Age, Death, and the Nucleus of Invention*

Christopher Esposito

Working Paper

Updated May 10, 2023

Abstract

This paper identifies the relationship between the age of inventors and knowledge spillovers that they generate. Linking age and death information from 13,112 patent inventors that died prematurely to 54,674 co-inventors, I show that inventors who lose an early-career collaborator to a premature death produce 0.1 fewer patents per year and 0.14 fewer highly cited patents per year than do inventors who lose mid-career collaborators. The effect size is greater for inventors who lose repeat collaborators and for inventors who work in fast-advancing technological fields. The results are consistent with a model in which knowledge obsolescence depreciate the value of vintaged human capital.

* University of Chicago (cresposito@uchicago.edu). I thank Luis Bettencourt, Eamon Duede, James Evans, Lee Fleming, Donghyun Kang, Jeffrey Lockhart, Olav Sorenson, Frank van der Wouden, and participants at the Mansueto Institute Colloquium Series for helpful comments. All errors are my own.

1. Introduction

Models of human capital accumulation typically argue that individuals accumulate human capital over time through experience-based learning and the diffusion of knowledge from senior mentors to junior mentees (Jones 2010). In technology, these models propose that new generations of inventors principally learn from senior mentors and managers. However, a growing body of research suggests that senior inventors' technological expertise may be too dated to train their mentees in the latest areas of technology. Aghion, et al. (2023) show that the ability for white-collar workers to appropriate the rents of new technologies decreases with age, and Kaltenberg, Jaffe, and Lachman (2023) demonstrate that inventors' patenting productivity peaks before the age of 40. Likewise, Esposito and van der Wouden (2022) find that the likelihood for inventors to create high-impact patents declines monotonically with career age, and that this relationship has strengthened in recent years. If, as these studies indicate, senior inventors are not particularly inventive, it begs the question of from whom young inventors learn their most valuable ideas.

This study proposes that inventors learn their most valuable ideas from junior collaborators. Junior collaborators are an important source for knowledge because they have recently completed their foundational training and thus have recently-vintaged human capital (Chari and Hopenhayn 1991). Because technological change is not a cumulative process – the development of new technological ideas often drives older ideas into obsolescence – new technological ideas are often more valuable than older ones (Katila 2002). As knowledge workers age, their search for new technological ideas narrows and they increasingly build on ideas that were developed years before (Cui, Wu, and Evans 2022; Esposito and Wouden 2022). Thus, younger collaborators tend to possess high-value and non-substitutable knowledge bases that are at the knowledge frontier. Through collaborations, their knowledge diffuses to their peers.

To study the relationship between the age of collaborators and the spillovers that they generate, I contrast the change in inventors' patenting productivity following the premature death of early-career collaborators (20-44 years old at time of death) with the change in patenting productivity following the premature death of mid-career collaborators (45-59 years old at time of death). When an inventor loses a collaborator to a premature death, their ability to learn knowledge in the deceased collaborator's area of expertise is impeded. When deceased collaborators die early in their career, their surviving collaborative partners lose a pipeline through which they could source recently-developed knowledge. This allows the relationship between biological age and knowledge spillovers to be studied in a causal setting. In my analysis, I omit deceased collaborators that die late in their careers (60+ years old at time of death), because such deaths could be more readily anticipated by their collaborators.

I find that inventors who lose early-career collaborators to premature deaths proceed to produce 0.1 fewer patents per year relative to inventors who lose mid-career collaborators. This difference becomes statistically significant the year after the death of the collaborator and endures for twelve years, after which it becomes indistinguishable from statistical noise. In percentage terms, the effect size translates into an 14% reduction in patent production relative to the median inventor. I find related and somewhat larger effect sizes for the production of highly cited patents: inventors produce 0.14 fewer patents per year in the top quartile of the citation distribution of their patent's grant year and CPC technology class field after losing an early-career collaborator. Pre-trend analyses confirm that these results satisfy the difference-in-difference assumption of parallel trends.

After these main analyses, I disaggregate the data by the age of collaborators at the time of collaborator death and the age of inventors at the time of collaborator death to identify the age at which inventors generate and absorb the most knowledge spillovers. I find that collaborator age at time of death as a U-shaped relationship with generated spillovers: very young collaborators (aged 20-29 at time of death) and relatively old collaborators (aged 50-59 at time of death) generate fewer knowledge spillovers than collaborators aged 30-49 at time of death. This U-shaped relationship suggests that very young collaborators lack sufficient experience to generate knowledge spillovers (Jones 2010), while older collaborators suffer from knowledge obsolescence (Aghion et al. 2023). Disaggregating the treatment effects by the age of the *focal inventor* at the time of the collaborators' death reveals that inventors are generally more reliant on collaborators that are younger than they are. For example, the productivity of inventors in their 40s is most negatively affected when they lose a collaborator in their 30s. This finding suggests that younger collaborators contribute more valuable, novel, and non-substitutable ideas to their teams.

Finally, analyze whether the knowledge obsolescence is a plausible mechanism that causes early-career collaborators to generate more spillovers than mid-career collaborators. I perform this analysis in two stages. First, I disaggregate the data to study how the treatment effect differs for inventors that principally patent in slow-advancing fields, where new inventions tend to combine old ideas, and for inventors that principally patent in fast-advancing fields, where new inventions tend to combine recently introduced knowledge. If early-career collaborators generate more spillovers than mid-career collaborators because their human capital is less obsolete, one would expect the treatment effect to be stronger for inventors that patent most often in fast-advancing knowledge fields. The results of the analysis supports this interpretation, showing that the differential effect of losing an early-career collaborator on subsequent productivity is negligible in slow-advancing fields, but it is substantial in fast-advancing knowledge fields.

In the second stage, I test whether the loss of a younger collaborator causes inventors to combine more dated ideas in their inventions. To perform this test, I compare how the vintage of the ideas that inventors combine in their patents changes for those who lose early-career collaborators with those who lose mid-career to premature deaths. I find that inventors that lose very young collaborators (collaborators in their 20s at time of death) proceed to combine ideas that are on average 2.5 years older than do inventors who lose older collaborators (collaborators in their 50s at time of death). The effect size is larger when the surviving inventor is old, suggesting that older inventors' collaborative networks are not resilient for sourcing recently developed ideas. In particular, I find inventors in their 50s proceed to combine ideas that are 8 years older following the death of a collaborator in her or his 20s. This change in the vintage of their human capital causes an inventor that was at the median of the knowledge vintage distribution pre-loss to the 70% percentile of the knowledge vintage distribution. Moreover, when older inventors lose young collaborators to premature deaths, 20% of the competition in the same technology field leapfrogs them in terms of patenting using recently-developed ideas.

The results of this study address a debate in the economics of innovation literature on the relationship between age and performance in knowledge-intensive occupations. Previous studies are deeply conflicted about whether the relationship between age and invention is positive, negative, or more complex. Jones (2009) found that the age at which Nobel Laureates and great inventors achieved their leading intellectual accomplishments increased over the 20^{th} century, but recent studies using more extensive datasets have questioned the generality of this finding. Using a dataset containing age information for 1.2 million U.S. patent inventors, Kaltenberg, Jaffe, and Lachman (2023) show that the patenting output of inventors peaks between ages 36 and 38 and that the modal age at which inventors receive their first patent declined during the 21^{th} century.

Esposito and van der Wouden (2022) generate related results by showing that inventors are most likely to create a high-impact patent in the first year of their patenting careers. In addition, research on academic scientists' hot streaks has also generated mixed evidence on the relationship between age and performance in creative and knowledge-intensive occupations: while most scientists experience at least one hot streak of elevated productivity during their careers, the onset of a hot streak is equally likely to occur during the beginning, middle, or end of a scientists' careers. This result indicates that there is no concrete relationship between age and onset of hot streaks (Sinatra et al. 2016; Liu et al. 2018; 2021). The present study contributes to this debate by identifying the biological age at which inventors generate the most knowledge spillovers. The analysis of spillovers by age is important, because it helps to uncover how long-run technological knowledge advance is sustained across generations of inventors (Romer 1990; Jones 2009).

By linking the age demographics of inventors to the vintage of their knowledge bases, this study also helps to resolve a broader question about whether technology progresses through incremental or radical advances (Kuhn 1962; Foster, Rzhetsky, and Evans 2015). If technology advanced incrementally, one would expect inventors to become more productive and generate more spillovers with age, because older inventors have had the time to accumulate more ideas that they can combine into technologies and diffuse to their mentees (Jones 2010). However, because junior inventors generate more spillovers than do older inventors, technological advance is evidently a disruptive enterprise, at least to the extent that new technological advances obsolesce the human capital of highly specialized individual inventors.

