Does relative age affect fame? Ask Wikipedia

Pablo A. Peña & Seth Stephens-Davidowitz

EDUCATION ECONOMICS
https://doi.org/10.1080/09645292.2021.1880548

Does relative age affect fame? Ask Wikipedia

Pablo A. Peña [a] and Seth Stephens-Davidowitz [b]

[a]Kenneth C. Griffin Department of Economics, University of Chicago, Chicago, IL, USA; [b]Independent Researcher

ABSTRACT
We analyze whether age relative to school classmates affects the likelihood of becoming famous. We measure such likelihood as the ratio of Wikipedia entries to births, by state and date of birth, among people born in 1969–1988 in the US. Using a reduced-form Regression Discontinuity Design, we find evidence that men born after the Kindergarten cutoff date (i.e., relatively older) are roughly between 5 and 10 percent more likely to become famous, by Wikipedia standards, in comparison to those born before the cutoff (i.e., relatively younger). We don't find evidence of a similar effect for women.

ARTICLE HISTORY
Received 10 July 2020
Accepted 13 January 2021

KEYWORDS
relative age; school-starting age; Wikipedia; fame

JEL Classification
I21; J11

1. Introduction

Over the last fifteen years, dozens of studies on the effects of relative age—how old a person is with respect to her school classmates—have been published in academic journals.[1] Most studies of relative age take a sample of subjects of different relative ages and compare outcomes of interest. The majority of those studies analyze outcomes among schoolchildren (test scores, ADHD diagnosis, socioemotional skills scales, etc.). Only a few focus on long-term effects, and some have found that relative age affects educational attainment (Kawaguchi 2011; Fredriksson and Ockert 2013; Peña 2017), occupation (Grenet 2011; Zweimüller 2013), professional performance (Bai et al. 2019), and earnings (Black, Devereux, and Salvanes 2011; Zweimüller 2013; Peña 2017).

Another type of relative-age study takes an indirect approach and looks at overrepresentation of relatively older people among a group deemed 'successful' in highly visible professional fields.[2] For instance, Muller and Page (2016) analyzed being elected to the US Congress as an indicator of success in politics. Fukunaga, Taguri, and Morita (2013) looked at receiving a Nobel prize in physics, chemistry, or physiology/medicine as an indicator of success among scientists. Du, Gao, and Levi (2012) examined being the CEO of a Standard & Poor's 500 company as an indicator of success among businesspeople. Those three studies found evidence of overrepresentation of individuals who were old relative to their school classmates.

Overrepresentation studies have several shortcomings. First, they define success in a professional field in a narrow way and therefore have small samples—no more than a few hundred individuals. Second, when separately looking at different professional fields, they may miss between-field substitution—from politics to business, from arts to sciences, etc. Third, the numerous ways one can define success may lead to publication bias. Researchers could easily gather information of successful professionals under different definitions—e.g. members of the National Academy of Sciences, Oscar-nominated movie directors, McArthur 'geniuses,' high-ranking military—to conduct overrepresentation analyses. It's conceivable that researchers have already done that using many different definitions of success, found no effect, and decided against writing or submitting a paper. [3] Thus, it's possible that published overrepresentation studies suffer from selection into

CONTACT Pablo A. Peña pablo@uchicago.edu

definitions of success where relative age effects are significant, leading to a biased view of the role of relative age. To be clear, we don't claim that published overrepresentation studies are biased. Our claim is that there may be studies that were not published *because* they didn't find overrepresentation. We don't provide evidence of that, but we believe it's a real possibility worth pondering by researchers on the topic.

In this study we try to overcome those shortcomings. We examine many professional fields simultaneously with a large sample leveraging a general indicator of success: fame. Of course, fame and success are not the same. Some people may be famous for the wrong reasons—e.g. committing a crime. At the same time, the threshold for fame may be quite different across contexts like sports, science, politics, or entertainment. Although we recognize the noise that fame involves, we also believe no proxy of success is perfect.

Famous people are very likely to have personal entries in Wikipedia. Those entries usually indicate date and state of birth. We combine that information with daily data from birth certificates in the US in 1969–1988 and compute the probability of having a Wikipedia personal entry by date and state of birth.[4] We analyze how that probability changes around the cutoff date for Kindergarten enrollment eligibility using a Regression Discontinuity Design. Before proceeding to the empirical analysis, we discuss a few points about the odds of success.

2. The odds of success

Notable success in any field is by definition a low probability event. How responsive that probability is to a better performance induced by relative age depends on the size of the gain as well as the distribution of performance. For simplicity, we assume the gain produced by relative age is fixed and consider different distributions of performance. Thus, the exact same gain in performance affects the probability of success differently.[5]

Let us define success as having a performance above a threshold denoted by a. The probability of success is $P[X^* \geq a]$, where the random variable X^* represents performance. Let R denote the relative age of a person, defined as the difference in years between the ages of that person and her youngest possible classmate given the cutoff date for enrollment eligibility. Assume performance can be decomposed into a component due to relative age, denoted by $\Delta \cdot R$, where $a > \Delta > 0$, and a random component X explained by other factors (e.g. intelligence, artistic talent, athletic ability, or looks):

$$X^* = X + \Delta \cdot R \tag{1}$$

Imagine two identical individuals (i.e. same X) born one day apart but at different sides of the cutoff date for enrollment eligibility. For the person born before the cutoff we have $R = 0$, whereas for the person born after the cutoff we have $R = 1$. The probability of success for the relatively young is $P[X \geq a]$, and the probability of success for the relatively old is $P[X + \Delta \geq a]$, or equivalently $P[X \geq a - \Delta]$. If we compare their odds of being successful, we obtain the ratio:

$$\frac{P[X \geq a - \Delta]}{P[X \geq a]} \geq 1 \tag{2}$$

We don't know a priori how close such odds ratio is to one, or whether it is increasing or decreasing in the threshold a. Holding constant the advantage Δ induced by relative age, the odds ratio depends on the threshold a and the distribution of X. Some probability distributions, like the normal, have thin tails that make the odds ratio increasing in a. Others, like the power-law distribution, have fat tails that make the odds ratio decreasing. Figure 1 presents numerical examples of odds ratios for three distributions: normal, exponential, and power law. In the exponential case, the odds ratio is $e^{\lambda\Delta}$, where $\lambda > 0$, and therefore it is constant across different levels of the performance threshold a. In the case of the power-law distribution, the odds ratio is $[a/(a - \Delta)]^{\alpha}$, which is decreasing in a. For the normal distribution we computed the odds ratio numerically using a

