The Slow Diffusion of Earnings Inequality

Isaac Sorkin
Stanford and NBER

Melanie Wallskog*
Stanford

October 2021

Abstract

Over the last several decades, rising pay dispersion between firms accounts for the majority of the dramatic increase in earnings inequality in the United States. This paper shows that a distinct cross-cohort pattern drives this rise: newer cohorts of firms enter more dispersed and stay more dispersed throughout their lives. A similar cohort pattern drives a variety of other closely related facts: increases in worker sorting across firms on the basis of pay, education, and age, and increasing productivity dispersion across firms. We discuss two important implications. First, these cohort patterns suggest a link between changes in firm entry associated with the decline in business dynamism and the rise in earnings inequality. Second, cohort effects imply a slow diffusion of inequality: we expect inequality to continue to rise as older and more equal cohorts of firms are replaced by younger and more unequal cohorts. Back of the envelope calculations suggest that this momentum could be substantial with increases in between-firm inequality in the next two decades almost as large as in last two.

*Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this data product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed under project DMS 6000576. (CBDRB-FY21-131 and CBDRB-FY21-217). This research uses data from the U.S. Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation. This research was also supported by the CenHRS project, funded by a Sloan Foundation grant to the University of Michigan, and by the Michigan Node of the NSF-Census Research Network (NSF SES 1131500). All errors are our own.

Over the past few decades, earnings inequality in the United States has increased dramatically, largely because of rising between firm pay dispersion (e.g., Song et al. (2019)). At the same time, there has been a decline in business dynamism, including a decrease in firm entry rates and a shift of employment towards older firms (e.g. Decker et al. (2020)).

In this paper, we combine the study of these two trends by exploring how labor market outcomes vary across the lifecycle of firms and across cohorts of firms. Using rich U.S. Census Bureau data, we track cohorts of firms over time. Our central finding is that firm cohort effects play an important role in accounting for a variety of important—and related—secular trends. We find an important role for firm cohort effects in explaining the rise of between-firm earnings inequality, the sorting of workers across firms, as well as rising productivity dispersion. This finding implies that these trends—including between-firm inequality—have diffused slowly as subsequent cohorts of firms have entered.

We use a variety of Census datasets that allow us to examine cohort effects starting in 1977. Briefly, our matched employer-employee data comes from the Longitudinal Employer Household Dynamics (LEHD) dataset, which runs from 1993 to 2013. We then link the LEHD to the Longitudinal Business Database (LBD) which dates firm entry starting in 1977. Hence, even though we only have matched employer-employee data from 1993 to 2013, we can learn something about the pre-1993 by looking at their behavior in our sample. For example, in 1994, the 1977 cohort is 17 years old and so if it behaves unusually relative to other 17 year old cohorts or other cohorts in 1994, we can learn something about the 1977 cohort.

We begin by presenting simple visual evidence that suggests the importance of cohort effects in these aggregate trends. We plot (see Figure 3) the between-firm variance of earnings within a cohort over time. We find two striking patterns. First, within a cohort, between firm earnings inequality declines as the cohort ages. Second, each subsequent cohort of firms enters with a higher level of between firm earnings inequality before declining on an approximately parallel path to previous cohorts. We find similar patterns in productivity dispersion at the cohort level.

We then develop an approach to formally bound the contributions of the cohort effects. We adopt an additively separable decomposition. As is well-known, without further restrictions, the age, time, and cohort effects are not separately identified (e.g., Hall (1968)). Logically, in the simple visual analysis, what we are inclined to label as cohort effects could instead be year effects. We leverage the fact that second differences of age, year, and cohort effects, i.e. the "shapes" of these effects, are identified (McKenzie (2006)). To understand the role of time and cohort effects, we consider two extreme assumptions: first, that the year effects explain none of the aggregate change in the variable; and second, that the year effects explain all of the aggregate change in the variables. Throughout, we report these as bounds on the contribution of age and cohort effects, and we focus on results where the sign of the cohort effects are consistent across these two normalizations.

Our more formal analysis confirms the basic findings from the visual analysis: as a cohort of firms ages, the between-firm dispersion of pay decreases. The within-cohort decline in between-firm variance of earnings is at odds with the aggregate increase. This finding highlights that attempts to

explore the source of the rise of between-firm pay inequality by following balanced panels of firms will tend to be misleading. Instead, one important reason between-firm pay dispersion has risen is because subsequent cohorts are fundamentally different than previous cohorts: they enter — and continue to be — more dispersed. Across our two normalizations, we find a quantitatively significant contribution of cohort effects to explaining the rise of between-firm earnings inequality—anywhere from 50 to 100%.[1]

We find similar cross-cohort patterns in a variety of other outcomes. Workers in firm cohorts with higher between-firm earnings inequality are more sorted on earnings, age, and college attainment. Additionally, cohorts with higher between-firm pay dispersion are more dispersed in productivity. On the other hand, pay at firms in these cohorts is less correlated with productivity and worker retention. The broad summary is that many of the important firm-level labor market trends in the last 30 years have an important cohort dimension.

At a high level, our results do not restrict the set of explanations for the rise of between-firm earnings inequality. Conceptually, the simplest versions of most explanations for the rise of between-firm earnings inequality—e.g., outsourcing, skill biased technical change, or deunionization—are often framed as "year" effects that affect all firms in the same way. For each of these explanations, however, there is a fairly natural extension where instead of happening to all firms at the same time, these changes diffuse slowly across subsequent cohorts. To take the simplest example, if unionization status is persistent within a firm, and subsequent cohorts are less unionized, then the decline of unions would more naturally be a cohort phenomenon than a year phenomenon. Or in the case of skill-biased technical change, while Katz and Murphy (1992) pose it as a year effect, it might be that it instead reflects subsequent cohorts of firms adopting different technologies.

What our results do imply is that the inequality "technology"—whatever that may be—has diffused slowly through the economy. Even if that technology does not change and so subsequent cohorts enter with the same "technology" as the most recent entrants, then we would expect between-firm inequality to continue to rise. We illustrate this force by computing forecasts through 2030 where we assume that post-2013 cohorts have the same cohort effects as the 2013 cohort. Depending on normalization, we find a rise in between-firm earnings inequality of 0.02 to 0.08 log points, which is large compared to the 0.09 log point rise from 1993 to 2013 in our sample. Similarly, the cohort perspective generates a different time path of the inequality "innovations" than looking at aggregate trends.

In the context of the business dynamics literature, a variety of papers (e.g., Sterk, Sedláček, and Pugsley (2021), Karahan, Pugsley, and Sahin (2021), and Hopenhayn, Neira, and Singhania (2019)) have emphasized changes in the entry process for firms. Our results suggest a connection between these changes in the firm entry process and the increase in between-firm earnings inequality. Several simple explanations are not consistent with the data. For example, rising within-cohort dispersion suggests a model of falling entry costs and so subsequent cohorts are less selected. But at the same

[1]The overall between firm variance of earnings also includes the between-cohort average pay differences. We keep track of these differences in our analysis. We normalize the cohort means separately from the cohort variances and so even when the year effects explain "everything" we still find a major role for the cohort effects.

time the rate of firm entry has fallen, which is inconsistent with falling entry costs. Similarly, one potential explanation is that our effects reflect "labor supply" factors through an aging workforce (e.g. Karahan, Pugsley, and Sahin (2019) and Engbom (2019)). We consider two exercises that suggest that this channel is not quantitatively important. First, we look at new hires, and find remarkably similar firm cohort patterns. Second, inspired by Guvenen et al. (2019), we look at cohorts by worker age and do not find analogous patterns. These findings point against a direct channel from "worker" cohort differences to "firm" cohort differences.

The paper unfolds as follows. Section 1 describes the data, variables, and samples. Section 2 presents the motivating aggregate trends in earnings inequality and business dynamism; we show that, as previous literature has found, much of the increase in aggregate earnings inequality is accounted for by increases in between-firm pay dispersion and there have been a substantial shift in employment from younger firms towards older firms over our sample time window. Here we also show that there are significant cross-cohort patterns: firms in newer cohorts are more dispersed both in their pay and productivity than older cohorts. We formally investigate these cohort patterns by estimating an age-time-cohort decomposition described in Section 3. We present the results in Section 4, where we show that between-firm earnings inequality declines over the lifecycle of a cohort of firms, and in Section 5, where we document robust cohort patterns that can account for much of the aggregate increase in between-firm earnings inequality. Section 6 considers several extensions and robustness analyses. Section 7 discusses our results in light of the existing literatures on earnings inequality and business dynamism. Section 8 concludes.

1 Data and variable definitions

At a high-level, we combine three data sources from the U.S. Census Bureau. We use matched employer-employee data to construct labor market statistics, including measures of earnings, employment, and demographics. We augment this data with information from a census of businesses in order to know when firms entered, which allows us to study cohorts of firms. Finally, in order to compute a notion of productivity, we additionally use revenue information from another dataset.

1.1 Data

We primarily use the Longitudinal Employer Household Dynamics (LEHD) data.[2] This dataset is constructed from firm-side unemployment insurance records and contains quarterly information on employment and earnings.

We use data from 1993 to 2013 for the nine states that have complete records over this span.[3] The unit of observation in the LEHD is a state-level unemployment insurance account, which can contain multiple establishments within a state; we aggregate the LEHD state-level unemployment

[2] See Abowd et al. (2009) for a detailed description of the LEHD.

[3] The states are: CA, FL, ID, LA, MD, NC, OR, WA, and WI. According to our calculations using the U.S. Census Bureau's Business Dynamics Statistics dataset (2018 vintage), these states accounted for 28.1% of 1993 national (50 states and DC) employment and 29.1% of 2013 national employment.

insurance accounts to the firm-level.[4] Since we do not have complete coverage of the United States, our unit of observation is thus approximately the firm (where each firm may have establishments in states outside our sample that we do not consider).

To date firm entry, we use the Longitudinal Business Database (LBD), an establishment-level dataset that starts in 1976. Following Haltiwanger, Jarmin, and Miranda (2013, pg. 353), the entry year of the firm is the entry year (i.e., the first year with positive employment in the payroll period that contains March 12) of the oldest establishment within the firm, meaning that we have non-censored entry dates for firms starting in 1977.[5] We call all firms that enter in the same year a cohort. By combining the LBD data with the LEHD, we track lifecycles of cohorts up to the age of 37. We one-index, such that we define the age of a firm in its first year as 1; a firm that enters in 1977 is aged 37 in 2013.

The LEHD covers almost all sectors of the economy. Specifically, it includes workers covered by the UI system, which in 1994 reflected about 96% of employment and 92.5% of wages and salaries (BLS (1997, pg. 42)). Important omissions include small non-profits, self-employed workers, as well as some agricultural workers and federal government workers.[6] We also omit firms whose primary sector is public administration (we discuss how we define a primary sector below), which includes local and state government.

We construct an annual dataset based on employment and earnings in the first quarter of each year (Q1); we do this to align with the LBD entry dates, which are based employment in Q1 (payroll period containing March 12).[7] We focus on workers where Q1 is a "full quarter" of employment at the employer. Following Abowd, Lengermann, and McKinney (2003), full quarter means that the worker also had positive earnings at the employer in Q4 of the previous year and Q2 of the current year.[8] Restricting to full quarter allows us to be confident that the worker's employment did not start in the middle of Q1.[9] We then convert to real 2017 dollars using the CPI-U and annualize by multiplying by four. We also impose an earnings floor of $3,250 (in $2017) in the annualized

[4]Note that some firms, particularly larger firms that span multiple industries, can have multiple stat-level unemployment insurance accounts within the same state. The Census variable that we use to identify a firm is "firmid"; in each year, we take each state-level unemployment insurance account's firmid from the LEHD's ECF-T26 dataset.

[5]We use the longitudinal links in the LBD to link establishments over time.

[6]For details see Kornfeld and Bloom (1999, pg. 173), BLS (1997, pg. 43) and `http://workforcesecurity.doleta.gov/unemploy/pdf/uilawcompar/2012/coverage.pdf`.

[7]Figure A3 shows that our main descriptive fact also holds using annual earnings. If we instead use true annual earnings for our sample, we get similar trends to earnings inequality, but higher levels, likely driven by workers changing firms in other quarters.

[8]UI earnings include the following components: "gross wages and salaries, bonuses, stock options, tips and other gratuities, and the value of meals and lodging" (BLS (1997, pg. 44)). UI earnings omit the following components: "employer contributions to Old-age, Survivors, and Disability Insurance (OASDI); health insurance; unemployment insurance; workers' compensation; and private pension and welfare funds" (BLS (1997, pg. 44)).

[9]If we include workers who start their employment at a firm in the middle of the quarter, we may conflate earnings variation with employment-spell length variation. In the absence of precise spell length data, we limit to full quarter to employment to isolate better the type of earnings variation that we wish to study. Note that variation in our measure of earnings may reflect a combination of both wage and hours variation. Figure A4 shows that there is a small upward trend in the share of employment accounted for by full-quarter employment over our sample period: from 1993 to 2013, the share of Q1 jobs that are full-quarter and earn above our earnings floor (see below) increases by 7%, i.e. from 70 percentage points to 75 percentage points.

earnings (i.e., the threshold for Q1 earnings is $3250/4$), which follows Sorkin (2018, pg. 1339) (who in turn follows Card, Heining, and Kline (2013)) and is similar to Song et al. (2019, pg. 13). Earnings are summed across all UI accounts within the firm within the state where the worker is located.[10] We conduct all of our analysis is in logs, and henceforth refer to log earnings as simply "earnings."

For essentially all worker we observe age, race and gender (a small share of these observations are imputed). We group the standard Census race and ethnicity categories into the following exhaustive and mutually exclusive groups: white, Black, hispanic, Asian, Native American and other. For the 8% of the sample that can be linked to the Census long-firm (see Vilhuber et al. (2018, pg. 5-2) for more extensive description), we also observe coarse education information; for this group, we create a dummy variable for those who have completed college.

We keep workers aged 25-60 (inclusive), where we measure age on December 31st of the year.[11] We pick 25 as the lower bound to be "post-schooling." Similarly, we pick age 60 in an attempt to minimize the role of retirement.

We assign time-varying 6-digit NAICS-based sector codes to the firms using sector information in the LEHD. We assign the sector code based on the sector with the maximum employment among the constituent UI accounts across the states in data.[12]

1.2 Variable definitions

We construct many variables.

Earnings: Throughout, we use log earnings, which is the log of real earnings in Q1, annualized by multiplying by four.

(Labor) Productivity: We bring in revenue information from the revenue-enhanced LBD (Haltiwanger et al. (2017)).[13] This dataset has firm-level revenue information starting in 1997.[14] We compute labor productivity as the firm-level revenue (in 2017 dollars, deflated using the CPI-U) divided by the firm-level number of workers, where these numbers include the sales and employment in the entire country; this definition means that the employment counts are taken from the LBD measures which reflects employment on the pay period including March 12.[15] We follow Decker

[10] Looking only within a state to define worker earnings reduces the computational burden. Note that individuals may have multiple employers, i.e. workers may appear in our sample at more than one firm in a given state, across states within the same firm, or across states at different firms; we treat these as distinct employment relationships.

[11] That is, worker age is equal to the year minus their year of birth.

[12] For the purposes of defining the firm's sector, if an individual is employed at multiple UI accounts within a firm within a state-year, then we assign the individual the industry of the UI account from which they earned the most in that quarter, with ties broken arbitrarily.

[13] This data is currently separate from the full LBD dataset (and thus we refer to it as a third data source) and is available to researchers on approved projects through the Federal Statistical Research Data Center (FSRDC) network, where additional documentation is available (Haltiwanger et al. (2019)).

[14] We use the 2015 vintage of the data.

[15] We compute an alternative measure of labor productivity as revenue divided by payroll, where national payroll is estimated as the product of the mean earnings in our sample at a firm and the employment count at the firm; this measure attempts to capture productivity in terms of the quality of labor input.

et al. (2018, pg. 9) and center this number around the national 6-digit NAICS by year productivity level.[16] Foster, Haltiwanger, and Krizan (2001) show that this measure is quite highly correlated with better measures of firm productivity that one can construct in the manufacturing sector. Using a firm-level measure has the benefit of combining all establishments. When we aggregate our productivity measure across firms, we weight by the LBD inverse propensity score weights, which are designed to make the estimates representative of the whole economy.

Sorting: Following the literature, we are interested in measures of how sorted workers are across firms on a variety of characteristics. We measure sorting in a way closely related to measures proposed in Kremer and Maskin (1996) and Borovičková and Shimer (2017) as the correlation of a worker's value of the outcome with her co-workers average value of that outcome. Formally, let $y_{i,j(i)}$ be a worker level variable when she is employed at firm i. Let $\bar{y}_{j(i)-i}$ be the leave-out mean of variable y at firm j. Then define:

$$\rho_y = \text{corr}(y_{i,j(i)}, \bar{y}_{j(i)-i}) \tag{1}$$

where this correlation weights each worker equally. This measure adjusts for mechanical sorting in small samples and is computationally cheap to compute.

We compute sorting on the basis of earnings, as well as college education, age, race/ethnic group, and being female.

Means of worker characteristics: We compute the within-firm shares of demographic groups as well as the mean age of workers.

Firm-level statistics: Besides productivity, we compute several other firm-level statistics:

- Employment: the number of full-quarter workers above the income threshold in Q1. We sometimes take log of the employment counts.
- Pay: the average of the worker-level log earnings.
- College-high school pay gap: the difference between the average of the worker-level log earnings for college-educated workers and the average of the worker-level log earnings for high school-educated workers.
- Exit: an indicator equal to 1 if a firm does not employ (full-quarter, above income threshold in Q1) workers in the following year, and 0 otherwise. Because this measure requires looking into the future, we can only compute this measure in 1993 through 2012.
- Reallocation dynamics measures:

[16] 6-digit NAICS productivity is the sum of revenue divided by the total number of bodies in the industry.

- Retention rate: the proportion of workers employed full-quarter at firm j in Q1 in year t who are also employed full-quarter at firm j in year $t+1$. Because this measure requires looking into the future, we can only compute this measure in 1993 through 2012. Note that if firm j exits in year $t+1$, it has a retention rate in year t of 0.
- Share new hires: the proportion of workers employed full-quarter at firm j in Q1 in year t who were not employed full-quarter at firm j in year $t-1$. Because this measure requires looking into the past, we can only compute this measure in 1994 through 2013.
- DHS employment growth: the difference between full-quarter employment in Q1 in year t and full-quarter employment in Q1 in year $t-1$, divided by their average (Davis, Haltiwanger, and Schuh (1996)). Because this measure requires looking into the past, we can only compute this and the following measures in 1994 through 2013. When we aggregate this and the following measures across firms, we worker-weight by the sum of employment in year t and year $t-1$.
- Job reallocation rate: the absolute value of the difference between full-quarter employment in Q1 in year t and full-quarter employment in Q1 in year $t-1$, divided by their sum. Because this measure requires looking into the past, we can only compute this measure in 1994 through 2013.
- Worker flows: the sum of the number of new hires (workers employed in year t but not in year $t-1$) and the number of separating workers (workers employed in year t but not in year $t+1$) (Burgess, Lane, and Stevens (2000)).
- Worker flow rate: the worker flows divided by the sum of employment in year t and year $t-1$ (Burgess, Lane, and Stevens (2000)).
- Excess reallocation: the difference between the worker flow rate and job reallocation rate.

We also compute the correlation between the firm-level statistics. For many of the firm-level statistics, we follow conventions in the literature and report firm-weighted estimates (as opposed to worker-weighted).

1.3 Samples

Throughout the paper, we use two different samples, which reflects the structure of our data and research questions.

Full sample: The full sample includes all firms (in our nine states) that are covered by our data from 1993 to 2013.

1977 cohorts and beyond: Based on the LBD, we can date firm entry dates starting in 1977. Thus, we create a sample of firms, which is a subset of the full sample, where we have non-censored

entry dates (and thus know firms' ages). We use this sample to estimate the age-time-cohort decomposition.

2 Aggregate trends in earnings inequality and business dynamism

We begin by documenting several aggregate trends that have been shown in the literature and serve as the backdrop of this paper. Specifically, we show that there has been an increase in earnings inequality and this increase has largely been driven by increasing between-firm differences. We also show that the age distribution of firms has changed, with more employment in older firms. We then show some simple plots of the data that emphasize striking patterns in between-firm earnings inequality within the lifecycle of a cohort of firms and across cohorts. Collectively, these patterns motivate our formal analysis that follows.

2.1 Rise in earnings inequality driven by increasing between-firm dispersion

The first major change is that earnings inequality has risen. Figure 1 (and Table 1) shows that in the last twenty years, earnings inequality has risen in the United States. Panel A shows that from 1993 to 2013 the total variance of earnings has increased from 0.61 to 0.78, an increase of 0.17 log points.[17] The Figure and Table also show that this pattern holds in our 1997 cohorts and beyond sample, albeit with slightly smaller magnitudes. In the 1977 cohorts and beyond sample, the variance of earnings rises by 0.13 log points.

The finding of rising earnings inequality is quantitatively consistent with the literature. For example, Song et al. (2019) find an increase of about 0.12 log points over the same period, though there are some differences in sample.[18] In the U.S. using the LEHD, Barth et al. (2016, Table 1) report an increase of about 0.08 log points from 1992 to 2007.

The second major fact is that most of the rise in inequality is between-firms. Let $y_{i,j(i)}$ be the earnings of worker i at firm j and $\bar{y}_j$ be the average earnings of workers at firm j. We can then write the variance of earnings as:

$$\underbrace{\text{var}(y_{i,j(i)})}_{\text{total dispersion}} = \underbrace{\text{var}_j(\bar{y}_j)}_{\text{between-firm dispersion}} + \underbrace{\sum_j s_j \times \text{var}(y_{i,j(i)}|j(i)=j)}_{\text{within-firm dispersion}}, \tag{2}$$

where s_j is the employment share of firm j. The firm term captures the between-firm dispersion in earnings. The second term captures the employment-weighted mean of the within-firm dispersion in earnings, i.e. to what degree the average worker's earnings varies from her coworkers' earnings.

[17] The total variance of earnings is the worker-weighted variance of log earnings across all individuals in our sample, in a given year.

[18] Song et al. (2019, Table II) find an increase from 1981 to 2013 from 0.652 to 0.846 of which the increase in the between component is 0.694 share. Their Figure II shows that this trend is quite linear. Extrapolating to our sample window, this implies an increase in total variance of earnings from 0.724 in 1993 to 0.846 in 2013. Note that their sample only includes firms with 20 or more workers, whereas we also include small firms.

Figure 1 shows that most of the rise of earnings inequality in the last twenty years has been between-firm rather than within-firm (Table 1 displays the tabular version of this result). Quantitatively, we find that from 1993 to 2013, 70% of the increase is between-firm.[19] This finding is quantitatively consistent with the literature. For example, Song et al. (2019, Table II) also find the between-firm share is about 70%.[20] This finding also holds in the subsample where we have non-censored age information. In particular, in the 1977 cohorts and beyond sample the between/within split is nearly identical to the full sample.

