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 using the local labor force projected based on historical births in each local labor market in the United States. Successive analyses of labor markets, firms, and inventors allow us to separate out effects such as firm and inventor life cycles. Corporate innovation activities reflect the innovative characteristics of younger labor forces. Additionally, firms in younger labor markets have higher valuations. Younger people as a group – inventors interacting with non-inventors – produce more innovation for firms through the labor supply channel rather than through a financing supply or consumer demand channel.

May 16, 2019

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

---

* Derrien is at HEC Paris, and Kecskés and Nguyen are at the Schulich School of Business, 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, Laurent Frésard, Louis Gagnon, Jim Goldman, Johan Hombert, 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, Miguel Palacios, Tianyue Ruan, Christophe Spaenjers, Johan Sulaeman, Lingna Sun, David Thesmar, Emiko Usui, Philip Valta, John Wald, Wei Wang, Zhao Wang, Dongling Xu, Bernie Yeung, Alminas Žaldokas, and Shan Zhao, and seminar participants at the 2018 Asian Development Bank Conference, City University of Hong Kong, the 2019 Eastern Finance Association Conference, the 2018 European Finance Association Conference, the 2018 Financial Intermediation Research Society Conference, the 2018 Financial Management Association Conference, the 2018 Financial Management Association Asia/Pacific Conference, Grenoble EM, the 2018 HEC Paris Workshop on Banking, Finance, Macroeconomics, and the Real Economy, the Hong Kong Polytechnic University, the 2019 Inter-Business School Finance Seminar, the 2019 International Conference on Technology, Demographics, and the Labor Market, the 2018 Irish Academy of Finance Conference, McMaster University, the 2019 Midwest Finance Association Conference, the 2018 Northern Finance Association Conference, Queen's University, the 2018 SFS Cavalcade Asia-Pacific, the 2018 Symposium of the Spanish Economic Association, 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, and the University of Toronto. 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. Indeed, the influence of the institutional and especially legal features of local labor markets on various corporate activities is the subject of a small but growing literature. For instance, unionization influences capital structure, and employment protections encourage corporate innovation while discouraging takeovers.[1] We consider a more fundamental but consequential aspect of labor markets: their demographics. The age structure of the local labor force, in particular, matters because young and old workers are heterogeneous inputs into the firm's production function and their effect on the firm's innovation output is plausibly different. In particular, 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 people, and younger labor forces, more innovative.

Corporate innovation activities benefit from a younger labor force in a number of ways. First, firms 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, and managers. 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 results in more individual, firm, and aggregate innovation. We hypothesize that it is through this labor supply channel that age structure affects corporate innovation.

[1] Matsa (2010); Acharya, Baghai, and Subramanian (2013, 2014); and Dessaint, Golubov, and Volpin (2017).
[2] See Liang, Wang, and Lazear (2017) for references.
[3] This is analogous to local agglomeration effects, for example, on 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).

Testing this 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, thus migration can affect age structure. Moreover, age structure may be driven by factors that it has in common with innovation but which are unobservable. Indeed, this endogeneity is obvious from our most basic empirical analysis. As we show, regressions of innovation on the actual age structure of the labor force lead to inferences that are unreliable. Since we are interested in the causal effect of local age structure on local innovation, we need to ensure that the latter is not confounded with the former.

Our novel approach is to use the age structure of the labor force projected based on historical births. We examine commuting zones, which are ideal units of analysis to study local labor markets because they are designed to capture largely self-contained areas in which people live and work. For every year and commuting zone in the United States, we reconstruct the population of 20-64 year olds using historical births in the prior 20-64 years, adjusted for survival. The resulting native born labor force captures the labor force in the absence of migration to and from the commuting zone, and it is plausibly exogenous to innovation today. It can also be viewed as an endowment of a given location that is not completely eliminated by migration across labor markets. In our empirical analysis, we use this local endowment to explain local innovation.[4]

Examination of the age structure that we use in our analysis shows that there is significant variation in age structure across locations. However, given our relatively short sample period (1990-2005), the age structure of the local labor force exhibits steady time trends but

[4] Our focus on the native born labor force does not require that innovation be produced only by the native born labor force. Rather, the only requirement is that the age structure of the native born labor force be the cause of such innovation as we capture in our analysis.

otherwise very little time-series variation. Our results on innovation are therefore largely identified based on cross-sectional variation in the age structure rather than time-series variation.

[Insert Figure 1 about here]

Figure 1 provides a simple but clear illustration of our empirical approach and main finding. We plot the mean age of the projected labor force and the number of patents per capita for every commuting zone in the U.S. in the year 2000. The results show a strong negative relationship between age and innovation: younger labor forces produce more innovation.

Turning to our regression analysis, we first examine innovation by public and private firms aggregated to the commuting zone level, as in Figure 1. Our regression specifications include state-year fixed effects as well as control variables that account for factors that are correlated with both local age structure and local innovation. Any such factors must be generated by shocks that were present 20-64 years ago, that directly affect historical births, and that also directly affect innovation today. To capture such factors, we control for the commuting zone's size, wealth, growth, government expenditures, educational attainment, and university patents.

For most of the rest of our analysis, we focus on firms. With detailed firm-level data, we can identify more precisely the channels through which innovation is affected by the age structure of the labor force at the headquarters location of firms. Our main specifications include firm-level control variables as well as state-year fixed effects, industry-year fixed effects, and firm age fixed effects. 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, our results must be interpreted as for firms of the same age, which ensures that they cannot be explained by firm life cycle effects. At both the commuting zone and firm levels, 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 the mean age causes a roughly 7% increase in innovation.

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 their innovation activities reflect these characteristics. To this end, we use measures of the creativity, riskiness, longevity, and interactivity of patents. We use patent citations, both those made by and made to the patents of our sample firms (backward and forward citations, respectively), to capture the characteristics of the firm's innovation activities. We find that the innovation activities of firms in younger labor markets do indeed reflect the innovative characteristics of younger people.

Our data allow us to move to an even higher level of granularity and focus on inventors who work for public and private firms. At this level of analysis, we can control for the ages of inventors and firms as well as firm fixed effects and inventor network scale effects. This allows us to distinguish between inventors and other workers in the local labor force. The results do show that younger inventors produce more innovation, but they also show that even after controlling for inventor age and other relevant factors, 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 their innovation. The general work environment appears to be important for the production of innovation.[5]

[5] 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 the age structure of the labor force are distinctly different constructs that cannot be directly compared. 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.

Having established that the age structure of the local labor force affects the innovation activities of local firms, we identify more precisely the channels through which this happens. Returning to firms, we distinguish between the labor supply channel and two other channels. In the financing supply channel, our age structure actually captures the local investor pool rather than the local labor force. Local investors who are younger are more willing to finance risky projects such as those that produce innovation, so firms in younger locations produce more innovation. We are able to identify the location at which firms produce innovation, their R&D hubs, for about half of our sample. When we include the age structure at both headquarters and R&D hubs, R&D hubs absorb most of the effect of headquarters. Additionally, we find that age structure has no effect on leverage. This is consistent with the labor supply channel but inconsistent with the pure financing supply channel.

In the consumer demand channel, our age structure actually captures the local consumer pool rather than the local labor force. Younger consumers demand more innovative products, so the firm produces more innovation to satisfy demand for it. Separating firms based on whether their workforce serves local versus non-local consumer demand, we find that the results are driven entirely by firms with local employees but non-local customers. This is inconsistent with the pure consumer demand channel.

We also find that our main result, i.e., that younger labor forces produce more innovation, is weakest for the youngest firms, it increases in firm age, and it is strongest for the oldest firms. This finding is inconsistent with a matching interpretation in which innovative firms start up in younger labor markets. Rather, this finding, together with the fact that 80% of our sample firms do not move during their lives, is more consistent with firms taking as given the age structure and other characteristics of the local labor market and innovating as a function thereof.

Finally, because innovation creates growth opportunities for firms, we examine the relationship between labor force age structure and firm valuations. We find that firms in younger labor markets have higher market-to-book ratios, with a one standard deviation decrease in the mean age causing a roughly 3% increase in valuations. These results reflect the value created by younger labor forces. We conclude that the age structure of the labor force has important consequences for inventors, firms, and the local economy.

Our main contribution is to the nascent literature on labor and finance, which studies the institutional and legal features of local labor markets on corporate activities. One pioneering paper finds that workforce unionization leads firms to use leverage strategically in bargaining over employment contracts (Matsa (2010)). Other papers in this literature study differences across jurisdictions in employment laws that provide protections or impose restrictions. Early papers find that lower employment risk, as captured by higher unemployment insurance benefits, increases leverage (Agrawal and Matsa (2013)) and entrepreneurship (Hombert, Schoar, Sraer, and Thesmar (2018)). A series of more recent papers finds that stronger wrongful discharge laws, as proxying for firing costs, increase innovation, but they also reduce takeover synergies and activity, leverage, and investment.[6] Additionally, non-compete agreements can increase investment by restricting labor mobility (Jeffers (2018)).

By contrast, our paper studies demographics, a more fundamental but consequential aspect of labor markets. 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 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

---

[6] Acharya, Baghai, and Subramanian (2013, 2014); John, Knyazeva, and Knyazeva (2015) and Dessaint, Golubov, and Volpin (2017); Simintzi, Vig, and Volpin (2015); and Bai, Fairhurst, and Serfling (2017).

have lower payouts (Becker, Ivkovic, and Weisbenner (2011)); and product market demographics influence stock returns (Dellavigna and Pollet (2007, 2013)).

We also make an important contribution to the literature on the macroeconomic consequences of demographics. We examine innovation whereas prior studies examine unemployment, aggregate output volatility, productivity, and aggregate stock returns.[7] 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 as an instrument for the age structure of the actual labor force (e.g., only the birth rates 20-39 years prior, or only the age structure 10 or 20 years prior). By contrast, we take a comprehensive and precise approach by using historical births data starting from the 1920s to reconstruct the age structure of the entire labor force in each commuting zone in the U.S.

The rest of this paper is organized as follows. Section 2 presents the methodology. Sections 3, 4, and 5 present the analyses at the commuting zone, firm, and inventor levels. Section 6 presents the results for firm value. Section 7 concludes.

2. Methodology

2.1. Measurement of Age Structure

Our objective is to explore and better understand the effect of the age structure of the labor force on corporate innovation. To this end, we run regressions principally of measures of innovation activity on measures of the age structure of the labor force 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, including innovation, can affect the local labor force because economic

[7] Respectively, Shimer (2001); Lugauer (2012) and Jaimovich and Siu (2009); Acemoglu and Restrepo (2018); and Poterba (2001) and Goyal (2004).

conditions affect the migration of people. This can be a serious problem if the population and the economy are measured contemporaneously.

For this reason, we do not use the actual labor force to measure age structure but rather the labor force projected based on historical births. 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. At the same time, the historical births over a long period of time are an integral component of the age structure today. However, the age structure of the projected labor force 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 historical births plus immigration minus emigration, thus the projected labor force is free from the effect of migration, which is likely correlated with contemporaneous economic activity.[8]

We construct our measures of age structure as follows. For every year and location in the United States, we use historical births in the prior 20-64 years to reconstruct the labor force that year, i.e., the population of 20-64 year olds. For example, in 1990, a person born in 1926 will be 64 years old, and a person born in 1970 will be 20 years old. We therefore use the historical births from 1926 to 1970 to project the labor force in 1990. We adjust for survival for each age group using national time-varying survival rates every decade. We then use the projected labor force in 1990 to construct our measures of the age structure in 1990. Our sample begins in 1990 because we need data on historical births to project the current population for up to 64 years of age, and births data at the county level only become well populated in the mid-1920s. Our

[8] We note that the native born labor force of a commuting zone necessarily excludes Americans born outside the commuting zone as well as foreigners.

sample ends in 2005 because we examine future innovation for up to five years, and the patent data end in 2010.

We use the age structure of the labor force projected based on historical births directly rather than as an instrument for the age structure of the actual labor force. Unlike the prior literature, we do not take an instrumental variables approach because the exogenous component of the actual labor force is the native born labor force, and this is precisely what we measure with the labor force projected based on historical births. Indeed, the projected labor force itself is our exogenous variable of interest, so we do not use it as an instrument.[9]

2.2. Empirical Examination of Age Structure

We examine the historical births data that we use to construct our measures of age structure. 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.[10] 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 projected labor force, i.e., the population of 20-64 year olds, during our entire sample period, i.e., 1990-2005). To allow us to make not only time-series but also cross-sectional comparisons, we focus on the

[9] There are a few studies of other consequences of the age structure (especially unemployment and aggregate output volatility) that instrument for the current age structure in various ways. 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 country. We are able to significantly improve upon existing approaches in a number of ways by using the entire native born local labor force in our analysis. First, we have births data on the entire labor force, not just a limited and recent number of years thereof. Second, we directly measure the age structure of the native born labor force, rather than having to instrument the age structure of the actual labor force with a small sample of births. Finally, we study local labor markets directly using commuting zones, instead of studying heterogeneous local labor markets aggregated into states or countries.

[10] 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.

census region level in our next analysis.[11] We detrend the data, for ease of exposition, by removing the annual average across census regions using year fixed effects.

[Insert Figure 2 about here]

Figure 2 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 the figure 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 projected age structure (i.e., based on historical births) are driven decades later by time-invariant heterogeneity across locations (e.g., persistent differences in economic activity).

