Supporting Information for

The Fragility of Artists' Reputations from 1795 to 2020

Letian Zhang, Mitali Banerjee, Shinan Wang and Zhuoqiao Hong

Corresponding Author Letian Zhang.
E-mail:letian.lt.zhang@gmail.com

This PDF file includes:

Supporting text
Figs. S1 to S8
Tables S1 to S7
SI References

Supporting Information Text

1. SUPPLEMENTARY MATERIALS: METHOD

A. A Machine-learning Approach: Word Embedding. We implemented a large-scale historical text analysis to understand changes in artists' reputations over time. When directly surveying respondents is not possible, text analysis can help us understand how society perceives a particular artist. In fact, text has served as a key data source for studying cultural history, as it is often the most semantically rich record a group leaves behind. Traditionally, scholars analyze text using interpretive close reading or systematic qualitative coding, both of which require intense human input and therefore limit the scope of analysis. Recent advances in natural language processing have enabled scholars to analyze massive text data systematically. The latest models have made great strides by representing relationships between words as vectors in a dense, continuous, high-dimensional space (1, 2). These vector-space models, known collectively as word embedding, have attracted widespread interest among computer scientists and computational linguists due to their ability to capture and represent complex semantic relations.

Word-embedding algorithms try to account for the different contexts in which words occur and use these contexts (embedding) to describe words in a dense vector space (with typically around 300 dimensions). This approach is based on the distributional hypothesis that "words that occur in similar contexts tend to have similar meanings" (3). In essence, these algorithms turn the embedding problem into a prediction problem: they try to optimize a vector such that it can be used within a predictive model (e.g., a logistic regression classifier) to predict which words are in the context—typically, all words within three to six words on either side of the target word. To measure the distance between words in an embedding space, scholars typically use cosine similarity, which is the cosine of the angle between two vectors. This is preferred to the Euclidean distance due to properties of high-dimensional space that violate intuitions formed in two or three dimensions (4).

In word-embedding models, words need not co-occur for their vectors to be positioned close together, since the optimal distance between two vectors is a function of shared context rather than strict co-occurrence. For instance, if both "basketball" and "tennis" appear near words such as "athletes," "tournaments," and "winning," then the vectors for "basketball" and "tennis" would be located near each other in the embedding, even if they never appear together in a text. Thus, a word's broader neighborhood in the embedding space is typically populated by terms with related meanings.

Simply by examining the shift in word vectors that surround an artist's name (which is treated as a word in the vector space), we can detect how people's perception of that artist changes over time. Scholars have used such methods to measure cultural changes. For example, Kulkarni et al. (5) have used word-embedding models in this way to trace shifts in the meaning of the word "gay" over the course of the twentieth century, from a location in the vector space beside "cheerful" and "frolicsome" to one near "lesbian" and "bisexual." Others have used the same strategy to examine shifting meanings of social class and changing gender and ethnic stereotypes (6, 7).

The precision of the word-embedding model has been well tested. In fact, it can even solve analogy problems by applying simply linear algebra to word vectors. For example, the analogy "man is to woman as king is to ______" can be solved with the word vectors $\overrightarrow{king} - \overrightarrow{man} + \overrightarrow{woman}$, with the resulting vector most proximate to the word vector for *queen*. This is because $(\overrightarrow{man} + \overrightarrow{woman})$ closely corresponds to a gender dimension. By adding $(\overrightarrow{man} + \overrightarrow{woman})$ to $\overrightarrow{king}$, we are essentially starting at king and moving one step on the gender dimension toward more femininity. Word-embedding algorithms can obtain as high as 74-percent success on challenging analogy tests (8).

B. Constructing Embedding Models. We picked twenty major newspapers in the United States: *New York Times*, *Wall Street Journal*,*Los Angeles Times*, *Washington Post*, *Chicago Tribune*, *Newsday*, *San Francisco Chronicle*, *Boston Globe*, *Philadelphia Inquirer*, *Detroit Free Press*, *St. Louis Post Dispatch*, *Baltimore Sun*, *Louisville Courier Journal*, *Cincinnati Enquirer*, *Hartford Courant*, *Nashville Tennessean*, *Austin American Statesman*, *Dayton Daily News*, *Christian Science Monitor*, and *New York Herald Tribune*. Our selection criteria are mainly based on circulation size, years of history, and the availability of digitized text. We included a few smaller newspapers to ensure some geographic variation and to provide a diverse range of political orientations.

Historic newspapers provide a massive amount of text data, which is crucial to reproducing accurate semantic relationships in word-embedding models. For example, a major newspaper in the US publishes over 50 thousand articles each year, which contain about one billion words. More importantly, the long time-span of these historical newspapers enables us to use the same data source across all periods. Previously, text analyses of cultural change had to rely on different sources for different periods. For example, the widely used Corpus of Historical American English dataset contains a combination of books, periodicals, magazines, and newspapers. Not only do the sources of these texts differ from period to period, but so does the genre composition. This creates a major concern that any observed cultural change simply reflects differences in the underlying sources. Although newspapers also change their ownership and writers, they provide a much more consistent data source.

We obtained digitized texts of these newspapers from ProQuest and Factiva, whose newspaper database has digitized a selection of historical newspapers. As Table S1 shows, we have digitized versions of these newspapers from 1795 to the present, although many of these newspapers have missing years prior to 1870. Most of these newspapers are daily news, thus providing an extraordinarily rich set of text data. Our entire corpus of newspapers contains more than 32 billion words, making it the largest historical text corpus currently available. In comparison, Corpus of Historical American English—currently the largest structured corpus of historical English from 1820 to 2010—has 475 million words; Corpus of US Supreme Court from 1790 to the present has 130 million words; and TIME Magazine Corpus from 1923 to 2006 has 100 million words.

[insert Table S1 about here]

We combined corpora from different newspapers and divided them into five-year periods. Table S2 shows that, at the higher end, many time periods have more than one billion words. At the lower end, even the period with the smallest corpus (1796–1800) has more than four million words. The large amount of text in each time period ensures that we can accurately map artist names and words onto the embedding vector space. In general, we have fewer observations from the early nineteenth century and from the most recent decade (2010–2020). To make sure that the high missing rate in these years is not driving our results, we ran separate analyses focusing on just the period from 1870 to 2010 with substantively similar results.

[insert Table S2 about here]

We trained separate models on texts from 1796–1800, 1801–1805, 1806–1810, and so on through 2016–2020, resulting in 45 independently constructed word-embedding models. By comparing these models side-by-side, we can trace how each artist's reputation has shifted over a long period. To make sure that our models correctly capture artists' names, we preprocessed artists' names (we discuss our sampling strategy below) in our models.

