Knowledge Spillovers Within and Across Age Groups of Inventors

Christopher Esposito
Anderson School of Management
UCLA
christopher.esposito@anderson.ucla.edu

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 – productivity differences, collaborator interdependence, and knowledge obsolescence – and find the strongest support for the knowledge obsolescence channel.

Acknowledgement: I thank Mike Andrews, Eamon Duede, James Evans, Lee Fleming, Ashvin Gandhi, Charu Gupta, Jennifer Kao, Donghyun Kang, Jeffrey Lockhart, Sarada, Olav Sorenson, Frank van der Wouden, and seminar participants at Berkeley Haas, UCLA Anderson, AOM Boston, the Mansueto Institute Colloquium Series, and ICSSI for helpful comments. All remaining errors are my own.

1. Introduction

Whom do inventors learn from? Inventors primarily compete on basis of their knowledge assets, making the individuals from which they learn fundamental inputs to their own ability. A growing literature identifies the attributes that lead certain groups of individuals to be particularly valuable sources of knowledge. This research has shown the significance of family members, community members, and teachers for providing early knowledge and encouragement to children and students that might pursue innovative careers (Delgado and Murray 2023; Bell et al. 2019; Azoulay, Liu, and Stuart 2017), the centrality of collaborators with unusually high levels of helpfulness, scientific preeminence, network brokerage, and commercialization experience for introducing academic scientists to new ideas and helping them translate those ideas into functioning technologies (Oettl 2012; Azoulay, Zivin, and Wang 2010; Mohnen 2022; Marx and Hsu 2022), and the influence of local inventors, immigrant collaborators, and tightly-connected teams on the learning conducted by other inventors (Bernstein et al. 2022; Balsmeier, Fleming, and Lück 2023; Jaravel, Petkova, and Bell 2018). However, a critical dimension that this literature has not yet examined is whether inventors learn more from their younger or older collaborators. This question is consequential for three reasons.

First, theory is ambiguous as to whether inventors learn more from their younger or older collaborators. Inventors absorb knowledge over time, which could make their older collaborators better teachers and mentors (Jones 2010). However, knowledge obsolescence also depreciates the human capital of knowledge-based workers, which could lead inventors' older collaborators to have smaller stocks of *effectual* knowledge, if they are specialized in older vintages of ideas (Aghion et al. 2023; Esposito and Wouden 2022). Therefore, the relationship between collaborator age and knowledge dissemination depends, in part, on whether collaborators accumulate knowledge faster through experience than they lose knowledge through obsolescence (Kredler 2014). Because theory alone cannot tell us which rate exceeds the other, empirical research is needed to fill the gap.

Second, the question of whether inventors learn more from their junior or senior collaborators has practical implications for organizations. Should organizations partner younger inventors with older collaborators, with the hope that knowledge will flow downward across generations? Or, if teams contain both junior and more senior inventors, could knowledge possibly flow *upwards*, to the older partners through reverse-intergenerational spillovers? Alternatively, could learning be improved by partnering inventors with similarly-aged collaborators, who might have greater capacity to absorb one-another's ideas? While previous research has studied the relationship between the age demographics of team members for their *static* innovation outcomes, such as whether mixed-age teams produce higher-quality inventions (Kaltenberg, Jaffe, and Lachman 2023), there is a paucity of research on how the different age distributions of teams relate to the dynamic learning of the team members.

Third, uncovering the relationship between collaborator age and knowledge spillovers would provide an informative datapoint on the extent to which technological change is a cumulative or disruptive process. Since at least Schumpeter (1962), innovation theorists, economists, and management scientists have sought to discern the extent to which new ideas build on existing ideas or drive existing ideas into obsolescence (Caballero and Jaffe 1993; Caballero and Hammour 1996; Ma 2021; Funk and Owen-Smith 2017; Park, Leahey, and Funk 2023; Esposito and Wouden 2022). Because the relationship between age and spillovers depends on the depreciation rate of existing knowledge, uncovering the age-spillovers relationship would provide crucial evidence on the extent to which technological change is generally a cumulative or disruptive process.

In this study, I measure the quantity of knowledge that inventors learn from their younger and older collaborators, by calculating the change in inventors' patenting output after losing collaborators to premature deaths. Premature deaths end the opportunity for surviving inventors to interact with their deceased collaborators, providing quasi-exogenous variation in inventors' learning opportunities. To identify the quantity of knowledge transmitted by collaborators of different ages, I contrast the change in patenting output of surviving inventors who lose junior collaborators (below age 45 at time of death) with those who lose mid-career collaborators (age 45 to 59). I omit late-career collaborators from the analysis (age 60 and above), because their deaths could be more anticipated and thus endogenous.

Pairing this identification strategy with USPTO data on patents and recently-collected records on inventors' birth and death years (Kaltenberg, Jaffe, and Lachman 2021), I find that junior collaborators generate more knowledge spillovers than do mid-career collaborators. Specifically, focal inventors who lose junior collaborators to premature deaths proceed to produce 8% fewer patents compared to those who lose mid-career collaborators. The differential treatment effect becomes statistically significant six years post-collaborator death and endures for twelve years. In addition, the effect size of the differential is larger for inventor-collaborator pairs with multiple collaborative patents, indicating that inventors with stronger relationships with junior collaborator are more affected by the losses. Finally, the effect is larger when the survivor's production of highly-cited patents is used as the outcome variable, suggesting that junior collaborators transmit knowledge not just of greater quantity, but also of greater quality.

Following this main analysis, I conduct three follow-on analyses to better understand how knowledge moves across different age groups of inventors and the mechanisms involved. In the first of these analyses, I identify the precise age at which collaborators generate the most spillovers. The greatest volume of spillovers are generated by moderately young collaborators (aged 30-44 at time of death), while fewer spillovers are generated by very young collaborators (aged 20-29 at time of death) and relatively old collaborators (aged 45-59 at time of death). These findings are consistent with the interpretation that knowledge growth entails both cumulative *and* disruptive processes, as very young collaborators may lack

sufficient experience to knowledge to generate extensive spillovers (Jones 2010), while older collaborators could suffer from knowledge obsolescence (Aghion et al. 2023; Esposito and Wouden 2022).

In the second of these analyses, I explore the interaction between survivor age, deceased collaborator age, and spillovers by decomposing the treatment effect by the collaborator's age *and* the survivor's age at the time of the collaborator's death. I find that the greatest volume of spillovers flows from deceased collaborators in their 30s to surviving inventors in their 30s. In addition, more spillovers flow from younger inventors to older inventors than from older inventors to their younger partners. Most notably, these findings demonstrate that downward spillovers across generations of collaborators are not particularly prevalent, which contrasts with the common viewpoint that younger knowledge workers mostly learn from their elders.

Finally, I present three pieces of evidence suggesting that the smaller volume of spillovers generated by older collaborators could be driven by knowledge obsolescence. First, I show that two competing hypotheses (differences in the productivity and differences in team interdependence) do not explain the main finding of the relatively larger spillovers generated by junior collaborators. Second, I show that the spillover differential associated with junior collaborators is larger in fast-advancing knowledge fields, in which new ideas are introduced quickly and old ideas quickly fall out of use. These fast-advancing fields plausibly have higher rates of knowledge obsolescence (Esposito and Wouden 2022). Third, I demonstrate that young collaborators are unique sources of recently-vintaged ideas for their partners: surviving inventors who lose junior collaborators to premature deaths proceed to cite older patents than they would have if they had lost older collaborators. Together, these findings suggest that young collaborators are beneficial for learning because they know recently-vintaged ideas, because young collaborators transmit these recently-vintaged ideas to their partners, and because inventors need to know recently-vintaged ideas if they hope to innovate in in fast-advancing knowledge environments.

In the remaining sections of this paper, I introduce the data sources, describe the methods, present the results, and discuss the implications of the findings. The appendix also contains two key materials. Appendix A presents a replication of the main analysis using a matched sample, wherein young deceased collaborators are matched to older deceased collaborators with similar patenting profiles. The matched sample generates similar results to the unmatched sample presented in the main text, affirming that the main results are not artifacts of an endogenous selection process. (The main text presents results from an unmatched sample in order to use a consistent sample across all analyses, which would be impossible using a matched sample). Appendix B presents a replication of the main analysis using a stacked regression, to omit prohibited contrasts from the staggered treatment estimation (Baker, Larcker, and Wang 2022; Strezhnev 2023). Again, the results from the stacked regression are similar to the two-way fixed effect estimates presented in the main text.