We examine in detail the age structure at the center of our analysis, which we measure at the commuting zone level. Commuting zones are ideal economic-geographic units to study the effect of demographics on innovation. Popular in labor economics, they are designed to capture the local labor market in which people live and work (Tolbert and Sizer (1996)). Indeed, they are constructed as clusters of counties characterized by strong labor market interactions within

[11] 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.

commuting zones and weak interactions across commuting zones (Autor and Dorn (2013)). There are 741 commuting zones that together cover the entire land area of the U.S.

We aggregate counties to commuting zones using the county-commuting zone crosswalk from Autor and Dorn (2013). We use two measures of age structure that are also widely used in the literature. The first is the mean age of the labor force, which has the advantage of being very simple. The second measure is the "young share" of the labor force, where the young are defined as the population of 20-39 year olds. This measure has the advantage of comparing the size of our specific group of interest, the young in the labor force, to the size of the entire labor force, i.e., the young plus the old.

We first examine the evolution of the raw age structure during our sample period. We focus on the mean age of the labor force rather than its young share. The results for both measures of the age structure are similar, but the mean age is perhaps more natural to interpret than the young share. Furthermore, while our analysis generally examines the labor force at the commuting zone level, in this case we focus on the census division level. Doing so allows us to visually represent the entire country with a manageable number of units compared to the 50 states that are sorted into the nine census divisions. We should note that our results are similar at more granular levels (e.g., at the state or commuting zone levels).

[Insert Figure 3 about here]

Figure 3 shows the mean age of the labor force for each census division. Panel A shows the labor force projected based on historical births, which is the focus of our analysis throughout the paper, and Panel B shows the actual labor force. Several patterns emerge. First, even at the division level, without examining more granular levels, there is considerable cross-sectional variation in the data. These patterns are similar for the projected and actual labor forces. Second, the U.S. is steadily becoming older over time, for every division of the country. The steadiness of this aging process implies that there is very little time-series variation in age structure at the division level (or even at the commuting zone level, as we verify), and such time-series variation as exists can be captured by a steady time trend.

Next, we compare the evolution of the mean age of the projected and actual labor forces for each census division. Given the findings of Figure 3, we detrend the mean age, removing the annual average across census divisions using year fixed effects, and standardize it at the census division level, subtracting the mean and dividing by the standard deviation. As a result, within each census division, we are comparing projected and actual labor forces both of which have a mean of 0 and a standard deviation of 1.

[Insert Figure 4 about here]

Figure 4 shows that the projected and actual labor forces generally follow the same direction over time. This provides validation for our use of the age structure projected based on historical births. Indeed, the correlation between the projected and actual age structures is approximately 0.5, which indicates that historical births explain a significant part of the current labor force. The difference between projected and actual labor forces is migration, which is modest (about 3% per annum across states since the late 1940s, the starting point of the data).

To summarize the results of our empirical examination of age structure, the projected and actual labor forces are similar in terms of age structure patterns, across both locations and time. Moreover, there is considerable cross-sectional variation, with age structure varying significantly from one location to another. However, the age structure exhibits steady time trends but otherwise very little time-series variation.

These results 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. This is especially the case for our firm- and inventor-level analyses since firms and inventors typically exist in the data for only a few years.[12] Second, it is impractical to identify entirely off the time-series variation in age structure. Instead, we identify off the cross-sectional variation because it accounts for most of the total variation. Since unobservable cross-sectional and time-series factors can potentially affect innovation, we always include state-year fixed effects to absorb them (as well as industry-year fixed effects, firm age fixed effects, and control variables, whenever appropriate). We thus identify off the residual variation across commuting zones within a given state in a given year (and considerably less residual variation in general since we include additional fixed effects and control variables).[13]

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). In our commuting zone-level analysis, we scale patents by the population of the commuting zone. However, in our firm-level analysis, we use unscaled patents following the large literature on firm-level innovation. We obtain data on patents, patent assignees (firms), and patent inventors during 1975-2010 from Li et al. (2014). 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.

[12] Our approach of using quinquennial or decennial observations is similar to that of Becker (2007) and Becker, Ivkovic, and Weisbenner (2011)).

[13] We would ideally like to have several decades of data on the projected age structure and innovation, and then study the effect of the former on the latter while controlling for time and location fixed effects. Unfortunately, the requisite historical births data do not exist before the 1920s, to our knowledge.

There are two important adjustments that we make. First, we adjust patent citations for the changing distribution of patent citations each year by scaling citations in a given year by the average number of citations per patent that year. This is necessary to solve the problem of patents granted later in time having fewer citations at the end of our sample period.

Second, in our inventor-level analysis, we adjust patent counts and citations by the number of inventors of the patent. This is necessary for apportioning credit because most patents are granted to more than one inventor (in fact, three inventors, on average).[14] 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. Levels of Analysis and Model Specifications

We perform our analysis principally at the commuting zone, firm, and inventor levels, which we explain in turn below. Our main analyses have a number of common features. We run regressions of innovation outcomes on age structure. Our measures of age structure are constructed using the projected labor force in the commuting zone. Since age structure is persistent, we use four quinquennial observations in our main regressions, from 1990 to 2005, rather than annually.[15] Since both age structure and innovation may have variation in common across space and time, we also include state-year fixed effects to remove such variation.

We examine the effect of the current age structure on future innovation, in the short run and the long run, measured the next year and the average of the next five years, respectively. This is necessary because innovation activities can take several years to produce results, so long-run regressions may be appropriate. Moreover, innovation outcomes are relatively infrequent at

[14] In commuting zone- and firm-level analyses, this is not an issue. It is rare for a patent to have inventors located in different commuting zones or 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.
[15] The results are similar regardless of which year from 1990 to 1994 we use to start the quinquennial periods.

the one-year horizon (e.g., in smaller commuting zones or for younger firms), so averaging the outcomes of several years generates more precise estimates than only using a single year.

3. Commuting Zone-Level Analysis

3.1. Model Specification

We begin our analysis at the commuting zone level. In this analysis, the equation for the baseline regressions is:

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

where c indexes commuting zones, s indexes states, and t indexes years. $X_{c,t}$ is a vector of commuting zone-level control variables, and $\delta_{s,t}$ is a state-year fixed effect. The location of innovation is determined by the address on the patent document. The state-year fixed effects in our regressions capture much of the variation across locations and time. Such variation as remains is within a given state in a given year but across commuting zones. Our results must therefore be interpreted as estimating the extent to which the variation in age structure across commuting zones explains the variation in innovation across commuting zones – in both cases, within a given state in a given year.

Since the age structure in all locations exhibits steady time trends but otherwise very little time-series variation, we do not include commuting zone fixed effects in our main specifications. As shown in Section 2.2 above, historical birth rates, which determine the age structure of the native born labor force, vary significantly over time in every location: it is not the case that the native born labor force in some locations is always younger than in other locations. Therefore, it is unlikely that the relationship between age structure and innovation is driven by time-invariant factors, with some locations being always younger and more innovative than others. However, as a robustness test, we do include commuting zone fixed effects in slightly modified regression specifications, in Section 3.3 below. We find that time-invariant commuting zone factors are unlikely to explain our results incrementally to our control variables.

We also account for factors that may be correlated with both the age structure of the native born labor force and innovation. In the context of our analysis, these factors can be time-varying or time-invariant, but they must meet very specific requirements. First, they must be generated by shocks that were present 20-64 years ago. Second, these shocks must directly affect historical births and hence age structure today. Third, these shocks must also affect innovation today (above and beyond any indirect effect they may have on innovation today through their effect on historical births and hence age structure today). Ideally, we would like to measure such shocks directly when they were present 20-64 years ago and then aggregate them to capture their effect today. In our main specifications, we use control variables that are contemporaneous to age structure and innovation but which may capture the effect of shocks that were present 20-64 years ago. We do so for simplicity and because our control variables are slow-moving. In robustness tests, we use their average value over 20 years and find similar results.

Two factors come to mind that may meet the aforementioned requirements: economic conditions and certain types of long-term investments. First, we account for local economic conditions using three standard control variables. We control for the size of the population today because it reflects the cumulative long-term effect of shocks to the local economy. 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 age structure today. 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)), which can affect both age structure and innovation. Motivated by the same reasoning, we also control for income per capita as a proxy for wealth, and the growth rate of total income to capture economic growth.

Second, certain types of long-term investment, particularly in infrastructure, education, and research, can affect contemporaneous economic conditions, hence migration and births, and ultimately age structure in future decades. Through its additional effect on future economic conditions, such investment can also affect future innovation. We account for such local investment using another three control variables. We control for government expenditures because, by way of example, greater spending, whether on infrastructure, education, research, or other goods and services, can directly lower the costs of having children. It can also create jobs, which indirectly encourages people to have children. In both cases, government expenditures can affect contemporaneous births and hence the future age structure of the commuting zone. They can also affect the future climate for business endeavors, including innovation.

The two other control variables that we use are the long-term consequence 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's degree or higher to the population aged 25 years or older, and local university patent counts per capita. In the same manner as we already described, investments in education and research made 20-64 years ago can affect historical births and hence the age structure today, and they can also affect the business atmosphere today and hence the production of innovation. However, investments in education also affect innovation today by providing people with the knowledge, skills, and training needed to produce innovation. Moreover, investments in research also affect innovation today through the typically substantial spillovers between educational institutions and the private sector. Overall, our control variables should capture most factors that may be correlated with both age structure and innovation.

We obtain data for our control variables from the Census Bureau and the patent database. We weight each commuting zone by the size of its labor force to account for the relative economic importance of commuting zones. We cluster standard errors by state-year, but the results are robust to clustering by commuting zone. We winsorize variables whenever appropriate at the 1st and 99th percentiles.

3.2. Sample and Descriptive Statistics

[Insert Table 1 about here]

The firms in our sample are the public and private firms in the patent database. The sample itself comprises 2,964 commuting zone-quinquennial period observations corresponding to 741 unique commuting zones. Requiring data on historical births decreases the sample size slightly, to 2,748 observations. This is because Alaska and Hawaii are completely missing the requisite data owing to their late admission to the Union (in 1959), and South Dakota and Texas are mostly missing the requisite data during the first quinquennial period. Table 1 provides descriptive statistics for the samples corresponding to our various levels of analysis starting with the commuting zone level of immediate interest (Panel A). The age structure is comparable across all three levels of analysis. The typical age of the labor force is about 40 years, and the young share is roughly 50%.

We report descriptive statistics on innovation not only here in the commuting zone-level analysis but also in our firm- and inventor-level analyses. However, we note here that, 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 commuting zone-years have no firms associated with them, while others have multiple firms. This is also the case for inventors. Second, only star inventors are used by construction at the inventor level, whereas all inventors are used at the other two levels. Third, the number of patents is adjusted for the number of inventors but only at the inventor level. Finally, patents are scaled by population but only at the commuting zone level. Within each level of analysis, whether we use a one-year or five-year horizon, the distribution of patents is similar on an annual basis. There are approximately 10 patents per annum on average (median 5) at the commuting zone-year level (per hundred thousand people).

3.3. Results

We consider regression specifications with various combinations of our six control variables to allow us to examine the sensitivity of the results to particular control variables. These combinations are as follows: the three control variables for economic conditions (population size, income per capita, and growth rate of total income), then each of three additional control variables (government expenditures, educational attainment, and university patent counts), and finally all six control variables.

[Insert Table 2 about here]

Table 2 presents the results, with the mean age measure of age structure in Panel A and the young share measure in Panel B. The results consistently show that younger labor forces produce more innovation. Panel A indicates that a one-year decrease in the mean age causes a roughly 10% increase in patents-to-population, whether counts or citations. For a one standard deviation decrease (2.3 years), patents increase by 25%-30% relative to their mean and by a somewhat smaller 20% or so compared to their standard deviation. Panel B indicates that a one percentage point increase in the young share causes a roughly 3% increase in patents. For a one standard deviation increase (8.5 p.p.), patents again increase by 25%-30% relative to their mean and about 20% compared to their standard deviation. Overall, the results for age structure are similar across all specifications (even in terms of the magnitude of the coefficient estimates), whether we examine the short run or the long run (the next year or the average of the next five years, respectively).

Like age structure and innovation, the control variables that we use are time-varying but slow-moving, so it should be sufficient to measure them in the single year at the beginning of each quinquennial period. Nevertheless, since the factors that we are trying to capture must have been generated by shocks that were present 20-64 years ago, we examine the robustness of our results to measuring our control variables over a longer period of time. We rerun the regressions in the last four columns of Table 2, but rather than measuring our control variables in a single year, we use their average value during the 20 years ending at the beginning of each quinquennial period. We use a 20-year window both because it is reasonably long and because the data are limited in time.[16] In untabulated results, the coefficient estimates decrease by about 15%-20% in magnitude, but they remain significant at the 1% level.

Overall, the plausibly exogenous projected age structure reliably leads to the inference that younger labor forces are more innovative. We examine whether this is the case for the endogenous actual age structure. To facilitate comparison, we include the control variables in the same sequence as before. We tabulate the results for the mean age in Appendix Table 2 Panel A; the untabulated results for the young share are similar. By contrast to the projected age structure,

[16] The data for size and wealth only begin in 1969, for growth in 1970, for government expenditures in 1972, for educational attainment in 1980, and for university patents in 1975. We go back as far as possible for up to 20 years, which is entirely possible by the time of our final two quinquennial periods.

which has a reliably negative effect on innovation in all specifications, the actual age structure leads to inferences that are unreliable. Depending on the control variables included, the partial correlation between the actual age structure and innovation is negative, positive, or insignificant.