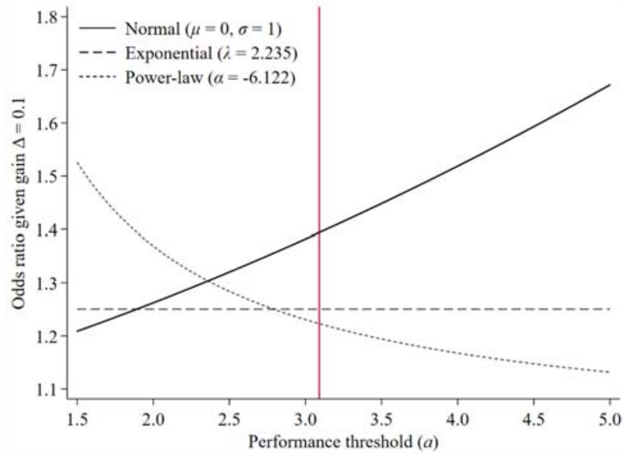

Figure 1. Odds ratios for different probability distributions of performance. The graph presents odds ratios according to equation (2) for three different performance distributions. At a performance threshold of 3.09, the probability of success is the same across the three distributions (10/10,000). However, the odds ratios and their relationship to the threshold differ (increasing in the normal case, constant in the exponential case, and decreasing in the power-law case).

mean equal to zero and a standard deviation equal to one. We chose the parameters α and λ so that $P[X \geq 3.09] = 0.001$ for the three distributions. As we describe below, that's the average probability of having a Wikipedia entry among males (10 per 10,000 births). We assumed an arbitrary gain of $\Delta = 0.1$.[6] Although the probability of success is the same given a threshold of 3.09 (illustrated by the vertical line), the odds ratios are different at that threshold. More importantly, ceteris paribus shifts in the threshold would impact the odds ratio differently, depending on the distribution. In other words, the effect of relative age on success responds differently to the stringency of the definition of success.

If performance is normally distributed and relative age boosts it uniformly, then we expect an increasing effect of relative age as we use more stringent definitions of success. Since having a Wikipedia entry is less stringent than winning a Nobel prize or being CEO of a Standard & Poor's 500 company, in that case we would expect smaller effects using Wikipedia entries as an indicator of success. If the distribution is exponential, we would expect the same effect. If performance follows a power-law distribution, we would expect a larger effect with our less stringent definition. We don't rule out a priori any possibility. Lastly, it should be clear there is no reason to expect similar estimates of the effects of relative age among overrepresentation studies. In all likelihood, they are based on different distributions, thresholds, and gains induced by relative age.

3. Data

Our data comes from three sources. The first is Wikipedia, 'an online free-content encyclopedia project […] based on a model of openly editable content.' [7] Wikipedia has a policy of allowing personal pages of 'notable' people only.[8] Being user-generated, Wikipedia likely involves differences or biases in the coverage of notable people by gender, geography, professional field, and other factors.[9] Our analysis inherits those coverage differences or bias, but we don't think there are reasons why they would invalidate a priori our identification strategy. Potential biases should be the same at each side of the discontinuities we analyze.

We downloaded from Wikidata the information from all Wikipedia personal pages of individuals born in the US in 1969–1988, which amounted to 56,791 unique records.[10] The place of birth is recorded in an irregular way in up to four fields. By parsing those fields, we identified the state of birth of 55,852 individuals (the rest had insufficient information). Of them, 51,394 had exact date of birth. All of those records also state the person's gender. Transgender individuals were classified according to their gender at birth. A group of 25 individuals whose gender wasn't specified (i.e. it was stated as 'intersex,' 'non-binary,' or 'fluid') were excluded from the analysis. In 586 cases, there was conflicting information on the date or state of birth or the gender. In those cases, we randomly selected one single observation. We only kept in our sample individuals born in one of the 50 states or the District of Columbia, leaving us with 50,568 individuals. Table 1 summarizes how we arrived at the sample of analysis starting from the initial sample downloaded from Wikidata.

Our second data source is the Natality Data from the National Vital Statistics System of the National Center for Health Statistics, which we downloaded from the website of the National Bureau of Economic Research (NBER). The data cover births that occurred in 1969–1988 in the US. According to the NBER website, 'Prior to 1972, data are based on a 50-percent sample of birth certificates from all states. Beginning in 1972, data are based on a 100-percent sample of birth certificates from some states and on a 50-percent sample from the remaining States. The number of States from which 100 percent of the records are used has increased from 6 in 1972 to all States and the District of Columbia in 1985.' [11] We expanded the number of births accordingly, to reflect the total number of births in each state on each date. Table 2 presents the total number of births. In each of the 7,305 days in the 1969–1988 period there was at least one person born in each state. Thus, in principle we have a total of 372,555 (= $51 \times 7{,}305$) unique combinations of state and date of birth in our sample. However, when births are split by gender, there are a few cases without any births, bringing the total of unique combinations of state and date of birth with non-zero births to 372,466 for females and 372,503 for males.[12]

We combined the data from Wikipedia entries and births by date and state of birth and gender.[13] The results are summarized in Figure 2, which displays the number of births and Wikipedia entries by quarter of birth. In every quarter of birth there are roughly between 300 and 500 thousand births of each gender, and between 100 and 200 female Wikipedia entries and between 400 and 600 male

Table 1. Description of Wikipedia entries data.

	Individuals
Initial sample (born in US in 1969–1988)†	56,791
With state of birth	55,852
With exact date of birth	51,394
With gender specified‡	51,369
With conflicting information*	586
Born in one of the 50 states or DC	50,568
Final sample	50,568
	%
With conflicting information ÷ Final sample	1.2
Final sample ÷ Initial sample	89.1
Female individuals ÷ Final sample	23.1

†Downloaded from Wikidata. ‡Excludes non-binary, gender-fluid, and intersex individuals. Transgender individuals were categorized according to gender at birth. *Individuals with more than one state of birth, more than one date of birth, or more than one gender.

Table 2. Description of the Natality Data.

	Individuals
Birth certificates in 50 states or DC, 1969–1988†	56,545,326
With exact date of birth	56,541,532
Total births‡	70,265,276
	%
Female births ÷ Total births	48.4

†Natality Data from the National Vital Statistics System of the National Center for Health Statistics downloaded from the NBER website. ‡At the beginning of the period, the data include a random sample of 50% of the birth certificates. For those years, the number of births was multiplied by two to represent their total number.

Wikipedia entries. Our outcome of interest is the number of Wikipedia entries per 10,000 births for each date of birth, by state. On average, among females there are roughly 3 entries per 10,000 births, whereas for males there are 10 entries per 10,000 births.

