Labor Force Demographics and Corporate Innovation

FRANÇOIS DERRIEN, AMBRUS KECSKÉS, and PHUONG-ANH NGUYEN*

Abstract

Firms in younger labor markets produce more innovation. We establish this by instrumenting the current labor force with historical births in each local labor market in the United States. Analyses of firms and inventors allow us to rule out unobservable heterogeneity across local labor markets and firms, life cycles, and other effects. Corporate innovation in younger labor markets reflects the innovative characteristics of younger labor forces and has greater market value. Younger workers as a group – inventors interacting with non-inventors – produce more innovation for firms through the labor force channel rather than through a financing or consumption channel.

August 22, 2021

JEL classification: G31, J11, J21, O31
Keywords: Innovation; Demographics; Labor markets; Inventors; Firm value

---

* Derrien is at HEC Paris; Kecskés is at the Schulich School of Business, York University; and Nguyen is at the School of Administrative Studies, York University. We greatly appreciate the comments of Pat Akey, Jean-Noël Barrot, Gennaro Bernile, Paul Calluzzo, Murat Çelik, Pierre Chaigneau, Peter Cziraki, Olivier Dessaint, Evan Dudley, Peter Egger, Makram El-Shagi, Laurent Frésard, Louis Gagnon, Samir Ghannam, Jim Goldman, Itay Goldstein, Johan Hombert, Wenyao Hu, Dalida Kadyrzhanova, Jun-Koo Kang, Kristoph Kleiner, Francis Kramarz, Olga Kuzmina, Tomislav Ladika, Jennifer Li, Si Li, Evgeny Lyandres, Song Ma, Roni Michaely, Fabio Moneta, Cal Muckley, Justin Hung Nguyen, Miguel Palacios, Ramesh Rao, Tianyue Ruan, Giorgo Sertsios, Christophe Spaenjers, Johan Sulaeman, Lingna Sun, David Thesmar, Şelale Tüzel, Patricio Valenzuela, Emiko Usui, Philip Valta, John Wald, Wei Wang, Zhao Wang, Dongling Xu, Bernie Yeung, Alminas Žaldokas, Balazs Zelity, Feida Zhang, and Shan Zhao, and seminar participants at the 2018 Asian Development Bank Conference, City University of Hong Kong, the 2018 European Finance Association Conference, the 2019 European Meeting of the Urban Economics Association, the 2018 Financial Intermediation Research Society Conference, Grenoble EM, the 2018 HEC Paris Workshop on Banking, Finance, Macroeconomics, and the Real Economy, the Hong Kong Polytechnic University, McMaster University, the 2019 Midwest Finance Association Conference, the 2018 Northern Finance Association Conference, Queen's University, the 2019 Santiago Finance Workshop, the 2018 SFS Cavalcade Asia-Pacific Conference, Universidad Carlos III de Madrid, the University of Calgary, the University of Fribourg, the 2018 University of Manchester Corporate Finance Conference, the University of Hong Kong, the University of Lausanne, the University of Piraeus, the University of Surrey, the University of Texas at San Antonio, the University of Toronto, and the 2019 Western Finance Association Conference. Derrien acknowledges financial support from the Investissements d'Avenir Labex (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047). This research was supported by the Social Sciences and Humanities Research Council of Canada and the Canadian Securities Institute Research Foundation.

1. Introduction

Labor markets are integral to the success of firms. Whereas a growing literature studies labor markets with a focus on their institutional and especially legal features,[1] we consider a more fundamental but highly consequential aspect of labor markets: their demographics. One important channel through which the age structure of the local labor force, in particular, matters is that young and old workers are heterogeneous inputs into the firm's production function, and they have plausibly different effects on the firm's innovation output. Specifically, as others have argued and shown, younger people are known to be more risk seeking, to have longer horizons, and to be more creative compared to older people.[2] These characteristics tend to make younger workers, and younger labor forces, more innovative.

There are a number of different ways in which individual worker characteristics aggregated to the firm level can affect corporate innovation activities. First, firms in younger labor markets can hire from a larger pool of productive inventors. Similarly, firms can hire younger workers who are not inventors themselves but who complement the firm's inventors in producing innovation (e.g., technicians, developers, managers, etc.). Furthermore, local knowledge spillovers arising from the interactions of younger workers across firm boundaries can also increase innovation in the local labor market as a whole.[3] Overall, younger labor markets can create a general work environment, for inventors and non-inventors alike, both within and across firms, that is more conducive to individual and corporate innovation. We

---

[1] For instance, unionization influences capital structure decisions (Matsa (2010)), and employment protections discourage takeovers (Dessaint, Golubov, and Volpin (2017)).
[2] See Liang, Wang, and Lazear (2018) for references.
[3] This is analogous to the effect of local agglomeration on, for example, corporate investment, as in Dougal, Parsons, and Titman (2015). For evidence on knowledge spillovers more generally, see Glaeser, Kallal, Scheinkman, and Shleifer (1992), Jaffe, Trajtenberg, and Henderson (1993), Audretsch and Feldman (1996), Bloom, Schankerman, and Van Reenen (2013), and Lychagin, Pinkse, Slade, and Van Reenen (2016).

hypothesize that it is through this labor force channel that age structure affects corporate innovation.

Testing our hypothesis is challenging because the direction of causality is difficult to establish. Age structure can affect innovation activity, but innovation can also attract migration to and from a given location, and migration can affect age structure. Moreover, innovation may be driven by factors that it has in common with age structure but which are unobservable.

What we are interested in is the causal effect of local age structure on local innovation through the labor force channel. Our novel approach to study this is to use the age structure based on historical births to instrument for the age structure of the current labor force. For every year from 1990 to 2005 and for every commuting zone in the United States, we construct the "births-based labor force" based on the number of births each year during the prior 20-64 years, adjusted for survival. That is, for each cohort of age X, we take the number of people born X years ago in the commuting zone. We use two measures of the age structure of this births-based labor force: its young share (the proportion of the 20-64 year olds that is 20-39 years old) and its mean age.

The resulting births-based local labor force is mechanically related to the actual local labor force because, as an identity, the latter equals the former plus local immigration minus local emigration. The births-based labor force is also plausibly exogenous to innovation today. It can be viewed as a historical endowment of a given location as long as it is not completely eliminated subsequently by migration across labor markets, a condition that is well supported by the data. In our empirical analysis, we use this local endowment to explain local innovation through the current local labor force.

We begin by examining innovation by publicly traded firms. With detailed firm-level data, we can identify quite precisely the channels through which innovation is affected by labor force age structure at the headquarters location of firms. The dependent variable in our baseline regressions is the quantity and quality of patents (counts and citations) and the main explanatory variable is the instrumented age structure. We include fixed effects for state-years, industry-years, and firm age groups in our main regression specifications. The industry-year fixed effects rule out the possibility that younger labor markets have a composition of industries with high innovation since the only remaining variation is within a given industry in a given year. With firm age fixed effects, the results must be interpreted as for firms of the same age, which ensures that they cannot be explained by firm life cycle effects. Since we use births-based age structure to instrument for actual age structure, reverse causality (births anticipating innovation 20-64 years later) is improbable. However, since historical births may be correlated with innovation in other ways than through the current age structure, our specifications also include control variables that account for potentially omitted factors that are correlated with both local age structure and local innovation. In our baseline specification, we include control variables at the levels of both the firm and the commuting zone, in the case of the latter, the commuting zone's size, wealth, government expenditures, educational attainment, university students, and university patents.

We find that younger labor forces produce more innovation, as indicated by higher patents counts and citations. For example, at the firm level, a one standard deviation decrease in age structure (higher young share or lower mean age) causes a roughly 8% increase in innovation. The results are similar when we identify entirely off the time-series variation in age structure by including fixed effects for commuting zones or firms.[4] We also obtain similar results when we add multiple lags of several historical control variables that are contemporaneous to historical births and even when we allow these variables to have different effects in each of our four quinquennial periods. Additionally, we perform numerous robustness tests, and among other findings, the balance between the young and the old in the labor force has no impact on the effect of age structure on innovation.

We then move to a higher level of granularity and focus on inventors who work for public and private firms. Doing so allows us to examine the general work environment as distinct from inventors. At this level of analysis, we can control for inventor age and experience as well as scale effects for both inventors and inventor networks. This allows us to distinguish between inventors and other workers in the local labor force. Our setting here also allows us to rule out time-varying firm-specific omitted factors (using firm-year fixed effects) and to identify off the variation across inventors within a given firm in a given year. The results show that inventors in younger labor markets produce more innovation. This suggests that the interactions of inventors, whether inside or outside of their firms, with fellow inventors or other workers in the local labor force, generate knowledge spillovers that affect the quantity and quality of these inventors' innovation. The general work environment appears to be important for the production of innovation, whether the inventor himself is young or old.

We then return to firms. Since we argue that firms in younger labor markets are more innovative as a result of the characteristics of younger people, we also examine whether corporate innovation activities reflect these characteristics. We use patent citations, both those

[4] Since age structure is a slow moving variable, it is generally not practical to identify entirely off its time-series variation. However, if we use the labor force of prime working age (20-54 year olds), we can increase our sample period by 50% and include fixed effects for commuting zones or firms. We implement this approach in a variation on our baseline analysis.

made by and made to the patents of our sample firms, to capture the characteristics of the firm's innovation activities, namely, their creativity, riskiness, longevity, and interactivity. We find that the innovation activities of firms in younger labor markets tend to reflect the innovative characteristics of younger individual workers aggregated to the firm level, which further supports the labor force channel.

Additionally, we also examine the effect of labor force age structure on the market value of innovations. Whereas citations to a patent capture its scientific value, the market value of a patent captures its commercial value. We find that firms in younger labor markets produce innovations that are more valuable, with a one standard deviation increase in the young share causing a roughly 12% increase in patent value. These results reflect the economic value created by younger labor forces, and they complement our results on the scientific value of innovations. We conclude that the age structure of the labor force has important consequences for inventors, firms, and the local economy as a whole.

We also distinguish between the labor force channel, according to which the age structure of the local labor force affects the innovation activities of local firms, and two other channels. First, in the financing channel, our age structure would capture the local investor pool, which is younger and more willing to finance risky projects that produce innovation. When we include the age structure at both the firm's headquarters and the R&D hubs at which the firm produces innovation, we find that R&D hubs have an incremental effect to headquarters that is similar to the magnitude of headquarters by itself. This is inconsistent with the pure financing channel. Second, in the consumption channel, our age structure would capture the local consumer pool, which is younger and demands more innovative products. When we separate firms based on whether their workforce serves local versus non-local consumers, we find that our results are driven entirely by firms with local employees but non-local customers. This is inconsistent with the pure consumption channel.

Finally, we perform sensibility checks of our main result using cross-sectional contrasts. We create intuitive contrasts at the industry and commuting zone levels. We find that the effect of age structure on innovation is stronger in more innovative industries, in more skilled labor markets, and in labor markets with more innovative universities.

Our main contribution is to extend the boundaries of the labor and finance literature. For instance, this literature studies differences across jurisdictions in employment laws that provide protections for workers or impose labor restrictions on firms. One set of papers finds that higher unemployment insurance benefits encourage entrepreneurship (Hombert, Schoar, Sraer, and Thesmar (2020)) and they increases leverage (Agrawal and Matsa (2013)). At the same time, other papers find that stronger wrongful discharge laws reduce takeover synergies and activity (John, Knyazeva, and Knyazeva (2015) and Dessaint, Golubov, and Volpin (2017)) and they lower leverage (Simintzi, Vig, and Volpin (2015)).[5]

By contrast, our paper studies demographics, a more fundamental but highly consequential aspect of labor markets, and it takes econometric advantage of the sizable native born labor force. Specifically, we are the first to show that the age structure of the local labor force has a causal effect on local innovation, and in so doing, we also contribute to the related literature on demographics and finance. Prior studies find that younger labor markets encourage firm creation and growth (Ouimet and Zarutskie (2014)); firms with a younger investor base have lower payouts (Becker, Ivkovic, and Weisbenner (2011)); and product market demographics influence stock returns (Dellavigna and Pollet (2007, 2013)). Our paper similarly

[5] For wrongful discharge laws, see also Acharya, Baghai, and Subramanian (2013, 2014) and Bai, Fairhurst, and Serfling (2020). For non-compete agreements, see Jeffers (2019).

sheds light on the debate about whether the optimal age for entrepreneurial success is relatively young or old (see Liang, Wang, and Lazear (2018) versus Azoulay, Jones, Kim, and Miranda (2020)).

We also make an important contribution to the literature on the macroeconomic consequences of demographics, which focuses on the country level. We examine innovation whereas prior studies examine unemployment, aggregate output volatility, productivity, and aggregate stock returns.[6] Similarly, Aksoy, Basso, Smith, and Grasl (2019) study demographics and innovation in a theoretical model which they estimate using panel VAR across countries and time. Likewise, Anelli, Basso, Ippedico, and Peri (2019) study migration and entrepreneurship for young versus old labor markets in a recent Italian setting. We also contribute significant methodological improvements to the broader literature on demographics. Some prior studies use births over a limited or recent number of years to instrument for actual labor force age structure (e.g., only birth rates 20-39 years prior, or only age structure 10 or 20 years prior). By contrast, we take a comprehensive and precise approach by using historical births data starting in the 1920s to construct the entire births-based labor force in each commuting zone in the U.S.

2. Methodology

2.1. Measurement of Age Structure

Our objective is to explore and better understand the effect of labor force age structure on corporate innovation. To this end, we run regressions of various measures of innovation activity on measures of labor force age structure at the same location as the firms and inventors that we examine. The difficulty is that people can choose where they live and work, and the firms at which people work can choose where they operate. Consequently, local economic activity,

[6] Respectively, Shimer (2001); Lugauer (2012) and Jaimovich and Siu (2009); Acemoglu and Restrepo (2021); and Poterba (2001) and Goyal (2004).

including innovation, can affect the local labor force because economic conditions affect the migration of people. This can be a serious problem if the population and the economy are measured contemporaneously.

Accordingly, we do not use the actual age structure itself in our analysis. Instead, we instrument it with the age structure of the births-based labor force, which we construct based on the number of births each year during the prior 20-64 years.[7] Births in a given time and place occur for a variety of reasons that may be correlated with or even caused by contemporaneous economic activity, but these reasons are plausibly exogenous to economic activity occurring several decades later. Further supporting this approach, the historical births over a long period of time in the past are an integral component of the age structure today. At the same time, the births-based age structure is plausibly uncorrelated with innovation (and indeed most economic activities) today, except through its effect on the actual age structure. As an identity, the current labor force equals births 20-64 years ago plus immigration of non-natives during the same timeframe minus emigration of natives during the same timeframe. It follows that the births-based labor force is free from the effect of migration, which would likely be correlated with contemporaneous economic activity.

We construct the births-based labor force as follows. For every year and location in the United States, we take the number of births each year during the prior 20-64 years, i.e., the population of 20-64 year olds that is native to the location. That is, for each cohort of age X, we

[7] There are several studies of other consequences of the age structure (especially unemployment and aggregate output volatility) that also instrument for the current age structure but in ways that are different from our instrumentation. Some studies use as an instrument the sum of state-level birth rates over a limited number of prior years: Shimer (2001) during the 16-24 years prior, and Lugauer (2012) during the 20-34 years prior. Somewhat differently, Jaimovich and Siu (2009) use country-level birth rates two to six decades before but at intervals of a decade (i.e., for a total of five observations). Additionally, prior studies measure the labor market at the level of the state or the country. We are able to significantly improve upon existing approaches in a number of ways. First, we use births data for the entire native born labor force in our analysis, not just for a limited and recent number of years. Second, we study local labor markets directly using commuting zones, instead of studying states or countries that necessarily combine many heterogeneous local labor markets into a single large labor market.

take the number of people born X years ago in the location. Note that this excludes Americans born outside the commuting zone as well as foreigners. The births-based young share is simply the number of 20-39 year olds born in a location as a proportion of the number of 20-64 year olds born in that location. The births-based mean age is the weighted average mean age of the 20-64 year olds born in the location, where the weight of each cohort of age X is the share of people born X years ago as a proportion of the people born 20-64 years ago. We perform analogous calculations for the actual labor force as for the births-based labor force. In the case of the births-based labor force, we additionally adjust for survival for each age group using national time-varying survival rates every decade.

We measure age structure at the commuting zone level. Commuting zones are ideal economic-geographic units to study the effect of demographics on innovation. They are designed to capture the local labor markets in which people live and work (Tolbert and Sizer (1996)). They are constructed as clusters of counties characterized by strong labor market interactions within commuting zones and weak interactions across commuting zones. We aggregate counties to commuting zones using the county-commuting zone crosswalk from Autor and Dorn (2013). There are 741 commuting zones that together cover the entire land area of the U.S.

We collect data for historical births at the county level starting in 1926. These data are from the Vital Statistics Yearbooks supplemented with and manually checked against data from Price Fishback for the early years of our sample period.[8] Similarly, we obtain population data at the county level from the Bureau of Economic Analysis. Altogether, the historical births data for our sample period span 1926-1985 (all the years needed to construct the births-based labor force,

[8] See Fishback et al. (2011) as well as Price Fishback, Jonathan Fox, Shawn Kantor, and Michael Haines, County births, deaths, infant deaths, and stillbirths, 1921-1929+A4, and also Price Fishback, Shawn Kantor, Trevor Kollman, Michael Haines, Paul Rhode, and Melissa Thomasson, Weather, demography, economy, and the New Deal at the county level, 1930-1940.

i.e., the population of 20-64 year olds, during our entire sample period, i.e., 1990-2005). We begin our sample in 1990 because we need data on historical births to construct the current population of 20-64 year olds (born 1926-1970), and births data at the county level only become well populated in the mid-1920s. We end our sample in 2005 because we examine future innovation for up to five years. By only including patents in our sample until 2010, we are able to accumulate citations to our sample patents for almost a decade, until 2019.