A detailed study of the actual age structure is beyond the scope of this paper. Nevertheless, we do directly compare the effects of the projected and actual age structures. This allows us to characterize the endogeneity of the actual age structure, and specifically the role of migration. To this end, we combine each of the corresponding regressions in Table 2 and Appendix Table 2 Panel A into a single regression. In so doing, we capture both the effect of the native born labor force on innovation as well as the incremental relationship between the migrant labor force and innovation. Appendix Table 2 Panel B shows that the coefficient estimates for the projected age structure are very similar to those in Table 2, while the coefficient estimates for the actual age structure are consistent with those in Appendix Table 2 Panel A. Once again, the projected age structure is reliably negative, whereas the actual age structure is unreliable. Based on the full specification (all control variables included), which shows coefficient estimates for the actual age structure that are positive, it appears that older migrants tend to move to more innovative locations.

4. Firm-Level Analysis

4.1. Model Specification

Our second analysis is at the firm level. We have detailed data on thousands of publicly traded firms, so we can control for a variety of firm-level characteristics and thus improve our model specification. We can also refine the interpretation of the results. Specifically, we can examine industry composition effects (more innovative industries matching to younger labor markets) and firm composition effects (more innovative firms matching to younger labor markets). We can also rule out firm life cycle effects (younger firms being more innovative rather than younger labor forces). We can also identify more precisely the channels through which age structure affects innovation.

In our firm level analysis, the equation for the baseline regressions is:

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

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, $\delta_{s,t}$ 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 together.

We measure age structure for firms using the commuting zone in which they are headquartered. We use headquarters location because firms tend to locate their R&D hubs, at which they produce innovation, close to their headquarters rather than dispersing them geographically (Howells (1990) and Breschi (2008)). We confirm this stylized fact for publicly traded firms in our sample. An R&D hub is any location in which there is at least one inventor working for the firm during the prior ten years. We find that about 50%-75% of a firm's innovation is produced in the commuting zone of its headquarters (details later). In most tests, we use headquarters rather than R&D hubs because, on balance, the former measure location more cleanly than the latter. In particular, given the definition of R&D hubs, their measurement may be stale and sparse. Moreover, since we are interested in the general work environment, we need to capture not just the location of inventors but also of non-inventors, both inside and outside the firm.

In firm-level regressions, we also include firm-level control variables, industry-year fixed effects, state-year fixed effects, and firm age fixed effects. The control variables are total assets, market-to-book, cash flow-to-total assets, stock returns, and stock return volatility. Both age structure and innovation may have common variation across industries and time, so we remove such variation using industry-year fixed effects, where industry is captured by two-digit SIC codes. If the effect of age structure on innovation at the commuting zone level is driven by more innovative industries tending to be in younger locations, then this effect should disappear at the firm level if we include industry-year fixed effects. If this effect does not disappear, then we can rule out the possibility of industry composition effects. Moreover, these fixed effects allow us to mitigate the biases in the patent data reported by Lerner and Seru (2017) along the dimensions of time, industry, and location.

Similarly, 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, our results must be interpreted as for firms of the same age. We measure firm age from the date the firm begins trading publicly in CRSP both because this is standard practice and because firm founding dates are not widely available. Finally, standard errors are clustered by industry-year, but the results are robust to clustering by firm.

4.2. 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 the patent database. Data on publicly traded firms are from CRSP and Compustat. Returning to our descriptive statistics in Table 1, there are about 5 patents per annum on average (median 0) at the firm-year level.

4.3. Baseline Results

[Insert Table 3 about here]

Table 3 presents the regression results, with the mean age in Panel A and the young share in Panel B. The results confirm that younger labor forces produce more innovation. A one standard deviation decrease in the mean age (2.2 years) causes at least a roughly 7% (=0.03×2.2) increase in patents, whether counts or citations (Panel A). Similarly, a one standard deviation increase in the young share (8.4 percentage points) causes at least a 7% (=0.8×0.084) increase in patents (Panel B). The results are once again similar in the short run and the long run (on the basis of patents per annum). Since the results at the firm level, with industry-year fixed effects, are similar to the results at the commuting zone level, which are necessarily without industry-year fixed effects, we can rule out the possibility of industry composition effects. Similarly, we can rule out a firm life cycle interpretation of our results because our firm age fixed effects force our results to be interpreted as for the firms of the same age.

As already explained, it is impractical to completely remove the cross-sectional variation in age structure and rely entirely on its time-series variation. The situation is obvious at the commuting zone level: both age structure and patent outputs are persistent at the annual frequency, and the sample period (1990-2005) is relatively short. To demonstrate this empirically, we rerun the regressions in Table 2 but at the annual frequency and using as explanatory variables only commuting zone fixed effects and year fixed effects. These two fixed effects by themselves explain 91%-93% of the variation in innovation at the one-year horizon and 96%-97% of the variation at the five-year horizon. The extreme explanatory power of these fixed effects is less likely to be an economic regularity than a statistical artifact of regressing one persistent variable on location and time fixed effects using a relatively short sample period. Our findings would likely be very different if we had a sample spanning at least several decades.

The situation is similar at the firm level but less dire because there is greater time-series variation in patent outputs for firms than for commuting zones. We rerun the regressions in Table 3 (including all control variables and fixed effects) but at the annual frequency and adding commuting zone fixed effects. Panels A and B of Appendix Table 3 show that our results are similar, if not as strong, if 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 replacing commuting zone fixed effects with even more demanding firm fixed effects (which subsume commuting zone fixed effects). Panels C and D of Appendix Table 3 show that our results always have the correct sign, their magnitude is smaller though still respectable, but they are not always statistically significant.

Furthermore, firm-level control variables are more precise at explaining corporate innovation than commuting zone-level control variables. This is why we use the former rather than the latter in our firm-level regressions. Nevertheless, to be conservative, we include our commuting zone-level control variables in our firm-level regressions. Appendix Table 4 Panels A and B show that our results are robust to their inclusion.

Turning to urban versus rural commuting zones, the population size control variable makes it unlikely that our results would be driven by the most populous commuting zones.

However, we examine this possibility directly by dropping the 10, 25, and 50 most populous commuting zones. These filters respectively eliminate approximately 25%, 40%, and 55% of the population of the country. Nevertheless, our results are robust to their exclusion, as illustrated by Panels C and D of Appendix Table 4 (dropping the 50 most populous commuting zones).

Regarding the measurement of firm age, we use listing dates rather than founding dates, but this is unlikely to generate very much measurement error because firms tend to go public when they are young. However, we can examine whether the accuracy of our firm age measure matters for 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 rerun the regressions in Table 3 using founding dates rather than listing dates. Appendix Table 4 Panels E and F show that the results are similar.

We also consider the possibility that managerial age affects innovation. Younger managers may be more prevalent in younger labor markets, and firms may produce more innovation as a result of managerial characteristics rather than the characteristics of the labor force. For example, Acemoglu, Akcigit, and Celik (2017) find a small positive association between executive age and creative corporate innovation. To isolate the effect of the age structure of the labor force from managerial age, we add CEO age as a control variable in our regressions using data from Execucomp. Panels G and H of Appendix Table 4 confirm that our results are robust to controlling for managerial age. Additionally, younger CEOs are associated with more innovation (not tabulated). Our results for age structure are consistent with firms producing more innovation as a result of innovative younger labor forces rather than innovative younger managers.

Finally, we examine whether our results are impacted by complementarity between younger and older people (e.g., between the creativity of the young and the experience of the old). It is possible that innovation is affected not only by the age of the typical worker in the labor force but also by the dispersion of the age of workers. While complementarity could be beneficial, it may also be detrimental, if a more dispersed age structure leads to coordination problems, culture clash, etc. between younger and older workers. We are therefore agnostic about the impact of age structure dispersion on innovation. We add a control variable in our regressions for age structure dispersion, which we measure as the coefficient of variation of age structure (the standard deviation divided by the mean). In Appendix Table 4 Panels I and J, the effect of age structure dispersion on innovation is negative (not tabulated), and the effect of the mean age and young share is similar to our main specification.

4.4. Innovation Characteristics

The premise of the labor supply channel is that younger labor forces are more innovative as a result of the various aforementioned characteristics of younger people: they are more creative, are willing to take more risk, have longer horizons, and are more socially interactive. If this is the case, these characteristics should be reflected in the innovation activities of firms in younger labor markets. We therefore examine the corresponding characteristics of patent outputs: 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 refer to citations made by a patent to previous patents, and forward citations, which refer to 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 requirement that inventions can be patented only if they are useful and novel. We have three 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.[17] The third is the proportion of a firm's patents in both the top 10% of forward citations (i.e., the most useful) and the bottom 10% of backward citations (i.e., the most novel).[18]

Riskiness, our measure of the second characteristic, is the 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 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

[17] The top 1% of patents by citations has roughly nine times as many citations as the average patent in the same grant year cohort.

[18] Patents in the top 10% of forward citations have roughly three times as many citations as the average patent in the same grant year cohort, and the average patent has roughly six times as many citations as patents in the bottom 10% of backward citations in the same grant year cohort. Our results are robust to independently lowering the thresholds for forward and backward citations (e.g., from 10% to 20% or 30%).

firm by the mean of the same variable calculated using all patents in the same 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 grant year cohort, of backward citations of all patents in all commuting zones to patents in the firm's commuting zone.

We construct our measures to reflect the characteristics of younger people that lead them to be more innovative. However, to check that it is only our measures that matter, not all measures based on patent citations, we use two other popular measures. These measures are the originality and generality of patents, and they are defined, following the literature, as the dispersion, respectively, of backward and forward citations across technology classes (see Appendix Table 1 for details). It is not clear whether they should increase or decrease in the age structure of the local labor force. Younger and thus less experienced inventors might choose to concentrate on a few technology fields in great depth, or they might instead diversify across many fields. Their inventions might correspondingly attract attention from the few fields on which they focus, or from a wide variety of fields across which they spread themselves. Even though these variables have ambiguous predicted effects, they are useful as placebos.

We run the same regressions as in Table 3 but using innovation characteristics 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 longevity, originality, and generality because they are symmetrically distributed. 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.

[Insert Table 4 about here]

Table 4 shows that younger labor forces produce innovation outputs that reflect the innovative characteristics of younger labor forces. In both panels, a one standard deviation change in age structure (i.e., a decrease in the mean age or an increase in the young share) causes at least a roughly 6% increase in creativity, a 10% increase in riskiness, a 5% increase in longevity (relative to the standard deviation of the dependent variable), and a 16% increase in interactivity. By contrast, the placebo measures, originality and generality, are not significant, which suggests that only our specific measures reflect the innovative characteristics of younger people, not all measures based on patent citations.

4.5. The Alternative Financing Supply Channel

Our results so far are consistent with the labor supply 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 supply channel in which the age structure measures capture the local investor pool rather than the local labor force. Younger people tend to invest in more risky assets than older people,[19] and investors in general tend to hold relatively more of

[19] 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

their portfolio in local firms (e.g., Coval and Moskowitz (1999)). This implies that younger local investors are more willing to finance the type of risky projects that produce innovation, so firms in younger locations produce more innovation. In this financing supply channel, it is younger investors who cause greater innovation.

To the extent that we can separate the locations of inventors and investors, we can test whether our results are explained by the labor supply or financing supply channel. We achieve this separation by using the firm's R&D hubs to proxy for the location of its inventors and by using the firm's headquarters to proxy for the location of its investors. The tradeoff that we have to make is to use a smaller sample because R&D hubs are identified based on the patent database, so we only have location data for R&D hubs for about half of our full sample of firm-years.

Comparing the number of inventors located at R&D hubs and headquarters in the same commuting zone, we find that the mean and median proportions are both approximately 50%. Assuming that firms without identified R&D hubs undertake their R&D activities at headquarters, the proportion rises to a mean of 75%. Overall, R&D hubs do tend to be located in the same commuting zone as headquarters, but there is meaningful dispersion of R&D hubs compared to headquarters.

[Insert Table 5 about here]

For the sample of firms for which we have location data for R&D hubs, we verify that our baseline results hold. We rerun the regressions in Table 3 for this subsample using the age structure at headquarters. Panels A and B of Table 5 show that the results are similar to our baseline results (Table 3).

---

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)).

We then examine the extent to which our results are driven by the labor supply and financing supply channels. We calculate the weighted average age structure of the firm across its R&D hubs just like we calculate its age structure at its headquarters but using the number of inventors at each R&D hub as weights. We rerun the regressions in Table 3 but include the age structures at both headquarters and R&D hubs.

The results are presented in Panels C and D of Table 5. The age structure at headquarters still affects innovation, but it is economically less significant and not always statistically significant. By comparison, the age structure at R&D hubs is similar in economic and statistical significance to our baseline results (Table 3) and indeed absorbs most of the effect of the age structure at headquarters (Table 5 Panels A and B). The fact that headquarters is still significant alongside 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.

Finally, we test the main prediction of the financing supply 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, firms in younger locations should have lower leverage. In fact, when we rerun the regressions in Table 3 but with leverage as the dependent variable, we find that age structure and leverage are not significantly related, neither economically nor statistically, which is inconsistent with the financing supply channel in our setting. Taken together, our results rule out the pure financing supply channel.

4.6. The Alternative Consumer Demand Channel

It is possible that our results are consistent with an alternative consumer demand channel in which the 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 produces more innovation to satisfy local consumer demand. Given that we use headquarters location to measure age structure, the consumer demand channel is much less plausible than the labor supply channel. For the smallest firms, employees are almost always more concentrated around headquarters than are customers. The operations of the largest firms may well be more dispersed in general, but the employees of such firms are almost always less dispersed than their customers, even when considering the various office locations of a given firm. The labor force at the firm's headquarters is therefore much more likely to capture its employees than its customers.

