Knowledge Transmission Within and Across Age Groups of Inventors*

Christopher Esposito

Working Paper

Updated February 9, 2024

Abstract

This paper identifies the relationship between the age of inventors and the transmission of knowledge spillovers. Linking age and death information from 13,305 patent inventors that died prematurely to 56,700 co-inventors, I show that inventors who lose early-career collaborators subsequently produce 8.5% fewer patents and 17% fewer highly-cited patents than do inventors who lose mid-career collaborators. Spillovers peak between ages 35 and 44, and mainly accrue to similarly-aged partners. I explore three potential mechanisms: younger inventors may be more productive, they may work more interdependently with their partners, and they may hold newer and more relevant knowledge.

* UCLA (christopher.esposito@anderson.ucla.edu). I thank Mike Andrews, Eamon Duede, James Evans, Lee Fleming, Ashvin Gandhi, Charu Gupta, Jennifer Kao, Donghyun Kang, Jeffrey Lockhart, Olav Sorenson, Frank van der Wouden, and participants at AOM Boston, the Mansueto Institute Colloquium Series, and ICSSI for helpful comments. All remaining errors are my own.

1. Introduction

This paper is the first to look at how knowledge moves between inventors of different ages. In particular, I measure the volume of knowledge that moves from older inventors to younger inventors, from younger inventors to older inventors, between younger inventors, and between older inventors. Previous studies have looked at how the volume of knowledge flows vary with the age of the knowledge recipient (Azoulay, Zivin, and Wang 2010; Jaravel, Petkova, and Bell 2018; Bernstein et al. 2022), but have not considered at how the volume varies with the age of the knowledge *transmitter,* nor the interaction between the age of the transmitter and the recipient.

Uncovering how knowledge flows vary with transmitter age is important for several economic issues. First, such information could help to reconfigure teams and networks in order to stimulate skill development, including skill redevelopment in older workers at risk of knowledge obsolescence (Jaravel, Petkova, and Bell 2018; Aghion et al. 2023). Second, such information could delineate institutional differences between technological change and scientific advance, which inefficiently rewards prominent (and presumably senior) scientists (Azoulay, Zivin, and Wang 2010; Azoulay, Fons-Rosen, and Zivin 2019). Finally, such information would help to establish the extent to which long-run knowledge growth is a cumulative process. In cumulative models of knowledge growth, ideas do not depreciate over time. An implication is that knowledge is disproportionately passed down across generations of knowledge workers, from the old to the young (Romer 1990; Jones 2009). However, if younger inventors generate large knowledge flows, then knowledge growth would appear to be a disruptive process (Aghion and Howitt 1992), with implications for R&D productivity growth over time.

To identify the volume of knowledge spillovers between age groups of inventors, I study how an inventor's own patent production changes after experiencing the untimely loss of a collaborator, using a difference-in-difference design that contrasts the effect of losing collaborators of different ages. I focus on knowledge transmitted through collaborative networks because these networks provide tangible records of intensive interaction. My analysis combines USPTO data on patent output and inventor collaboration together with records on the birth and death year of patent inventors, with information for 13,505 inventors that died prematurely (before age 60) and their 56,700 co-inventors.

I begin the analysis by showing that young collaborators generate more knowledge spillovers than do older ones. Focal inventors who lose early-career collaborators (age 20-44 at time of death) proceed to produce 8% fewer patents relative to inventors who lose mid-career collaborators (age 45-59 at time of death). This difference becomes statistically significant the year after the death of the collaborator and endures for twelve years. The effect size of the differential is larger for inventor-collaborator dyads with stronger collaborative relationships, and the effect size is larger when the survivor's production of highly-cited patents is the outcome variable.

After these initial analyses, I disaggregate the treatment effect by the age group of the deceased collaborator. The relationship between collaborator age and generated spillovers is U-shaped, with the death of very young collaborators (aged 20-29 at time of death) and relatively old collaborators (aged 45-59 at time of death) inducing a smaller patenting decline than the death of moderately young collaborators (aged 30-44 at time of death). This U-shaped relationship suggests that very young collaborators may lack sufficient experience and knowledge to generate a high volume of spillovers (Jones 2010), while older collaborators may lack knowledge of recently-developed ideas that are useful for generating new inventions in their contemporary technological and economic environments (Aghion et al. 2023; Esposito and Wouden 2022). Therefore, knowledge growth appears to occur through processes of both accumulation and depreciation.

I go on to decompose the treatment effects by the age of both the deceased collaborator *and the focal inventor* at the time of the collaborators' death. This decomposition shows that the greatest quantity of spillovers flow from collaborators in their 30s to inventors in their 30s. In addition, more spillovers flow from younger inventors to older inventors (i.e. from collaborators in their 30s to co-inventors in their 40s) than the other way around. These results demonstrate that junior inventors are at the forefront of driving the social learning process.

Finally, I explore three potential mechanisms for the main finding that junior collaborators generate more spillovers than mid-career collaborators. The first potential mechanism is differences in productivity. Inventors' patenting output peaks at a relatively young age, which could create patenting spillovers that are not related to knowledge transmission (Kaltenberg, Jaffe, and Lachman 2023). The second is collaborator interdependence. Inventors may have more interdependent relationships with their junior collaborators, either because of greater knowledge complementarities, or because of a larger incentive to invest in relationship-building with young collaborators due to the potentially longer time horizon of the relationship (Jaravel, Petkova, and Bell 2018). The third is human capital relevance. Younger inventors have recently-vintaged human capital, which may be more useful for producing patentable technologies in the current technological environment. To explore these mechanisms, I decompose the main effect by collaborators' pre-death patenting productivity, the intensity of the collaborative pair's co-patenting activity, and two variables that capture the recency and relevance of the deceased collaborator's human capital: their human capital vintage (measured by the age of the citations on their patents), and the "fertility" of their human capital for creating subsequent inventions (measured by the number of additional patents that cite the same patents that their own patents cite).

My decomposition exercise generates no evidence that the greater spillovers produced by junior collaborators is driven by sheer differences in productivity or by more interdependent relationships. On the other contrary, the recent vintage and high fertility of junior collaborators' knowledge is more powerful in explaining why junior collaborators generate more spillovers than do senior collaborators.

An implication of these findings is that younger collaborators should be particularly important sources of knowledge spillovers in fast-advancing knowledge fields, where the knowledge frontier is quickly expanding and the set of ideas that are most useful for creating new technologies changes quickly. I test this proposition in Appendix A4, where I measure the rate advance of each knowledge field by the average age of the citations made in those fields. I find that differential in the spillovers generated by junior collaborators is larger in fast-advancing knowledge fields. Finally, if young collaborators transmit new-to-the-world ideas to their partners, then inventors who lose younger collaborators should also lose access to these new-to-the-world ideas. I also test this prediction in Appendix A4, where I show that inventors who lose junior collaborators proceed to cite older patents than they would have otherwise.

Analytically, this study most closely resembles the work of Jaravel, Petkova, and Bell (2018), Balsmeier, Fleming, and Lück (2023), Oettl (2012), and Azoulay, Graff-Zivin, and Wang (2010), which use the premature deaths of inventors and scientists to identify knowledge spillovers generated by the deceased. A distinguishing feature of this paper is its focus on heterogeneous treatment effects, in particular the differential spillovers of collaborators based on their biological age. Therefore, the difference-in-difference term of interest is the contrast between the impact of losing a junior collaborator from that of losing an older collaborator. This contrast raises a question of whether inventors who lose mid-career collaborators are appropriate counterfactuals for inventors who lose early-career collaborators. I take three steps to address this. First, I show pre-trend analyses demonstrating that inventors who lose early-career collaborators to premature deaths have parallel pre-trends to inventors who lose mid-career collaborators (Figure 1). Second, I show that collaborators' pre-death patenting productivity does not explain differences in the knowledge spillovers generated by junior collaborators (Table 2). Third, I demonstrate that collaborators who die early in their careers have similar career patenting profiles as those who die in their mid-careers, except at very young ages. I further show that dropping collaborators that died at a very young age does not materially change the results (Figure A5).

Finally, because the treatment events are staggered, two-way fixed effects estimates can be biased by phased-in treatment effects (Baker, Larcker, and Wang 2022). The difference-in-difference framework I use, which conducts a within-event-time contrast between inventors who lose junior and mid-career collaborators, does not resolve this issue because biased estimates from the phase-in treatment effects can load onto the first-difference estimator and the individual and time fixed effects, similar to the bias introduced by staggered treatments in triple-difference designs (Strezhnev 2023). Therefore, in Appendix B, I develop a stacked regression that omits prohibited contrasts. I find that the results of the stacked regression are not meaningfully different from the two-way fixed effect estimates presented in the main text.

In the following sections, I introduce the data sources, describe the methods, present the results, and discuss the implications of the findings.

2. Data and Methods

I collect patent data from two sources. The first source is PatentsView, from which I collect the set of all utility patents granted by the U.S. Patent and Trademark Office between 1976 and

2020. These data patent ID numbers, current CPC classification information codes (at the class level, for which there are 123 classes), front-page patent citations records, disambiguated inventor IDs, and patent application year information. Because PatentsView lacks inventor, application, and citation information for patents granted before 1976, I restrict the study to patents applied for in 1976 or later. In addition, because I use forward citation counts to identify high-impact inventions, and because citations accumulate over time, I restrict the sample to patents which were applied for no later than 2013. This cutoff allows for 7 years for patent applications to be granted and receive citations. I use 5-year windows to count the forward citations received by patents, and I define "high impact" patents as those in the top quartile of their grant year and CPC technology class in terms of their number of citations received.

My second data source are records on the year of birth and year of death records for 1.9 million inventors recently made available by Kaltenberg, Jaffe, and Lachman (2023). The records were compiled by scraping three websites that aggregate birth and death records for the general population, and by matching them to USPTO patent inventors by name and residential location. The data collection procedures and descriptive statistics are shared in Kaltenberg, Jaffe, and Lachman (2023) and in a 2021 working paper (Kaltenberg, Jaffe, and Lachman 2021). In Appendix A1 of this paper, I describe steps I take to clean the data and to match inventors to their deceased collaborators. The Kaltenberg, Jaffe, and Lachman (2023) data have been used by Balsmeier, Fleming, and Lück (2023) to analyze how premature deaths affect the geographic spillovers of patent citations, but to my knowledge have not been used to analyze how the externalities generated by patent collaborators vary by the age of the deceased collaborator.