2. Data and Methods

I collect 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 contain patent ID numbers, current CPC classification information codes at the class level, for which there are 123 classes, front-page patent citations, disambiguated inventor IDs, and patent application years. 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 patents accumulate citations 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 collect birth and death records for the general population, and by matching these records 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). To my knowledge, these data have not been used to analyze how the spillovers generated by collaborators varies with collaborator age.

Because the Kaltenberg, Jaffe, and Lachman (2023) data contain birth and death records that have varying likelihoods of being accurate, I clean the data using the following steps: first, I omit all birth records with accuracy scores of 0 (as reported by Kaltenberg, Jaffe, and Lachman (2023)); 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, and in the most consequential step, 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. This final step ensures that I only retain death records that can be matched with high confidence to a birth record.

To link inventors to their deceased collaborators, I first identify each inventor that died between 20 and 59 years old, indexed by c, and the focal inventors, i, that co-invented a patent with each of the deceased collaborators 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. In the appendix, Figure D2 shows that the general results are robust to using a 3-year cutoff to define collaborators.

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

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 Table 1. 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.

Table 1: Summary Statistics of Focal Inventors and Deceased Collaborators

		Quantile		
	Variable	**25%**	**50%**	**75%**
Focal (Surviving) Inventors	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
Deceased Collaborators	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.

3. Empirical Analysis

3.1 Main Analysis

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 in terms of their forward citations. 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 2. 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 generate even more spillovers for high-impact patenting than 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 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
$PostCollabDeath_{i,t}$	-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 Observations	296,901	209,204	96,604	74,041
Inventor Observations	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. The model is estimated with a quasi-Poisson. Standard errors are clustered at the deceased collaborator. *** p < 0.01; ** p < 0.05; * p < 0.01.

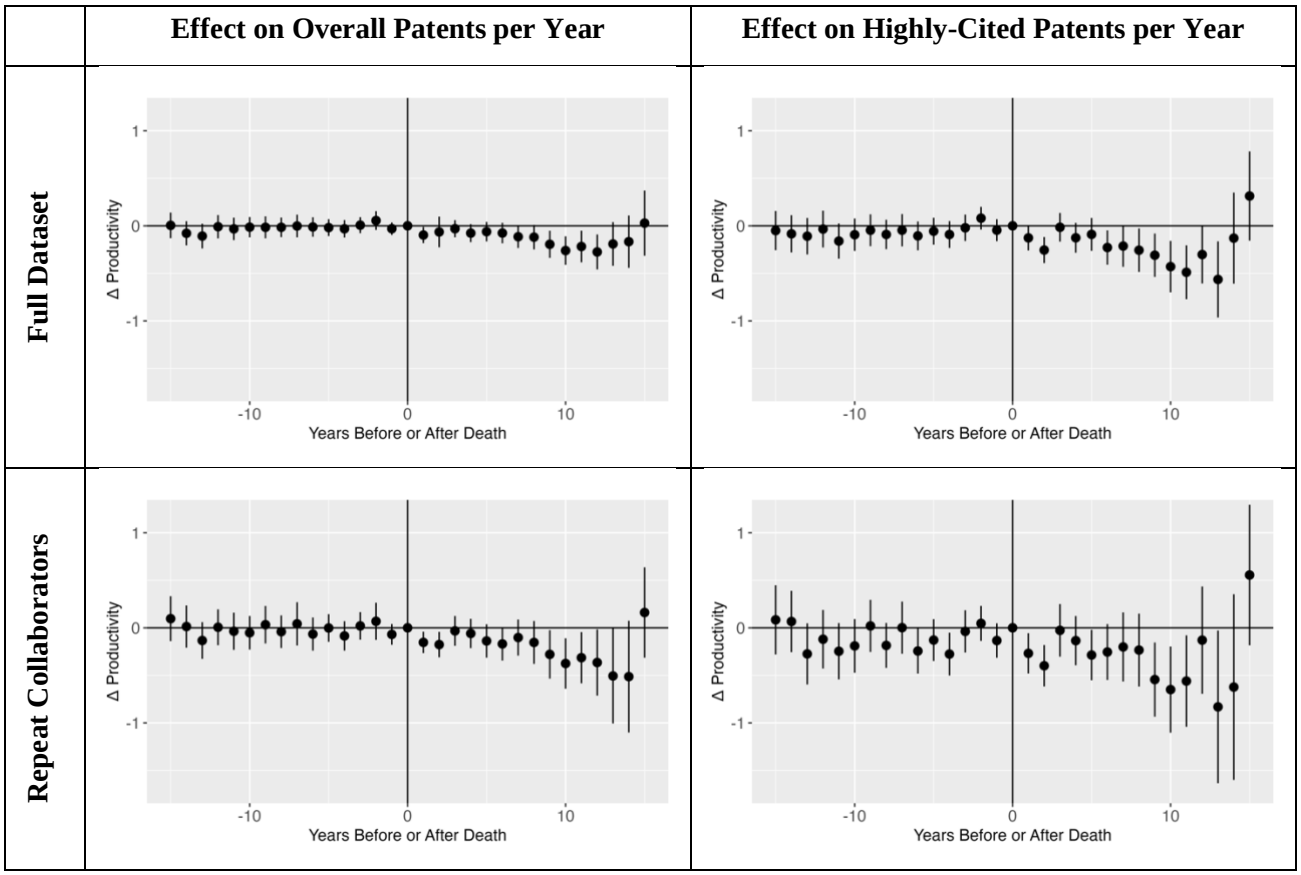
Next, I explore the dynamics of the treatment effect by estimating the regression model described by Equation 2, with a set of 15 forward and 15 backward lags to the treatment event. Because these lag variables are perfectly colinear with the inventor and year fixed effects, the base terms of these lags (without the interaction with $JuniorDeceasedCollab_i$) drop out of the regression (Hall, Mairesse, and Turner 2005).

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

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 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 B, I show that the results are similar when a stacked regression is used. Therefore, the results in Figure 1 are not driven by the staggered design.

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) using a Quasi-Poisson model. All models include focal inventor and year*technology class fixed effects. Standard errors are clustered at the deceased collaborator.

The time dynamics of the treatment effects in Figure 1 provide some suggestions as to mechanisms. In particular, the largest impact is not felt until 10 years after the loss of a junior collaborator, which may indicate 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.

3.2 Effects by Collaborator Age

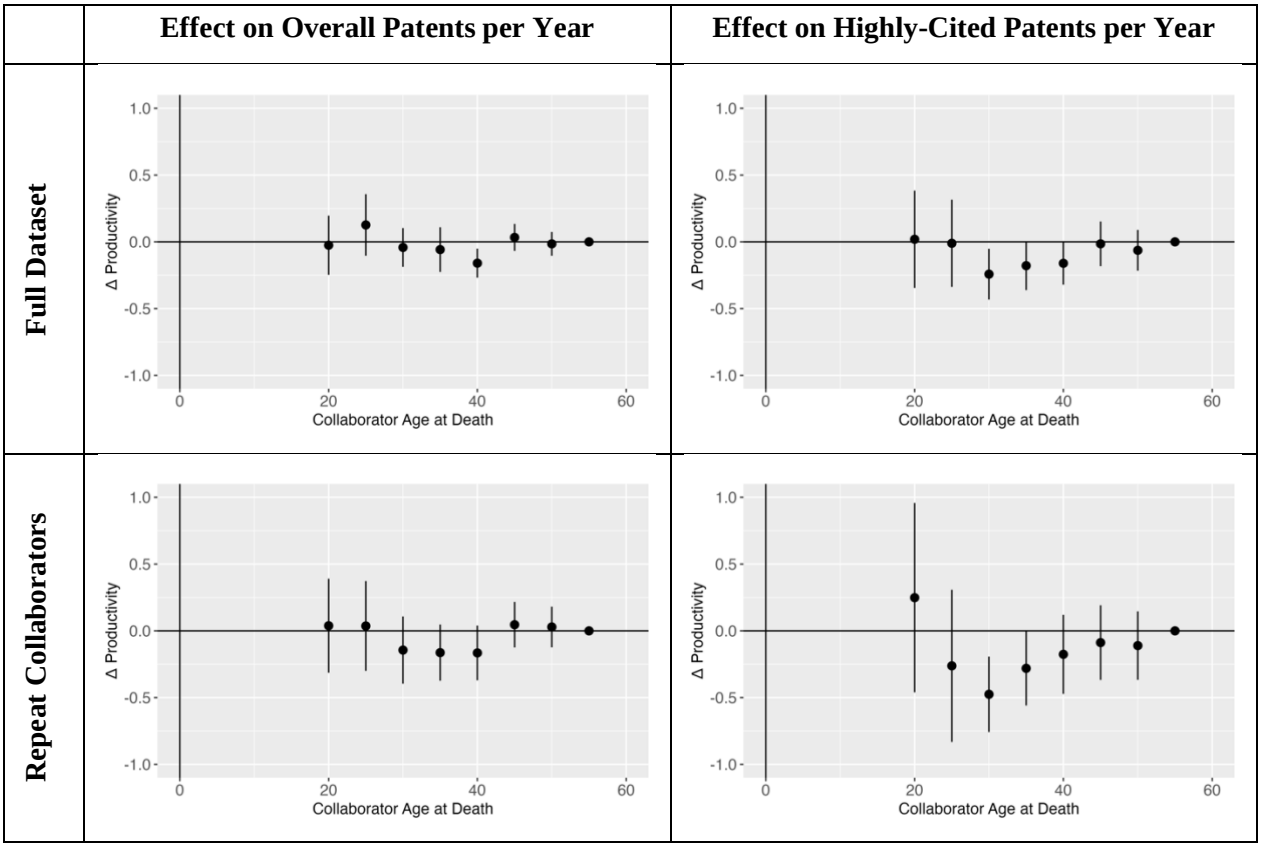
To identify the specific age at time of death of a collaborator that has the largest effect on the subsequent patenting of a focal inventor, inventor, I estimate a regression model described by Equation 3:

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

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 than 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 is consistent with the interpretation that both experience (Jones 2010) and knowledge obsolescence (Aghion et al. 2023) play important roles in determining the quantity and quality of the spillovers that collaborators generate, with the experience mechanism prevailing for very young collaborators and the obsolescence mechanism prevailing for older collaborators.

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. The model is estimated with a Quasi-Poisson and standard errors are clustered at the deceased collaborator.

3.3 Effects by Inventor and 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:

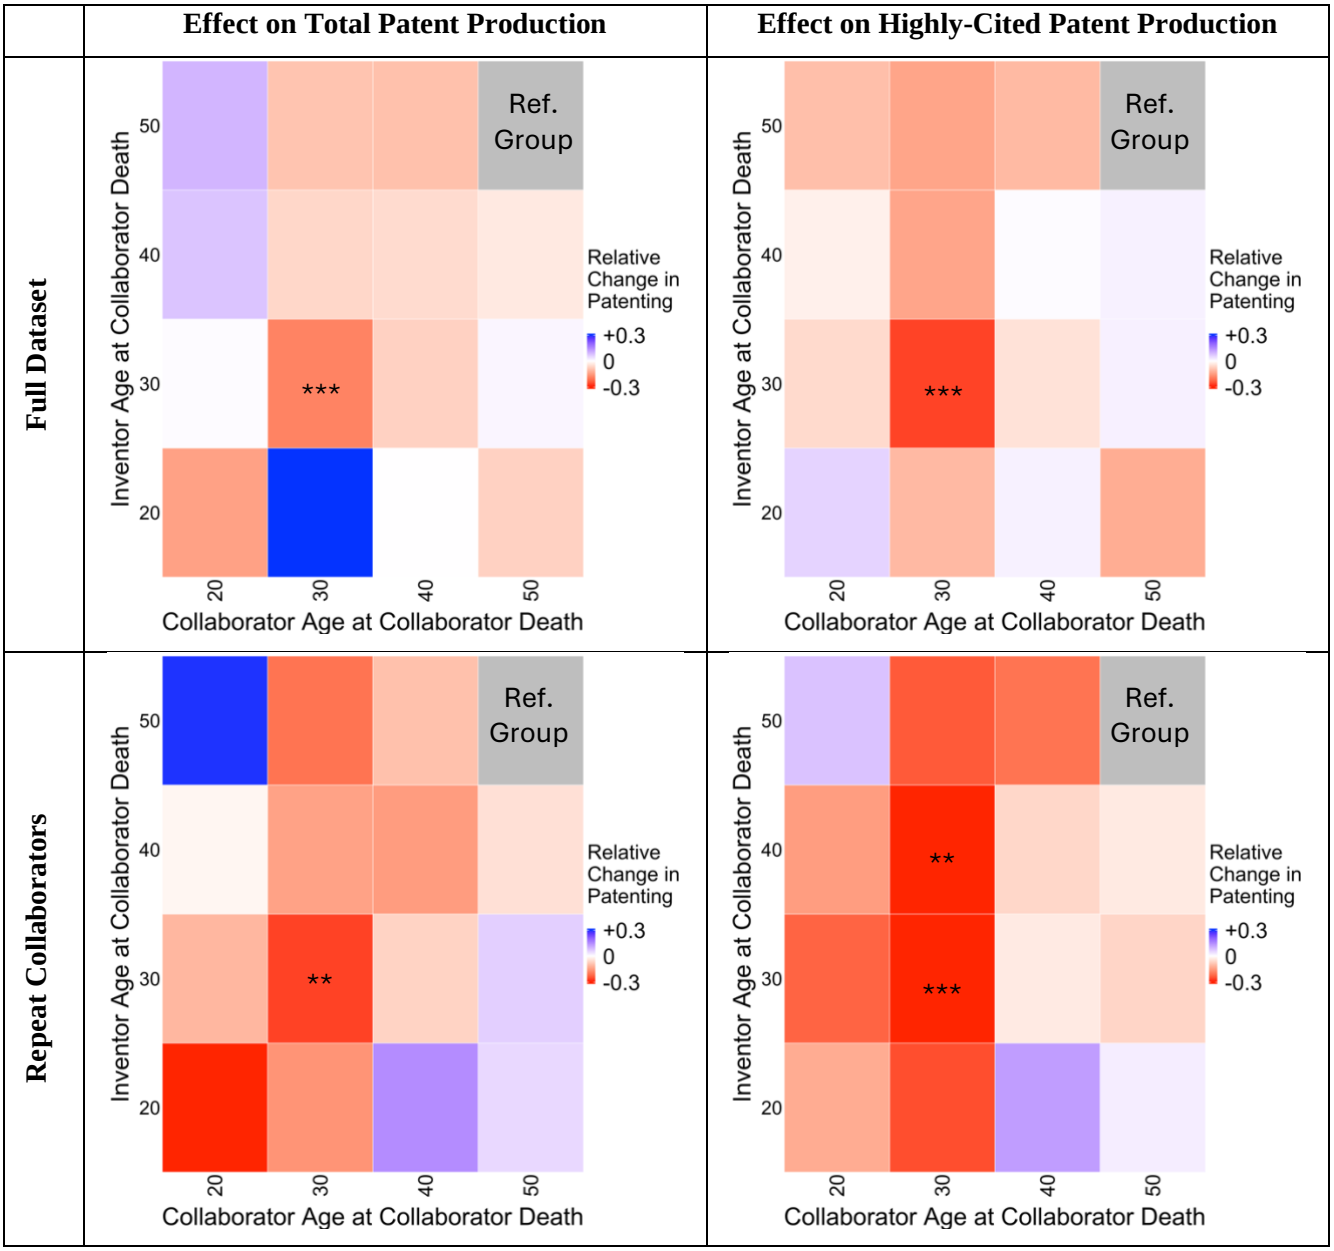
$$(4) \quad PatProd_{it} = \exp\left[\beta_0 PostCollabDeath_{it} + \sum_{m=20}^{59}\sum_{n=20}^{59} \beta_{mn} PostCollabDeath_{it} * AgeCollabAtDeath_i^m * AgeInventorAtDeath_i^n + \alpha_i + \tau_{modalfield*t} + \varepsilon_{it}\right]$$

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 of their collaborator, when the collaborator also dies in the 50-59 age group. I estimate Equation 4 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 the one associated with inventors in their 30s who lose collaborators in their 30s. There are two main reasons that could explain this result. First, older inventors may lack the absorptive capacity to learn the ideas known by their younger collaborators (Cohen and Levinthal 1990). Second, older inventors may continue to receive sufficient returns from their existing human capital, reducing the incentive to learn their young collaborators' ideas (March 1991). One potential explanation that *can* be rejected based on the evidence in Figure 3 is that young inventors have few useful ideas to teach their older collaborative partners. The strong spillovers that circulate between young collaborators demonstrates that the knowledge held by young inventors is highly valuable for producing new inventions.