How important are sectoral shifts? One simple possible explanation is that the rise of the between-firm earnings inequality reflects sectoral shifts: i.e., the employment shares of more unequal sectors has expanded. In Table 1 we show that this is not the case. Specifically, we perform a shift-share decomposition and find that reallocation of employment across sectors accounts for a very small share (less than 5 percent) of the increase of between-firm earnings inequality.[21]

The new aspect of this paper is that we connect these two facts about earnings inequality to facts about firm dynamics. This focus on firm dynamics differs from the direction the recent literature motivated by the two earnings inequality facts has gone. Most prominently, Card, Heining, and Kline (2013) and Song et al. (2019) have focused on repeated pooled cross-sections through the lens of a two-way fixed effects regression and sought to understand how much this increase reflects changing patterns of between-firm differences in pay policies and how much reflects the sorting of workers to firms.[22] Building on this theme, there has been a revival of interest in models of imperfect competition in the labor market (e.g., Berger, Herkenhoff, and Mongey (2019), Card et al. (2018), Jarosch, Nimczik, and Sorkin (2019), Kline et al. (2019), and Lamadon, Mogstad, and Setzler (2019)).

2.2 Decline in business dynamism

The first important firm dynamics fact is that employment has shifted towards older firms. Figure 2 contrasts the share of employment at young (under the age of five) versus older (over the age of fifteen) firms. The Figure shows that, from 1993 to 2013, the share of employment accounted for by young firms falls by over a third (or over 7 percentage points, from a base of 17 percentage points), while the share of employment at older firms rises by over 13 percentage points[23]. This pattern is similar to that documented in other work. For example, the change in employment share

[19] If instead of studying annualized (i.e. Q4 multiplied by four) earnings we use true annual earnings, this value is slightly larger: 75% of the increase in annual earnings inequality is between-firm.

[20] Barth et al. (2016, Table 1) reports an increase in the between component is from 0.219 to 0.275 (growth of 0.056, 67% of their reported total increase).

[21] Spletzer and Haltiwanger (2020)) raise a related point that the between sector variance has risen at the same time that the overall between-firm variance has risen. In Figures A1 and A2, we show that our main facts also exist within sector.

[22] More recently, Lachowska et al. (2020) and Engbom, Moser, and Sauermann (2021) augment this framework by directly allowing firm effects to vary by year.

[23] Figure 2 also shows that the share of firms that are young falls by a little under a third (or 10 percentage points, from a base of 38 percentage points).

among firms five and under is comparable to that documented in Decker et al. (2018, Figure 2).[24] It is also part of a broader pattern of changes in business dynamics which has been referred to as the "decline in business dynamism" (see, e.g., Haltiwanger, Decker, and Jarmin (2015) on declines in start-ups, job-to-job mobility, and firm responsiveness to shocks and Akcigit and Ates (2019) on the decline of knowledge diffusion across firms).

The second important firm dynamics fact, which emphasizes a tight connection between-firm dynamics and earnings inequality, is that the between-firm component of earnings inequality both evolves over the lifecycle of a firm cohort and has changed across subsequent cohorts. Table 1 begins to show this pattern by looking at the 1993 cohort. Similar to the full sample, there is a sharp rise in earnings inequality within the cohort from 1993 to 2013. The composition of this change in inequality within the cohort, however, is very different than the overall sample. In particular, whereas in the overall sample it is between-firm inequality that rises, in this single cohort between-firm inequality actually falls. This finding turns out to hold more systematically. Thus, the pattern of between-firm inequality among continuing firms is quite different than among all firms, emphasizing the importance of thinking about firm entry.

To see this pattern across all cohorts, Figure 3 plots the patterns of the between-firm variance of earnings for the cohorts that entered starting in 1993. One can see that the general pattern for the 1993 cohort holds: over the life of the cohort, the between-firm variance of earnings declines. In contrast, across cohorts there is a general pattern that younger cohorts enter with higher initial levels of earnings inequality. This Figure suggests the importance of cohort effects in explaining the rise of earnings inequality: newer cohorts appear to be fundamentally different than older cohorts in their level of between-firm earnings inequality, and these differences persist across the lifecycles of these cohorts.[25]

The bottom panel of Figure 3 shows that this cross-cohort pattern is not confined to earnings. The Figure shows that within a cohort the variance of productivity falls over time, while across cohorts the variance rises. This emphasizes that subsequent cohorts of firms are not different just in terms of how they pay workers, but also in terms of one fundamental determinant of pay: productivity. In our analysis in the following sections, we explore more systematically how cohort variation in productivity, along with other outcomes, co-moves with cohort variation in earnings inequality.

We now turn to more formal investigations of these patterns that acknowledge the difficulties in separating age, cohort and year effects.

[24] Note that Decker et al. (2018) zero-index age, such that new firms are age 0; thus, their "young" firms aged $<$ 5 are our firms aged $\leq$ 5.

[25] In concurrently written work, Engbom, Moser, and Sauermann (2021) find similar cohort patterns to firm pay dispersion in Sweden, taking worker fixed effects into account. We take their findings as supportive evidence that the trends we document are neither unique to the United States nor driven by worker selection (which we explore in greater detail below).

3 The age-time-cohort decomposition

To more formally investigate the patterns documented in Figure 3, we now develop an age-time-cohort decomposition. We first explain what quantities are identified. We then discuss how in estimation we impose a range of normalizations.

3.1 What is identified

We adopt the additively separable age-time-cohort decomposition, where we can write a cohort-year outcome, e.g. the between-firm variance of earnings, as:

$$y_{c,a,t} = \chi_c + \alpha_a + \tau_t + \epsilon_{c,a,t}, \tag{3}$$

where y is a cohort outcome, c is a cohort, a is an age, and t is a time period; that is, we decompose the outcome into cohort, age, and year "effects" (χ_c, α_a, and τ_t respectively) and a residual ($\epsilon_{c,a,t}$). Naturally, since Hall (1968) (see Deaton (1997, pg. 123-128) for a textbook treatment) it has been understood that the effects are not separately identified without further restrictions. The fundamental identification issue is that the age, time, and cohort effects are linearly dependent; e.g., an age-time is also a cohort-age.

Fortunately, there are certain quantities that can be identified. Specifically, as McKenzie (2006) shows, second differences (i.e., "shapes") of the age, year, and cohort profiles are identified. To see an example of this fact, take first differences within a cohort over time:

$$\Delta_t y_{c_{1994},a_2,t_{1995}} = (\alpha_{a_2} - \alpha_{a_1}) + (\tau_{t_{1995}} - \tau_{t_{1994}}) + \Delta_t \epsilon_{c_{1994},a_2,t_{1995}}, \tag{4}$$

where Δ_t indicates that we difference with respect to time. Then time-difference within the next cohort in the same time window:

$$\Delta_t y_{c_{1993},a_3,t_{1995}} = (\alpha_{a_3} - \alpha_{a_2}) + (\tau_{t_{1995}} - \tau_{t_{1994}}) + \Delta_t \epsilon_{c_{1993},a_3,t_{1995}}. \tag{5}$$

Take differences across the cohorts (where Δ_c indicates a cohort difference):

$$\Delta_c \Delta_t y_{c_{1993},a_3,t_{1995}} = (\alpha_{a_3} - \alpha_{a_2}) - (\alpha_{a_2} - \alpha_{a_1}) + \Delta_c \Delta_t \epsilon_{c_{1993},a_3,t_{1995}}. \tag{6}$$

We can see that the curvature in the age profile is identified; that is, comparing within-cohort across-time changes in the outcome across cohorts isolates the change-in-the-change of age effects (i.e. the curvature of the age profile). This logic generalizes such that the second differences in age, year, and cohort in fact are identified.

Naturally, estimating the age, cohort, and year effects still requires that we impose additional normalizations. In particular, levels are not identified, and one slope parameter is not identified.[26]

[26] That is, conditional on levels for one cohort, one age, and one year effect as well as one slope (first difference) of either a cohort, age, or year effect (in addition to the identified second differences), the remaining levels and slopes

Because we are interested in explaining changes over time, the level normalizations are uninteresting; that is, though we pick a level for the age, time, and cohort effects, we end up differencing out the levels for all results that we report.

In terms of slope parameters, the central question for this paper is how much of the aggregate effects are driven by year vs. cohort effects: i.e. do the cohort differences in Figure 3 reflect true cohort effects or rather year patterns? We report all results with two separate extreme normalizations; first, we normalize the slope parameters by picking year effects such that the change in the year effects from the first to the last time period explain the entire aggregate trend in a given outcome; second, we alternatively normalize the slope parameters by picking year effects such that there is no change in the year effects from the first to the last time period and so year effects explain none of the aggregate trend. We offer formal details below.

Beyond allowing us to estimate equation (3), how do we interpret these normalizations? We interpret there to be "robust" evidence of cohort effects if cohort effects contribute to explaining an aggregate pattern under both year-slope normalizations. The reason there is scope for there to be "robust" evidence of cohort effects is that the age effects also change when we change the year normalization. Suppose that the additively separable model fit the data perfectly. Then each data point would be explained by a combination of age, year and cohort effects. When we decrease the year effect, then either the cohort effect or the age effect has to increase. Thus, to the extent that the age profile is sensitive to the normalization of the year effects, then we will find robust evidence of cohort effects.

A substantive limitation of this decomposition is that while the additive separability is a natural way to attempt to disentangle age, year and cohorts effects it is a perhaps unnatural assumption in many economic models. For instance, we argue below that younger cohorts are more dispersed in terms of productivity than older cohorts. If cohorts compete in the same markets, then we may expect spillovers across cohorts. To the extent that these spillovers are common across all cohorts, then it would end up in the year effects. But if these spillovers affect "closer" cohorts more, then this decomposition would miss these effects. Nonetheless, we believe that the simple decomposition, in addition to having tractability, sheds at least some light on meaningful patterns in the data. Additionally, the additively separable model is supported by visual evidence in the aggregate patterns: in Figure 3, cohorts appear to enter at different levels and then remain at different levels throughout their lifecycles. This "intercept" difference, paired with approximately parallel trajectories across subsequent cohorts, is broadly consistent with additively separable year and cohort effects. Below we offer additional evidence that the decomposition fits the data well in that we re-restimate some coefficients on subsets of the data and find similar values, which is inconsistent with important interaction terms.

follow. See Section A for complete details.

3.2 Implementing the age-time-cohort decomposition

We begin by noting that in the age-time-cohort decomposition an observation is an outcome within an age-cohort cell. In the context of the variances, then, we are estimating the between-firm variance of earnings *within* an age-cohort cell and these do not capture the between age-cohort cell terms. To address this issue, when we aggregate our estimates in Section 5, we additionally keep track of the difference of the age-cohort means to arrive at the overall between-firm variance of earnings.

We implement our decomposition by estimating a worker-weighted and constrained version of equation (3) on the post-1977 sample (for whom we know all firms' ages). We estimate a total of 37 age effects (age 1 to 37), 37 cohort effects (1977 to 2013), and 21 year effects (1993 to 2013). Recall that an observation is a cohort-year (or an age-year). For weights, we use the employment share of the cohort-year in a given year. We impose four sets of constraints. The first set of constraints are three level normalizations. We fix the 1977 cohort effect and the 1993 year effect to zero. Hence, we estimate 93 coefficients. We also need to pick a level for the age effects and we set age 1 to zero.

The second constraint is that our coefficient estimates "explain" the overall change in our sample. Formally, letting $s_{c,t}$ be the cohort c share of employment in year t and and $y_{c,t}$ be the cohort c outcome in year t, we define the aggregate change $\Delta y_{agg} \equiv \sum_c s_{c,2013} y_{c,2013} - \sum_c s_{c,1993} y_{c,1993}$. Note that we use the post-1977 sample in this estimation and so the sample in which the aggregate is calculated is the post-1977 cohort sample. The reason to impose this constraint is that we subsequently want to impose restrictions that year effects explain different shares of the overall change, and so we need to fix the overall change ex-ante.

The third constraint is that the year effects capture a particular share of the aggregate change, Δy_{agg}. Because we normalize the first year effect to zero, $\tau_{1993} = 0$, this amounts to picking the last year effect, τ_{2013}. We use two alternatives: that the year effects explain none of the aggregate change ($\tau_{2013} = 0$) and that the year effects explains all of the aggregate change ($\tau_{2013} = \Delta y_{agg}$).

The final set of constraints are the second differences. Specifically, we estimate the second differences using the average of all the data; i.e., in the above example of the age curvature from age 1 to 3, we use not only the comparison of the 1993 to 1995 cohorts, but the 1994 to 1996 cohorts, the 1995 to 1997, and etc., until the end of our data. We weight each of these comparisons equally.

We estimate standard errors using the residual bootstrap (see, e.g., MacKinnon (2006)). Bootstrapping allows us to take into account the estimation of the constraints in constructing our standard errors.[27] The residual bootstrap allows us to maintain the same set of age-time-cohort observations in each estimation sample and so we are able to compute the constraints in each bootstrap repetition. Formally, we estimate equation (3) and store the residuals. We then build a bootstrap repetition by resampling these residuals with replacement. To compute standard errors we generate 50 sets of bootstrap estimates and take the square root of the variance of the bootstrap

[27] We note that this procedure treats the cohort-year observations as being estimated without noise and only account for the uncertainty in the age, cohort, and year effects. We view this assumption as reasonable since we essentially have the population and so the standard error on the estimate of any observation is essentially zero.

estimates. In Appendix B, we provide Monte Carlo evidence on the performance of our procedure.

4 Lifecycle patterns in earnings inequality

The key novelty of this paper is to follow cohorts of firms over time and to document facts about earnings inequality. We report all results with the two alternative normalizations of year effects explain "all" and year effects explain "none" of the overall change.[28]

Figure 4 reports the basic lifecycle patterns for earnings variables. Panel A shows that there is not robust evidence of a trend in overall earnings inequality as the cohort ages. Specifically, depending on normalization, we either find significant increases or decreases in the total variance of earnings within a cohort. To understand magnitudes, note that the final coefficient (age 37) is about 0.17 when we normalize the year effects to zero, which is large compared to the overall variance of earnings of 0.61 in 1993.

In contrast, Panels B and C shows that there is robust evidence of off-setting changes over the lifecycle for between and within-firm earnings inequality: over the lifecycle, between-firm inequality declines within a cohort while within-firm inequality rises, regardless of the year slope normalization. This result confirms and quantifies the intuitive finding we might have gleaned from Table 1 for the 1993 cohort and from Figure 3 which showed patterns of between-firm earnings inequality for the non-censored cohorts.

Panel D of Figure 4 shows that mean earnings rise as the cohort ages. We display mean earnings because when we aggregate across cohorts to compute the overall between-firm variance we need the patterns in the mean earnings as well. The pattern here is quite sensitive to the normalization. When year effects explain all of the rise in earnings, then there are essentially no increases in earnings after the first few years. But when year effects explain none of the aggregate trend, we find a steady rise throughout the lifecycle.

In Figure 5 we revisit the lifecycle pattern of between-firm inequality by looking at balanced panels. We construct this Figure by restricting to the cohorts for which we observe the first 10 years (i.e., from age 1 to 10; the 1993 to 2004 cohorts) or the second 10 years (i.e., from age 11-20; the 1983-1992 cohorts) of their lifecycles. We then run the same age-time-cohort decomposition as before. We do two versions: one where we construct the cohorts using only the balanced set of firms and one with the unbalanced set of firms. Panels A and B contrast age effects when we change the sample to estimate the age effects on just either age 1 to 10 firms or age 11 to 20 firms, but without restricting to the balanced set of firms. The take-away from these figures is that changing the sample does not have a large effect on estimates. (This exercise has a dual interpretation as a specification test of the additively separable model; this figure shows that results are similar on a

[28] The literature documenting firm dynamics facts adopts different normalizations. E.g., Hsieh and Klenow (2014, Figure 1) present synthetic cohorts which normalize cohort effects to zero; Haltiwanger, Jarmin, and Miranda (2013) assume that age and cohort effects each sum to zero in their baseline specifications, and explore robustness to only normalizing cohort effects to zero; in their main results, Pugsley, Sedlacek, and Sterk (2019) assume that cohort and year effects sum to zero; when they explore changes over time they split the sample and thus adopt an analogous assumption in each subsample.

subset of the data).

Turning to the balanced panels, the central message is that in the balanced panel the lifecycle patterns are much weaker in the first 10 years of a firm's life. In Panel C and D we see that from age 1 to 10 in the balanced panel there is essentially no trend. That is, if we pick the set of firms that will survive until age 10, then there is no decrease in the between-firm variance of earnings as the cohort ages. The difference between the balanced and unbalanced panels thus implies that selection is important in explaining the *lifecycle* patterns in the between-firm variance of earnings. Loosely, the between-firm variance of earnings declines as a cohort of firms ages because "outlier" firms exit. In contrast, for age 11 to 20, restricting to the balanced panel makes much less of a difference. Hence, selection that is relevant for thinking about between-firm earnings inequality occurs in the first 10 years of a firm's lifecycle.

The decline of between-firm earnings inequality as a cohort ages has important implications for wanting to study where rising between-firm earnings inequality comes from. In particular, among a continuing set of firms, the trend in between-firm earnings inequality is on average the opposite of the aggregate trend (or, in the balanced version, the trend is flat at the beginning of the lifecycle). Given that the average worker is increasingly likely to work at an older firm (Figure 2), understanding the aggregate increase in between-firm earnings inequality requires either that there are important year effects or important cohort effects, which we turn to next.

5 What has changed: Cohort and year effects

So far we have established that the between-firm component of earnings inequality plays a large role in explaining the increase in earnings inequality. Similarly, we have established a downward slope in between-firm inequality over a cohort's lifecycle. Since employment has shifted towards older firms, changes in the age composition of firms do not explain the increase in between firm earnings inequality. Hence, in this section we explore cohort and year effects.

5.1 Cohort and year effects in between-firm earnings inequality

Figure 6 shows the cohort and year effects under our two extreme normalizations. Consistent with what might be surmised from Figure 3, the Figure displays the central novel finding of this paper: regardless of normalization, there is a significant rise of cohort effects in between-firm earnings inequality. That is, more recent cohorts have higher between-firm earnings inequality than older cohorts.

The reason the importance of cohort effects persists across the year normalizations can be intuited in Panel B of Figure 4: as we increase the role of year effects, the age effects rotate down, which gives scope for the cohort effects to continue to have explanatory power. In contrast, there is much less clear evidence of cohort effects in mean earnings: the sign of the cohort effects in earnings are sensitive to the normalization we choose.

How quantitatively important are the cohort effects? In order to aggregate our results, we need

to define some intermediate notation. Let $\text{var}_{a,t}(\bar{y})$ be the between-firm variance of earnings of firms of age a in year t and let $\text{var}_t(\bar{y})$ be the overall between-firm variance of earnings in year t. Let $\bar{y}_{a,t}$ be the mean of earnings of firms aged a in year t. Finally, let $s_{a,t}$ be the employment share of firms aged a in year t.

We use the following identity:

$$\text{var}_t(\bar{y}) = \sum_a s_{a,t}\text{var}_{a,t}(\bar{y}) + \sum_a s_{a,t}\left(\bar{y}_{a,t} - (\sum_a s_{a,t}\bar{y}_{a,t})\right)^2. \tag{7}$$

The right hand side consists of two terms. The first term measures the within-cohort inequality across firms, while the second term measures the between-cohort inequality.

Our approach to quantifying the overall contribution of cohort and age effects to the change in the between-firm variance of earnings takes equation (7) and replaces each term with predicted values from our estimates of equation (3), under the two different year normalizations. In order to set a fair basis of comparison, we first compute the between-firm variance of earnings using the estimated values of the age, year, and cohort effects. That is, let $\widehat{\text{var}}_{a,t}(\bar{y})^{a,c,t}$ be the fitted value from using the age, year and cohort effects in the estimation of equation (3), and similarly for $\widehat{\bar{y}}_{a,t}^{a,c,t}$. We then have estimates of the overall between-firm variance of earnings using the fitted values:[29]

$$\widehat{\text{var}}_t(\bar{y})^{a,c,t} = \sum_a s_{a,t}\widehat{\text{var}}_{a,t}(\bar{y})^{a,c,t} + \sum_a s_{a,t}\left(\widehat{\bar{y}}_{a,t}^{a,c,t} - (\sum_a s_{a,t}\widehat{\bar{y}}_{a,t}.^{a,c,t})\right)^2. \tag{8}$$

We compute this value in 1993 and 2013. In 1993 we use firms aged 1 to 17; in 2013 we use firms aged 1 to 37. The reason to change the set of ages we consider is that this matches the sample we use in the estimates underlying Figure 6. By taking the difference between the two years, we have the fitted growth in between-firm earnings inequality:

$$\widehat{\text{var}}_{2013}(\bar{y})^{a,c,t} - \widehat{\text{var}}_{1993}(\bar{y})^{a,c,t}. \tag{9}$$

Given this change in between-firm earnings inequality predicted by the combination of age, cohort, and year effects (and composition, i.e. shares), we subsequently consider the roles of age and cohort effects on their own.

First, we compute the contribution of age effects to the rise of between-firm earnings inequality where we use the fitted values from only the age effects (i.e. omitting cohort and year effects) to compute the same difference:

$$\widehat{\text{var}}_{2013}(\bar{y})^{a} - \widehat{\text{var}}_{1993}(\bar{y})^{a}. \tag{10}$$

[29]Note that since the age, year, and cohort effects for (within-cohort) between-firm earnings inequality and those for mean earnings are estimated separately, they are subject to distinct year slope normalizations. Thus, under the normalizations where the year slopes account for all of the aggregate change in the either the (within-cohort) between-firm earnings inequality or the mean earnings, it is not guaranteed that the implied contribution of year effects to the overall trend in the between-firm variance of earnings perfectly matches the change in the overall between-firm variance.

What varies between these two terms is the shares attached to firms of different ages and the estimated values themselves (since they are based on different ages in the two years). This difference tells us to what degree the changing firm age composition predicts a change in between-firm earnings inequality. Because we include a different set of ages, this term mixes the change in the age distribution and the ages that we include. The reason we might think that this comparison would yield large effects is that we showed in Figure 4 that there are strong lifecycle patterns in the between-firm variance of earnings and we showed in Figure 2 that there are large changes in the age composition of firms: older firms have different inequality than younger firms, and older firms employ an increasing share of workers (Alon et al. (2018) document that such a term is large for productivity). Thus, we might expect that this expression would be large.