Finally, to compute inventors' patenting productivity, I count the number of total patents and high-impact patents invented by each inventor in each year. Inventors' patenting careers span from their first recorded patent to their final patent. I include interim years with zero patents in the dataset. I do not down-weight inventors' patents if they were co-invented by two or more inventors. The average patent was co-invented by 1.8 co-inventors, suggesting that inventors tend to contribute substantially to each patent on which they are listed as inventors.

Following the data construction, my dataset contains 56,700 focal inventors and 13,305 collaborators that died prematurely. I provide summary statistics in Appendix A1. The summary statistics show considerable variation in the age of surviving and deceased inventors, with the 25th and 75th percentiles of surviving inventor age at time of collaborator death ranging from 39 to 55 years and the associated figures for the age of deceased collaborators ranging from 43 to 55 years. Because senior deceased inventors (those that die at age 60 or above) are omitted from the data, the distribution of age at collaborator death is capped at a value of 59. In terms of career length and patent production, the 25th and 75th percentiles of surviving inventors had careers that lasted 7 to 22 years, during which they produced 4 to 22 patents.

3. Empirical Analysis

3.1 Main Effects

To estimate the effect of the age at time of death of the collaborator on a surviving co-inventor's patenting productivity, I estimate the difference-in-difference model described by Equation 1:

$$(1)\quad PatProd_{i,t} = \exp\left[\beta_0 PostCollabDeath_{i,t} + \beta_1 PostCollabDeath_{i,t} * JuniorDeceasedCollab_i + \alpha_i + \tau_{modalfield*t} + \varepsilon_{i,t}\right]$$

In Equation 1, $PatProd_{i,t}$ is the number of patents produced by focal inventor i in the application year t and $PostCollabDeath_{i,t}$ is a binary variable that equals 0 for the 10 years before the year of the collaborator's death and equals 1 for the 10 years following the collaborator's death. $JuniorDeceasedCollab_i$ is a binary variable that equals 1 if the deceased collaborator was between ages 20 and 44 at the time of death and equals 0 if the deceased collaborator was between ages 45 and 59 at time of death. As discussed in the data section, I omit all collaborator deaths where the collaborator died outside ages 20-59 because such deaths may have been easier to anticipate and thus endogenous to the dependent variable. α_i are inventor fixed effects. Because the $JuniorDeceasedCollab_i$ term is constant within inventors, its base term is subsumed into the α_i fixed effects and so it only appears as an interaction in the model. τ_{f*t} are fixed effects for a surviving inventors' modal CPC technological class, defined as the most frequent technology class assigned to each inventor's patents, interacted with year indicators. There are 123 unique CPC classes at this level of aggregation. Because the τ_{f*t} fixed effects contain unique intercepts for each year, inventor age effects (which are collinear with the inventor and year-specific fixed effect terms) are projected out of the variation in the dependent variable. Therefore, the remaining variation that may load onto the β_0 and β_1 coefficient terms is deflated for the mean values of surviving inventors, class*year pairs, and the linear trends within inventors (Hall, Mairesse, and Turner 2005). Because the dependent variable is a count variable, I estimate Equation 1 using a Poisson Quasi-Maximum Likelihood estimator.

I additionally test whether the loss of a collaborator and a junior collaborator causes a decline in a surviving inventor's rate of producing high-impact patents. High impact patents are defined as those in the top quartile of their grant year and CPC class. Finally, the effect of losing a collaborator is likely to be stronger for inventors that lose repeat collaborators with whom they have developed stronger relationships (Jaravel, Petkova, and Bell 2018). Therefore, I also run the model after restricting the dataset to inventor-collaborator pairs that co-invented two or more patents in the five years leading up to the collaborator's death.

Regression results for Equation 1 are shown in Table 1. The first column shows that $\beta_0 = -0.448$, indicating that inventors produce $e^{-0.448} - 1 = 36\%$ fewer patents per year following the death of a collaborator. The coefficient on the interaction term β_1 is -0.0898, indicating that inventors that lose junior collaborators subsequently produce 8.6% fewer patents per year than surviving inventors that lose mid-career collaborators. The qualitative significance of this effect size can be interpreted by considering that the median surviving inventor in the dataset produces 9 patents over the course of a 13-year career. The second column of Table 1 shows that the effects on high-impact patenting are larger compared to those on overall patenting. Inventors who lose a collaborator prematurely proceed to produce 33% fewer high-impact patents per year, and those who lose junior collaborators proceed to produce an additional 13% fewer high-impact patents. Thus, early-career collaborators even more important for high-impact patenting than they are for overall patenting.

As expected, the effect sizes are larger in the models that only consider repeat collaborators. An inventor who loses an early-career repeat collaborator proceeds to produce 16% fewer patents per year than do inventors who lose repeat mid-career collaborators (column 3). The effect size on high-impact patenting is also larger for the loss of a repeat early-career collaborator than for the full set of early-career collaborators (column 4).

Table 1: Fixed Effect Quasi-Poisson Estimates of Loss of Junior Collaborator on Patenting

	All Collaborators		Repeat Collaborators	
	Total Patenting	High-Impact Patenting	Total Patenting	High-Impact Patenting
$PostCollabDeath_{i,t}$	-0.448*** (0.0248)	-0.400*** (0.0378)	-0.579*** (0.0422)	-0.513*** (0.0629)
$PostCollabDeath_{i,t}$ $* JuniorDeceasedCollab_i$	-0.0898** (0.0399)	-0.141*** (0.0522)	-0.174*** (0.0613)	-0.198** (0.0808)
Inventor and Class*Year Fixed Effects?	Y	Y	Y	Y
Inventor*Year Obs	296,901	209,204	96,604	74,041
Inventor Obs	40,129	22,900	11,599	8,038

Notes: The table presents regression estimates for Equation 1. Junior deceased collaborators are defined as those that die between ages 20 and 45. The reference set of deceased collaborators are those that die between ages 46 and 59. Repeat collaborators are those which surviving inventors co-invented 2+ patents.

Next, I explore the dynamics of the treatment effect by estimating the regression model described by estimating the model with a set of 15 forward and 15 backward lags to the treatment event. Because these lag variables are perfectly colinear with the inventor and year fixed effects, the base terms of these lags (without the interaction with $JuniorDeceasedCollab_i$) cannot be estimated (Hall, Mairesse, and Turner 2005). The regression is described by Equation 2:

$$(2) \quad PatProd_{it} = \exp\left[\sum_{\tau=-15}^{15} \beta_\tau YearsToTreatment_{it}^{\tau} * JuniorDeceasedCollab_i + \alpha_i + \tau_{modalfield*t} + \varepsilon_{it}\right]$$

I plot the resulting coefficients and their 95% confidence intervals in Figure 1. There are no obvious pre-trends in the data. The negative relationship between losing a junior begins one year after a collaborator death, becomes statistically significant after 7 years, and continues for 13 years following the collaborator death. In both charts, the effect size is at its maximum 10-12 years following a collaborator's death. The effect dissipates 13-15 years after the collaborator death. The effect sizes are about twice as large for high-impact patenting and for inventors who lose repeat collaborators. In Appendix A2, I show that the results are similar when a stacked regression is used. Therefore, the results in Figure 1 are not driven by potential contamination from heterogeneous treatment effects in the staggered design.

The time dynamics of the treatment effects in Figure 1 provide some suggestions as to mechanisms. In particular, the full impact is not felt until 10 years after the loss of a junior collaborator, suggesting that the loss of junior collaborators affects focal inventors' productivity not by disrupting existing projects, but by slowing down the initiation of new projects or by reducing networking and learning opportunities. In this regard, it is important to keep in mind that the coefficients show the differential effect of losing a junior collaborator relative to a mid-career collaborator. Losing both types of collaborators may impact surviving inventors equally in the short term (1-5 years) by interrupting ongoing projects (Jaravel, Petkova, and Bell 2018), while in the longer run (6-12 years), the impact of losing a junior collaborator is greater, possibly because the junior collaborator provided access to networking or learning opportunities that would have taken several years to mature.

Figure 1: Change in Patenting Productivity Following Death of Junior Collaborator Relative to Mid-Career Collaborator

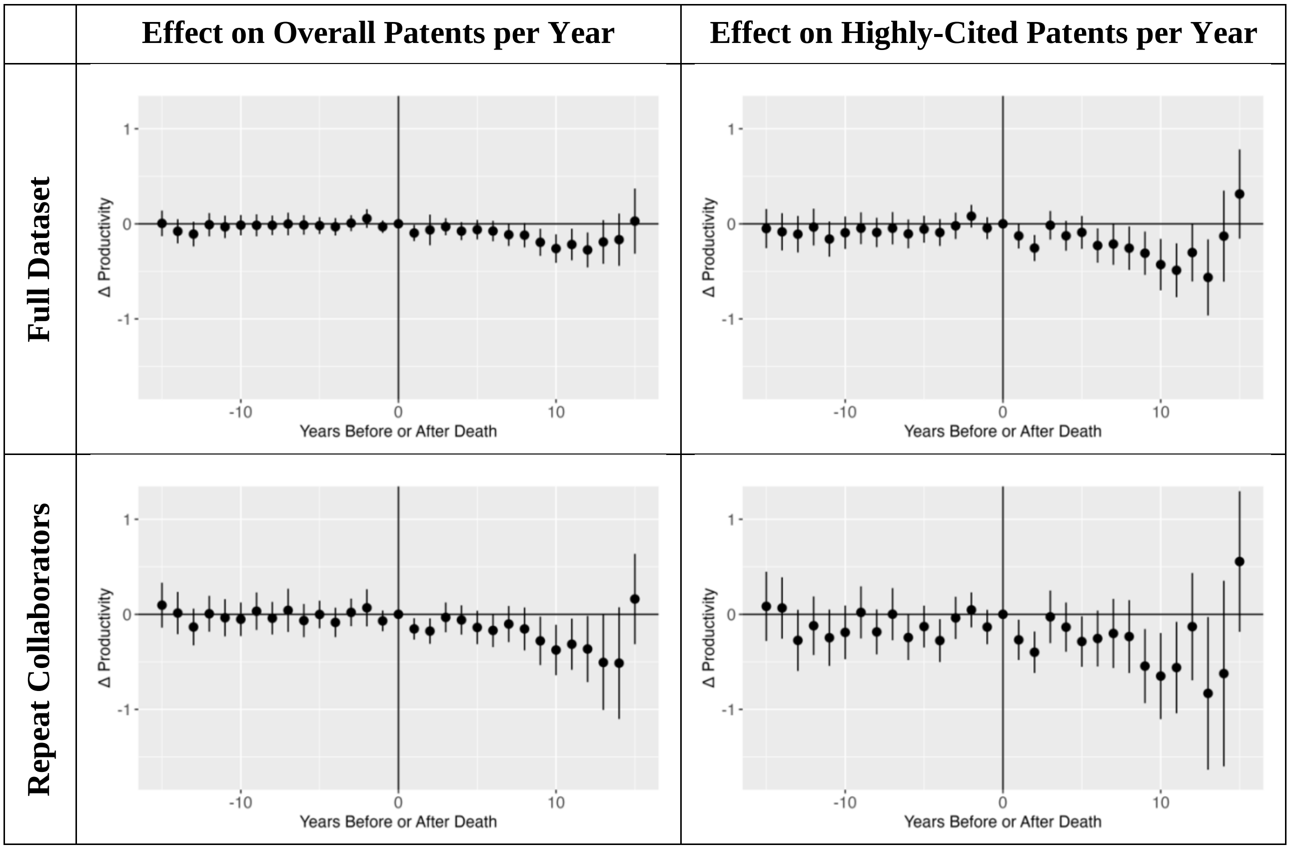

*Notes: Figures shows coefficients associated with losing an early-career collaborator (age 20-44 at time of death) relative to losing a mid-career collaborator (aged 45-59 at death). All models include focal inventor and year*technology class fixed effects. Standard errors are clustered at the deceased collaborator.*