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. The model is estimated with a Quasi-Poisson and standard errors are clustered at the deceased collaborator. Dyads with surviving inventors aged 50-59 at the time of death and collaborators aged 50-59 are the reference group. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.01$.

3.4 Decomposition of Main Effect by Potential 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 (Sarada,

Andrews, and Ziebarth 2019; Kaltenberg, Jaffe, and Lachman 2023). The second 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 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 the 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 also include a third control variable, the number of patents produced the focal inventor before the collaborator's death. 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 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.[1] 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 knowledge obsolescence.

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 $JuniorDeceasedCollab_i$, $PreDeathCollabPatents_i$, $PreDeathCoinventedPatents_i$, $PreDeathSurvivorPatents_i$, $PreDeathCollabCiteAge_i$, and $PreDeathCollabCiteFertility_i$. Because collaborators' pre-death mean citation age is strongly right-

[1] For example, assume an inventor dies in 2000 produced one patent (Patent X) in the three years before death. Patent X 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.

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 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.

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

I show estimates for surviving inventors' overall patent output in Table 3. Similar results for surviving inventors' high-impact patenting output are shown in Appendix Table D1. In Table 3, 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 3: Fixed Effect Quasi-Poisson Estimates of Mechanisms for Patenting Decline

	Mechanism Tested by Model				
	Biological Age	Collaborator Productivity	Collaborative Relationship Intensity	Collaborator Knowledge Obsolescence	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
$* CollabPreDeathCollabPatents_i$		(0.00132)			(0.00151)
$PostCollabDeath_{i,t}$			0.00591		0.00383
$*$ PreCollabDeathCoinventedPatents$_i$			(0.00738)		(0.00806)
$PostCollabDeath_{i,t}$			-0.00294***		-0.00312***
$*$ PreCollabDeathSurvivorPatents$_i$			(0.000862)		(0.000933)
$PostCollabDeath_{i,t}$				-0.0912**	-0.0659*
$* LogCollabPreDeathCollabCiteAge_i * -1$				(0.0421)	(0.0394)
$PostCollabDeath_{i,t}$				-0.0570***	-0.0637***
$* CollabPreDeathCollabKnowledgeFertility_i$				(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

Notes: The table presents regression estimates for Eq. 5. Junior deceased collaborators are those that died age 20 to 45. The reference set are those that died age 46 to 59. The equation is estimated with a quasi-Poisson and standard errors are clustered at the deceased collaborator. 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 pre-death. Mean knowledge fertility is standardized by the application year of the cited patents. Focal inventor pre-death patents are 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 dyad in the 5 years pre-death. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.01$.

Table 3 shows that the recency and relevance of junior collaborator's knowledge are the strongest predictors of the treatment outcome. The coefficient on the biological age 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 outcome of the change in patent output.

Table 3 generates no support for the collaborator productivity and 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 insignificant, 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 is interpreted as the effect of the relationship's intensity.

Finally, the horse race model, which simultaneously tests all of the explored mechanisms, affirms the predictive strength of knowledge obsolescence for the treatment outcome. In the horse race model, the collaborator citation age and collaborator knowledge fertility 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.

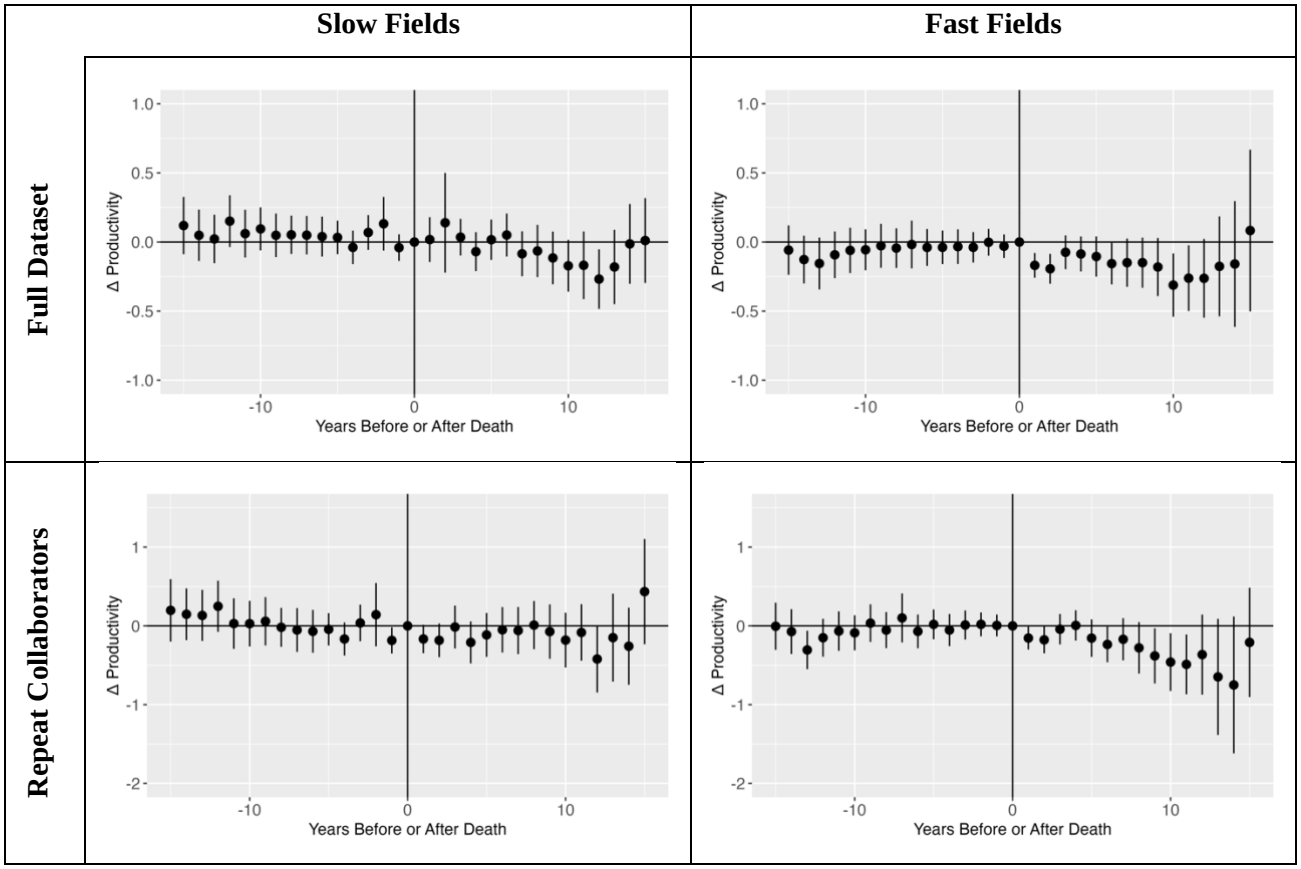
3.5 Heterogeneity of Main Effect by Field Speed

If knowledge obsolescence induces early-career inventors to generate more patenting spillovers than mid-career inventors, then one would expect the treatment effect to be larger for inventors that work in fast-advancing knowledge fields. I test this proposition by studying heterogeneity in the treatment for surviving inventors that primarily patent in fast and slow-advancing CPC subclasses fields. There are 123 unique technology classes at this level of the CPC classification scheme. 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, I use grant years to compute citation age. Descriptively, in 2010, the fastest-advancing subclasses were Information Storage, Electric Communication Techniques, and Semiconductor Devices. The slowest-advancing subclasses were Saddlery and Upholstery, Skins, Hides, Pelts, and Leather, and Brushware, which refers to the design and manufacture of brushes for painting. The regression model is similar to the one described by Equation 2. The β_1 coefficients are plotted in Figure 4.

Figure 4 shows that the negative effect of losing an early-career collaborator on subsequent patenting is more immediate and 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, which one might expect if repeat collaborations indicate stronger relationships.

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



Notes: The 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.