Similarly, we compute predicted changes in the between-firm variance of earnings using just the variation in the cohort effects (i.e. omitting age and year effects):

$$\widehat{\text{var}}_{2013}(\bar{y})^c - \widehat{\text{var}}_{1993}(\bar{y})^c. \tag{11}$$

Here, there are two differences between the terms: the shares attached to firms of different cohorts (equivalently, firms of different ages) and the estimated values themselves (since they are based on different cohorts in the two years).[30]

Table 4 reports the results of these various exercises. The first row shows that our fitted values generate an increase in the between-firm variance of earnings of 0.865, which is slightly smaller than the 0.915 we reported for the "1977 cohorts and beyond" sample in Panel B of Table 1. The reason for the slight misalignment is that we estimate the we estimate the cohort mean earnings separately from the between-firm variance of earnings. The next row shows that as we change the year normalization from explaining 0% of the aggregate trend to 100%, not surprisingly, the year effects go from explaining none to (nearly) all of the aggregate change. (When we change the year normalization to explain 100% of the change, the Table shows that year effects only explain 99.2% of the change; the reason for this discrepancy is that we estimate the between-firm variances and mean-cohort earnings separately and so these do not perfectly add up to the aggregate.)

The Table then shows that as we change the year normalization, the contribution of the age effects shifts dramatically. When year effects are constrained to explain none of the aggregate trend, the age effects also explain none of the aggregate trend. Once we increase the explanatory power of the year effects, however, the explanatory power of the age effects turns *negative*. This negative contribution explains why there is room for cohort effects to continue to explain a large share of

[30] We can subsequently decompose the contribution of the cohort effects into a "shift" and a "share" contribution, where the first captures the importance of the fact that newer cohorts are different from older cohorts and the second measures how the relative change in employment at different cohorts matters. Specifically, the "shift" contribution comes from taking $\widehat{\text{var}}_{2013}(\bar{y})^c$ and replacing $s_{a,2013}$ with $s_{a,1993}$ (for the cohorts in 2013 for which we have 1993 employment shares). This simply asks how much of the growth of between-firm variance comes from the changing cohort values, holding fixed the relative employment shares of firms of each age. Meanwhile, the "share" contribution comes from taking $\widehat{\text{var}}_{1993}(\bar{y})^c$ and replacing $s_{a,1993}$ with $s_{a,2013}$; i.e., from Figure 2 we know that the age composition of firms has changed over time. Therefore, it might be simply changing the weights on cohorts would generate large changes.

the aggregate trend.

The final row of Table 4 shows that cohort effects explain a significant portion of the rise of between-firm earnings inequality, and this large role is not dependent on the normalization that we adopt. The Table shows that the contribution of cohort effects to the aggregate trend are either 55% or 107%, depending on which normalization we adopt.

The bottom line is that regardless of the normalization, the change in cohort effects explains a significant fraction of the rise of between-firm earnings inequality: the fact that newer cohorts have higher between-firm pay dispersion accounts for a substantial fraction of the overall trend. We note that this finding is related to Card, Heining, and Kline (2013, Figure 9), though they do not attempt to quantify its importance or explain it. We now turn to trying to shed some light on what is different about subsequent cohorts that could explain these patterns.

5.2 What has changed

We have shown that newer cohorts are substantially different from older cohorts in their between-firm pay dispersion. We now want to understand how else these cohorts are different, with the aim of understanding how these cross-cohort changes in earnings inequality have arisen.

To set the stage for our analysis, in Tables 2 and 3 we present a variety of summary statistics about the labor market and firm dynamics in 1993 and 2013. There are several broad aggregate changes that stand out. First, we can see in terms of racial composition of the workforce that the U.S. has become less white and slightly more female. Second, as has been noted elsewhere, workers are more sorted on the basis of earnings. Commensurate with this sorting on earnings, workers are also more sorted on the basis of college education and age. In contrast, workplaces have become more integrated along lines of race/ethnicity (the one exception is that Black workers are if anything less integrated), as well as gender.

There are several notable trends about between-firm differences: the correlation of earnings and log employment and the variance of log employment have increased. Similarly, the variance of productivity has also increases. Finally, while the mean retention rate of workers has increased, the correlation of the retention rate and pay has decreased; higher paying firms are less likely to retain workers. Following the logic of Sorkin (2018), this last fact is consistent with a rise in the importance of compensating differentials.

To understand the role of cohort effects in driving these patterns, we implement the same age-time-cohort decomposition presented in Section 3 under both normalizations (year slope captures none or all of the aggregate change). We then correlate the outcome-specific cohort effects with the cohort effects for the between-firm variance of earnings under the same normalization. We compute standard errors around the correlation using our residual bootstrap.

Figure 7 summarizes the correlations of the cohort effects of the between-firm variance of earnings and the cohort effects in other variables. The Figure emphasizes a few factors for which cohort effects are significantly correlated with the cohort effects in the between-firm variance of earnings

under both normalizations.[31] First, a number of factors that reflect sorting in the labor market have similar cohort effect patterns to those of between-firm earnings inequality. Specifically, the cohort effects in sorting on the basis of earnings, education (college), and age are all strongly positively correlated with the cohort effects in the between-firm variance of firms: newer cohorts exhibit both stronger sorting and larger between-firm earnings inequality.

Second, a few productivity related variables are strongly related: the cohort effects in the variance of productivity (which is positively correlated) and the correlation of productivity and pay (which is negatively with the cohort effects). These patterns mean that newer cohorts exhibit both higher dispersion in productivity and smaller correlations between productivity and pay, in addition to being more dispersed in terms of pay. In addition, for our alternative measure of productivity—revenue per payroll—we find positive correlation with the mean as well as the variance.

Third, we find that the patterns of selection into exit are related. The cohorts with higher between-firm variance of earnings also have higher—that is, less negative—correlations between exiting and size and exiting and pay. Put differently, the cohorts that are more unequal are also less selected in who continues, which is consistent with the more unequal cohorts also being the cohorts with greater variance of productivity. While looking just at the variance of productivity it might be that these firms are just more different, looking also at patterns in exit is consistent with the view that this represents something closer to misallocation.

Finally, the correlation of the retention rate and pay is negatively correlated with the cohort effects in the between-firm variance of earnings.

Figure 8 presents the cohort effects for a set of factors that are significantly correlated with the between-firm variance of earnings. The Figure shows that the "sorting" variables and the variance of productivity exhibit an upward trend, while the correlation of pay and productivity and the correlation of retention and pay exhibit a downward trend. The magnitudes of these trends are also large. For example, recall that the sorting coefficients are correlations and so an increase of 0.1 to 0.2 is large in correlation space and is larger than the aggregate change documented in Table 2. Similarly, the change in estimated cohort effects in the variance of productivity more than accounts for the aggregate change documented in the summary statistics. The same is true for the large cohort-level decline in the correlation of pay and the retention rate.

To formalize this accounting, we conduct decompositions for these outcomes similar to those in Section 5.1, decomposing the change in the outcomes to the change predicted by year, age, and cohort effects alone.[32] As Table 5 shows, changes in cohort effects account for large shares of

[31] Meanwhile, some outcomes, including within-firm earnings inequality and earnings gaps between college- and high school-educated workers, demographic shares, and most reallocation dynamics measures, are not significantly correlated with the cohort effects in the between-firm variance of earnings under both normalizations.

[32] For the sorting and correlation outcomes, the estimated aggregate outcome depends on five sets of estimates of within-cohort estimates: the covariance of the two underlying variables (e.g., age and mean coworkers' age), the means of the variables, and the variances of the variables. Specifically, for variables x and y, the aggregate correlation is $\text{corr}_t(x,y) = \frac{\text{cov}_t(x,y)}{\sqrt{\text{var}_t(x)}\sqrt{\text{var}_t(y)}}$, where the aggregate variance is defined as in equation (7) and the aggregation covariance is $cov_t(w_i, e_i) = \sum_a s_a cov_{a,t}(x,y) - \sum_a s_a(\bar{x}_t\bar{y}_t - \bar{x}_{a,t}\bar{y}_{a,t})$.

the aggregate changes of these outcomes, regardless of the normalization. For example, the Table shows that the contribution of cohort effects to the aggregate trend in sorting on age are either 41% or 94%, depending on which normalization we adopt. Similarly, the Table shows that the contribution of cohort effects to the aggregate trend in productivity dispersion is between 65% and 119%, depending on the normalization. Collectively, these results demonstrate that the cohort patterns to these outcomes account for significant shares of the aggregate changes.

In summary, many of the broad changes in the U.S. labor market over the last several decades—namely, rising between-firm variance of earnings, increasing sorting workers across firms, and rising productivity dispersion—are all accounted for in part by cohort effects. This fact is novel.

6 Extensions and robustness

6.1 The role of exit and selection into exit

We now explore features of the data that allow us to understand the role of selection into exit in driving the patterns we find.

In Figure 9, we explore the role of selection in generating the cohort patterns for the between-firm variance of earnings. The Figure is constructed in the same way as Figure 5: in Panels A and B we plot the cohort effects in our main sample, in restricted samples where we only use cohorts in their first ten years of life, and in their second ten years of life. The basic message of these panels is that using cohorts aged 11 to 20 gives nearly identical answers to the overall sample. Using cohorts when they are aged 1 to 10 gives slightly steeper cohort profiles.

Turning to Panels C and D, we look at the balanced panel. Looking first at the aged 11 to 20 year line we see that these are very similar to the main estimates, which implies that selection at these ages does not play an important role in these cohort trends. Looking at the aged 1 to 10 line, we see that there is some suggestive evidence that selection at younger ages plays some role in explaining the cohort patterns. Specifically, if we estimate the cohort effects using the balanced panel of younger firms, then we see that the cohort patterns are slightly shallower than our main estimates. This change is evidence that differential patterns of selection across cohorts plays some role in explaining our findings. If we compare to the unbalanced panel aged 1 to 10—rather than to the main estimates—we can see that the magnitude of the gap is larger (since the cohort effects based on the unbalanced panel lie above the main estimates). Thus, the bottom-line from Panels C and D is that there has been some change in selection into exit across cohorts that is concentrated in the first ten years of firms' lifecycles.[33]

By contrasting the patterns of cohort effects in the balanced and unbalanced panels, we have concluded that selection into exit has changed across cohorts. This statement reflects a combination of two factors: is it changing probabilities of exit or changing selection into exit? Figure 10 presents a variety of evidence to shed light on this question. When we look at worker-weighted regressions

[33]Interestingly, we find no evidence of selection driving the patterns we find in the variance of productivity. See Appendix Figure A5.

(see Panel A), we find some evidence that firms in younger cohorts are more likely to exit. Similarly, in Panels C and E we find that there is some evidence of a positive trend in the cohort effects in the worker-weighted correlation between exit and pay and exit and size. Higher-paying and larger firms are less likely to exit, so the positive trend says that in younger cohorts this positive correlation has weakened. Turning to the firm-weighted patterns in Panels B, D and F we find that some of these trends are weaker or reversed. Conceptually, we prefer the worker-weighted analysis in that it parallels the rest of the analysis and aggregates to the experience of the "typical" worker.

In summary, we find some role for changing patterns of selection in driving our cohort results, and some evidence that this combines changes in the probability of exit and the selection into exit, where we define selection in terms of correlations of exit and pay and size.

6.2 New hires vs. all workers

One important hypothesis for why firm cohorts have changed is that the workers they employ were hired at different point in time. Specifically, younger cohorts of firms hired workers more recently than did older firms and so what we are labelling *firm* cohorts might instead reflect differences across *worker* cohorts. A first way to assess this possibility is to restrict attention to new hires (that is, workers who have earnings at the employer in year t and not in year $t-1$). Figure A6 presents the estimated cohort and year effects for between-firm variance and mean earnings. As the figure shows, the results for new hires are very similar to those for all workers. In particular, the cohort and year effects for between-firm earnings inequality are extremely similar across the two samples. This finding implies that the phenomenon we document is not about changes in the composition of workers available to firms.

6.3 Worker cohorts

A second way of assessing the hypothesis that the firm cohort patterns we document are about worker cohorts is to see whether there are analogous patterns in the worker cohorts. If newer cohorts of workers exhibit higher earnings inequality *and* are segregated from older cohorts across firms, it is possible that worker cohort patterns could underly firm cohort patterns.

Guvenen et al. (2019) document substantial worker cohort patterns to median lifetime earnings, with relatively newer cohorts of women having higher median earnings but relatively newer cohorts of men having lower median earnings; they also find that newer cohorts of both men and women exhibit higher lifetime earnings inequality.

We replicate our regression analysis at the worker level, where we estimate worker cohort, worker age, and year effects for two outcomes: mean earnings and variance of earnings. We study workers aged 25-60 (i.e. those in our sample), who are born in (and thus belong to cohorts) 1933 through 1988.[34] Figure A7 presents the estimated effects. Life cycle patterns appear to dominate both

[34] Guvenen et al. (2019) study cohorts that enter the labor market, i.e. are 25, between 1957 and 1983, and track workers until they are 55. This means that they study individuals born (by our definition of age) between 1932 and 1968.

mean and variance of earnings: older workers tend to earn more and have higher inequality. We find some evidence that newer worker cohorts also have higher mean and variance of earnings, but these patterns are not consistent across the two year slope normalizations.

Inspired by Guvenen et al. (2019), we also estimate these models separately for men and women. We present these results in Figures A8 and A9, respectively. For both men and women, life cycle effects account for a lot of the patterns in average earnings and earnings inequality. However, we find larger cohort effects for women (under the normalization in which the year effects capture none of the change). Consistent with Guvenen et al. (2019), newer female cohorts tend to have both higher average earnings and inequality.

Taken together, these results demonstrate that worker-level analysis of earnings inequality *is* informative: older workers tend to have higher earnings inequality, and it is possible (but inconclusive) that newer worker cohorts do as well. However, these worker cohort patterns do not explain our results for firm cohorts. Rather, as Figure 7 shows, we do not find any systematic relationship between firm cohorts that exhibit higher between-firm earnings inequality and those that hire proportionally either more older workers (as captured by the mean worker age) or more women.[35]

7 Discussion

We have documented that much of the between-firm earnings inequality can be accounted for by the fact that firms in newer cohorts are more disparate in their pay than those in older cohorts. Further, these cohort patterns correlate with cohort patterns in worker sorting and retention and productivity. What do these results mean for our understanding of inequality and business dynamics?

7.1 Firms and inequality

Let us first consider the literature on firms and inequality. The broad theme that emerges from considering particular hypotheses for the increase in the between-firm earnings inequality is that while the simplest versions of most explanations for rising between-firm earnings inequality do not feature cohort effects, it is fairly easy to imagine natural extensions of these explanations/models that would generate cohort effects. The key feature of the "cohort-effect-augmented" version of these explanations is that they feature slow diffusion of the relevant change.

For example, Song et al. (2019, pg. 6-7) discuss several mechanisms for the rise in between-firm earnings inequality. First, there could be an increase of outsourcing which leads to increasing segregation and sorting of workers, as diverse workplaces split into homogeneous sub-companies. Second, rising returns to skill would lead to rising sorting and segregation of workers on the basis

[35] We find virtually no cohort effects for the female shares at firms. We find inconclusive evidence for cohort effects for the mean worker age at firms: under the normalization in which the year slope accounts for all of the aggregate trend in mean worker age at firms, we estimate virtually zero cohort effects; for the normalization in which the year slope accounts for none of the aggregate trend, we find large effects for newer cohorts.

of income if workers were already somewhat sorted. Third, declines in unionization or changes in corporate culture could lead to a decline in access to high-paying jobs for low-wage workers.

We can generate versions of each of these mechanisms that are consistent with cohort effects. If we conceive of outsourcing, as it is typically studied in empirical papers, as a set of existing companies outsourcing an existing stock of workers (Goldschmidt and Schmieder (2017)), then this mechanism does not explain the presence of cohort effects. Of course, if instead part of the effect of outsourcing is not through changes in existing companies but in how new companies are designed and run (i.e., from the beginning a firm outsources its janitorial staff), then this would show up as a cohort effect. In its simplest form—e.g., Katz and Murphy (1992)—a general rise in the return to skill would affect all workers at all firms in the same way and so would show up as a time effect. Of course, one could develop extended versions of models of skill biased technical change where the technology is embodied in the firm from its founding and so technical change comes out as a cohort effect (e.g., in vintage capital models). Finally, in terms of unionization and corporate culture a similar theme emerges. If we conceive of this hypothesis as operating through all unions declining in sync, then this would not generate cohort effects. In contrast, if younger firms are less likely to be unionized, then this could generate these types of cohort patterns (as Card, Heining, and Kline (2013) emphasize).

What is distinctive in thinking about cohort effects versus year effects is thus that these explanations have different implications for the future because cohort effects imply slow diffusion of changes. While year effects affect all firms in the same way, cohort effects imply something distinct happening when firms enter.

We explore the implications of the cohort patterns on future inequality with simple projections of our estimated effects, which we show in Figure 11. In these projections, we estimate the level of overall between-firm earnings inequality (see the decomposition in equation (7)) after 2013, under several assumptions. First, we assume that the employment distribution across firms' ages remains fixed at the 2013 level.[36] Second, we assume that after 2013, each year has the same year effect as 2013, and each new cohort has the same cohort effect as the 2013 cohort (and age effects are unchanged). By doing this, we isolate the effect of slowly replacing the older cohorts with newer cohorts, who behave like the 2013 cohort. Given these assumed employment shares and year, cohort, and age effects, we predict the overall between-firm earnings inequality in each year after 2013, as in equation (8), for both year slope normalizations. We compare the resulting patterns in earnings inequality to the observed time series (from Figure 1); we normalize each series' level such that the 2013 between-firm earnings inequality is 0, in order to improve ease of comparison.

Figure 11 highlights the descriptive "fit" of the additively separable model, in that our estimated effects fit the data well.

Figure 11 also shows that as we replace older cohorts with newer cohorts, between-firm earnings inequality continues to rise. In our projections, the between-firm variance of earnings increases by an additional 0.08 or 0.02 log from 2013 to 2030 depending on year normalization. These projected

[36] This means that, e.g., in 2014 we consider firms aged 1 to 37, i.e. cohorts 1978 to 2014.

increases are large relative to the the rise in between-firm earnings inequality in the 1977 cohorts and beyond sample of 0.09 log points from 1993 to 2013 (see Table 1).

The bottom line is that the slow diffusion of the inequality "technology" implies that all else equal we would expect inequality to continue to rise even without further innovation in this technology. In a retrospective sense, the cohort perspective also sheds new light on when "innovation" in the inequality technology was fastest, which potentially differs from when inequality itself was changing fastest. This distinction arises because cohorts enter small relative to the whole economy and gradually grow, and so technology innovation across cohorts takes time to contribute to aggregate inequality.

7.2 Business dynamism

Our results also have implications for the literature on business dynamism, which is interested in why entry and other measures of reallocation have declined.[37] A typical view in this literature (e.g., Decker et al. (2020)) is that, from a static perspective, a natural measure of misallocation is to look at productivity dispersion. Insofar as productivity dispersion increases, then this pushes towards wanting to interpret the decline of dynamism as a negative development in the U.S. economy. Indeed, Decker et al. (2016) and Barth et al. (2016) find evidence of increased productivity dispersion overall. What is striking in our results relative to this literature is that we find cohort effects in increased productivity dispersion: newer cohorts exhibit higher variances of labor productivity than older cohorts. We believe this finding is new in the literature.

Our findings of cohort patterns in productivity dispersion, in addition to pay dispersion, suggest that there is something important about when firms enter. This finding is consistent with the literature that emphasizes changes in the process governing entry (e.g., Sterk, Sedláček, and Pugsley (2021), Karahan, Pugsley, and Sahin (2021), and Hopenhayn, Neira, and Singhania (2019)); some "technology" has changed about how new firms enter that has made newer cohorts different from older ones.

Our results suggest that rising between-firm earnings inequality is not driven by declining firm entry itself. The reason is that the trend of employment shifting towards older firms (see Figure 2) actually predicts a decline in between-firm variance of earnings. This point is illustrated in Table 4: all else equal, a shift in employment towards older firms (e.g. because of declining firm entry) predicts a decline in between-firm earnings inequality, both because older firms exhibit lower pay dispersion (age effects trend down across age, as seen in Figure 4) and because older cohorts exhibit lower pay dispersion (cohort effects trend up from older to newer cohorts, as seen in Figure 6). Thus, our results suggest that declining firm entry, by itself, actually offsets the rise in inequality. Instead, some aspect of the process by which new firms enter has changed, such that the firms that enter as newer cohorts are different from those who entered previously. The decline in firm entry is perhaps a symptom of this change, but does not by itself explain the rise in between-firm earnings

[37] Decker et al. (2020, pg. 3952) point to declines in the "rate of business entry, the prevalence of high-growth firm outcomes, the rate of internal migration and the rate of job and worker reallocation."

inequality.

What exactly has changed about the firm entry process is not obvious. If it were simply easier to enter (e.g. lower entry or operating costs), then newer firms would be less selected and thus more dispersed in terms of their productivity. In this case, however, we would expect higher rates of firm entry, which is not consistent with aggregate firm aging. Further, general changes in entry and operating costs might be expected to affect incumbent firms as well, as argued by Karahan, Pugsley, and Sahin (2019).[38] One possibility is that newer cohorts face different labor markets, for instance due to labor supply aging (see e.g., Karahan, Pugsley, and Sahin (2019) and Engbom (2019)). If hiring at entry is particularly important for the path a firm takes, e.g. because of high adjustment costs, the state of the aggregate labor market may affect newer firms differently from older firms. If this were the case, then we may expect the cohort effects to be driven by the workers who are hired at entry. In fact, as we showed, if we restrict our analysis to workers who are new employees to firms in each year, we observe very similar cohort effects for between-firm earnings inequality. Another possibility is that firms in different cohorts exit at different rates across the lifecycle, such that incumbent firms may be differentially selected across cohorts. We have shown some evidence along these lines.

The crucial takeaway is that whatever "technology" has changed over time must affect firms of different ages (i.e. cohorts in a given year) differently in order to be consistent with our cohort effects; i.e. aggregate shocks diffuse slowly across cohorts, rather than affecting all firms at once. Overall, our findings suggest that there is a connection between firm entry, productivity dispersion, and labor market outcomes, since we see correlative cohort patterns across productivity dispersion, between-firm earnings inequality, and sorting of workers across firms. This connection tells us that the decline in business dynamism (and change in firm entry) is likely connected with the rise of earnings dispersion across firms.

8 Conclusion

In this paper, we document some important new facts about the rise of earnings inequality in the United States. Our novel lens is to look at cohorts of firms. First, we find that between-firm pay dispersion declines within a cohort. Second, there are striking cohort patterns: more recent cohorts are more dispersed than older cohorts. This pattern accounts for a large fraction of the aggregate rise in between-firm earnings inequality. Third, many other notable labor market facts also exhibit these cohort patterns: workers in more recent firm cohorts are more sorted on the basis of earnings, age, and college attainment. Furthermore, these firm cohorts have more dispersed productivity.

Taken together, our results emphasize that earnings inequality has diffused slowly through new cohorts. So, as older cohorts are replaced with cohorts with inequality "technology" of the more recent vintage, we expect inequality to continue to rise, even without a change in that underlying

[38] It is conceivable that incumbent firms may be affected less than new firms, which could produce patterns consistent with our cohort effects. The inconsistency arises if all firms are affected equally in a given year or at a given age.

technology.