There are important differences across states in the number of Wikipedia entries per 10,000 births. Figure 3 shows the number of entries per 10,000 births across all years of birth by gender and state. Among males, the maximum difference is observed between West Virginia (5.4) and the District of Columbia (20.4). Among females, the maximum difference is between West Virginia (1.2) and Alaska (7.6).

Our third data source is the studies of Bedard and Dhuey (2007, 2012), which provide a compilation of cutoff dates for Kindergarten enrollment eligibility from the fifty states in the school years 1964–2005.[14] We complemented that source with information for the District of Columbia from the D.C. Code Ann. § 38–202. We excluded from our sample combinations of state and school year without a cutoff for Kindergarten enrollment in the period of analysis. Colorado, Massachusetts, and New Jersey (marked in red in Figure 3) fall in that category and were therefore excluded (the cutoff in those states was determined by local educational authorities). To simplify the analysis, throughout we express cutoffs *in terms of dates of birth*. To do it, we subtract five years from each actual cutoff date. For instance, the cutoff date to start Kindergarten in Florida in the school year 1990 was September 1. Thus, children could enroll in Kindergarten in the Fall of that year only if they were five years old by September 1, 1990. The same cutoff expressed in terms of dates of birth would be September 1, 1985, that is, in order to be permitted to enroll in Kindergarten in the Fall of 1990, a child must have been born before September 1, 1985.

For each person born in 1969–1988 in state *i* we define the 'relevant cutoff' as the closest cutoff date (expressed in terms of birthdates) in state *i* to the date when that person was born. This definition allows us to focus on birthdates around the cutoffs. As an example, Figure 4 shows the relevant cutoffs for Florida. The step function illustrates how the relevant cutoff depends on the date of birth. In absence of shifts in cutoffs, people born in a ±182-day neighborhood of a cutoff would share the same relevant cutoff. That is the case of Florida prior to the birth year 1975 (when the cutoff was Jan. 1) or after the birth year 1978 (when the cutoff was Sep. 1). In the birth years 1975–1978 the cutoff was shifted one month per year (from Jan. 1 to Sep. 1), making the 'steps' shorter than twelve months and not symmetric around the cutoff date.

We defined a 'cohort' as the group of individuals born in the same state who share the same relevant cutoff (e.g. the same 'step' in Figure 4). We calculated the distance in days between the relevant cutoff and the birthdate of each person (birthdate minus cutoff.) Some states had cutoffs too close to the end or the start of the calendar year, making it impossible to use all births from the first and last years of the sample. We kept only 'relatively balanced' cohorts, that is, cohorts for which the relevant cutoff leaves at least 90 birthdates at each side. As a result, our sample of analysis includes 48 states, 892 unique combinations of state and cutoff, and 317,744 combinations of state and date of birth for females and 317,775 for males. In that sample, we have 29,110,424 female births and 30,621,253 male births. In terms of Wikipedia entries per 10,000 births, on average we have 3.1 for females and 9.7 for males. Table 3 summarizes the most relevant aspects of the sample once we combined births, Wikipedia entries, and Kindergarten enrollment cutoffs.

We are interested in the variation of Wikipedia entries per 10,000 births around the cutoffs. Figure 5 presents the total number of births and the number of Wikipedia entries per 10,000 births by distance to the relevant cutoff only for the cohorts that had 180 birthdates at each side of the cutoff. For both genders the number of births is rather flat and, most importantly, doesn't show any abrupt change at the cutoff. This is consistent with the results of Dickert-Conlin and Elder (2010), who analyzed whether there was strategic timing of births around cutoff dates and found no evidence of such behavior. The number of Wikipedia entries per 10,000 births shows more variation and—to the naked eye—no clear discontinuity.