Nevertheless, to the extent that we can separate firms based on whether their workforce serves local versus non-local consumer demand, we can test whether our results are explained by the labor supply or consumer demand channel. Following Mian and Sufi (2014), our central insight is that production and consumption must coincide in space and time for firms in non-tradable industries, i.e., labor supply and consumer demand are local. However, for tradable industries, production requires specialization and scale and therefore must be local, i.e., labor supply is local, whereas consumption is non-local. Consequently, the local labor force may capture both employees and customers in non-tradable industries, but it can only capture 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). 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 a 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. We rerun the regressions in Table 3. The specifications include the non-tradable industry dummy variable and its interaction with all other variables. Firms in tradable industries are the base group.

[Insert Table 6 about here]

Table 6 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). This is the case for both measures of tradable versus non-tradable industries. By contrast, the effect of age structure on innovation for firms in non-tradable industries is never significantly different from zero, with all of the coefficient estimates being statistically insignificant at conventional levels.[20] In other words, our results are clearly driven by firms in tradable industries, for which the labor force can only capture employees, not customers. On the whole, the results allow us to rule out the pure consumer demand channel.

[20] The level of patent outputs is, on average, very low for firms in non-tradable industries (10% or 25% of our sample firms, depending on the measure). These industries may in fact be less innovative generally, or their innovation activities may not be well captured by patent outputs. Whatever the case may be, the results in Table 6 show that our baseline results, in which innovation activities are measured by patent outputs, are inconsistent with the pure consumer demand channel.

4.7. Firm Composition

The finding that age structure affects innovation is consistent with two slightly different alternative interpretations of the labor supply channel: firms innovating more in younger labor markets taking as given the age structure and other characteristics of the local labor market, and more innovative firms matching to younger labor markets. Both interpretations are consistent with age structure having a causal effect on innovation, and they can mutually coexist. We can better understand the two interpretations based on how the baseline results depend on firm age. The assumption underlying the following analysis is that the youngest firms have the easiest time to choose their location with respect to the age structure of the labor force, to the extent that the age structure affects their innovation activities. At the very least, new startups have zero costs of moving to their first location. In this analysis, matching should generate the strongest results for the youngest firms.

We perform two simple variations on the regressions in Table 3. In the first, we sort firms into four age groups: [0,5), [5,10), [10,20), and [20,∞). We pick these particular cutoffs because they are intuitive, and they generate groups of firms with roughly the same number of observations. The specifications include the first three of four firm age group dummy variables and their interaction with all other variables. The fourth and oldest firm age group is the base group. We continue to measure firm age from the date the firm begins trading publicly in CRSP, and we measure firm location at the time of the firm's final appearance in Compustat.

[Insert Table 7 about here]

Panels A and B of Table 7 show that our baseline results are not driven by the youngest firms. For the youngest firms (the first group), the effect of age structure on innovation is significantly weaker than for the oldest firms, both in economic and statistical terms. In fact, the effect for the youngest firms is generally not significantly different from zero, with two of the eight coefficient estimates being statistically significant at the 10% level and the rest being insignificant at conventional levels. The effect is gradually increasing in strength for the progressively older firms in the second and third groups. For the oldest firms (the fourth and base group), the effect of age structure on innovation is strongest, even stronger than for the average firm in our baseline results (Table 3), by a factor of roughly two.

The second variation is only different from the first in that we only use firms with the same location throughout their lives. We impose this restriction because measuring location at the time of the firm's final appearance in Compustat may lead to mismeasurement of location for firms that move across commuting zones during their lives. To this end, we use firms that are located in the same commuting zone at the times of their IPO and final appearance in Compustat. The tradeoff that we have to make is to use a smaller sample because we only have IPO location data from SDC for about 60% of our full sample of firm-years. Within this sample, 80% of firms do not move during their lives. Additionally, since the IPO location data only begin in 1970, the sample of firms is now younger. To ensure that each group of firms has roughly the same number of observations, we combine the third and fourth firm age groups into a single group: $[10,\infty)$. The base group is now the third and still oldest firm age group.

The results in Panels C and D of Table 7 are similar to the results in Panels A and B. Once again, the effect of age structure on innovation is strongest for the oldest firms. While the magnitude of the results is greater than in our firm-level results (Table 3), it is comparable to our commuting zone-level results (Table 2). The greater magnitude is also consistent with measuring age structure more accurately. Moreover, it is reasonable that the local labor pool may be more important for the oldest firms because they may need it to regularly replenish their human capital. For the youngest firms, by contrast, a stable group of a few core employees comprise most of their human capital, so the local labor pool may be of less importance to them.

5. Inventor-Level Analysis

5.1. Model Specification

Our final analysis is at the inventor level. At this level of analysis, we have thousands of inventors in public and private firms. We can control for inventor- and firm-level characteristics. We can also isolate the effect of the age structure of the labor force from inventor age, firm age, and network scale effects within firms. In this analysis, the equation for the baseline regressions is:

$$Innovation_{i,j,k,c,s,t+1} = \alpha \cdot Age\_Structure_{c,t} + \beta \cdot X_{i,k,t} + \delta_{s,t} + \delta_{j,t} + \delta_i + \varepsilon \quad (3)$$

where i indexes firms, j indexes industries, k indexes inventors, c indexes commuting zones, s indexes states, and t indexes years. $X_{i,k,t}$ is a vector of firm- and inventor-level control variables, $\delta_{s,t}$ is a state-year fixed effect, $\delta_{j,t}$ is an industry-year fixed effect, and δ_i is a firm 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.[21]

We focus on these star inventors for a number of reasons. The top several percent of inventors account for the majority of patents and citations (see Akcigit, Baslandze, and Stantcheva (2016) and Moretti and Wilson (2017)). Moreover, we can determine the location of

[21] Compare Moretti and Wilson (2017), who use the top 5% of inventors, and Akcigit, Baslandze, and Stantcheva (2016), who use the top 1%.

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 year and the next five years. Additionally, star inventors have high moving costs, just like firms. We can confirm in our data that star inventors move at a similar rate to firms changing the location of their headquarters: about 1.5% per annum. This stability alleviates concerns that our results may be driven by more productive inventors choosing younger locations.

Returning to the inventor-level regressions, we measure age structure for inventors using the commuting zone in which they are located. The data are not as rich at the inventor level as at the firm level because the sample includes both public and private firms. However, we are able to control for the age of both the inventor and the firm. Younger inventors and firms are likely to be more innovative, so we control for the age of both. Doing so also allows us to isolate their effect from the age structure of the labor force. We measure inventor age from the date of the inventor's first patent, and analogously for the firm. We do so because we neither have inventor birth dates, nor are firm founding dates widely available. 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 R&D hub as a given inventor, the number of inventors working at any of the firm's R&D hubs, and the firm's number of R&D hubs.

Since the inventors working for a given firm can be located in different commuting zones, we also include firm fixed effects and thus identify off the variation across commuting zones but within firms. Additionally, we use industry-year fixed effects to remove variation across industries and time that is common to both age structure and innovation. Similar to the inventor-level analyses performed in many studies in the literature and as suggested by Lerner and Seru (2017), we use the technology classes of inventors to create industry-year fixed effects. 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. Finally, standard errors are clustered by industry-year, but the results are robust to clustering by inventor.

5.2. Sample and Descriptive Statistics

The inventors in our sample are stars (as defined above), accounting 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 the 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, there is roughly 1 patent per annum on average (median 0.3-0.5) at the inventor-year level. The typical sample inventor has been working for approximately 16 years relative to the date of his first patent. He produces an average of 2.4 patents per year (median of 2.0) (inventor patent stock divided by inventor age).

The last four control variables in Table 1 Panel C are tabulated at the inventor level like the other variables. This is necessary for correctly interpreting our regression results. However, these 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. Rather than repeating the tabulated figures, we report the corresponding aggregated figures. At the firm R&D hub level, there are a mean of 119 inventors (rather than the 599 tabulated). Further aggregating to the firm level, the average firm has been producing patents for 16 years. Additionally, firms have an average of 310 inventors working for them, including both stars and non-stars. The average firm has 20 R&D hubs, although at least half of inventors are located at headquarters.

5.3. Results

[Insert Table 8 about here]

The regression results are presented in Table 8, with the mean age in Panel A and the young share in Panel B. Once again, the results confirm that younger labor forces produce more innovation. In Panel A, a one standard deviation decrease in the mean age (2.0 years) causes an increase in patents of roughly 5%-11%. In Panel B, a one standard deviation increase in the young share (7.8 percentage points) causes a similar 6%-13% increase in patents. Whether for patent counts and citations, or the short run and the long run, the results are similar. Since we control for the ages of inventors and firms, we can rule out inventor and firm life cycle interpretations of our results. We can also rule out firm-specific time-invariant omitted factors because we include firm fixed effects.

Furthermore, two of our control variables provide a nice complement to the age structure of the labor force. In particular, both inventor age and firm age have a negative relationship with innovation. For comparison with the causal effect of age structure, a one standard deviation decrease in inventor age is associated with a 17%-33% increase in patents (e.g., in Column 1, $-0.846\times(-6.1/15.8)=33\%$). For firm age, the corresponding increase in patents is 53%-63% (e.g., in Column 1, $-0.880\times(-18.2/26.7)=60\%$). While these results cannot be interpreted causally, they do suggest that individuals and organizations, much like labor markets, are more innovative when they are younger. More importantly, it suggests that labor markets have both individual and social effects on innovation. In a younger labor market, the average inventor is younger, and he may be more innovative as a result of his youth. This individual effect is captured by the inventor age variable in our regressions.

However, our regressions also 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, the age structure of the local labor force may not even operate primarily through the age of the inventor identified on the patent document, who is a relatively senior employee. Instead, it may operate mainly through the relatively junior graduate students, post-doctoral researchers, scientists, engineers, etc. who work for the inventor, like in the academic setting examined by Bowen, Frésard, and Taillard (2017). Such knowledge spillovers generated by the interactions of inventors are captured by the age structure of the labor force.

6. Value Implications

Finally, returning to the firm level, we examine whether the effect of age structure on innovation is reflected in firm valuations. If the increase in innovation that results from hiring younger workers generates growth opportunities for firms, then firms in younger labor markets should have higher valuations. We use market-to-book of equity to measure valuation. We examine the effect of age structure on valuation in the short run and the long run, measured the next year and the average of the next five years, respectively. We run firm-level regressions for valuation similar to those for innovation (Table 3), but we use market-to-book as the dependent variable rather than patents, and we exclude market-to-book from our control variables.

[Insert Table 9 about here]

Table 9 presents the results, which show that firms in younger labor markets have higher valuations, both in the short run and the long run. A decrease of one standard deviation in the mean age (2.2 years) causes a roughly 3% (=0.015×2.2) increase in valuations (Panel A). An increase of one standard deviation in the young share (8.4 percentage points) causes a similar 3% increase in valuations (=0.35×0.084) (Panel B). As for innovation, so for valuation the results are similar in the short run and the long run. Overall, the results suggest that the growth opportunities generated by younger labor forces are reflected in firm valuations.

We also examine one of the fundamental drivers of firm valuations: firm productivity. If the increase in innovation that results from hiring younger workers improves firm productivity, then firms in younger labor markets should, all else equal, be more productive. We obtain total factor productivity (TFP) data from İmrohoroğlu and Tüzel (2014). They estimate TFP for a sample of publicly traded firms using a Cobb-Douglas production function in which value added is explained by the labor and capital of the firm as well as its productivity. These data only cover about 70% of our full sample of firm-years. All of the variables in the model, including TFP, are measured in natural logarithms, and the model is estimated using industry-year fixed effects. We run firm-level regressions similar to those for innovation (Table 3), but we use productivity as the dependent variable, and we exclude industry-year fixed effects because these are already removed from the TFP estimates.

The results in Appendix Table 5 show that firms in younger labor markets have higher productivity in the short run. In both panels, a one standard deviation change in age structure (i.e., a decrease in the mean age or an increase in the young share) causes a roughly 2% increase in productivity. The results in the long run are similar in magnitude but are no longer statistically significant. Overall, the results provide suggestive evidence that improved firm productivity reflects the greater innovation generated by younger labor forces.

7. Conclusion

We study the effect of age structure of the labor force on corporate innovation. We argue that because of their creativity, risk tolerance, horizons, and interactivity, younger people are instrumental in the production of innovation. We therefore hypothesize that a younger local labor force produces more innovation.

We measure locations using commuting zones and innovation using patent outputs. We reconstruct the labor force using historical births, and use this plausibly exogenous native born labor force to measure the age structure of the labor force. We perform our analysis at three levels, each with their own economic and statistical advantages: commuting zones, firms, and inventors.

At each level of analysis, we consistently find that a younger age structure causes more innovation. Our results are not driven by firm or inventor life cycle effects, nor can they be explained by financing supply or consumer demand effects. Finally, we find that firm valuations reflect the effect of age structure on innovation.

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. These findings support at least three types of public policies that can counter the effects of an aging population: improving the education and training of the native born population; encouraging young and skilled immigration; and incentivizing domestic population growth.

References

Acemoglu, Daron, and Pascual Restrepo, 2018, Demographics and automation, working paper.