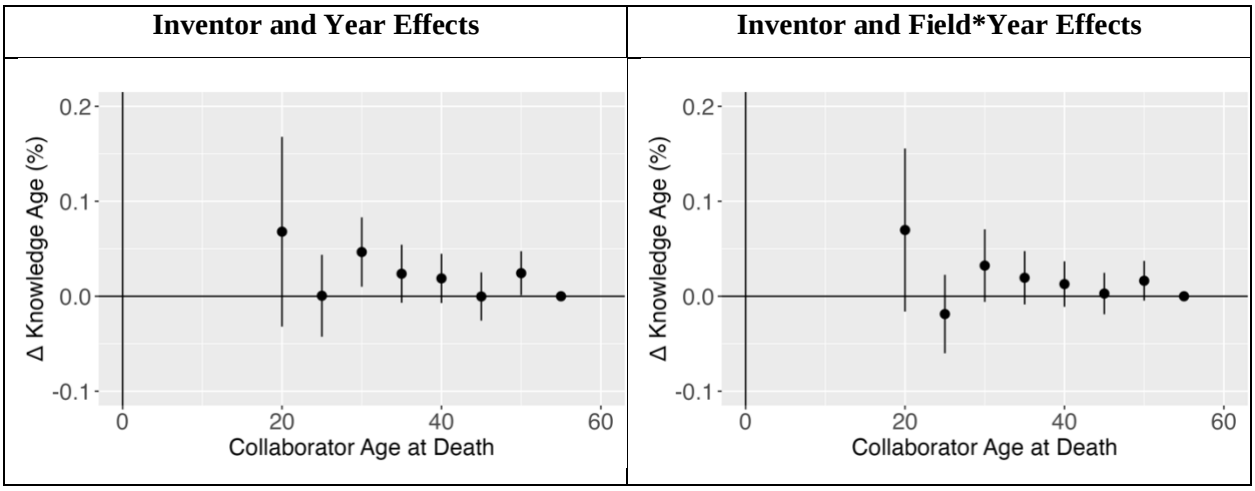
3.6) Effect of Losing a Junior Collaborator on Reference Age

A final test of whether knowledge obsolescence induces younger collaborators more valuable sources for learning is whether the loss of an early-career collaborator impedes their surviving partners' ability to learn new ideas at the knowledge frontier. 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 shift the unit of analysis from the inventor-year to the inventor-patent, and calculate the mean age of the patents cited by each patent. 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 6:

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

In Equation 6, patents are indexed by *p*. The reference group for the $AgeDeceasedCollab_i$ variable 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 6 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 5.

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



Notes: The 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 5 shows that inventors combine older ideas after losing younger collaborators. The effects are similar when field*year effects are introduced in the model (Figure 5 right panel), suggesting that inventors generally do not switch to slower-advancing technology classes following the loss of junior collaborators, but instead combine knowledge that is further from the forefront of those fields.

Discussion

This paper studied the relationship between the age of inventors and the spillovers that they absorb and generate. Data on the birth and death years of inventors were combined with patent information to identify the age at which collaborators generate the greatest spillovers for their partners, the age groups to whom those spillovers accrue, and the underlying mechanisms. 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 mostly accrue to similarly-young partners. No evidence was found in support of the view that knowledge is primarily passed down from older generations to younger ones. On the other hand, some evidence was found that knowledge is transmitted up across generations, although this channel was relatively weak. The differential effect of the death of an early-career collaborator on subsequent patenting was most strongly predicted by the type of knowledge known by young collaborators. Finally, young collaborators were shown to transmit knowledge of recently-developed ideas to their partners, and to be particularly valuable sources for knowledge in fast-advancing fields.

These findings generate five main contributions. The first is that they advance the understanding of the relationship between biological age and innovation. For several writers, age is positively associated with innovative outcomes, because age is a necessary condition for the accumulation of knowledge through experience and education (Jones 2009; 2010; Azoulay et al. 2020). On the other hand, a handful of studies have documented disadvantages associated with age for knowledge-intensive workers, including the aging of vintaged human capital (Aghion et al. 2023; Esposito and Wouden 2022; Cui, Wu, and Evans 2022) and the slower uptake of new technologies (Choi et al. 2023). More broadly, recent studies that have matched patentee information to age records have shown that the average age of inventors is only slightly higher than that of the U.S. population as a whole, and that inventors are becoming younger over time, relative to the age distribution of the full population (Kaltenberg, Jaffe, and Lachman 2023; Sarada, Andrews, and Ziebarth 2019). Furthermore, the benefits of learning from young teachers may also extend to the knowledge that inventors learn inside classrooms, as suggested by college course syllabi: universities with course syllabi that contain recently-vintaged references are associated with higher patenting performance of their students post-graduation, and younger university instructors are more likely than senior instructors to include recently-vintaged references in their syllabi (Biasi and Ma 2023). The results of the present study contribute to this literature by showing that younger individuals are critical for inventor learning not just during their youth, but also as adults.

The second contribution of this study pertains to the extent to which technological advance is a disruptive process. Because the relationship between collaborator age and knowledge dissemination depends, in part, on whether knowledge advance is disruptive (Schumpeter 1962; Caballero and Hammour 1996; Caballero and Jaffe 1993; Funk and Owen-Smith 2017; Park, Leahey, and Funk 2023), the findings

presented here suggest that knowledge advance is generally a disruptive process. Moreover, inventors appear to lose knowledge to obsolescence faster than they gain knowledge through experience, leading them to generate fewer spillovers as they age (with the exception of inventors under the age of 30). This observation has implications for not only for individual inventors, but also for their organizations, which must adopt strategies to avert obsolescence at the organizational level (Sørensen and Stuart 2000).

With regard to organizational-level obsolescence, the results presented here provide some insight into how organizations may allocate inventors across their teams in order to retain non-obsolete organizational-level knowledge. Because the strongest spillover effects were found between junior collaborators, junior inventors may benefit from the opportunity to work in close and sustained interaction with their similarly-aged peers. Furthermore, the relatively minimal spillovers that junior inventors receive from their older collaborators suggests that junior inventors may learn quickest when given some degree of autonomy. Finally, while it may be tempting to attempt to spur skill re-development in older inventors by partnering older inventors with younger partners, it is important to observe that relatively few knowledge spillovers flow from younger collaborators to their older partners (See Figure 3). Therefore, reverse-intergenerational knowledge spillovers may be insufficiently powerful to instigate meaningful skill re-development, in the absence of broader institutional or behavioral changes. In the absence of such changes, organizations may need to engage in employee turnover (Jain 2016), the acquisition of inventor teams (Marx and Timmermans 2017), or the movement of older inventors to managerial positions in order to maintain inventor-teams that work at the knowledge frontier.

The fourth implication of these findings is that they help to delineate institutional differences between academic research, technological innovation, and entrepreneurship. In academic research, scientists with seniority and prominence generate large spillovers for their direct collaborators, but their status also dissuades newcomers from entering and competing in their fields, which slows scientific progress (Azoulay, Zivin, and Wang 2010; Azoulay, Fons-Rosen, and Zivin 2019). It is possible that this mechanism found in academic research does not extend to technological invention, because inventors are more directly exposed to market-based competition than are academic scientists. Moreover, the finding that younger inventors generate more spillovers than do older inventors suggests that older inventors in technology are generally unable to stifle new technologies that threaten their aging knowledge assets. With regard to entrepreneurship, recent research has demonstrated that successful high-growth entrepreneurs tend to be somewhat older than what is imagined in the public discourse, with the most successful entrepreneurs founding companies at age 45 (Azoulay et al. 2020). In contrast, inventors generate their greatest spillovers in their 30s. These results may be reconciled by observing the different type of knowledge required of inventors and entrepreneurs. To be a successful entrepreneur requires more than a product design (which might involve technological knowledge); it also requires capital acquisition and team

organization. If individuals accumulate entrepreneurial skills more slowly than they accumulate inventor skills, or if entrepreneurial skills depreciate slower over time, then one would expect the ability of entrepreneurs to peak at a later age than the age at which inventors generate their greatest spillovers.

Finally, an important limitation of this study is its use of collaborative networks to identify how knowledge moves across ages. This data source limits the findings' generalizability to the tacit knowledge transmitted through collaborators, as opposed to the codified knowledge found in articles, software, and books. It is plausible that inventors increasingly codify their knowledge as they grow older. If so, older inventors may also provide critical contributions to the social learning process, less by benefitting their collaborators, but more through the books, documents, and theorems that bear their names. This suggests a mechanism for why the contributions of junior inventors to technology have been relatively underacknowledged in the broader discourse: with their knowledge uncodified and locked to their teams, they have less opportunity to capture value from their ideas.