Finally, these results also help to identify the potential drivers of the ongoing R&D productivity decline (Jones 2010; Bloom et al. 2020). A leading hypothesis for the R&D productivity decline is that the accumulation of knowledge over time creates an educational burden, as more knowledge needs to be absorbed for new inventors to be brought up-to-pace with the knowledge frontier (Jones 2009). Notably, the education burden hypothesis assumes that technological knowledge growth is a cumulative process. If, as the results of this study suggest, new technological advances occasionally drive the knowledge known by incumbents into obsolescence, the educational burden is less significant than is commonly assumed: some of the stock of accumulated knowledge that could potentially have created an educational burden for new inventors is in fact obsolete knowledge. In a regime of innovation that includes technological disruption, an R&D productivity decline would emerge not only from a rising education burden, but from a rising rate of knowledge obsolescence, as increased R&D expenditures lead to increased knowledge obsolescence and thus reduce the societal returns to R&D.

Because this study compares the spillovers generated by early-career inventors with those generated by mid-career inventors, its identification strategy contrasts between surviving inventors that are not identical. An advantage to this approach is that the lag time to treatment is matched for inventors who lose juniors collaborators with those who lose mid-career collaborators, eschewing several of the identification issues regarding the study of the effects of premature deaths described by Azoulay, Zivin, and Wang (2010). A downside of this identification strategy is that the age of a deceased collaborator may be associated with unobserved confounding factors. Some, but not all, of the potential confounding factors correlated with age at time of a collaborators' death are dealt with by the modeling strategy. For example, inventors who lose mid-career collaborators to premature deaths may themselves be more mature in their careers, which prior studies have shown to be associated with changes in inventive performance (Esposito and Wouden 2022; Kaltenberg, Jaffe, and Lachman 2023). I address this concern by using inventor and year-specific fixed effects. Because linear inventor age is perfectly collinear with these two fixed effect terms, linear age effects are projected out of the variation in the dependent variable. The resulting regression coefficients are interpreted as changes in the productivity of inventors, conditional on the mean productivity of an inventor, conditional on the mean productivity in a given year, and conditional on all unobserved trends that are linear within inventors. This approach does not fully eliminate the concern that a contrast between different collaborator ages at time of death can pick up on confounding factors, but it does account for the confounding factors that are explained by the linear age of the surviving inventor.

Through this methodology, the current study also generates a methodological contribution to the literature on the career dynamics of inventors and scientists. Studies of inventor and scientist career dynamics frequently analyze how knowledge production changes across careers, but these processes cannot be fully identified because of the collinearity between age, period, and cohort (APC) effects. This issue is well-documented in demographic-related literatures and is discussed with an application to scientific knowledge workers by Hall, Mairesse, and Turner (2005). The core issue is that when year and individual fixed effects are included in a model, variables that are linear over time such as age or experience cannot be included in the model. This study circumvents the APC issue by analyzing the effect of a collaborator's age on patent productivity, rather than the age of the focal inventor. This methodological approach can be useful to learn how age and experience affects individuals in other domains, so long as there are interactions between agents and there is a plausible reason to expect the performance or behaviors of the interacting agents to be positively correlated. For example, studies that try to discern how preferences change with age can study how the age at time of deaths of interacting partners changes the preferences of survivors using models with period and cohort effects.

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 three sources. The first source is the set of all patents granted by the U.S. Patent and Trademark Office between 1976 and 2018, sourced through PatentsView. There are 6.2 million patents in this dataset granted to 3.2 million inventors. For each patent, I extract the patent number, application year, CPC technological class and technological group, and disambiguated inventor IDs. The second data source is a dataset recently made available by Kaltenberg, Jaffe, and Lachman (2023) that contains information on the year of birth and year of death records for 1.9 million inventors in the PatentsView dataset. Kaltenberg, Jaffe, and Lachman (2023) compiled the dataset by scraping three websites that collect birth information from birth records and two websites that collect death information from obituaries, and by matching these scraped records to USPTO patent inventors using inventor names residential locations. The data collection procedures, descriptive statistics, and limitations of the Kaltenberg, Jaffe, and Lachman (2023) dataset are thoroughly discussed in their 2023 article as well as in a 2021 NBER working paper (Kaltenberg, Jaffe, and Lachman 2021). The 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 has not been used to analyze how the externalities generated by patent collaborators vary by the age of the deceased collaborator. Using a different dataset on the deaths of star scientists, Azoulay, Graff Zivin, and Wang (2010) show in their supplemental materials that superstar deaths have a larger effect on younger surviving scientists than on older ones, but they do not analyze how the impact of superstar deaths varies with the age of the deceased.

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, Kaltenberg, Jaffe, and Lachman (2023) note that their 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.

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 unique 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 between 1975 and 2018. There are 38,997 such "premature deceased collaborators" (indexed by c). 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. Because my inventor fixed effects models cannot be estimated for single-patent inventors, I omit all single-patent inventors from the dataset and the descriptive statistics.

To study whether the loss of early-career collaborators affects inventors' ability to create high-impact patents, I combine my dataset with counts on the number of forward citations received by patents. I count the forward citations that patents receive within 5 years of their grant year. 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.

Finally, to compute inventors' patenting productivity, I aggregate inventors' count of total patents and their count of high-impact patents by inventor and application year. Because patents need to be applied for several years before they receive a significant share of their citations, I restrict the dataset that I analyze to the application years 1975-2013. Because my full dataset runs through the 2018 grant year, the 1975-2013 application year subset allows for patents that were applied for in 2013 to still receive patents granted up to 5 years later.

Following the data subsets described above, I am left with 54,673 focal inventors and 13,112 collaborators that died prematurely during the study period. I provide summary statistics for these data in Table 1. The summary statistics show considerable variation in the age of surviving and deceased inventors across the dataset, 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 frame, the distribution of age at collaborator death is capped at a value of 59.

Table 1: Descriptive Statistics

	Quantile		
Variable	**25%**	**50%**	**75%**
Surviving Inventor Birth Year	1952	1960	1968
Surviving Inventor Cohort Year	1990	1998	2003
Surviving Inventor Career Length (Years)	6	13	20
Surviving Inventor Total Patents	4	9	21
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

Note: Summary statistics are for 54,673 focal inventors that created two or more patents between 1975 and 2013 and lost a collaborator of age 20-59 to a premature death. There were 13,112 inventors that died between these ages during the study period.

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:

$$(1) PatProd_{i,t} = \beta_0 PostCollabDeath_{i,t} + \beta_1 PostCollabDeath_{i,t} * JuniorDeceasedCollab_i + \alpha_i + \tau_{f*t} + \varepsilon_{i,t}$$

In Equation 1, $PatProd_{i,t}$ is the number of patents produced by focal inventor i in the application year *t*, $PostCollabDeath_{i,t}$ is a binary variable that equals 0 for the 10 years leading up to 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 inventors outside the 20-59 age range 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' primary CPC technological class, defined as the most frequent technology class each inventor patents, interacted with year indicators. 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 inventor ages. 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 inventors' rate of producing high-impact patents. High impact patents are defined as those in the top quartile of their grant year and CPC class distribution, as described in the data section.

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. 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 2. The first column shows that inventors produce 0.443 fewer patents per year following the death of a collaborator. The coefficient on the interaction term, -0.0998, indicates that inventors that lose junior collaborators subsequently produce 0.0998 fewer patents per year than surviving inventors that lose mid-career collaborators. The size of this effect can be interpreted by considering that the median surviving inventor in the dataset produces 9 patents over the course of a 13-year career, or 0.69 patents per year (Table 1). Thus, the median inventor who loses a junior will produce 14% fewer patents each year over the remainder of her or his or career. The second column of Table 2 shows that the effects on high-impact patenting are slightly larger compared to those on overall patenting. Inventors who lose a collaborator prematurely proceed to produce 0.392 fewer high-impact patents per year, and those who lose junior collaborators proceed to produce an additional 0.141 fewer high-impact patents. Thus, early-career collaborators even more important for high-impact patenting than they are for overall patenting.