To support our use of the native born labor force, we provide evidence that a large proportion of Americans in the labor force (20-64 year olds) reside near their place of birth. Using decennial census data from the year 2000 (which produces results similar to 1990 and 2010), we estimate that 60% of native born American workers reside in their state of birth. Furthermore, 16% of native born American workers reside in their commuting zone of birth. This figure increases to 33% if we restrict attention to the 10 largest commuting zones[9] (which contain 27% of the U.S. population and a heavy concentration of publicly traded firms). The corresponding figures are similar for workers with at least a college education, which is a group that includes the vast majority of inventors.

2.2. Empirical Examination of Age Structure

[Insert Figure 1 about here]

We empirically examine the age structure of the labor force beginning with the historical births data that we use to construct the births-based age structure. To allow us to make not only time-series but also cross-sectional comparisons, we focus on the census region level in our next

[9] In descending order, New York, Los Angeles, Chicago, Philadelphia, Detroit, Boston, Washington, D.C., San Francisco, Houston, and Atlanta.

analysis.[10] We detrend the data, for ease of exposition, by removing the annual average across census regions using year fixed effects. Figure 1 provides a graphical description of the annual birth rates for each census region. The birth rate is measured per thousand people. As a basis of comparison for the regions, the national birth rate fluctuates considerably over the decades, ranging from roughly 15 to 26 births with a mean of approximately 20.

Several features of Figure 1 have important implications for our analysis. First, there is considerable time-series variation in birth rates, from one decade to the next, in a given location. There is also considerable cross-sectional variation in birth rates at any point in time. This implies that, decades later, there should be significant variation in the age structure across locations. Second, since birth rates vary considerably over time in a given location, periodically reverting to the cross-sectional mean (i.e., the mean across all regions in a given year), the proportion of the young versus the old (which completely determines the age structure) also varies over time. It is therefore unlikely that the cross-sectional differences in the births-based age structure are driven decades later by time-invariant heterogeneity across locations (e.g., persistent differences in economic activity).

We also examine the evolution of the age structure during our sample period (as opposed to historical births examined previously). We focus on the young share of the labor force as opposed to its mean age, but the results for both measures of the age structure are similar. We compare the evolution of the young share of the births-based and actual labor forces for each census division. Since the U.S. is steadily becoming older over time, for every division of the country, we detrend the young share, removing the annual average across census divisions using

[10] The census regions allow us to visually represent the entire country with a manageable number of units compared to the 50 states that are sorted into the four census regions. The nine census divisions lead to the same inferences, but their visual representation is more complicated.

year fixed effects, and standardize it at the census division level, subtracting the mean and dividing by the standard deviation.

[Insert Figure 2 about here]

Figure 2 shows that the births-based and actual labor forces generally follow the same direction over time. This provides validation for our use of births-based age structure to instrument for actual age structure. Indeed, the simple pairwise correlation between the births-based and actual age structures is close to 0.5, which indicates that historical births explain a significant part of the current labor force.

These features of Figure 2 also have important implications for our regression analysis. First, given the steady time trends in age structure, it seems reasonable to use one observation for every five-year period, i.e., quinquennially, since firms and inventors typically exist in the data for only a few years.[11] Second, it is generally not practical to identify entirely off the time-series variation in age structure. Instead, in our main analyses, we mainly identify off the cross-sectional variation because it accounts for most of the total variation. Since omitted cross-sectional and time-series factors can potentially affect innovation, we remove them by always including fixed effects for state-years and industry-years as well as control variables at the levels of both the firm and the commuting zone. We thus identify off the residual variation across commuting zones within a given state in a given year and within a given industry in a given year (and in fact considerably less residual variation since we include additional fixed effects and control variables). Nevertheless, in robustness tests, we do identify entirely off the time-series variation in age structure by including fixed effects for commuting zones or firms.

[11] This is similar to Becker (2007) and Becker, Ivkovic, and Weisbenner (2011) both of whom use decennial observations.

2.3. Measurement of Innovation

We use two main measures of innovation corresponding to the quantity and quality of innovation: patent counts and patent citations. These measures are defined as in the literature (see Appendix Table 1 for details). We obtain data on patents, patent assignees (firms), and patent inventors during 1975-2019 from Li et al. (2014) (with updates). To identify firms that are publicly traded, we merge these data over the same period with data from Kogan, Papanikolaou, Seru, and Stoffman (2017). The primary source of both databases is the patent data of the USPTO.

There are two important adjustments that we make. First, we adjust patent citations for the changing distribution of patent citations each technology class and year by scaling citations in a given year by the average number of citations per patent for that technology class-year. Doing so is necessary to solve the problem of those patents that are granted closer to the end of our sample period accumulating fewer citations, a problem that may vary across both technology classes and time.

Second, in our inventor-level analysis, we adjust patent counts and citations by the number of inventors on each patent. This is necessary for apportioning credit because most patents are granted to more than one inventor (in fact, three inventors, on average).[12] Specifically, for a patent with N inventors, we give each inventor credit for 1/N patents and 1/N of the patent's citations.

2.4. Model Specifications

We perform our analysis principally at the firm and inventor levels, which we explain in turn below. Our main analyses have a number of common features. We run regressions of

[12] In our firm-level analysis, this is not an issue. It is rare for the same patent to have inventors working for different publicly traded firms in our sample, so there are few patents that would be duplicated if we did not adjust for the number of inventors.

innovation outcomes on age structure. We instrument the actual age structure with the births-based age structure. Since age structure is persistent, we measure it every five years, starting in 1990, rather than annually, for a total four quinquennial period observations in our main regressions.[13]

We examine the effect of the current age structure on future innovation measured over the next five years. The first of our four adjacent and non-overlapping quinquennial periods spans 1991-1995. Since innovation activities can take several years to produce results, long-run regressions are appropriate. Moreover, innovation outcomes are relatively infrequent at the one-year horizon, so averaging the outcomes of several years generates more precise estimates than using only a single year.

Additionally, both age structure and innovation may have variation in common across space and time as well as across industry and time (e.g., because of time-varying spatial as well as industrial agglomeration). We therefore include fixed effects for state-years and industry-years in our regressions to remove these common sources of variation. An added benefit of these fixed effects is that they mitigate the biases in the patent data reported by Lerner and Seru (2017) along the dimensions of time, industry, and location.

2.5. Conditions for Instrumental Variable Validity

A valid instrumental variable needs to satisfy both relevance and exogeneity. It seems plausible that historical births affect innovation through the current age structure of the labor force. Satisfying the relevance condition does not appear to pose any difficulties, and we verify it in the first stage of our analysis.

The exclusion restriction is more difficult to satisfy, i.e., it is harder to be sure that historical births affect innovation only through the current age structure. This is because

[13] The results are similar regardless of which year from 1990 to 1994 we use to start the quinquennial periods.

historical births may affect innovation decades later through channels other than the current age structure. Alternatively, historical births may be caused by other historical factors that also affect innovation decades later but not through either historical births or the current age structure. In this case, historical births would be correlated with innovation because of their common origins.

Our task is to build regression specifications that minimize these potential violations of the exclusion restriction. We begin by offering examples of the aforementioned "other channels" and "common origins" violations. Since these types of violations are difficult to distinguish empirically, we proceed by motivating and describing the explanatory variables that we use to convincingly satisfy the exclusion restriction. We begin with contemporaneous commuting zone-level control variables, which can also capture much of the historical differences between commuting zones because such differences tend to be highly persistent. We continue with controlling for time-invariant differences using fixed effects for commuting zones and firms. We end with historical commuting zone-level control variables.

We offer examples of potential violations of the exclusion restriction, starting with "other channels". Suppose that a labor force that is more educated is, consequently, more innovative. Suppose further that governments spend more on education when the population is younger, or that younger people seek more education. Now, consider a location that experienced an explosion of births 20-39 years ago compared to births 40-64 years ago. As a consequence, such a location has a younger current labor force (assuming that migration does not completely eliminate the pattern of historical births), which has been better educated and is more innovative as a result. In this example, the proximate cause of innovation is education, not the current age structure. Education can easily be generalized to other examples of investments in local goods and services, public and private.

We also offer examples of "common origins" violations of the exclusion restriction. Suppose that a more prosperous local economy encourages younger workers to have children and at the same time increases the returns to corporate investments that produce innovation. Suppose further that the local economy becomes increasingly prosperous over time. Historical prosperity will have generated a rising birth rate over time and therefore a younger current labor force, and it will also, over time, have increased investments that produce more future innovation. In this example, it is in fact greater historical prosperity that causes a younger current labor force through its effect on historical births as well as causing more future innovation. Historical births themselves have no causal effect on anything except the current age structure.

To deal with such threats to our identification, we include in our regression specifications contemporaneous commuting zone-level control variables that we can categorize into local economic conditions and local investments. We use contemporaneous variables, even if some of the corresponding omitted factors are more naturally viewed from an historical perspective, because these variables and factors are persistent and because the requisite historical data are not always available. We capture local economic conditions using two standard contemporaneous control variables. We control for the size of the population today because it reflects the cumulative effect of shocks to the local economy over time. These same shocks can also affect both historical births and innovation today. By way of example, better economic performance 20-64 years ago attracts more people to the commuting zone and also encourages people to have more children. This increases the size of the population and it also affects the births-based age structure. At the same time, the greater scale of the commuting zone today may stimulate more business activity, including the production of innovation. Furthermore, population size is fairly persistent over time, so it also captures time-invariant features such as urban versus rural nature of the commuting zone, its vibrancy, diversity, etc. (e.g., Dougal, Parsons, and Titman (2015)), and these features can affect both the births-based age structure and innovation. Motivated by similar reasoning, we also control for income per capita as a proxy for wealth.

We capture local investments using four contemporaneous control variables. Investment in assets like infrastructure, education, and research can, over time, affect economic conditions, hence migration and births, and ultimately the births-based age structure. Through their additional effect on contemporaneous economic conditions, such investments can also affect future innovation. Our control variables include government expenditures because, by way of a specific example, historical spending on education and research can affect historical births and hence the births-based age structure, and it can also affect the business atmosphere today and provide people with the knowledge, skills, and training needed to produce innovation. Education and research spending can also affect innovation through the usual spillovers between educational institutions and the business sector. The three other control variables that we use measure important consequences of public and private investments in education and research. Specifically, we control for educational attainment, as measured by the ratio of people with a bachelor degree to the population aged 25 years or older; local university students per capita; and local university patent counts per capita.

Our contemporaneous commuting zone-level control variables plausibly capture the causes or consequences of historical births that affect future innovation – but without necessarily operating through the current age structure. This is especially likely for omitted factors that are sufficiently persistent, which appears to be the case for all of our control variables (population size, income per capita, government expenditures, educational attainment, university students per capita, and university patents per capita). However, it is possible that we still omit some important factors that are correlated with both historical births and future innovation, for example, local prosperity may not be perfectly captured by our control variables.

Such threats to our identification can be usefully grouped as time-invariant and time-varying. Trends in some omitted factors may be so persistent over a long period of time (e.g., some locations may be increasingly prosperous over time, causing an increasing number of births and more innovation) that they will appear to be time-invariant as we move through our four sample quinquennial periods (which always share at least 30 overlapping years out of 45). Such omitted factors can be captured by time-invariant fixed effects, which is what we include for commuting zones and firms. Alternatively, variation in other omitted factors may be too complex to be captured by a fixed effect even during our four sample quinquennial periods (e.g., prosperity may first rise and then fall, all within the span of the 16 years from 1990 to 2005). These types of omitted factors require time-varying control variables to capture.

We include fixed effects for commuting zones and firms to account for the possibility that some locations have consistently increasing historical births, so they have a consistently younger current labor force, and they also consistently produce more innovation. To account for time-varying economic conditions and investments rather than time-invariant differences across locations, we include, in two variations, four historical control variables that are contemporaneous to historical births (as opposed to our already included control variables that are contemporaneous to the current labor force). In the first variation, we add each of these control variables with five lags such that each lag is one decade apart from the previous lag. In the second variation, we allow these historical control variables to have different effects over time. We accomplish this by interacting each of the historical control variables and their five lags with four time fixed effects corresponding to our four quinquennial periods.

The four historical control variables that we include are broadly motivated by the earlier examples in this section, subject to the availability of the requisite data going back a century at the commuting zone level. We continue to categorize our control variables into local economic conditions and local investments. In addition to our earlier example of local economic prosperity as the causal historical factor, consider also the example of historical education spending. Greater spending on education may encourage younger workers, who are better educated, more productive, and hence wealthier, to have children, even as it increases the return to producing innovation. Decades later, the labor force is consequently both younger and more innovative, but these outcomes are not caused by historical births, which themselves are the consequence of greater historical education spending. These examples of economic prosperity and investments motivate the historical commuting zone-level control variables that we use: income per capita, the unemployment rate, educational attainment, and university students per capita. As for data sources for our control variables, we use Census Bureau data via IPUMS (with public use microdata areas (PUMAs) converted to commuting zones using weights provided by IPUMS) and our patent database.

One final approach is worth considering: exogenous shocks to the current age structure or to historical births. For instance, we could identify exogenous variation in contemporaneous immigration and use it to shock the current age structure, or we could use plausibly exogenous historical events to shock historical birth rates (e.g., wars, laws, or diseases). Unfortunately, even a priori, this approach has a number of serious limitations. Even as far as the relevance condition is concerned, it is difficult to find shocks that affect historical births in just a couple of years but which have a detectable effect on the current age structure. Consequently, we face the problem of a weak instrument. Another limitation is that for shocks to be more credibly exogenous, they need to occur at the country level or at least at the state level. Applying state- or country-level shocks to the commuting zone level runs the serious risk of the local consequences of these shocks being the result of interactions with persistent differences across locations, and of course such differences cannot be swept out mechanically without also eliminating the key source of variation at the commuting zone level.

Perhaps most importantly, it is almost impossible to argue that shocks that differentially affect locations do not have a variety of direct effects decades later on innovation. Whether wars, laws, or diseases, if such shocks leave enough of a mark on historical births to show up in the current age structure, then it is very likely that they would also damage the local economy and thus violate the exclusion restriction. Nevertheless, we did attempt to exploit immigration shocks to the current age structure as well as sex imbalance ratios driven by historical events as shocks to historical births. However, as anticipated, these instruments turned out to be too weak to generate conclusive inferences. Ultimately, we cannot use shocks to identify the effect of age structure on innovation and instead use our births-based age structure instrument.

3. Main Firm-Level Analysis

3.1. Model Specification

We begin our analysis of age structure and innovation at the firm level. We have detailed data on thousands of publicly traded firms, which allows us to control for a variety of firm-level characteristics. The equation for the baseline regressions is:

$$Innovation_{i,j,a,c,s,t+1} = \alpha \cdot Age\_Structure_{c,t} + \beta \cdot X_{i,t} + \gamma \cdot X_{c,t} + \delta_{s,t} + \delta_{j,t} + \delta_{a} + \varepsilon \quad (1)$$

where i indexes firms, j indexes industries, a indexes firm age, c indexes commuting zones, s indexes states, and t indexes years. $X_{i,t}$ is a vector of firm-level control variables, $X_{c,t}$ is a vector of commuting zone-level control variables, $\delta_{s,t}$ is a state-year fixed effect, $\delta_{j,t}$ is an industry-year fixed effect, and δ_a is a firm age fixed effect. We use publicly traded firms, which account for roughly half of all patents granted to public and private firms combined. The location of innovation is determined by the address on the patent document.

We measure age structure for firms using the commuting zone in which they are headquartered. While we would ideally like to measure age structure in a manner that reflects the actual location of the entire workforce of the firm in commuting zones across the country, these data are not available for our sample firms. Our choices are either using the firm's headquarters location, which we can measure precisely, or using R&D hubs, which we, like the literature, are limited to inferring with considerable noise from our patent database. Moreover, starting from our baseline specifications, we include state-year fixed effects, which we could not implement for inferred R&D hubs weighted by the number of inventors. Similarly, in one important test, we include fixed effects for commuting zones, which we also could not implement using inferred R&D hubs because these data are not available for a sufficiently long period of time. For these reasons, which we discuss in detail in Appendix 1, we use age structure at the firm's headquarters in most of our tests. In robustness tests, we verify that our results are not sensitive to this choice.

We include control variables at the levels of both the firm and the commuting zone. The firm-level control variables are total assets, market-to-book, cash flow-to-total assets, stock returns, and stock return volatility. For reasons explained in Section 2.2, we also include state-year fixed effects and industry-year fixed effects, where industry is captured by two-digit SIC codes.

Additionally, both age structure and innovation may have variation in common across different levels of firm age (e.g., Adelino, Ma, and Robinson (2017)). For instance, firms and their locations may become older and less innovative over time. Firm age as a control variable may not completely capture firm life cycle commonalities because the relationship may not be linear, so we instead include firm age fixed effects as captured by five-year groups of firm age in a piecewise linear fashion. Consequently, the results must be interpreted as for firms of the same age. Following standard practice, we measure firm age from the date the firm begins trading publicly according to CRSP.