We used *Word2vec*, the most widely used word-embedding algorithm, to train our newspaper corpus. *Word2vec* has two model architectures to produce a distributed representation of words: continuous bag-of-words and continuous skip-gram. In the bag-of-words architecture, the model predicts a focal word from a window of nearby words. This approach does not consider the order of the surrounding words. Models using the skip-gram architecture use the focal word to predict the surrounding window of nearby words. Such an approach weighs nearby context words more heavily than more distant context words. We used the skip-gram architecture, since it can better capture infrequent words and phrases.

A newspaper corpus is by no means a culturally representative sample of the US general public. Newspapers are written mostly by journalists and writers and target a relatively elite, literary public. They are thus poorly suited for identifying subcultural views and minority voices. Nonetheless, they are a widely used source for studying public perception, as they generally capture mainstream values and beliefs.

C. Artist Sample. To identify a list of artists, we focus on the well-known ones because they are more likely to be consistently mentioned by newspapers across time periods. Our definition of artists is rather broad, including painters, sculptors, musicians, composers, filmmakers, architects, and writers. We chose this broad definition because we also want to understand if the posthumous change in reputation differs across artistic fields.

We used the Pantheon 2.0 database—covering all individuals with a Wikipedia page—to create a list of well-known artists from around the world. Our final list contains 3,394 artists in the Pantheon dataset who satisfy the following criteria: (a) they are mostly known for their achievement in art, music, film, or writing; (b) they died no later than 2010; and (c) their full name is mentioned at least five times in at least one of our embedding models. This list includes a few artists prior to the Renaissance period; we excluded them because an ancient Greek poet like Homer differs from a seventeenth-century artist in too many dimensions. For each artist, the Pantheon database includes information on the year of birth and death, birthplace, place of death, gender, and occupation. We complemented this information by manually coding each artist's race, sexual orientation, socioeconomic status at birth, education level, cause of death, and wealth at time of death.

Our list of artists spans the seventeenth to the twenty-first century. As Figure S4 shows, about half were born in the nineteenth century; 46 percent were born in the twentieth century. The median lifespan is 74 years but there is significant variation. Age at death ranged from 22 to 108; the 25th percentile of our sample has a lifespan of 62 years and the 75th percentile has a lifespan of 84 years. Geographically, although these artists' birthplaces span 95 countries, they are heavily Western-centric: about 46 percent were born in the United States, 17 percent in the United Kingdom, 9 percent in France, and 13 percent in the rest of western Europe. While part of this Western bias is likely a reflection of Wikipedia's user base, some of it is also attributable to our embedding models, which rely on major US newspapers.

[insert Figure S4 about here]

Table S3 shows the occupational breakdown of artists in our sample. 39 percent are writers, 17 percent musicians, 11 percent film directors, 12 percent composers, and 10 percent painters. These are quite different forms of art. In the analysis, we will show occupational as well as aggregate trends.

[insert Figure S3 about here]

D. Measuring and Validating Artist Reputation. We measured artist reputation in each period using our trained embedding models. First, we identified a list of keywords—such as "successful," "gifted," and "accomplished"—that represent artist reputation. To do so, we surveyed 200 online participants and asked each to list a set of words (both adjectives and nouns) that would describe an accomplished artist and another set that would describe a failed artist. A few of the words listed, such as "old" and "rich," are descriptions of an artist's personal characteristics; some respondents listed specific artist names; a few words are highly ambiguous, such as "good" and "great." We excluded these words, as well as those with misspellings. We tallied the remaining words and Table S4 shows the most commonly mentioned words in order. There is more consensus on positive than negative reputation words, hence a longer list for positive words.

[insert Table S4 about here]

We tested a wide range of combinations of these keywords, using different numbers of positive and negative words and different combinations of these words. We also replaced many of these words with close synonyms to see whether that makes a striking difference. In total, we tested more than 300 combinations of these words. Using different versions of the word list produces slightly different values for artists' reputation (we discuss the detailed process of variable construction and analytical strategy below), but these different versions of reputation are highly correlated—with a pairwise correlation above 0.85 in all cases—and the posthumous reputation changes are strikingly similar in direction and magnitude. For example, Table S1 shows posthumous reputation change using a different number of reputation keywords: they produce qualitatively similar conclusions. For this analysis, our reputation keywords are the 20 most frequently mentioned positive words and the 10 most frequently mentioned negative words.

After establishing a set of reputation keywords, we calculated the average geometric distance between the vector of an artist's name and the vector of each of these reputation keywords. The closer an artist's name is to a positive (negative) reputation word in the vector space, the more (less) accomplished the artist is perceived to be. We calculated the cosine distances between an artist's name and each positive reputation keyword and took the average, then subtracted the average cosine distance between an artist's name and each negative reputation word. We treated each reputation keyword in an equal manner and, by taking the average across all of them, we minimized the impact of any unexpected connections between an artist and a particular word (e.g., an artist may have used a certain keyword in his or her work). The resulting variable reflects an artist's perceived reputation in each time period. Our list of artists includes only accomplished artists, hence the cosine distances in our sample are generally positive.

One potential concern with this approach is that different time periods may use reputation words in slightly different ways. For example, the word "accomplished" may be used for above-average artists in the twenty-first century but only for exceptional artists in the nineteenth century. When examining across time periods, temporal changes in artists' cosine distance to reputation words may reflect such changes in the use of those words. Given the large number of artists in the sample, we address this concern by simply standardizing—across all artists in each time period—the average cosine distance between an artist and the reputation words. The resulting measure follows a standard normal distribution for each time period.

To validate our measure of artists' perceived reputation, we compared it to an online survey that we conducted in October 2021. We picked the 50 artists in our sample with the highest number of Wikipedia views—a measure of their current popularity—because the public is more likely to be able to identify them. We recruited 500 respondents from the platform Prolific to evaluate each of these 50 artists. We recruited our respondents so that our sample corresponds to the general US population in terms of education level, age, gender, and race. Each respondent was given a list of 20 artists and was asked "how would you assess the accomplishment of the following artists" on a 10-point scale. Thus, each artist is evaluated by 25 respondents. We then simply averaged the survey score for each artist and compared it to the average estimated reputation score derived from the embedding models from 2000 to 2020. Figure S5 shows a correlation of 0.75 between the survey score and our embedding score, suggesting that our embedding model generally captures the public perception of these artists.