Acemoglu, Daron, Ufuk Akcigit, and Murat Alp Celik, 2017, Young, restless and creative: Openness to disruption and creative innovations, working paper.

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.

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, 2017, Employment protection, investment, and firm growth, working paper.

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.

Breschi, Stefano, 2008, Innovation-specific agglomeration economies and the spatial clustering of innovative firms, in Charlie Karlsson, ed.: *Handbook of Research on Innovation and Clusters* (Edward Elgar Publishing, Cheltenham).

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, 2018, Can unemployment insurance spur entrepreneurial activity? Evidence from France, working paper.

Howells, Jeremy, 1990, The location and organisation of research and development: New horizons, *Research Policy* 19, 133-46.

İmrohoroğlu, Ayşe, and Şelale Tüzel, 2014, Firm-level productivity, risk, and return, *Management Science* 60, 2073-2090.

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, 2018, 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.

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, 2017, Demographics and entrepreneurship, forthcoming *Journal of Political Economy*.

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.

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.

Trajtenberg, Manuel, Rebecca Henderson, and Adam Jaffe, 1997, University versus corporate patents: A window on the basicness of invention, *Economics of Innovation and New Technology* 5, 19-50.

Table 1
Descriptive Statistics

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

Panel A: Commuting Zone-Level Sample					
	Mean	Standard deviation	25th percentile	Median	75th percentile
Independent variables					
- Projected mean age (years)	41.6	2.3	40.3	41.9	43.1
- Projected young share (%)	44.6	8.5	38.7	43.4	49.4
- Actual mean age (years)	40.7	1.7	39.6	40.7	41.8
- Actual young share (%)	47.8	6.6	43.5	48.1	52.3
Dependent variables					
- Patent counts-to-population: Next year	10.3	15.5	0.9	4.8	12.2
- Patent counts-to-population: Next 5 years	11.0	16.3	1.9	5.0	12.7
- Patent citations-to-population: Next year	12.0	20.2	0.0	4.3	13.7
- Patent citations-to-population: Next 5 years	12.9	20.7	1.4	5.1	14.6
Control variables					
- Population size (thousands)	330.0	692.3	35.4	99.1	281.7
- Income per capita ($ thousands)	26.2	5.3	22.6	25.3	28.9
- Growth rate of total income (%)	1.3	3.4	-0.6	1.3	3.3
- Government spending-to-total income (%)	10.3	3.6	7.8	9.4	11.8
- Educational attainment (%)	16.4	5.9	12.1	15.0	19.4
- University patent counts-to-population	0.4	1.1	0.0	0.0	0.0

Panel B: Firm-Level Sample					
	Mean	Standard deviation	25th percentile	Median	75th percentile
Independent variables					
- Projected mean age (years)	38.8	2.2	37.3	38.9	40.4
- Projected young share (%)	54.7	8.4	48.5	54.6	59.8
- Actual mean age (years)	39.6	1.1	38.9	39.5	40.3
- Actual young share (%)	52.3	4.7	49.1	52.6	55.4
Dependent variables					
- Patent counts: Next year	4.5	18.9	0.0	0.0	1.0
- Patent counts: Next 5 years	4.8	19.6	0.0	0.0	1.0
- Patent citations: Next year	5.7	23.5	0.0	0.0	0.5
- Patent citations: Next 5 years	6.0	24.4	0.0	0.0	1.0
- Citations per patent	1.4	1.6	0.5	0.9	1.7
- Proportion of extremely useful patents	1.8	7.4	0.0	0.0	0.0
- Proportion of useful and novel patents	0.2	0.8	0.0	0.0	0.0
- Volatility of citations per patent	1.2	1.3	0.4	0.8	1.6
- Longevity of citations per patent	1.0	0.3	0.9	1.0	1.1
- Proportion of local citations	5.2	15.5	0.0	1.3	3.7
- Originality (%)	46.3	18.2	36.0	47.3	58.5
- Generality (%)	40.5	21.0	25.0	43.0	56.1
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

Panel C: Inventor-Level Sample					
	Mean	Standard deviation	25th percentile	Median	75th percentile
Independent variables					
- Projected mean age (years)	39.4	2.0	38.0	39.4	41.1
- Projected young share (%)	52.2	7.8	45.9	52.4	58.1
- Actual mean age (years)	40.0	1.1	39.2	40.0	40.8
- Actual young share (%)	50.4	4.7	46.6	50.1	53.8
Dependent variables					
- Patent counts: Next year	1.00	1.39	0.13	0.50	1.25
- Patent counts: Next 5 years	0.86	1.05	0.24	0.51	1.03
- Patent citations: Next year	1.26	2.67	0.00	0.31	1.19
- Patent citations: Next 5 years	1.07	2.10	0.10	0.34	1.02
Control variables					
- Inventor age (years)	15.8	6.1	11.0	15.0	20.0
- Inventor patent stock	38	25	22	30	44
- R&D hub number of inventors	599	895	28	165	708
- Firm age (years)	26.7	18.2	12.0	25.0	35.0
- Firm number of inventors	2,368	4,108	133	770	2,831
- Firm number of R&D hubs	65	59	16	52	95

Table 2
The Effect of Age Structure on Innovation: Baseline Commuting Zone-Level Analysis

This table shows the results of regressions of innovation on age structure. The regressions follow Equation 1. The unit of observation is the commuting zone-quinquennial period. The sample and specifications are described in the text. Age structure is measured for the labor force projected based on historical births. 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 Mean Age								
	Dependent variable is ln(1+patents/population per annum)							
	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.112***	-0.108***	-0.111***	-0.133***	-0.121***	-0.112***	-0.124***	-0.110***
	(-3.60)	(-3.80)	(-4.31)	(-4.72)	(-5.18)	(-4.89)	(-5.30)	(-4.94)
ln(Population size)	0.014	0.008	-0.027	-0.021	-0.038	-0.066**	-0.024	-0.049
	(0.34)	(0.20)	(-0.85)	(-0.71)	(-1.32)	(-2.37)	(-0.75)	(-1.57)
ln(Income per capita)	2.488***	2.106***	0.780	1.892***	0.733	0.793	0.815	0.758
	(4.47)	(3.78)	(1.38)	(4.22)	(1.27)	(1.41)	(1.41)	(1.39)
Growth rate of total income	0.257	-0.336	-1.337	-1.441	-2.314*	-1.884	-1.852	-1.511
	(0.11)	(-0.16)	(-0.77)	(-0.88)	(-1.68)	(-1.47)	(-1.29)	(-1.18)
Government spending-to-total income		-8.566**			-7.759***	-7.907***	-8.071***	-8.405***
		(-2.50)			(-2.71)	(-2.75)	(-2.67)	(-2.84)
Educational attainment			7.292***		4.414***	4.682***	5.261***	5.694***
			(6.61)		(4.29)	(4.67)	(4.92)	(5.59)
ln(1+University patent counts/population)				0.606***	0.382***	0.390***	0.379***	0.379***
				(8.95)	(6.01)	(6.18)	(5.25)	(5.49)
State-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,748	2,748	2,748	2,748	2,748	2,748	2,748	2,748
Adjusted R^2	0.625	0.640	0.669	0.673	0.695	0.718	0.693	0.717

Panel B: Age Structure Measured Using Young Share								
	Dependent variable is ln(1+patents/population per annum)							
	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	3.024***	2.910***	2.992***	3.549***	3.230***	2.991***	3.232***	2.905***
	(3.77)	(3.92)	(4.47)	(4.91)	(5.29)	(5.01)	(5.25)	(4.98)
ln(Population size)	0.017	0.011	-0.024	-0.016	-0.034	-0.063**	-0.019	-0.044
	(0.42)	(0.28)	(-0.74)	(-0.55)	(-1.19)	(-2.25)	(-0.59)	(-1.45)
ln(Income per capita)	2.491***	2.111***	0.783	1.899***	0.738	0.797	0.823	0.764
	(4.49)	(3.79)	(1.38)	(4.25)	(1.28)	(1.41)	(1.41)	(1.39)
Growth rate of total income	0.268	-0.319	-1.324	-1.416	-2.289	-1.863	-1.815	-1.485
	(0.11)	(-0.15)	(-0.76)	(-0.86)	(-1.65)	(-1.44)	(-1.25)	(-1.15)
Government spending-to-total income		-8.526**			-7.720***	-7.870***	-8.039***	-8.372***
		(-2.49)			(-2.69)	(-2.73)	(-2.65)	(-2.83)
Educational attainment			7.289***		4.433***	4.699***	5.289***	5.714***
			(6.60)		(4.28)	(4.65)	(4.91)	(5.57)
ln(1+University patent counts/population)				0.604***	0.379***	0.387***	0.375***	0.376***
				(8.95)	(5.94)	(6.13)	(5.16)	(5.43)
State-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,748	2,748	2,748	2,748	2,748	2,748	2,748	2,748
Adjusted R^2	0.625	0.640	0.669	0.672	0.695	0.718	0.692	0.717

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 follow Equation 2. The unit of observation is the firm-quinquennial period. The sample and specifications are described in the text. Age structure is measured for the labor force projected based on historical births. 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 Mean Age				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.030***	-0.033***	-0.040***	-0.045***
	(-2.99)	(-3.33)	(-3.62)	(-4.32)
ln(Total assets)	0.243***	0.258***	0.259***	0.273***
	(10.86)	(11.30)	(11.01)	(11.49)
Market-to-book	0.019***	0.023***	0.022***	0.028***
	(7.44)	(7.96)	(7.92)	(8.47)
Cash flow-to-total assets	-0.239***	-0.224***	-0.260***	-0.222***
	(-5.32)	(-4.31)	(-5.40)	(-4.27)
Stock returns	0.033**	0.061***	0.043***	0.066***
	(2.16)	(3.95)	(2.66)	(3.92)
Stock return volatility	0.112***	0.114***	0.136***	0.125***
	(3.23)	(3.23)	(3.67)	(3.18)
Industry-year fixed effects?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	14,086	14,086	14,086	14,086
Adjusted R^2	0.429	0.454	0.374	0.402

Panel B: Age Structure Measured Using Young Share				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	0.833***	0.915***	1.132***	1.260***
	(3.04)	(3.33)	(3.64)	(4.26)
ln(Total assets)	0.243***	0.258***	0.259***	0.273***
	(10.87)	(11.31)	(11.02)	(11.51)
Market-to-book	0.019***	0.023***	0.022***	0.028***
	(7.45)	(7.97)	(7.93)	(8.49)
Cash flow-to-total assets	-0.240***	-0.224***	-0.261***	-0.223***
	(-5.34)	(-4.32)	(-5.41)	(-4.28)
Stock returns	0.033**	0.061***	0.043***	0.066***
	(2.18)	(3.97)	(2.68)	(3.95)
Stock return volatility	0.111***	0.114***	0.135***	0.125***
	(3.23)	(3.23)	(3.67)	(3.18)
Industry-year fixed effects?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	14,086	14,086	14,086	14,086
Adjusted R^2	0.429	0.454	0.375	0.403

Table 4
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 follow Equation 2. The unit of observation is the firm-quinquennial period. The sample and specifications are described in the text. Age structure is measured for the labor force projected based on historical births. 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 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	Other #1: Originality	Other #2: Generality
Age structure	-0.026**	-0.032***	-0.028**	-0.046***	-0.008**	-0.072**	-0.285	-0.070
	(-2.00)	(-3.01)	(-2.19)	(-2.75)	(-2.33)	(-2.41)	(-1.38)	(-0.27)
ln(Total assets)	0.128***	0.080***	0.112***	0.218***	0.007	0.194***	0.039	-0.230
	(7.12)	(6.74)	(11.87)	(9.64)	(1.24)	(8.86)	(0.24)	(-1.18)
Market-to-book	0.016***	0.011***	0.009***	0.015***	0.002*	0.007	-0.009	-0.037
	(4.38)	(3.35)	(3.68)	(3.48)	(1.77)	(1.21)	(-0.19)	(-0.62)
Cash flow-to-total assets	-0.170*	-0.027	0.025	-0.002	-0.024	-0.364**	-2.102	-1.587
	(-1.71)	(-0.43)	(0.50)	(-0.02)	(-0.97)	(-2.25)	(-1.40)	(-1.05)
Stock returns	-0.027	0.003	-0.001	0.012	0.001	0.162***	-0.005	-0.682*
	(-0.77)	(0.11)	(-0.04)	(0.39)	(0.14)	(4.02)	(-0.01)	(-1.83)
Stock return volatility	0.068	0.086	0.149***	0.052	0.017	-0.117	-0.337	1.574
	(1.06)	(1.58)	(4.26)	(0.79)	(0.73)	(-1.16)	(-0.23)	(1.16)
Industry-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,768	5,768	5,768	4,361	4,365	5,768	5,754	5,454
Adjusted R^2	0.228	0.112	0.098	0.197	0.039	0.121	0.101	0.384