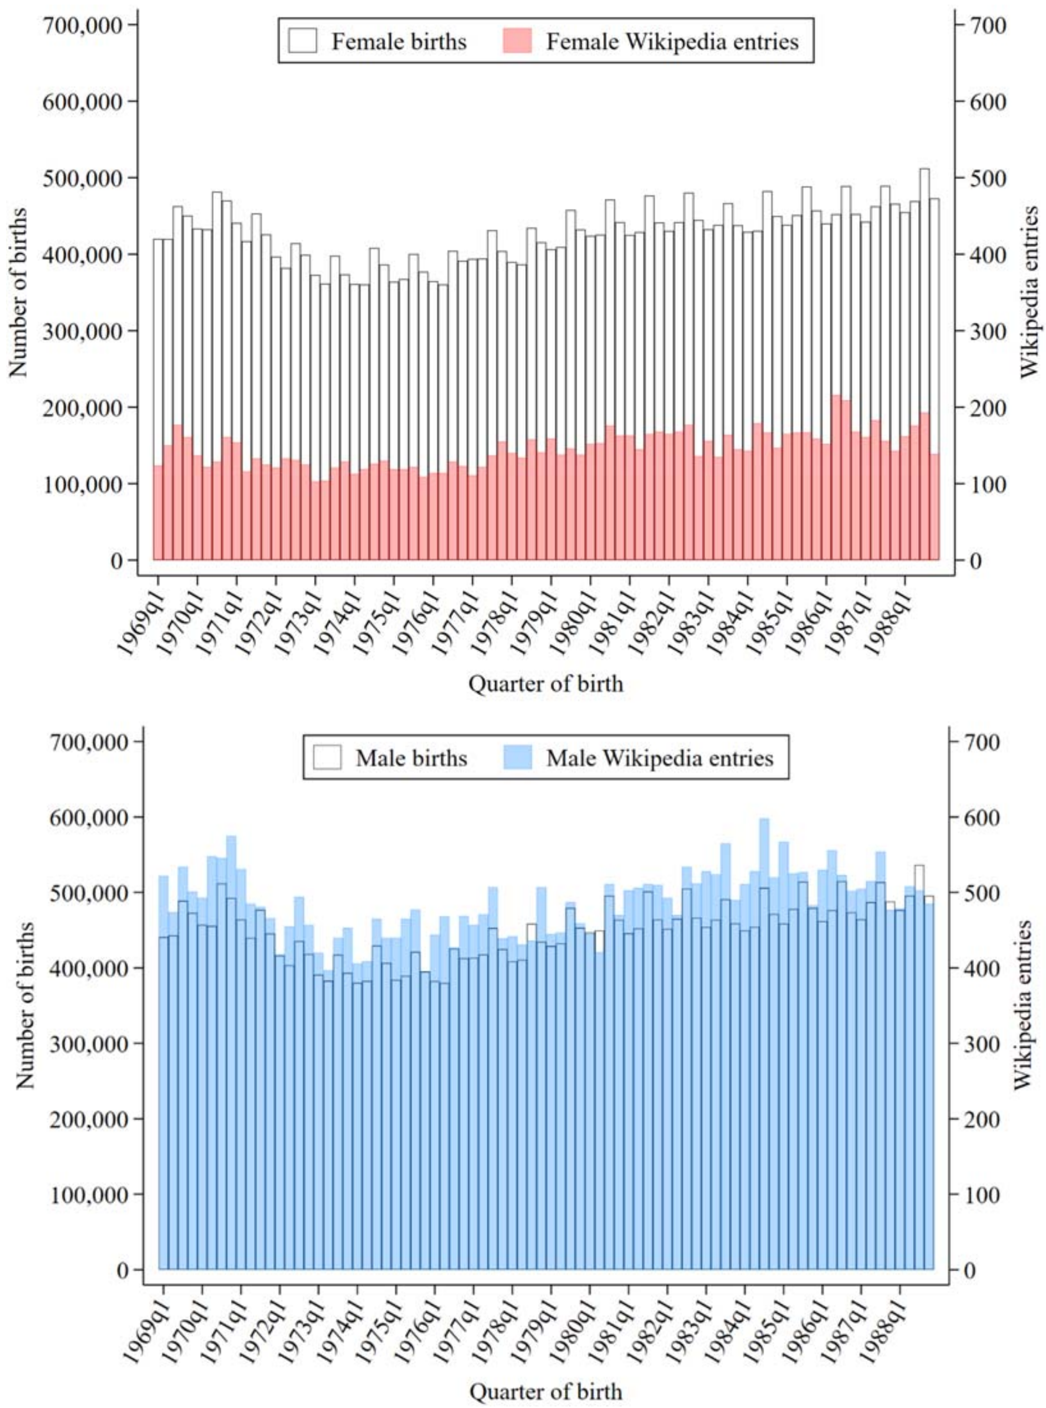

Figure 2. Number of births and Wikipedia entries by quarter of birth. It includes the 50 states and the District of Columbia. Source: Natality Data from the National Vital Statistics System of the National Center for Health Statistics, and Wikidata.

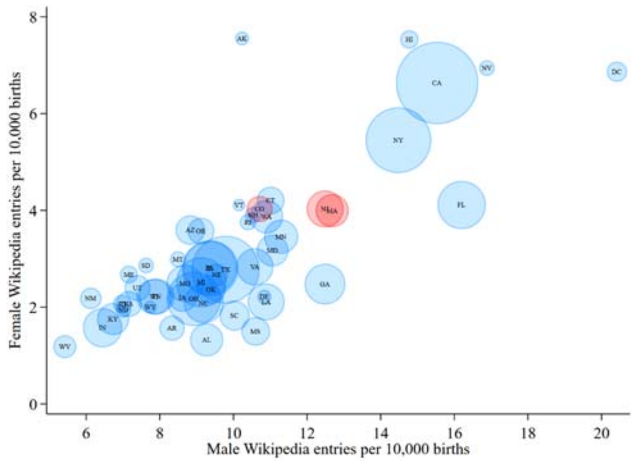

Figure 3. Wikipedia entries per 10,000 births by gender and state of birth. The size of the markers represents the total number of births (male and female) in the period 1969–1988. Source: Natality Data from the National Vital Statistics System of the National Center for Health Statistics, and Wikidata.

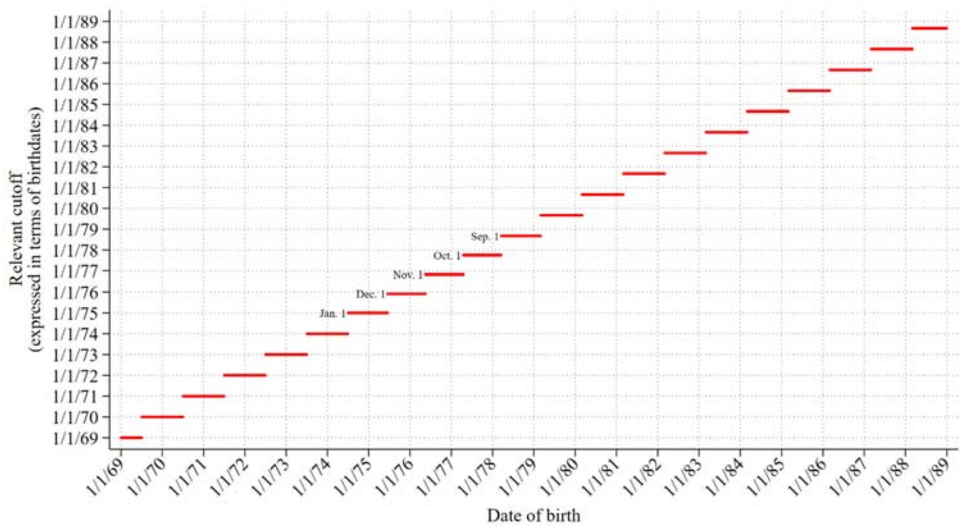

Figure 4. Example of relevant cutoffs in Florida by date of birth. The cutoff in a given school year is expressed in terms of birthdates by subtracting five years. The relevant cutoff is defined as the cutoff (expressed in terms of birthdates) that is closest to the date of birth.

4. Empirical strategy

Other than through relative age, birthdates are unlikely to affect the odds of becoming famous by Wikipedia standards in a discontinuous way. We adopted a reduced-form approach and estimated the effect of the discontinuity at the cutoff on the number of Wikipedia entries per 10,000 births.[15] To do it, we aligned all cutoffs in the running variable k, which is the number of days between a birthdate and the relevant cutoff (birthdate minus cutoff). Our econometric specification is:

$$y_{ijk} = \theta_{ij} + \beta d_k + P(k) + d_k Q(k) + \delta X_{ijk} + \varepsilon_{ijk} \tag{3}$$

Where y_{ijk} is the number of Wikipedia entries per 10,000 births for cohort i (individuals sharing the same relevant cutoff) born in state j, and for which the distance between the cutoff and the birthdate

Table 3. Description of the sample combining births, Wikipedia entries, and Kindergarten enrollment cutoffs.

	Females				Males			
	Number of state-birthdate combinations	Number of births	Wikipedia entries per 10,000 births Mean (s.d.)		Number of state-birthdate combinations	Number of births	Wikipedia entries per 10,000 births Mean (s.d.)	
Initial sample (50 states and DC, 1968–1988)	372,466	34,242,255	3.2	(33.9)	372,503	36,023,021	10.1	(54.4)
Excluding state and date of birth combinations without a cutoff†	326,252	30,064,547	3.0	(34.2)	326,288	31,624,965	9.8	(55.2)
Only cohorts with relative balance around the cutoff‡	317,744	29,110,424	3.1	(34.2)	317,775	30,621,253	9.7	(54.8)

Source of cutoff information: Bedard and Dhuey (2007, 2012). †Colorado, New Jersey, and Massachusetts are excluded because they did not have a state-wide cutoff. Other states are partially excluded because they did not have a cutoff some years. ‡To have relative balance, a cohort must include at least 90 birthdates immediately before the relevant cutoff, and 90 birthdates immediately after it.