[insert Figure S5 about here]

E. Coding Unexpected Death. We defined "unexpected death" as death due to an accident, a stroke or heart attack (with no prior health issues), suicide or homicide (including war), drug/alcohol overdose, or infection. We defined "regular death" as death due to long-term illness or to old age. Two trained research assistants manually coded each artist's cause of death, with an intercoder reliability of 0.9. For those few cases with disagreement, we discussed each case and made a decision. We categorized 61 artists' deaths as unexpected and 894 as regular. For the remaining artists, we could not find sufficient information to make a verdict.

F. Measuring Artist Visibility. Besides reputation, we also examined each artist's visibility in each time period. Intuitively, visibility should decline after an artist's death. We examine this intuition but mostly focus on visibility as a potential mechanism of reputation change.

Capturing an artist's visibility in any given period is much more straightforward than measuring reputation: we simply used the number of times an artist is mentioned in newspapers and books. For newspaper mentions, we used newspapers.com, the most comprehensive databases of historical newspapers, covering more than one million US newspapers since the nineteenth century. We wrote a simple algorithm to count the number of times an artist's name appears across all newspapers in each of our five-year periods. We focused on the total count and did not differentiate where the name appeared. For example, if an artist's name appears five times in one newspaper and three times in another, we consider that artist to have appeared eight times. Book mentions are based on Google books, one of the largest collections of full-text books in English. We used the Google Books Ngram Viewer, which allow us to enter a phrase and observe its frequency of occurrence in each decade. As with our newspaper count, we used the total number of times an artist's name appears, regardless of whether these appearances are in the same book or different books.

Both the number and the length of instances of media coverage vary significantly over time. Thus, a simple count of an artist's name could be misleading, since some periods could have more (or longer) newspaper articles and books available than others. We used two strategies to alleviate this concern. The first is simply adding time-period fixed effects in our estimation, which account for macro-level changes in newspaper and book coverage. Alternatively, we calculated percentile rank—based on the total number of mentions—across all artists in the sample within the same time period. The two measures are correlated at 0.85 and both produce substantially similar conclusions.

An artist's appearances in newspapers and in books are correlated. Using both logged total counts and percentile rank, we found that correlations between the number of newspaper mentions and of Google Book mentions are around 0.6. The difference between the two may reflect the fact that most newspapers target the general public while many books have more specific target audiences. Nevertheless, analyses using both newspapers and books produce qualitatively similar patterns.

Interestingly, the resulting measures of artist visibility have only weak correlations with our measures of reputation, ranging from 0.01 to 0.18 depending on the measure we use. The low correlation confirms the common intuition that the most frequently mentioned artists may not be the most highly regarded. This suggests that although visibility and reputation are related, they are driven by different factors. For instance, a scandal could damage an artist's reputation but improve his or her visibility.

G. Measuring Taste. To measure taste, we focus on painters, who are more easily categorizable into genres. We hired a research associate experienced in art history to manually categorize painters in our sample into 32 genres, such as romanticism, post-impressionism, and minimalism. Some painters can be categorized into one genre, others are associated with two or three. A few painters were so ambiguous in their style that we decided to remove them from the sample. In the end, we were able to identify the genre of 128 painters: 45 percent are associated with one genre, 23 percent with two, and 32 percent with three or more.

We then calculated each genre's reputation in the same way that we calculated artist reputation. Using the same set of keywords, we used our embedding models to calculate the cosine similarity between a genre and reputation keywords. We then mapped genre reputation to each painter in our sample. For painters associated with multiple genres, we used the one with the highest reputation score. Given the large number of artists per genre, genre reputation and artist reputation are only weakly correlated at -0.05, alleviating any collinearity concern.

2. SUPPLEMENTARY MATERIALS: ADDITIONAL RESULTS

A. Detailed Results on Demographic Traits and Posthumous Reputation. This section provides more detail on how the posthumous decline in reputation varies across demographic groups. We considered three demographic dimensions: race, gender, and country of residence. Figure S6 shows these patterns. In Figure S6a, we find that the posthumous decline in reputation does not apply to the 195 racial minority artists in our sample. Our US-centric sample has few Asian or Hispanic artists; the majority of our racial minority artists are Black Americans. While the small number of racial minority artists in our sample makes our analysis more speculative, it is possible that earlier Black artists' work was underappreciated due to cultural bias and then gained more traction over the years as the mainstream US public became more open to minority culture.

[insert Figure S6 about here]

Figure S6b shows that there is little difference in reputation decline across gender: both male and female artists experienced a similar reputation decline after death. Figure S6c breaks down the sample by artists' location (at the time of death). We find that non-US artists experienced a considerably larger reputation decline than US artists. The cause of this gap is unclear: it may be that American artists have more visibility than non-American artists in US newspapers after their death, or that US audiences have more appreciation for earlier US art, or simply that American and European artistic taste have diverged over the years.

B. Detailed Results on Reduced Visibility. This section provides more detail on how visibility explains some of the posthumous reputation decline. Artists can lose visibility after they die and thus have fewer opportunities for recognition and acclaim. In Figure S7, we test artists' declining visibility after death. As in our main model, we used logged number of appearances in newspapers as the dependent variable and estimated it using years after death, with fixed effects on both time periods and artists. As the figure shows, the amount of newspaper coverage increases over time as an artist ages, a pattern that differs from that of reputation, which tends to be stable during the last years of an artist's life. Consistent with our intuition, after an artist's death, newspaper coverage drops precipitously. This pattern applies to painters, writers, musicians, composers, and other types of artist. In Figure S7, we show that results using book coverage instead of newspapers and using percentile rank instead of logged counts produce substantively similar patterns.

[insert Figure S7 about here]

After establishing the posthumous decline in visibility, we included visibility as a control in our model, which reduces the posthumous decline in reputation by about a third. To illustrate this mediating process, Table S6 uses percentile rank of newspaper mentions, which is comparable across periods. As the table shows, as more years pass after death, an artist's visibility decreases. Without controlling for visibility, the annual rate of reputation decline is about 0.0084, but this drops to 0.0057 after accounting for visibility. Given the importance of visibility in promoting reputation, this result gives some indication that reduced visibility might be a mechanism underlying artists' posthumous reputation decline.

[insert Table S6 about here]

C. Detailed Results on Changing Taste. This section discusses in greater detail how changing taste explains artists' posthumous reputational decline. As aesthetic taste shifts, what was considered accomplished art in the past may no longer be perceived in the same way. To examine this mechanism, we focus on painters, who can be more easily categorized into genres. The main idea is that our appreciation of a particular genre may change over time, which could explain why we appreciate an artist less as time passes.