Panel B: 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	Other #1: Originality	Other #2: Generality
Age structure	0.727**	0.881***	0.708**	1.178***	0.192**	1.869**	8.176	2.141
	(2.18)	(3.06)	(2.11)	(2.78)	(2.16)	(2.32)	(1.53)	(0.32)
ln(Total assets)	0.128***	0.080***	0.112***	0.218***	0.007	0.195***	0.038	-0.231
	(7.11)	(6.74)	(11.85)	(9.65)	(1.24)	(8.86)	(0.23)	(-1.19)
Market-to-book	0.016***	0.011***	0.009***	0.015***	0.002*	0.007	-0.009	-0.037
	(4.38)	(3.35)	(3.67)	(3.47)	(1.77)	(1.21)	(-0.19)	(-0.62)
Cash flow-to-total assets	-0.171*	-0.028	0.025	-0.003	-0.024	-0.366**	-2.113	-1.590
	(-1.72)	(-0.45)	(0.48)	(-0.03)	(-0.97)	(-2.26)	(-1.40)	(-1.06)
Stock returns	-0.027	0.003	-0.000	0.012	0.001	0.162***	-0.004	-0.681*
	(-0.76)	(0.11)	(-0.03)	(0.40)	(0.15)	(4.03)	(-0.01)	(-1.83)
Stock return volatility	0.068	0.086	0.150***	0.052	0.017	-0.117	-0.341	1.571
	(1.05)	(1.57)	(4.27)	(0.80)	(0.74)	(-1.15)	(-0.23)	(1.16)
Industry-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,768	5,768	5,768	4,361	4,365	5,768	5,754	5,454
Adjusted R^2	0.228	0.112	0.098	0.197	0.039	0.121	0.101	0.384

Table 5
The Effect of Age Structure on Innovation: Firm-Level Analysis for Firms with R&D Hubs

This table shows the results of regressions of innovation on age structure. The regressions follow Equation 2. The unit of observation is the firm-quinquennial period. The regressions are the same as in Table 3, but the sample is restricted to firms with R&D hubs. The specifications include control variables and fixed effects for industry-year, state-year, and firm age. Panels A and B use only the age structure at the firm's headquarters. Panels C and D use the age structures at both the firm's headquarters and its R&D hubs. Age structure is measured for the labor force projected based on historical births. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Headquarters Only: Age Structure Measured Using Mean Age

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure at headquarters	-0.026*	-0.029**	-0.046***	-0.051***
	(-1.93)	(-2.44)	(-3.06)	(-4.04)
Observations	6,259	6,259	6,259	6,259
Adjusted R^2	0.347	0.384	0.284	0.335

Panel B: Headquarters Only: Age Structure Measured Using Young Share

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure at headquarters	0.732**	0.819***	1.296***	1.415***
	(2.07)	(2.62)	(3.17)	(4.17)
Observations	6,259	6,259	6,259	6,259
Adjusted R^2	0.347	0.384	0.285	0.336

Panel C: Both Headquarters and R&D Hubs: Age Structure Measured Using Mean Age

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure at headquarters	-0.017	-0.021*	-0.029*	-0.034**
	(-1.27)	(-1.66)	(-1.81)	(-2.57)
Age structure at R&D hubs	-0.025**	-0.024*	-0.052***	-0.048***
	(-2.05)	(-1.95)	(-3.77)	(-3.47)
Observations	6,259	6,259	6,259	6,259
Adjusted R^2	0.347	0.385	0.286	0.337

Panel D: Both Headquarters and R&D Hubs: Age Structure Measured Using Young Share

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure at headquarters	0.486	0.587*	0.811*	0.962***
	(1.37)	(1.80)	(1.90)	(2.67)
Age structure at R&D hubs	0.740**	0.697**	1.455***	1.359***
	(2.24)	(2.06)	(3.93)	(3.56)
Observations	6,259	6,259	6,259	6,259
Adjusted R^2	0.348	0.385	0.287	0.338

Table 6
The Effect of Age Structure on Innovation: Firm-Level Analysis 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 follow Equation 2. The unit of observation is the firm-quinquennial period. 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 all other variables. The specifications also include control variables and fixed effects for industry-year, state-year, and firm age. In Panels A and B, 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 Panels C and D, 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. Age structure is measured for the labor force projected based on historical births. The regressions include a consumer orientation dummy variable and its interaction with age structure. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Measure #1 of Industry Type: Age Structure Measured Using Mean Age

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.032**	-0.031**	-0.047***	-0.047***
	(-2.33)	(-2.45)	(-3.19)	(-3.52)
Age structure × Non-tradable industry dummy variable	0.035**	0.036**	0.056***	0.059***
	(2.12)	(2.45)	(2.94)	(3.54)
Observations	9,056	9,056	9,056	9,056
Adjusted R^2	0.498	0.531	0.439	0.475

Panel B: Measure #1 of Industry Type: Age Structure Measured Using Young Share

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	0.928**	0.905**	1.348***	1.376***
	(2.43)	(2.53)	(3.28)	(3.59)
Age structure × Non-tradable industry dummy variable	-1.012**	-1.060**	-1.591***	-1.695***
	(-2.27)	(-2.58)	(-3.10)	(-3.67)
Observations	9,056	9,056	9,056	9,056
Adjusted R^2	0.498	0.531	0.440	0.476

Panel C: Measure #2 of Industry Type: Age Structure Measured Using Mean Age				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.044*	-0.037*	-0.059**	-0.050**
	(-1.93)	(-1.74)	(-2.38)	(-2.19)
Age structure × Non-tradable industry dummy variable	0.036	0.026	0.055**	0.040
	(1.48)	(1.14)	(2.10)	(1.63)
Observations	7,546	7,546	7,546	7,546
Adjusted R^2	0.499	0.534	0.455	0.492

Panel D: Measure #2 of Industry Type: Age Structure Measured Using Young Share				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	1.329**	1.130*	1.786***	1.532**
	(2.16)	(1.93)	(2.65)	(2.40)
Age structure × Non-tradable industry dummy variable	-1.121*	-0.866	-1.708**	-1.316*
	(-1.74)	(-1.39)	(-2.40)	(-1.93)
Observations	7,546	7,546	7,546	7,546
Adjusted R^2	0.499	0.534	0.455	0.492

Table 7
The Effect of Age Structure on Innovation Conditional Upon Firm Age: Firm-Level Analysis

This table shows the results of regressions of innovation on age structure conditional upon firm age. The regressions follow Equation 2. The unit of observation is the firm-quinquennial period. The regressions are the same as in Table 3, but the specifications include the firm age group dummy variables and their interaction with all other variables. In Panels A and B, firm location is measured at the time of the firm's final appearance in Compustat (IPO). In Panels C and D, firm location is the same at the times of IPO and final appearance in Compustat. In Panels A and B (C and D), the first three (two) firm age groups are used in interactions, and the fourth (third) and oldest firm age group is the base group. The sample and specifications are described in the text. Age structure is measured for the labor force projected based on historical births. Variables are defined in Appendix Table 1. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Location at Final Appearance: Age Structure Measured Using Mean Age

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.059***	-0.055**	-0.083***	-0.076***
	(-2.73)	(-2.59)	(-3.59)	(-3.27)
Age structure × Firm age ∈ [0,5) years dummy variable	0.048**	0.036*	0.070***	0.048**
	(2.27)	(1.70)	(2.95)	(2.01)
Age structure × Firm age ∈ [5,10) years dummy variable	0.033	0.018	0.042*	0.029
	(1.42)	(0.77)	(1.70)	(1.12)
Age structure × Firm age ∈ [10,20) years dummy variable	0.005	0.005	0.017	0.013
	(0.19)	(0.22)	(0.64)	(0.48)
Control variables?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	13,909	13,909	13,909	13,909
Adjusted R^2	0.465	0.484	0.399	0.420

Panel B: Location at Final Appearance: Age Structure Measured Using Young Share

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	1.716***	1.598***	2.389***	2.177***
	(2.85)	(2.71)	(3.73)	(3.39)
Age structure × Firm age ∈ [0,5) years dummy variable	-1.419**	-1.100*	-2.035***	-1.453**
	(-2.50)	(-1.94)	(-3.21)	(-2.27)
Age structure × Firm age ∈ [5,10) years dummy variable	-1.003	-0.577	-1.257*	-0.854
	(-1.60)	(-0.90)	(-1.91)	(-1.22)
Age structure × Firm age ∈ [10,20) years dummy variable	-0.193	-0.223	-0.551	-0.444
	(-0.28)	(-0.33)	(-0.74)	(-0.59)
Control variables?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	13,909	13,909	13,909	13,909
Adjusted R^2	0.465	0.484	0.399	0.421

Panel C: Same Location at Initial and Final Appearance: Age Structure Measured Using Mean Age				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.116***	-0.102***	-0.163***	-0.142***
	(-4.96)	(-3.92)	(-5.56)	(-4.91)
Age structure × Firm age ∈ [0,5) years dummy variable	0.103***	0.076***	0.150***	0.105***
	(3.83)	(3.00)	(4.17)	(3.52)
Age structure × Firm age ∈ [5,10) years dummy variable	0.087***	0.056	0.105***	0.074*
	(2.75)	(1.63)	(2.88)	(1.80)
Control variables?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	5,959	5,959	5,959	5,959
Adjusted R^2	0.371	0.400	0.316	0.362

Panel D: Same Location at Initial and Final Appearance: Age Structure Measured Using Young Share				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	3.290***	2.872***	4.597***	4.012***
	(5.01)	(4.08)	(5.65)	(5.08)
Age structure × Firm age ∈ [0,5) years dummy variable	-2.978***	-2.247***	-4.302***	-3.063***
	(-4.01)	(-3.24)	(-4.35)	(-3.82)
Age structure × Firm age ∈ [5,10) years dummy variable	-2.568***	-1.679*	-3.078***	-2.213**
	(-3.03)	(-1.86)	(-3.14)	(-2.05)
Control variables?	Yes	Yes	Yes	Yes
Industry-year fixed effects?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Firm age fixed effects?	Yes	Yes	Yes	Yes
Observations	5,959	5,959	5,959	5,959
Adjusted R^2	0.371	0.400	0.316	0.362

Table 8
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 follow Equation 3. The unit of observation is the inventor-quinquennial period. The sample and specifications are described in the text. Age structure is measured for the labor force projected based on historical births. 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 Mean Age				
	Dependent variable is ln(0.01+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.047**	-0.027**	-0.055**	-0.033**
	(-2.33)	(-2.40)	(-2.55)	(-2.06)
ln(Inventor age)	-0.846***	-0.439***	-0.830***	-0.514***
	(-14.22)	(-12.55)	(-14.03)	(-10.10)
ln(Inventor patent stock)	1.226***	0.712***	1.280***	0.869***
	(23.79)	(22.41)	(20.91)	(17.69)
ln(R&D hub number of inventors)	0.055***	0.046***	0.049**	0.040***
	(2.99)	(4.50)	(2.53)	(2.80)
ln(Firm age)	-0.880***	-0.780***	-0.929***	-0.889***
	(-6.44)	(-12.60)	(-6.95)	(-9.47)
ln(Firm number of inventors)	0.216	0.157	0.305	0.233*
	(1.10)	(1.61)	(1.51)	(1.76)
ln(Firm number of R&D hubs)	0.199	0.041	-0.074	-0.086
	(0.78)	(0.31)	(-0.28)	(-0.47)
Industry-year fixed effects?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Firm fixed effects?	Yes	Yes	Yes	Yes
Observations	14,078	14,078	14,078	14,078
Adjusted R^2	0.210	0.295	0.262	0.337

Panel B: Age Structure Measured Using Young Share				
	Dependent variable is ln(0.01+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	1.353**	0.792**	1.606***	0.978**
	(2.46)	(2.51)	(2.70)	(2.23)
ln(Inventor age)	-0.846***	-0.439***	-0.830***	-0.515***
	(-14.22)	(-12.55)	(-14.03)	(-10.11)
ln(Inventor patent stock)	1.226***	0.712***	1.280***	0.870***
	(23.79)	(22.40)	(20.91)	(17.69)
ln(R&D hub number of inventors)	0.055***	0.046***	0.049**	0.040***
	(2.97)	(4.48)	(2.50)	(2.78)
ln(Firm age)	-0.881***	-0.781***	-0.930***	-0.890***
	(-6.44)	(-12.60)	(-6.96)	(-9.48)
ln(Firm number of inventors)	0.215	0.157	0.305	0.233*
	(1.10)	(1.61)	(1.51)	(1.76)
ln(Firm number of R&D hubs)	0.199	0.041	-0.074	-0.085
	(0.79)	(0.31)	(-0.28)	(-0.47)
Industry-year fixed effects?	Yes	Yes	Yes	Yes
State-year fixed effects?	Yes	Yes	Yes	Yes
Firm fixed effects?	Yes	Yes	Yes	Yes
Observations	14,078	14,078	14,078	14,078
Adjusted R^2	0.210	0.295	0.262	0.337

Table 9
The Effect of Age Structure on Valuation: Firm-Level Analysis

This table shows the results of regressions of valuation on age structure. The unit of observation is the firm-quinquennial period. The sample and specifications are described in the text. Age structure is measured for the labor force projected based on historical births. 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 Mean Age		
	Dependent variable is ln(Market-to-book)	
	Next year	Average of next 5 years
Age structure	-0.016***	-0.015***
	(-2.89)	(-2.76)
ln(Total assets)	-0.007	0.007*
	(-1.47)	(1.68)
Cash flow-to-total assets	-0.604***	-0.739***
	(-12.95)	(-19.42)
Stock returns	0.261***	0.172***
	(16.49)	(11.94)
Stock return volatility	-0.172***	-0.076***
	(-7.01)	(-3.40)
Industry-year fixed effects?	Yes	Yes
State-year fixed effects?	Yes	Yes
Firm age fixed effects?	Yes	Yes
Observations	13,617	13,887
Adjusted R^2	0.213	0.181