3.2 Effects by Collaborator Age

To identify the specific age at time of death of a collaborator that has the largest effect on the subsequent patenting of a focal inventor, inventor, I estimate a regression model described by Equation 3:

(3) $$PatProd_{i,t} = \exp\left[\beta_0 PostCollabDeath_{i,t} + \sum_{m=20}^{59} \beta_m PostCollabDeath_{it} * AgeCollabAtDeath_i^m + \alpha_i + \tau_{modalfield*t} + \varepsilon_{i,t}\right]$$

In Equation 3, patenting productivity is a function of $PostCollabDeath_{it}$, which records a value of 0 for the 10 years leading up to the death of the collaborator and 1 for the 10 years after, and the interaction term $PostCollabDeath_{it} * AgeCollabAtDeath_i^m$. The second variable of the interaction term, $AgeCollabAtDeath_i^c$, records the age of the collaborator at time of death. Because of the relatively small number of observations of collaborator deaths at each specific age, I group $AgeCollabAtDeath_i^m$ into 5-year bins: 20-24 years old, 25-29 years old, and so on, up to the reference group of collaborators that die between ages 55 and 59. I plot the β_n coefficients and their 95% confidence intervals in Figure 2.

Figure 2 shows that collaborator age at time of death has a U-shaped relationship with the subsequent patenting productivity of the focal inventor. Very young collaborators (aged 20-29 at death) generate no more spillovers do older collaborators (aged 45-59 at death). The greatest spillovers are generated by collaborators between the ages of 30 and 44. This U-shaped relationship is more pronounced for high-impact patenting than for patenting overall, and for inventors who lose repeat collaborators. The U-shaped relationship suggests that both experience (Jones 2010) and knowledge obsolescence (Aghion et al. 2023) may play important roles in determining the quantity of spillovers that collaborators generate.

Figure 2: Change in Patenting Productivity by Collaborator Age at Death

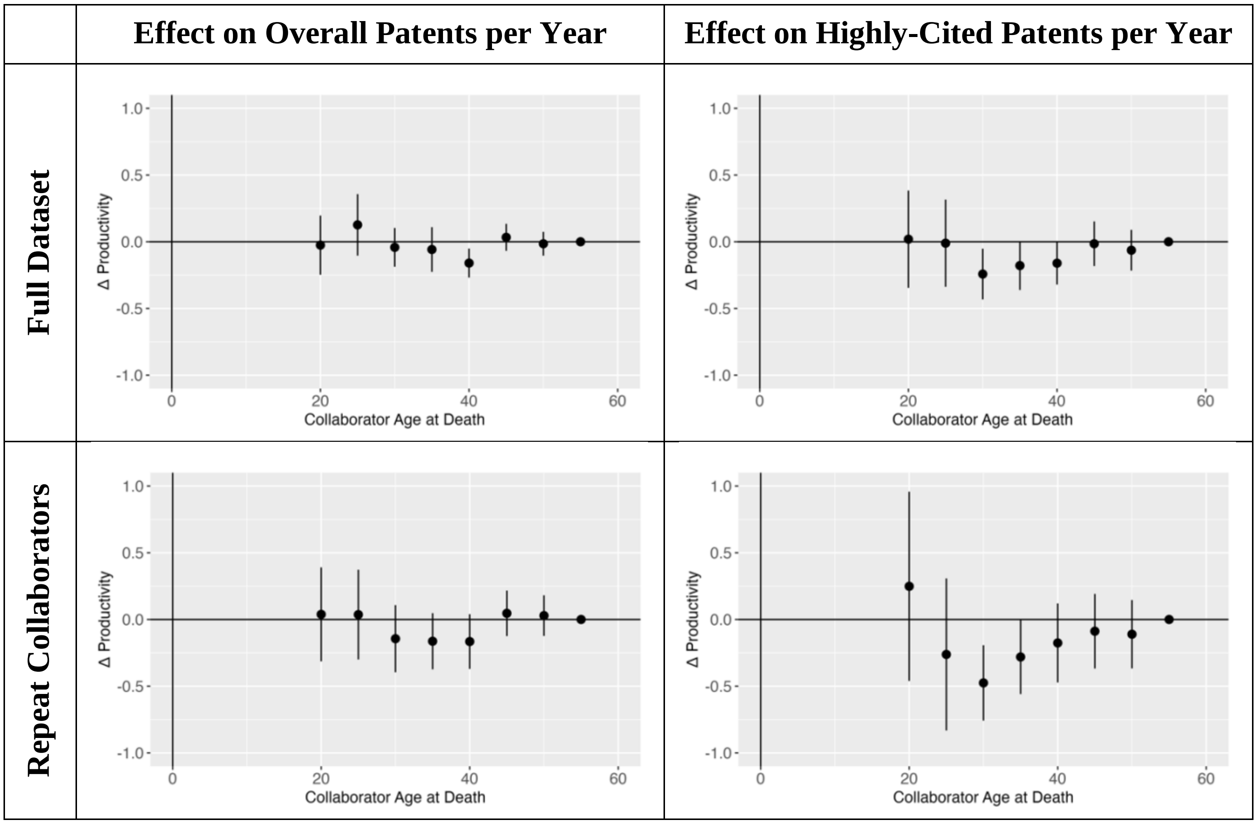

Notes: Figures show the differential effect of losing a collaborator at a specific age at of death relative time (grouped into 5 year bins) to losing a collaborator aged 55-59 at time of death on patenting output.

3.3 Effects by Collaborator Age

I now analyze age at which *focal inventors* are most affected by the loss of collaborators of different ages. I do so by estimating Equation 4 using a Quasi-Poisson:

(4) $$PatProd_{it} = \exp\left[\beta_0 PostCollabDeath_{it} + \sum_{m=20}^{59} \sum_{n=20}^{59} \beta_{mn} PostCollabDeath_{it} * AgeCollabAtDeath_i^m * AgeInventorAtDeath_i^n + \alpha_i + \tau_{modalfield*t} + \varepsilon_{it}\right]$$

The coefficient matrix β_{mn} records the change in patenting productivity for inventors at age *n* at the time of death of their collaborators at age *m*. Because the pairwise disaggregation leads to relatively few observations in any pairwise cell, I aggregate inventors to 10-year age based on their age at time of collaborators' deaths. The reference group in the regression are inventors aged 50-59 at the time of death the death of their collaborator, when the collaborator also dies in the 50-59 age group. I estimate Equation 3 for overall patenting, high-impact patenting, and for inventors who lose repeat collaborators and plot heat maps of the β_{mn} coefficient matrix in Figure 3. Deeper red indicates larger negative coefficients, and thus greater spillovers. Asterisks are placed in the cells of the heatmaps for values with coefficients that are statistically different from 0.

If patenting spillovers primarily flowed from older collaborators to their younger co-inventors, then the bottom-right triangles of each of the heatmaps in Figure 3 would be shaded in deep red. However, this is not the case. Instead, the most consistently negative and significant coefficient across the four panels of the figure is associated with inventors in their 30s who lose collaborators in their 30s. In addition, the top-left triangle of each heatmap is more deeply shaded in red than the bottom-right triangle, and in the fourth panel of the figure, the coefficient associated with inventors in their 40s who lose collaborators in their 30s is statistically different from 0. This asymmetry shows that more patenting spillovers flow from younger collaborators to older inventors than from the old to the young.

Figure 3: Change in Highly-Cited Patent Productivity Following Death of a Collaborator