Finally, since firms in similar lines of business tend to behave in a correlated fashion, we cluster standard errors by industry-year in our baseline regressions. However, we show that the results are robust to alternative forms of clustering. Before taking the logarithm of a variable that takes on zero values, we add a small constant that approximately equals the smallest increment of the values of the variable. We verify that our results are robust to adding a constant that is higher or lower by up to two orders of magnitude. We winsorize variables whenever appropriate at the 1^{st} and 99^{th} percentiles.

3.2. Sample and Descriptive Statistics

The firms in our sample are publicly traded U.S. operating firms excluding financials and utilities. The sample itself comprises 15,730 firm-quinquennial period observations corresponding to 8,002 unique firms and 321 unique commuting zones. Roughly 40% of firms have at least one patent sometime during the next five years. At this level of analysis, we do not restrict the sample to firms in our patent database. Data on publicly traded firms are from CRSP and Compustat.

[Insert Table 1 about here]

Table 1 provides descriptive statistics for the samples corresponding to our firm and inventor levels of analysis in Panels A and B, respectively. Age structure is comparable across all both levels of analysis. The typical age of the labor force is a bit under 40 years, and the young share is a bit over 50%. There are about 5 patents per annum on average (median 0) at the firm-year level (Panel A).

3.3. Results

3.3.1. Births-Based Age Structure and Actual Age Structure

[Insert Table 2 about here]

Table 2 presents the results of regressions of actual age structure on births-based age structure, i.e., the first stage of the instrumental variables regressions. The simplest, univariate regression specification (Column 1) shows that births-based age structure significantly predicts actual age structure. For a range of reasonable specifications most of which we use for our second stage regressions (i.e., Columns 4 to 8), the first stage results are similar. Column 2 adds state-year fixed effects, and Column 3 further adds commuting zone-level control variables.

Further additions to the specification in Column 3 do not meaningfully affect the coefficient on births-based age structure. Column 4 adds firm-level control variables as well as fixed effects for industry-years and firm age groups. The estimate of the elasticity of actual age structure with respect to births-based age structure stabilizes at approximately 0.10, it is highly statistically significant, and the R^2 is a high 90% or so for both young share and mean age. Columns 5 and 6 show that the results are similar when we include fixed effects for commuting zones and firms, respectively. The sample size increases in these cases because we measure age structure slightly differently (adding 10 years to our sample period by measuring the labor force until age 54 rather than 64, as we explain in Section 3.3.3). Columns 7 and 8 also show similar results when we include five decennial lags of historical commuting zone-level control variables and also when we interact these variables with quinquennial period fixed effects.

To choose our baseline regression specification, we face a tradeoff: convincingly satisfying the exclusion restriction versus including so many other explanatory variables that we eliminate the variation in births-based age structure that exogenously affects innovation through actual age structure. We take a step back from our first stage regressions and examine the extent to which variation in births-based age structure can be explained by the rest of our other explanatory variables in Table 2 (i.e., other than births-based age structure). Starting with Table 2 Column 4, when we regress births-based age structure on these other explanatory variables (results not tabulated), the R^2 is already 83%-84%. This is consistent with our argument in Section 2.5 that our contemporaneous control variables capture a large proportion of the variation in age structure.

By comparison, in Columns 5 and 6, i.e., when we include the corresponding additional fixed effects, the R^2 rises to 99%. That these fixed effects leave almost no residual variation in births-based age structure is not surprising because age structure is a slow moving variable, and our sample period is relatively short. However, these results indicate that it is difficult to identify the effect of age structure on innovation in our baseline specification only off time-series variation. Instead, we more reasonably identify also off cross-sectional variation. As for Columns 7 and 8, i.e., when we add historical commuting zone control variables, the R^2 again rises compared to Column 4, to 86%-88%. With the specifications in Columns 7 and 8, given that there is only 16%-17% of variation in age structure remaining in Column 4, we would give up 3-4 percentage points, or about 20%-25%, of the residual variation, which is significant. In these specifications, we include many additional control variables (four variables, with up to five lags, and interacted with four quinquennial periods). In light of the aforementioned tradeoff we face, we use the more parsimonious specification in Column 4 as our baseline, but we do check that our results are robust to using the specifications in Columns 5 to 8.

3.3.2. Age Structure and Innovation

[Insert Table 3 about here]

Table 3 presents the results of regressions of innovation on age structure, i.e., the second stage of the instrumental variables regressions. These baseline firm-level regression results show that younger labor forces produce more innovation. A typical increase in young share (0.80 percentage points)[14] causes a roughly 8% (= 10.4 × 0.80 p.p.) increase in patent counts and a roughly 10% (= 11.9 × 0.80 p.p.) increase in patent citations. A typical decrease in mean age (0.22 years) causes similar magnitude increases in patent counts and citations (= -0.37 × -0.22 years and = -0.43 × -0.22 years, respectively). Appendix Table 2 shows that the results are robust to various alternative forms of clustering: by commuting zone-year (Panel A), commuting zone (Panel B), firm (Panel C), industry-year and commuting zone-year (Panel D), and industry-year and commuting zone (Panel E).[15]

Furthermore, our estimates of the effect of age structure on innovation seem reasonable in magnitude. The local labor force not only affects the firm itself directly but also indirectly, through its effect on the general work environment in the local labor market. This combination of effects on innovation can be powerful. The evidence from the literature also indicates that labor

[14] We translate a typical change in births-based age structure to the equivalent change in actual age structure as follows. The elasticity of the actual and births-based age structures is roughly 0.10 (Table 2), and a one standard deviation change in births-based young share is 8.4 percentage points (Table 1 Panel A), so a typical increase in births-based young share equals a roughly 0.80 p.p. (= 0.10 × 8.4 p.p.) increase in actual young share. For mean age, the corresponding decrease in the births-based measure equals a roughly 0.22 year (= 0.10 × 2.2 years) decrease in the actual measure. These are the typical changes in actual age structure (0.80 p.p. for young share and 0.22 years for mean age) that we use to interpret its effect on innovation in most of the paper.

[15] We do not cluster by time alone (i.e., without interacting it with industry, commuting zone, etc.) because doing so would impose a very strong assumption, economically and econometrically. Specifically, it would assume that all firms in all five years of a given quinquennial period could behave in correlated fashion across industry, location, and indeed any other possible dimension. Additionally, four time clusters (i.e., four quinquennial periods) would violate the well known rule of thumb for the minimum number of clusters required to avoid issues with drawing inferences.

force demographics have a powerful effect on numerous economic outcomes. For example, examining labor markets at the state level, Lugauer (2012) finds that a 10 percentage point increase in the young share causes a 58% decrease in GDP volatility relative to its mean. Similarly, at the country level, Liang, Wang, and Lazear (2018) find that a one standard deviation decrease in the median age of the labor force causes a 40% increase in new business formation. We study local labor markets and find an effect on corporate innovation of comparable magnitude to the literature from a one standard deviation increase in age structure.

Finally, we examine the effect of age structure on innovation using the uninstrumented actual age structure. We rerun the regressions in Table 3, but we do not instrument the actual age structure with the births-based age structure. The results in Table 4 lead to similar inferences as the results in Table 3, but the estimates in the former are smaller in economic magnitude.

We consider why the OLS estimates may be smaller than the IV estimates. Neither prior literature nor first principles provide clarity about the direction of the bias of endogenous OLS estimates relative to the true effect of age structure on innovation. The OLS bias depends on which of the violations of the conditions for instrumental variable validity are dominant (for example, whether younger locations tend to attract younger migrants rather than older ones). However, we do know that, as an identity, the difference between the current labor force and the labor force based on historical births is the net migration (immigration minus emigration) of the working age population. Therefore, any difference between estimates based on actual age structure (Table 4) and births-based age structure (Table 3) must arise from net migration during the prior 20-64 years.

We provide a simple illustration of why OLS could lead to biased inferences. Factually, part of the current population includes recent immigrants (people from outside the commuting zone) who may not instantaneously achieve the same productivity as the rest of the labor force. This could happen because we measure current age structure and innovation relatively contemporaneously (with at most a five year lag), whereas in practice adaptation to the local work environment takes time. In this illustration, we could have a problem if, for example, migration patterns vary across commuting zones as a function of local innovation.

Census data support this possibility. If we split cities into high and low innovation groups based on patent counts per capita in the year 2000 and compare migration patterns for the two groups, we find that high versus low innovation cities have a higher proportion of recent (last five years) immigrants (24% of the labor force versus 20%) as well as immigrants who are younger (mean age of 35.7 years old versus 36.8). Consequently, it appears that recent immigrants reduce the mean age of high versus low innovation cities. Applying this illustration to our regressions, the variation in innovation between high versus low innovation cities is the same in both the OLS and IV analyses, but it will be explained by variation in age structure (between the corresponding high versus low innovation cities) that is greater for the current labor force (and hence the actual age structure) than the labor force based on historical births (and hence the births-based age structure). It follows that sensitivity of innovation to age structure will be estimated to be smaller in OLS than IV regressions.

This foregoing simple illustration covers only a single bias based on a static comparison between actual and births-based age structures. In fact, the differences between these age structures and the resulting effects on the OLS versus IV estimates occur over longer periods of time (up to 64 years) and are influenced by a multitude of factors (including those for which we control in our regressions). Our approach to regression analyses is to use the native born labor force to construct the births-based age structure as an instrument for the actual age structure, thereby removing the migration component of the actual labor force. We likewise use demanding regression specifications that are designed to mitigate remaining econometric concerns. This approach can plausibly avoid the problems associated with the actual age structure.

3.3.3. Alternative Specifications

While we argue that our baseline specification already captures potentially omitted factors, we use our baseline specification as a point of departure and examine alternative specifications. We do this to address the possibility that historical births are correlated with omitted factors that affect innovation and thereby threaten our identification. Categorizing such factors as time-invariant and time-varying, we first examine whether our results can be explained by time-invariant heterogeneity across commuting zones. Since our sample has four adjacent and non-overlapping quinquennial period observations, and age structure is a slow moving variable that we measure from 1990 to 2005, it is generally not practical to identify entirely off the time-series variation in age structure. Nevertheless, this may be possible if we simultaneously extend our sample period and shorten the age window that we use to measure age structure. In particular, rather than using the entire labor force (ages 20-64), we can use the labor force of prime working age (ages 20-54). By subtracting 10 years from measuring age structure, we can add 10 years to our sample period in exchange for a little less accuracy in measuring age structure. Consequently, we can have six quinquennial periods starting in 1980 rather than four quinquennial periods starting in 1990. We rerun the regressions in Table 3 with these slight modifications to our sample period and measures of age structure.

[Insert Table 5 about here]

Table 5 Panel A shows that the results are broadly similar to those in Table 3 when we identify entirely off the time-series variation within commuting zones (and within state-years, industry-years, and firm age groups). We also rerun the same regressions but replace commuting zone fixed effects with even more demanding firm fixed effects (which subsume commuting zone fixed effects). Table 5 Panel B shows that the results are again broadly similar. In both panels, the results are slightly stronger for patent counts and fairly similar for patent citations. Overall, we can rule out the possibility that our results are driven by unobservable heterogeneity across local labor markets or firms.

We then examine whether our results can be explained by time-varying heterogeneity across commuting zones. We first include four historical control variables that capture time-varying local economic conditions and local investments. Second, we allow these four variables to have different effects in each of our four sample quinquennial periods. The four variables are income per capita and the unemployment rate (representing economic conditions) as well as educational attainment and university students per capita (representing investments). We include five lags of each of these variables, each lag is one decade apart from the previous lag, and the lags represent the middle of the decade to which they correspond (e.g., for the 20-29 year olds, we use a 25 year lag).[16]

Table 5 Panels C and D show that the results are similar to those in Table 3 not only when we include five lags of the four historical control variables but even when we allow these variables and their lags to vary in each quinquennial period (i.e., when we interact the control variables with quinquennial period fixed effects). If omitted factors were driving our results, we would have expected that demanding specifications that add key historical control variables would substantially weaken or eliminate the results in Table 3, whereas in fact they hold up quite well. Ultimately, our results survive the inclusion of fixed effects for firms and commuting zones

[16] Owing to insufficient historical data, we are forced to omit the fourth and fifths lags of income per capita and educational attainment, and we must likewise omit the fifth lag of the unemployment rate.

as well as the inclusion of historical commuting zone-level control variables the effects of which are allowed to vary over time.

3.3.4. Robustness Checks

For reasons that we articulated previously, we use the firm's headquarters location to measure age structure. This choice may be problematic if the firm produces innovation across commuting zones in which age structure has a different effect on innovation compared to the firm's headquarters. Therefore, we test whether our results are sensitive this choice. First, we modify the outcome to be innovation produced only at the firm's headquarters, and otherwise we redo our baseline regressions in Table 3. That is, we consider the effect of age structure on innovation produced in the exact same location, the firm's headquarters, and we ignore the firm's other R&D hubs. To get a sense of the importance of the firm's headquarters relative to all of its R&D hubs, we calculate the proportion of the firm's inventors located in the commuting zone of the firm's headquarters. This proportion turns out to be roughly 50% for both the mean and median firm with at least one inferred R&D hub, suggesting that a firm's headquarters is typically also its largest R&D hub.

Second, again in line with Table 3, we consider the firm's innovation as a whole regardless of where it is produced, but instead we modify age structure and the commuting zone-level control variables to reflect the locations of the firm's R&D hubs. Specifically, we construct weighted averages of these variables using as weights the number of inventors at each inferred R&D hub. We treat firms with no inferred R&D hubs as if they had a single R&D hub located at their headquarters. We continue to use headquarters location for the state-year fixed effects.

[Insert Table 6 about here]

For both of these robustness tests, we first verify that the results of the first stage regressions (not tabulated) are similar to those of our baseline. We then turn to second stage regressions. The results presented in Panels A and B of Table 6 show that our baseline results (Table 3) do not depend on using age structure at the firm's headquarters, as we do in most of our tests.

Turning to further robustness checks, we first examine whether the results are impacted by complementarity between the young and the old. In the presence of complementarity, a more dispersed age structure, i.e., both more young and more old, should increase innovation. This could be the case unconditionally, or conditional upon age structure, i.e., the better the balance between the young and the old, the more innovation should increase as a result of a younger labor force. We test these predictions by starting with the baseline regressions in Table 3 and first including only age structure dispersion, which we measure as the coefficient of variation of age structure (the standard deviation divided by the mean). As we do for age structure itself, we instrument actual dispersion with births-based dispersion. Appendix Table 3 Panel A shows that the effect of age structure dispersion is not statistically significant. Turning to our conditional prediction, we also include the interaction of age structure with its dispersion. Appendix Table 3 Panel B shows that age structure dispersion has a negative effect on innovation, and conditional upon the increase in innovation resulting from a younger labor force, greater dispersion in labor force age structure does not have a reliably significant relationship with innovation.

Next, we examine whether our results depend on the measurement of age structure using firm location that is time-invariant rather than time-varying. Instead of Compustat header headquarters location, we use EDGAR historical headquarters location based on data obtained from Bill McDonald. Unlike Compustat location data, EDGAR location data are only available starting in the mid-1990s. We begin by examining our baseline results (Table 3) for the roughly half of the full sample of firm-years for which time-varying location is available. In this restricted sample and using time-invariant location, the results are in fact roughly 40% larger and otherwise similar (not tabulated). We then examine our baseline results using time-varying location. Appendix Table 3 Panel C shows that the results are broadly similar to the restricted sample with time-invariant location. Since the choice of time-invariant versus time-varying location does not materially change our results, we use time-invariant firm location throughout the paper because data availability would otherwise reduce the sample size by roughly half.

We also explore whether our results can be explained by firms changing their location so as to match their innovation production to the local age structure. To address this possibility, we restrict our sample to firms which we can determine as not changing location at least during our sample period. We construct this sample by supplementing EDGAR historical headquarters location with IPO location from SDC. The results for "non-mover" firms are shown in Appendix Table 3 Panel D, and they are comparable to the results (not tabulated) for the combined sample of non-mover and mover firms (the latter identified based on EDGAR). These finding suggest firms changing their location does not explain our main results.

Turning from the measurement of age structure to the measurement of firm age, we note that we use listing dates to measure firm age throughout the paper rather than founding dates. To directly examine whether this makes a difference to our results, we obtain data on founding dates for some publicly traded firms from Jovanovic and Rousseau (2001) and for many IPOs from the Field-Ritter database (Field and Karpoff (2002) and Loughran and Ritter (2004)). These data cover about 80% of the full sample of firm-years. We begin by verifying our baseline results (Table 3) for the restricted sample with founding dates, and we find that the results are about

20% smaller and otherwise similar (not tabulated). We then examine our baseline results using founding dates to measure firm age rather than listing dates. The results, shown in Appendix Table 3 Panel E, are similar.

We also consider the possibility that the age structure of the labor force may be capturing the effect of managerial age on innovation. Younger corporate managers may be more prevalent in younger labor markets, and firms may produce more innovation as a result of managerial characteristics (see, e.g., Acemoglu, Akcigit, and Celik (2020)) rather than the characteristics of the labor force more broadly. To isolate the effect of labor force age structure from managerial age, we use data from Execucomp and add CEO age as a control variable in our regressions. Panel F of Appendix Table 3 shows that the results for labor force age structure are robust to controlling for managerial age, and indeed these results are similar to our baseline results restricted to the sample for which we have managerial age (not tabulated). Additionally, younger CEOs are associated with more innovation (results not tabulated).

Since both people and firms in the U.S., in correlated fashion, are more concentrated in some commuting zones that others, we also examine whether our results are driven by the largest commuting zones. Specifically, we successively drop the top 3, 5, and 10 the most populated commuting zones and correspondingly examine our baseline results. In each case, the results are similar for the restricted samples (not tabulated) and the full sample.