Panel B: Age Structure Measured Using Young Share		
	Dependent variable is ln(Market-to-book)	
	Next year	Average of next 5 years
Age structure	0.410***	0.355**
	(2.70)	(2.44)
ln(Total assets)	-0.007	0.007*
	(-1.47)	(1.69)
Cash flow-to-total assets	-0.605***	-0.740***
	(-12.96)	(-19.44)
Stock returns	0.261***	0.172***
	(16.51)	(11.95)
Stock return volatility	-0.172***	-0.075***
	(-7.01)	(-3.39)
Industry-year fixed effects?	Yes	Yes
State-year fixed effects?	Yes	Yes
Firm age fixed effects?	Yes	Yes
Observations	13,617	13,887
Adjusted R^2	0.213	0.181

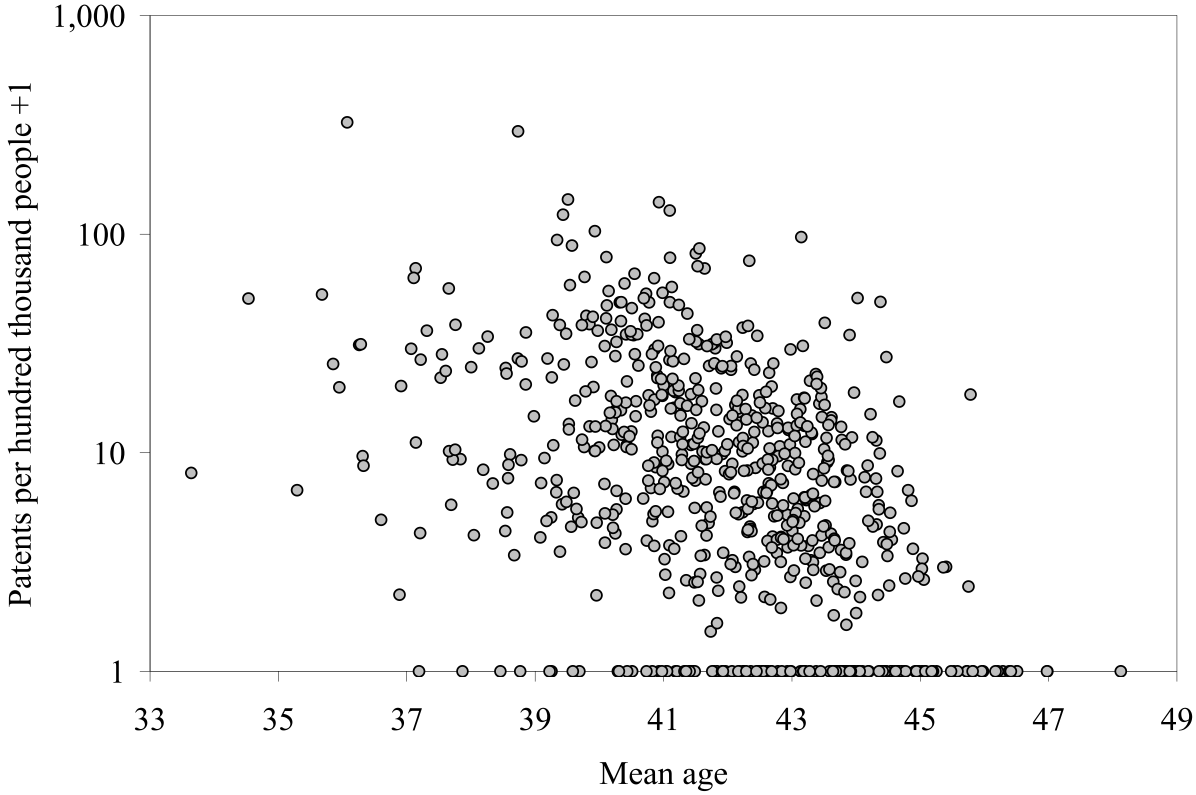

Figure 1. The effect of age structure on innovation. This figure shows the ratio of patent counts to population as a function of the age structure of the projected labor force for all commuting zones in the U.S. in the year 2000. The population of the commuting zone is measured in hundred thousands. The projected labor force refers to the labor force projected based on historical births.

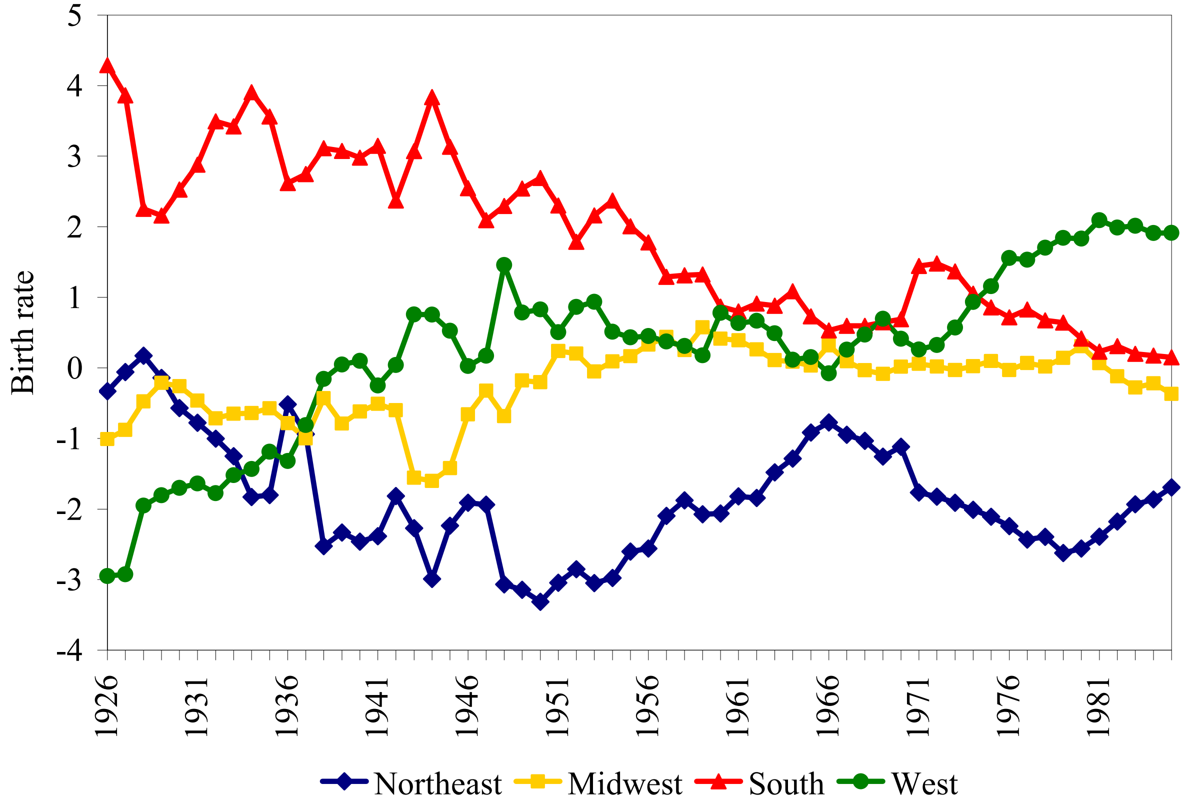

Figure 2. Evolution of the historical births used to construct the projected age structure of the labor force by location. This figure shows the evolution of the historical births from 1926, or 64 years before the beginning of the sample period in 1990, to 1985, or 20 years before the end of the sample period in 2005. The census regions shown in the figure span the country. The birth rate is measured per thousand people and detrended.

Panel A: Projected Labor Force

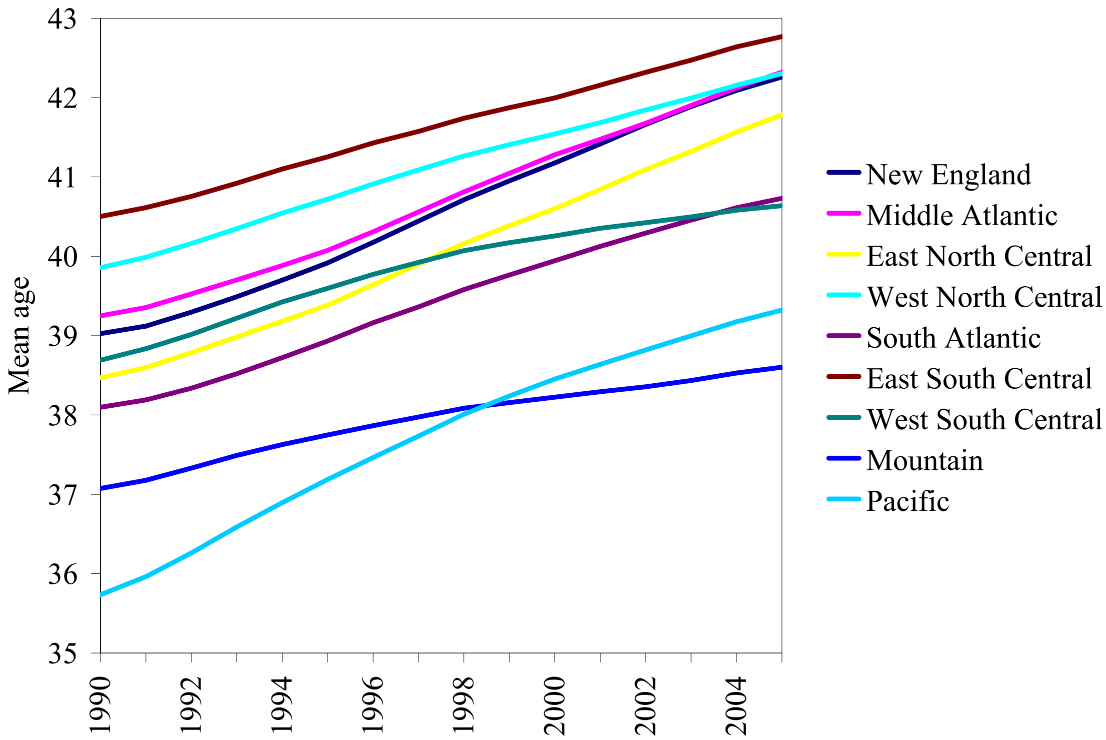

Panel B: Actual Labor Force

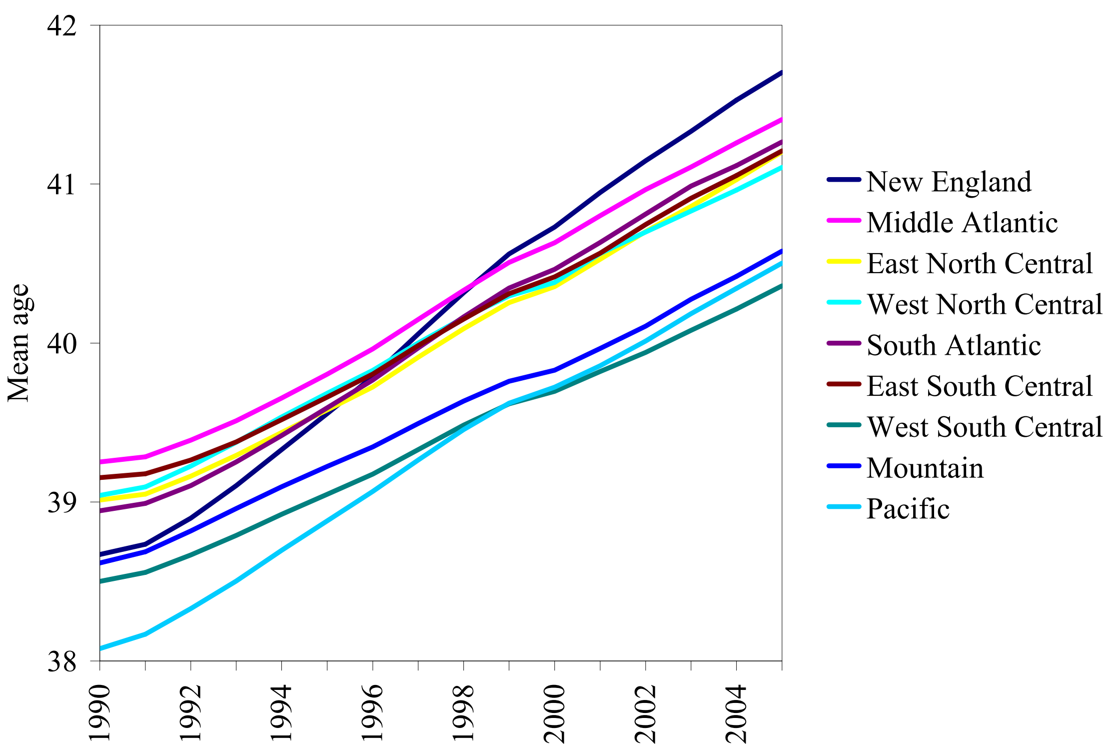

Figure 3. Evolution of the mean age of the projected and actual labor forces by location. This figure shows the mean age of the projected and actual labor forces during the sample period (1990-2005). The census divisions shown span the country. The projected labor force refers to the labor force projected based on historical births.