Our results also emphasize a connection between changes in firm entry emphasized by the decline in dynamism literature and the rise of earnings inequality. The exact details of this connection await future research.

References

Abowd, John M., Paul Lengermann, and Kevin L. McKinney. 2003. "The Measurement of Human Capital in the U.S. Economy." Working paper.

Abowd, John M., Bryce E Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L McKinney, Marc Roemer, and Simon Woodcock. 2009. "The LEHD infrastructure files and the creation of the Quarterly Workforce Indicators." In *Producer dynamics: New evidence from micro data*. University of Chicago Press, 149–230.

Akcigit, Ufuk and Sina T Ates. 2019. "What Happened to US Business Dynamism?" Tech. rep., National Bureau of Economic Research.

Alon, Titan, David Berger, Robert Dent, and Benjamin Pugsley. 2018. "Older and slower: The startup deficits lasting effects on aggregate productivity growth." *Journal of Monetary Economics* 93:68–85.

Barth, Erling, Alex Bryson, James C. Davis, and Richard Freeman. 2016. "Its Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States." *Journal of Labor Economics* 34 (S2):S67–S97.

Berger, David W., Kyle F. Herkenhoff, and Simon Mongey. 2019. "Labor market power." Tech. rep., National Bureau of Economic Research.

BLS. 1997. *U.S. Bureau of Labor Statistics Handbook of Methods*.

Borovičková, Katarína and Robert Shimer. 2017. "High wage workers work for high wage firms." Tech. rep., National Bureau of Economic Research.

Burgess, Simon, Julia Lane, and David Stevens. 2000. "Job flows, worker flows, and churning." *Journal of labor economics* 18 (3):473–502.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and labor market inequality: Evidence and some theory." *Journal of Labor Economics* 36 (S1):S13–S70.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace heterogeneity and the rise of West German wage inequality." *Quarterly Journal of Economics* 128 (3):967–1015.

Davis, Steven J, John Haltiwanger, and Scott Schuh. 1996. "Small business and job creation: Dissecting the myth and reassessing the facts." *Small business economics* 8 (4):297–315.

Deaton, Angus. 1997. *The Analysis of Household Surveys : A Microeconometric Approach to Development Policy*.

Decker, Ryan A, John C. Haltiwanger, Ron S Jarmin, and Javier Miranda. 2016. "Declining business dynamism: What we know and the way forward." *American Economic Review* 106 (5):203–07.

———. 2018. "Changing business dynamism and productivity: Shocks vs. responsiveness." Tech. rep., National Bureau of Economic Research.

———. 2020. "Changing business dynamism and productivity: Shocks versus responsiveness." *American Economic Review* 110 (12):3952–90.

Engbom, Niklas. 2019. "Firm and worker dynamics in an aging labor market." *unpublished paper* .

Engbom, Niklas, Christian Moser, and Jan Sauermann. 2021. "Firm pay dynamics." .

Foster, Lucia, John C. Haltiwanger, and Cornell John Krizan. 2001. "Aggregate productivity growth: lessons from microeconomic evidence." In *New developments in productivity analysis*. University of Chicago Press, 303–372.

Goldschmidt, Deborah and Johannes F Schmieder. 2017. "The rise of domestic outsourcing and the evolution of the German wage structure." *The Quarterly Journal of Economics* 132 (3):1165–1217.

Guvenen, Fatih, Greg Kaplan, Jae Song, and Justin Weidner. 2019. "Lifetime incomes in the United States over Six Decades." Working paper.

Hall, Robert E. 1968. "Technical Change and Capital from the Point of View of the Dual." *Review of Economic Studies* 35 (1):35–46.

Haltiwanger, John, Ron Jarmin, Robert Kulick, Javier Miranda, Veronika Penciakova, Cristina Tello-Trillo et al. 2019. "Firm-level Revenue Dataset." Tech. rep., Center for Economic Studies, US Census Bureau.

Haltiwanger, John C., Ryan Decker, and Ron Jarmin. 2015. "Top ten signs of declining business dynamism and entrepreneurship in the US." In *Kauffman Foundation New Entrepreneurial Growth Conference*. Citeseer.

Haltiwanger, John C., Ron S. Jarmin, Robert Kulick, and Javier Miranda. 2017. "High Growth Young Firms: Contribution to Job, Output, and Productivity Growth." In *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*, edited by John C. Haltiwanger, Erik Hurst, Javier Miranda, and Antoinette Schoar. University of Chicago Press, 11–62.

Haltiwanger, John C., Ron S. Jarmin, and Javier Miranda. 2013. "Who creates jobs? Small versus large versus young." *Review of Economics and Statistics* 95 (2):347–361.

Hopenhayn, Hugo, Julian Neira, and Rish Singhania. 2019. "From population growth to firm demographics: Implications for concentration, entrepreneurship and the labor share." Working paper.

Hsieh, Chang-Tai and Peter J. Klenow. 2014. "The life cycle of plants in India and Mexico." *Quarterly Journal of Economics* 129 (3):1035–1084.

Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin. 2019. "Granular search, market structure, and wages." Tech. rep., National Bureau of Economic Research.

Karahan, Fatih, Benjamin Pugsley, and Ayseguil Sahin. 2021. "Demographic Origins of the Startup Deficit." Tech. rep.

Karahan, Fatih, Benjamin Pugsley, and Aysegul Sahin. 2019. "Demographic origins of the startup deficit." Working paper.

Katz, Lawrence F and Kevin M Murphy. 1992. "Changes in relative wages, 1963–1987: supply and demand factors." *The quarterly journal of economics* 107 (1):35–78.

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar. 2019. "Who profits from patents? rent-sharing at innovative firms." *The Quarterly Journal of Economics* 134 (3):1343–1404.

Kornfeld, Robert and Howard S. Bloom. 1999. "Measuring Program Impacts on Earnings and Employment: Do Unemployment Insurance Wage Reports from Employers Agree with Surveys of Individuals." *Journal of Labor Economics* 17 (1):168–197.

Kremer, Michael and Eric Maskin. 1996. "Wage inequality and segregation by skill." Tech. rep., National bureau of economic research.

Lachowska, Marta, Alexandre Mas, Raffaele D Saggio, and Stephen A Woodbury. 2020. "Do firm effects drift? Evidence from Washington administrative data." Tech. rep., National Bureau of Economic Research.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler. 2019. "Imperfect competition, compensating differentials and rent sharing in the US labor market." Tech. rep., National Bureau of Economic Research.

MacKinnon, James G. 2006. "Bootstrap methods in econometrics." *Economic Record* 82:S2–S18.

McKenzie, David J. 2006. "Disentangling age, cohort and time effects in the additive model." *Oxford Bulletin of Economics and Statistics* 68 (4):473–495.

Pugsley, Benjamin, Petr Sedlacek, and Vincent Sterk. 2019. "The Nature of Firm Growth." .

Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter. 2019. "Firming up inequality." *The Quarterly Journal of Economics* 134 (1):1–50.

Sorkin, Isaac. 2018. "Ranking firms using revealed preference." *Quarterly Journal of Economics* 133 (3):1331–1393.

Spletzer, James R. and John C. Haltiwanger. 2020. "Between firm changes in earnings inequality: the dominant role of industry effects." Working Paper 26786, NBER.

Sterk, Vincent, Petr Sedláček, and Benjamin Pugsley. 2021. "The nature of firm growth." *American Economic Review* 111 (2):547–579.

Vilhuber, Lars et al. 2018. "LEHD Infrastructure S2014 files in the FSRDC." *US Census Bureau, Center for Economic Studies Discussion Papers, CES* .

Table 1: Earnings inequality: 1993 and 2013

	Panel A: Full sample			1993 firm cohort		
	1993 Value	2013 Value	Change	1993 Value	2013 Value	Change
Total	0.6140	0.7809	0.1669	0.6495	0.7789	0.1294
Between-Firm	0.2672	0.3844	0.117	0.4002	0.3829	-0.0173
Share of total			*0.7021*			*-0.1337*
Within-Firm	0.3468	0.3965	0.0497	0.2493	0.3960	0.1467
Share of total			*0.2979*			*1.1337*
Between-firm: Δ sector share only	0.2672	0.2727		0.4003	0.4224	
N(workers)	14,440,000	19,340,000		233,000	256,000	
N (firms)	996,000	1,233,000	65,000	19,000		

	Panel B: 1977 cohorts and beyond		
	1993 Value	2013 Value	Change
Total	0.6344	0.7626	0.1282
Between-Firm	0.3020	0.3935	0.0915
Share of total		*0.7137*	
Within-Firm	0.3324	0.3690	0.0366
Share of total		*0.2854*	
N(workers)	7,391,000	13,590,000	
N (firms)	829,000	1,170,000	

Notes: This table shows statistics about earnings inequality in 1993 and 2013. Panel A presents statistics for all workers at firms appearing in 1993 or 2013, with the subset of these firms who entered in 1993 in the right margin; Panel B restricts to firms entering in 1977 or after. Each panel presents the decomposition of the variance of earnings into between-firm dispersion in average pay and within-firm dispersion (equation (7)). Panel A additionally presents the "share" from a shift-share decomposition of the between-firm dispersion across sectors (where between-firm dispersion equals the weighted average of within-sector between-firm dispersion plus variation in sectors' mean earnings); the 2013 value is based on the between-firm variance of earnings and mean earnings within sectors in 1993, where sector weights are employment shares in 2013.

Table 2: Descriptive statistics on pay and demographics: 1993 and 2013

	1993 Value	2013 Value
Mean earnings	10.48	10.59
Mean college-high school pay gap	0.4356	0.5818
Mean age	39.38	42.17
Mean share college	0.2637	0.2905
Mean share White	0.7164	0.5947
Mean share Black	0.0927	0.1086
Mean share Hispanic	0.1226	0.1891
Mean share Asian	0.0544	0.0898
Mean share Native American	0.0044	0.0040
Mean share female	0.4574	0.4753
Sorting on earnings	0.5895	0.6568
Sorting on college (non-imputed)	0.5482	0.5962
Sorting on age	0.2200	0.2532
Sorting on White	0.4558	0.4304
Sorting on Black	0.3870	0.3900
Sorting on Hispanic	0.4728	0.4112
Sorting on Asian	0.4266	0.4161
Sorting on Native American	0.1284	0.1153
Sorting on female	0.4465	0.4206
N(workers)	14,440,000	19,340,000
N (firms)	996,000	1,233,000

Notes: This table presents summary statistics for firms in the post-1977 sample in 1993 and 2013. These statistics include mean earnings at firms, mean demographics, and sorting on demographics (measured by equation (1)), all worker-weighted.

Table 3: Descriptive statistics on size, productivity, exit, and dynamics: 1993 and 2013

	1993 Value	2013 Value
Mean firm employment (firm-weighted)	14.85	16.20
Corr. of earnings and log employment (“”)	0.3433	0.2954
Variance of log employment (“”)	1.376	1.428
Mean productivity (firm-weighted, 1997 and 2013)	-0.2038	-0.2806
Variance of productivity (“”)	0.8720	1.073
Corr. productivity and size (“”)	-0.0407	-0.0021
Corr. productivity and pay(“”)	0.2293	0.2130
Mean log(revenue/payroll)	-0.2038	-0.2806
Variance of log(revenue/payroll) (“”)	0.8720	1.073
Corr. log(revenue/payroll) and size (“”)	-0.0407	-0.0021
Corr. log(revenue/payroll) and pay(“”)	0.2293	0.2130
Exit rate (1993 and 2012)	0.0857	0.0535
Corr. exit and size	-0.1421	-.01198
Corr. exit and pay	-0.0939	-0.0689
Exit rate (firm-weighted, 1993 and 2012)	0.1789	0.1322
Corr. exit and size (“”)	-0.2070	-0.1922
Corr. exit and pay (“”)	-0.1128	-0.1038
Corr. exit and productivity (firm-weighted, 1997 and 2012)	-0.1059	-0.1100
Corr. exit and log(revenue/payroll) (“”)	-0.0016	-0.0115
Mean share new hires (1994 and 2013)	0.3805	0.3267
Mean retention rate (1993 and 2012)	0.6938	0.7356
Corr retention rate and pay (“”)	0.2812	0.2568
Mean DHS employment growth (1994 and 2013)	0.1392	0.1134
Mean job reallocation rate (“”)	0.1492	0.1156
Mean workers flows (“”)	465.8	1557
Mean worker flow rate (“”)	0.3374	0.2885
Mean excess reallocation (“”)	0.1882	0.1729
N(workers)	14,440,000	19,340,000
N (firms)	996,000	1,233,000

Notes: This table presents summary statistics for firms in the post-1977 sample in 1993 and 2013 (or other years, when noted). The table presents means, variances, and correlations of firm size (employment), productivity, exit, and dynamics. Values are worker-weighted unless noted otherwise.

Table 4: Contribution of year, age, and cohort effects to changes in the between-firm variance of earnings

	Year effects explain ___ of aggregate change: none			*Year effects explain ___ of aggregate change:* all		
	1993 Level	2013 Level	Share of total change (%)	1993 Level	2013 Level	Share of total change (%)
Total	0	0.0865	100.0	0	0.0866	100.0
Year effects (Δ year)	-0.0317	-0.0317	0.0	0.03698	0.1229	99.2
Age effects (Δ ages and age distribution)	-0.0834	-0.0845	-1.3	-0.0493	-0.0981	-56.5
Cohort effects (Δ cohorts and age distribution)	0.0548	0.1476	107.2	0.0853	0.1336	55.7

Notes: This table presents aggregations based on the main regression estimates (all ages, 1-37). We decompose the predicted change in between-firm earnings inequality from the age, cohort, and year effects for within-cohort between-firm earnings variance and within-cohort mean earnings, estimated on the post-1977 sample under the two year slope normalizations. The aggregate change between 1993 and 2013 is 0.0915; we underestimate this change because we separately estimate the between-firm earnings variance and mean earnings.

The first row presents the total predicted between-firm earnings variance in each year (equation (8)) and the change (equation (9)). In this row and all subsequent ones, we normalize levels by subtracting off the 1993 total value (within panel), such that the 1993 value is 0. "Share of total change" refers to the 1993 to 2013 change for a given row, divided by the total change (first row), for each normalization.

The second row presents values when we only use year effects for prediction. Note that the year effects change does not quite account for 100% in the second panel (where year effects account for all the change), because we are summing across the two separate sets of estimates (between-firm earnings variance and mean earnings).

The third row presents the values when we only use age effects for prediction. Note that both the ages present in each year and the employment shares at each age change.

The fourth row presents the values when we only use cohort effects for prediction. Note that both the cohorts present in each year and the employment shares for each cohort change.

Table 5: Contribution of year, age, and cohort effects to aggregate changes

	Year effects explain ___ of aggregate change: none			*Year effects explain ___ of aggregate change:* all		
	1993 Level	2013 Level	Share of total change (%)	1993 Level	2013 Level	Share of total change (%)
Panel A: Sorting on earnings (actual change: 0.074)						
Total	0	0.0736	100.0	0	0.0737	100.0
Year effects (Δ year)	-0.0678	-0.0678	0.0	0.0011	0.0737	98.6
Age effects (Δ ages and age distribution)	-0.0689	-0.0582	14.5	-0.0320	-0.0686	-50.1
Cohort effects (Δ cohorts and age distribution)	0.0036	0.0728	94.0	0.0307	0.0610	41.1
Panel B: Sorting on age (actual change: 0.0306)						
Total	0	0.0246	100.0	0	0.0284	100.0
Year effects (Δ year)	0.0031	0.0031	0.0	0.0243	0.0470	80.0
Age effects (Δ ages and age distribution)	-0.0103	0.0125	92.2	-0.0097	-0.0263	-58.5
Cohort effects (Δ cohorts and age distribution)	0.0253	0.0601	141.2	0.0326	0.0482	54.8
Panel C: Sorting on college (actual change: 0.014)						
Total	0	0.0142	100.0	0	0.0131	100.0
Year effects (Δ year)	-0.0006	-0.0006	0.0	0.0102	0.0235	101.0
Age effects (Δ ages and age distribution)	-0.0225	-0.0290	-45.5	-0.0145	-0.0275	-98.6
Cohort effects (Δ cohorts and age distribution)	0.0219	0.0421	141.7	0.0247	0.0366	91.0
Panel D1: Variance of productivity (1997-2013, actual change: 0.190)						
Total	0	0.1314	100.0	0	0.1313	100.0
Year effects (Δ year)	0.3638	0.3638	0.0	0.5270	0.6576	99.4
Age effects (Δ ages and age distribution)	-0.0671	-0.0859	-14.3	0.0179	-0.0677	-65.1
Cohort effects (Δ cohorts and age distribution)	0.4338	0.5899	118.8	0.5092	0.5951	65.4
Panel D2: Variance of productivity (1997-2013, actual change: 0.190; firm-weight for aggregation)						
Total	0	0.2282	100.0	0	0.2281	100.0
Year effects (Δ year)	0.2962	0.2962	0	0.4594	0.5899	57.2
Age effects (Δ ages and age distribution)	-0.0873	-0.1121	-10.9	0.0154	-0.0439	-26.0
Cohort effects (Δ cohorts and age distribution)	0.3855	0.6420	112.4	0.4441	0.6005	68.6
Panel E: Correlation of retention and pay (1993-2012, actual change: -0.037)						
Total	0	-0.0371	100.0	0	-0.0372	100.0
Year effects (Δ year)	0.0226	0.0226	0	0.0149	-0.0007	41.9
Age effects (Δ ages and age distribution)	0.0158	0.0469	-83.9	0.0125	0.0920	-213.5
Cohort effects (Δ cohorts and age distribution)	0.0140	0.0043	26.1	0.0074	0.0009	17.5
Panel F: Correlation of exit and pay (1993-2012, actual change: 0.025)						
Total	0	0.0248	100.0	0	0.0204	100.0
Year effects (Δ year)	-0.0042	-0.0042	0	0.0223	0.0453	87.6
Age effects (Δ ages and age distribution)	-0.0569	-0.1201	-255.1	-0.0284	-0.1098	-310.7
Cohort effects (Δ cohorts and age distribution)	0.0226	0.0540	126.7	0.0377	0.0586	80.1

Notes: This table presents aggregations based on the main regression estimates (all ages, 1-37). We decompose the predicted change in outcomes from the age, cohort, and year effects for within-cohort outcomes, estimated on the post-1977 sample under the two year slope normalizations with worker-weighting unless noted. The aggregate change between the first (1993 unless noted) and last (2013 unless noted) year is listed in the panel headers; we imperfectly estimate these changes because we separately estimate the underlying variables (covariances, variances, and mean).

The first row presents the total predicted outcome in each year (equation (8)) and the change (equation (9)). In this row and all subsequent ones, we normalize levels by subtracting off the 1993 total value (within panel), such that the first year value is 0. "Share of total change" refers to the first to last year change for a given row, divided by the total change (first row), for each normalization.

The second row presents values when we only use year effects for prediction. Note that the year effects change does not quite account for 100% in the second panel (where year effects account for all the change), because we are summing across the multiple separate sets of estimates (between-firm earnings variance and mean earnings).

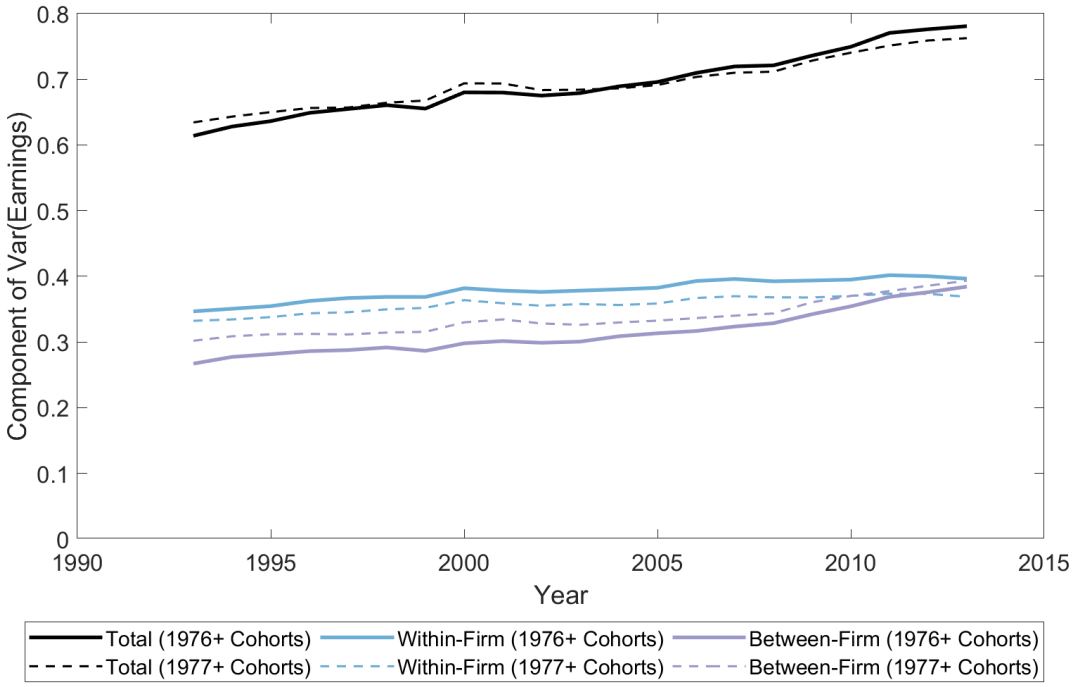
The third row presents the values when we only use age effects for prediction. Note that both the ages present in each year and the employment shares at each age change.

The fourth row presents the values when we only use cohort effects for prediction. Note that both the cohorts present in each year and the employment shares for each cohort change.

For correlations, aggregate correlation is the ratio of predicted covariance to the square root of the product of predicted variances. We normalize the levels of the predicted variances by adding constants (within each year slope normalization) to the estimates to minimize the difference between the predicted correlation change and the actual change (while keeping the constants of similar orders of magnitude).