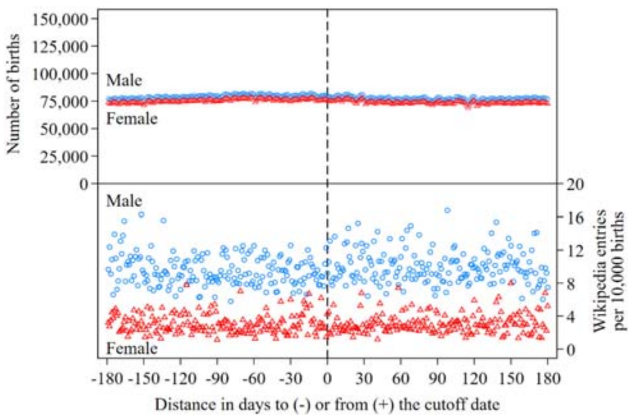

Figure 5. Number of births and Wikipedia entries per 10,000 births by distance to the relevant cutoff. The graph presents different cohorts aligned by distance to the cutoff. For illustrative purposes, it includes only cohorts with 180 birthdates at each side of the cutoff (they represent 92.3% of the births in the sample of analysis).

is k days. The variable d_k is a dummy indicating a birthday after the cutoff ($k > 0$). The coefficient θ_{ij} is a fixed effect for cohort i in state j. $P(k)$ and $Q(k)$ stand for polynomials in the running variable k. We used polynomials of different degrees (zero to three). X_{ijk} is a vector of covariates that include average age of the mother at birth, birthweight of the newborn, and day-of-week dummies. Lastly, ε_{ijk} is the error term, which we clustered by the combination of state and cohort. The coefficient of interest is β, which is the difference in Wikipedia entries per 10,000 births between people born at different sides of the cutoff. We analyzed different bandwidths around the cutoff, from ±15 to ±180 days, in increments of five days at each side. We estimated separate regressions for females and males.

Our estimates of β are comparable to the estimates from overrepresentation studies but not to estimates of causal effects from studies that rely on Instrumental Variable methods. IV methods help address potential biases resulting from redshirting—i.e. voluntarily postponing school enrollment—and grade retention. Since we don't observe year of enrollment in Kindergarten or school, we cannot apply such methods. By adopting a reduced-form approach, if a large fraction of children born right before the cutoff are redshirted, then our estimates may understate the causal effect of relative age. This issue would be more important if redshirting occurred among children whose background makes them more likely to be famous in the first place—perhaps having highly-educated or wealthy parents. Redshirting is more prevalent among children born right before the cutoff and it increases with parental education (Dickert-Conlin and Elder 2010). It is therefore likely that our reduced-form estimates are biased downward in comparison to the causal effect. However, our sample includes cohorts for whom redshirting was less common (Deming and Dynarski 2008), possibly limiting the magnitude of such downward bias.

It's important to acknowledge the role of measurement error in our estimation. First, the state of birth isn't always the state where people attend school—people move across states. Thus, there is noise in the running variable that can result in attenuation bias. Second, Wikipedia entries are noisy indicators of success themselves, perhaps due to different thresholds applied across different professional fields. That noise would reduce the accuracy of our estimates.

Birth certificates provide information about the age of the mother and the birthweight of the newborn.[16] As covariates in our regression model, we use average age of the mother and average birthweight calculated for each combination of date and state of birth. There are no missing values for average age of the mother, and only a handful of observations (combinations of state and date of birth) are missing average birthweight (10 for males and 17 for females). Table 4 presents averages for those variables and the number of daily births in a ±30-day neighborhood of the cutoff. There is hardly any economically relevant difference between the averages for the 30 days before and after the cutoff, supporting the validity of the comparison around the cutoff. Table 4 also shows tests for smoothness of those covariates around the cutoffs using different linear polynomials at each side of the cutoff and including day-of-week and state-and-cohort fixed effects. Although the estimates at the discontinuity are significant in some cases, their magnitudes are minuscule and, if anything, their signs suggest positive sorting into birthdates *before* the cutoff. On average, people born after the cutoff had younger mothers and weighed less at birth than people born before the cutoff. Similarly, there are fewer daily births after the cutoff. [17] Based on those averages alone we would hypothesize that individuals born before the cutoff would be more likely to be famous, for their mothers are older, they were healthier at birth, and their births were more likely intentional.

5. Results

Figures 6 and 7 present the estimates according to specification (3) using polynomials of different degrees and different bandwidths around the cutoff. Figure 6 shows the results of unweighted regressions, that is, giving each date and state of birth the same importance. In Figure 7 observations are weighted by the number of births in each combination of date and state of birth.[18] All regressions include fixed effects for each combination of cohort and state, day-of-week fixed effects, and averages of maternal age and birthweight. The shaded areas represent the 95

Table 4. Smoothness of covariates and births across the cutoff.

		Age of mother (years)		Birthweight (grams)		Number of daily births	
	Obs.†	Mean (std. dev.) [std. error]					
Females							
30 days before cutoff	26,757	24.96	(1.34)	3,387.0	(138.0)	98.73	(110.55)
30 days after cutoff	26,757	24.96	(1.37)	3,391.0	(140.6)	97.48	(110.68)
Discontinuity at cutoff‡		−0.05	[0.02]	−6.9	[2.2]	−1.13	[0.25]
Males							
30 days before cutoff	26,754	24.96	(1.36)	3,262.4	(130.2)	93.87	(105.04)
30 days after cutoff	26,751	24.97	(1.37)	3,265.2	(131.8)	93.10	(105.77)
Discontinuity at cutoff ‡		−0.03	[0.02]	−2.9	[2.1]	0.00	[0.00]

†Each observation corresponds to a unique combination of date and state of birth. ‡Regression Discontinuity estimates obtained from separate regressions ±30 days around the cutoff that include linear polynomials in the distance to the cutoff, with different slopes at each side of the cutoff. They also include day-of-week fixed effects and state-year fixed effects. Robust standard errors clustered by state-cohort in brackets.

percent confidence intervals using standard errors clustered by the combination of cohort and state. Since standard errors increase closer to the cutoff, we present results starting at different bandwidths, depending on the degree of the polynomials (±15 days for the constant specification, ±30 days for the linear specification, ±45 days for the quadratic specification, and ±60 days for the cubic specification). There is a noticeable difference between estimates for males and females. All point estimates for males in Figure 6 are positive, in some cases significant at 95% confidence, and have a magnitude around one additional Wikipedia entry per 10,000 births. In contrast, the estimates for females vary in sign and aren't significant in general. The estimates from the weighted regressions in Figure 7 are qualitatively similar to those in Figure 6. For males, they are positive in all cases, and they are significant at 95 percent confidence with the constant and linear specifications. In all cases the magnitude is roughly around 0.5 Wikipedia entries per 10,000 births. For females, most estimates are negative, and none is statistically significant.