As anticipated, 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 0.188 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 2: 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	-0.443*** (0.0246)	-0.392*** (0.365)	-0.574*** (0.0424)	-0.505*** (0.0618)
PostCollabDeath* JuniorDeceasedCollab	-0.0998** (0.0392)	-0.141*** (0.0496)	-0.188** (0.0605)	-0.202** (0.0780)
Inventor and Class*Year Fixed Effects?	Y	Y	Y	Y
Inventor*Year Obs	299,282	213,193	96,246	74,547
Inventor Obs	40,112	23,123	11,385	7,990

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. To do so, I estimate the regression model described by Equation 1 but replace the binary $PostCollabDeath_{i,t}$ variable with a vector of indicator variables that reflect the number of years pre or post the treatment. I plot the resulting coefficients and their associated 95% confidence intervals in Figure 1A. In Figure 1B, I generate a similar plot for the effect of losing an early-career collaborator on high-impact patenting. Figures 1A and 1B show that there are no obvious pre-trends in the data, and that the negative relationship between losing a junior emerges the year after a collaborator death and continues for 13 years following the collaborator death. In both figures, the effect size is at its maximum 10-12 years following a collaborator's death. This time dynamic indicates that the impact of losing a junior collaborator neither dissipates nor compounds. Moreover, inventors who lose junior collaborators suffer a substantial one-time shock to their patenting productivity that they never fully recover from. The results are similar, but with larger effect sizes, for inventors who lose repeat collaborators (Appendix Figure A3).

Figure 1: Change in Patenting Productivity by Age of Deceased Collaborator

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

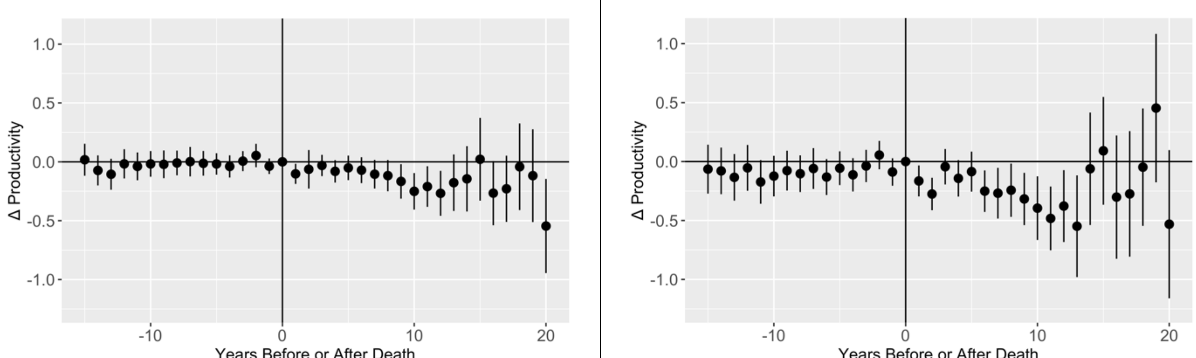

Notes: Figures shows effect of losing an early-career collaborator (age 20-44 at time of death) relative to losing a mid-career collaborator (aged 45-59 at death) on patenting productivity. All models include focal inventor and year*technology class fixed effects. Standard errors are clustered at the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

3.2 Heterogeneous Effects by Inventor and 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 2:

(2) $PatProd_{i,t} = \beta_0 PostCollabDeath_{i,t} + \beta_1 PostCollabDeath_{i,t} * AgeDeceasedCollab_i + \alpha_i + \tau_{f*t} + \varepsilon_{i,t}$

In Equation 2, the patenting productivity of a focal inventor is a function of $PostCollabDeath_{i,t}$, 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 $PostCollabDeath_{i,t} * AgeDeceasedCollab_i$. The second variable of the interaction term, $AgeDeceasedCollab_i$, records the age of the collaborator at time of death. Because of the relatively small number of observations of collaborator deaths at a specific age, I group $AgeDeceasedCollab_i$ into 10-year bins: 20-29 years old, 30-39 years old, 40-49, and 50-59. In the regression, collaborators that die between 50 and 59 are the reference group. As before, the model contains focal inventor and CPC technology class*year indicator variables. I estimate Equation 2 using a Quasi Poisson Maximum Likelihood estimator and plot the β_1 coefficients and their 95% confidence intervals for total patenting and high-impact patenting in Figure 2.

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

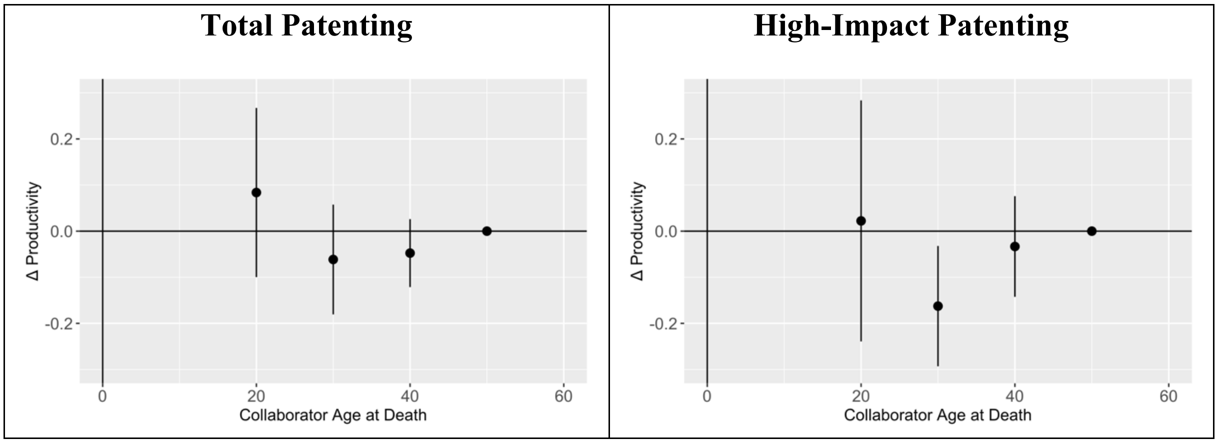

Notes: Figures show the differential effect of losing a collaborator at a specific age at of death relative time (grouped into 10 year bins) to losing a collaborator aged 50-59 at time of death on patenting output.

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 than the oldest subset of collaborators (aged 50-59 at death). The greatest spillovers are generated by collaborators between the ages of 30 and 49. This U-shaped relationship is more pronounced for high-impact patenting than for patenting overall. The U-shaped relationship indicates that the greatest volume of patenting spillovers flow from inventors that are sufficiently experienced (Jones 2010), but not too old to be at severe risk of knowledge obsolescence (Aghion et al. 2023). Appendix Figure A4 shows that when the sample is restricted to repeat collaborators. the results are similar but with larger effect sizes.

Next, I analyze age at which *focal inventors* are most affected by the loss of collaborators of different ages. I do so by estimating Equation 3:

(3) $$PatProd_{i,t} = \beta_0 PostCollabDeath_{i,t} + \beta_1 PostCollabDeath_{i,t} * AgeInventorAtCollabDeath\_AgeDeceasedCollabAtDeath_i + \alpha_i + \tau_{f*t} + \varepsilon_{i,t}$$

The coefficient matrix β_1 records the change in patenting productivity for inventors of a specific age at the time of death of their collaborators who lose a collaborator at a specific age. As before, 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 and high-impact patenting using Quasi Poisson Maximum Likelihood model and plot heat maps of the β_1 coefficients for overall patenting in Figure 3A and for high-impact patenting in Figure 3B. I place asterisks in the cells of the heatmaps for values where the coefficients are statistically different from 0.

Figure 3A shows considerable variation across inventor-collaborator age dyads in terms of the quantity of spillovers that flow between them. Relative to the reference group, more knowledge spillovers flow between inventor-collaborator dyads in the [30s, 30s], [40s, 30s], [30s, 40s], and [50s, 40s] cells, though only the first of these cells (30s, 30s) is statistically significant at the 95% level. Nonetheless, Figure 3A suggests two inferences: (a) collaborators aged 30-49 generate the most spillovers, and (b) these spillovers tend to flow to inventors that are of similar age or slightly older than their collaborators. Moreover, in contrast to the standard assumption that knowledge is passed down through generations (Jones 2009; 2010), there is no evidence in Figure 3A to suggest that most patenting spillovers flow from older mentors and managers to younger their younger mentees.

Figure 3B shows a similar heatmap but for inventors' high-impact patenting rate. The trends in Figure 3B are similar to those in Figure 3A, except that the [40s-30s] cell is now statistically significant at the 99% level. This result suggests that slightly older inventors (those in their 40s) are strongly reliant on slightly younger inventors (those in their 40s) for creating high-impact patents. Appendix Figure A5 shows similar results with larger effect sizes for the repeat collaborator sample.