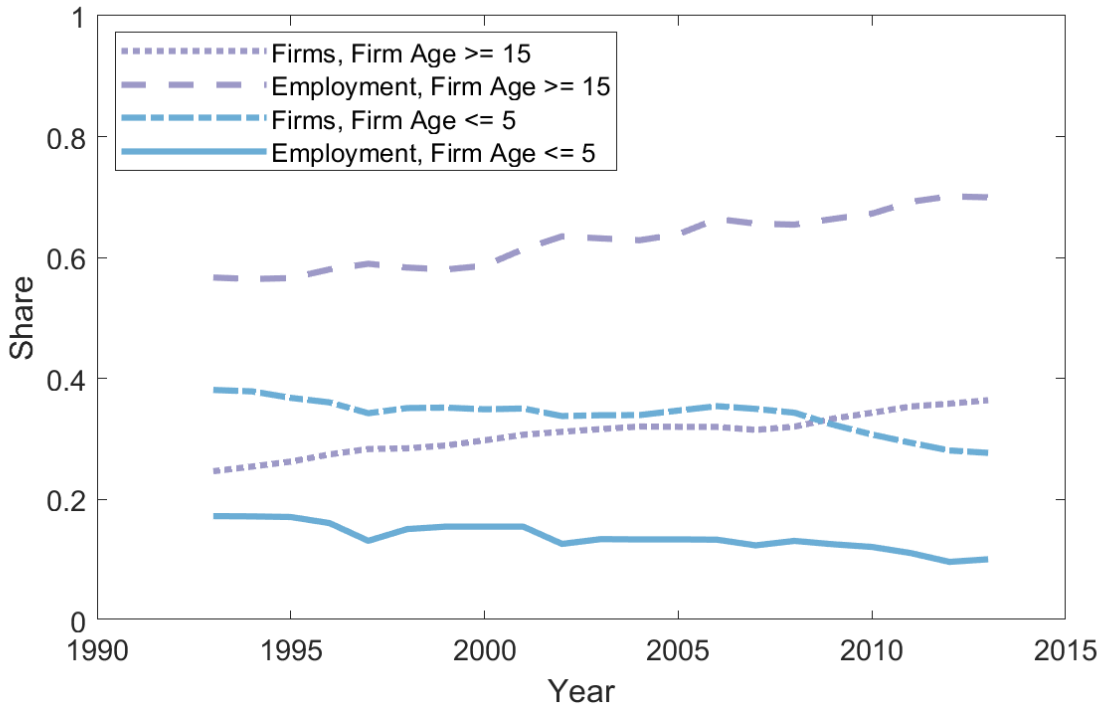
Panels D1 and D2 show the decomposition of aggregate productivity dispersion based on worker-weighting (D1) and firm-weighting (D2) across cohorts. While the outcome is firm-weighted within-cohort, our estimated effects are based on employment-weighting across cohorts. We include both Panels for completeness.

Figure 1: Aggregate Trends in Earnings Inequality



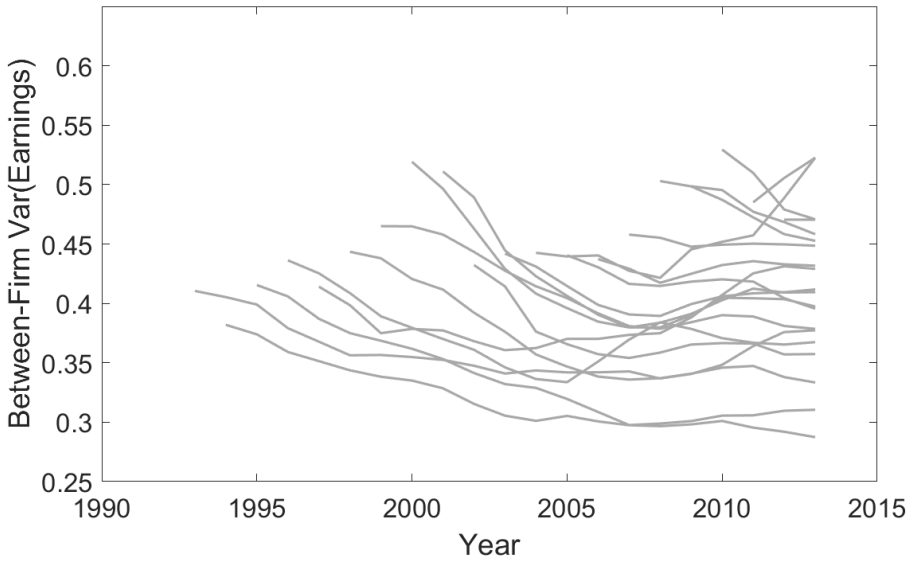
Notes: This figure presents the aggregate trends in earnings inequality for two samples: workers at all firms in our full sample ("1976+ Cohorts") and workers at firms in the 1977 cohort and beyond sample ("1977+ Cohorts"). We decompose the total variance of earnings into the dispersion in average pay at firms (between-firm) and the dispersion of pay Within-firms (within-firm) (see equation (2)).

Figure 2: Changing age composition of employment

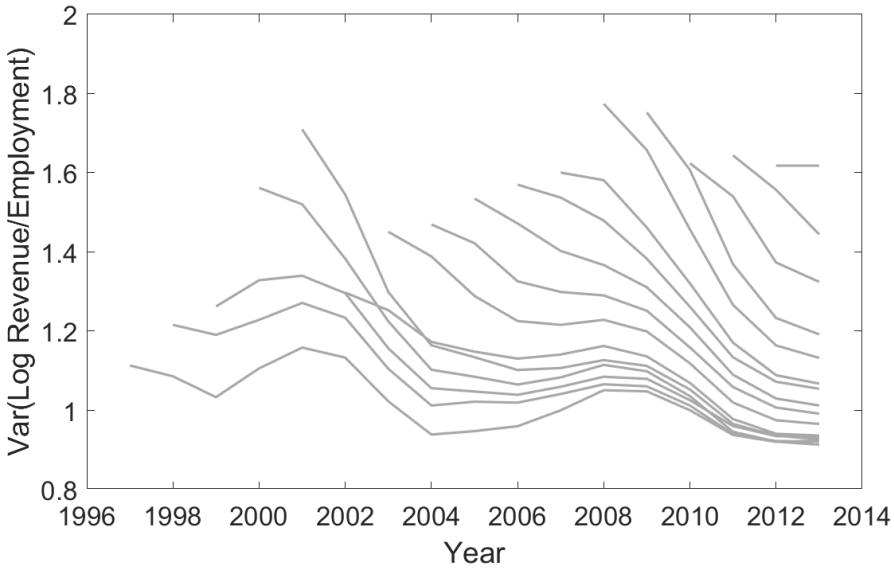


Notes: This figure presents the aggregate trends in the share of firms or employment accounted for by young (aged 1 to 5) and old (aged 15+) firms, for our full sample. Note that young firms tend to be disproportionately small, and so the share of firms that are young is larger than the share of employment that is accounted for by young firms; meanwhile, old firms tend to be disproportionately large, and so the reverse pattern holds.

Figure 3: Across cohort variation



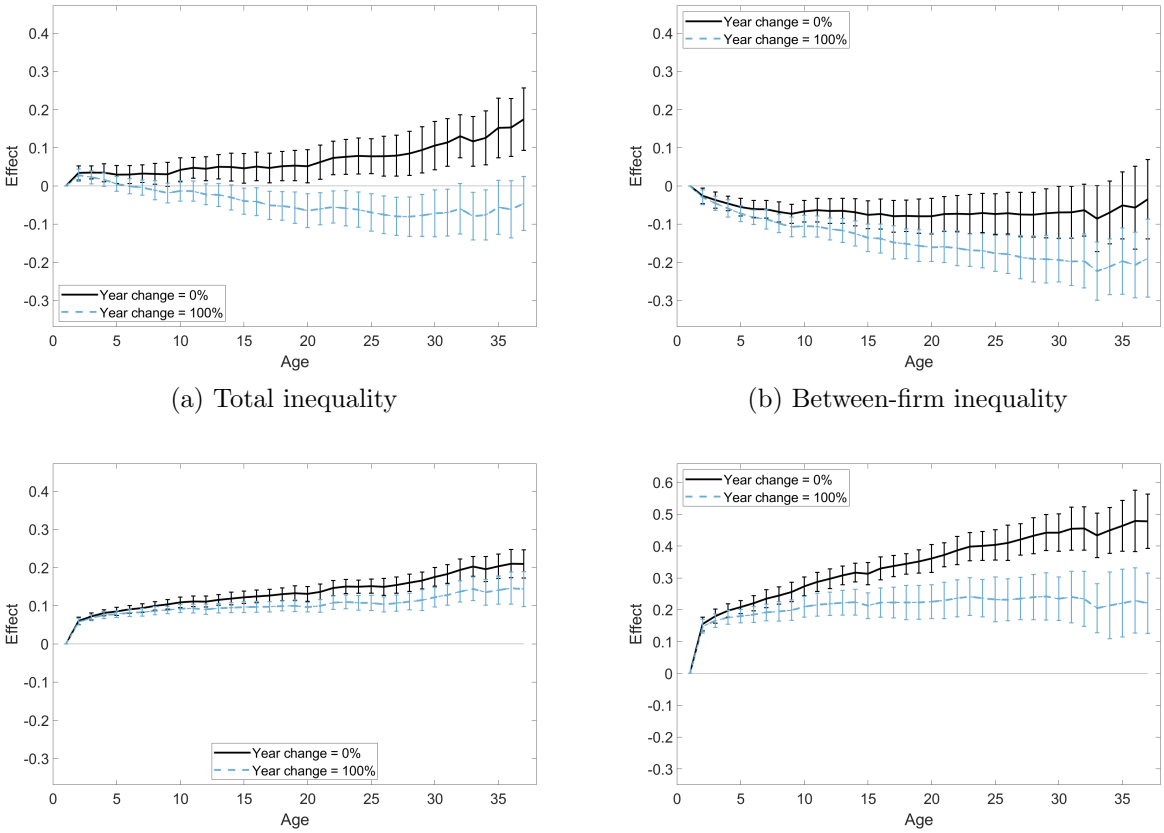
(a) Between-firm variance of earnings



(b) Variance of productivity

Notes: This figure presents the cohort (three-year moving average) trends in between-firm variance of earnings (panel a) and productivity dispersion (panel b), for workers in our 1977 cohort and beyond sample. Each line represents a cohort as it ages over time.

Figure 4: Age effects: earnings variables



(a) Total inequality

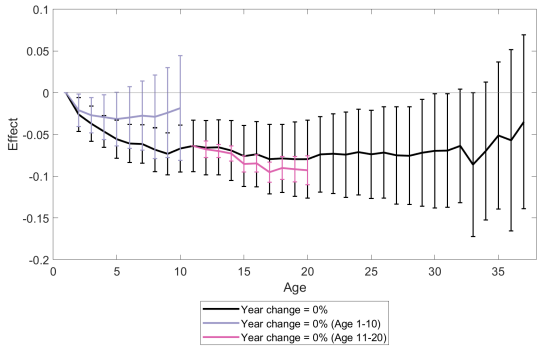
(b) Between-firm inequality

(c) Within-firm inequality

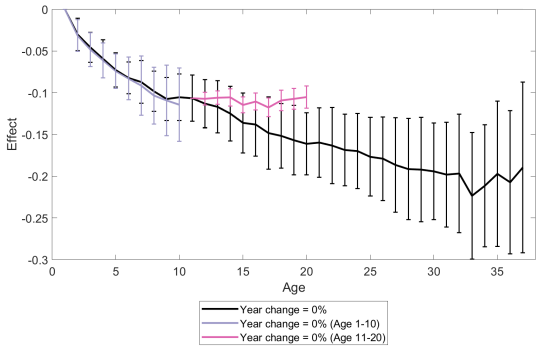
(d) Mean earnings

Notes: This figure presents estimated age effects for several outcomes (total variance of earnings in panel a, within-cohort between-firm variance of earnings in panel b, within-firm variance of earnings in panel c, and mean earnings in panel d) from the age-time-cohort decomposition in equation (3). These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for the 1977 cohort and beyond sample. Each panel shows two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%") or all of the aggregate change ("Year change = 100%").

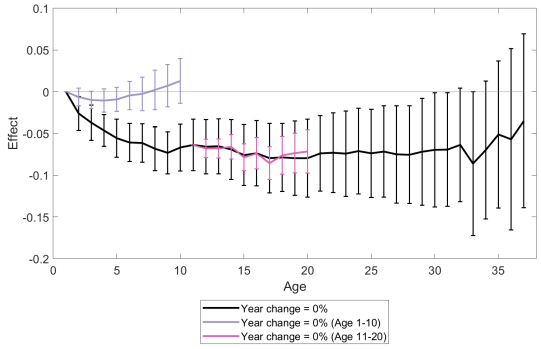
Figure 5: Age effects for between-firm earnings inequality: unbalanced vs. balanced vs. main samples



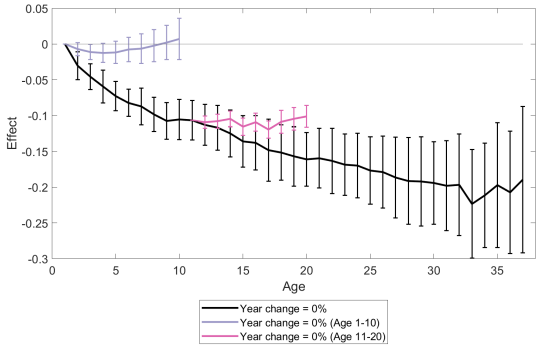
Between-firm variance: age,
Unbalanced vs. main,
(a) Year change = 0%



Between-firm variance: age,
Unbalanced vs. main,
(b) Year change = 100%



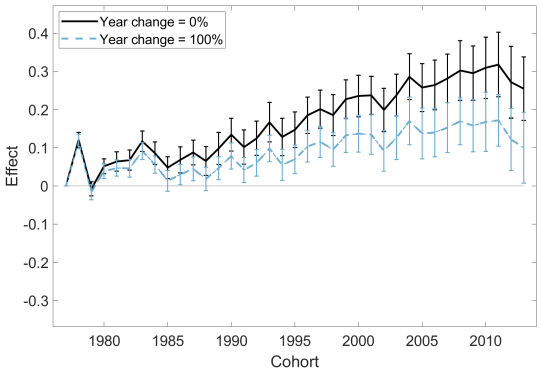
Between-firm variance: age,
Balanced vs. main,
(c) Year change = 0%



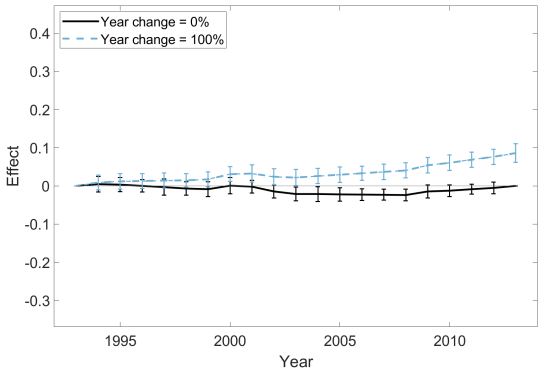
Between-firm variance: age,
Balanced vs. main,
(d) Year change = 100%

Notes: This figure presents estimated age effects for within-cohort between-firm variance of earnings from the age-time-cohort decomposition in equation (3), for five different samples. First, we consider our main sample (in black in all panels). In panels a and b, we consider the sample of firms aged 1-10 ("unbalanced," in purple) and the sample of firms aged 11-20 ("unbalanced," in pink). In panels c and d, we consider the sample of firms aged 1-10 that we observe at every age in that window ("balanced," in purple) and the sample of firms aged 11-20 that we observe at every age in that window ("balance," in pink). We normalize the estimates from the latter four samples so that they match the first sample for the first relevant age. These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for the 1977 cohort and beyond sample. The panels present two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%", panels a and c) or all of the aggregate change ("Year change = 100%", panels b and d).

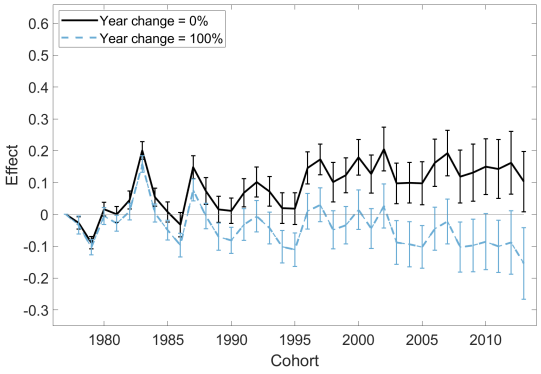
Figure 6: Cohort and year effects for between-firm earnings inequality and mean earnings



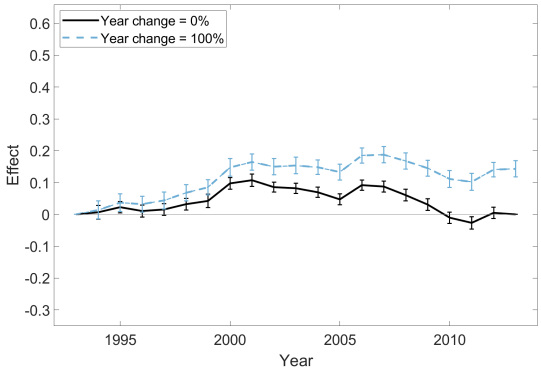
(a) Between-firm variance: cohort



(b) Between-firm variance: year



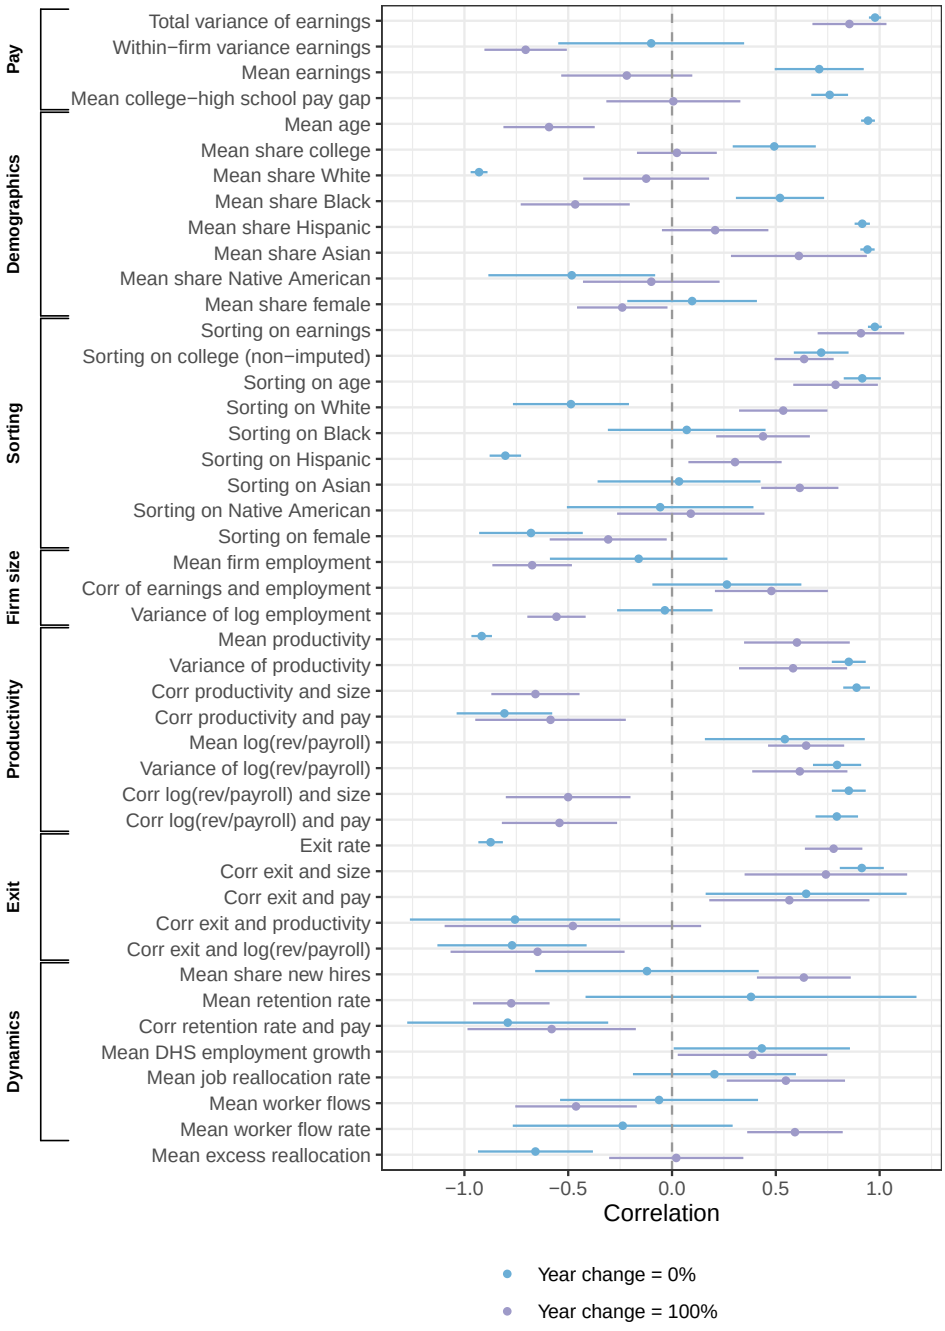
(c) Mean earnings: cohort



(d) Mean earnings: year

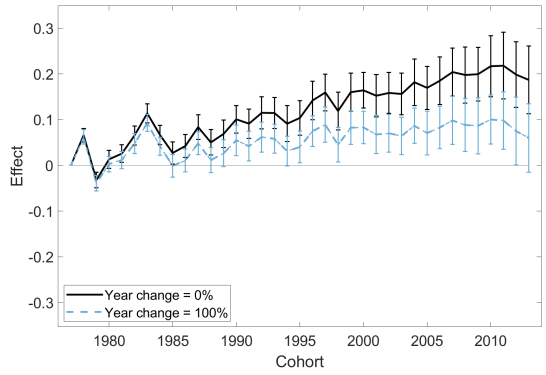
Notes: This figure presents estimated cohort and year effects for several outcomes (within-cohort between-firm variance of earnings in panels a and b and mean earnings in panels c and d) from the age-time-cohort decomposition in equation (3). These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for the 1977 cohort and beyond sample. Each panel shows two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%") or all of the aggregate change ("Year change = 100%").

Figure 7: Correlates of cohort effects in between-firm variance of earnings

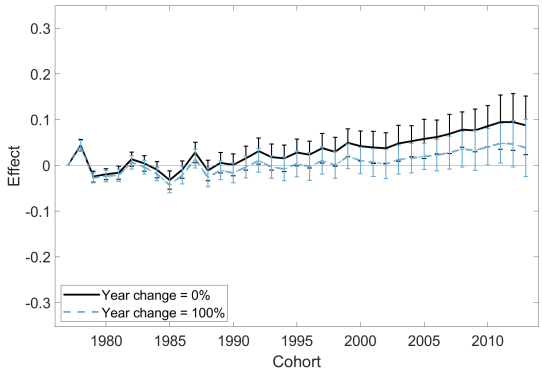


Notes: This figure summarizes the correlations of the cohort effects of the between-firm variance of earnings and the cohort effects in other variables, from estimation of the age-time-cohort decomposition in equation (3) for each variable using the procedure described in Section 3. Results are presented for the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%") or all of the aggregate change ("Year change = 100%"); 95% confidence intervals are based on bootstrapped standard errors.