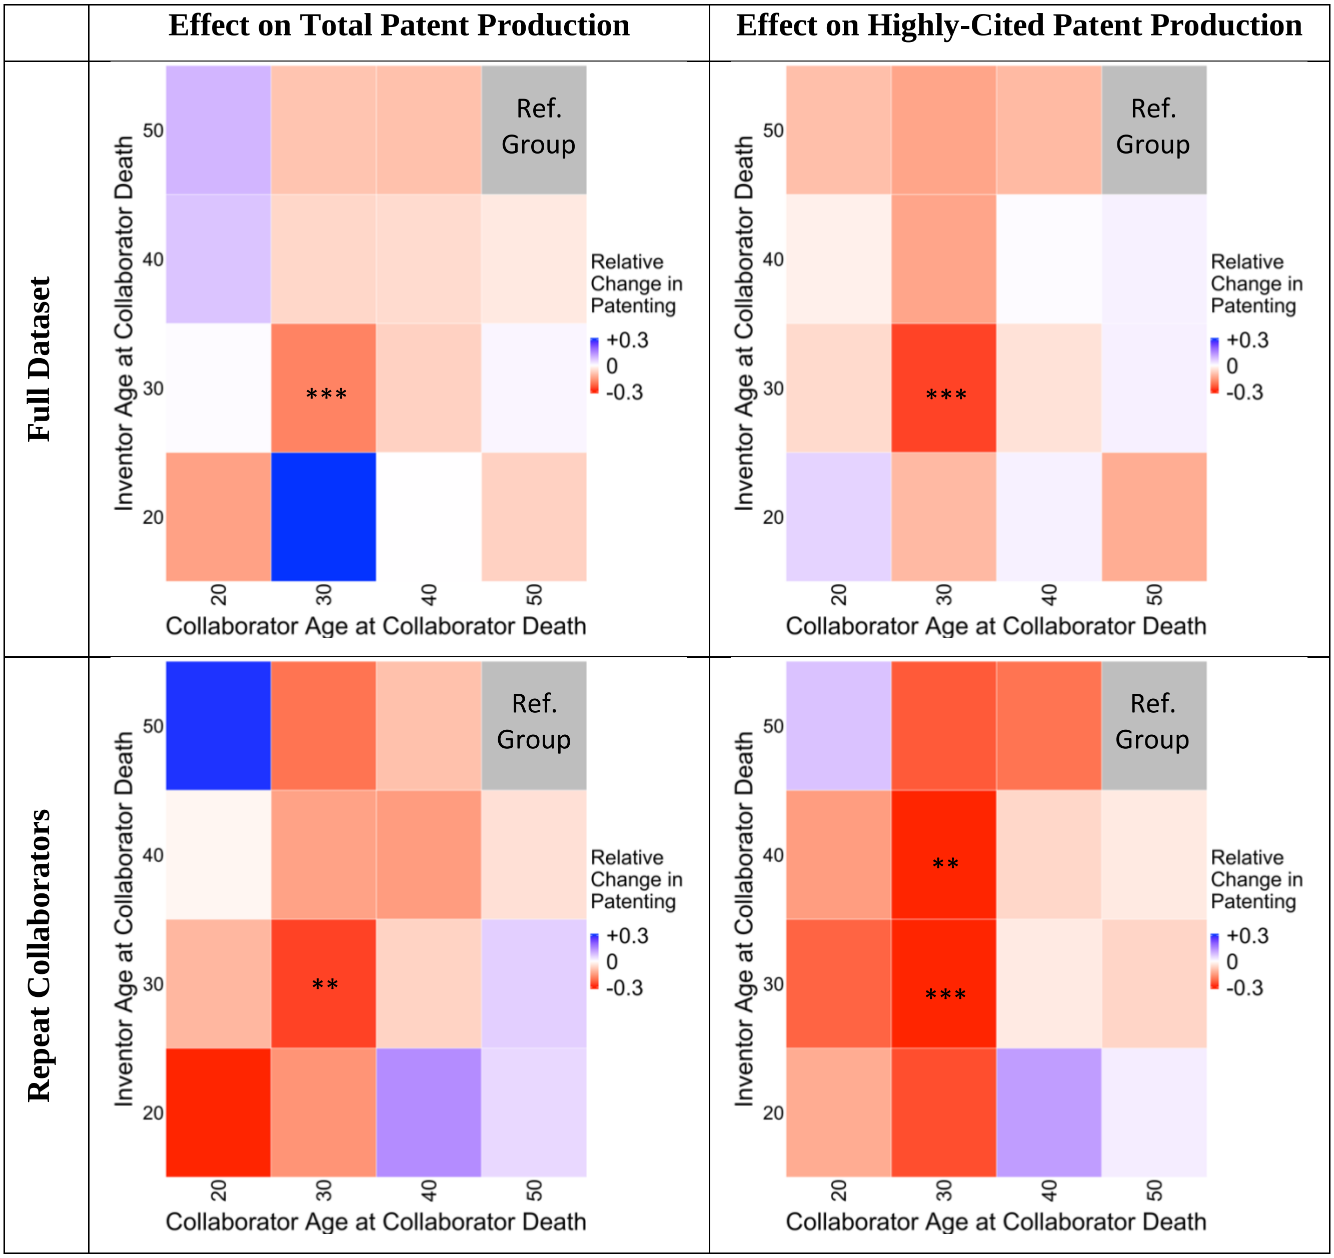

*Notes: Heatmaps show the change in patent productivity of surviving inventors following a collaborator's premature death, broken outs by inventor age at time of collaborator death and collaborator age at time of death. Dyads with surviving inventors aged 50-59 at the time of death and collaborators aged 50-59 are the reference group. Asterisks indicate statistical significance (*** $p < 0.01$; ** $p < 0.05$; * $p < 0.01$).*

3.4 Mechanisms

The literature suggests three potential explanations for why early-career inventors generate more patenting spillovers than do older inventors. The first is that early-career inventors may be more productive, and thus the spillovers detected by the earlier models may not reflect knowledge transmission. The second potential mechanism is team interdependence. The intensity of collaboration is heterogenous across dyads of collaborators. This heterogeneity is plausibly correlated with the age of collaborators, because the incentive to invest in building strong collaborative relationships with junior inventors could be larger due to their longer expected careers (Jaravel, Petkova, and Bell 2018). The third potential mechanism is knowledge obsolescence. Younger inventors have more recently-vintaged human capital, which may be better-oriented for innovating in the current knowledge environment (Aghion et al. 2023; Esposito and Wouden 2022).

To test these mechanisms, I decompose the main effect of the loss of a junior collaborator on subsequent patent output using five key variables. The first variable is collaborator's pre-collaborator-death patenting productivity, which captures the intensity of the labor that inventors contribute to their teams. The second variable is the number of co-patents produced by the focal inventor and the deceased collaborator before the collaborator's death. To ensure that this variable isolates a dyad's collaboration intensity from the surviving inventor's patent production, I include a third control variable, the number of patents produced the focal inventor before the collaborator's death, in the regression.[2] The fourth variable is the collaborators' pre-death human capital vintage, which indicates whether collaborators have knowledge of new-to-the-world ideas. I measure

[2] By including the focal inventor's pre-collaborator-death patent output in the regression, the number of co-patents produced by the focal inventor and the deceased collaborator is interpreted as the intensity of the collaborative relationship. A co-patent share variable is unnecessary because the numerator and denominator of a such a variable are already included as separate terms.

collaborator's human capital vintage as the mean age of the citations on their patents in the three years leading up to their death. Inventors whose patents cite more recent patents are determined to have more recently-vintaged human capital. The fifth term is collaborators' pre-death human capital "fertility", which indicates whether a collaborator has knowledge of ideas that can be used to make many subsequent inventions. I measure a collaborator's knowledge fertility as the mean number of additional patents that cite the same patents cited by a collaborator in the three years before the collaborator's death.[3] Inventors that cite patents that can be leveraged to produce many more patents possess highly fertile knowledge. Together, these five variables test whether the greater spillovers generated by junior collaborators are driven by potential differences in productivity, collaborative relationship intensity, or human capital recency and relevance.

The difference-in-difference model with six second-difference terms is given by Equation 5. The model contains the base term, $PostCollabDeath_{it}$, and interactions between this variable and the decomposition terms $JuniorDeceasedCollab_i$, $PreDeathCollabPatents_i$, $PreDeathCoinventedPatents_i$, $PreDeathPatents_i$, $PreDeathCollabCiteAge_i$, and $PreDeathCollabCiteFertility_i$. Because collaborators' pre-death mean citation age is strongly right-skewed, I take the variable's natural log. I also multiply the variable by -1, so higher values indicate more recently-vintaged human capital. Finally, to isolate knowledge fertility from calendar year effects, I standardize collaborators' knowledge fertility by the calendar year of their patent applications. This standardization controls for a possible mechanical relationship where inventors that cite older patents could appear to have more fertile knowledge, because the older

[3] For example, assume an inventor dies in 2000 produced one patent (Patent A) in the three years before death. Patent A cites two earlier patents, Patent α and Patent β. By 2000, Patent α was cited by 10 other patents, and Patent β was cited by 20 other patents. In this case, the inventors' knowledge fertility is 15, because on average 15 other patents drew knowledge from the patents that the inventor drew knowledge from.

patents that they cite have had more time to accumulate citations. I estimate the model with individual and field*year fixed effects using a Quasi-Poisson.

$$(5) \qquad PatProd_{i,t} = \exp\big[\beta_0 PostCollabDeath_{it} + \beta_1 PostCollabDeath_{it} * JuniorDeceasedCollab_i + \beta_2 PostCollabDeath_{it} * PreDeathCollabPatents_i + \beta_3 PostCollabDeath_{it} * PreDeathCoinventedPatents_i + \beta_4 PostCollabDeath_{it} * PreDeathPatents_i + \beta_5 (PostCollabDeath_{it} * \log(CollabPreDeathCiteAge)_i * -1) + \beta_6 PostCollabDeath_{it} * PreDeathCollabCiteFertility_i + \alpha_i + \tau_{ft} + \varepsilon_{it}\big]$$

In Table 2, I show Estimates for surviving inventors' overall patent output in Panel A, and estimates for surviving inventors' high-impact patent output in Panel B. The first column shows the relationship between collaborator biological age and subsequent productivity of the survivor, the second column tests the collaborator productivity hypothesis, the third column tests the relationship intensity hypothesis, the fourth column tests the collaborator human capital recency hypothesis, and the fifth column includes all 6 interaction terms in a "horse race" model.

Table 2: Fixed Effect Quasi-Poisson Estimates of Mechanisms for Patenting Decline

Panel A: Effects on Overall Patent Production of Surviving Inventor	Mechanism Tested by Model				
	Biological Age	Collaborator Productivity	Collaborative Relationship Intensity	Collaborator Knowledge Relevance	Horse Race
$PostCollabDeath_{i,t}$	-0.453***	-0.463***	-0.413***	-0.716***	-0.621***
	(0.0255)	(0.0269)	(0.0341)	(0.114)	(0.109)
$PostCollabDeath_{i,t}$	-0.0878**	-0.0864**	-0.0774**	-0.0630	-0.0468
$* JuniorDeceasedCollab_i$	(0.0398)	(0.0395)	(0.0387)	(0.0455)	(0.0436)
$PostCollabDeath_{i,t}$		0.00133			0.00177
$* CollabPreDeathPatents_i$		(0.00132)			(0.00151)
$PostCollabDeath_{i,t}$			0.00591		0.00383
$* \text{PreCollabDeathCoinventedPatents}_i$			(0.00738)		(0.00806)
$PostCollabDeath_{i,t}$			-0.00294***		-0.00312***
$* \text{PreCollabDeathPatents}_i$			(0.000862)		(0.000933)
$PostCollabDeath_{i,t}$				-0.0912**	-0.0659*
$* LogCollabPreDeathCiteAge_i * -1$				(0.0421)	(0.0394)
$PostCollabDeath_{i,t}$				-0.0570***	-0.0637***
$* CollabPreDeathKnowledgeFertility_i$				(0.0201)	(0.0201)
Inventor + Class*Year Fixed Effects?	Y	Y	Y	Y	Y
Inventor*Year Observations	292,856	292,856	292,856	220,312	220,312
Inventor Observations	33,051	33,051	33,051	24,367	24,367