Figure 3: Change in Patenting Productivity Following Death of a Collaborator

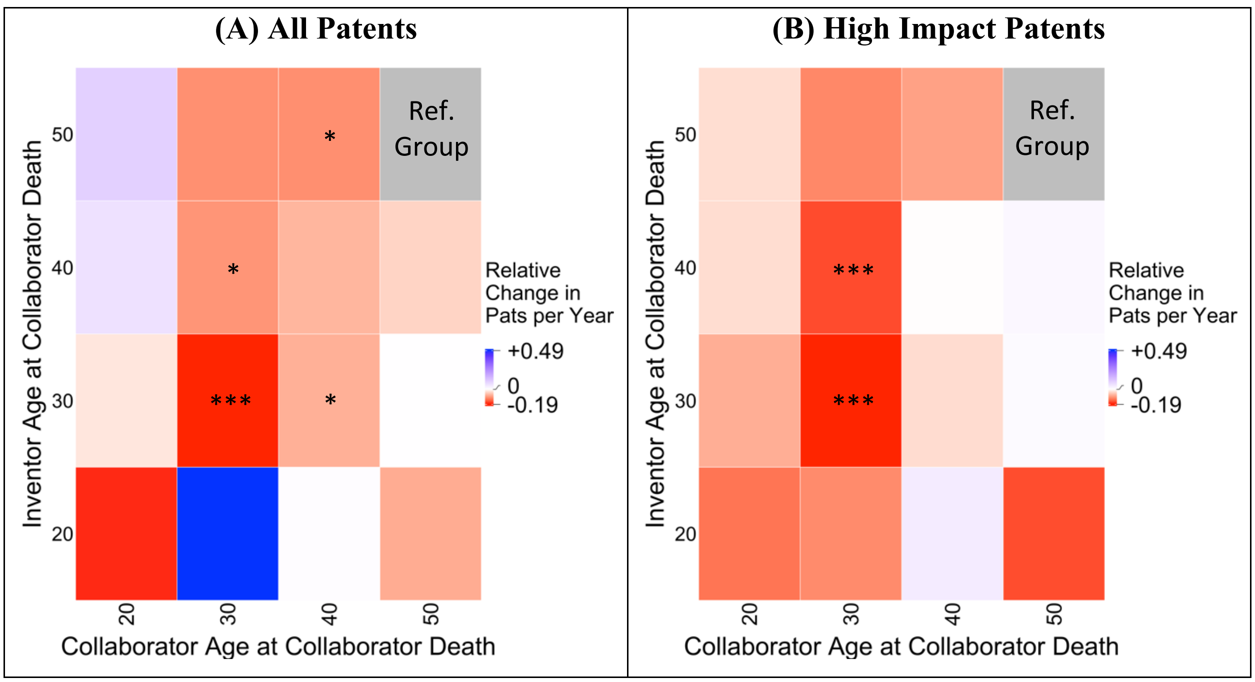

Notes: Heatmaps show the change in patent productivity of surviving inventors following a collaborator's premature death, broken outs by inventor age at time of collaborator death and collaborator age at time of death. Dyads of inventors aged 50-59 at the time of death of collaborators aged 50-59 are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

3.3 Interpretation

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 ideas that inventors combine in each field and in each year. I then 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.

Computing the average age of subclasses in a technology field requires long-run harmonized data on the ideas combined in patents. Moreover, to compute the age of ideas combined in fields, one must observe the first year in which each idea was developed. Therefore, I make use of the historical USPC patent classification file, available through PatentsView, to identify the ideas combined in patents and year they they were first introduced. The historical USPC patent classification file contains patent numbers and fine-grained USPC subclasses for each patent granted between 1836 and 2014. The USPC classification system is no longer updated and has been replaced by the CPC schema as the international standard for classifying patents, but it has the advantage of extending back to 1836, while to my knowledge CPC classification information is not available before 1975.

To compute the rate of advance of inventors' technology fields, I begin by identifying the initial year that each USPC subclass (6 digit) first appeared on a granted patent. Second, I group patents by their primary USPC class (3 digit) and grant year and compute the mean age of the USPC subclasses (6 digit) on patents within each 3-digit class and year group. Third, I compute the median average subclass age across all patents, and group technology fields based on whether their own mean average subclass age is above or below the annual median. Finally, I record whether the majority of an inventors' patents were produced in technological fields with above-average or below-average mean subclass age. Inventors that produce the majority of their patents in USPC classes that have a mean subclass age below the annual median are defined as inventors that primary work in fast-advancing knowledge fields.

Using the resulting measure of the speed of advance in inventors' primary technology fields, I run the regression described by Equation 1 separately for surviving inventors that primarily patent in fast-advancing and slow-advancing knowledge fields. The results are shown in Figure 4.

Figure 4: Change in Patenting Productivity in Slow and Fast-Advancing Knowledge Fields

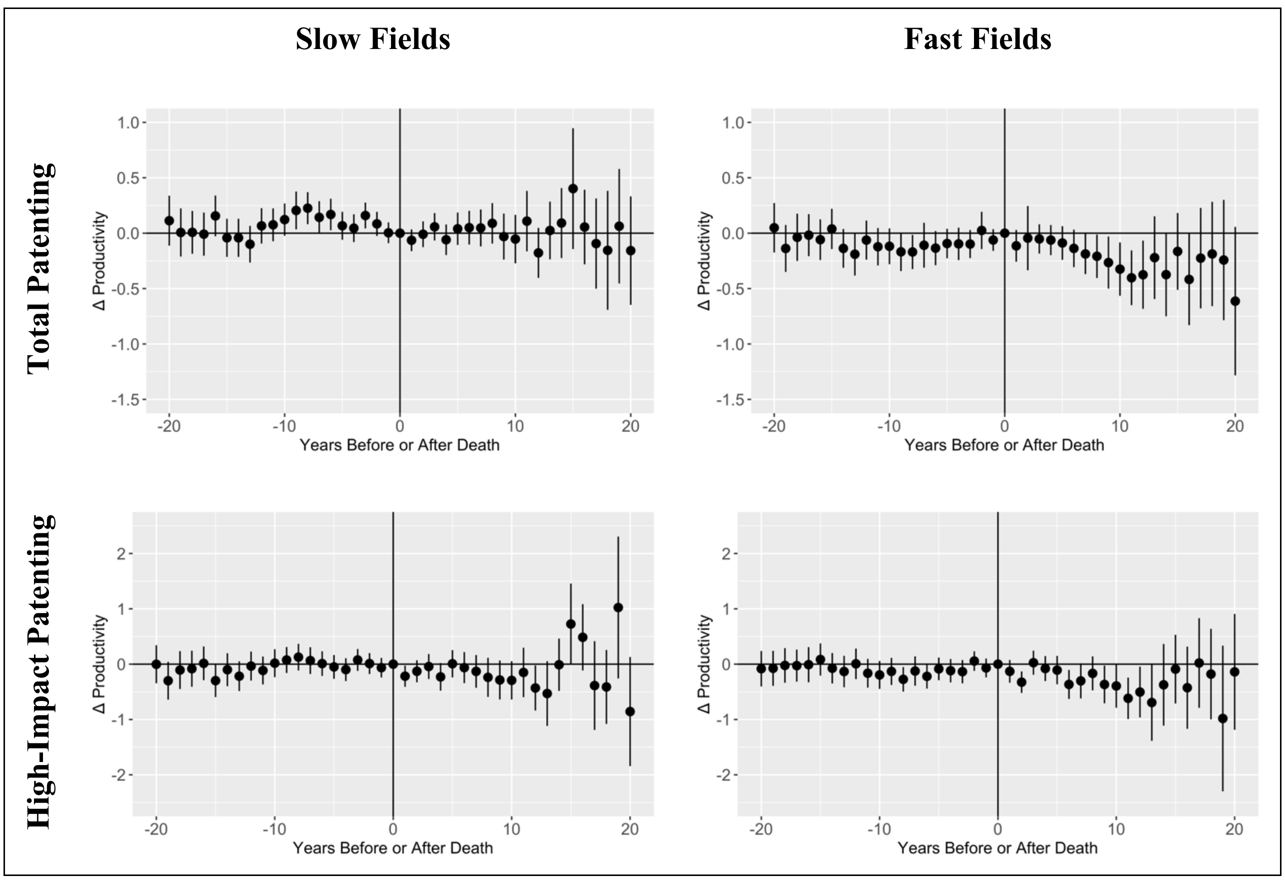

Notes: Figures show effect of losing an early-career collaborator on subsequent patenting rate, relative to losing a mid-career collaborator. Slow and fast-advancing knowledge fields are identified based on the average age of the subclass codes in each inventor's modal technological field, as described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

Figure 4 shows that the negative effect of losing an early-career collaborator on subsequent patenting is more pronounced for inventors that primarily patent in fast-advancing knowledge fields. Inventors in fast-advancing knowledge fields that lose early-career collaborators experience a significant decline in their subsequent overall patenting rate, while inventors in slow-advancing knowledge fields experience no such decline. For high-impact patenting, inventors in both slow and fast-advancing knowledge fields experience a decline after losing a junior collaborator, but the effect size is larger for those in fast-advancing knowledge fields. Appendix Figure A6 shows that the size of these heterogeneous effects are larger when the sample is restricted to repeat collaborators.