Finally, we also create intuitive cross-sectional contrasts to check the sensibility of our main result. To save space here, this analysis and the results are described in Appendix 2.

4. Inventor-Level Analysis

4.1. Model Specification

We continue our analysis of age structure and innovation at the more granular level of the inventor rather than the firm. Among other advantages, this setting allows us to distinguish between inventors and other workers in the local labor force. We have thousands of inventors in public and private firms, which allows us to isolate the effect of labor force age structure from inventor age, time-varying firm-specific omitted factors, and network scale effects within firms. The equation for the baseline regressions is:

$$Innovation_{i,j,k,c,s,t+1} = \alpha \cdot Age\_Structure_{c,t} + \beta \cdot X_{k,t} + \gamma \cdot X_{c,t} + \delta_{s,t} + \delta_{j,t} + \delta_{i,t} + \varepsilon \quad (2)$$

where i indexes firms, j indexes technology classes, k indexes inventors, c indexes commuting zones, s indexes states, and t indexes years. $X_{k,t}$ is a vector of inventor-level control variables, $X_{c,t}$ is a vector of commuting zone-level control variables, $\delta_{s,t}$ is a state-year fixed effect, $\delta_{j,t}$ is a technology class-year fixed effect, and $\delta_{i,t}$ is a firm-year fixed effect. The location of inventors, like that of innovation, is determined by the address on the patent document. To ensure that inventor locations are reasonably accurate relative to the beginning of our quinquennial period observations, we only use locations that are no more than five years old. Furthermore, to allow inventors enough time to establish a track record, we count the number of patents per inventor during the prior 10 years, and we select the top 5% of inventors based on number of patents.[17]

[17] Moretti and Wilson (2017) use the top 5% of inventors, and Akcigit, Baslandze, and Stantcheva (2016) use the top 1%. We focus on star inventors for a number of reasons. The top several percent of inventors account for the majority of patents and citations (details in the previous two papers). Moreover, we can determine the location of star inventors and measure their innovation output with good temporal precision. Since location is determined based on patent documents, it is only star inventors who have enough patents to allow us to reliably determine inventor location at least once every few years and to measure inventor innovation outputs over the next five years. We can confirm in our data that star inventors move at a similar rate to firms changing their headquarters location: about 1.5% per annum. This stability alleviates concerns that our results may be explained by more productive inventors choosing younger locations.

We measure age structure for inventors using the commuting zone in which they are located. While our inventor-level analysis covers both public and private firms, our data on private firms are not as rich as our data on public firms as far as control variables are concerned. However, we are able to include firm-year fixed effects and thereby identify off the variation across inventors within a given firm in a given year. We are also able to control for inventor age, which is important because younger inventors are likely to be more innovative. Controlling for inventor age allows us to isolate its effect from labor force age structure. Our data on inventor birth dates are from Kaltenberg, Jaffe, and Lachman (2021). Additionally, since inventors with more experience are likely to produce more patents, we control for inventor experience measured from the date of the inventor's first patent. Furthermore, inventors that produced more patents in the past are likely to produce more patents in the future, so we control for the patent stock of the inventor. Moreover, inventors working in larger groups may be more innovative. We account for such network scale effects using the number of inventors working for the firm at the same inferred R&D hub as a given inventor.

Additionally, we include technology class-year fixed effects at the inventor level as analogous to industry-year fixed effects at the firm level. The technology class of an inventor is the single technology field (out of roughly 500 possible fields) in which he has the largest number of patents. We also include state-year fixed effects as before. Finally, since inventors working for similar businesses tend to behave in correlated fashion, we cluster standard errors by technology class-year, but we show that the results are robust to alternative forms of clustering.

4.2. Sample and Descriptive Statistics

Since the inventors in our sample are stars by construction, they unsurprisingly account for over 20% of all patents and over 25% of all patent citations. The firms at which these inventors work are the public and private firms in our patent database. They break down into roughly one-third public firms and two-thirds private firms. The sample itself comprises 14,541 inventor-quinquennial period observations corresponding to 9,843 unique inventors, 1,728 unique firms, and 303 unique commuting zones.

Returning to the descriptive statistics in Table 1, and focusing now on Panel B, there is roughly 1 patent per annum on average (median of about 0.5) at the inventor-year level. The typical sample inventor is about 49 years old and has approximately 16 years of experience relative to the date of his first patent. He produces an average of 2.5 patents per year (median of 2.1) (inventor patent stock divided by inventor experience).[18]

4.3. Results

[Insert Table 7 about here]

The baseline inventor-level regression results are presented in Table 7, and they once again confirm that younger labor forces produce more innovation.[19] A typical increase in young share (0.92 percentage points) causes a roughly 7% (= 7.3 × 0.92 p.p.) increase in patent counts and a roughly 9% (= 9.4 × 0.92 p.p.) increase in patent citations. A typical decrease in mean age (0.24 years) causes similar magnitude increases in patent counts (= -0.24 × -0.24 years) and

[18] The control variables in Table 1 Panel B are tabulated at the inventor level, like the other variables. This is necessary for correctly interpreting the regression results. However, these control variables are not very meaningful to interpret as descriptive statistics because the sample is duplicated for firm R&D hubs and firms with multiple inventors. The results are therefore driven by those firm R&D hubs and firms that have the most inventors. For instance, at the firm R&D hub level, there are a mean of 161 inventors (rather than the 693 inventors tabulated at the inventor level). Additionally, for a number of reasons, innovation is difficult to compare across levels of analysis. First, each level of analysis corresponds to a different degree of aggregation of patents. Some firm-years have no inventors associated with them, while others have multiple inventors. Second, only star inventors are used by construction at the inventor level, whereas all inventors are used at the firm level. Third, patent counts and citations are adjusted by the number of inventors on each patent, but only at the inventor level.

[19] We add a smaller increment of 0.01 to patents per annum before taking logarithms in our inventor-level analyses rather than 1 as in our firm-level analyses because inventors have about two orders of magnitude fewer patents than firms (c.f. Panels A and B of Table 1).

citations (= -0.32 × -0.24 years).[20] Appendix Table 4 shows that the results are robust to alternative forms of clustering.

The inventor age control variable provides a plausible complement to labor force age structure. Inventor age has a negative relationship with patent citations, but the results are not significant, economically or statistically, for patent counts. While this evidence is neither causal nor strong, it does suggest that individuals, like labor markets, are more innovative when they are younger.[21]

In our inventor-level analyses, we are comparing two inventors with similar characteristics who work for the same firm at the same time but in different commuting zones with different age structures. The results show that a given inventor is more innovative when other workers around him are younger, even after taking into account his own age. These other workers include not just those within his firm but also outside of it, and not just fellow inventors but also other workers in the local labor force. Indeed, local labor force age structure may not even operate primarily through the age of the inventor identified on the patent document, who is typically a relatively experienced and senior employee. Instead, it may operate mainly through relatively junior graduate students, post-doctoral researchers, scientists, engineers, etc. who work for the inventor, much like in the academic setting examined by Bowen, Frésard, and Taillard (2017). Such knowledge spillovers generated by the interactions of inventors are captured by labor force age structure.

---

[20] In this analysis, the elasticity of the actual and births-based age structures is 0.12 compared to 0.10 in our baseline analysis, and the standard deviation of the births-based age structure measures is roughly 10% smaller.

[21] There is a small literature that documents the relationship between individual inventor age and innovation (Jones (2010) and Jones and Weinberg (2011)). However, individual inventor age and labor force age structure are constructs that are distinctly different from each other and which cannot be directly compared. To illustrate, with each passing year, a given inventor ages by exactly one year while accruing seniority, resources, and managerial responsibilities. This inventor-level life cycle of innovation does not translate to the level of the labor force, which is an ever changing social group comprising both inventors and non-inventors.

5. Innovation Characteristics

The premise of the labor force channel is that younger labor forces are more innovative as a result of the various aforementioned characteristics of younger people: such workers are more creative, are willing to take more risk, have longer horizons, and are more socially interactive. If this is the case, then these characteristics should be reflected in the innovation activities of firms in younger labor markets. We therefore examine the corresponding characteristics of patent outputs at the firm level: creativity, riskiness, longevity, and interactivity. To this end, we develop a number of novel measures of innovation activities.

To capture these characteristics, we use patent citations. This restricts our sample to the roughly 40% of firms that have at least one patent sometime during the next five years. We use a five-year horizon because many firms do not have a patent every year. Furthermore, we use two types of citations: backward citations, which are citations made by a patent to previous patents, and forward citations, which are citations made to a patent by future patents. Moreover, all of our measures are suitably adjusted (see Appendix Table 1 for details) to ensure their precision (especially so that they do not mechanically increase in the number of citations) and to address the limitations of the patent data (particularly truncation).

Our measures of the first characteristic of younger people, creativity, are based on the legal requirement that inventions can be patented only if they are useful and novel. We have three firm-level measures of creativity, of which the first two capture usefulness and the third captures both usefulness and novelty. The first measure is the mean number of forward citations per patent. The second is the proportion of a firm's patents in the top 1% of forward citations.

The third is the proportion of a firm's patents in both the top 20% of forward citations (i.e., the most useful) and the bottom 20% of backward citations (i.e., the most novel).[22]

Riskiness, our measure of the second characteristic of younger people, is the firm-level volatility of forward citations per patent. The underlying intuition is that the dispersion of citations per patent of the firm's patent portfolio provides an estimate of the ex ante risk of the firm's innovation activities. Longevity, our measure of the third characteristic of younger people, is the firm-level mean age of the newest forward citation per patent, where citation age is measured relative to the grant year. The underlying intuition is that longer lived inventions should continue to be cited for many years. The length of time over which a firm's patent portfolio is cited provides an estimate of the ex ante horizon of the firm's innovation activities. To ensure that this measure does not simply capture the number of forward citations of a patent, we scale this variable calculated for a given firm by the mean of the same variable calculated using all patents in the same technology class and grant year cohort and citation decile.

Interactivity, our measure of the final characteristic of younger people, is the mean proportion of backward citations of a firm's patents to patents in the firm's commuting zone. The underlying intuition is that greater interaction within local labor markets should generate more local knowledge spillovers. The extent to which a firm's patent portfolio makes local citations provides an estimate of the ex ante degree of local knowledge incorporated into the firm's innovation activities. To ensure that our measure does not simply capture the general popularity of the firm's commuting zone, we exclude self citations, and we scale this variable by the mean proportion, in the same technology class and grant year cohort, of backward citations of all patents in all commuting zones to patents in the firm's commuting zone.

---

[22] For the second measure, our inferences are similar if we raise the threshold for forward citations to 0.1% or lower it to 10%. For the third measure, we also draw similar inferences if we use thresholds for forward citations in the top 10%-25% and thresholds for backward citations in the bottom 10%-25%.

We run the same regressions as in Table 3 but using innovation characteristics as the dependent variable instead of innovation outputs. Here as elsewhere, the dependent variables are bounded below by zero and typically right skewed, so we take their natural logarithms. We make an exception for riskiness and longevity because they are symmetrically distributed. Firm-year observations with no patents are treated as missing.

[Insert Table 8 about here]

Table 8 shows that younger labor forces produce innovation outputs that tend to reflect the innovative characteristics of younger labor forces. According to the results, the patents of firms in younger labor markets are more creative (although citations per patent is not statistically significant at conventional levels), more volatile, and involve more local knowledge spillovers. This is the case whether age structure is measured using young share or mean age. Appendix Table 5 shows that the results are generally robust to alternative forms of clustering.

6. Value Implications

Having provided evidence showing that firms in younger labor markets produce more innovation, we examine whether such firms produce more valuable innovations. If firms are able to successfully commercialize the innovations resulting from a younger labor force, then this should be reflected in the market value of these innovations. We distinguish here between a patent's commercial value and its scientific value, the latter of which is reflected by the citations to the patent.

We measure the market value of innovations using the patent value estimates of Kogan, Papanikolaou, Seru, and Stoffman (2017) based on data from Noah Stoffman. We follow the same procedure for innovation value as for innovation outputs (Table 3). We first sum up patent values to the firm level, and then we take the annual average over the next five years. Finally, we run firm-level regressions using patent value as the dependent variable instead of patent outputs. Firm-year observations with no patents are treated as missing.

[Insert Table 9 about here]

Table 9 presents the results, which show that firms in younger labor markets produce innovations that are more valuable. A typical increase in young share (1.18 percentage points) causes a roughly 12% increase in patent value (=10.2 × 1.18 p.p.). A typical decrease in mean age (0.32 years) causes a similar magnitude increase in patent value (= -0.38 × -0.32 years).[23]

We also perform a number of robustness checks of the results. In each case, we measure the market value of innovations in an alternative way. First, rather than treating firm-year observations with no patents as missing (e.g., as in Table 8), we treat them as having zero patent value (e.g., as in Table 3). Second, we scale the baseline patent value (Table 9) by total assets. Third, we both treat missing patents as zero patents and scale patent value by total assets. In all three cases, the results are similar to those in Table 9 (not tabulated). Overall, the successful commercialization by firms of the innovations resulting from a younger labor force appear to be reflected in the market value of these innovations.

7. Alternative Channels

7.1. The Alternative Financing Channel

Our results thus far are consistent with the labor force channel. In this channel, younger people are more innovative, and firms in younger labor markets are able to hire younger workers who in turn produce more innovation for firms. Nevertheless, it is possible that our results are consistent with an alternative financing channel in which our age structure measures capture the local investor pool rather than the local labor force. There is evidence that younger people tend

[23] In this analysis, the elasticity of the actual and births-based age structures is 0.14 compared to 0.10 in our baseline analysis.

to invest in more risky assets than older people,[24] and investors in general tend to hold relatively more of their portfolio in local firms (e.g., Coval and Moskowitz (1999)). This implies that younger local investors are more willing to finance risky projects that produce innovation. In this case, it is younger investors who cause firms to produce more innovation.

We can test this alternative channel by separating the locations of inventors and investors. We use the firm's inferred R&D hubs to proxy for the location of its inventors and the firm's headquarters to proxy for the location of its investors. Since R&D hubs are identified based on inference from patents that are granted, we only have location data for inferred R&D hubs for about half of our full sample of firm-years.

[Insert Table 10 about here]

We construct the weighted average age structure of the firm across its inferred R&D hubs just like we construct its age structure at its headquarters, but we use the number of inventors at each inferred R&D hub as weights. We rerun the regressions in Table 3 but include the age structures at headquarters only or at both headquarters and inferred R&D hubs. Panel A of Table 10 shows that the results for headquarters only are similar to our baseline results (Table 3) allowing for some difference in magnitudes. Panel B of Table 10 shows that the age structures at headquarters and inferred R&D hubs both affect innovation, and the magnitude of each effect in Panel B is similar to that in Panel A. The fact that headquarters is significant alongside inferred R&D hubs is likely due to the fact that the firm's innovation output is not produced by the firm's inventors in complete isolation but rather in collaboration with the firm's other employees.

[24] Fagereng, Gottlieb, and Guiso (2017) find that, as people age, they reduce the share of their portfolio invested in stocks. Betermier, Calvet, and Sodini (2017) find that older people shift their stock portfolio from higher risk to lower risk stocks. Others find that the share of old people in the population is associated with a greater supply of low risk investments and financing such as bank deposits (Becker (2007)) as well as greater dividend payments to shareholders by local firms (Becker, Ivkovic, and Weisbenner (2011)).

Additionally, we test the main prediction of the financing channel for capital structure: firms with younger investors should have lower leverage. Since equity is a more risky financial claim than debt, if younger investors are more willing to finance risky projects, then they should provide firms with more financing in the form of equity rather than debt. Therefore, all else equal, firms in younger locations should have lower leverage. We rerun the regressions in Table 3 but with leverage as the dependent variable, and we find that age structure is not significantly related to leverage, economically or statistically (results not tabulated). While we want to be very careful in our interpretation because financing policy may also be determined by investment policy, this non-result for the financing alternative hypothesis is reassuring for the labor force hypothesis. Taken together, our results allow us to rule out the pure financing channel.

7.2. The Alternative Consumption Channel

It is also possible that our results are consistent with an alternative consumption channel in which our age structure measures capture the local consumer pool rather than the local labor force. The assumption underlying this channel is that younger local consumers demand more innovative products, so the firm responds by producing more innovation. We can test this alternative channel by separating firms based on whether their workforce serves local versus non-local consumers. Following Mian and Sufi (2014), our central insight is that production and consumption coincide in space and time for firms in non-tradable industries, i.e., the labor force and consumption are both local. However, for tradable industries, the labor force is local whereas consumption is not. Consequently, the local labor force may capture both employees and customers in non-tradable industries, but it only captures employees, not customers, in tradable industries.

We examine the extent to which our results are driven by firms in tradable versus non-tradable industries. We sort firms into tradable and non-tradable industries based on the two measures used by Mian and Sufi (2014),[25] and we rerun the regressions in Table 3. The specifications include the non-tradable industry dummy variable and its interaction with age structure and all control variables. Firms in tradable industries are the base group.

[Insert Table 11 about here]

Table 11 presents the results. For firms in tradable industries, the effect of age structure on innovation is similar to our baseline results for all firms (Table 3). By contrast, for firms in non-tradable industries, the effect of age structure on innovation is significantly smaller than for firms in tradable industries. In other words, our results are clearly driven by firms in tradable industries, for which the labor force only captures employees, not customers.[26] On the whole, the results allow us to rule out the pure consumption channel.