Table 2 (continued)

Panel B: Effects on High-Impact Patent Production of Surviving Inventor	Mechanism Tested by Model				
	Biological Age	Collaborator Productivity	Collaborative Relationship Intensity	Collaborator Knowledge Relevance	Horse Race
$PostCollabDeath_{i,t}$	-0.411***	-0.406***	-0.349***	-0.706***	-0.627***
	(0.0376)	(0.0377)	(0.052)	(0.141)	(0.144)
$PostCollabDeath_{i,t}$	-0.136***	-0.136***	-0.131**	-0.0854	-0.0741
$* JuniorDeceasedCollab_i$	(0.0503)	(0.0503)	(0.0509)	(0.0560)	(0.0562)
$PostCollabDeath_{i,t}$		0.000626			0.00156
$* CollabPreDeathPatents_i$		(0.000666)			(0.00132)
$PostCollabDeath_{i,t}$			-0.00230		-0.000728
$* \mathrm{PreDeathCoinventedPatents}_i$			(0.00747)		(0.00700)
$PostCollabDeath_{i,t}$			-0.00276**		-0.00308**
$* \mathrm{PreDeathPatents}_i$			(0.00128)		(0.000134)
$PostCollabDeath_{i,t}$				-0.106**	-0.0925*
$* LogCollabPreDeathCiteAge_i * -1$				(0.0533)	(0.0537)
$PostCollabDeath_{i,t}$				-0.136***	-0.133***
$* CollabPreDeathKnowledgeFertility_i$				(0.0314)	(0.0327)
Inventor + Class*Year Fixed Effects?	Y	Y	Y	Y	Y
Inventor*Year Observations	212,070	212,070	212,070	161,411	161,411
Inventor Observations	21,446	21,446	21,446	16,031	16,031

Notes: The table presents regression estimates for Eq. 5. Panel A shows effects for overall patent output of survivors; Panel B shows effects on high-impact patents (top-25% cited within-year). Junior deceased collaborators are those that died age 20 to 45. The reference set are those that died age 46 to 59. Collaborator pre-death productivity is the count of patents produced in the three years leading up to their death, pre-death knowledge age is the mean age of the citations made in their patents in the three years pre-death, and pre-death knowledge fertility is the mean number of patents that cite the same patents cited by their pre-death patents. Mean knowledge fertility is standardized by the application year of the cited patents to remove temporal effects. Focal inventor pre-death patents is the number of patents produced by the focal inventor in the 5 years pre-death, and pre-death coinvented patents is the number of patents co-invented by the survivor and deceased collaborator in the 5 years pre-death.

Table 2 shows that the recency and relevance of junior collaborator's knowledge is the strongest predictor of the treatment effect. The coefficient on the biological of the collaborator shrinks and loses significance when the collaborator's citation age and fertility are included in the model. In addition, both of these terms are negative and significantly associated with the treatment outcomes of overall patenting (Panel A) and high-impact patenting (Panel B).

Table 2 does not generate evidence in support of the collaborator productivity nor the relationship intensity mechanisms. The variable measuring a collaborator's pre-death patenting productivity is statistically insignificant and the coefficient size is small (exponentiating the coefficient indicates that survivors produce 0.1% fewer patents for each additional pre-death patent produced by their deceased collaborators). The coefficient on the dyad's number of co-invented patents is also small and significant, indicating that survivors who lose collaborators with whom they had more intensive pre-death collaborative relationships do not suffer larger patenting declines, once the biological age of the collaborator is controlled for. The inclusion of the control variable for the number of patents produced by the survivor in the years leading up to the collaborator's death implies that the co-patents variable should be interpreted as the effect of the relationship's intensity; a co-patent share variable would be redundant.

Finally, the horse race model, which includes all 6 interaction terms, affirms the predictive strength of the collaborator's citation age and knowledge fertility. In the horse race model, both terms are significant and negative (albeit at the 90% level for citation age), the terms associated with the collaborator productivity and collaborative intensity mechanisms are insignificant, and the coefficient on biological age is insignificant.

An implication of the findings from Table 2 is that the differential in spillovers generated by junior collaborators may be greater for inventors that work in fast-moving knowledge fields, where knowledge of new-to-the-world ideas could be particularly valuable. In Appendix A4, I present evidence in support for this hypothesis. In addition, in the same appendix, I show that inventors who lose junior collaborators proceed to cite older knowledge. This latter result suggests that the loss of a junior collaborator impedes the survivor from learning new-to-the-world technological ideas.

Discussion

This paper studied the relationship between the age of inventors and the spillovers that they absorb and generate, using the premature deaths as an exogenous shock to inventors' collaborative networks. Data on the birth and death years of inventors were merged to patent data to identify the age at which inventors generate the greatest spillovers, the age groups to whom those spillovers accrue, and potential mechanisms regarding inventors' productivity, relationship intensity, and knowledge recency and usefulness.

The results show that early-career collaborators generate more spillovers than mid-career collaborators, with spillovers peaking between the age 30 and 44, and that these spillovers accrue mostly to similarly-young partners. There was no evidence that knowledge transmission is passed down from older generations to younger ones. Evidence was found that knowledge is transmitted "up" across generations, in reverse-intergenerational knowledge spillovers; however, this channel was weak. Finally, the differential effect of the death of an early-career collaborator was most strongly predicted by the type of knowledge that young collaborators know: young collaborators tend to know more recently-developed ideas that are particularly useful for generating new inventions in the current economic environment. When an inventor loses a young collaborator to a premature death, their ability to access new ideas diminishes,

These findings speak to research on skill redevelopment for knowledge workers, the institutional differences between science and technological innovation, and the process of long-run knowledge growth. A central objective in skill redevelopment is the transmission of new skills to older workers (Aghion et al. 2023). Unfortunately, the results of this study show that older inventors' patenting output is only minimally affected by the presence (or disappearance) of collaborators of any age, with the implication that reverse-intergenerational spillovers may be insufficiently powerful to instigate meaningful skill redevelopment, in the absence of broader institutional or behavioral changes.

With regard to institutional differences between science and technological innovation, junior collaborators were shown to generate more spillovers for patent production than mid-career collaborators. A similar test of the effect of collaborator age at death has not been carried out for scientists, but it is plausible that that junior scientists may not generate greater publication spillovers than older scientists, because of the friction that preeminent scientists exert against the introduction of new ideas to their fields (Azoulay, Zivin, and Wang 2010; Azoulay, Fons-Rosen, and Zivin 2019). Relative to science, in technology market-based competition may be stronger, which allows young inventors to diffuse their novelties more readily.

Finally, with regard to long-run knowledge growth, the larger spillovers generated by younger collaborators, and in particular collaborators with recently-vintaged human capital, suggests that the value of ideas depreciates over time (Aghion and Howitt 1992). Knowledge depreciation is not considered in canonical models of endogenous growth (Romer 1990; Jones 2009; 2010), with the implication they may overstate the size of the stock of knowledge in the economy, as well as the educational burden associated with bringing new generations of inventors to the research frontier.

References

Aghion, Philippe, Ufuk Akcigit, Ari Hyytinen, and Otto Toivanen. 2023. "A Year Older, A Year Wiser (and Farther from Frontier): Invention Rents and Human Capital Depreciation." *The Review of Economics and Statistics*, March. https://doi.org/10.3386/W29863.

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth Through Creative Destruction." *Econometrica* 60 (2): 323–51.

Azoulay, Pierre, Joshua S. Graff Zivin, and Jialan Wang. 2010. "Superstar Extinction." *Quarterly Journal of Economics* 125 (2): 549–89. https://doi.org/10.1162/QJEC.2010.125.2.549.

Baker, Andrew, David Larcker, and Charles Wang. 2022. "How Much Should We Trust Staggered Difference-In-Differences Estimates?" *Journal of Financial Economics* 144 (2): 370–95.

Balsmeier, Benjamin, Lee Fleming, and Sonja Lück. 2023. "Isolating Personal Knowledge Spillovers: Co-Inventor Deaths and Spatial Citation Differentials." *American Economic Review: Insights*.

Bloom, Nicholas, Charles I. Jones, John van Reenen, and Michael Webb. 2020. "Are Ideas Getting Harder to Find?†." *American Economic Review* 110 (4): 1104–44. https://doi.org/10.1257/aer.20180338.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs" 134 (3): 1405–54.

Chari, V. V., and Hugo Hopenhayn. 1991. "Vintage Human Capital, Growth, and the Diffusion of New Technology." *Journal of Political Economy* 99 (6): 1142–65. https://doi.org/10.1086/261795.

Esposito, Christopher, and Frank Van Der Wouden. 2022. "Learning Amid Technological Disruption: Evidence from U.S. Inventors, 1836-1999." *Https://Doi.Org/10.5465/AMBPP.2022.17898abstract* 2022 (1). https://doi.org/10.5465/AMBPP.2022.17898ABSTRACT.

Hall, Bronwyn H., Jacques Mairesse, and Laure Turner. 2005. "Identifying Age, Cohort and Period Effects in Scientific Research Productivity: Discussion and Illustration Using Simulated and Actual Data on French Physicists." http://www.nber.org/papers/w11739.

Jones, Benjamin F. 2009. "The Burden of Knowledge and the 'Death of the Renaissance Man': Is Innovation Getting Harder?" *Review of Economic Studies* 76 (1). https://doi.org/10.1111/j.1467-937X.2008.00531.x.

———. 2010. "Age and Great Invention." *Review of Economics and Statistics* 92 (1). https://doi.org/10.1162/rest.2009.11724.

Kaltenberg, Mary, Adam B. Jaffe, and Margie E. Lachman. 2023. "Invention and the Life Course: Age Differences in Patenting." *Research Policy* 52 (1): 104629. https://doi.org/10.1016/J.RESPOL.2022.104629.

Kaltenberg, Mary, Adam Jaffe, and Margie Lachman. 2021. "The Age of Invention: Matching Inventor Ages to Patents Based on Web-Scraped Sources." *NBER Working Paper Series*.

Liu, Lu, Yang Wang, Roberta Sinatra, C. Lee Giles, Chaoming Song, and Dashun Wang. 2018. "Hot Streaks in Artistic, Cultural, and Scientific Careers." *Nature 2018 559:7714* 559 (7714): 396–99. https://doi.org/10.1038/s41586-018-0315-8.