Table S7 shows how genre reputation mediates the posthumous reputation decline. First, using individual fixed effects, we found that the reputation of an artist's associated genre declines after artist death: genre reputation declines by half a standard deviation in about 60 years. This is consistent with our intuition that taste in genre gradually shifts and what was considered a highly reputable genre during an artist's lifetime may no longer be considered so decades later.

[insert Table S7 about here]

Next, we estimated an artist's reputation: we first estimated it only using year after death and then successively added artist visibility and genre reputation as controls. As Table S7 shows, the addition of artist visibility in this case only decreases the coefficient for *year after death* by about four percent. But the addition of genre reputation decreases the coefficient by another 32 percent. Thus, for this group of painters, we find that reduced visibility explains only a small part of their posthumous reputation decline, while a changing perception toward genres explains almost a third of it.

D. Reputations of a Few Famous Artists Might Shape Beliefs about Enduring Reputations. In this section, we discuss in greater detail what might be driving the popular glorifying-the-dead belief. We think two related dynamics may be responsible. First, the media and popular discourse love rags-to-riches stories and frequently feature artists who received little recognition for years but suddenly gained acclaim either late in life or after death. Although these artists are few in number, their stories may have received disproportional coverage. The second dynamic may be selection bias. The artists who lost reputation over time may no longer be widely discussed; hence the artists that come to mind today are more likely those who have maintained their reputation. For instance, many people today would not have recognized Patrick Dennis, who was once regarded as one of the best writers of his generation. These two processes could have fostered the false perception that we glorify the dead.

To shed some light on this, we analyzed the most visible dead artists today, based on the number of times they appear across all the newspapers from 2016 to 2020. Given their high visibility, these artists' stories could make a strong impression on the general public. Figure 5(a) plots the reputation trajectory of the 50 most frequently mentioned dead artists today. Unlike the general pattern, we find that this group of artists did not experience a reputation decline after death. If anything, their reputation a century later is slightly higher than it was during their lifetime. However, we should keep in mind that this is an extremely small group. In Figure 5(b), when we expand the sample to the 500 most frequently mentioned dead artists today, their reputation trajectory is similar to what we saw in the main analysis: their reputations significantly declined after death. These analyses suggest that our glorification of the dead may be based on a very select group of highly visible artists.

[insert Figure 5 about here]

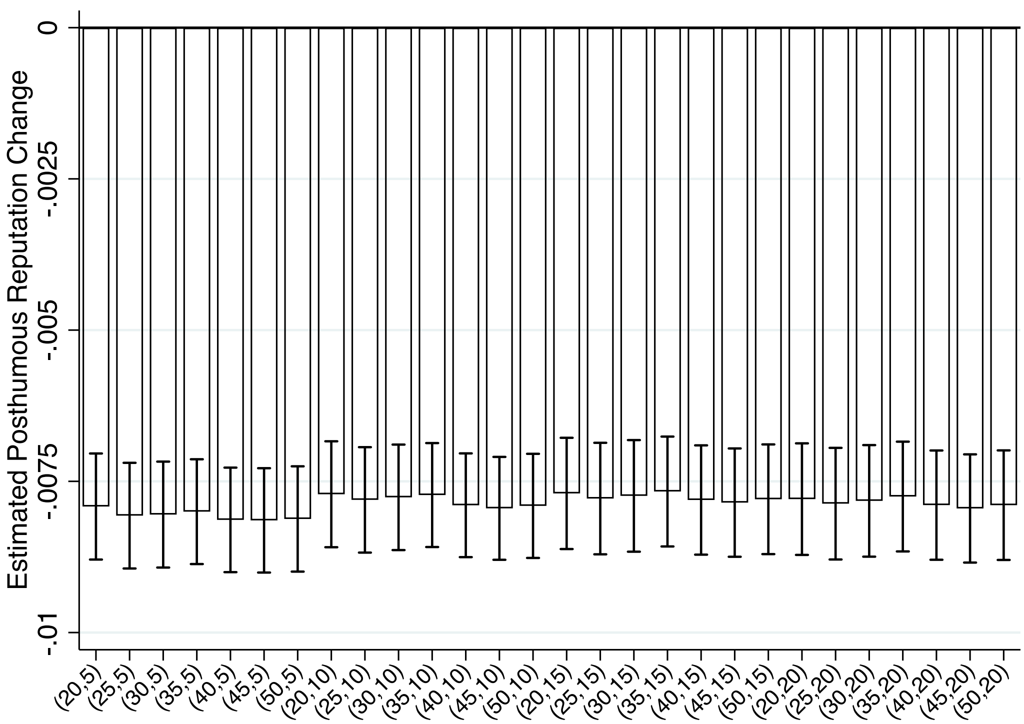

Fig. S1. Testing Different Combinations of Reputation Keywords

Notes: This figure shows the estimated posthumous reputation decline using different combination of reputation keywords. On the x-axis, (a, b) indicates using the first a positive words and first b negative words in Table S4. Keywords words in Table S4 are generated from an online survey. The y-axis is the estimated coefficient for year after death in predicting reputation.

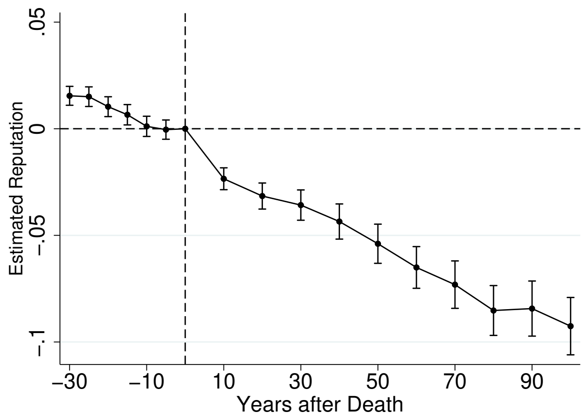

(a) Reputation: Not Standardizing by Year

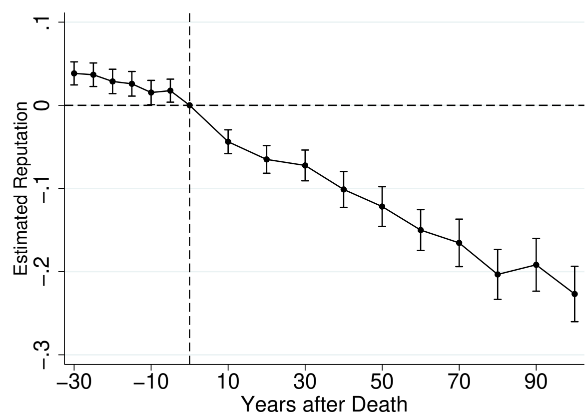

(b) Reputation: Using Percentile Rank

Fig. S2. Alternative Measures of Artist Reputation

Notes: Figures a and b show estimated change in an artist's reputation years before/after death, similar to Figure 1, except using alternative measures of reputation. Figure a does not standardize the reputation measure by year and Figure b uses the percentile rank of all artists in the same year. Figure c uses the same reputation measure as in Figure 1 but controlling for artists' newspaper count percentile. All models use individual artist fixed effects and are clustered by artist.