We interpret the results in Figures 6 and 7 as weak evidence of a relative age effect on fame. Using the average of ten Wikipedia entries per 10,000 births among males, a point estimate of one additional entry implies an increase in the likelihood of being famous by Wikipedia standards of 10 percent. Using the results from the weighted regressions, the increase would be equivalent to roughly 5 percent. Thus, the point estimates for males are not negligible. In the case of females, there is no evidence of an effect of relative age on fame. The estimates

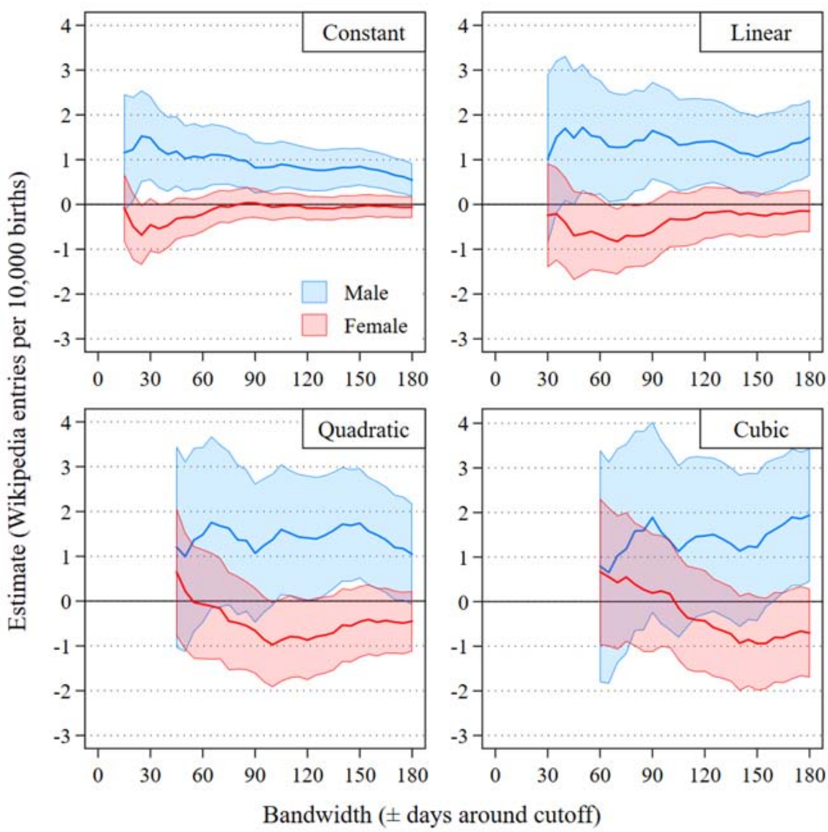

Figure 6. Reduced-form Regression Discontinuity estimates of the effect of relative age on the probability of becoming famous, unweighted. The thick line represents the estimates of the discontinuity in Wikipedia entries per 10,000 births for different bandwidths. The shaded areas represent 95% confidence intervals using standard errors clustered by cohort and state of birth. Each panel presents a different degree of the polynomial in the running variable.

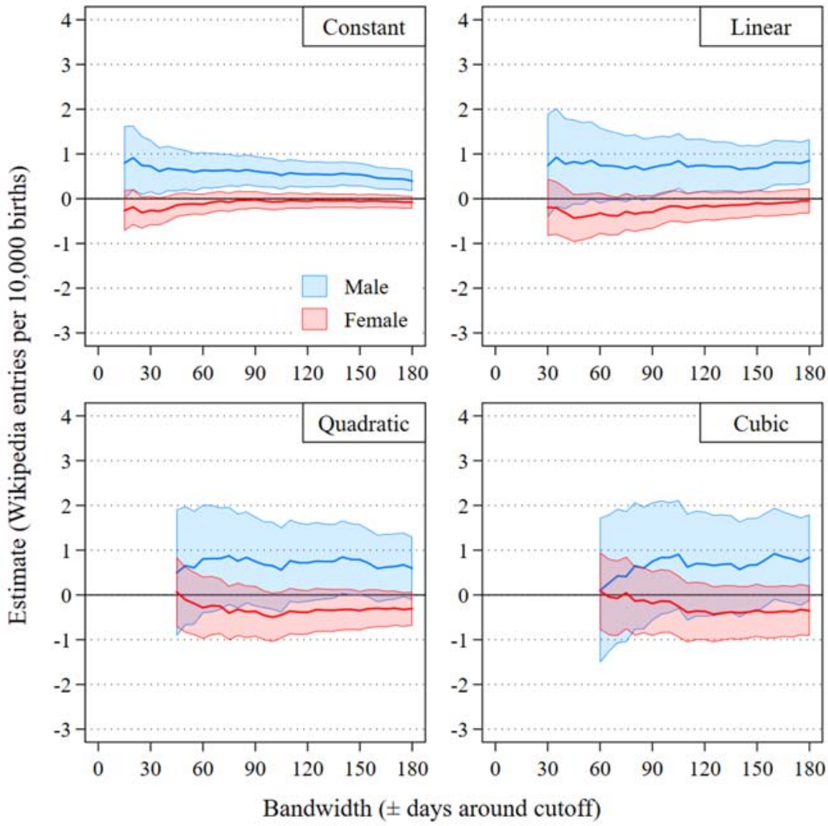

Figure 7. Reduced-form Regression Discontinuity estimates of the effect of relative age on the probability of becoming famous, weighted by the number of births. The thick line represents the estimates of the discontinuity in Wikipedia entries per 10,000 births for different bandwidths. The shaded areas represent 95% confidence intervals using standard errors clustered by cohort and state of birth. Each panel presents a different degree of the polynomial in the running variable. Observations are weighted by the number of births in each state and date of birth.

don't have a consistent sign, many point estimates are negative, and almost all are statistically insignificant.

6. Discussion

We looked at Wikipedia entries for people born in the US and found weak evidence of an effect of relative age on fame. The effect is limited to men, and the point estimates are positive and non-negligible. For males, being born right after the cutoff date for Kindergarten enrollment increases the likelihood of having a Wikipedia entry roughly 5–10 percent. For females, there is no evidence of an effect.