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 knowledge fields. To test this proposition directly, I proceed to examine whether inventors proceed to combine older knowledge in their patents following the loss of an early-career collaborator. I use the age of the USPC subclass codes assigned to patents, as described above, to calculate the mean age of the subclasses assigned to a patent *p.* Using this measure of the age of the knowledge embedded each patented invention, I develop a regression model that tests whether inventors combine older ideas following the loss of an early-career collaborator. The model is given by Equation 4:

$$(4) \quad MeanSubclassAge_p = \beta_0 PostCollabDeath_{i,t} + \beta_1 PostCollabDeath_{i,t} * AgeDeceasedCollab_i + \alpha_i + \tau_{USPCClass*t} + \varepsilon_{i,t}$$

In Equation 4, the dependent variable is a continuous variable. Therefore, I estimate Equation 4 with OLS. I show estimates of β_1 using the full dataset in Figure 5A, and for the restricted dataset of repeat collaborators in Figure 5B. The reference group for the variable $AgeDeceasedCollab_i$ is a collaborator that dies between ages 50 and 59. Therefore, the β_1 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 50 and 59. $\tau_{USPCClass*t}$ are indicator variables for the primary USPC 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_1) is interpreted in relation to the mean age of the knowledge produced in the same field and year.

Figure 5: Change in Knowledge Vintage by Collaborator Age at Death

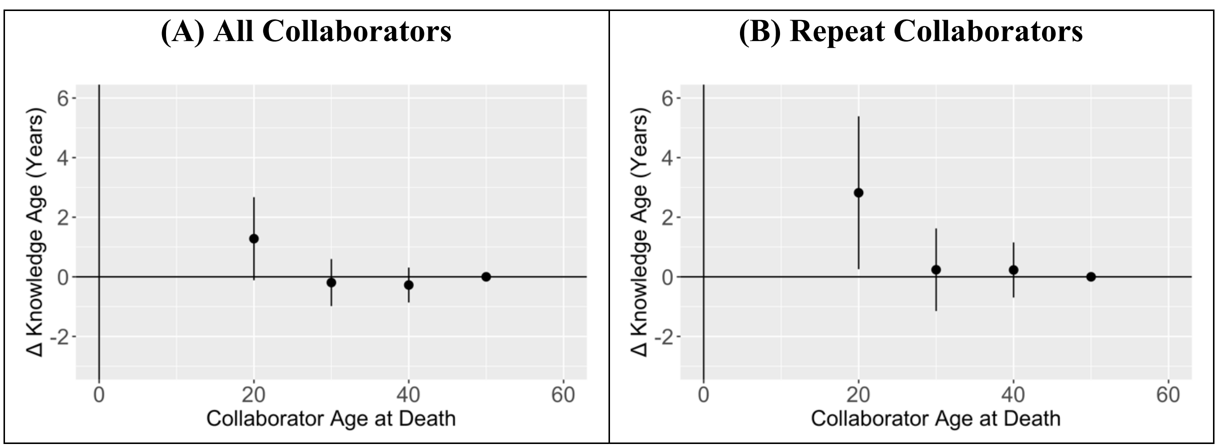

Notes: Figures show the differential effect of losing a collaborator at a specific age at of death relative time (grouped into 10 year bins) to losing a collaborator aged 50-59 at time of death on the vintage of the surviving inventors' primary subclasses on patents.

Figure 5 shows that inventors combine older ideas following the loss of collaborators that die in their 20s, relative to the loss of collaborators that die between ages 50 and 59. On average, inventors combine knowledge that is 1 year older. The effect size is larger for inventors who lose repeat collaborators: the loss of repeat collaborators in their 20s causes their surviving partners to combine knowledge that is 2.8 years older than the reference group of collaborators aged 50-59. Pretend analyses show that these effects emerge only after the death of the collaborator (Appendix Figure A7).

To test whether the effect of the loss of a collaborator depends on the age of the *surviving inventor*, I administer a similar regression model that decomposes the treatment effect by the age group of the survivor. The results for the full dataset and the dataset of repeat collaborators are shown in the heatmap in Figure 6. Red shading indicates larger (positive) coefficients, and asterisks denote statistical significance. The reference group for the regressions are inventors that are aged 50-59 at the time of death of a collaborator who dies between ages 50-59. This reference group is shaded in grey.

Figure 6: Change in Knowledge Vintage Following Death of a Collaborator

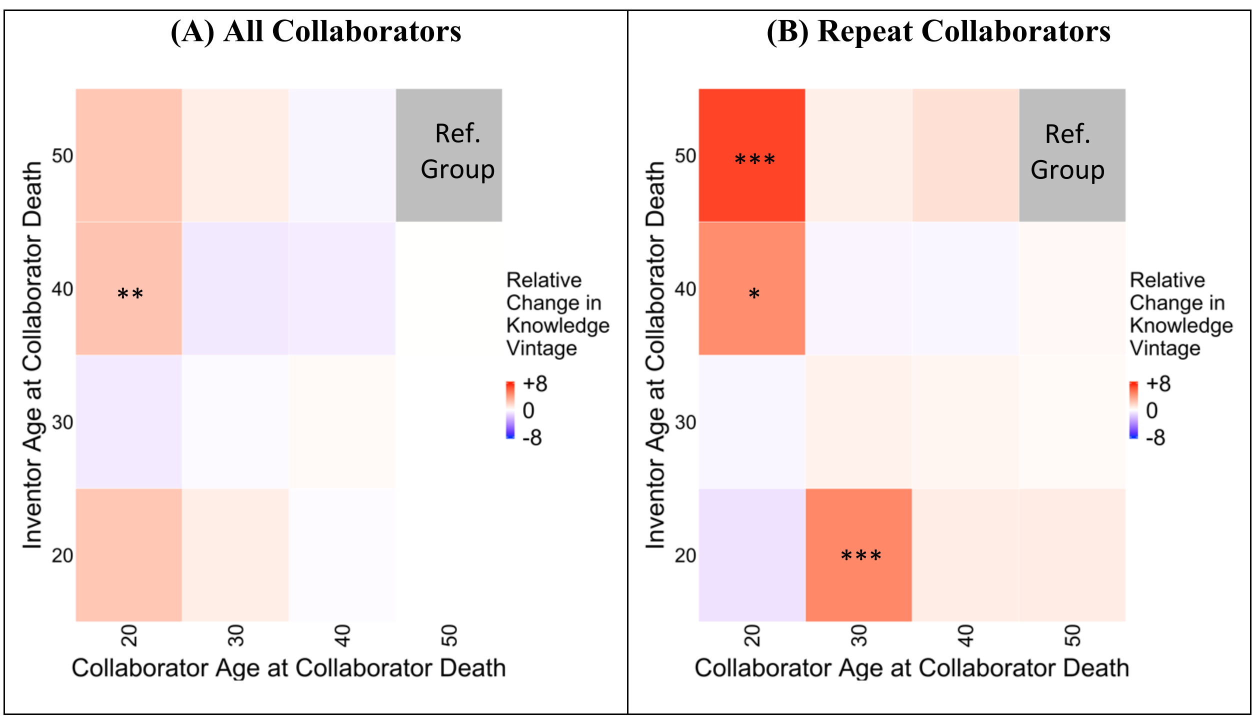

Notes: Heatmaps show the change in knowledge vintage of subclasses on surviving inventors' patents following a repeat collaborator's premature death, broken outs by the inventor age at time of collaborator death and the collaborator age at time of death. Dyads of inventors aged 50-59 at the time of death of collaborators aged 50-59 are the reference group, shaded in grey. Cell values are interpreted as changes in the age of subclasses relative to this reference group. Asterisks indicate statistical significance (*** $p < 0.01$; ** $p < 0.05$; * $p < 0.01$).