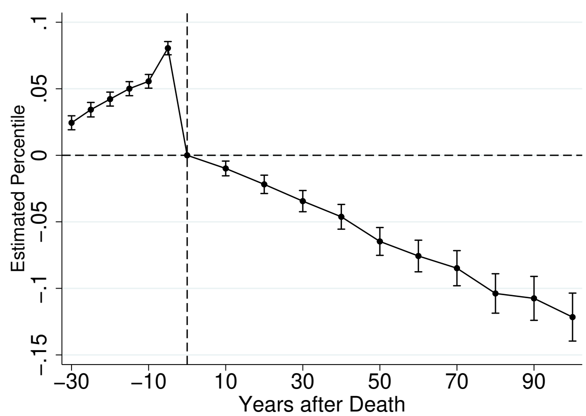

(a) Percentile Rank of Newspaper Mentions

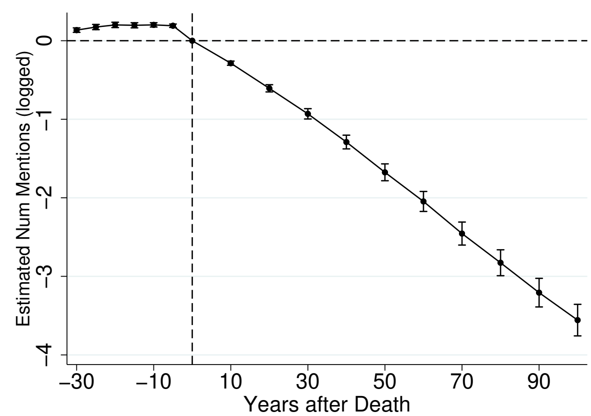

(b) Num Mentions on Google Books (logged)

Fig. S3. Alternative Measures of Artist Visibility

Notes: The figure shows estimated change in an artist's reputation years before/after death, similar to that in Figure S7 except not standardizing the reputation measure. All models use individual artist fixed effects and are clustered by artist.

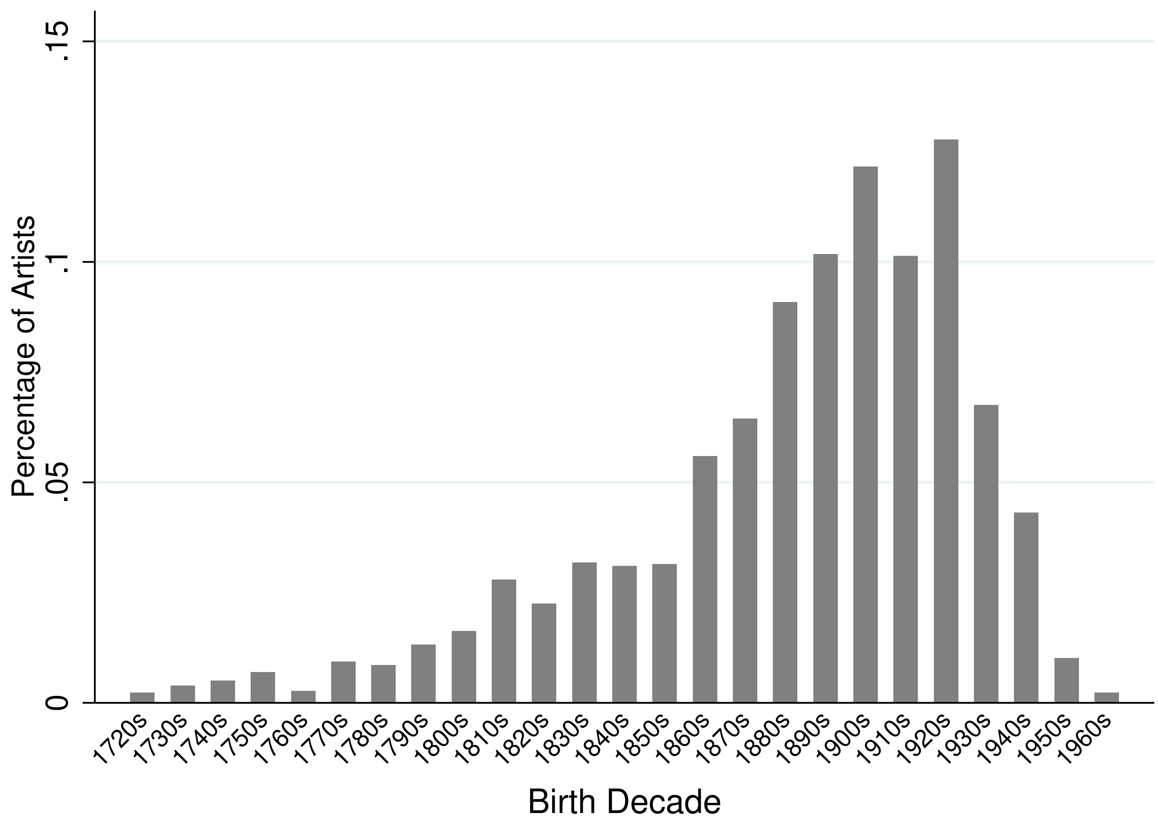

Fig. S4. Distribution of Artists' Birth Years

Notes: This figure shows the distribution of birth years for all artists in our sample.