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

Sinatra, Roberta, Dashun Wang, Pierre Deville, Chaoming Song, and Albert László Barabási. 2016. "Quantifying the Evolution of Individual Scientific Impact." *Science* 354 (6312). https://doi.org/10.1126/SCIENCE.AAF5239/SUPPL_FILE/SUPPLEMENTARYDATA.ZIP.

Strezhnev, Anton. 2023. "Decomposing Triple-Differences Regression Under Staggered Adoption." *arXiv*. https://doi.org/10.48550/arXiv.2307.02735.

Appendices for Online Publication

A1) Data Construction

In this section, I describe steps taken to clean the data on inventor birth and death information and to link inventors to their deceased collaborators.

The Kaltenberg, Jaffe, and Lachman (2023) dataset contains birth and death records that have varying likelihoods of being accurate. Because their data matches inventors to online-scraped records and patent data by inventor name and residential location, inventors that have patented while living in different places may contain multiple records in their dataset. In addition, the raw scraped records contain errors, and the match between the scraped data and the patent data can falsely link individuals. To help to alleviate these problems, Kaltenberg, Jaffe, and Lachman (2023) score the probable accuracy of each record in their dataset using a points-based system. In addition, as Kaltenberg, Jaffe, and Lachman (2023) note, the scraped death records often contain information on deceased inventors' year of birth. If an inventor's birth and death records list the same year of birth, then both records are more likely to be accurate. Thus, to create a dataset of inventor births and deaths birth with the highest possible accuracy, I follow the following rules. First, I omit all birth records with accuracy scores of 0. Second, for each unique inventor, I keep only the birth record (year of birth) with the highest accuracy score. Third, I omit all death records with accuracy scores of 0. Fourth, for each inventor, I keep only the death records (year of death) with the highest accuracy score. Fifth, I keep only the death records that (a) contain birth years, and (b) can be matched to a birth record with birth year within 2 years of the death record's birth year. Following these steps, I am left with 1,309,669 "high confidence" inventor birth years and 125,338 "high confidence" inventor death years.

To link inventors to their deceased collaborators, I first identify each inventor that died between 20 and 59 years old. There are 38,997 such premature deceased collaborators (indexed by *c*). Finally, to remove right-truncation in focal inventors' high-impact patenting (patents must be created before they can be cited), I analyze patents from 1976-2013. This brings the number of premature deaths in the dataset to 13,305.

To identify the focal inventors *i* that co-invented a patent with each of the deceased collaborators, I record each inventor that co-invented a patent with each deceased collaborator within 5 years of the collaborator's death. I use a 5-year cutoff to exclude older collaborations because they may not be associated with active relationships. Moreover, a 5-year cutoff is the norm used in academia to define active collaborations and thus potential conflicts of interest when requesting reviewers for a journal submission. I use the most recent collaboration between an inventor and a collaborator when computing this time lag. For example, I include an inventor-collaborator pair in the dataset if they collaborated twice during their careers, 7 years before and 3 years before the collaborator's death. In Appendix Figure A1, I show that the general results are robust to excluding collaborators that die between ages 55 and 59, and in Appendix Figure A2, I show that the main results are robust to using a 3-year cutoff value when identifying inventors' collaborators.

An inventor *i* can experience multiple premature deaths of collaborators *c* during her or his career. In such cases, I omit all deaths of collaborators *c* after the first experienced by an inventor *i*. In addition, inventors that co-invent with deceased collaborators can also die during the study timeframe. In this case, I keep the focal inventors *i* in the dataset up to the year of their death.

Table A1: Descriptive Statistics

	Quantile		
Variable	**25%**	**50%**	**75%**
Surviving Inventor Birth Year	1952	1961	1969
Surviving Inventor Cohort Year	1991	1998	2004
Surviving Inventor Career Length (Years)	7	14	22
Surviving Inventor Total Patents	4	9	22
Surviving Inventor Total Top Quartile Patents	1	2	7
Surviving Inventor Age at Collaborator Death	39	46	55
Year of Collaborator Death	2004	2009	2013
Age of Collaborator at Death	43	50	55
Age of Collaborator References at Death	9	13	18
Fertility of Collaborator References at Death	11	20	38

Note: Summary statistics are for 56,700 focal inventors that created two or more patents applied for between 1976 and 2013 and lost a collaborator of age 20-59 to a premature death. There were 13,305 inventors that died between these ages during the study period.

A2) Stacked Regression

In this section, I develop a stacked regression to test for robustness after removing any "prohibited comparisons" from the staggered difference-in-difference regression model. Standard two-way fixed effects difference-in-difference models with staggered treatments can violate the diff-in-diff assumption of parallel trends when the treatment effect is nonconstant over time, because recently treated units are used as comparison groups before the full treatment effect has taken time to stabilize. Stacked regressions can eliminate these prohibited comparisons by setting time windows around each treatment event to omit comparison groups shortly before or after each treatment event.

I follow the general outline of Cengiz et al. (2019) to develop the data stack. First, I compute the calendar year an inventors' collaborator dies, defined as $CollabDeathYear_i$. Second, from the full dataset analyzed in the main text, I extract all inventor panels for which $CollabDeathYear_i = t$. I label these inventors as the treated units in "sub-experiment d", so I create the variable, $treat_{i,d} = 1$, for the inventors in this data subset. Third, from the full dataset I subset all inventor*year observations for which $t - 5 < CollabDeathYear_i < t + 12$, meaning that these observations were taken more than 5 years before or 12 years after their collaborator died. These are the control observations for sub-experiment d, so I set $treat_{i,d} = 0$ for these observations. This ensures that the control group does not contain invalid comparisons due to pre-or-post trends associated with a loss of a collaborator. Fourth, I append the treated observations dataset to the control observations dataset. Fifth, I omit from this appended dataset all observations that fall outside the treatment window (5 years before or 12 years after year t). Finally, I repeat the above steps for all treatment years, 1980 to 2001, to create a data stack containing 5,026,225 inventor*year observations with 12,232 unique "treated" inventors and 65,915 unique "control" inventors.

To estimate the effect of the loss of a junior collaborator on surviving inventors' patenting productivity, I estimate the model described by Equation A1.

$$(A1)\quad PatProd_{i,t} = \exp\left[\sum_{\tau=-5}^{12} \Phi_\tau TimeToSubExperiment_{idt}^{\tau} + \sum_{\tau=-5}^{12} \Psi_{\tau i} TimeToSubExperiment_{idt}^{\tau} * Treat_{id} + \sum_{\tau=-5}^{12} \Omega_{\tau i} TimeToSubExperiment_{idt}^{\tau} * Treat_{id} * JuniorDeceasedCollab_i + \alpha_{id} + \rho_t + \varepsilon_{it}\right]$$

In the model, τ indexes event time (-4 to 12) and t indexes calendar time (1975 to 2013). $TimeToSubExperiment_{idt}^{\tau}$ contains indicators for event time, with 0 (the year of the event) taken as the reference group. Thus, Φ_τ gives the change in patent production over event time for the control observations in sub-experiment *d*, relative to the event year 0. $\Psi_{\tau i}$ gives the patent differential for the inventors treated in sub-experiment *d*. The coefficients of interest are $\Omega_{\tau i}$, which gives the differential for losing a junior collaborator, for treated units in sub-experiment *d*. The model also includes individual*sub-experiment and calendar time fixed effects. Standard errors are clustered at the sub-experiment*deceased collaborator level.

As in the main analysis in the paper's full text, I estimate Equation A1 four times to explore robustness of the results. The coefficient values are plotted in Figure A1. The results are similar to the ones presented in the main analysis. Specifically, inventors who lose junior collaborators to premature deaths suffer a larger decline in subsequent patenting than do inventors who lose mid-career collaborators. The treatment effect becomes significant 7 years after the collaborators' death.

Figure A1: Change in Patenting Productivity Following Death of Junior Collaborator Relative to Mid-Career Collaborator using Stacked Regression

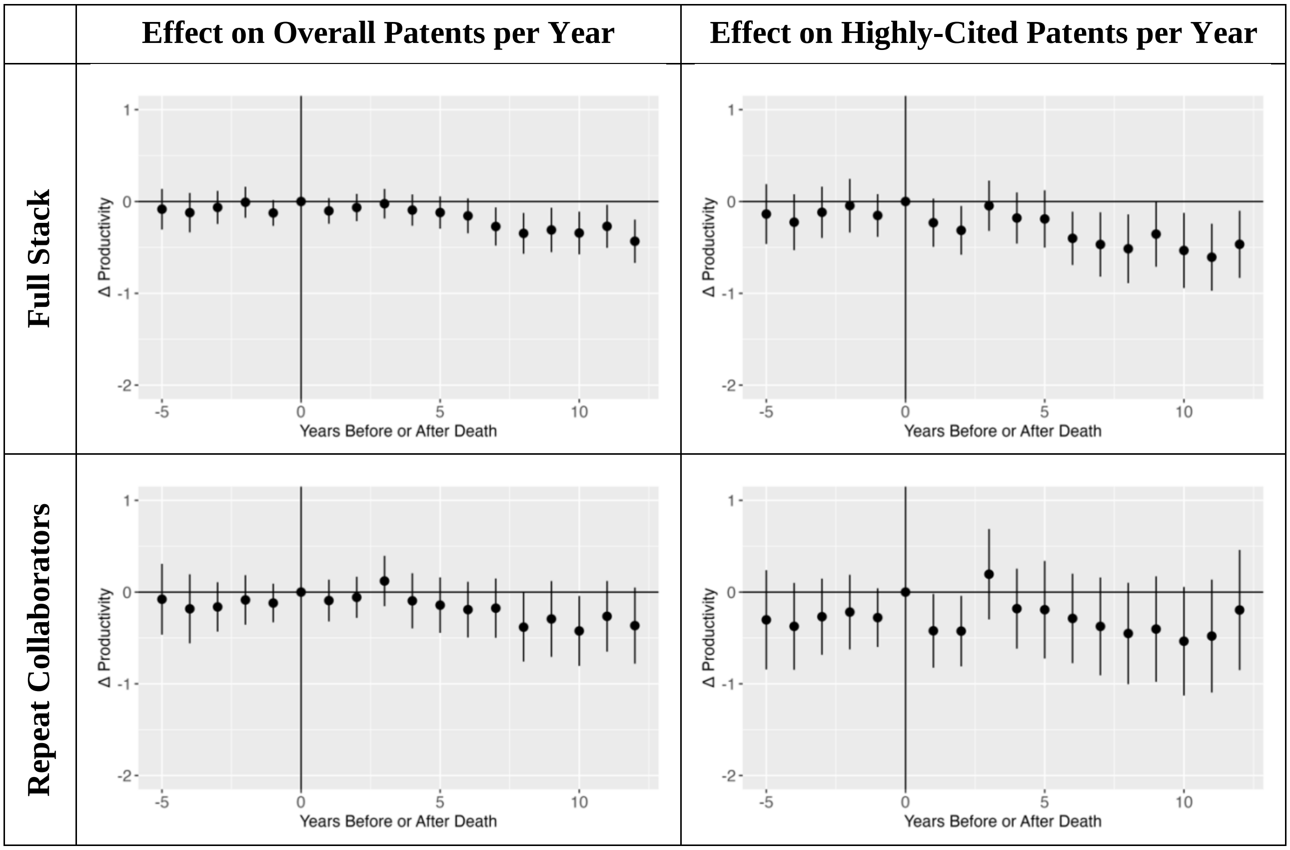

Notes: Figures shows effect of losing an early-career collaborator (age 20-44 at time of death) relative to losing a mid-career collaborator (aged 45-59 at death) on patenting productivity using a stacked regression. Standard errors are clustered at the deceased collaborator * sub-experiment level.