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, Christian Fons-Rosen, and Joshua S. Graff Zivin. 2019. "Does Science Advance One Funeral at a Time?" *American Economic Review* 109 (8): 2889–2920. https://doi.org/10.1257/aer.20161574.

Azoulay, Pierre, Benjamin Jones, Daniel Kim, and Javier Miranda. 2020. "Age an High-Growth Entrepreneurship." *American Economic Review: Insights* 2 (1): 65–82.

Azoulay, Pierre, Christopher Liu, and Toby Stuart. 2017. "Social Influence Given (Paritally) Deliberate Matching: Career Imprints in the Creation of Academic Entrepreneurs." *American Journal of Sociology* 122 (4).

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*.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen. 2019. "Who Becomes an Inventor in America? The Importance of Exposure to Innovation." *Quarterly Journal of Economics* 134 (2): 647–713.

Bernstein, Shai, Rebecca Diamond, Abhisit Jiranaphawiboon, Timothy McQuade, and Beatriz Pousada. 2022. "The Contribution of High-Skilled Immigrants to Innovation in the United States." *NBER Working Paper Series* 30797.

Biasi, Barbara, and Song Ma. 2023. "The Education-Innovation Gap." *NBER Working Paper Series*, no. 29853.

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.

Cohen, Wesley M, and Daniel A Levinthal. 1990. "Absorptive Capacity: A New Perspective on Learning and Innovation." *Administrative Science Quarterly* 35 (1): 128–52.

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

Delgado, Mercedes, and Fiona Murray. 2023. "Faculty as Catalysts for Training New Inventors: Differential Outcomes for Male and Female PhD Students." *Proceedings of the National Academy of Science* 120 (36).

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.

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.

Jain, Amit. 2016. "Learning by Hiring and Change to Organizational Knowledge: Countering Obsolescence as Organizations Age." *Strategic Management Journal* 37 (8): 1667–87.

Jaravel, Xavier, Neviana Petkova, and Alex Bell. 2018. "Team-Specific Capital and Innovation." *American Economic Review* 108 (4–5): 1034–73.

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*.
Kogut, Bruce, and Udo Zander. 1992. “Knowledge of the Firm, Combinative Capabilities, and the Replication of Technology.” *Organization Science* 3 (3): 383–97.
Kredler, Matthias. 2014. “Experience vs. Obsolescence: A Vintage-Human-Capital Model.” *Journal of Economic Theory* 150:709–39.
March, James G. 1991. “Exploraion and Exploitation in Organizational Learning.” *Organ. Sci.* 2 (1): 71–87. https://doi.org/10.1287/orsc.2.1.71.
Marx, Matt, and David Hsu. 2022. “Revisiting the Entrepreneurial Commercialization of Academic Science: Evidence from ‘Twin’ Discoveries.” *Management Science* 68 (2): 1330–52.
Marx, Matt, and Bram Timmermans. 2017. “Hiring Molecules, Not Atoms: Comobility and Wages.” *Organization Science* 28 (6): 1115–33.
Mohnen, Myra. 2022. “Stars and Brokers: Knowledge Spillovers Among Medical Scientists.” *Management Science* 68 (4): 2513–32.
Oettl, Alexander. 2012. “Reconceptualizing Stars: Scientist Helpfulness and Peer Performance.” *Management Science* 58 (6): 1122–40.
Sarada, Sarada, Michael J. Andrews, and Nicolas L. Ziebarth. 2019. “Changes in the Demographics of American Inventors, 1870–1940.” *Explorations in Economic History* 74 (October). https://doi.org/10.1016/j.eeh.2019.05.003.
Singh, Jasit, and Ajay Agrawal. 2011. “Recruiting for Ideas: How Firms Exploit the Prior Inventions of New Hires.” *Management Science* 57 (1): 129–50.
Sørensen, Jesper B., and Toby E. Stuart. 2000. “Aging, Obsolescence, and Organizational Innovation.” *Source: Administrative Science Quarterly*. Vol. 45.
Strezhnev, Anton. 2023. “Decomposing Triple-Differences Regression Under Staggered Adoption.” *arXiv*. https://doi.org/10.48550/arXiv.2307.02735.

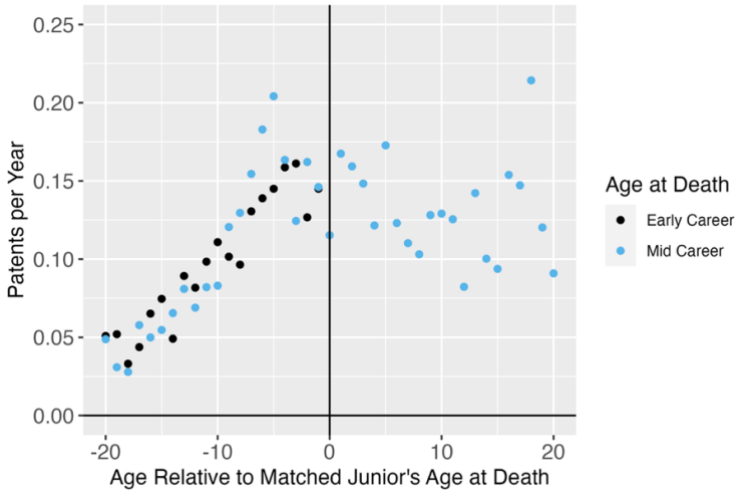
Appendices

Appendix A: Regression with Matched Deaths

To ensure that my results are not driven by unobserved differences in the ability of early-career and mid-career deceased collaborators, I repeat my main analyses using a matched sample. I exactly match deceased early-career collaborators to deceased mid-career collaborators based on three criteria: they must share the same modal technology class (or technological specialization), they must have produced the same total number of patents up to the age of death of the early-career collaborator, and they must have produced the same number of patents in the 3 years leading up to the age of death of the early-career collaborator. Together, the second and third criterion ensure that the matched pairs to have similar cumulative output and similar trends in output, up until they reach the age at death of the junior deceased collaborator. I match each deceased junior collaborator to one deceased mid-career collaborator. When a junior deceased collaborator has more than one potential match, I randomly select one from the viable set, with replacement. This process results in a dataset of 3,961 deceased junior collaborators matched top 3,060 unique deceased mid-career collaborators.

Figure A1 shows the mean number of patents produced per year by living collaborators that died early in their careers and matched collaborators that died mid-career by age, relative to the age of death of death of the early-career collaborator. Patent production is similar for both groups through the age at death of the early-career collaborators die.

Figure A1: Patent Production by Age, Benchmarked to Junior Deceased Collaborators' Age at Death



Note: The mean calculations are conditional on being alive at a given age, so collaborators that die mid-career are not included in the calculation for the years after their death.

To test whether the loss of a junior collaborator induces a larger decline in patenting productivity than does the loss of a mid-career collaborator, I estimate the difference-in-difference model specified in the main text, but using the matched sample. The model is given by Equation A1:

$$(A1)\quad PatProd_{i,t} = \exp\left[\beta_0 PostCollabDeath_{i,t} + \beta_1 PostCollabDeath_{i,t} * JuniorDeceasedCollab_i + \alpha_i + \tau_{modalfield*t} + \varepsilon_{i,t}\right]$$

I estimate Equation A1 with a quasi-Poisson estimator and cluster the errors at the deceased collaborator level. I estimate the model five times. I first estimate the model without the $PostCollabDeath_{i,t} * JuniorDeceasedCollab_i$ interaction term, to compare the effect size of $PostCollabDeath_{i,t}$ with the existing literature. After this, I estimate the complete model (including the interaction term) using four specifications. The first specification uses patents produced per year as the dependent variable; the second specification uses highly-cited patents as the dependent variable; the third specification patents produced per year as the dependent variable and restricts the sample to inventor-collaborator dyads that co-invented two or more patents (that is, they are repeat collaborators); the fourth uses highly-cited patents as the dependent variable also restricts the sample to repeat collaborators. Estimates are given in Table A1.