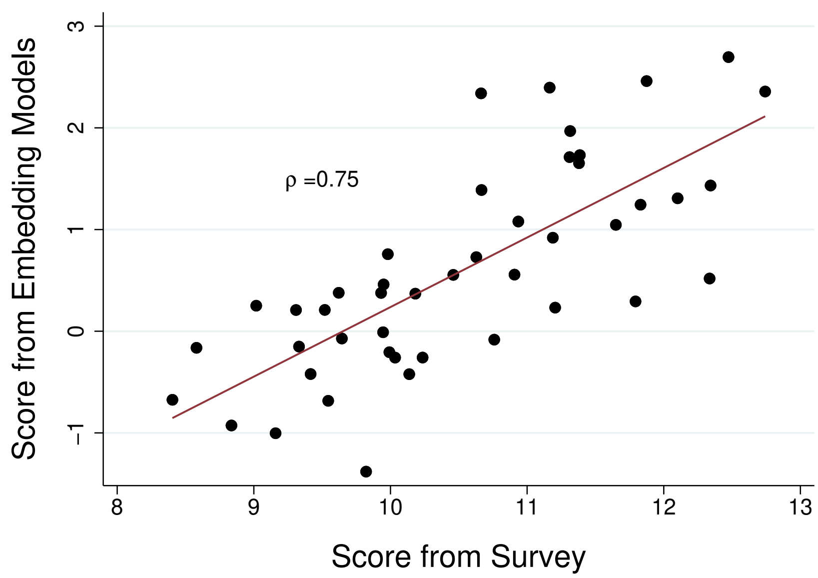

Fig. S5. Validating Reputation Measure using Survey Data

Notes: We picked the 50 artists in our sample with the most Wikipedia views to validate our embedding score. We conducted an online survey in which 500 respondents were asked to assess these artists' accomplishments. These scores were averaged per artist and plotted on the x-axis. We then took the average reputation score for these artists from our embedding models 2000–2020 and plotted it on the y-axis. The two scores have a correlation of 0.75.

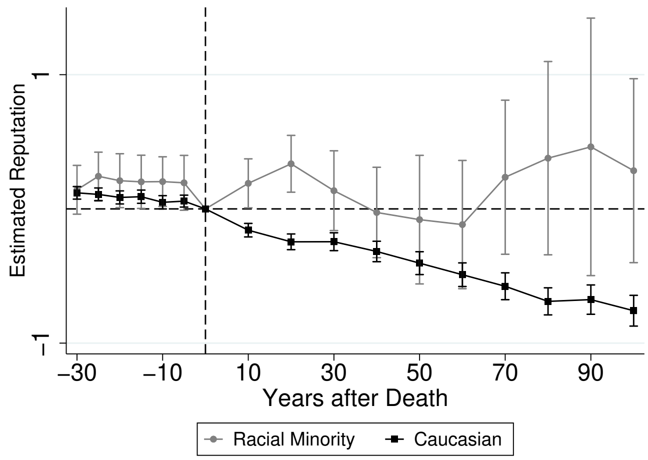

(a) Sorted by Race

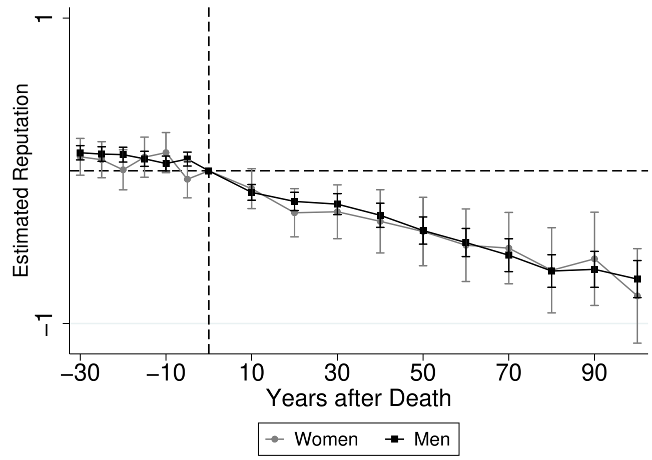

(b) Sorted by Gender

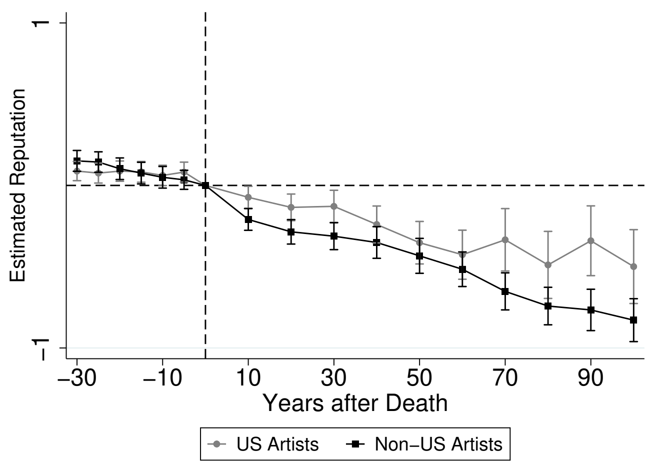

(c) Sorted by Country of Death

Fig. S6. Demographics and Posthumous Reputation

Notes: Figure (a) plots reputation change separately for White and non-White artists; figure (b) plots separately for women and men artists; figure (c) plots separately artists who died in the US and outside the US. In all figures, we set the y-axis to zero at the time of death. All models use individual artist fixed effects and are clustered by artist.