8. Commuting Zone-Level Analysis

In our final analysis, we consider the effect of age structure on innovation for the full spectrum of innovative entities. We conduct this analysis at the commuting zone level to allow us both to generalize our results and implications as well as to easily compare our results across innovative entities. Our sample here comprises 2,964 commuting zone-quinquennial period

[25] Using the first measure, industries are classified as non-tradable if they are retail- or restaurant-related, and industries are classified as tradable if they exceed some minimum level of U.S. imports plus exports. According to the first measure, roughly 10% of our sample firms are in non-tradable industries, and about half are in tradable industries. Using the second measure, industries are classified as non-tradable if they are in the top quartile of the geographic dispersion of employment across counties in the U.S. (i.e., employment is most dispersed). Industries in the bottom quartile of dispersion (i.e., employment is most concentrated) are classified as tradable. Non-tradable industries are less innovative than tradable industries, but non-tradable industries still produce a positive amount of innovation outputs.

[26] We also check whether the results in Table 11 are sensitive to comparing industries that are consumer oriented to those that are not. To this end, we use the five industry grouping from the Fama-French industry classification system, and we compare the first group (about 25% of our sample) to the other groups to test whether the effect of age structure on innovation is different for consumer oriented industries. In fact, we find that this effect is significantly smaller for consumer oriented industries (results not tabulated), which similarly does not support a local consumption interpretation of our results.

observations corresponding to 741 unique commuting zones. In our sample, public and private firms respectively produce 45% and 49% of all patents. Government entities produce about 1.5% and universities produce 4%. Individuals produce well under 1%.

[Insert Table 12 about here]

We begin by examining total patent production. We use regression specifications similar to our baseline. However, the unit of observation is the commuting zone-quinquennial period, and we only include commuting zone-level control variables and state-year fixed effects. Panel A of Table 12 shows that the effect of age structure on innovation is economically and statistically significant, consistent with our previous results.

We then break out total commuting zone-level patent production for various innovative entities. Starting with public firms, Panel B shows that the effect is four times larger at the commuting zone level than at the firm level (Table 3). However, we are cautious about such comparisons because at the commuting zone level it is impossible to include all of the control variables and fixed effects that we were able to include at the firm level of analysis. Further results indicate that the age structure of the local labor force affects innovation for private firms (Panel C) and government entities (Panel D) as it does for the commuting zone in its entirety (Panel A). Local universities appear to be unaffected, perhaps because they are relatively self contained ecosystems insofar as their dependence on the labor force of the commuting zone as a whole is concerned.

Finally, still at the commuting zone level, we compare firms based on their age. We are interested in whether age structure affects innovation differently at more entrepreneurial firms compared to more mature firms. Since we do not have data on firm founding dates across the board, we measure firm age from the date of the firm's first patent. We split firms into two groups based on whether they are younger or older than five years of age (which produces a roughly 25%-75% split).

The results in Table 12 show that both groups of firms are affected by age structure, but the effect is roughly double for older versus younger firms (Panel G versus Panel F). A possible explanation for this finding is that the local labor pool is more important for more mature firms because such firms may need it to regularly replenish their human capital. For more entrepreneurial firms, by contrast, a small, stable group of core employees comprises most of their human capital, so the local labor pool is less important to them.

9. Conclusion

We study the effect of labor force age structure on corporate innovation. We argue that because of their creativity, risk tolerance, horizons, and interactivity, younger workers are instrumental in the production of innovation. In a number of different ways, younger labor markets can create a general work environment that is more conducive to innovation. We therefore hypothesize that a younger labor force produces more innovation.

We measure locations using commuting zones and innovation using patent outputs. We instrument the current labor force with historical births in each local labor market in the United States. In analyses at the firm, inventor, and commuting zone levels, we find that a younger age structure causes more innovation. Our regression specifications allow us to rule our results being explained by unobservable heterogeneity across local labor markets and firms, life cycles for firms and inventors, and financing or consumption effects. We also find that corporate innovation activities tend to reflect the innovative characteristics of younger labor forces. Additionally, we find that the market value of innovations is higher for firms in younger labor markets.

Taken as a whole, our findings indicate policy recommendations for the demographic challenges confronting the world today. We find that not only do younger labor forces produce more innovation, they also create more wealth. This suggests that there are potential net benefits to public policies that can counter the effects of an aging labor force: improving education and training, encouraging young and skilled immigration, and incentivizing domestic population growth.

Appendix 1

Ideally, we would like to measure the age structure of the labor force such that our measure reflects the actual location of the firm's entire workforce, inventors and non-inventors alike, in commuting zones across the country. However, we face the practical limitation that there is no widely available data for our samples on the location and size of the firm's workforce in each commuting zone. Even an R&D hub, which is a location at which there are inventors working for the firm, can only be inferred by us, like the literature, based on observing whether our patent database contains at least one inventor who has had at least one patent granted to him during the prior 10 years.

The inference of R&D hubs from patent grants creates numerous measurement errors. First, firms may no longer have future innovation production capacity at locations inferred based on past innovation output. Consider a firm that had a single patent granted 10 years ago and thus would be inferred to have an R&D hub at that location. However, the firm could have closed that R&D hub nine years ago, and therefore it would be impossible for the firm to produce another patent at that location during the next five years.

Second, firms may have future innovation production capacity at locations that could not be inferred from past innovation output. Suppose a firm at a given location opens a new and extremely productive R&D hub at that location in any year during the next five years. Such a firm would retrospectively be inferred to have no R&D hub and therefore no innovation at that location.

Inferring R&D hubs from prospective patent grants in our patent database during the next five years is not done in the literature and with good reason. Such inference would mechanically classify locations as R&D hubs if they are successful at producing innovation in the future and similarly mechanically overlook locations of R&D hubs that fail to produce future innovation. This classification scheme could mechanically induce a correlation between current age structure (and also other location characteristics) and future innovation that would undermine our careful identification based on historical births.

Third, allowing a weighted average of different locations for a given firm would make it difficult to deal with the threat to our identification of omitted factors. Starting from our baseline specification onwards, we would like to sweep out variation across states and time using state-year fixed effects. However, this is not something we could do for inferred R&D hubs weighted by the number of inventors. Similarly, we would like to be able to sweep out time-invariant differences across locations by including commuting zone fixed effects in our regression specifications. However, the interpretation of these fixed effects would be unclear because the "location" they would capture would be the time-invariant component of a weighted average of locations such that both the weights (i.e., the number of employees) and the locations (i.e., of these employees) change over time. Moreover, in order to include time-invariant fixed effects, we need a minimal number of observations for each firm across time. For this reason, in our corresponding analysis, we need to add two quinquennial periods to our sample by starting in 1980 rather than 1990 (at the expense of having to measure the labor force using 20-54 year olds rather than 20-64 year olds). This is problematic using R&D hubs because they are inferred by looking backward at patents granted during the prior 10 years. Since the patent data only start in 1976, we would not be able to include two additional quinquennial periods and to start our sample in 1980.

The measurement errors resulting from inferring R&D hubs are increasingly severe for younger and smaller firms. Such firms, by definition, have a shorter history during which we can potentially observe patent grants. Similarly, the smaller scale of these firms means that they have fewer patent grants from which we must necessarily infer, with greater noise, the times at which they opened or closed R&D hubs. We should avoid using a more noisy measure of age structure for the very types of firms that might have more difficulty adjusting to the characteristics of the local labor force.

In contrast to inferred R&D hubs, we can measure a firm's headquarters location precisely. Furthermore, for firms with at least one patent, we find that roughly half of a firm's innovation is produced in the commuting zone in which the firm is headquartered. Accordingly, a firm's headquarters is, on average, its single most productive inferred R&D hub, consistent with the stylized fact that firms tend to locate their R&D hubs close to their headquarters. Moreover, it appears that a firm's headquarters location is representative of its inferred R&D hubs in terms of commuting zone characteristics. Specifically, again for firms with at least one patent, we find that the age structure and patent output of the commuting zone are economically indistinguishable when we compare headquarters and inferred R&D hubs. For all of these reasons, we use a firm's headquarters as its location in our baseline analysis.

Appendix 2

We examine how various intuitive factors moderate the effect of age structure on innovation. Our general motivation is that the labor force should naturally be a more important driver of innovation activities in the following circumstances: for businesses that themselves produce more innovation, when the local labor force itself has higher skill levels, and when the businesses in the local labor market can take advantage of innovation spillovers from local universities. To this end, we create contrasts at the industry, commuting zone, and firm levels.

First, we sort firms into innovative and non-innovative industries, and we perform simple variations on the regressions in Table 3. The specifications include the innovative industry dummy variable and its interaction with age structure and all control variables. Firms in innovative industries are the base group. Industries are classified as innovative if their three-digit SIC code is one of the following: 283 (drugs), 357 (computer hardware), 366 (communications equipment), 367 (electronics), 381 (navigation equipment), 382 (measuring and controlling instruments), 384 (medical instruments), 481 (telephone equipment), 489 (communications services), 737 (software), 873 (research services), and 874 (management services). Appendix Table 6 shows that our results are stronger for innovative industries than non-innovative industries (Panel A).

Second, we take the same approach to local labor market skill, which is measured by educational attainment. We classify commuting zones as more skilled and less skilled based on whether their educational attainment is above or below the median. Appendix Table 6 shows that our results are stronger for more skilled labor markets than less skilled ones (Panel B).

Finally, we take a similar approach to the innovativeness of local universities as measured by their patent counts. We use the top 15 commuting zones, which roughly splits our sample of firm-years into two halves, but the results are robust to using the top 10 or 20 commuting zones. The following commuting zones are in the top 15 for at least three quinquennial periods: Baltimore, Boston, Chicago, Detroit, Houston, Los Angeles, New York, Philadelphia, Raleigh-Durham, San Francisco, San Jose, and Washington, D.C. Appendix Table 6 shows that our results are stronger for more innovative labor markets than less innovative ones (Panel C).

References

Acemoglu, Daron, and Pascual Restrepo, 2021, Demographics and automation, forthcoming *Review of Economic Studies*.

Acemoglu, Daron, Ufuk Akcigit, and Murat Alp Celik, 2020, Radical and incremental innovation: The roles of firms, managers and innovators, forthcoming *American Economic Journal: Macroeconomics*.

Acharya, Viral A., Ramin P. Baghai, and Krishnamurthy V. Subramanian, 2013, Labor laws and innovation, *Journal of Law and Economics* 56, 997-1037.

Acharya, Viral A., Ramin P. Baghai, and Krishnamurthy V. Subramanian, 2014, Wrongful discharge laws and innovation, *Review of Financial Studies* 27, 301-346.

Adelino, Manuel, Song Ma, and David Robinson, 2017, Firm age, investment opportunities, and job creation, *Journal of Finance* 72, 999-1038.

Agrawal, Ashwini K., and David A. Matsa, 2013, Labor unemployment risk and corporate financing decisions, *Journal of Financial Economics* 108, 449-470.

Akcigit, Ufuk, Salomé Baslandze, and Stefanie Stantcheva, 2016, Taxation and the international mobility of inventors, *American Economic Review* 106, 2930-2981.

Aksoy, Yunus, Henrique S. Basso, Ron P. Smith, and Tobias Grasl, 2019, Demographic structure and macroeconomic trends, *American Economic Journal: Macroeconomics* 11, 193-222.

Anelli, Massimo, Gaetano Basso, Giuseppe Ippedico, and Giovanni Peri, 2019, Youth drain, entrepreneurship and innovation, working paper.

Audretsch, David B., and Maryann P. Feldman, 1996, R&D spillovers and the geography of innovation and production, *American Economic Review* 86, 630-640.

Autor, David H., and David Dorn, 2013, The growth of low-skill service jobs and the polarization of the US labor market, *American Economic Review* 103, 1553-1597.

Bai, John, Douglas Fairhurst, and Matthew Serfling, 2020, Employment protection, investment, and firm growth, *Review of Financial Studies* 33, 644-688.

Becker, Bo, 2007, Geographical segmentation of US capital markets, *Journal of Financial Economics* 85, 151-178.

Becker, Bo, Zoran Ivkovic, and Scott Weisbenner, 2011, Local dividend clienteles, *Journal of Finance* 66, 655-683.

Betermier, Sebastien, Laurent E. Calvet, and Paolo Sodini, 2017, Who are the value and growth investors?, *Journal of Finance* 72, 5-46.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, 2013, Identifying technology spillovers and product market rivalry, *Econometrica* 81, 1347-1393.

Bowen, Donald E. III, Laurent Frésard, and Jérôme P. Taillard, 2017, What's your identification strategy? Innovation in corporate finance research, *Management Science* 63, 2529-2548.

Coval, Joshua D., and Tobias J. Moskowitz, 1999, Home bias at home: Local equity preference in domestic portfolios, *Journal of Finance* 54, 2045-2073.

DellaVigna, Stefano, and Joshua M. Pollet, 2007, Demographics and industry returns, *American Economic Review* 97, 1167-1702.

DellaVigna, Stefano, and Joshua M. Pollet, 2013, Capital budgeting versus market timing: An evaluation using demographics, *Journal of Finance* 68, 237-270.

Dessaint, Olivier, Andrey Golubov, and Paolo Volpin, 2017, Employment protection and takeovers, *Journal of Financial Economics* 125, 369-388.

Dougal, Casey, Christopher A. Parsons, and Sheridan Titman, 2015, Urban vibrancy and corporate growth, *Journal of Finance* 70, 163-210.

Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso, 2017, Asset market participation and portfolio choice over the life cycle, *Journal of Finance* 72, 705-750.

Field, Laura C., and Jonathan Karpoff, 2002, Takeover defenses of IPO firms, *Journal of Finance* 57, 1857-1889.

Fishback, Price, Werner Troesken, Trevor Kollmann, Michael Haines, Paul Rhode, and Melissa Thomasson, 2011, Information and the impact of climate and weather on mortality rates during the Great Depression, in Gary Libecap and Richard Steckel, eds. *The Economics of Climate Change* (University of Chicago Press).

Glaeser, Edward L., Hedi D. Kallal, José A. Scheinkman, and Andrei Shleifer, 1992, Growth in cities, *Journal of Political Economy* 100, 1126-1152.

Goyal, Amit, 2004, Demographics, stock market flows, and stock returns, *Journal of Financial and Quantitative Analysis* 39, 115-142.

Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg, 2005, Market value and patent citations, *RAND Journal of Economics* 36, 16-38.

Hombert, Johan, Antoinette Schoar, David Sraer, and David Thesmar, 2020, Can unemployment insurance spur entrepreneurial activity? Evidence from France, *Journal of Finance* 75, 1247-1285.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson, 1993, Geographic localization of knowledge spillovers as evidenced by patent citations, *Quarterly Journal of Economics* 108, 577-598.

Jaimovich, Nir, and Henry E. Siu, 2009, The young, the old, and the restless: Demographics and business cycle volatility, *American Economic Review* 99, 804-826.

Jeffers, Jessica, 2019, The impact of restricting labor mobility on corporate investment and entrepreneurship, working paper.

John, Kose, Anzhela Knyazeva, and Diana Knyazeva, 2015, Employee rights and acquisitions, *Journal of Financial Economics* 118, 49-69.

Jones, Benjamin F., 2010, Age and great invention, *Review of Economics and Statistics* 92, 1-14.

Jones, Benjamin F., and Bruce A. Weinberg, 2011, Age dynamics in scientific creativity, *Proceedings of the National Academy of Sciences* 108, 18910-18914.

Jovanovic, Boyan, and Peter L. Rousseau, 2001, Why wait? A century of life before IPO, *American Economic Review* 91, 336-341.

Kaltenberg, Mary, Adam B. Jaffe, and Margie E. Lachman, 2021, The age of invention: Matching inventor ages to patents based on web-scraped sources, working paper.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, Technological innovation, resource allocation, and growth, *Quarterly Journal of Economics* 132, 665-712.

Lerner, Josh, and Amit Seru, 2017, The use and misuse of patent data: Issues for corporate finance and beyond, working paper.

Li, Guan-Cheng, Ronald Lai, Alexander D'Amour, David M. Doolin, Ye Sun, Vetle I. Torvik, Amy Z. Yu, and Lee Fleming, 2014, Disambiguation and co-authorship networks of the U.S. patent inventor database (1975-2010), *Research Policy* 43, 941-955.

Liang, James, Hui Wang, and Edward P. Lazear, 2018, Demographics and entrepreneurship, *Journal of Political Economy* 126, S140-S196.

Loughran, Tim, and Jay R. Ritter, 2004, Why has IPO underpricing changed over time?, *Financial Management* 33, 5-37.

Lugauer, Steven, 2012, Estimating the effect of the age distribution on cyclical output volatility across the United States, *Review of Economics and Statistics* 94, 896-902.

Lychagin, Sergey, Joris Pinkse, Margaret Slade, and John Van Reenen, 2016, Spillovers in space: Does geography matter?, *Journal of Industrial Economics* 64, 295-335.

Matsa, David A., 2010, Capital structure as a strategic variable: Evidence from collective bargaining, *Journal of Finance* 65, 1197-1232.

Mian, Atif, and Amir Sufi, 2014, What explains the 2007-2009 drop in employment?, *Econometrica* 82, 2197-2223.

Moretti, Enrico, and Daniel J. Wilson, 2017, The effect of state taxes on the geographical location of top earners: Evidence from star scientists, *American Economic Review* 107, 1858-1903.

Ouimet, Paige, and Rebecca Zarutskie, 2014, Who works for startups? The relation between firm age, employee age, and growth, *Journal of Financial Economics* 112, 386-407.

Poterba, James M., 2001, Demographics structure and asset returns, *Review of Economics and Statistics* 83, 565-584.