If we interpret having a Wikipedia entry as a noisy indicator of fame, the standard errors of the estimate of the effect of relative age on fame are overstated. In absence of noise, the significance of the results would increase. Also, since the state where people attended school may be different from their state of birth, the relevant cutoff is likely measured with error for many individuals. Measurement error in the explanatory variable leads to attenuation bias, understating the effect of relative age and reducing the significance of the estimates. If we could observe and use in our estimation the state where people attended school, we would expect estimates of greater magnitude and significance. Lastly, for the same reason, imperfect compliance with the cutoffs is likely biasing downward our reduced-form estimates. Although academic redshirting was less frequent among people born in the years analyzed than it is today, based on anecdotal evidence it seems that bending the rules to enroll in school earlier than supposed was more common in the past. If we had exact data on year of enrollment in Kindergarten or first grade, it would be possible to estimate the causal effect of relative age using an Instrumental Variable or Fuzzy Regression Discontinuity approach, which in principle could produce estimates of greater magnitude.

In terms of the mechanisms that could explain the results, there is a vast literature about the positive effects of relative age on academic performance. Those findings are complemented by studies of the effects of relative age on self-concept and aspirations (e.g. Peña 2020). In that sense, the results are compatible with fame being the result of greater investment in human capital among individuals who were older relative to their classmates. However, the difference between the results for males and females is puzzling. Based on the smaller number of Wikipedia entries per 10,000 births among women relative to men (3 versus 10), we speculate that the bar for fame is higher for women, leading to a less relevant role for arbitrary differences such as those created by relative age. In other words, if only truly exceptional women have Wikipedia entries, relative age may be less of a factor for women. This hypothesis is consistent with the power-law distribution of performance described in section 2 and women facing a higher success threshold than men.

Another possible explanation for the asymmetry between females and males could be the relevance of sports. The most important professional sport leagues—by revenues and fans—in the US are for male athletes (NBA, NFL, NHL, and MLB). Those leagues have college drafts, and colleges recruit students from high schools. In this process the advantages conferred by relative age given the cutoffs for Kindergarten enrollment may spill over into athletic settings. Further research using Wikipedia could look into the differential effects of relative age by occupations and test this kind of hypotheses.

The lack of evidence of strong relative age effects on success as proxied by Wikipedia entries contrasts with the results of other overrepresentation studies (Du, Gao, and Levi 2012; Fukunaga, Taguri, and Morita 2013; Muller and Page 2016). In our view, this apparent inconsistency could in principle be explained by definitions of success with different levels of stringency. After all, there are tens of thousands of Wikipedia entries and only a few hundred Nobel laureates or CEOs of Standard & Poor's 500 companies. Another reason could be a 'file drawer problem' among overrepresentation studies. Researchers may decide not to submit for publication studies in which no significant effects are found.[19] Using Wikipedia as a dragnet for many definitions of success we implicitly average relative-age effects across multiple fields. In some of them relative age may not be important, leading to a lower estimate of the effect.

Based on our results, we echo the recommendation of previous studies of age-adjusting performance on tests that can have lasting effects on human capital investment (Crawford, Dearden, and Greaves 2014; Peña 2020). A positive effect of relative age on human capital accumulation and the odds of success implies a cost in efficiency (people with more talent get fewer opportunities to succeed) and equity (people with equal talent are treated differently). Age-adjusting tests may help level the playing field, boosting efficiency and equity at the same time.

Lastly, we have two recommendations for researchers and journal editors. First, editors should make a deliberate effort to encourage the publication of non-significant results to avoid the potentially false impression that relative-age effects are ubiquitous. Holding quality constant, the sign and significance of the estimates shouldn't be a factor in the decision to publish. Simultaneously, authors should submit studies for publication regardless of the statistical significance of their results. We are at a point where over one hundred published studies have found significant effects of relative age. It's time to pay attention to the cases where no relative age effects are found. Our second recommendation is to focus on results that matter to larger sectors of the population. Very narrow definitions of success are irrelevant for most people. The effects of relative age on health, earnings, or incarceration among regular people are more important are than the effects on success among the upper echelons. If the ultimate goal is to spur a policy change regarding relative age, the case should be made in terms of its widespread effects.

Notes

1. A partial summary can be found in Peña (2017).
2. That is the approach taken in the study of professional hockey players (Barnsley and Thompson 1988) made famous by Malcolm Gladwell in *Outliers*.
3. The kind of publication bias alluded here has been documented by Franco, Malhotra, and Simonovits (2014) in a more general context.
4. Wikipedia has been used to study economic phenomena, such as the determinants of content growth (Aaltonen and Seiler 2016), incentives to contribute (Chen and Yeckehzaare 2020), and the effects of content on economic variables (Hinnosaar et al. 2019).
5. Alternatively, we could fix the distribution of performance and assume relative age produces gains of different magnitudes depending on the baseline level of performance. The insights would be the same.
6. Our purpose is to show that the relationship between the threshold a and the odds ratio can be increasing, decreasing, or constant. That property doesn't depend on the choice of the magnitude of Δ. The power-law and exponential distributions clearly illustrate this point. In the case of the normal distribution, we must assume a magnitude of Δ in order to compute numerically the odds ratio, but the specific value chosen is inconsequential.
7. Source: https://en.wikipedia.org/wiki/Wikipedia:About (last accessed Oct. 30, 2020).
8. On Wikipedia, 'notability is a test used by editors to decide whether a given topic warrants its own article. For people, the person who is the topic of a biographical article should be 'worthy of notice' or 'note'—that is, 'remarkable' or 'significant, interesting, or unusual enough to deserve attention or to be recorded' within Wikipedia as a written account of that person's life. 'Notable' in the sense of being 'famous' or 'popular'—although not irrelevant—is secondary." Source: https://en.wikipedia.org/wiki/Wikipedia:Notability_(people) (last accessed Oct. 30, 2020).
9. Some of those differences or biases have been documented by Reagle and Rhue (2011), Greenstein and Zhu (2012, 2018), and Graham et al. (2014).
10. The download was done on Oct. 18, 2020. We used the Python package 'qwikidata' to collect the information from Wikidata through SPARQL queries. For individuals whose exact date of birth is unknown, Wikidata queries conducted this way return January 1 as their birth date. To correctly categorize those cases as missing the birth date we used the Python packages 'requests' and 'bs4.' We web-scraped the Wikidata webpage of each person whose birth date was originally retrieved as January 1. We then verified whether, according to the Wikidata webpage, the month and day of birth were actually January 1.
11. Source: http://www2.nber.org/data/vital-statistics-natality-data.html (last accessed Oct. 30, 2020).
12. Forty-two states have births of each gender in all 7,305 days. The state with the fewest days with births is Wyoming, which has 7,261 days (99.4%) with female births and 7,280 days (99.7%) with male births.
13. We found only one person (a woman born in Alaska in 1973) with a Wikipedia entry whose date and state of birth were not matched by the birth data. This is likely due to the random sampling of birth certificates in the earlier years of the period covered by the sample.
14. The missing information in Bedard and Dhuey (2012) for the school years 1989–1994 and for Alaska and Hawaii was complemented with information from Bedard and Dhuey (2007).
15. A similar approach was taken by Peña (2019) using the same birthdate data combined with prison records to analyze relative age effects on incarceration rates in Florida.
16. Educational attainment and marital status of the mother are observed only on the more recent years of the sample. Using that information would imply a reduction the statistical power of our analysis.
17. The small magnitudes we find are roughly in line with the results of Dickert-Conlin and Elder (2010), who use birth certificates from 1999–2004 to analyze intentional timing and parental sorting around the cutoff date for school eligibility in the US, and find no evidence of such behaviors.
18. For instance, in the unweighted regressions California and Wyoming are given the same importance on each date. In contrast, in the weighted regressions their importance is proportional to the number of births in each state on that particular date.
19. For evidence on this type of behavior in the social sciences see Franco, Malhotra, and Simonovits (2014).