A3) The Relationship between Collaborator Age, Knowledge Age, and Knowledge Fertility

In the main text, Table 2 showed that inventors suffer large declines in patent production when they lose collaborators that have knowledge of newly-introduced and fertile ideas. In addition, the table showed that these relationships can fully erode the explanatory power of a collaborators' age at time of death on the patenting productivity of their co-inventors.

In this section of the appendix, I show that younger inventors have knowledge of newer and more fertile ideas. To conduct this test, I administer a regression where the mean age and mean fertility of citations made in collaborators' patents in the three years pre-death are a function of the collaborators' age at time of death and death year fixed effects. Specifically, the OLS regression model is given by Equation A2:

$$\text{(A2)} \quad MeanCiteAgePreDeath_j = \beta_0 + \beta_1 AgeAtDeath_j + \lambda_t + \varepsilon_j$$

In the model, j indexes deceased collaborators. Each collaborator appears only once in the regression. λ_t are calendar year effects for the year of the collaborator's death. The regression for $MeanCiteFertilityPreDeath_j$ is similar. Estimates for Equation A2 using the full set of deceased collaborators are provided in Table A2. The results in the table show that older collaborators cite older and less-fertile patents.

Table A2: Regression Results for Equation A2 on Collaborator Age, Knowledge Age, and Knowledge Fertility

	Dependent Variable	
	$PreDeathMeanCiteAge_j$	$PreDeathMeanCiteFertility_j$
$AgeAtDeath_j$	0.0427*** (0.00884)	-0.00221** (0.00116)
Death Year Fixed Effects?	Yes	Yes
R^2	0.021	0.004

Observations (Collaborators)	9,757	10,135

Notes: PreDeathMeanCiteAge and PreDeathMeanCiteFertility are computed using the patents applied for by the deceased collaborator in the 3 years before their death.

A4) Field Speed Analyses

If knowledge obsolescence causes early-career inventors to generate more patenting spillovers than mid-career inventors, one would expect two patterns. First, the effect size should be larger for inventors that work in fast-advancing knowledge fields. Second, the loss of an early-career collaborator should impede the ability of their surviving partners from learning new ideas in recently developed areas of technology. To test the first proposition, I measure the rate of advance in inventors' technological fields by calculating the average age of the citations made in each CPC technology field and year. Decomposing the effects by the rate of field advance in which each inventor primarily patents, I test whether inventors that primarily work in fast-advancing fields experience a greater decline in patenting productivity following the loss of an early-career collaborator than do inventors who work in slower-advancing fields. There are 123 unique technology classes in my dataset at this level of the CPC classification scheme.

To test whether the differential effect of losing a junior collaborator is greater in fast-advancing knowledge fields, I run the main regression described by Equation 1 separately for surviving inventors that primarily patent in fast-advancing and slow-advancing knowledge fields. I define fast and slow fields in each year based on whether the mean age of the citations made in the class are above or below the yearly median. Because some patents cite pre-1976 patents, for which I do not have application year data available, I use grant years to compute citation age. I plot the full host of β_1 coefficients for each regression in Figure A3.

Figure A3 shows that the negative effect of losing an early-career collaborator on subsequent patenting is more pronounced in fast-advancing knowledge fields. Inventors in fast-advancing knowledge fields that lose early-career collaborators experience a significant and immediate decline in their subsequent overall patenting rate, while inventors in slow-advancing knowledge fields experience a smaller decline that takes longer to take effect. This difference is greater for inventors who lose repeat collaborators.

Figure A3: Change in Patenting Productivity in Slow and Fast-Advancing Knowledge Fields

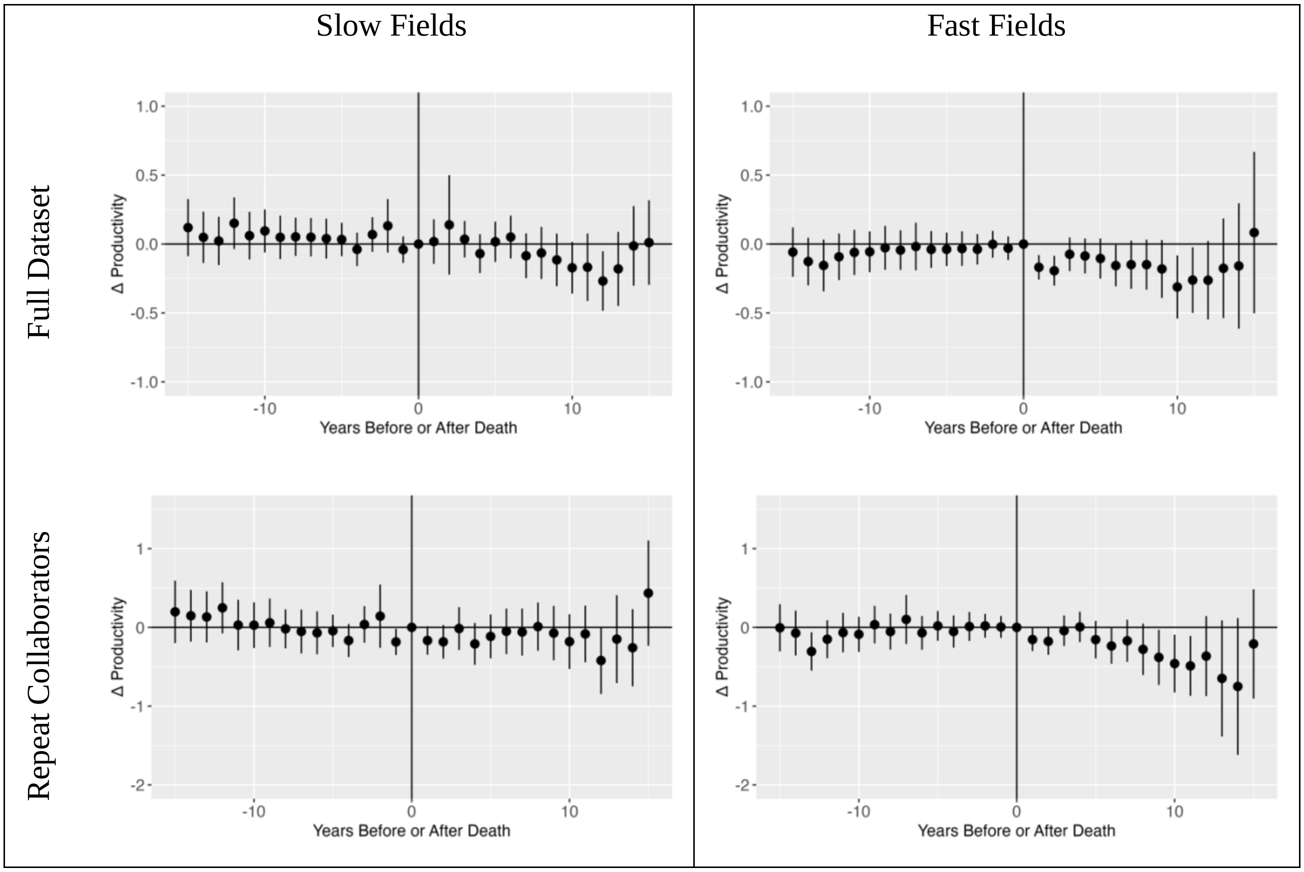

Notes: Figure shows effect of losing an early-career collaborator on subsequent patenting rate, relative to losing a mid-career collaborator. Slow and fast-advancing knowledge fields are identified based on the average age of the patent citations in each inventor's modal technological field, as described in the text.

The above results suggest that the loss of an early-career collaborator impedes inventors' ability to learn new ideas that are at the frontier of their fields. To test this proposition directly, I examine whether inventors proceed to cite older patents following the loss of an early-career collaborator. To do so, I calculate the mean age of the patents cited by a patent *p*. Using the resulting measure of the age of the knowledge cited by each patented invention, I develop a regression model that tests whether inventors combine older ideas following the loss of an early-career collaborator. The OLS model is given by Equation A3:

$$\text{(A3)} \quad \log\left(MeanCitationAge\right)_p = \beta_0 PostCollabDeath_{i,t} + \sum_{m=20}^{59} \beta_m PostCollabDeath_{it} * AgeCollabAtDeath_i^m + \alpha_i + \tau_{ft} + \varepsilon_{i,t}$$

The reference group for the variable $AgeDeceasedCollab_i$ is a collaborator that dies between ages 55 and 59. Therefore, the β_m coefficients are interpreted as the change mean subclass age of the USPC subclasses on an inventors' patents following the loss of an collaborators in a specific age range, relative to collaborators that die between 55 and 59. τ_{f*t} are indicator variables for the primary technology class on patent *p*, interacted with year *t*. These indicator variables adjust for the mean subclass age in the primary technology class of patent p in the year in which it is applied for. Thus, the coefficient of interest (β_m) is interpreted in relation to the mean age of the knowledge produced in the same field and year. I estimate Equation A2 twice, first using inventor and year effects, and second using inventor and field*year effects. This separate estimation allows me to test whether inventors respond to the loss of younger collaborators by switching to slower-advancing technology fields. The plots are in Figure A4.