Figure 6 shows that the effect of losing a young (aged 20-29) collaborator on the vintage of the ideas they combine is larger when the surviving inventor is old. Across the full dataset (Figure 6A), the loss of a collaborator in their 20s causes the age of the knowledge to increase by 2.45 years for inventors in their 40s and by 2.33 years for inventors in their 50s, though the latter estimate was not statistically significant. The effect sizes are larger when the dataset is restricted to repeat collaborators (Figure 6B). The loss of a repeat collaborator in their 20s causes the age of the knowledge to increase by 4.65 years for inventors in their 40s and by 7.02 years for inventors in their 50s. The former estimate is statistically significant at the 90% level while the latter estimate is significant at the 99% level.

Are these effect sizes large enough to cause inventors to fall behind the state of the art in their knowledge fields? To investigate this question, in Figure 7 I plot the distribution of patents by the mean age of their subclasses, demeaned the mean age of the subclasses on each patent by its USPC class*year average. Thus, the histogram in Figure 7 gives the distribution of mean subclass ages on patents *within* knowledge fields. To assess whether the effect size is sufficient to cause inventors to fall behind the state-of-the-art in their fields, I overlay two vertical lines on the histogram. The first vertical line (dotted blue) shows the effect of losing a collaborator aged 20-29 on the knowledge vintage of the median inventor aged 50-59 (as in Figure 7A). This coefficient has a value of 2.33 years. The second vertical line (solid blue) is the effect of losing a repeat collaborator aged 20-29 on the median inventor aged 50-59 (as in Figure 7B). This coefficient has a value of 7.02.

Figure 7: Loss of Young Collaborators Causes Older Inventors to Fall Behind their Fields

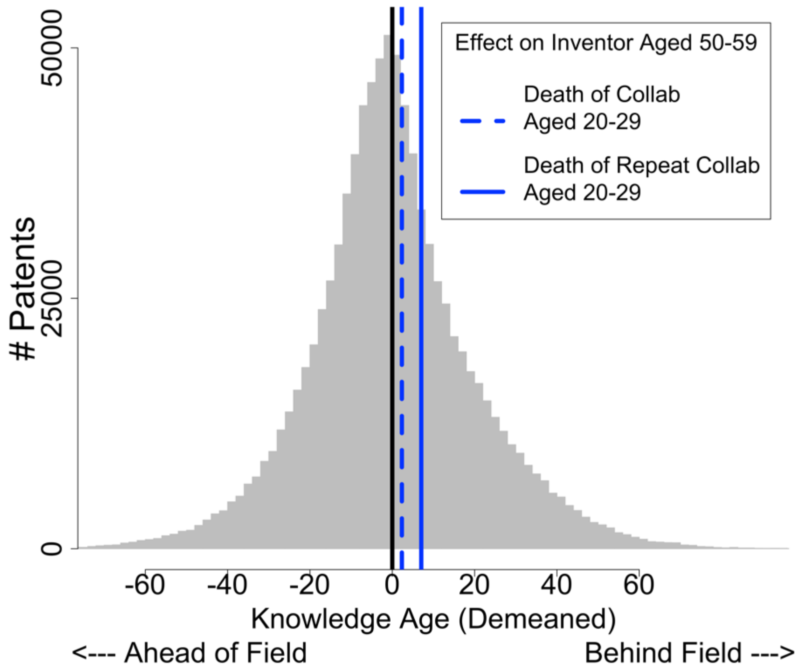

Notes: Distribution of knowledge age is demeaned by the USPC class*year of patents.

Figure 7 shows that the distribution of subclass ages within class*year pairs is normally distributed with a (demeaned) median value of 0 years and a standard deviation of 20.3 years. While the loss of a 20-29 year old collaborator in general does not shift the age of a 50-59 year old inventors' knowledge very much, relative to the breadth of the distribution (dotted blue line), the loss of a 20-29 year old *repeat* collaborator shifts the age of the knowledge of the surviving 50-59 year old partners substantially (solid blue line). Moreover, an inventor aged 50-59 who loses a repeat collaborator aged 20-29 to a premature death moves from the median of the distribution to the 70% percentile.

Discussion

This paper studied the relationship between the age of inventors and the spillovers that they generate, using the premature deaths as an exogenous shock to inventors' collaborative networks.

Early-career collaborators were shown to generate more spillovers than mid-career collaborators, with spillovers peaking between the ages of 30 and 49. The differential effect of the death of an early-career collaborator was also greatest for inventors that work in fast-advancing knowledge fields, and the loss of young collaborators was shown to cut off their surviving partners from novel areas of technological knowledge. These results indicate that early-career collaborators are particularly important for inventors to learn novel ideas at the technological frontier.

While collaborators aged 30-39 were found to contribute the most to the patenting productivity of their partners, collaborators aged 20-29 were found to contribute most to their partners' ability to learn new, frontier technological ideas. The novelty content of collaborators' spillovers may reach their maximum at a younger age than their overall patenting productivity spillovers because learning novel ideas and deploying those ideas into useful, patentable technologies are non-identical processes. Novel ideas are only deployed into functioning technologies after they have been debugged and integrated into larger technological systems. This process is time-consuming and uncertain (Adler and Clark 1991). Inventors that die in their 20s may not have lived long enough to refine their ideas.

Together, these findings suggest that technological knowledge advance is a process that has both cumulative and disruptive properties (Romer 1990; Jones 2009; Aghion et al. 2023). The cumulative rendering of knowledge advance was supported by the finding that very young collaborators (aged 20-29) generate relatively few patenting spillovers for their partners, presumably because they lack the necessary experience to absorb the large stock of knowledge produced by prior generations of inventors and scientists (Figure 3). Support for the disruptive vision of technological advance was generated by three results. First, collaborators aged 30-49 generate more knowledge spillovers than do collaborators aged 50-59 (Figure 2). Second, the differential number of spillovers produced by early-career collaborators is greatest for inventors that work in fast-advancing knowledge fields (Figure 4). Three, very young collaborators transmit new, cutting-edge ideas to their partners, which allows their partners continuously learn novel ideas as they grow older (Figures 5-7). By demonstrating that technological knowledge is not only passed down through generations, but also emerges from the interactions between early-career inventors, these results suggest that the education burden for sustained innovation across generations is not as onerous as is commonly believed (Jones 2009; 2010; Bloom et al. 2020), and they suggest the costs of knowledge obsolescence borne on the economy are high (Aghion and Howitt 1992; Aghion et al. 2023).

References

Adler, Paul, and Kim Clark. 1991. "Behind the Learning Curve: A Sketch of the Learning Process." *Management Science* 37 (3): 267–81.

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.

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.

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.

Cui, Haochuan, Lingfei Wu, and James A. Evans. 2022. "Aging Scientists and Slowed Advance."

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

Foster, Jacob G., Andrey Rzhetsky, and James A. Evans. 2015. "Tradition and Innovation in Scientists' Research Strategies." *American Sociological Review* 80 (5): 875–908. https://doi.org/10.1177/0003122415601618.

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*.

Katila, Ritta. 2002. "New Product Search over Time: Past Ideas in Their Prime?" *Academy of Management Journal* 45 (5): 995–1010.

Kuhn, Thomas. 1962. *The Structure of Scientific Revolutions*. University of Chicago Press.

Liu, Lu, Nima Dehmamy, Jillian Chown, C. Lee Giles, and Dashun Wang. 2021. "Understanding the Onset of Hot Streaks across Artistic, Cultural, and Scientific Careers." *Nature Communications 2021 12:1* 12 (1): 1–10. https://doi.org/10.1038/s41467-021-25477-8.

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.