Table A1: Quasi-Poisson Estimates of Collaborator Loss on Patenting (Matched Sample)

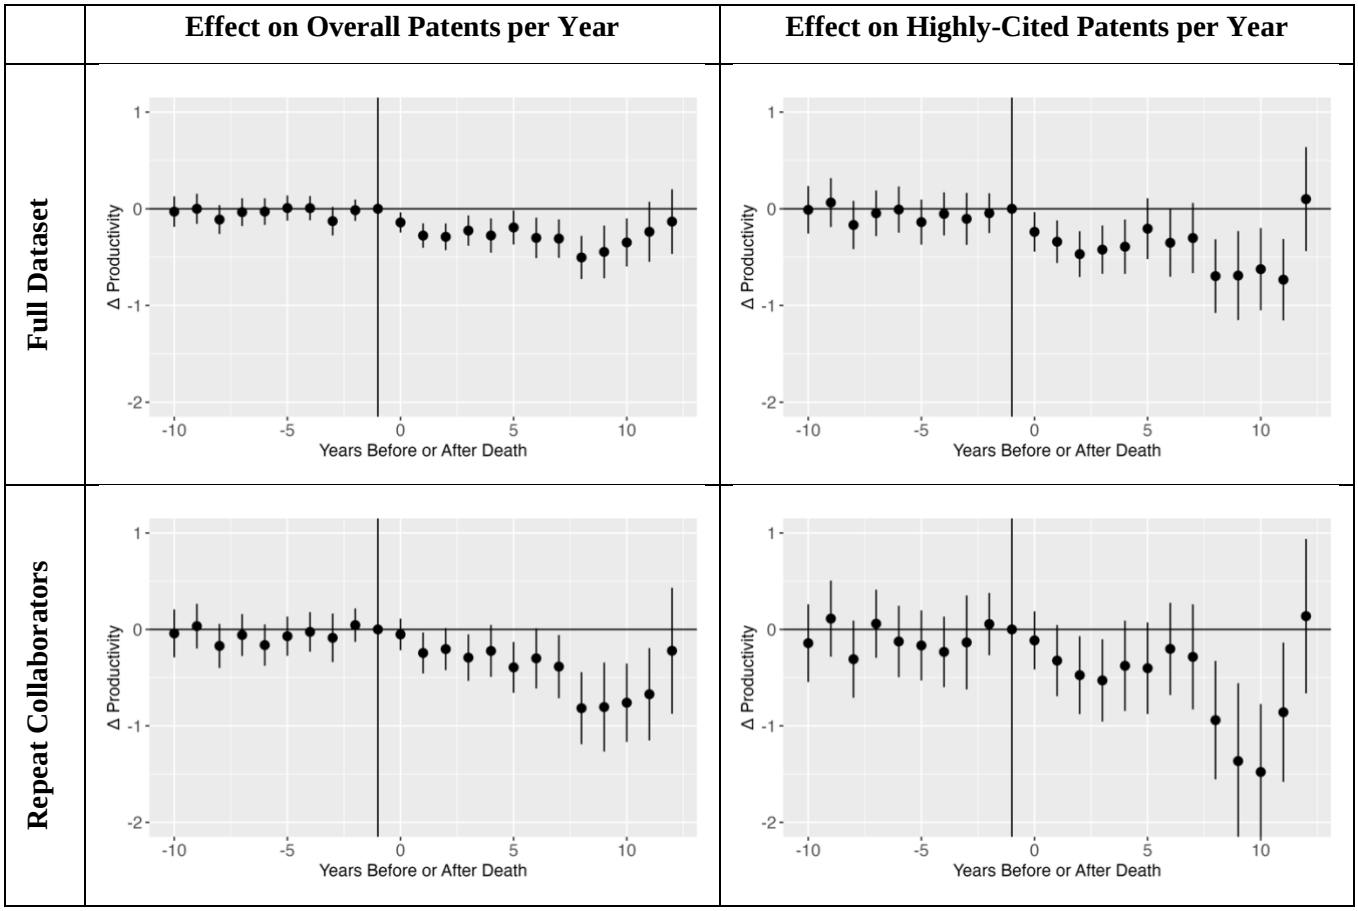
	Complete Matched Sample			Repeat Collaborators Sample	
Dependent variable:	Total Patenting	Total Patenting	High-Impact Patenting	Total Patenting	High-Impact Patenting
$PostCollabDeath_{i,t}$	-0.178***	-0.0206	0.0895	-0.00827	0.0270
	(0.0323)	(0.0455)	(0.0726)	(0.0672)	(0.106)
$PostCollabDeath_{i,t}$		-0.241***	-0.351***	-0.274***	-0.407***
$* JuniorDeceasedCollab_i$		(0.0490)	(0.0764)	(0.0766)	(0.117)
Inventor and Class*Year Fixed Effects?	Y	Y	Y	Y	Y
Inventor*Year Observations	115,042	114,167	75,297	43,430	29,127
Inventor Observations	16,703	16,578	8,882	5,949	3,546

Notes: The table presents regression estimates for Equation A1. 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 that have been matched to early-career deceased collaborators (see text). Repeat collaborators are those which surviving inventors co-invented 2+ patents. The model is estimated with a quasi-Poisson. Standard errors are clustered at the deceased collaborator. *** p < 0.01; ** p < 0.05; * p < 0.01.

The results in Table A1 show that the loss of a junior collaborator induces a larger decline in patenting output than does the loss of a mid-career collaborator with a similar career patenting profile. The first regression, which does not include the interaction term with *JuniorDeceasedCollab*, shows that inventors in the sample experience an exp(-0.178)-1 = 16% decline in patent output following the loss of a collaborator. This effect size is slightly larger, but in the same vicinity, of prior work using a different inventor death sample (Jaravel, Petkova, and Bell, 2018). The second through fifth regressions show that effect of losing a collaborator on patenting can be fully accounted for by whether that collaborator was young at time of death.

Finally, to further confirm that the mid-career deaths in the matched sample are valid counterfactuals for the early-career deaths, I show event plots using the matched sample in Figure A2. The figure shows that the pretends are parallel. The regression model is similar to Equation 1. The pre-trends are parallel across all four specifications.

Figure A2: Event Plots of Differential Effect of Death of Junior Collaborator on Patent Production (Matched Sample)



Note: The plotted values are coefficients on the $\text{YearsToCollabDeath}_{i,t} * \text{JuniorDeceasedCollab}_i$ term in a regression similar to the one specified by Equation A1.

Appendix B: 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. (2017) 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.

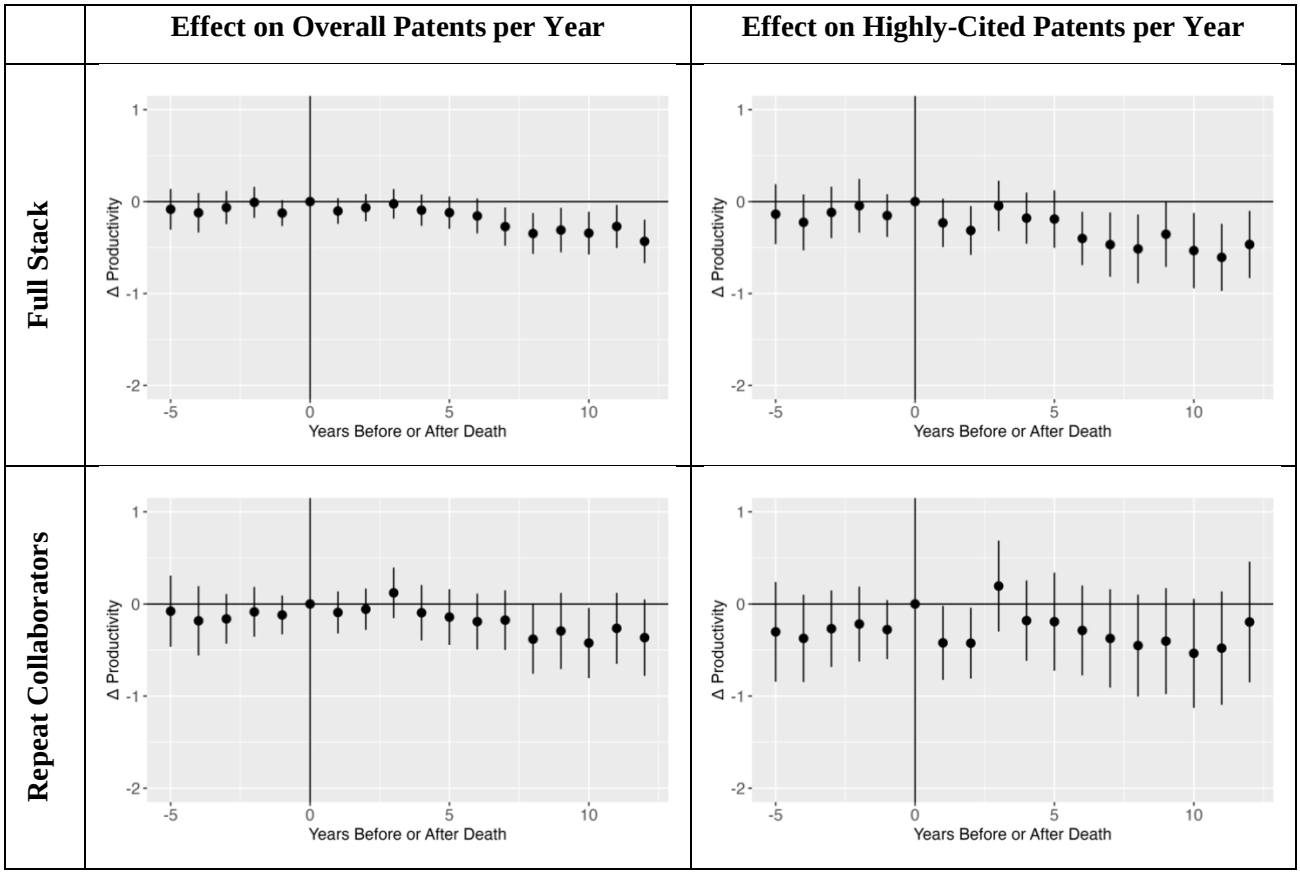
To estimate the effect of the loss of a junior collaborator on surviving inventors' patenting productivity, I estimate the model described by Equation B1.