Ruggles, Steven, Sarah Flood, Sophia Foster, Ronald Goeken, Jose Pacas, Megan Schouweiler, and Matthew Sobek. IPUMS USA: Version 11.0 [dataset]. Minneapolis, MN: IPUMS, 2021. https://doi.org/10.18128/D010.V11.0

Shimer, Robert, 2001, The impact of young workers on the aggregate labor market, *Quarterly Journal of Economics* 116, 969-1007.

Simintzi, Elena, Vikrant Vig, and Paolo Volpin, 2015, Labor protection and leverage, *Review of Financial Studies* 28, 561-591.

Tolbert, Charles M., and Molly Sizer, 1996, U.S. commuting zones and labor market areas: A 1990 update, working paper.

Table 1
Descriptive Statistics

This table presents descriptive statistics for the main variables used in this paper. The samples in the panels are described in the text. Variables are defined in Appendix Table 1.

Panel A: Firm-Level Sample

	Mean	Standard deviation	25th percentile	Median	75th percentile
Independent variables					
- Births-based young share (%)	54.7	8.4	48.5	54.6	59.8
- Births-based mean age (years)	38.8	2.2	37.3	38.9	40.4
- Actual young share (%)	52.3	4.7	49.1	52.6	55.4
- Actual mean age (years)	39.6	1.1	38.9	39.5	40.3
Dependent variables					
- Patent counts	4.8	19.6	0.0	0.0	1.0
- Patent citations	6.0	24.1	0.0	0.0	1.1
- Citations per patent	1.3	1.1	0.6	1.0	1.6
- Proportion of extremely useful patents (%)	1.2	4.7	0.0	0.0	0.0
- Proportion of useful and novel patents (%)	1.5	5.0	0.0	0.0	0.0
- Volatility of citations per patent	1.1	0.9	0.5	0.9	1.5
- Longevity of citations per patent	1.0	0.2	0.9	1.0	1.1
- Proportion of local citations	2.5	5.1	0.0	0.9	2.3
- Patent value ($ millions)	288.4	1,259.4	1.1	4.7	31.1
Firm-level control variables					
- Total assets ($ millions)	1,076	3,362	29	113	509
- Market-to-book	3.3	4.4	1.1	2.0	3.6
- Cash flow-to-total assets (%)	5.1	23.4	1.7	11.0	17.2
- Stock returns (%)	4.0	67.1	-31.0	3.6	38.6
- Stock return volatility (%)	68.7	40.8	38.4	57.5	88.6
Commuting zone-level control variables					
- Population size (thousands)	4,379	4,061	1,616	3,050	5,244
- Income per capita ($ thousands)	30.4	8.8	23.4	29.2	36.7
- Government spending-to-total income (%)	8.5	1.9	7.2	8.3	9.5
- Educational attainment (%)	26.2	6.1	22.0	26.0	30.1
- University students per capita (%)	6.7	1.4	5.8	6.6	7.7
- University patent counts per capita	1.4	1.5	0.5	1.0	1.7

Panel B: Inventor-Level Sample					
	Mean	Standard deviation	25th percentile	Median	75th percentile
Independent variables					
- Births-based young share (%)	52.1	7.7	45.4	52.4	58.1
- Births-based mean age (years)	39.4	2.0	38.0	39.4	41.1
- Actual young share (%)	50.1	4.5	46.5	49.9	53.3
- Actual mean age (years)	40.1	1.1	39.3	40.1	40.9
Dependent variables					
- Patent counts	0.87	1.08	0.23	0.51	1.05
- Patent citations	1.08	1.88	0.14	0.42	1.13
Inventor-level control variables					
- Inventor age (years)	48.5	9.7	42.0	48.0	55.0
- Inventor experience (years)	15.5	6.1	11.0	15.0	19.0
- Inventor patent stock	39	26	23	31	46
- Inferred R&D hub number of inventors	695	938	55	251	973
Commuting zone-level control variables					
- Population size (thousands)	3,525	3,393	1,169	2,398	4,986
- Income per capita ($ thousands)	37.5	7.6	31.9	36.4	41.5
- Government spending-to-total income (%)	8.6	1.9	7.1	8.4	9.9
- Educational attainment (%)	26.9	6.2	23.0	26.8	31.3
- University students per capita (%)	6.6	1.3	5.8	6.4	7.5
- University patent counts per capita	1.8	1.8	0.6	1.1	1.9

Table 2
The Effect of Births-Based Age Structure on Actual Age Structure

This table shows the results of regressions of actual age structure on births-based age structure. The regressions correspond to the first stage of Equation 1. The unit of observation is the firm-quinquennial period. Independent and dependent variables are generally measured contemporaneously at the beginning of every adjacent and non-overlapping quinquennial period starting in 1990. The sample and specifications are described in the text. The actual labor force in a given year is self-explanatory. The births-based labor force in a given year is the labor force constructed based on the number of births each year during the prior 20-64 years. In Columns 1 to 4 and 7 to 8, the sample comprises four quinquennial periods starting in 1990, and age structure is measured using the population of 20-64 year olds. In Columns 5 and 6, the sample comprises six quinquennial periods starting in 1980, and age structure is measured using the population of 20-54 year olds. Column 4 corresponds to Table 3, and Columns 5 to 8 correspond to Table 5 Panels A to D, respectively. Variables are defined in Appendix Table 1. The partial R^2 pertains to the births-based age structure variable. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Age Structure Measured Using Young Share								
	Dependent variable is actual age structure							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Births-based age structure	0.356***	0.154***	0.096***	0.096***	0.102***	0.110***	0.099***	0.100***
	(23.72)	(19.45)	(12.63)	(13.15)	(14.17)	(11.49)	(15.02)	(14.38)
Firm-level control variables?	No	No	No	Yes	Yes	Yes	Yes	Yes
Contemporaneous CZ-level control variables?	No	No	Yes	Yes	Yes	Yes	Yes	Yes
State-year fixed effects?	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-year fixed effects?	No	No	No	Yes	Yes	Yes	Yes	Yes
Firm age fixed effects?	No	No	No	Yes	Yes	Yes	Yes	Yes
CZ fixed effects?	No	No	No	No	Yes	No	No	No
Firm fixed effects?	No	No	No	No	No	Yes	No	No
Historical CZ-level control variables?	No	No	No	No	No	No	Yes	Yes
Historical CZ-level control variables × Quinquennial period fixed effects?	No	No	No	No	No	No	No	Yes
Observations	15,333	15,327	15,327	14,087	19,743	15,867	14,087	14,087
Adjusted R^2	0.432	0.873	0.928	0.929	0.990	0.991	0.951	0.957
Partial R^2	0.432	0.108	0.069	0.065	0.038	0.044	0.079	0.081

Panel B: Age Structure Measured Using Mean Age								
	Dependent variable is actual age structure							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Births-based age structure	0.283***	0.145***	0.097***	0.097***	0.070***	0.082***	0.105***	0.108***
	(21.79)	(16.30)	(11.27)	(11.59)	(9.66)	(8.62)	(14.17)	(13.93)
Firm-level control variables?	No	No	No	Yes	Yes	Yes	Yes	Yes
Contemporaneous CZ-level control variables?	No	No	Yes	Yes	Yes	Yes	Yes	Yes
State-year fixed effects?	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-year fixed effects?	No	No	No	Yes	Yes	Yes	Yes	Yes
Firm age fixed effects?	No	No	No	Yes	Yes	Yes	Yes	Yes
CZ fixed effects?	No	No	No	No	Yes	No	No	No
Firm fixed effects?	No	No	No	No	No	Yes	No	No
Historical CZ-level control variables?	No	No	No	No	No	No	Yes	Yes
Historical CZ-level control variables × Quinquennial period fixed effects?	No	No	No	No	No	No	No	Yes
Observations	15,333	15,327	15,327	14,087	19,743	15,867	14,087	14,087
Adjusted R^2	0.316	0.815	0.901	0.903	0.992	0.993	0.933	0.940
Partial R^2	0.316	0.083	0.063	0.060	0.022	0.028	0.082	0.086

Table 3
The Effect of Age Structure on Innovation: Baseline Firm-Level Analysis

This table shows the results of regressions of innovation on age structure. The regressions correspond to Equation 1. The unit of observation is the firm-quinquennial period. Observations temporally are adjacent and non-overlapping. Independent variables are generally measured starting in 1990, and dependent variables are measured starting during 1991-1995. The sample and specifications are described in the text. The actual age structure is instrumented with the births-based age structure. The actual labor force in a given year is self-explanatory. The births-based labor force in a given year is the labor force constructed based on the number of births each year during the prior 20-64 years. Variables are defined in Appendix Table 1. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	10.352***	11.903***	-0.370***	-0.429***
	(4.04)	(4.47)	(-3.96)	(-4.38)
Firm-level control variables				
ln(Total assets)	0.253***	0.264***	0.253***	0.264***
	(11.33)	(11.53)	(11.32)	(11.53)
Market-to-book	0.022***	0.026***	0.022***	0.026***
	(7.50)	(7.82)	(7.48)	(7.80)
Cash flow-to-total assets	-0.213***	-0.233***	-0.211***	-0.230***
	(-4.17)	(-4.39)	(-4.14)	(-4.36)
Stock returns	0.060***	0.069***	0.058***	0.067***
	(3.84)	(4.14)	(3.74)	(4.02)
Stock return volatility	0.113***	0.103***	0.115***	0.105***
	(3.21)	(2.70)	(3.25)	(2.74)
CZ-level control variables				
ln(Population size)	-0.202***	-0.236***	-0.191***	-0.225***
	(-8.37)	(-8.83)	(-8.37)	(-8.78)
ln(Income per capita)	1.017***	1.230***	1.140***	1.382***
	(5.73)	(6.26)	(5.44)	(5.97)
Government spending-to-total income	5.646***	6.103***	5.757***	6.272***
	(4.43)	(4.17)	(4.40)	(4.17)
Educational attainment	-0.725	-0.913	-0.685	-0.891
	(-1.32)	(-1.56)	(-1.22)	(-1.47)
University students per capita	-13.203***	-15.041***	-14.754***	-16.962***
	(-4.45)	(-4.48)	(-4.36)	(-4.42)
ln(1+University patent counts per capita)	0.205***	0.214***	0.191***	0.199***
	(4.89)	(4.40)	(4.81)	(4.27)
Fixed effects				
State-year	Yes	Yes	Yes	Yes
Industry-year	Yes	Yes	Yes	Yes
Firm age	Yes	Yes	Yes	Yes
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	172.8	172.8	134.2	134.2

Table 4
Replication of Baseline Firm-Level Results using Uninstrumented Actual Age Structure

This table shows the results of regressions of innovation on age structure. The regressions are the same as in Table 3, but the actual age structure is not instrumented with the births-based age structure. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	1.734***	2.408***	-0.053***	-0.076***
	(3.06)	(3.88)	(-2.70)	(-3.59)
Firm-level control variables				
ln(Total assets)	0.253***	0.264***	0.253***	0.264***
	(11.25)	(11.44)	(11.24)	(11.43)
Market-to-book	0.023***	0.026***	0.023***	0.027***
	(7.45)	(7.78)	(7.45)	(7.77)
Cash flow-to-total assets	-0.213***	-0.232***	-0.212***	-0.232***
	(-4.15)	(-4.41)	(-4.14)	(-4.40)
Stock returns	0.058***	0.067***	0.058***	0.067***
	(3.77)	(4.02)	(3.75)	(4.00)
Stock return volatility	0.108***	0.098**	0.108***	0.098**
	(3.08)	(2.55)	(3.07)	(2.56)
CZ-level control variables				
ln(Population size)	-0.125***	-0.151***	-0.121***	-0.147***
	(-8.38)	(-8.63)	(-8.40)	(-8.61)
ln(Income per capita)	0.402***	0.553***	0.401***	0.559***
	(3.33)	(3.97)	(3.30)	(3.99)
Government spending-to-total income	2.642***	2.793***	2.566***	2.722***
	(2.83)	(2.74)	(2.78)	(2.70)
Educational attainment	0.944**	0.925**	1.001**	0.986**
	(2.40)	(1.97)	(2.53)	(2.08)
University students per capita	-4.483***	-5.433***	-4.436***	-5.480***
	(-3.81)	(-4.06)	(-3.69)	(-4.02)
ln(1+University patent counts per capita)	0.138***	0.140***	0.134***	0.135***
	(4.50)	(3.84)	(4.43)	(3.74)
Fixed effects				
State-year	Yes	Yes	Yes	Yes
Industry-year	Yes	Yes	Yes	Yes
Firm age	Yes	Yes	Yes	Yes
Observations	14,087	14,087	14,087	14,087
Adjusted R^2	0.207	0.194	0.207	0.193

Table 5
Replication of Baseline Firm-Level Results Accounting for Additional Factors

This table shows the results of regressions of innovation on age structure. The regressions are the same as in Table 3 but with slight modifications, as indicated. In Panels A and B, the baseline four quinquennial periods starting in 1990 are replaced with six quinquennial periods starting in 1980. Also in these two panels, age structure is measured using the population of 20-54 year olds rather than the population of 20-64 year olds. Panel A of this table adds commuting zone fixed effects to Table 3. Panels B adds firm fixed effects to Table 3. In Panel C, historical commuting zone-level control variables are added to Table 3. In Panel D, these control variables are interacted with quinquennial period fixed effects. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Adding Commuting Zone Fixed Effects

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	12.977***	9.332*	-0.694**	-0.432
	(2.96)	(1.96)	(-2.52)	(-1.40)
Firm-level control variables?	Yes	Yes	Yes	Yes
Contemporaneous CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
CZ fixed effects?	Yes	Yes	Yes	Yes
Observations	19,743	19,743	19,743	19,743
F-statistic for instrument	200.8	200.8	93.28	93.28

Panel B: Adding Firm Fixed Effects

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	15.207***	11.822**	-0.671**	-0.478*
	(3.04)	(2.35)	(-2.40)	(-1.65)
Firm-level control variables?	Yes	Yes	Yes	Yes
Contemporaneous CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Firm fixed effects?	Yes	Yes	Yes	Yes
Observations	15,867	15,867	15,867	15,867
F-statistic for instrument	131.9	131.9	74.38	74.38

Panel C: Adding Historical Commuting Zone-Level Control Variables				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	9.781***	11.730***	-0.335***	-0.406***
	(4.48)	(5.15)	(-4.41)	(-5.12)
Firm-level control variables?	Yes	Yes	Yes	Yes
Contemporaneous CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Historical CZ-level control variables?	Yes	Yes	Yes	Yes
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	225.5	225.5	200.8	200.8

Panel D: Adding Historical Commuting Zone-Level Control Variables Interacted with Quinquennial Period Fixed Effects				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	10.519***	12.794***	-0.348***	-0.428***
	(4.42)	(5.02)	(-4.38)	(-4.96)
Firm-level control variables?	Yes	Yes	Yes	Yes
Contemporaneous CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Historical CZ-level control variables × Quinquennial period fixed effects?	Yes	Yes	Yes	Yes
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	206.8	206.8	194.0	194.0

Table 6
Replication of Baseline Firm-Level Results with Alternative Definitions of Firm Location

This table shows the results of regressions of innovation on age structure. The regressions are the same as in Table 3 but with slight modifications, as indicated. In Panel A, the left hand side of the regression equation is modified such that only innovation produced at headquarters is measured. In Panel B, the right hand side of the regression equation is modified such that age structure and the commuting zone-level control variables are measured across the inferred R&D hubs of the firms and weighted by the number of inventors at each R&D hub. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Measuring Innovation Produced Only at Headquarters

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	9.319***	10.899***	-0.336***	-0.399***
	(3.53)	(3.70)	(-3.50)	(-3.67)
Firm-level control variables?	Yes	Yes	Yes	Yes
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	172.8	172.8	134.2	134.2

Panel B: Measuring Firm Location as the Weighted Average of Inferred R&D Hubs

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	14.161***	15.714***	-0.568***	-0.631***
	(4.64)	(4.72)	(-4.37)	(-4.43)
Firm-level control variables?	Yes	Yes	Yes	Yes
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	150.2	150.2	102.5	102.5

Table 7
The Effect of Age Structure on Innovation: Baseline Inventor-Level Analysis

This table shows the results of regressions of innovation on age structure. The regressions correspond to Equation 2. The unit of observation is the inventor-quinquennial period. Observations temporally are adjacent and non-overlapping. Independent variables are generally measured starting in 1990, and dependent variables are measured starting during 1991-1995. The sample and specifications are described in the text. The actual age structure is instrumented with the births-based age structure. The actual labor force in a given year is self-explanatory. The births-based labor force in a given year is the labor force constructed based on the number of births each year during the prior 20-64 years. Variables are defined in Appendix Table 1. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is ln(0.01+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	7.336**	9.413**	-0.244**	-0.323**
	(2.42)	(2.18)	(-2.34)	(-2.15)
Inventor-level control variables				
ln(Inventor age)	-0.053	-0.197**	-0.053	-0.196**
	(-0.81)	(-2.48)	(-0.80)	(-2.47)
ln(Inventor experience)	-0.418***	-0.408***	-0.418***	-0.408***
	(-10.58)	(-7.56)	(-10.61)	(-7.56)
ln(Inventor patent stock)	0.748***	0.871***	0.747***	0.870***
	(23.94)	(17.88)	(23.96)	(17.88)
ln(Inferred R&D hub number of inventors)	0.045***	0.034**	0.045***	0.035**
	(4.11)	(2.48)	(4.14)	(2.50)
CZ-level control variables				
ln(Population size)	-0.069**	-0.114**	-0.057*	-0.100*
	(-1.98)	(-2.03)	(-1.78)	(-1.95)
ln(Income per capita)	0.584*	1.126**	0.675*	1.268**
	(1.68)	(2.30)	(1.79)	(2.37)
Government spending-to-total income	0.629	2.775	0.811	3.078
	(0.31)	(0.97)	(0.40)	(1.06)
Educational attainment	-1.425	-2.829*	-1.368	-2.833*
	(-1.13)	(-1.69)	(-1.08)	(-1.68)
University students per capita	-7.660*	-4.911	-7.860*	-5.530
	(-1.72)	(-0.80)	(-1.70)	(-0.86)
ln(1+University patent counts per capita)	-0.035	0.023	-0.042	0.018
	(-0.44)	(0.21)	(-0.53)	(0.16)
Fixed effects				
State-year	Yes	Yes	Yes	Yes
Technology class-year	Yes	Yes	Yes	Yes
Firm-year	Yes	Yes	Yes	Yes
Observations	11,975	11,975	11,975	11,975
F-statistic for instrument	229.5	229.5	243.4	243.4