Figure A4: Change in Knowledge Vintage by Collaborator Age at Death

Inventor and Year Effects	**Inventor and Field*Year Effects**
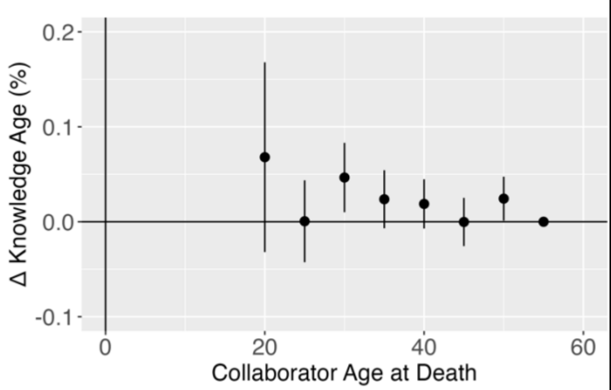	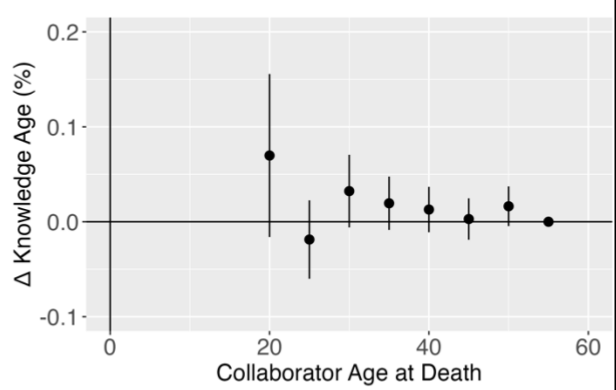

Notes: Figures show the differential effect of losing a collaborator at a specific age at of death relative time (grouped into 5 year bins) to losing a collaborator aged 55-59 at time of death on the vintage of the surviving inventors' primary subclasses on patents.

Figure A4 shows that inventors combine ideas that are 25-30% older if they lose a collaborator that dies in her or his 20s instead of a collaborator that dies in her or his 50s. The coefficients associated with collaborators that die in their 20s are noisy, but generate some evidence that very young collaborators may also act as conduits for inventors to access very recently-developed ideas. The effects are similar when field*year effects are introduced in the model, suggesting that inventors generally do not switch to slower-advancing technology classes following the loss of junior collaborators. Instead, inventors who lose junior collaborators continue to patent in technological fields that advance at similar rates, but they combine knowledge that is further from the forefront of those fields.

A5) Viability of Early-Career Deaths as Counterfactual for Mid-Career Deaths

To determine whether collaborators who die in their early careers are valid counterfactuals for inventors that die mid-career, in the left panel of Figure A5 I plot the patents per year produced by deceased collaborators at each age. The data are broken out by collaborators that die early-career and mid-career, with 45 years old at death as the breakpoint.

Collaborators that die in their early careers produce more patents per year than collaborates who die in the mid careers at very low ages (below age 34). This difference implies that collaborators who die early are highly productive at a young age, which leads them to have many collaborators before they die young. However, the patenting advantage of the early-career deceased collaborators stops after age 35, when their patenting output converges to the rate of collaborators that die mid-career. This convergence implies that the early-career deceased do not reach higher peak productivity levels than do inventors who die in their mid-careers. Therefore, a focal inventor who loses an early-career collaborator does not lose a collaborator that would have reached higher peak productive than an inventor who loses a mid-career collaborator.

Despite their similar peak productivity, collaborators that die early-career have higher patenting productivity in their *very early careers* (before age 35) than do collaborators that die mid-career. This difference implies that an inventor who loses a collaborator at a very young age (before age 35) could suffer a loss of a more productive collaborator. To ensure that the productivity differences of early-career collaborator deaths before age 35 do not bias the analysis, I also reproduce the main event study, omitting all collaborators that die before age 35. Therefore, in the robustness check, I define early-career collaborator deaths as those that occur between ages 35-44, and mid-career collaborator deaths as those that occur between ages 45-59. I exclude other collaborator deaths from the analysis. The regression is otherwise identical to the one described by Equation 1 in the main text. I plot the differential treatment effect of losing a junior collaborator, relative to a mid-career collaborator, in the right panel of Figure A5. The event plot is very similar to the one in the main analysis that does not exclude collaborates that die between ages 20 and 34: the pre-trends are parallel, and the treatment effect becomes significant and negative 7 years after the collaborator’s death.

Figure A5: Analysis of Viability of Mid-Career Collaborator Deaths as a Counterfactual for Early-Career Collaborator Deaths

(A) Mean Patents per Year for Collaborators that Die Mid-Career and Early Career

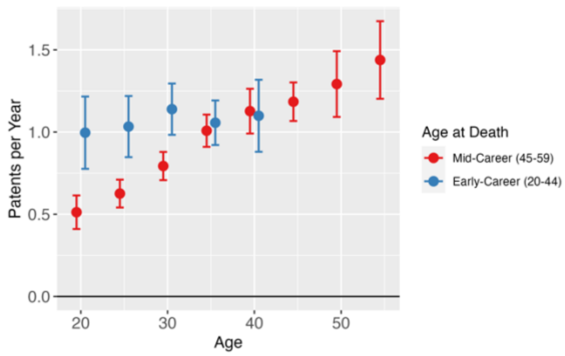

(B) Differential Change in Patents per Year Following Death of Junior Collaborator, Omitting Collaborator Deaths Below Age 35

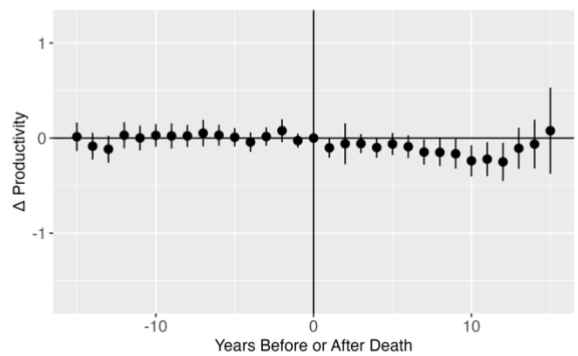

A6) Addition Results

This appendix shows results where inventors that die between ages 55-59 are omitted from the data, and where a window of 3 years is used to define inventors' past collaborators. The results conform to the results presented in the main analysis.

Figure A6: Change in Patenting Productivity by Age of Deceased Collaborator, Deaths Under Age 55

(A) Effect on Overall Patenting	(B) Effect on High-Impact Patenting
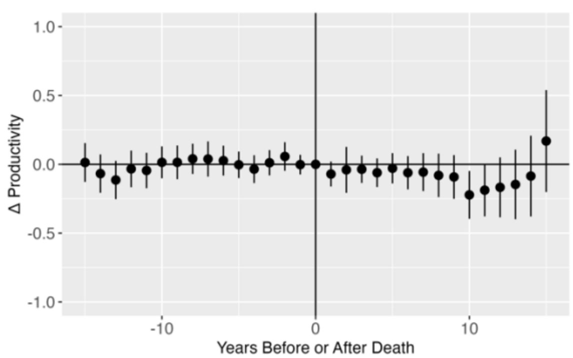	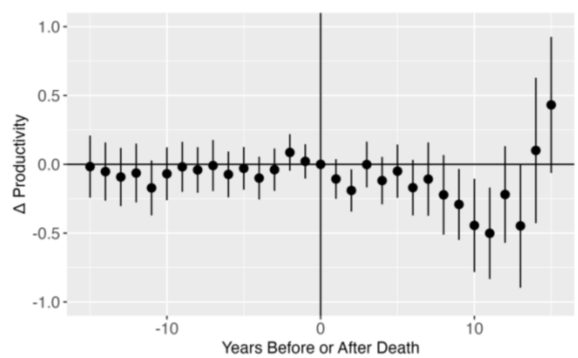

Notes: Figures shows effect of losing an early-career collaborator (age 20-44 at time of death) relative to losing a mid-career collaborator (aged 45-54 at death) on patenting productivity. All models include focal inventor and year*technology class fixed effects. Standard errors are clustered at the deceased collaborator.

Figure A7: Change in Patenting Productivity by Age of Deceased Collaborator, Max 3 Years Between Collaboration and Death

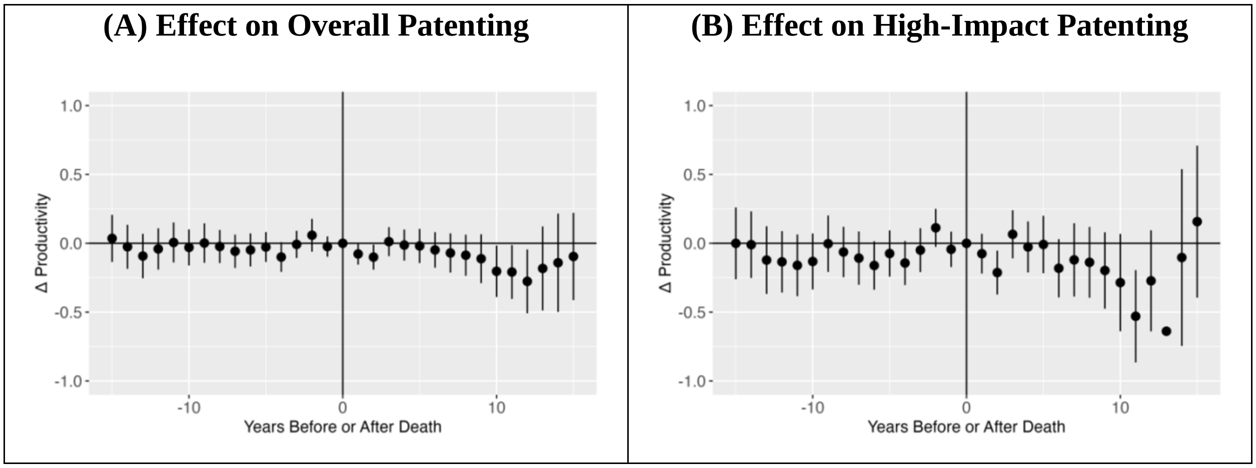

Notes: Figures shows effect of losing an early-career collaborator (age 20-44 at time of death) relative to losing a mid-career collaborator (aged 45-59 at death) on patenting productivity. A focal inventor's deceased collaborators are defined as those that a focal inventor co-invented a patent with in the 3 years leading up to the collaborator's premature death. This 3-year threshold contrasts with the 5 year threshold used in the other analyses. All models include focal inventor and year*technology class fixed effects. Standard errors are clustered at the deceased collaborator.