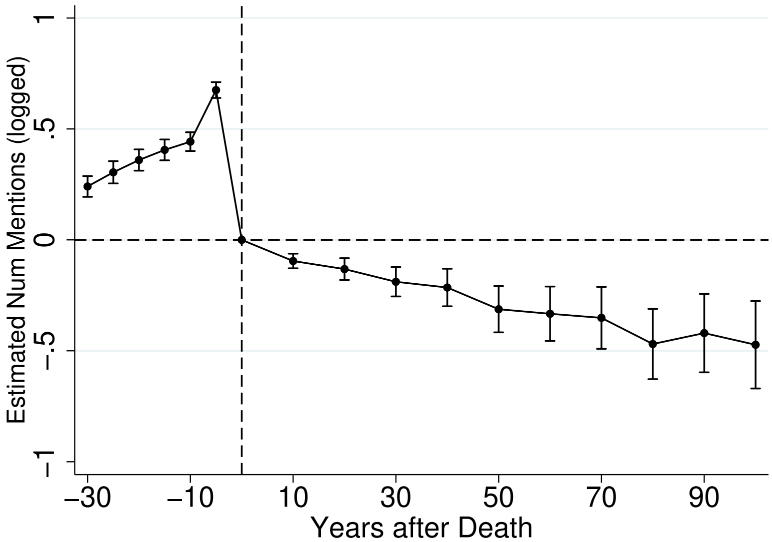

(a) All Artists

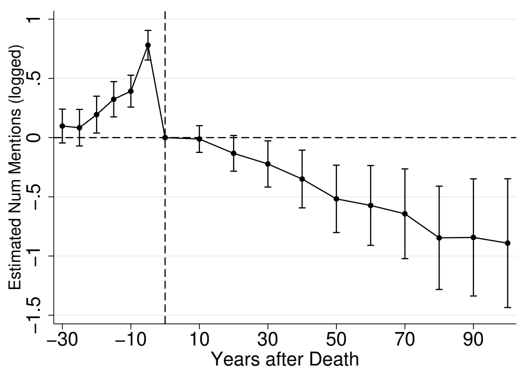

(b) Painters

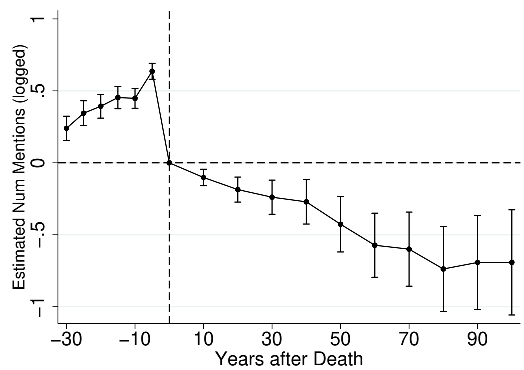

(c) Composers/Musicians

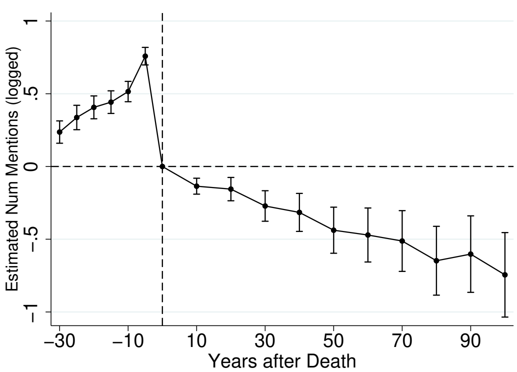

(d) Writers

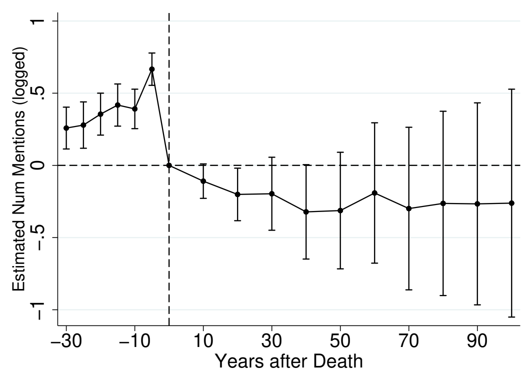

(e) Other Artists

Fig. S7. Predicted Visibility: Number of Newspaper Mentions (logged)

Notes: The figure shows estimated change in an artist's newspaper mentions years before/after death. The dependent variable is the number of times (logged) an artist's name is mentioned in all newspapers in a year. All models use individual artist fixed effects and year fixed effects and are clustered by artist.

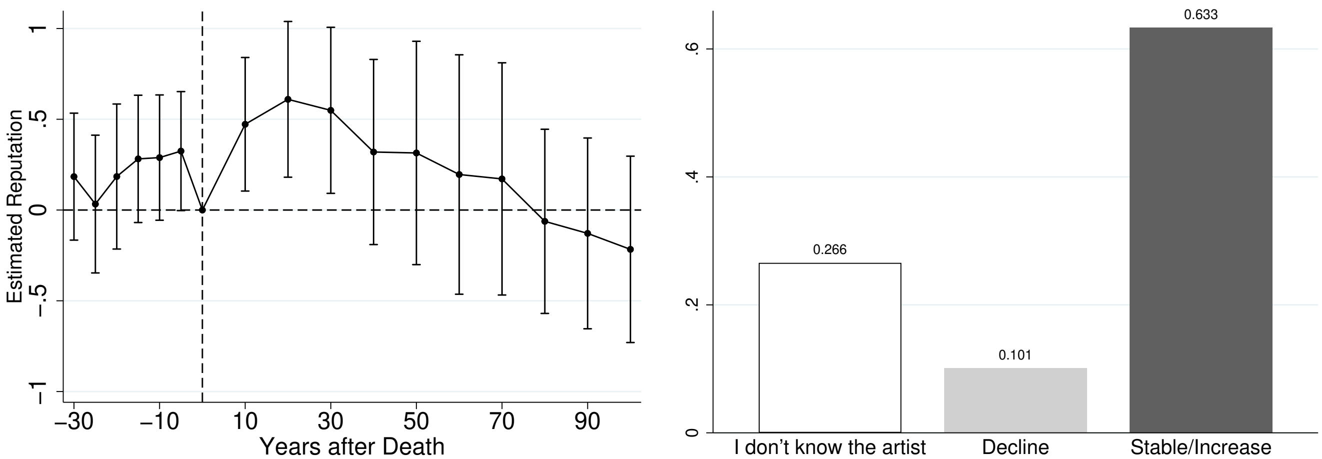

(a) Newspaper Corpus

(b) Prolific Survey

Fig. S8. Fig. S8. Comparing Popular View and Artistic Reputation in our Embedding Models

Notes: Figure (a) plots the reputation trend of 50 artists ranked in the 51–500 most-mentioned artists in our newspaper corpus 2016–2020. In Figure (b), we surveyed 500 respondents on Prolific and asked about their perceptions of these artists. Respondents were asked, "On a scale of 1 to 5, how do you perceive the change in the following artist's reputation after their death over a period of 100 years?" and were given the options, "I don't know the artist," "significant decline," "moderate decline," "stable," "moderate increase," and "significant increase."

Table S1. List of Newspapers

NO.	Newspaper Name	Availability	Circulation Rank
1	New York Times New York Times (Factiva)	1851_1934 1969_2020	3
2	Wall Street Journal Wall Street Journal (Factiva)	1889_1933 1979_2020	2
3	Los Angeles Times	1881_1950	4
4	Washington Post Washington Post (Factiva)	1877_1937 1977_2020	7
5	Chicago Tribune	1849_1950	8
6	Newsday	1940_1990	12
7	San Francisco Chronicle	1865_1922	13
8	Boston Globe	1872_2020	14
9	Philadelphia Inquirer	1860_2020	17
10	Detroit Free Press	1831_1999	20
11	St. Louis Post Dispatch	1874_2020	26
12	Baltimore Sun	1837_1932	33
13	Louisville Courier Journal	1830_2000	41
14	Cincinnati Enquirer	1841_2009	46
15	Hartford Courant	1764_1934	56
16	Nashville Tennessean	1812_2002	58
17	Austin American Statesman	1871_2020	59
18	Dayton Daily News	1898_1922	82
19	Christian Science Monitor	1908_1995	100+
20	New York Herald Tribune	1841_1962	N/A (suspended in 1962)

Notes: We list all newspapers used in our embedding models. Digitized newspapers were provided by Proquest and Factiva.