(B1)
$$PatProd_{i,t} = \exp\left[\sum_{\tau=-5}^{12} \Phi_\tau TimeToSubExperiment_{idt}^{\tau} + \sum_{\tau=-5}^{12} \Psi_{\tau i} TimeToSubExperiment_{idt}^{\tau} * Treat_{id} + \sum_{\tau=-5}^{12} \Omega_{\tau i} TimeToSubExperiment_{idt}^{\tau} * Treat_{id} * JuniorDeceasedCollab_i + \alpha_{id} + \rho_t + \varepsilon_{it}\right]$$

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 B1 four times to explore robustness of the results. The coefficient values are plotted in Figure B1. 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 B1: Change in Patenting Productivity Following Death of Junior Collaborator Relative to Mid-Career Collaborator using Stacked Regression



Notes: The figure 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.

Appendix C: The Relationship between Collaborator Age, Knowledge Age, and Knowledge Fertility

In the main text, Table 3 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 C1:

(C1) $$MeanCiteAgePreDeath_j = \beta_0 + \beta_1 AgeAtDeath_j + \lambda_t + \varepsilon_j$$

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 C1 using the full set of deceased collaborators are provided in Table C1. The results in the table show that older collaborators cite older and less-fertile patents.

Table C1: Regression Results for Equation A2 on Collaborator Age, Knowledge Age, and Knowledge Fertility

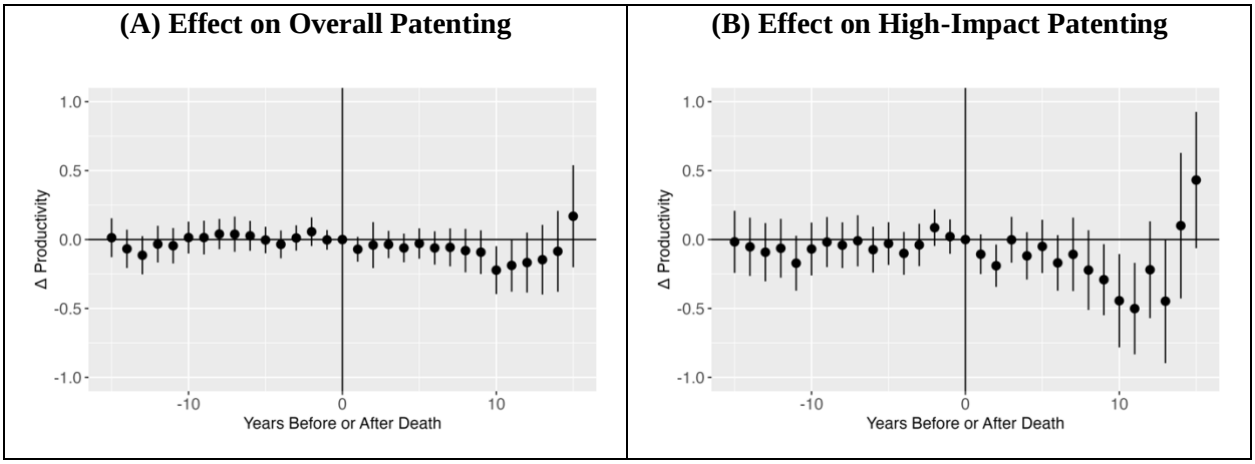
	Dependent Variable	
	$PreDeathMeanCiteAge_j$	$PreDeathMeanCiteFertility_j$
$AgeAtDeath_j$	0.0427*** (0.00884)	-0.00221* (0.00116)
Death Year Fixed Effects?	Yes	Yes
R^2	0.021	0.004
Observations (Collaborators)	9,757	10,135

Notes: PreDeathMeanCiteAge and PreDeathMeanCiteFertility are computed using the patents applied for by the deceased collaborator in the 3 years before their death. *** p < 0.01; ** p < 0.05; * p < 0.01.

Appendix D: Additional Results

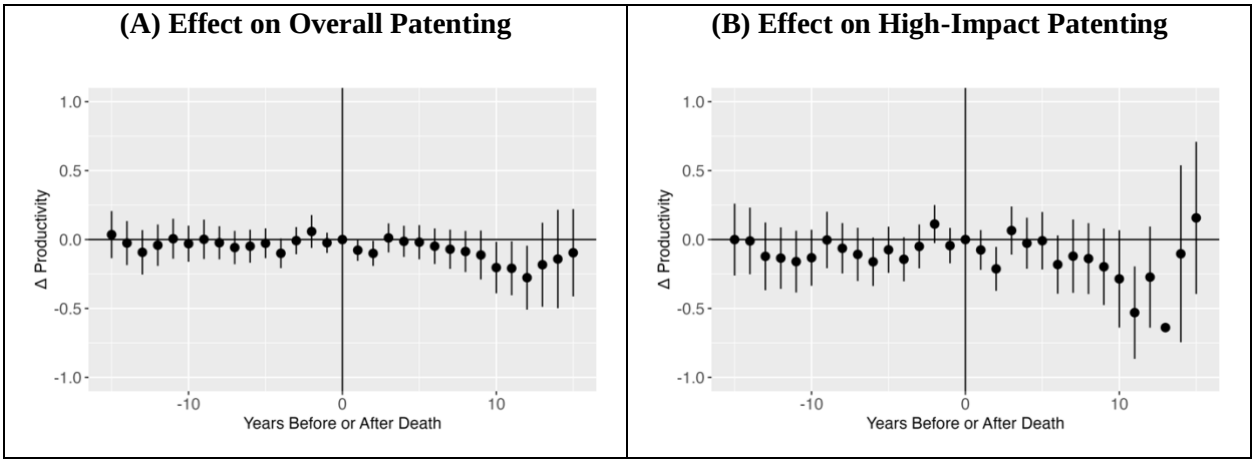
This appendix shows results omitting inventors that die between ages 55, omitting inventor-collaborator dyads that did not collaborate within 3 years pre-death, and a decomposition of mechanisms with high-impact patenting as the outcome variable. The findings are similar to those in the main analysis.

Figure D1: Change in Patenting Productivity by Deceased Collaborator Age, Deaths Under Age 55



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 D2: Change in Patenting Productivity by Deceased Collaborator Age, 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.

Table D1: Regression Results for High-Impact Patenting

	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
$*$ PreDeathCoinventedPatents$_i$			(0.00747)		(0.00700)
$PostCollabDeath_{i,t}$			-0.00276**		-0.00308***
$*$ 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 of the main text. The outcome variable is the number of high-impact patents (top-25% cited within-year) produced by a focal inventor in a 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. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.01$.