by the evidence).* For distant-cognitive partners in the same broad field, the deceased competed for journal slots and citation share without sharing tacit knowledge (Borjas and Doran, 2015; Azoulay et al., 2019), suggesting a relief margin at the distant band: $R(d_{\text{distant}}) > 0$ with $R(d_{\text{close}}) \approx R(d_{\text{middle}}) \approx 0$. The linear band-specific BJS estimates in Table A.4 initially appear to display this pattern, returning positive long-run distant-band coefficients. The band-specific Poisson event studies in Table A.10, however, show that the positivity does not survive the count-data link, so the reversal is attributable to the linear functional form rather than to a relief channel (Section 6.3). We therefore retain this mode as a theoretical possibility only.
- *Same-institution resource-competition relief.* For same-institution close-cognitive pairs, the deceased and survivor competed for laboratory space, students, and internal grants (Azoulay et al., 2010; Borjas and Doran, 2012). When j dies, this channel contributes a positive shock that offsets the combined loss from the spillover channel and the smaller residual complementarity channel at close cognitive distance. We posit $R(d_{\text{close}}) > 0$ and $R(d_{\text{middle}}) \approx R(d_{\text{distant}}) \approx 0$ in the same-institution subsample, with $w_y^R R(d_{\text{close}})$ large enough to offset most of $w_y^S S(d_{\text{close}}) + w_y^C C(d_{\text{close}})$, so that $\tilde{L}_y(d_{\text{close}}) \approx 0$. This is the

pattern Figure 5 displays: all four same-institution close-band estimates shift toward zero relative to the different-institution subsample, and none is statistically distinguishable from zero. The strict inequality $w_y^R R(d_{\text{close}}) > w_y^S S(d_{\text{close}}) + w_y^C C(d_{\text{close}})$, which would deliver an outright sign reversal, is needed only for the most stringent quality measure, the one outcome whose point estimate turns positive. By contrast, same-institution distant-cognitive pairs are rarer matches: the search frictions that make distant complementarity hard to assemble also make replacement within the same campus difficult (Fernandez et al., 2016; Singh and Marx, 2013; Sosa and Maoret, 2023; van der Wouden and Youn, 2023), so w_y^C is elevated (a weight shift under the first mapping) and $L_y(d_{\text{distant}})$ is amplified rather than attenuated. Because w_Q^C is small under (A1), the mechanism as stated would put that amplification largest in quantity. Figure 5 instead shows amplification of similar size on all four outcomes, so we read the mechanism as speaking to the sign of the two tails rather than to their relative size across outcome families. The two same-institution effects therefore point in opposite directions at the two tails.

These mappings are illustrative rather than identifying. The empirical specifications in Equations (4) and (5) estimate average effects within each subgroup but do not separately identify the channel weights from the curvature parameters α, β . The same-institution split of Figure 5 reports the pattern organized by the auxiliary competition-relief mechanism, while the remaining splits of Figure 5 (elite/non-elite, senior/junior, recent/former, regular/infrequent, deceased first/last) trace the channel-weight reading in the preceding bullets.

C.6 Robustness to alternative functional forms

The qualitative predictions of Propositions 1 and 2 do not depend on the specific power-function form $(1-d)^\alpha, d^\beta$. To illustrate, consider the exponential pair

$$\tilde{S}(d) = e^{-\alpha'd}, \quad \tilde{C}(d) = 1 - e^{-\beta'd}, \quad \alpha', \beta' > 0, \quad (13)$$

in which $\tilde{S}$ is strictly decreasing on $(0, 1)$ from 1 to $e^{-\alpha'}$, and $\tilde{C}$ is strictly increasing on $(0, 1)$ from 0 to $1 - e^{-\beta'}$. Substituting into Equation (8) and writing $\check{L}_y$ for the resulting loss yields

$$\begin{aligned} \check{L}_q(d) &= w_q^S e^{-\alpha'd} + w_q^C (1 - e^{-\beta'd}), \\ \check{L}'_q(d) &= -\alpha' w_q^S e^{-\alpha'd} + \beta' w_q^C e^{-\beta'd}. \end{aligned}$$

Boundary values of the first derivative are $\check{L}'_q(0) = -\alpha' w_q^S + \beta' w_q^C$ and $\check{L}'_q(1) = -\alpha' w_q^S e^{-\alpha'} + \beta' w_q^C e^{-\beta'}$. The boundary signs $\check{L}'_q(0) < 0$ and $\check{L}'_q(1) > 0$ both hold whenever the loading ratio satisfies

$$\frac{\beta'}{\alpha'} < \frac{w_q^S}{w_q^C} < \frac{\beta'}{\alpha'} e^{\alpha' - \beta'},$$

a non-empty range whenever $\alpha' > \beta'$. By the intermediate value theorem there is an interior $d^* \in (0, 1)$ with $\check{L}'_q(d^*) = 0$. At any such zero, the first-order condition $\alpha' w_q^S e^{-\alpha'd^*} = \beta' w_q^C e^{-\beta'd^*}$ implies $\check{L}''_q(d^*) = (\alpha' - \beta') \beta' w_q^C e^{-\beta'd^*} > 0$, so $\check{L}'_q$ can cross zero only once; hence d^* is unique. Combined with the boundary signs $\check{L}'_q(0) < 0 < \check{L}'_q(1)$, $\check{L}_q$ is strictly decreasing on $(0, d^*)$ and strictly increasing on $(d^*, 1)$, so d^* is a strict global minimum on $[0, 1]$. For the example $\alpha' = 2$, $\beta' = 1$, $w_q^S = w_q^C = 1$, this gives $d^* = \ln 2 \approx 0.693$ with $\check{L}_q(d^*) = 0.75 < \min\{\check{L}_q(0), \check{L}_q(1)\}$.

Quality monotonicity is preserved under the analogous small-loading-ratio condition. The quality counterpart of $\check{L}_q$ is $\check{L}_Q(d) = w_Q^S e^{-\alpha'd} + w_Q^C (1 - e^{-\beta'd})$, with derivative $\check{L}'_Q(d) =$

$-\alpha' w_Q^S e^{-\alpha'd} + \beta' w_Q^C e^{-\beta'd}$, which is negative whenever $w_Q^C/w_Q^S < (\alpha'/\beta') e^{(\beta'-\alpha')d}$. The right-hand side has a strictly positive infimum on $(0, 1)$, so the inequality is uniformly satisfied when the loading ratio is small enough; in the limit $w_Q^C \rightarrow 0$, $\check{L}_Q$ reduces to $w_Q^S e^{-\alpha'd}$, which is strictly decreasing on $(0, 1)$ for any $\alpha' > 0$.

Hence the qualitative predictions of Propositions 1 and 2 do not depend on the specific functional form of the channel decomposition. Three ingredients matter. The spillover and complementarity channels must be monotonic in opposite directions. Curvature and loading ratios must jointly make the quantity-loss derivative negative at $d = 0$ and positive at $d = 1$ with a single crossing in between, so that a unique interior minimum exists. The channel loadings must differ across outcomes in the way captured by assumptions (A1) and (A2). The exponential family is in fact the more demanding case: its U-shape requires the loading ratio to fall inside the window above, whereas the power-function case delivers the interior minimum for any strictly positive loadings under the strict-curvature condition. The power-function parameterization in Section 2.2 is therefore chosen for analytical convenience, not because the qualitative predictions are sensitive to it.