Appendix

Figure A1: Change in Patenting Productivity by Age of Deceased Collaborator, Deaths Under Age 55

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

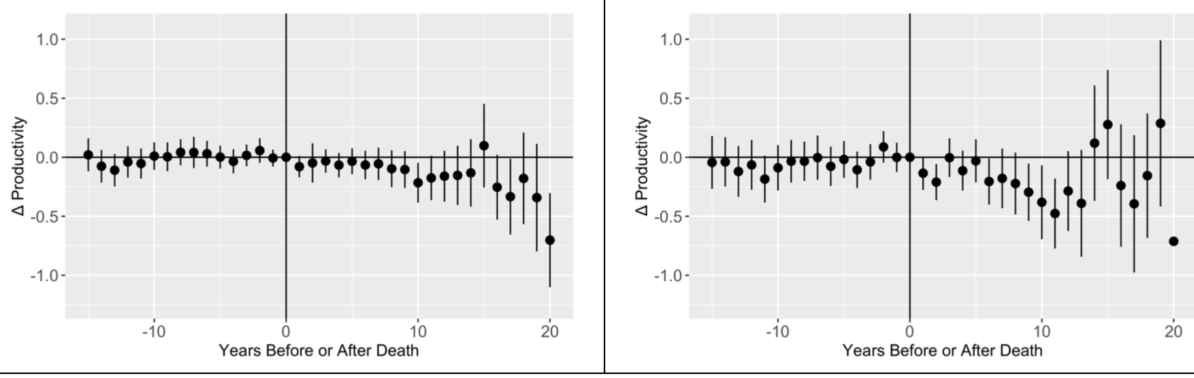

Notes: Figures shows effect of losing an early-career collaborator (age 20-44 at time of death) relative to losing a mid-career collaborator (aged 45-54 at death) on patenting productivity. All models include focal inventor and year*technology class fixed effects. Standard errors are clustered at the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

Figure A2: Change in Patenting Productivity by Age of Deceased Collaborator, Max 3 Years Between Collaboration and Death

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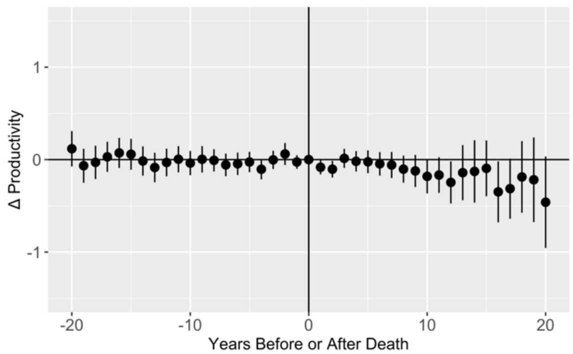	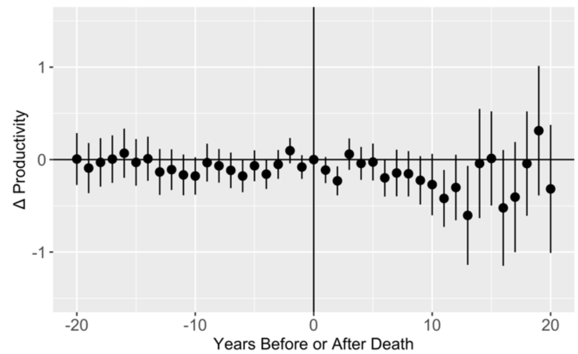

Notes: Figures shows effect of losing an early-career collaborator (age 20-44 at time of death) relative to losing a mid-career collaborator (aged 45-59 at death) on patenting productivity. A focal inventor's deceased collaborators are defined as those that a focal inventor co-invented a patent with 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.

Figure A3: Change in Patenting Productivity by Age of Deceased Collaborator Following Death of a Repeat Collaborator

(A) Effect on Overall Patenting

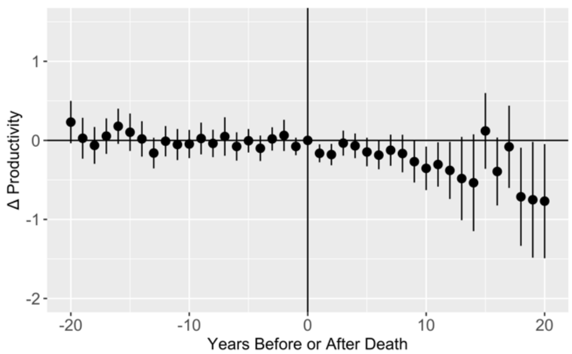

(B) Effect on High-Impact Patenting

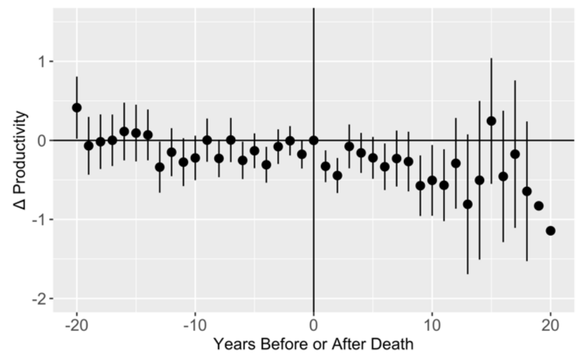

Notes: Figures shows effect of losing an early-career collaborator (age 20-44 at time of death) relative to losing a mid-career collaborator (aged 45-59 at death) on patenting productivity. All models include focal inventor and year*technology class fixed effects. Standard errors are clustered at the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

Figure A4: Change in Patenting Productivity by Collaborator Age at Death Following Death of a Repeat Collaborator

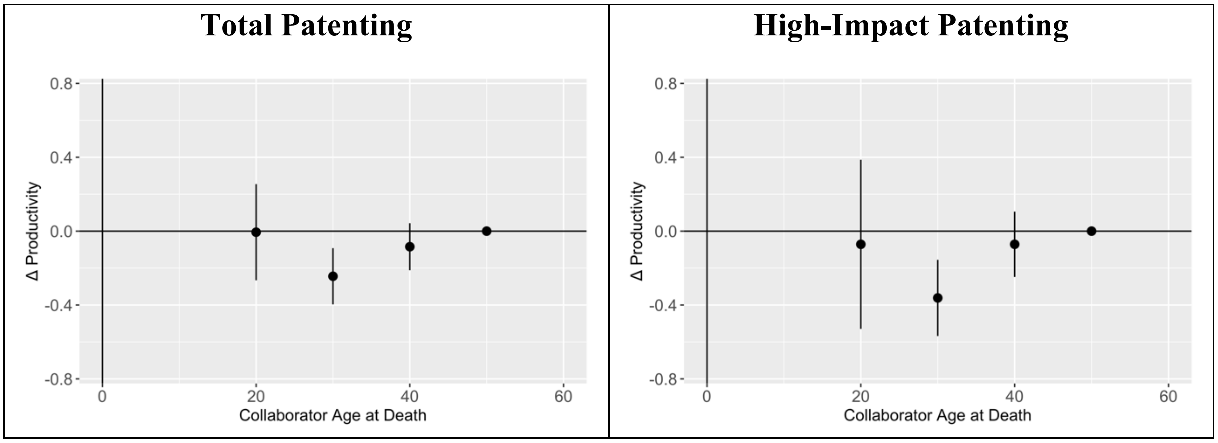

Notes: Figures show the differential effect of losing a collaborator at a specific age at of death relative time (grouped into 10 year bins) to losing a collaborator aged 50-59 at time of death.

Figure A5: Change in Patenting Productivity Following Death of a Repeat Collaborator