Table 8
The Effect of Age Structure on Innovation Characteristics: Firm-Level Analysis

This table shows the results of regressions of innovation characteristics on age structure. The regressions correspond to Equation 1. The unit of observation is the firm-quinquennial period. Observations temporally are adjacent and non-overlapping. Independent variables are generally measured starting in 1990, and dependent variables are measured starting during 1991-1995. The sample and specifications are described in the text. The actual age structure is instrumented with the births-based age structure. The actual labor force in a given year is self-explanatory. The births-based labor force in a given year is the labor force constructed based on the number of births each year during the prior 20-64 years. Variables are defined in Appendix Table 1. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Age Structure Measured Using Young Share						
	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	3.070	9.443**	12.893***	5.345*	0.488	9.681**
	(1.36)	(2.46)	(2.92)	(1.82)	(0.63)	(2.38)
Firm-level control variables?	Yes	Yes	Yes	Yes	Yes	Yes
CZ-level control variables?	Yes	Yes	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	230.8	230.8	230.8	207.2	230.9	230.8

Panel B: Age Structure Measured Using Mean Age						
	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	-0.117	-0.341**	-0.486***	-0.207*	-0.022	-0.361**
	(-1.40)	(-2.44)	(-2.96)	(-1.85)	(-0.74)	(-2.47)
Firm-level control variables?	Yes	Yes	Yes	Yes	Yes	Yes
CZ-level control variables?	Yes	Yes	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	205.3	205.3	205.3	180.1	204.6	205.3

Table 9
The Effect of Age Structure on the Market Value of Innovation: Firm-Level Analysis

This table shows the results of regressions of the market value of innovation on age structure. The unit of observation is the firm-quinquennial period. Observations temporally are adjacent and non-overlapping. Independent variables are generally measured starting in 1990, and dependent variables are measured starting during 1991-1995. The sample and specifications are described in the text. The actual age structure is instrumented with the births-based age structure. The actual labor force in a given year is self-explanatory. The births-based labor force in a given year is the labor force constructed based on the number of births each year during the prior 20-64 years. Variables are defined in Appendix Table 1. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is ln(value of mean patents per annum)	
	Age structure = Young share	Age structure = Mean age
Age structure	10.217***	-0.384***
	(2.64)	(-2.69)
Firm-level control variables?	Yes	Yes
CZ-level control variables?	Yes	Yes
State-year fixed effects?	Yes	Yes
Industry-year fixed effects?	Yes	Yes
Firm age fixed effects?	Yes	Yes
Observations	5,607	5,607
F-statistic for instrument	218.5	196.0

Table 10
The Effect of Age Structure on Innovation: Firm-Level Analysis Testing the Financing Channel by Comparing Age Structure at Headquarters and Inferred R&D Hubs

This table shows the results of regressions of innovation on age structure. The regressions correspond to Equation 1. The regressions are the same as in Table 3, but the sample is restricted to firms with inferred R&D hubs. The actual age structure is instrumented with the births-based age structure. The actual labor force in a given year is self-explanatory. The births-based labor force in a given year is the labor force constructed based on the number of births each year during the prior 20-64 years. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Only Age Structure at Headquarters

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure at headquarters	6.157***	8.153***	-0.218**	-0.296***
	(2.69)	(2.77)	(-2.55)	(-2.66)
Firm-level control variables?	Yes	Yes	Yes	Yes
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	6,259	6,259	6,259	6,259
F-statistic for instrument	227.7	227.7	190.0	190.0

Panel B: Age Structures at Both Headquarters and Inferred R&D Hubs

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure at headquarters	4.746**	6.051*	-0.164*	-0.214*
	(1.98)	(1.95)	(-1.83)	(-1.84)
Age structure at inferred R&D hubs	4.150**	6.183***	-0.156**	-0.237***
	(2.20)	(3.05)	(-2.09)	(-2.99)
Firm-level control variables?	Yes	Yes	Yes	Yes
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	6,259	6,259	6,259	6,259
F-statistic for instrument	113.3	113.3	92.93	92.93

Table 11
The Effect of Age Structure on Innovation: Firm-Level Analysis Testing the Consumption Channel by Comparing Tradable and Non-Tradable Industries

This table shows the results of regressions of innovation on age structure for firms in non-tradable industries compared to firms in tradable industries. The regressions correspond to Equation 1. The regressions are the same as in Table 3, but firms are classified as being in either tradable or non-tradable industries, and the specifications include the non-tradable firm dummy variable and its interaction with age structure and all control variables. In Panel A, industries are classified as non-tradable if they are retail- or restaurant-related, and as tradable if they exceed a minimum level of U.S. imports plus exports. In Panel B, industries are classified as non-tradable if they are in the top quartile of the geographic dispersion of employment, and as tradable if they are in the bottom quartile of dispersion. The actual age structure is instrumented with the births-based age structure. The actual labor force in a given year is self-explanatory. The births-based labor force in a given year is the labor force constructed based on the number of births each year during the prior 20-64 years. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Measure #1 of Tradable versus Non-Tradable Industry

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	8.637***	10.021***	-0.310***	-0.365***
	(2.78)	(2.98)	(-2.66)	(-2.83)
Age structure × Non-tradable industry dummy variable	-7.316**	-9.757***	0.353**	0.473***
	(-2.51)	(-3.18)	(2.44)	(3.10)
Firm-level control variables?	Yes	Yes	Yes	Yes
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	9,083	9,083	9,083	9,083
F-statistic for instrument	62.36	62.36	52.28	52.28

Panel B: Measure #2 of Tradable versus Non-Tradable Industry

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	11.738***	11.489***	-0.434**	-0.423**
	(2.64)	(2.72)	(-2.38)	(-2.43)
Age structure × Non-tradable industry dummy variable	-8.052**	-8.804***	0.368**	0.398**
	(-2.45)	(-2.76)	(2.12)	(2.36)
Firm-level control variables?	Yes	Yes	Yes	Yes
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	7,581	7,581	7,581	7,581
F-statistic for instrument	37.14	37.14	43.53	43.53

Table 12
The Effect of Age Structure on Innovation: Commuting Zone-Level Analysis

This table shows the results of regressions of innovation on age structure. The regressions are similar to those in Table 3, but the unit of observation is the commuting zone-quinquennial period, and only commuting zone-level control variables and state-year fixed effects are included. In Panels F and G, firms are split into two groups based on whether they are younger or older than five years of age. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Total Innovation Produced in the Commuting Zone				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	31.410***	32.424***	-1.084***	-1.135***
	(3.45)	(3.45)	(-3.41)	(-3.42)
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Observations	2,740	2,740	2,740	2,740
F-statistic for instrument	16.46	16.46	15.64	15.64

Panel B: Innovation Produced in the Commuting Zone by Publicly Trade Firms				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	44.168***	43.679***	-1.533***	-1.535***
	(3.53)	(3.60)	(-3.48)	(-3.57)
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Observations	2,740	2,740	2,740	2,740
F-statistic for instrument	16.46	16.46	15.64	15.64

Panel C: Innovation Produced in the Commuting Zone by Private Firms				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	21.655***	25.090***	-0.746***	-0.879***
	(3.16)	(3.20)	(-3.10)	(-3.14)
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Observations	2,740	2,740	2,740	2,740
F-statistic for instrument	16.46	16.46	15.64	15.64

Panel D: Innovation Produced in the Commuting Zone by Government Entities				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	24.256***	23.460***	-0.885***	-0.852***
	(3.06)	(3.09)	(-3.15)	(-3.18)
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Observations	2,740	2,740	2,740	2,740
F-statistic for instrument	16.46	16.46	15.64	15.64

Panel E: Innovation Produced in the Commuting Zone by Universities				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	-2.593	-5.584	0.113	0.209
	(-0.68)	(-1.21)	(0.87)	(1.30)
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Observations	2,740	2,740	2,740	2,740
F-statistic for instrument	18.29	18.29	17.19	17.19

Panel F: Innovation Produced in the Commuting Zone by Young Firms				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	21.719***	23.382***	-0.747***	-0.815***
	(3.18)	(3.01)	(-3.13)	(-2.95)
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Observations	2,740	2,740	2,740	2,740
F-statistic for instrument	16.46	16.46	15.64	15.64

Panel G: Innovation Produced in the Commuting Zone by Old Firms				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	38.008***	39.752***	-1.309***	-1.382***
	(3.50)	(3.58)	(-3.46)	(-3.56)
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Observations	2,740	2,740	2,740	2,740
F-statistic for instrument	16.46	16.46	15.64	15.64

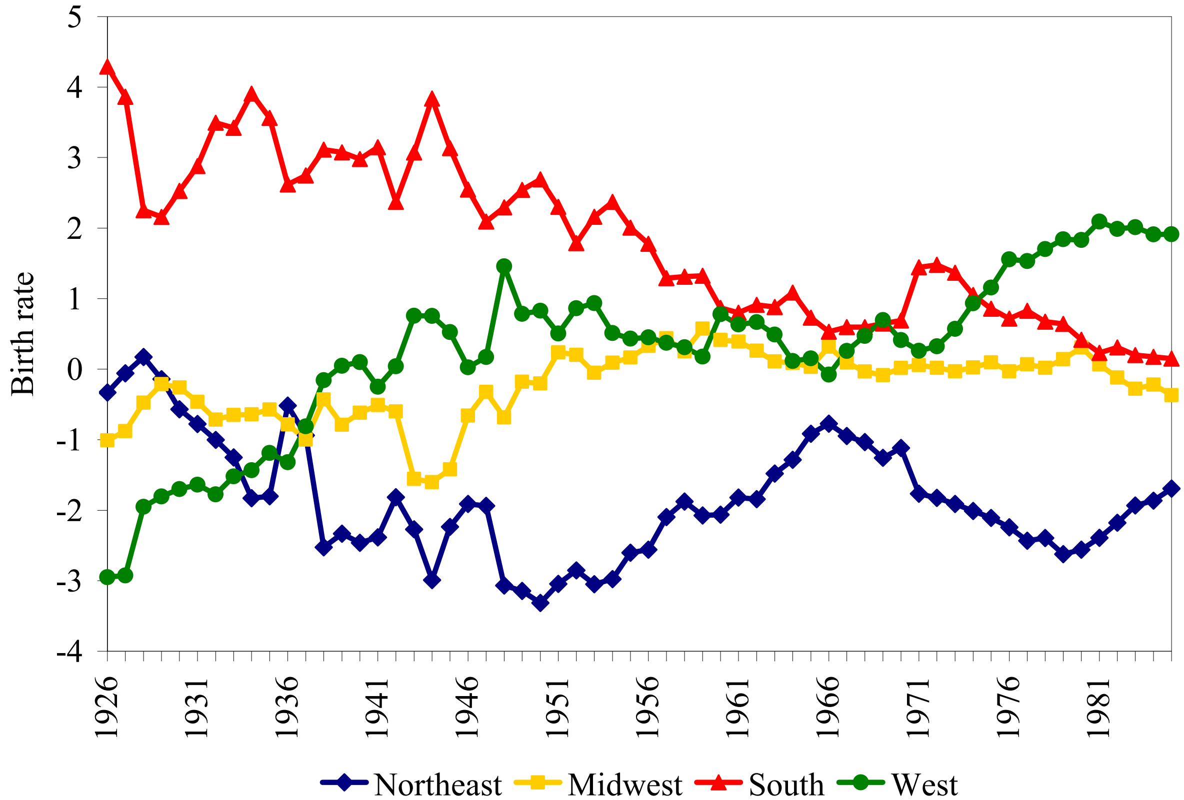

Figure 1. Evolution of the historical births used to construct the births-based labor force age structure by location. This figure shows the evolution of the births each year from 1926, or 64 years before the beginning of the sample period in 1990, to 1985, or 20 years before the end of the sample period in 2005. The census regions shown in the figure span the country. The birth rate is measured per thousand people and detrended.

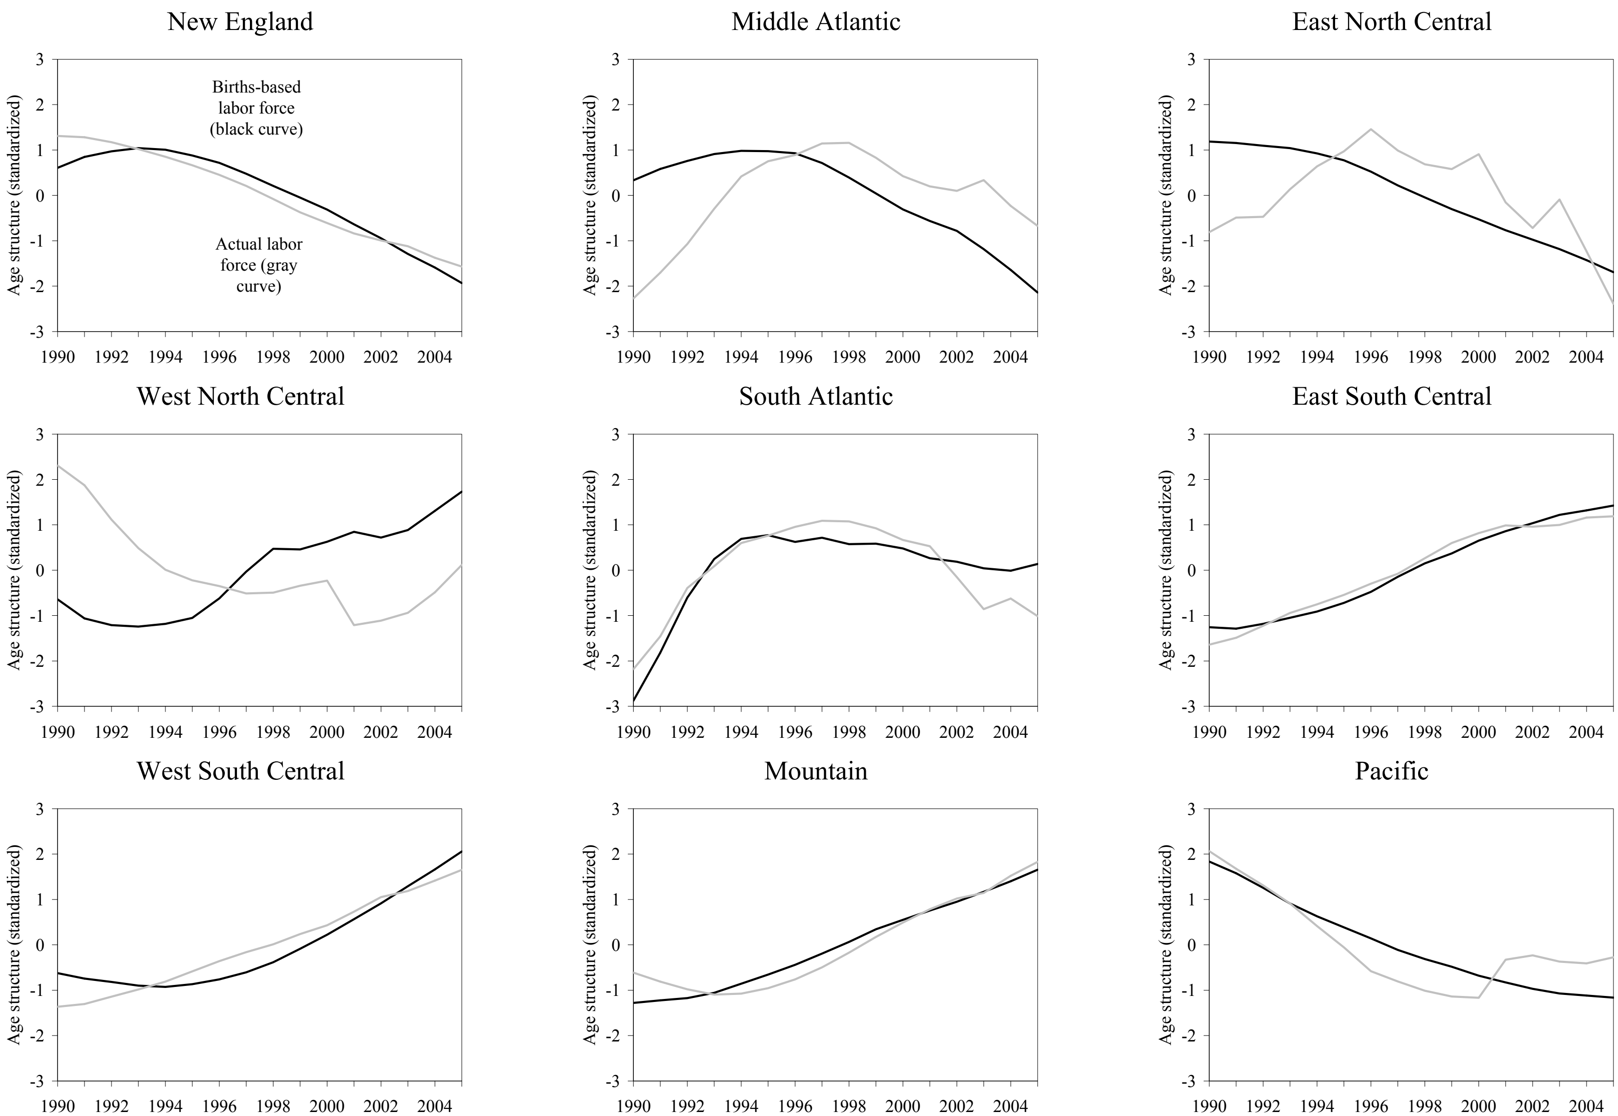

Figure 2. Evolution of the young share of the births-based and actual labor forces by location. This figure shows the young shares of the births-based and actual labor forces during the sample period (1990-2005). The census divisions shown span the country. The young share is first detrended and then standardized within each location. The births-based labor force in a given year is the labor force constructed based on the number of births each year during the prior 20-64 years.