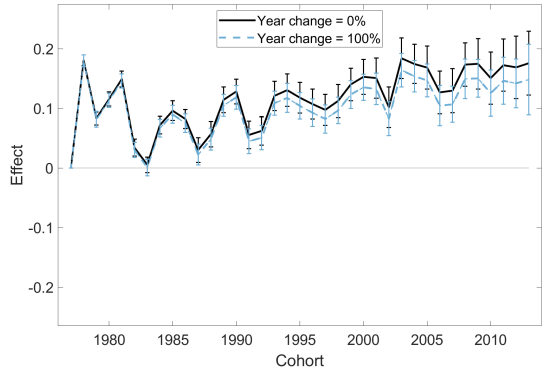
Figure 8: Cohort effects that are correlated with cohort effects in between-firm earnings variance



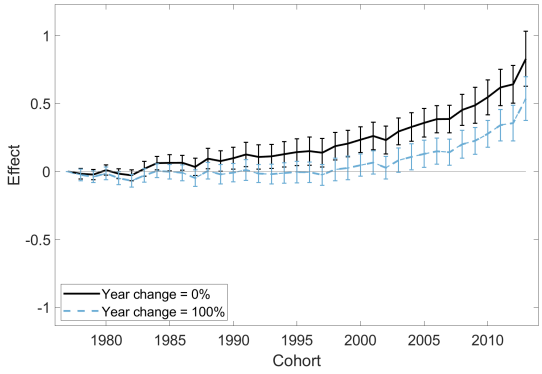
(a) Sorting on earnings



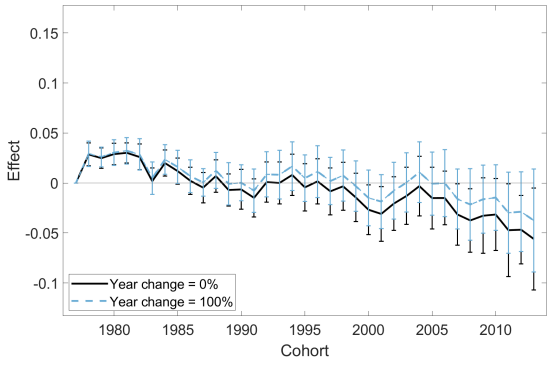
(b) Sorting on age



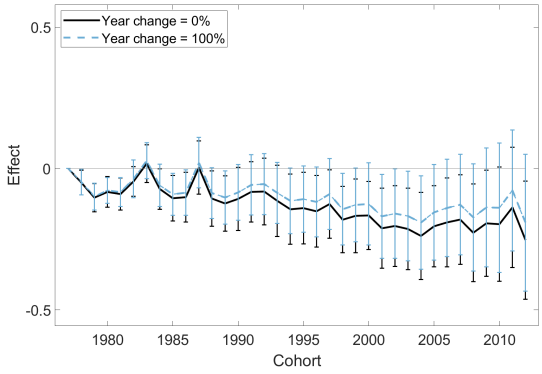
(c) Sorting on college



(d) Variance of productivity



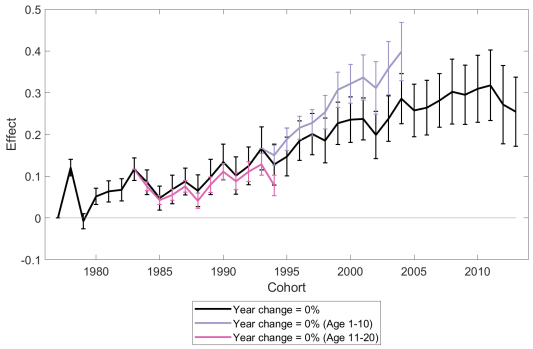
(e) Correlation of pay and productivity



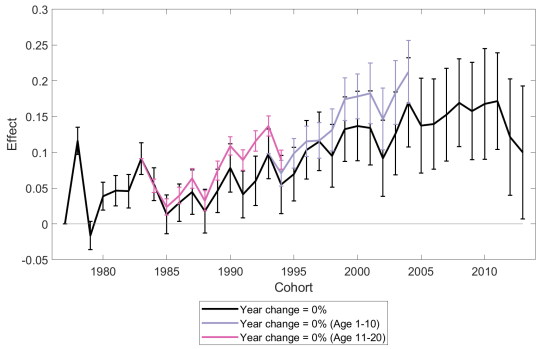
(f) Correlation of retention and pay

Notes: This figure presents estimated cohort effects for several outcomes (worker-weighted sorting on earnings, age, and college in panels a, b, and c, respectively; the firm-weighted variance of productivity and correlation of pay and productivity in panels d and e, respectively; and the worker-weighted correlation of worker retention and pay in panel f) from the age-time-cohort decomposition in equation (3). These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for the 1977 cohort and beyond sample. Each panel shows two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%") or all of the aggregate change ("Year change = 100%").

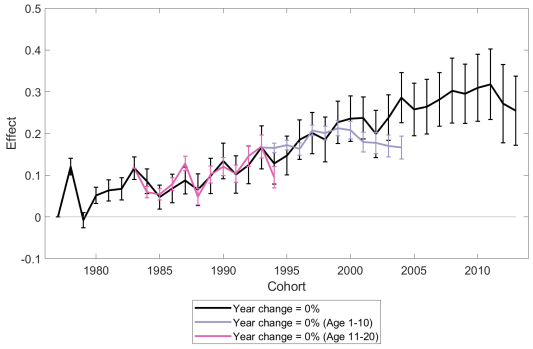
Figure 9: Cohort effects for between-firm earnings inequality: unbalanced vs. balanced vs. main samples



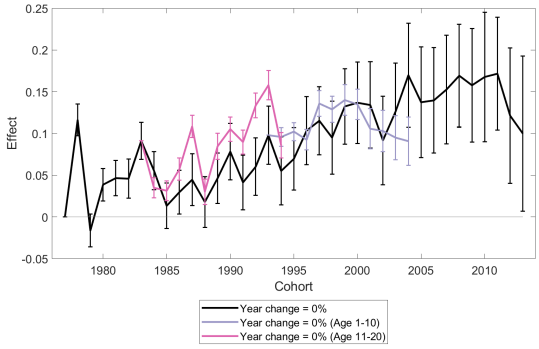
Between-firm variance: cohort, Unbalanced vs. main, (a) Year change = 0%



Between-firm variance: cohort, Unbalanced vs. main, (b) Year change = 100%



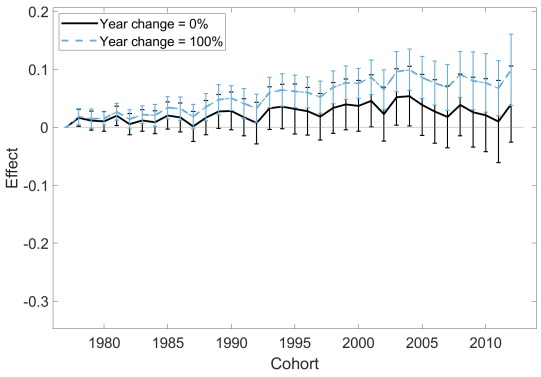
Between-firm variance: cohort, Balanced vs. main, (c) Year change = 0%



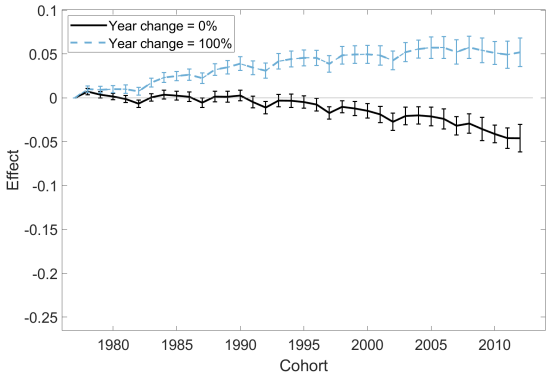
Between-firm variance: cohort, Balanced vs. main, (d) Year change = 100%

Notes: This figure presents estimated cohort effects for within-cohort between-firm variance of earnings from the age-time-cohort decomposition in equation (3), for five different samples. First, we consider our main sample (in black in all panels). In panels a and b, we consider the sample of firms aged 1-10 ("unbalanced," in purple) and the sample of firms aged 11-20 ("unbalanced," in pink). In panels c and d, we consider the sample of firms aged 1-10 that we observe at every age in that window ("balanced," in purple) and the sample of firms aged 11-20 that we observe at every age in that window ("balance," in pink). We normalize the estimates from the latter four samples so that they match the first sample for the first relevant cohort. These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for the 1977 cohort and beyond sample. The panels present two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%", panels a and c) or all of the aggregate change ("Year change = 100%", panels b and d).

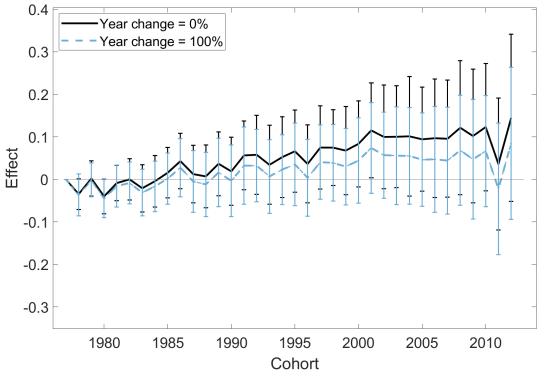
Figure 10: Cohort effects of exit and predictors of exit



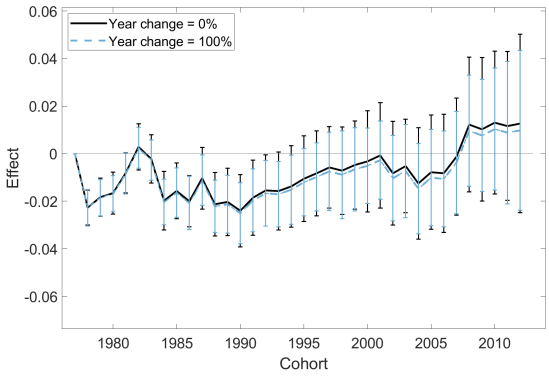
(a) Exit, worker-weighted



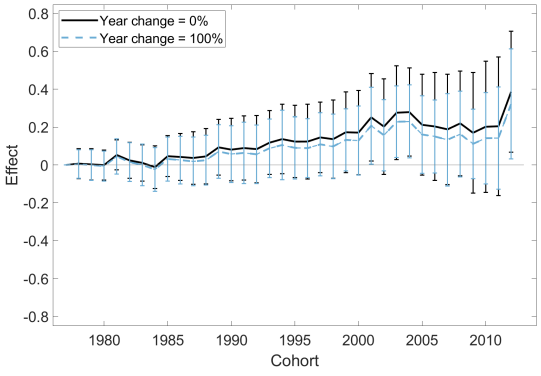
(b) Exit, firm-weighted



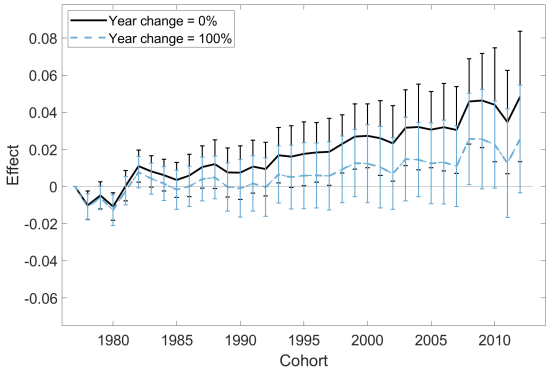
(c) Correlation of exit and pay, worker-weighted



(d) Correlation of exit and pay, firm-weighted



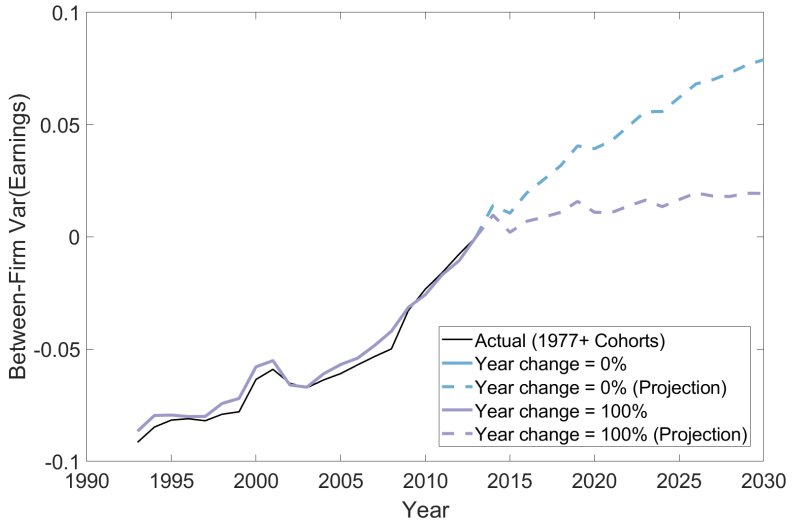
(e) Correlation of exit and size, worker-weighted



(f) Correlation of exit and size, firm-weighted

Notes: This figure presents estimated cohort effects for several outcomes (worker-weighted and firm-weighted exit, correlation of exit with pay, and correlation of exit with firm size) from the age-time-cohort decomposition in equation (3). These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for the 1977 cohort and beyond sample. Each panel shows two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%") or all of the aggregate change ("Year change = 100%").

Figure 11: Projecting forward between-firm variance of earnings



Notes: This figure shows projections of future overall between-firm variance of earnings, based on our estimated age, cohort, and year effects. From 1993 to 2013, we plot our predicted values (equation (8)) under the two year slope normalizations, compared to the actual time series (from Figure 1). From 2014 to 2030, we plot predicted values, assuming the following. First, the employment distribution across firm ages is fixed to the 2013 value. Second, the year effects for years after 2013 are equal to the 2013 year effect. Third, the cohort effects for cohorts after 2013 are equal to the 2014 cohort effect. (This means that the only thing changing across the years is that older cohorts are replaced by newer cohorts, who behave like the 2013 cohort.) We normalize each line so that the 2013 value equals zero (by demeaning by the 2013 value).

A Details of the age-time-cohort decomposition

Here we provide additional details on how levels of age, year, and cohort effects are identified from an additively separably decomposition, given several normalizations. Note that in our implementation of the overall estimation, the effects are jointly estimated in a constrained regression, described in the main text; this section lays out how the levels are identified, under the normalizations.

We adopt the additive model where we can write an outcome y (say, the between-firm variance of earnings inequality) of a cohort c at age a and time period t as

$$y_{c,a,t} = \chi_c + \alpha_a + \tau_t + \epsilon_{c,a,t}, \tag{A1}$$

We note one naming convention for this section: the cohorts are ordered in terms of their reverse entry dates. That is, c_1 is a cohort that enters one year later than when c_2 enters.

A.1 Estimating second differences

As McKenzie (2006) shows, what is identified are second differences of each of the age, period, and cohort effects.

A.1.1 Age effects

We want to estimate second derivatives of age effects, i.e.

$$\tilde{\alpha}_{a_{j+2}} = (\alpha_{a_{j+2}} - \alpha_{a_{j+1}}) - (\alpha_{a_{j+1}} - \alpha_{a_j}).$$

In order to estimate these second derivatives, we can take second differences of the outcome y. First, subtract the outcome of a given cohort in a given year from the same cohort in the next year (one year older), e.g.:

$$\begin{aligned} \Delta_t y_{c_1,a_2,t_2} &= y_{c_1,a_2,t_2} - y_{c_1,a_1,t_1} = \big(\chi_{c_1} + \alpha_{a_2} + \tau_{t_2} + \epsilon_{c_1,a_2,t_2}\big) - \big(\chi_{c_1} + \alpha_{a_1} + \tau_{t_1} + \epsilon_{c_1,a_1,t_1}\big) \\ &= \big(\alpha_{a_2} - \alpha_{a_1}\big) + \big(\tau_{t_2} - \tau_{t_1}\big) + \Delta_t \epsilon_{c_1,a_2,t_2}. \end{aligned} \tag{A2}$$

Second, subtract the first difference of a younger cohort from the first difference of a one-year-older cohort, across the same years (and thus one age apart):

$$\begin{aligned} \Delta_c \Delta_t y_{c_2,a_3,t_2} &= \Delta_t y_{c_2,a_3,t_2} - \Delta_t y_{c_1,a_2,t_2} \\ &= \Big(\big(\alpha_{a_3} - \alpha_{a_2}\big) + \big(\tau_{t_2} - \tau_{t_1}\big) + \Delta_t \epsilon_{c_2,a_3,t_2}\Big) - \Big(\big(\alpha_{a_2} - \alpha_{a_1}\big) + \big(\tau_{t_2} - \tau_{t_1}\big) + \Delta_t \epsilon_{c_1,a_2,t_2}\Big) \\ &= \big(\alpha_{a_3} - \alpha_{a_2}\big) - \big(\alpha_{a_2} - \alpha_{a_1}\big) + \Delta_c \Delta_t \epsilon_{c_2,a_3,t_2} \equiv \tilde{\alpha}_{a_3} + \Delta_c \Delta_t \epsilon_{c_2,a_3,t_2}. \end{aligned} \tag{A3}$$

Given second differences for all ages, we can estimate the second derivatives $\hat{\tilde{\alpha}}_{a_3}$; we observe multiple observations for each age, and so we can average across these observations to get the estimates. Specifically, the regression implementation of this averaging is to run the following regression (where we drop the c subscripts for simplicity):

$$\Delta_c \Delta_t y_{a_3,t_k} = \tilde{\alpha}_{a_3} + \Delta_c \Delta_t \epsilon_{a_3,t_k}, \tag{A4}$$

for $k \in [2, ...K]$ (in our example, where we have data from 1993 to 2013, we would have $t_2 = 1994$ so that $t_K = 2013$), where the coefficient of interest is just the estimate of the dummy variable. For the age effects we then have regressions of the above form for $\{a_3, ..., a_{37}\}$.

A.1.2 Year effects

We want to estimate second derivatives of year effects, i.e.

$$\tilde{\tau}_{t_{j+2}} = (\tau_{t_{j+2}} - \tau_{t_{j+1}}) - (\tau_{t_{j+1}} - \tau_{t_j}).$$

We can use the first differences from the age effects section, but instead of forming the second difference across cohorts in the same years, we consider adjacent cohorts in adjacent years (i.e. at the same age). Note that in order to cancel out the age effects and leave in the year effects, we need to subtract off an older cohort from a younger one (so, the opposite from as in the age effects). That is, e.g.:

$$\begin{aligned}\Delta_{-c,t}\Delta_t y_{c_0,a_2,t_3} &= \Delta_t y_{c_0,a_2,t_3} - \Delta_t y_{c_1,a_2,t_2} \\ &= \Big((\alpha_{a_2} - \alpha_{a_1}) + (\tau_{t_3} - \tau_{t_2}) + \Delta_t \epsilon_{c_0,a_2,t_3}\Big) - \Big((\alpha_{a_2} - \alpha_{a_1}) + (\tau_{t_2} - \tau_{t_1}) + \Delta_t \epsilon_{c_1,a_2,t_2}\Big) \\ &= (\tau_{t_3} - \tau_{t_2}) - (\tau_{t_2} - \tau_{t_1}) + \Delta_{-c,t}\Delta_t \epsilon_{c_0,a_2,t_3} \equiv \tilde{\tau}_{t_3} + \Delta_{-c,t}\Delta_t \epsilon_{c_0,a_2,t_3}.\end{aligned} \tag{A5}$$

Given second differences for all time years, we can estimate the second derivatives $\hat{\tilde{\tau}}_{t_j}$; we observe multiple observations for each year, and so we can average across these observations to get the estimates. We can also estimate the first differences too, which we will do in the final estimation of levels.

The regression implementation of this averaging is to run (where we drop the cohort subscript for simplicity):

$$\Delta_{-c,t}\Delta_t y_{a_l,t_3} = \tilde{\tau}_{t_3} + \Delta_{-c,t}\Delta_t \epsilon_{a_l,t_3} \tag{A6}$$

for $l \in [2, ..., 37]$. Rather than running separate regressions, we can stack the regressions and run:

$$\Delta_{-c,t}\Delta_t y_{a_l,t_k} = \tilde{\tau}_{t_k} + \Delta_{-c,t}\Delta_t \epsilon_{a_l,t_k} \tag{A7}$$

for $l \in [2, ..., 37]$ and $k \in [3, ..., 37]$.

A.1.3 Cohort effects

We want to estimate second derivatives of cohort effects, i.e.

$$\tilde{\chi}_{c_{j+2}} = (\chi_{c_{j+2}} - \chi_{c_{j+1}}) - (\chi_{c_{j+1}} - \chi_{c_j}).$$

Unlike the age and year effects, we have a different first difference. That is, first subtract the outcome of a younger cohort from the outcome of an older cohort, in the same year, e.g.:

$$\begin{aligned}\Delta_c y_{c_2,a_3,t_2} &= y_{c_2,a_3,t_2} - y_{c_1,a_2,t_2} = \big(\chi_{c_2} + \alpha_{a_3} + \tau_{t_2} + \epsilon_{c_2,a_3,t_2}\big) - \big(\chi_{c_1} + \alpha_{a_2} + \tau_{t_2} + \epsilon_{c_1,a_2,t_2}\big) \\ &= \big(\chi_{c_2} - \chi_{c_1}\big) + \big(\alpha_{a_3} - \alpha_{a_2}\big) + \Delta_c \epsilon_{c_2,a_3,t_2}.\end{aligned} \tag{A8}$$

Second, subtract the first difference of a younger cohort (in a later year) from the first difference of an older cohort (in an earlier year), at the same ages, e.g.:

$$\begin{aligned}\Delta_{-t}\Delta_c y_{c_3,a_3,t_1} &= \Delta_c y_{c_3,a_3,t_1} - \Delta_c y_{c_2,a_3,t_2} \\ &= \Big(\big(\chi_{c_3} - \chi_{c_2}\big) + \big(\alpha_{a_3} - \alpha_{a_2}\big) + \Delta_c \epsilon_{c_3,a_3,t_1}\Big) - \Big(\big(\chi_{c_2} - \chi_{c_1}\big) + \big(\alpha_{a_3} - \alpha_{a_2}\big) + \Delta_c \epsilon_{c_2,a_3,t_2}\Big) \\ &= \big(\chi_{c_3} - \chi_{c_2}\big) - \big(\chi_{c_2} - \chi_{c_1}\big) + \Delta_{-t}\Delta_c \epsilon_{c_3,a_3,t_1} \equiv \tilde{\chi_{c_3}} + \Delta_{-t}\Delta_c \epsilon_{c_3,a_3,t_1}.\end{aligned} \tag{A9}$$

Given second differences for all cohorts, we can estimate the second derivatives $\hat{\tilde{\chi_{c_3}}}$; we observe multiple observations for each cohort, and so we can average across these observations to get the estimates. We can also estimate the first differences too, which we will do in the final estimation of levels.