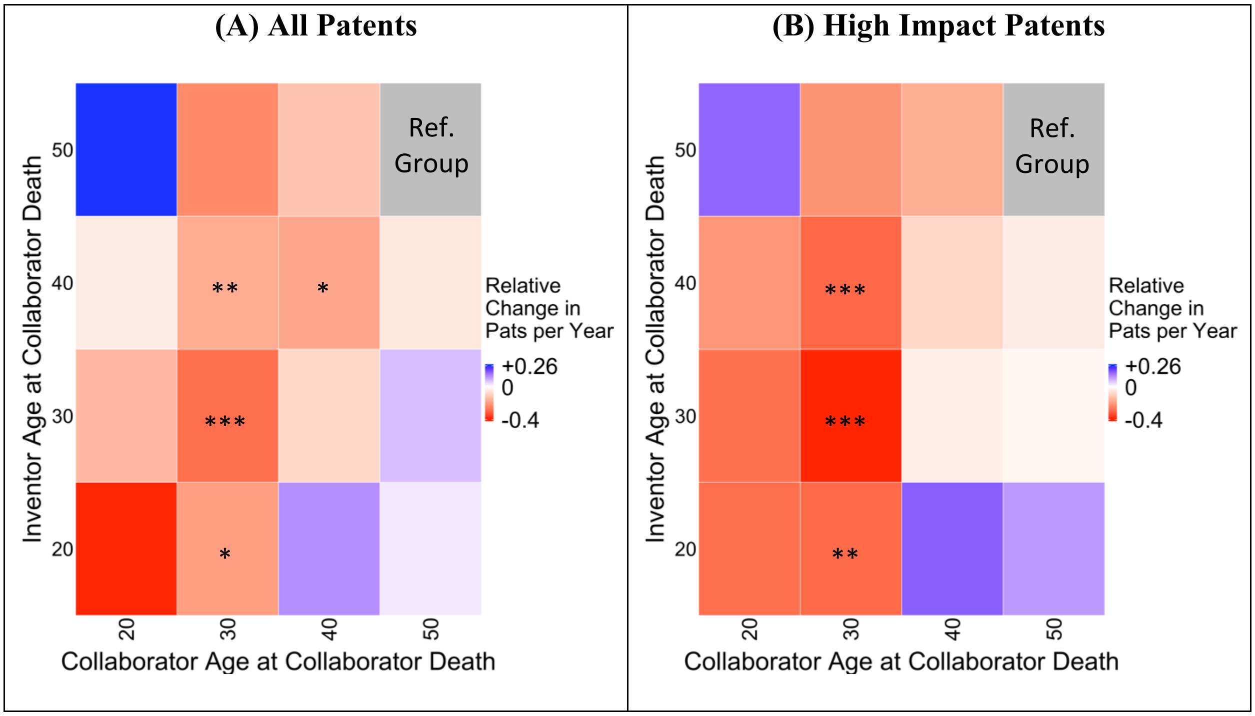

Notes: Heatmaps show the change in patent productivity of surviving inventors following a repeat collaborator's premature death, broken outs by the inventor age at time of collaborator death and the collaborator age at time of death. Dyads of inventors aged 50-59 at the time of death of collaborators aged 50-59 are the reference group, shaded in grey. Cell values are interpreted as changes in patenting productivity relative to this reference group. Asterisks indicate statistical significance (*** p < 0.01; ** p < 0.05; * p < 0.01).

Figure A6: Change in Patenting Productivity in Slow and Fast-Advancing Knowledge Fields Following Death of a Repeat Collaborator

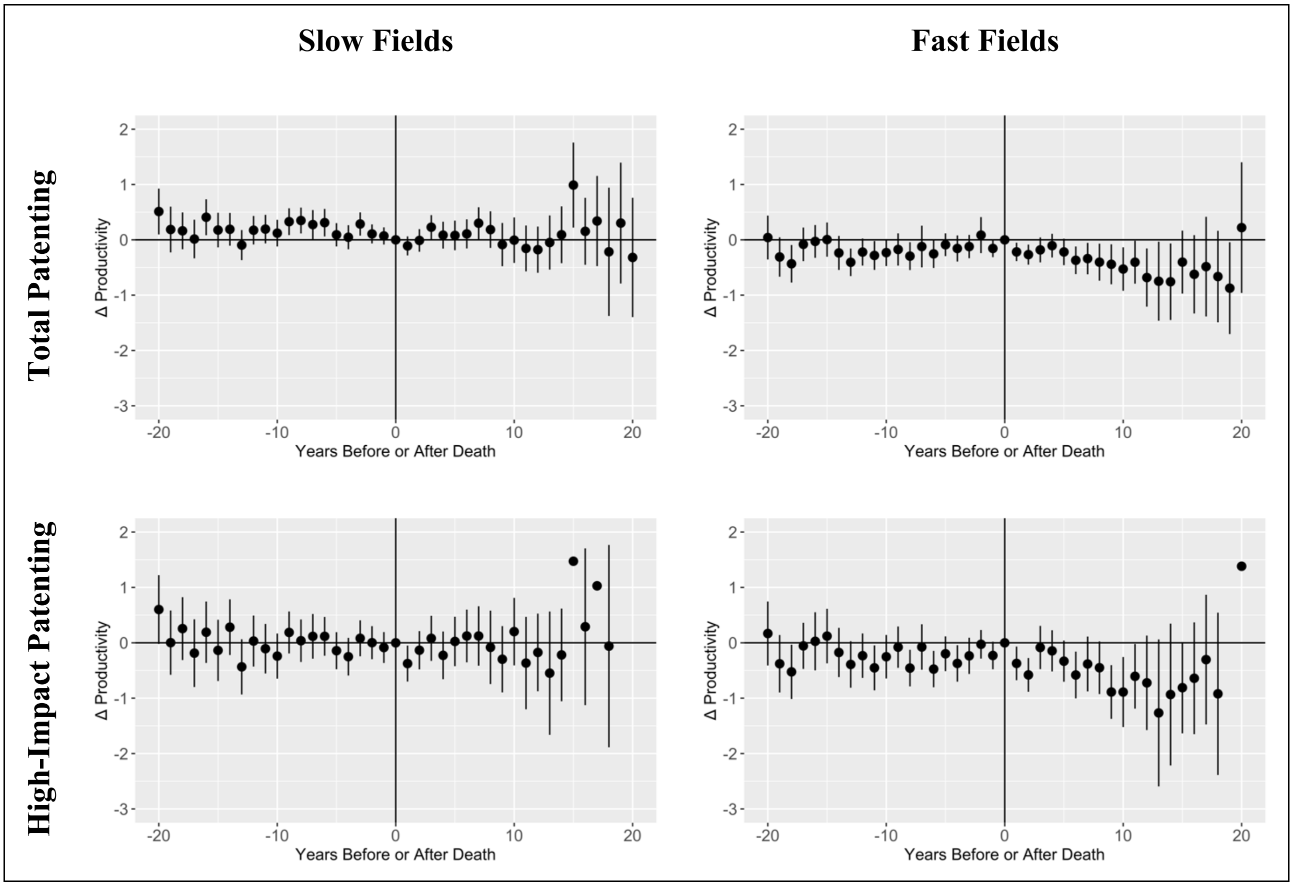

Notes: Figures show effect of losing an early-career collaborator on subsequent patenting rate, relative to losing a mid-career collaborator. Slow and fast-advancing knowledge fields are identified based on the average age of the subclass codes in each inventor's modal technological field, as described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

Figure A7: Pre-trend Analysis of Change in Knowledge Age Following Death of Collaborator Aged 20-29 Relative to Death of Collaborator Aged 50-59

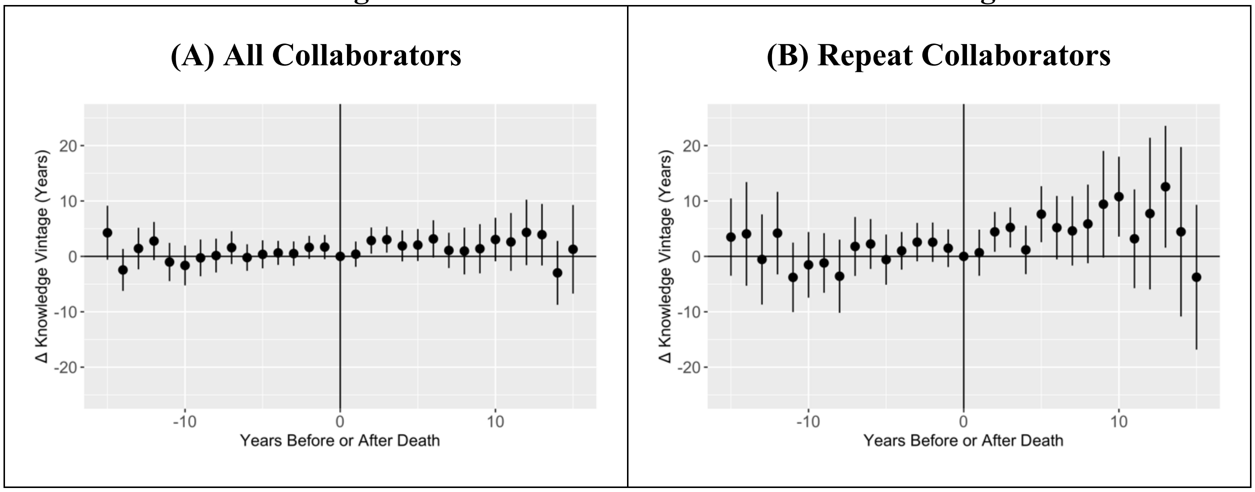

Notes: Figures shows effect of losing a very young collaborator (age 20-29 at time of death) relative to losing an older collaborator (aged 50-59 at death) on the vintage of the knowledge combined in the surviving partner's patents. All models include focal inventor and year*technology class fixed effects. Standard errors are clustered at the deceased collaborator.