Table S2. Word Counts for Each Model Year

Model Year	Word Count	Model Year	Word Count
1796-1800	4.40E+06	1911-1915	1.50E+09
1801-1805	5.60E+06	1916-1920	2.00E+09
1806-1810	5.70E+06	1921-1925	1.70E+09
1811-1815	4.10E+06	1926-1930	1.60E+09
1816-1820	1.30E+07	1931-1935	1.40E+09
1821-1825	1.70E+07	1936-1940	1.20E+09
1826-1830	1.90E+07	1941-1945	7.30E+08
1831-1835	1.90E+07	1946-1950	9.10E+08
1836-1840	2.00E+07	1951-1955	5.00E+08
1841-1845	2.20E+07	1956-1960	5.90E+08
1846-1850	3.20E+07	1961-1965	7.10E+08
1851-1855	4.90E+07	1966-1970	1.40E+09
1856-1860	2.10E+08	1971-1975	1.60E+09
1861-1865	2.80E+08	1976-1980	1.80E+09
1866-1870	3.70E+08	1981-1985	1.60E+09
1871-1875	4.60E+08	1986-1990	1.30E+09
1876-1880	8.00E+08	1991-1995	5.90E+08
1881-1885	9.50E+08	1996-2000	6.90E+08
1886-1890	9.30E+08	2001-2005	3.70E+08
1891-1895	1.20E+09	2006-2010	2.50E+08
1896-1900	1.30E+09	2011-2015	1.40E+08
1901-1905	1.50E+09	2016-2020	8.30E+07
1906-1910	1.50E+09		

Table S3. Types of Artists Included

Occupation	Frequency	Percent
Architect	82	2.42
Artist	30	0.88
Comedian	15	0.44
Comic Artist	22	0.65
Composer	402	11.84
Dancer	42	1.24
Designer	29	0.85
Fashion Designer	10	0.29
Film Director	384	11.31
Musician	582	17.15
Painter	343	10.11
Photographer	50	1.47
Producer	26	0.77
Sculptor	47	1.38
Writer	1330	39.19
Total	3394	100

Notes: Our sample of artists came from the Pantheon 2.0 database, including all artists who have a Wikipedia page. We restricted our sample to artists (a) who are mostly known for their achievement in art, music, film, or writing; (b) who died no later than 2010; and (c) whose full name is mentioned at least five times in at least one of our embedding models.

Table S4. Keywords Lists

Positive Keywords		Negative Keywords	
creative	acclaimed	**failure**	common
talented	legendary	**unsuccessful**	unrecognized
successful	extraordinary	**untalented**	infamous
famous	exceptional	**unskilled**	mundane
skilled	recognition	**unoriginal**	obsolete
artistic	reputation	**uninspired**	
gifted	esteemed	**talentless**	
brilliant	excellence	**unimaginative**	
renowned	fame	**incompetent**	
accomplished	celebrity	**bland**	
inspired	star	ordinary	
amazing	outstanding	worthless	
respected	thriving	uninspiring	
admired	wonderful	unknown	
fantastic	virtuoso	unpopular	
genius	insightful	undistinguished	
inspiring	achiever	unremarkable	
skillful	classy	unskillful	
visionary	competent	incompetent	
impressive	creator	ungifted	
influential	prestigious	mediocre	
recognized	groundbreaking	obscure	
remarkable	masterful	insignificant	
success	noteworthy	terrible	
talent	phenomenal	worthless	

Notes: These reputation-related keywords were compiled from an online survey of 200 respondents, who were asked to list words that describe an accomplished and a failed artist. We tallied the words and listed the 50 most commonly mentioned positive words and the 30 most commonly mentioned negative words, in order of frequency with which they were mentioned.

Table S5. Predicting Artists' Reputation Score

	Full Sample	Split by Artist Type				
		Painters	Composers/Musicians	Writers	Directors	Other Artists
Years after Death	-0.00843***	-0.00624***	-0.00897***	-0.00897***	-0.00847**	-0.00740***
	(0.000449)	(0.00125)	(0.000943)	(0.000576)	(0.00271)	(0.00185)
Observations	26523	2888	7437	11192	2364	2642
R^2	0.451	0.484	0.425	0.442	0.373	0.452
Individual Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S6. Mechanism: the Mediating Role of Artist Visibility

	Artist Visibility	Artist Reputation	
Years after Death	-0.00172***	-0.00843***	-0.00569***
	(0.0000788)	(0.000449)	(0.000475)
Artist Visibility			1.591***
			(0.0963)
Observations	26523	26523	26523
R^2	0.811	0.451	0.464
Individual Leader Fixed Effects	Yes	Yes	Yes

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S7. Mechanism: the Mediating Role of Genre Change for Painters

	Artist Visibility	Genre Reputation	Artist Reputation		
Years after Death	-0.000156	-0.00705***	-0.00510*	-0.00488	-0.00333
	(0.000347)	(0.00105)	(0.00253)	(0.00251)	(0.00248)
Artist Visibility				1.423***	1.476***
				(0.422)	(0.418)
Genre Reputation					0.219**
					(0.0801)
Observations	1217	1217	1217	1217	1217
R^2	0.842	0.651	0.489	0.496	0.501
Individual Leader Fixed Effects	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

References

1. T Mikolov, Wt Yih, G Zweig, Linguistic regularities in continuous space word representations in *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. pp. 746–751 (2013).
2. J Pennington, R Socher, CD Manning, Glove: Global vectors for word representation in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. pp. 1532–1543 (2014).
3. ZS Harris, Distributional structure. *Word* **10**, 146–162 (1954).
4. JC Eichstaedt, et al., Closed-and open-vocabulary approaches to text analysis: A review, quantitative comparison, and recommendations. *Psychol. Methods* **26**, 398 (2021).
5. V Kulkarni, R Al-Rfou, B Perozzi, S Skiena, Statistically significant detection of linguistic change in *Proceedings of the 24th International Conference on World Wide Web*. pp. 625–635 (2015).
6. N Garg, L Schiebinger, D Jurafsky, J Zou, Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proc. Natl. Acad. Sci.* **115**, E3635–E3644 (2018).
7. AC Kozlowski, M Taddy, JA Evans, The geometry of culture: Analyzing the meanings of class through word embeddings. *Am. Sociol. Rev.* **84**, 905–949 (2019).
8. S Ji, N Satish, S Li, PK Dubey, Parallelizing word2vec in shared and distributed memory. *IEEE Transactions on Parallel Distributed Syst.* **30**, 2090–2100 (2019).