The regression version of this is to run (dropping the age subscripts for simplicity):

$$\Delta_{-t}\Delta_c y_{c_3,t_k} = \tilde{\chi}_{c_3} + \Delta_{-t}\Delta_c \epsilon_{c_3,t_k} \tag{A10}$$

for $k \in [1, ..., 21]$. We can then stack these regressions and run them all at once:

$$\Delta_{-t}\Delta_c y_{c_l,t_k} = \tilde{\chi}_{c_l} + \Delta_{-t}\Delta_c \epsilon_{c_l,t_k} \tag{A11}$$

for various $l's$ and $k's$.

A.2 Estimating slopes and levels

Given the estimated second differences, if we make some normalizations, we can estimate slopes and levels of the effects.

A.2.1 Year effects

We choose to make main normalization in the year effects. That is, we suppose that the year effects over the full year span $[1, T]$ capture some proportion of the aggregate change in the outcome. Denote the aggregate change $\Delta(Y)$, which is the difference in the weighted average of outcomes for all cohorts appearing in the first versus final time year. That is, suppose

$$\tau_{t_T} - \tau_{t_1} = x,$$

where we normalize x in two alternative ways: in the first, $x = 0$, while in the second, $x = \Delta(Y)$. Given this normalization and our estimates of the second derivatives $\hat{\tilde{\tau}}_t$, we want to estimate the slopes, and then the levels, of the year effects. We can start this process by estimating the initial year slope $\tau_{t_2} - \tau_{t_1} = s$. To solve for s as a function of x, note that:

$$
\begin{aligned}
\tilde{\tau}_{t_{j+2}} &= \big(\tau_{t_{j+2}} - \tau_{t_{j+1}}\big) - \big(\tau_{t_{j+1}} - \tau_{t_j}\big) \\
&\rightarrow \tau_{t_{j+2}} = \tilde{\tau}_{t_{j+2}} + 2\tau_{t_{j+1}} - \tau_{t_j} \\
\text{and } \tilde{\tau}_{t_{j+3}} &= \big(\tau_{t_{j+3}} - \tau_{t_{j+2}}\big) - \big(\tau_{t_{j+2}} - \tau_{t_{j+1}}\big) \\
&\rightarrow \tau_{t_{j+3}} = \tilde{\tau}_{t_{j+3}} + 2\tau_{t_{j+2}} - \tau_{t_{j+1}} = \tilde{\tau}_{t_{j+3}} + 2\tilde{\tau}_{t_{j+2}} + 2\tau_{t_{j+1}} + \tau_{t_j} \\
&\dots \\
&\rightarrow \tau_{t_T} = \sum_{p=3}^{T} \tilde{\tau}_{t_p}(T-p+1) + (T-1)\tau_{t_2} + (T-2)\tau_{t_1} \\
&\rightarrow x + \tau_{t_1} = \sum_{p=3}^{T} \tilde{\tau}_{t_p}(T-p+1) + (T-1)(s+\tau_{t_1}) + (T-2)\tau_{t_1}
\end{aligned}
\tag{A12}
$$

If we make the additional normalization of setting the initial year effect $\tau_{t_1} = 0$, then this simplifies:

$$
s = \frac{1}{T-1}\big(x - \sum_{p=3}^{T} \tilde{\tau}_{t_p}(T-p+1)\big) \tag{A13}
$$

We can thus estimate the initial year effect slope as

$$
\widehat{(\tau_{t_2} - \tau_{t_1})} = \frac{1}{T-1}\big(x - \sum_{p=3}^{T} \hat{\tilde{\tau}}_{t_p}(T-p+1)\big)
$$

Note too that once we have estimates for the first slope and the normalization of the first level, we can iterate to estimate all year effect slopes and levels, given the relationships:

- Slope in year = slope in previous year + second derivative in year

$$
(\tau_{t_{j+2}} - \tau_{t_{j+1}}) = (\tau_{t_{j+1}} - \tau_{t_j}) + \tilde{\tau}_{t_{j+2}} \tag{A14}
$$

- Level in year = level in previous year + slope in year

$$
\tau_{t_{j+2}} = \tau_{t_{j+1}} + (\tau_{t_{j+2}} - \tau_{t_{j+1}}) \tag{A15}
$$

A.2.2 Age and cohort effects

We now have estimates of all τ_{t_j}. As our final step, we want to estimate all α_{a_j} and χ_{c_j}. We can do this using the year effects estimates and the first differences above, conditional on normalizing initial levels $\alpha_{a_1} = 0$ and $\chi_{c_{37}} = 0$ (that is, age 1 and cohort 1977).

First, consider age effects. Recall the first difference in equation (A2):

$$
\Delta_t y_{c_1,a_2,t_2} = \big(\alpha_{a_2} - \alpha_{a_1}\big) + \big(\tau_{t_2} - \tau_{t_1}\big) + \Delta_t \epsilon_{c_1,a_2,t_2} \tag{A16}
$$

Estimates of this first difference are equal to the age slope plus the year slope (that matches with the relevant year), the latter of which we have already estimated. This means that we can estimate

age slopes by averaging over cohorts (where we drop the cohort subscript in the Δ_t expression for notational compactness):

$$\widehat{(\alpha_{a_2} - \alpha_{a_1})} = \frac{1}{T-1} \sum_{k=t_2}^{T} \left(\Delta_t y_{a_2,t_k} - \widehat{(\tau_k - \tau_{k-1})} \right)$$

We can do this estimation for all age slopes, and then follow the year effect procedure to estimate all age effect levels.

Second, consider cohort effects. Recall the first difference in equation (A8):

$$\Delta_c y_{c_2,a_3,t_2} = \left(\chi_{c_2} - \chi_{c_1} \right) + \left(\alpha_{a_3} - \alpha_{a_2} \right) + \Delta_c \epsilon_{c_2,a_3,t_2}. \tag{A17}$$

Estimates of this first difference are equal to the cohort slope plus the age slope, the latter of which we have already estimated. This means that we can estimate cohort slopes:

$$\widehat{(\chi_{c_2} - \chi_{c_1})} = \frac{1}{A-2} \sum_{k=a_3}^{A} \left(\Delta_c y_{c_2,k} - \widehat{(\alpha_k - \alpha_{k-1})} \right)$$

We can do this estimation for all cohort slopes, and then follow the year effect procedure to estimate all cohort effect levels.

B Monte Carlo evidence

In order to evaluate the performance of our regression estimation procedure, i.e. estimate the coverage of our standard errors, we conduct a Monte Carlo analysis. In this analysis, we take the estimated (within-cohort) between-firm variance of earnings (i.e. age plus cohort plus year effects), under the two year slope normalizations. Given these predicted values, we conduct the following nested loop procedure. In the outer loop (of 100 iterations), we draw new between-firm earnings inequality values for each cohort-age cell, by adding to each cohort-age cell's predicted value a randomly drawn (with replacement) residual from the set of all cohort-age cells (this is analogous to the procedure by which we estimate standard errors via residual bootstrapping).

Then, assuming these new values of between-firm earnings inequality are "the truth," we estimate our constrained regression on these new values. Next, within this outer loop, we procedure to the inner loop (of 50 iterations), where, for the new "true" values, we estimate confidence intervals via residual bootstrapping (by drawing new values from predicted outcomes and randomly drawn residuals, as in the main regression estimation). That is, after this inner loop, we can construct 95% confidence intervals for each age, cohort, and year effects by looking at the middle 95% of the bootstrapped estimates; within the outer loop, we then check whether the estimated age, cohort, and year effects based on the "true" between-firm earnings inequality lie within the 95% confidence intervals. We conduct this inner procedure for all iterations of the outer loop. In the end, we measure what share of the estimated effects based on the "true" between-firm earnings inequality measures lie within their corresponding 95% confidence interval. The closer this share is to 95%, which would be the expected value under a well-performing estimation, the more confidence we place in our main regression estimation procedure.

Table A1 present the results of the Monte Carlo analysis. As the table shows, the shares of estimates based on the "true" (outer loop) values that lie within their corresponding confidence intervals are close to 0.95, across the age, cohort, and year effects and normalizations. The coverage is the worst for the cohort effects under the normalization in which year effects explain none of the aggregate change; here, the estimate cohort effects lie in their related 95% confidence intervals 92.89% of the time. We expect that our coverage would improve if we increased in the number of bootstrap loops in our main regressions, where we have limited the loops due to computational restrictions. Overall, with coverage close to 95%, we are confident in our estimation procedure.

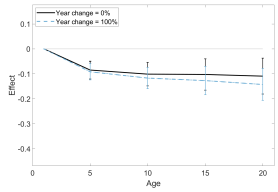
Table A1: Monte Carlo coverage

	Age effects		Cohort effects		Year effects	
Year effects explain ___ of aggregate change:	none	all	none	all	none	all
Share of Monte Carlo confidence intervals that contain "true" estimates	0.9586	0.9608	0.9289	0.9349	0.9514	0.941

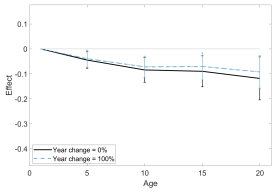
Notes: This table presents the Monte Carlo evidence of the performance of our estimation of age, cohort, and year effects for (within-cohort) between-firm earnings inequality, under the two year slope normalization. The values listed are the share of the Monte Carlo 95% confidence intervals that contain the "true" estimates obtained by adding noise to our main regression estimates. The fact that the values are all close to 0.95, which would be the expected value under a well-performing estimation, we are confident in our estimation procedure.

C Additional Figures

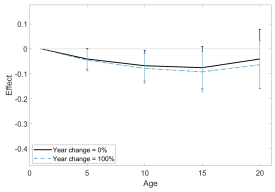
Figure A1: Age effects: Between-firm inequality



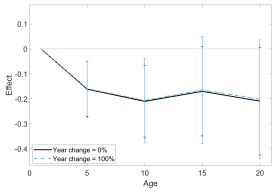
(a) Average across sectors



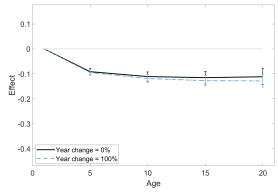
(b) Agriculture



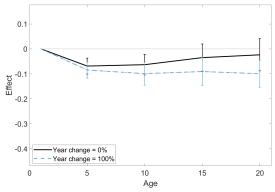
(c) Mining



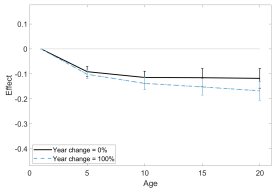
(d) Utilities



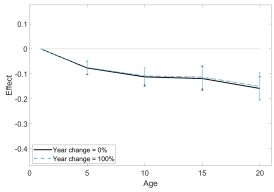
(e) Construction



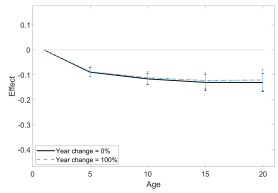
(f) Manufacturing



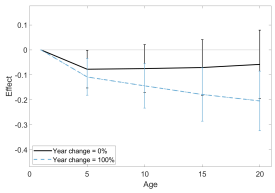
(g) Wholesale Trade



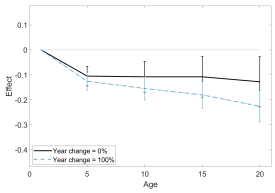
(h) Retail Trade



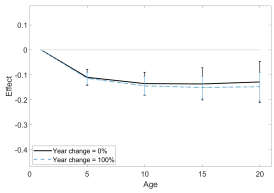
(i) Transport/Warehousing



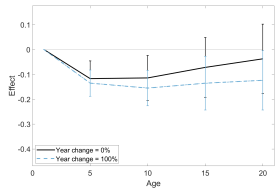
(j) Information



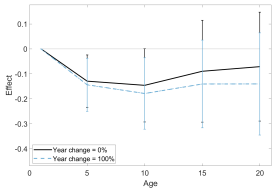
(k) Finance and Insurance



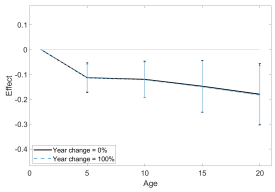
(l) Real Estate



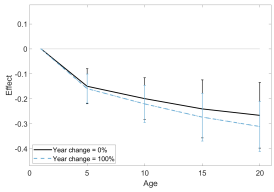
(m) Professional, Science, and Technical Services



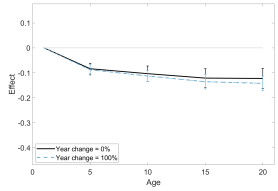
(n) Management



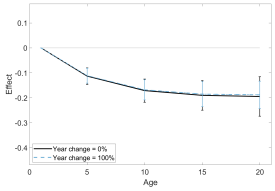
(o) Admin, Support, and Waste Management



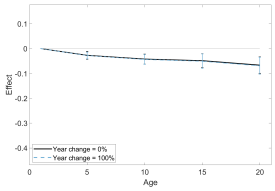
(p) Education



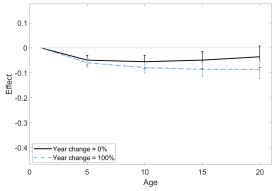
(q) Health



(r) Arts and Entertainment



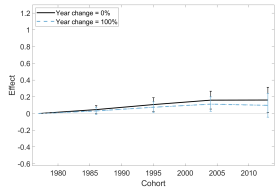
(s) Accommodation and Food Services



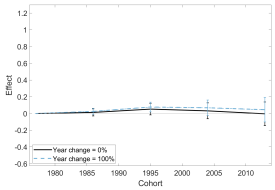
(t) Other Services

Notes: This figure presents select estimated age effects for within-cohort between-firm variance of earnings by sector from the age-time-cohort decomposition in equation (3). These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for the 1977 cohort and beyond sample. Each panel shows two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%") or all of the aggregate change ("Year change = 100%"). Panel a shows the employment-weighted average of panels c-t, using 2013 national employment shares from the Business Dynamics Statistics dataset (which does not contain information on the Agriculture sector).

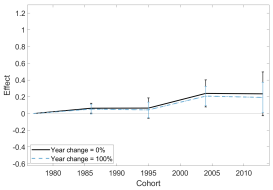
Figure A2: Cohort effects by sector: Between-firm inequality



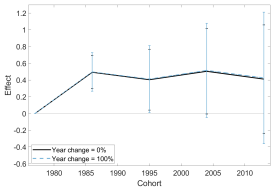
(a) Average across sectors



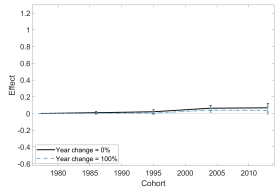
(b) Agriculture



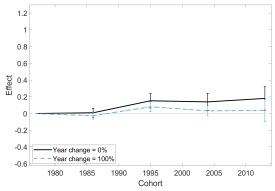
(c) Mining



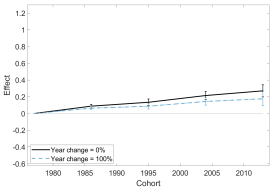
(d) Utilities



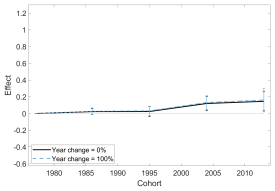
(e) Construction



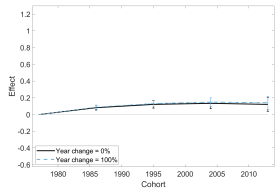
(f) Manufacturing



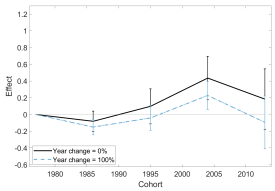
(g) Wholesale Trade



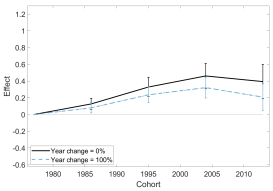
(h) Retail Trade



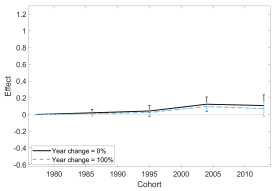
(i) Transport/Warehousing



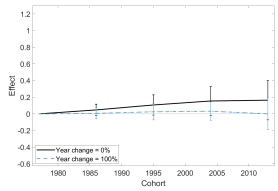
(j) Information



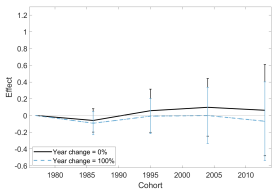
(k) Finance and Insurance



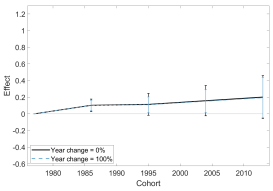
(l) Real Estate



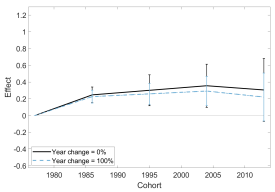
(m) Professional, Science, and Technical Services



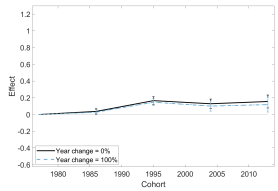
(n) Management



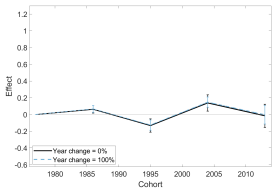
(o) Admin, Support, and Waste Management



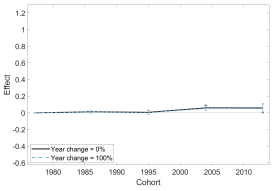
(p) Education



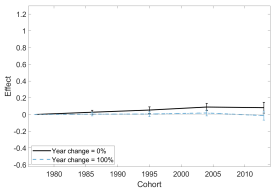
(q) Health



(r) Arts and Entertainment



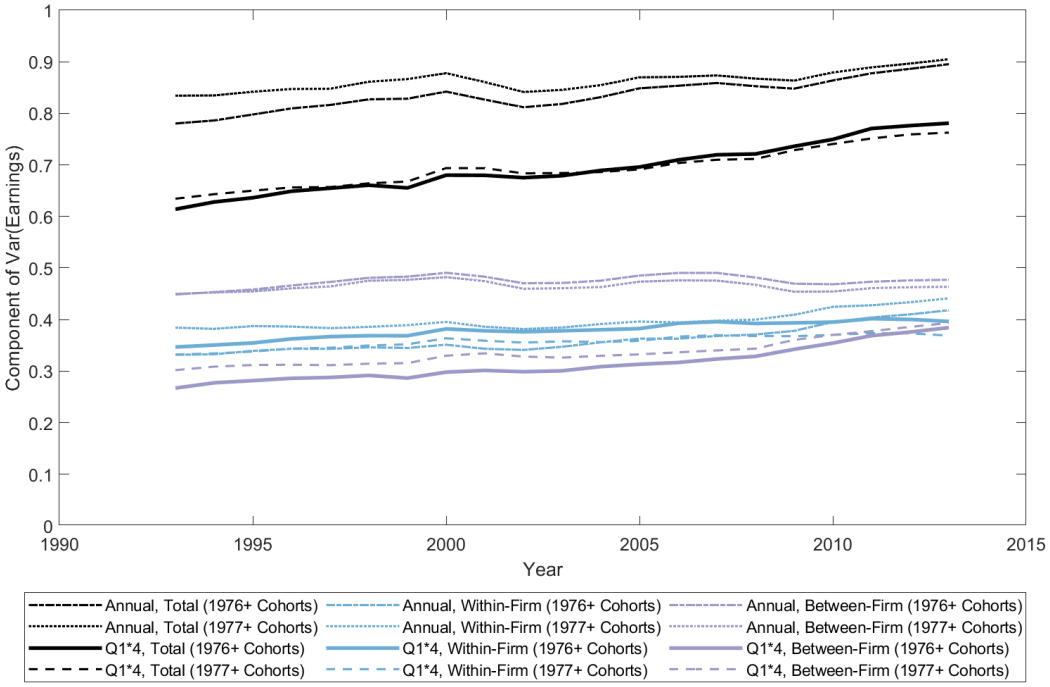
(s) Accommodation and Food Services



(t) Other Services

Notes: This figure presents select estimated cohort effects for within-cohort between-firm variance of earnings by sector from the age-time-cohort decomposition in equation (3). These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for the 1977 cohort and beyond sample. Each panel shows two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%") or all of the aggregate change ("Year change = 100%"). Panel a shows the employment-weighted average of panels c-t, using 2013 national employment shares from the Business Dynamics Statistics dataset (which does not contain information on the Agriculture sector).

Figure A3: Aggregate Trends in Earnings Inequality: Robustness to measurement of earnings

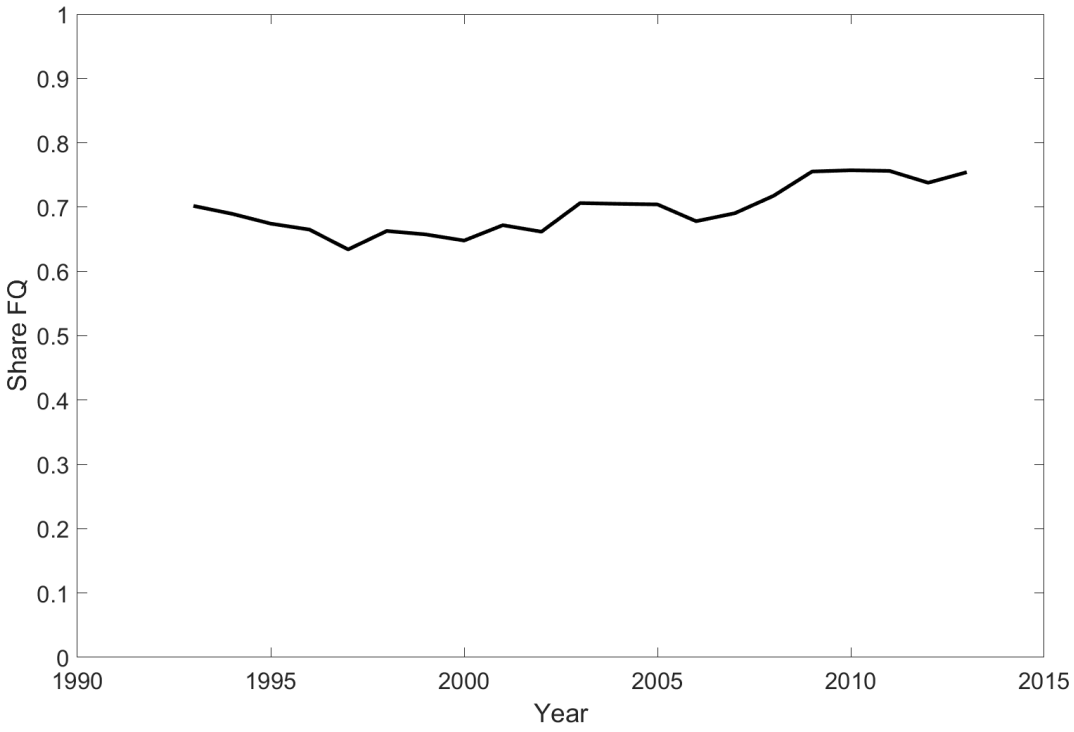


Notes: This figure presents robustness to our measurement of earnings by demonstrating the aggregate trends in earnings inequality for two different measures of earnings: our standard measure (*annualized* earnings, i.e. Q1 earnings multiplied by 4) and (log) true *annual* earnings (i.e. the observed annual earnings in the LEHD). We do this for two samples: workers at all firms in our full sample ("1976+ Cohorts") and workers at firms in the 1977 cohort and beyond sample ("1977+ Cohorts"). We decompose the total variance of earnings into the dispersion in average pay at firms (between-firm) and the dispersion of pay Within-firms (within-firm) (see equation (2)).