Acknowledgement

We are grateful for the valuable comments of two anonymous reviewers and the excellent research assistance of Jun Ho Choi. All errors and omissions are our responsibility.

Disclosure statement

No potential conflict of interest was reported by the authors.

ORCID

Pablo A. Peña http://orcid.org/0000-0002-1327-5412

References

Aaltonen, A., and S. Seiler. 2016. "Cumulative Growth in User-Generated Content Production: Evidence from Wikipedia." *Management Science* 62 (7): 2054–2069.

Bai, J., L. Ma, K. A. Mullally, and D. H. Solomon. 2019. "What a Difference a (Birth) Month Makes: The Relative age Effect and Fund Manager Performance." *Journal of Financial Economics* 132 (1): 200–221.

Barnsley, R. H., and A. H. Thompson. 1988. "Birthdate and Success in Minor Hockey: The Key to the NHL." *Canadian Journal of Behavioral Science* 20 (2): 167–176.

Bedard, Kelly, and Elizabeth Dhuey. 2007. "Is September Better than January? The Effect of Minimum School Entry Age Laws on Adult Earnings." Unpublished manuscript.

Bedard, Kelly, and Elizabeth Dhuey. 2012. "School-Entry Policies and Skill Accumulation Across Directly and Indirectly Affected Individuals." *Journal of Human Resources* 47 (3): 643–683. doi:10.3368/jhr.47.3.643.

Black, S. E., P. J. Devereux, and K. G. Salvanes. 2011. "Too Young to Leave the Nest? The Effect of School Starting age." *The Review of Economics and Statistics* 93 (2): 455–467.

Chen, Y., and I. Yeckehzaare. 2020. Motivating Experts to Contribute to Digital Public Goods: A Personalized Field Experiment on Wikipedia. Available at SSRN 3588132.

Crawford, C., L. Dearden, and E. Greaves. 2014. "The Drivers of Month-of-Birth Differences in Children's Cognitive and non-Cognitive Skills." *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, doi:10.1111/rssa.12071.

Deming, D., and S. Dynarski. 2008. "The Lengthening of Childhood." *Journal of Economic Perspectives* 22 (3): 71–92.

Dickert-Conlin, S., and T. Elder. 2010. "Suburban Legend: School Cutoff Dates and the Timing of Births." *Economics of Education Review* 29 (5): 826–841.

Du, Qianqian, Huasheng Gao, and Maurice D. Levi. 2012. "The Relative-age Effect and Career Success: Evidence from Corporate CEOs." *Economics Letters* 117: 660–662. doi:10.1016/j.econlet.2012.08.017.

Franco, A., N. Malhotra, and G. Simonovits. 2014. "Publication Bias in the Social Sciences: Unlocking the File Drawer." *Science* 345 (6203): 1502–1505.

Fredriksson, P., and B. Ockert. 2013. "Life-cycle Effects of age at School Start." *The Economic Journal* 124 (579): 977–1004.

Fukunaga, Hisanori, Masataka Taguri, and Satoshi Morita. 2013. "Relative age Effect on Nobel Laureates in the UK." *Journal of the Royal Society of Medicine Short Reports* 4: 1–2. doi:10.1177/2042533313492514.

Graham, M., B. Hogan, R. K. Straumann, and A. Medhat. 2014. "Uneven Geographies of User-Generated Information: Patterns of Increasing Informational Poverty." *Annals of the Association of American Geographers* 104 (4): 746–764.

Greenstein, S., and F. Zhu. 2012. "Is Wikipedia Biased?" *American Economic Review* 102 (3): 343–348.

Greenstein, S., and F. Zhu. 2018. "Do Experts or Crowd-Based Models Produce More Bias? Evidence from Encyclopedia Britannica and Wikipedia." *MIS Quarterly* 42 (3): 945–959.

Grenet, J. 2011. Academic performance, educational trajectories and the persistence of date of birth effects. Evidence from France. Manuscript.

Hinnosaar, M., T. Hinnosaar, M. E. Kummer, and O. Slivko. 2019. Wikipedia matters. Available at SSRN 3046400.

Kawaguchi, D. 2011. "Actual age at School Entry, Educational Outcomes, and Earnings." *Journal of the Japanese and International Economies* 25 (2): 64–80.

Muller, Daniel, and Lionel Page. 2016. "Born Leaders: Political Selection and the Relative age Effect in the US Congress." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 179 (3): 809–829. doi:10.1111/rssa.12154.

Peña, Pablo A. 2017. "Creating Winners and Losers: Date of Birth, Relative age in School, and Outcomes in Childhood and Adulthood." *Economics of Education Review* 56: 152–176. doi:j.econedurev.2016.12.001.

Peña, Pablo A. 2019. "Relative age and Incarceration: Born on the Wrong Side of the Calendar." *Education Economics* 27: 588–607. doi:10.1080/09645292.2019.1653826.

Peña, Pablo A. 2020. "Relative age and Investment in Human Capital." *Economics of Education Review* 78: 102039. doi:10.1016/j.econedurev.2020.102039.

Reagle, J., and L. Rhue. 2011. "Gender Bias in Wikipedia and Britannica." *International Journal of Communication* 5: 21.

Zweimüller, M. 2013. "The Effects of School Entry Laws on Educational Attainment and Starting Wages in an Early Tracking System." *Annals of Economics and Statistics* 111/112: 141–169.