Appendix Table 1
Variable Definitions

Commuting Zone-Level Variables Common to All Regressions	
Name	Definition
Age structure variables	
- Births-based young share	The young share (ages 20-39) of the labor force (ages 20-64) in a commuting zone and in a given year, constructed based on the number of births each year during the prior 20-64 years and adjusted for survival. The weight of a given year during the prior 20-64 years is the number of births that year.
- Births-based mean age	Mean age of the labor force (ages 20-64) in a commuting zone and in a given year, constructed based on the number of births each year during the prior 20-64 years and adjusted for survival. The weight of a given year during the prior 20-64 years is the number of births that year.
- Actual young share	The young share (ages 20-39) of the labor force (ages 20-64) in a commuting zone and in a given year. The weight of a given year during the prior 20-64 years is the number of people that year.
- Actual mean age	Mean age of the labor force (ages 20-64) in a commuting zone and in a given year. The weight of a given year during the prior 20-64 years is the number of people that year.
Control variables	
- Population size	The population of a commuting zone
- Income per capita	The income per capita of a commuting zone
- Government spending-to-total income	Local government expenditures divided by the total income of a commuting zone
- Educational attainment	The ratio of people with a bachelor degree to the population aged 25 years or older in a commuting zone
- University students per capita	The ratio of people enrolled in a bachelor degree program to the population of a commuting zone
- University patent counts per capita	The number of patents of all universities in a commuting zone. Scaled by the population of the commuting zone measured in hundred thousands.

Firm-Level Regressions	
Name	Definition
Innovation variables	
- Patent counts	The number of patents of a firm. Annual average over the next five years. Patent counts are adjusted for truncation (see Hall, Jaffe, and Trajtenberg (2005) and Kogan, Papanikolaou, Seru, and Stoffman (2017)).
- Patent citations	The weighted number of patent citations of a firm. Annual average over the next five years. Patent citations in a given year are weighted by the average number of citations per patent in the same technology class and grant year (for details, see Hall, Jaffe, and Trajtenberg (2005) and Kogan, Papanikolaou, Seru, and Stoffman (2017)).
- Citations per patent	Mean number of forward citations per patent. Forward citations per patent are scaled by mean of the same variable calculated using all patents in the same technology class and grant year cohort.
- Proportion of extremely useful patents	The proportion of a firm's patents in the top 1% of forward citations. Citations are ranked relative to patents in the same technology class and grant year cohort.
- Proportion of useful and novel patents	The proportion of a firm's patents in the top 20% of forward citations and the bottom 20% of backward citations. Citations are ranked relative to patents in the same technology class and grant year cohort.
- Volatility of citations per patent	Standard deviation of the number of forward citations per patent. Forward citations per patent are scaled by mean of the same variable calculated using all patents in the same technology class and grant year cohort.
- Longevity of citations per patent	Mean age of the newest forward citation per patent, where citation age is measured relative to the grant year. Scaled by the mean of the same variable calculated using all patents in the same technology class and grant year cohort and citation decile.
- Proportion of local citations	Mean proportion of backward citations of a firm's patents to patents in the firm's commuting zone (excluding self citations). Scaled by mean of the same variable calculated using all patents in the same technology class and grant year cohort.
- Patent value	The value of patents of a firm. Annual average over the next five years. Patent value estimates are from Kogan, Papanikolaou, Seru, and Stoffman (2017)).
Firm-level control variables	
- Total assets	AT from Compustat
- Market-to-book	(PRCC_F×CSHO)/(TXDITC+CEQ) from Compustat
- Cash flow-to-total assets	OIBDP/AT from Compustat
- Stock returns	Annualized mean daily stock returns
- Stock return volatility	Annualized standard deviation of daily stock returns

Inventor-Level Regressions	
Name	Definition
Innovation variables	
- Patent counts	The number of patents of an inventor constructed as described above for firms. Annual average over the next five years. For each patent with N inventors, each inventor is credited with 1/N patents.
- Patent citations	The weighted number of patent citations of an inventor constructed as described above for firms. Annual average over the next five years. For each patent with N inventors, each inventor is credited with 1/N patent citations.
Inventor-level control variables	
- Inventor age	Time elapsed since the inventor's date of birth
- Inventor experience	Time elapsed since the date of the inventor's first patent
- Inventor patent stock	The number of patents of the inventor
- Inferred R&D hub number of inventors	The number of inventors working for the firm in the same commuting zone as a given inventor

Appendix Table 2
Replication of Baseline Firm-Level Results with Alternative Forms of Clustering

This table shows the results of regressions of innovation on age structure. The regressions are the same as in Table 3 but with slight modifications, as indicated. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Standard Errors Clustered By Commuting Zone-Year				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	10.352***	11.903***	-0.370***	-0.429***
	(4.11)	(4.13)	(-3.85)	(-3.91)
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	14.03	14.03	12.91	12.91

Panel B: Standard Errors Clustered By Commuting Zone				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	10.352***	11.903***	-0.370***	-0.429***
	(2.97)	(2.92)	(-2.66)	(-2.64)
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	4.175	4.175	3.961	3.961

Panel C: Standard Errors Clustered By Firm				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	10.352***	11.903***	-0.370***	-0.429***
	(3.06)	(3.17)	(-2.94)	(-3.08)
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	179.1	179.1	152.6	152.6

Panel D: Standard Errors Double Clustered By Industry-Year and Commuting Zone-Year				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	10.352***	11.903***	-0.370***	-0.429***
	(3.93)	(4.11)	(-3.76)	(-3.94)
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	15.30	15.30	13.87	13.87

Panel E: Standard Errors Double Clustered By Industry-Year and Commuting Zone				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	10.352***	11.903***	-0.370***	-0.429***
	(2.91)	(2.93)	(-2.65)	(-2.66)
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	4.334	4.334	4.095	4.095

Appendix Table 3
Replication of Baseline Firm-Level Results for Various Robustness Tests

This table shows the results of regressions of innovation on age structure. The regressions are the same as in Table 3 but with slight modifications, as indicated. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Adding Age Structure Dispersion				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	12.418**	15.246***	-0.401**	-0.500***
	(2.40)	(3.01)	(-2.52)	(-3.24)
Age structure dispersion	-18.865	-30.518	-8.179	-18.599
	(-0.50)	(-0.75)	(-0.25)	(-0.54)
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	4.774	4.774	5.911	5.911
Panel B: Adding Interaction of Age Structure and Age Structure Dispersion				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	14.042***	16.600***	-0.529***	-0.621***
	(2.89)	(3.48)	(-3.15)	(-3.65)
Age structure dispersion	-58.179**	-63.281**	-63.380**	-70.726**
	(-1.98)	(-2.08)	(-2.08)	(-2.21)
Age structure	3.628	3.023	0.078*	0.073
× Age structure dispersion	(1.30)	(0.84)	(1.90)	(1.37)
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	9.500	9.500	8.134	8.134
Panel C: Using Time-Varying Firm Location				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	15.638***	18.287***	-0.558***	-0.659***
	(4.59)	(4.73)	(-4.55)	(-4.72)
Observations	7,945	7,945	7,945	7,945
F-statistic for instrument	206.6	206.6	206.6	206.6
Panel D: Using Only Firms that Do Not Change Location				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	12.760***	15.849***	-0.467***	-0.588***
	(5.20)	(6.09)	(-5.26)	(-6.15)
Observations	8,936	8,936	8,936	8,936
F-statistic for instrument	313.6	313.6	278.1	278.1

Panel E: Firm Age Measured from Founding Date				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	7.772***	9.508***	-0.278***	-0.343***
	(3.18)	(3.84)	(-3.06)	(-3.69)
Observations	11,558	11,558	11,558	11,558
F-statistic for instrument	203.0	203.0	158.5	158.5

Panel F: Adding a Control Variable for Managerial Age				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	11.828**	15.399***	-0.452**	-0.598***
	(2.37)	(2.70)	(-2.38)	(-2.70)
Observations	3,591	3,591	3,591	3,591
F-statistic for instrument	122.1	122.1	119.6	119.6

Appendix Table 4
Replication of Baseline Inventor-Level Results with Alternative Forms of Clustering

This table shows the results of regressions of innovation on age structure. The regressions are the same as in Table 7 but with slight modifications, as indicated. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Standard Errors Clustered By Commuting Zone-Year

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	7.336**	9.413***	-0.244**	-0.323***
	(2.02)	(2.60)	(-2.01)	(-2.68)
Observations	11,975	11,975	11,975	11,975
F-statistic for instrument	18.69	18.69	22.71	22.71

Panel B: Standard Errors Clustered By Commuting Zone

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	7.336**	9.413***	-0.244**	-0.323***
	(2.02)	(2.68)	(-2.10)	(-2.81)
Observations	11,975	11,975	11,975	11,975
F-statistic for instrument	8.091	8.091	10.14	10.14

Panel C: Standard Errors Clustered By Firm

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	7.336**	9.413**	-0.244**	-0.323**
	(2.49)	(2.33)	(-2.38)	(-2.30)
Observations	11,975	11,975	11,975	11,975
F-statistic for instrument	80.35	80.35	94.10	94.10

Panel D: Standard Errors Double Clustered By Technology Class-Year and Commuting Zone-Year

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	7.336**	9.413**	-0.244**	-0.323***
	(1.99)	(2.49)	(-1.99)	(-2.61)
Observations	11,975	11,975	11,975	11,975
F-statistic for instrument	20.36	20.36	24.52	24.52

Panel E: Standard Errors Double Clustered By Technology Class-Year and Commuting Zone

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	7.336**	9.413**	-0.244**	-0.323***
	(1.99)	(2.56)	(-2.07)	(-2.72)
Observations	11,975	11,975	11,975	11,975
F-statistic for instrument	8.783	8.783	10.96	10.96

Appendix Table 5
Replication of Firm-Level Innovation Characteristics Results with Alternative Forms of Clustering

This table shows the results of regressions of innovation characteristics on age structure. The regressions are the same as in Table 8 but with slight modifications, as indicated. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Standard Errors Clustered By Commuting Zone-Year, and Age Structure Measured Using Young Share

	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	3.070**	9.443***	12.893***	5.345***	0.488	9.681**
	(2.08)	(3.69)	(3.08)	(2.70)	(0.98)	(2.48)
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	31.30	31.30	31.30	33.06	30.92	31.30

Panel B: Standard Errors Clustered By Commuting Zone-Year, and Age Structure Measured Using Mean Age

	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	-0.117**	-0.341***	-0.486***	-0.207***	-0.022	-0.361**
	(-2.12)	(-3.43)	(-3.22)	(-2.80)	(-1.17)	(-2.46)
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	28.87	28.87	28.87	29.72	28.49	28.87

Panel C: Standard Errors Clustered By Commuting Zone, and Age Structure Measured Using Young Share

	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	3.070*	9.443***	12.893***	5.345***	0.488	9.681**
	(1.70)	(2.99)	(3.39)	(3.35)	(0.95)	(2.47)
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	10.17	10.17	10.17	10.99	10.05	10.17

Panel D: Standard Errors Clustered By Commuting Zone, and Age Structure Measured Using Mean Age						
	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	-0.117*	-0.341***	-0.486***	-0.207***	-0.022	-0.361**
	(-1.73)	(-2.87)	(-3.34)	(-3.39)	(-1.11)	(-2.42)
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	10.18	10.18	10.18	10.68	10.05	10.18

Panel E: Standard Errors Clustered By Firm, and Age Structure Measured Using Young Share						
	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	3.070	9.443**	12.893***	5.345*	0.488	9.681**
	(1.25)	(2.02)	(2.62)	(1.67)	(0.69)	(2.00)
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	219.0	219.0	219.0	191.3	217.6	219.0

Panel F: Standard Errors Clustered By Firm, and Age Structure Measured Using Mean Age						
	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	-0.117	-0.341**	-0.486***	-0.207*	-0.022	-0.361**
	(-1.29)	(-1.99)	(-2.70)	(-1.72)	(-0.80)	(-2.02)
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	201.6	201.6	201.6	177.9	200.4	201.6

Panel G: Standard Errors Double Clustered By Industry-Year and Commuting Zone-Year, and Age Structure Measured Using Young Share						
	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	3.070*	9.443***	12.893***	5.345**	0.488	9.681**
	(1.70)	(3.77)	(2.99)	(2.27)	(0.79)	(2.48)
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	35.24	35.24	35.24	36.93	34.85	35.24

Panel H: Standard Errors Double Clustered By Industry-Year and Commuting Zone-Year, and Age Structure Measured Using Mean Age						
	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	-0.117*	-0.341***	-0.486***	-0.207**	-0.022	-0.361**
	(-1.71)	(-3.56)	(-3.09)	(-2.26)	(-0.94)	(-2.53)
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	32.10	32.10	32.10	32.62	31.70	32.10

Panel I: Standard Errors Double Clustered By Industry-Year and Commuting Zone, and Age Structure Measured Using Young Share						
	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	3.070	9.443***	12.893***	5.345**	0.488	9.681**
	(1.48)	(3.05)	(3.27)	(2.61)	(0.78)	(2.48)
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	10.79	10.79	10.79	11.65	10.66	10.79

Panel J: Standard Errors Double Clustered By Industry-Year and Commuting Zone, and Age Structure Measured Using Mean Age						
	Dependent variables are innovation characteristics					
	Creativity #1: Citations per patent	Creativity #2: Proportion of extremely useful patents	Creativity #3: Proportion of useful and novel patents	Riskiness: Volatility of citations per patent	Horizon: Longevity of patent citations	Local spillovers: Proportion of local citations
Age structure	-0.117	-0.341***	-0.486***	-0.207**	-0.022	-0.361**
	(-1.50)	(-2.96)	(-3.19)	(-2.53)	(-0.90)	(-2.50)
Observations	5,768	5,768	5,768	4,361	5,757	5,768
F-statistic for instrument	10.77	10.77	10.77	11.27	10.64	10.77

Appendix Table 6
The Effect of Age Structure on Innovation: Firm-Level Analysis with Cross-Sectional Contrasts

This table shows the results of regressions of innovation on age structure with various interactions. The regressions correspond to Equation 1. In Panel A, the regressions are the same as in Table 3, but industries are classified as being in either innovative or non-innovative industries, and the specifications include the innovative industry dummy variable and its interaction with age structure and all control variables. Industries are classified as innovative based on their three-digit SIC codes, and otherwise they are classified as non-innovative. Panels B and C are analogous to Panel A. However, in Panel B, commuting zones are classified as more skilled or less skilled based on educational attainment, and in Panel C, commuting zones are classified as innovative or non-innovative based on local university patent counts. The actual age structure is instrumented with the births-based age structure. The actual labor force in a given year is self-explanatory. The births-based labor force in a given year is the labor force constructed based on the number of births each year during the prior 20-64 years. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Conditional Upon Industry Innovativeness

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	4.189	4.209	-0.073	-0.065
	(1.53)	(1.35)	(-0.60)	(-0.48)
Age structure × Innovative industry dummy variable	7.414**	9.488***	-0.436**	-0.536***
	(2.46)	(3.06)	(-2.54)	(-3.02)
Firm-level control variables?	Yes	Yes	Yes	Yes
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	76.86	76.86	49.34	49.34

Panel B: Conditional Upon Skill Level of Local Labor Market

	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	4.197*	4.736*	-0.114	-0.126
	(1.71)	(1.65)	(-1.09)	(-1.00)
Age structure × Skilled labor market dummy variable	6.940***	8.227***	-0.319**	-0.384***
	(2.79)	(3.10)	(-2.59)	(-2.89)
Firm-level control variables?	Yes	Yes	Yes	Yes
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	58.68	58.68	35.70	35.70

Panel C: Conditional Upon Innovativeness of Local Universities				
	Dependent variable is ln(1+mean patents per annum)			
	Age structure = Young share		Age structure = Mean age	
	Patent counts	Patent citations	Patent counts	Patent citations
Age structure	5.118*	5.409**	-0.157	-0.170
	(1.92)	(1.99)	(-1.51)	(-1.59)
Age structure × Innovative labor market dummy variable	7.934***	10.118***	-0.321***	-0.399***
	(3.86)	(4.47)	(-3.69)	(-4.15)
Firm-level control variables?	Yes	Yes	Yes	Yes
CZ-level control variables?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	14,087	14,087	14,087	14,087
F-statistic for instrument	83.01	83.01	50.98	50.98