Two patterns emerge when comparing the trends to inequality for annualized ("Q1*4") and annual earnings. First, the trends are quite similar. For instance, 70% of the increase in total *annualized* earnings inequality (for our full sample) is accounted for by the rise in between-firm earnings inequality. For *annual* earnings, the comparable share is 75%.

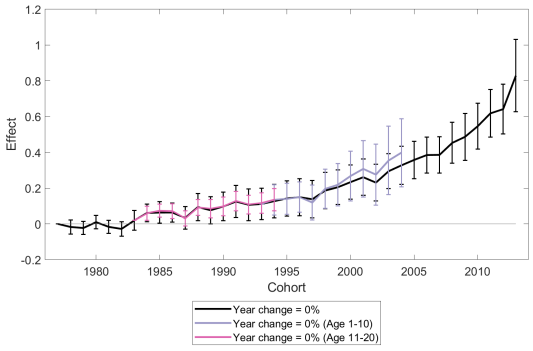
Second, the level of earnings inequality is higher when we use annual earnings, and this is particularly true for between-firm earnings inequality. This is likely driven by two factors. First, and arguably most importantly, inequality in *annual* earnings includes some variation from workers changing firms; recall that our sample restricts to full-quarter jobs in Q1 (such that these workers are still employed at the same jobs in Q2) but these jobs could end within Q2 or in Q3 or Q4, generating more variation. This feature of annual earnings is what inspires us to use annualized earnings, since we do not want to conflate job moves with pay variation. Second, seasonal variation in pay, for instance through annual bonuses, may generate more variation when we use annual earnings.

Figure A4: Share of full-quarter jobs by year



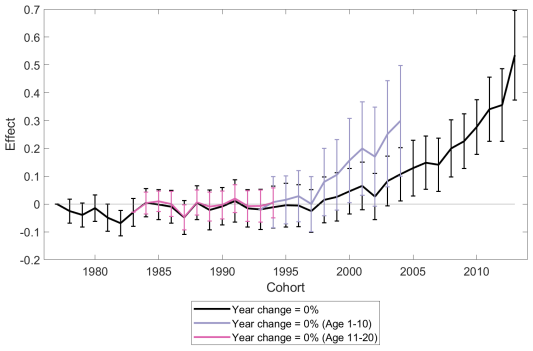
Notes: This figure presents the share of jobs that are full-quarter (and earn above our earnings threshold, i.e. the Q1 earnings is at least $^{3250}/_{4}$) and are thus eligible for our full sample, over time. From 1993 to 2013, this share increases by 7%, a modest increase.

Figure A5: Cohort effects for variance of productivity: unbalanced vs. balanced vs. main samples



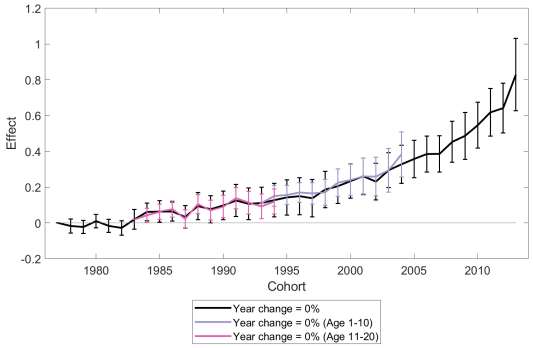
Variance of productivity: cohort, Unbalanced vs. main, Year change = 0%

(a)



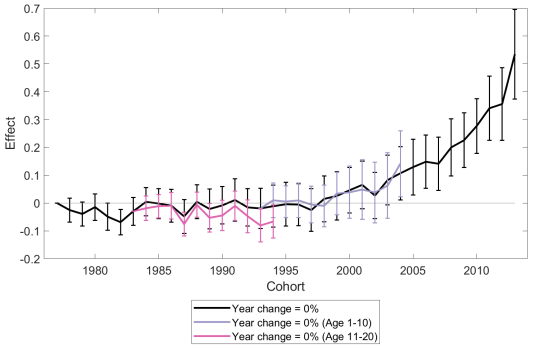
Variance of productivity: cohort, Unbalanced vs. main, Year change = 100%

(b)



Variance of productivity: cohort, Balanced vs. main, Year change = 0%

(c)

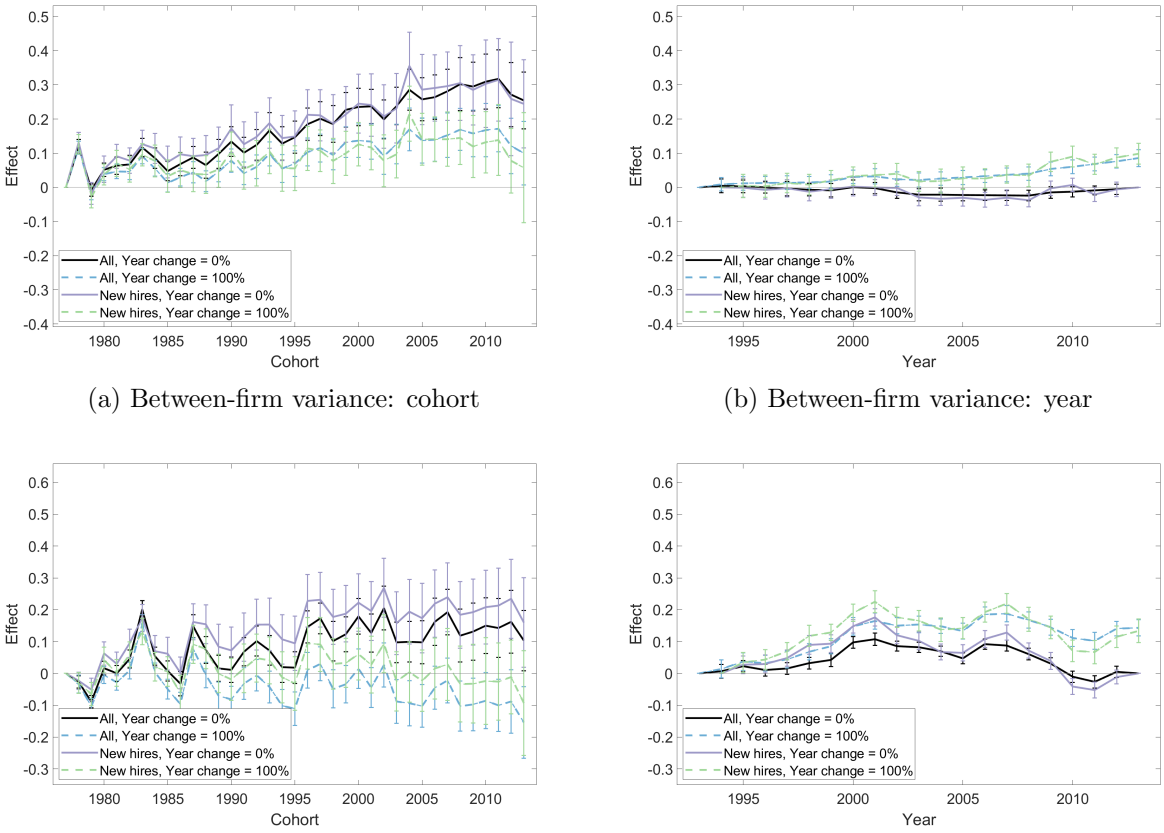


Variance of productivity: cohort, Balanced vs. main, Year change = 100%

(d)

Notes: This figure presents estimated cohort effects for within-cohort variance of productivity from the age-time-cohort decomposition in equation (3), for five different samples. First, we consider our main sample (in black in all panels). In panels a and b, we consider the sample of firms aged 1-10 ("unbalanced," in purple) and the sample of firms aged 11-20 ("unbalanced," in pink). In panels c and d, we consider the sample of firms aged 1-10 that we observe at every age in that window ("balanced," in purple) and the sample of firms aged 11-20 that we observe at every age in that window ("balance," in pink). We normalize the estimates from the latter four samples so that they match the first sample for the first relevant cohort. These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for the 1977 cohort and beyond sample. The panels present two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%", panels a and c) or all of the aggregate change ("Year change = 100%", panels b and d).

Figure A6: New hires vs. all workers



(a) Between-firm variance: cohort

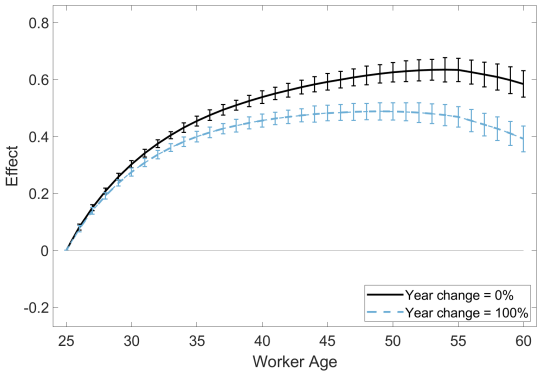
(b) Between-firm variance: year

(c) Mean earnings: cohort

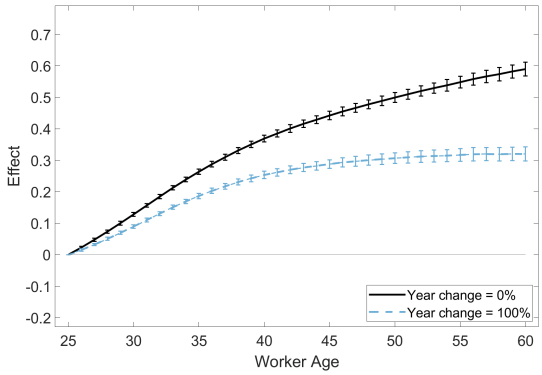
(d) Mean earnings: year

Notes: This figure compares the results in Figure 6 to cohort and year effects estimated on the subsample of workers who are new hires in a given year (i.e. were not employed at their firm in the prior year). That is, the figure presents estimated cohort and year effects for several outcomes (within-cohort between-firm variance of earnings in panels a and b and mean earnings in panels c and d) from the age-time-cohort decomposition in equation (3). These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for the 1977 cohort and beyond sample, as well as the subsample of workers who are new hires. Each panel shows two sets of estimates from the two year slope normalizations (for each of the two samples), where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%") or all of the aggregate change ("Year change = 100%"). As the figure shows, results are similar across all workers and new hires.

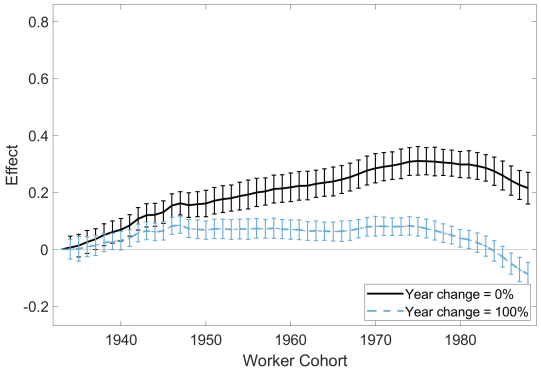
Figure A7: Worker age, worker cohort, and year effects



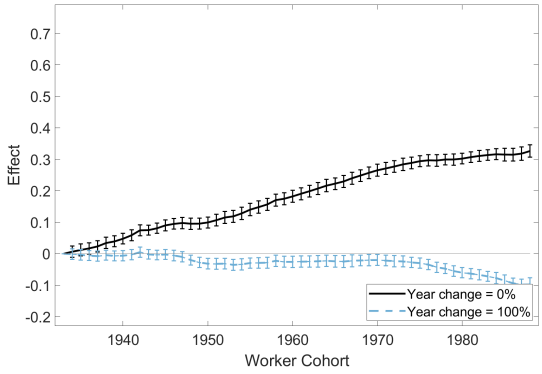
(a) Mean earnings, age effects



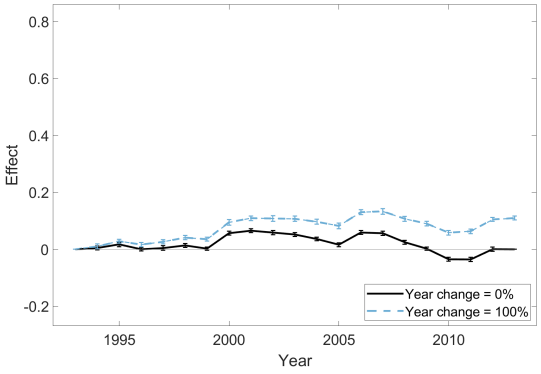
(b) Variance of earnings, age effects



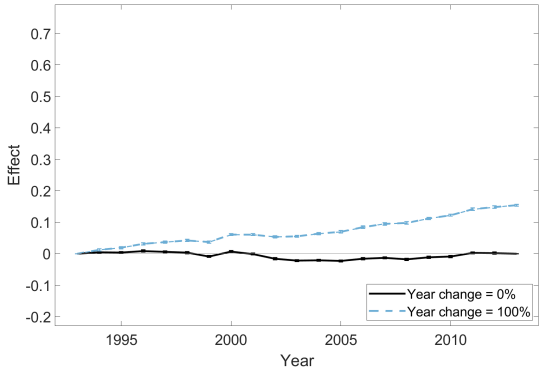
(c) Mean earnings, cohort effects



(d) Variance of earnings, cohort effects



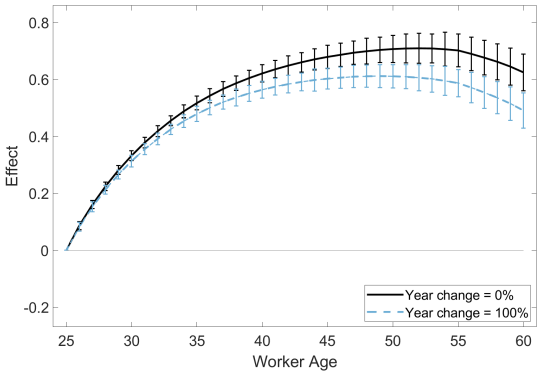
(e) Mean earnings, year effects



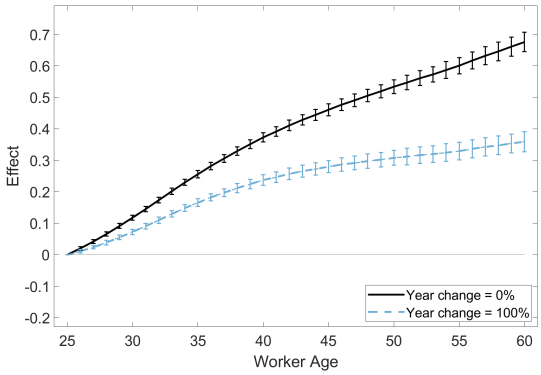
(f) Variance of earnings, year effects

Notes: This figure presents estimated worker age, worker cohort, and year effects for the mean and variance of earnings from the age-time-cohort decomposition in equation (3). These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for workers in the 1976 (firm) cohort and beyond sample. Each panel shows two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%") or all of the aggregate change ("Year change = 100%").

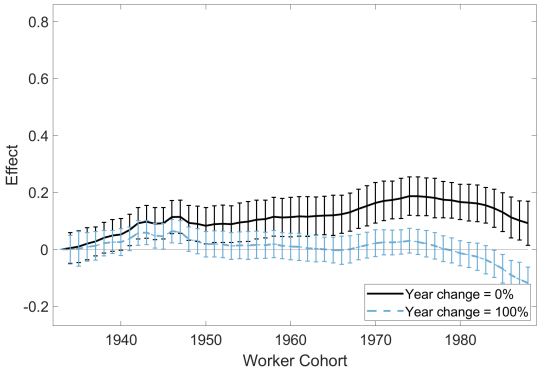
Figure A8: Worker age, worker cohort, and year effects: Men



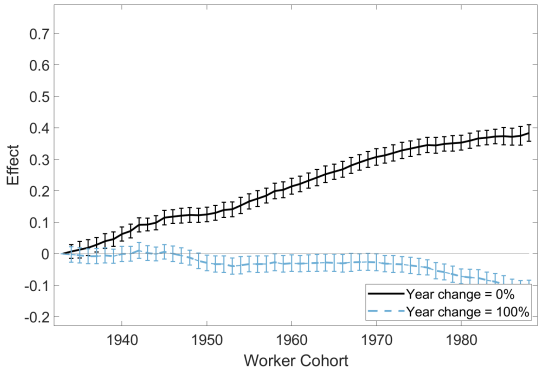
(a) Mean earnings, age effects



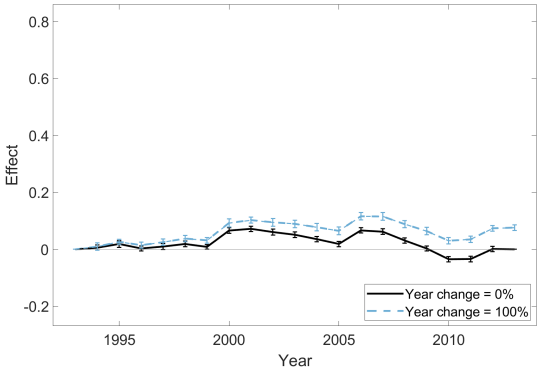
(b) Variance of earnings, age effects



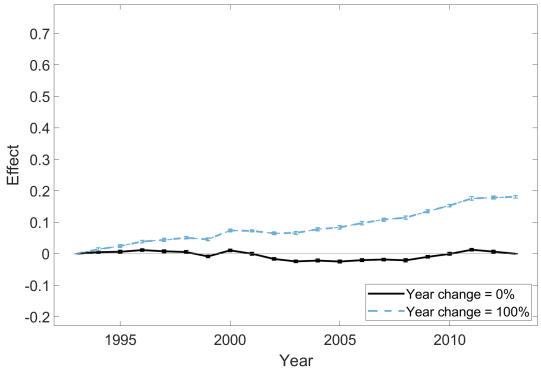
(c) Mean earnings, cohort effects



(d) Variance of earnings, cohort effects



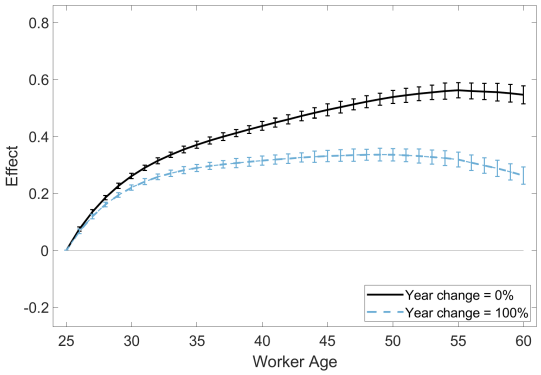
(e) Mean earnings, year effects



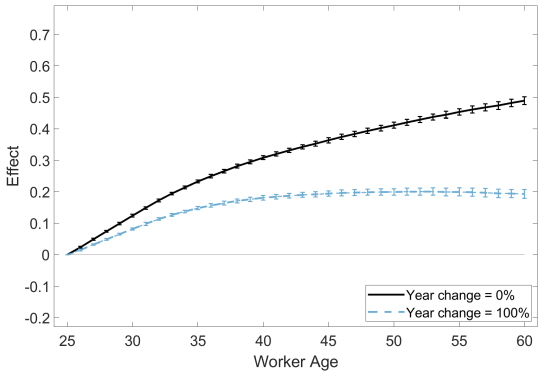
(f) Variance of earnings, year effects

Notes: This figure presents estimated worker age, worker cohort, and year effects for the mean and variance of earnings from the age-time-cohort decomposition in equation (3). These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for male workers in the 1976 (firm) cohort and beyond sample. Each panel shows two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%") or all of the aggregate change ("Year change = 100%").

Figure A9: Worker age, worker cohort, and year effects: Women



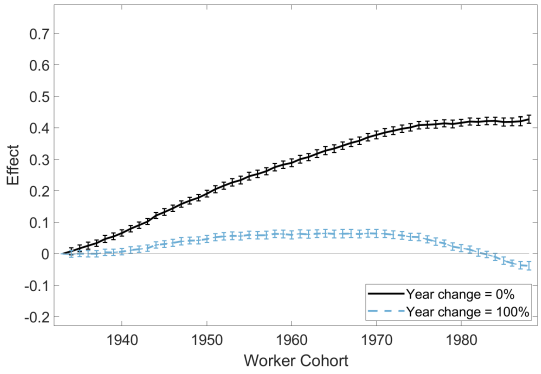
(a) Mean earnings, age effects



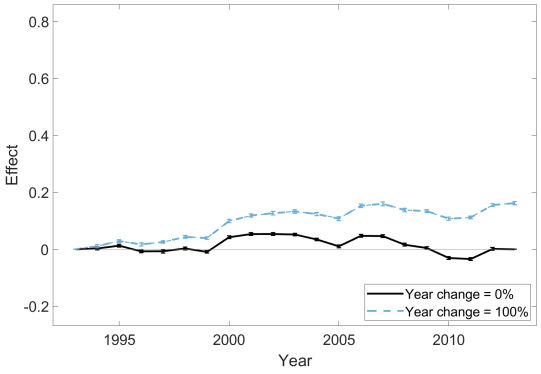
(b) Variance of earnings, age effects



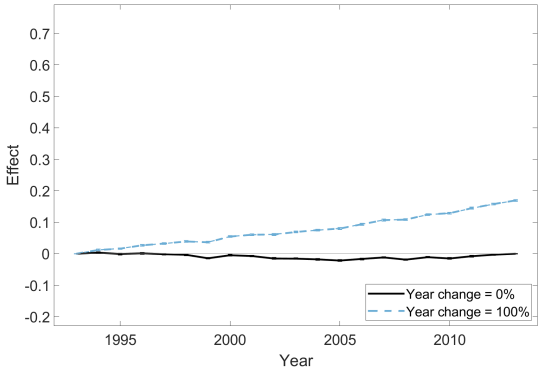
(c) Mean earnings, cohort effects



(d) Variance of earnings, cohort effects



(e) Mean earnings, year effects



(f) Variance of earnings, year effects

Notes: This figure presents estimated worker age, worker cohort, and year effects for the mean and variance of earnings from the age-time-cohort decomposition in equation (3). These effects, and their 95% confidence intervals (1.96 times standard error), are estimated using the procedure described in Section 3 for female workers in the 1976 (firm) cohort and beyond sample. Each panel shows two sets of estimates from the two year slope normalizations, where we either constrain the year slope to capture none of the aggregate change in the outcome ("Year change = 0%") or all of the aggregate change ("Year change = 100%").