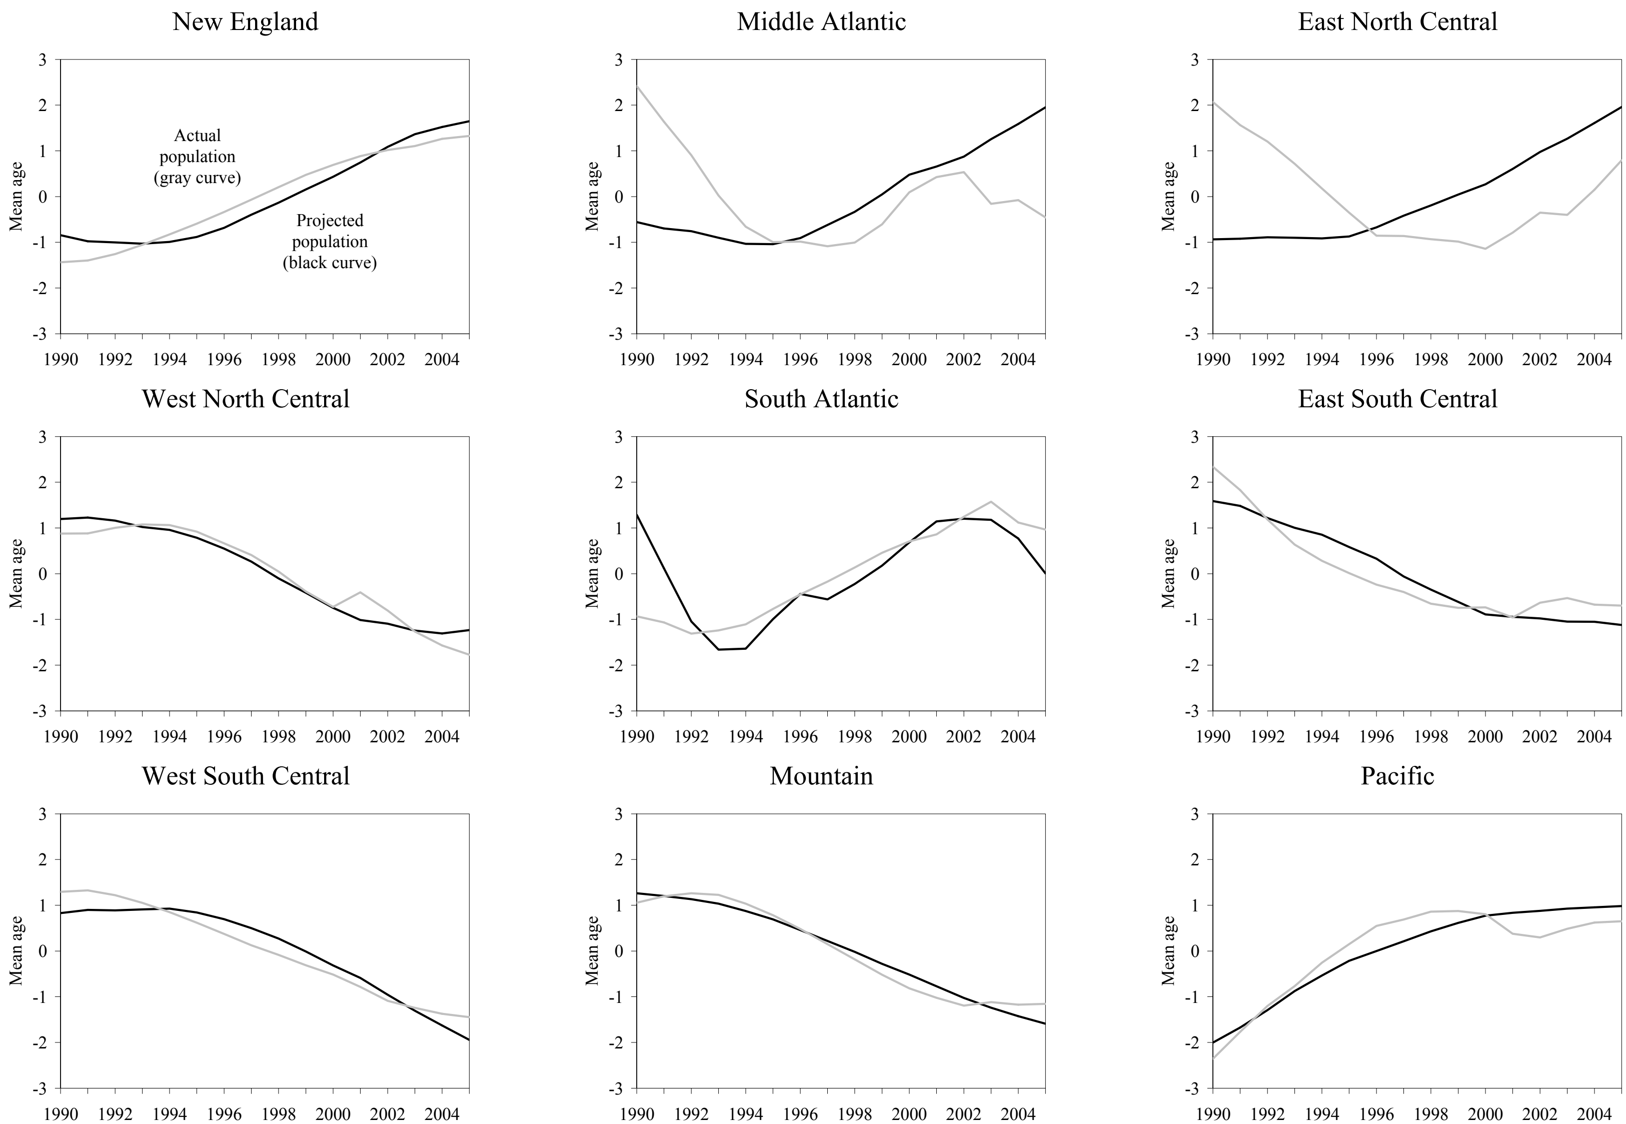

Figure 4. Evolution of the mean age of the projected and actual labor forces by location. This figure shows the mean age of the projected and actual labor forces during the sample period (1990-2005). The census divisions shown span the country. The mean age is first detrended and then standardized within each location. The projected labor force refers to the labor force projected based on historical births.

Appendix Table 1
Variable Definitions

Name	Definition
Demographic Variables Common to All Regressions	
- Projected mean age	Mean age of the labor force (ages 20-64) in a commuting zone, projected based on historical births and adjusted for survival
- Projected young share	The young share (ages 20-39) of the labor force (ages 20-64) in a commuting zone, projected based on historical births and adjusted for survival
- Actual mean age	Mean age of the labor force (ages 20-64) in a commuting zone
- Actual young share	The young share (ages 20-39) of the labor force (ages 20-64) in a commuting zone
Commuting Zone-Level Regressions	
Innovation variables	
- Patent counts-to-population	The number of patents of all firms in a commuting zone, adjusted for truncation (see Hall, Jaffe, and Trajtenberg (2005) and Kogan, Papanikolaou, Seru, and Stoffman (2017)). Scaled by the population of the commuting zone measured in hundred thousands.
- Patent citations-to-population	The weighted number of patent citations of all firms in a commuting zone. Patent citations in a given year are weighted by the average number of citations per patent in the same year (for details, see Hall, Jaffe, and Trajtenberg (2005) and Kogan, Papanikolaou, Seru, and Stoffman (2017)). Scaled by the population of the commuting zone measured in hundred thousands.
Control variables	
- Population size	The population of a commuting zone
- Income per capita	The income per capita of a commuting zone
- Growth rate of total income	The growth rate of the total income 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's degree or higher to the population aged 25 years or older in a commuting zone
- University patent counts-to-population	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 constructed as described above
- Patent citations	The weighted number of patent citations of a firm constructed as described above
- 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 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 grant year cohort.
- Proportion of useful and novel patents	The proportion of a firm's patents in both the top 10% of forward citations and the bottom 10% of backward citations in the same patent year cohort. Citations are ranked relative to patents in the same 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 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 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 grant year cohort.
- Originality and generality	Mean dispersion of citations across technology classes. Dispersion of citations for a patent is measured as one minus the Herfindahl index of citations. (See Trajtenberg, Henderson, and Jaffe (1997).) Citations are backward citations for originality and forward citations for generality.
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 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 each patent with N inventors, each inventor is credited with 1/N patent citations.
Control variables	
- Inventor age	The inventor's age measured from the date of his first patent
- Inventor patent stock	The number of patents of the inventor
- R&D hub number of inventors	The number of inventors working for the firm in the same commuting zone as a given inventor
- Firm age	The firm's age measured from the date of the firm's first patent
- Firm number of inventors	The number of inventors working for the firm in any commuting zone
- Firm number of R&D hubs	The number of commuting zones in which there is at least one inventor working for the firm

Appendix Table 2
Replication of Baseline Commuting Zone-Level Results Using Actual Age Structure

This table shows the replication of the results in Table 2 Panel A. Panel A of the current table uses the mean age of the actual labor force instead of the labor force projected based on historical births. Panel B of the current table uses the mean age of both the projected and actual labor forces. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Actual Labor Force Only: Age Structure Measured Using Mean Age								
	Dependent variable is ln(1+patents/population per annum)							
	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.058**	-0.055**	0.128***	0.033	0.132***	0.135***	0.126***	0.124***
	(-2.00)	(-2.03)	(2.75)	(1.03)	(3.00)	(3.23)	(2.83)	(2.97)
ln(Population size)	0.038	0.031	0.071*	0.055	0.066**	0.035	0.079**	0.048
	(0.75)	(0.66)	(1.88)	(1.43)	(1.99)	(1.12)	(2.20)	(1.45)
ln(Income per capita)	2.770***	2.368***	0.330	2.086***	0.223	0.259	0.338	0.272
	(5.09)	(4.17)	(0.44)	(4.41)	(0.30)	(0.36)	(0.45)	(0.39)
Growth rate of total income	0.435	-0.174	-0.481	-0.588	-1.375	-0.983	-0.917	-0.647
	(0.20)	(-0.09)	(-0.35)	(-0.40)	(-1.16)	(-0.84)	(-0.69)	(-0.48)
Government spending-to-total income		-8.783**			-8.151***	-8.271***	-8.469***	-8.760***
		(-2.43)			(-2.70)	(-2.75)	(-2.65)	(-2.83)
Educational attainment			9.187***		6.743***	7.018***	7.513***	7.877***
			(5.75)		(4.33)	(4.70)	(4.74)	(5.26)
ln(1+University patent counts/population)				0.590***	0.330***	0.342***	0.325***	0.332***
				(8.19)	(5.15)	(5.44)	(4.41)	(4.74)
State-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,748	2,748	2,748	2,748	2,748	2,748	2,748	2,748
Adjusted R^2	0.612	0.628	0.661	0.653	0.685	0.711	0.685	0.711

Panel B: Both Projected and Actual Labor Forces: Age Structure Measured Using Mean Age								
	Dependent variable is ln(1+patents/population per annum)							
	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Projected age structure	-0.108***	-0.105***	-0.131***	-0.144***	-0.143***	-0.134***	-0.145***	-0.130***
	(-3.28)	(-3.46)	(-4.86)	(-4.75)	(-5.86)	(-5.59)	(-5.87)	(-5.46)
Actual age structure	-0.032	-0.031	0.171***	0.078**	0.179***	0.178***	0.173***	0.167***
	(-1.04)	(-1.03)	(3.57)	(2.21)	(4.07)	(4.24)	(3.89)	(3.96)
ln(Population size)	0.001	-0.005	0.028	0.008	0.020	-0.008	0.032	0.006
	(0.01)	(-0.11)	(0.85)	(0.24)	(0.67)	(-0.30)	(0.99)	(0.19)
ln(Income per capita)	2.545***	2.161***	-0.102	1.699***	-0.176	-0.115	-0.067	-0.092
	(4.33)	(3.69)	(-0.14)	(3.45)	(-0.26)	(-0.17)	(-0.10)	(-0.14)
Growth rate of total income	0.083	-0.499	-0.967	-1.186	-1.946*	-1.517	-1.495	-1.167
	(0.04)	(-0.23)	(-0.66)	(-0.75)	(-1.69)	(-1.38)	(-1.19)	(-0.98)
Government spending-to-total income		-8.548**			-7.730***	-7.878***	-8.043***	-8.377***
		(-2.49)			(-2.86)	(-2.90)	(-2.80)	(-2.98)
Educational attainment			9.784***		6.896***	7.162***	7.668***	8.017***
			(6.34)		(5.00)	(5.39)	(5.46)	(6.02)
ln(1+University patent counts/population)				0.664***	0.399***	0.407***	0.395***	0.396***
				(8.40)	(6.74)	(6.94)	(5.74)	(6.02)
State-year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,748	2,748	2,748	2,748	2,748	2,748	2,748	2,748
Adjusted R^2	0.625	0.640	0.680	0.676	0.708	0.731	0.703	0.727

Appendix Table 3
Replication of Baseline Firm-Level Results Using Annual Frequency and Fixed Effects for Commuting Zones or Firms

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. The four quinquennial periods from 1990 to 2005 are replaced with every year from 1990 to 2005. Panels A and B of this table add commuting zone fixed effects to Table 3. Panels C and D add firm fixed effects to Table 3. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Commuting Zone Fixed Effects: Age Structure Measured Using Mean Age

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.035**	-0.036***	-0.029*	-0.029**
	(-2.48)	(-2.77)	(-1.89)	(-1.99)
Observations	58,504	58,504	58,504	58,504
Adjusted R^2	0.441	0.473	0.394	0.430

Panel B: Commuting Zone Fixed Effects: Age Structure Measured Using Young Share

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	1.222***	1.118***	1.055***	1.017***
	(3.87)	(3.86)	(3.01)	(3.09)
Observations	58,504	58,504	58,504	58,504
Adjusted R^2	0.441	0.473	0.394	0.430

Panel C: Firm Fixed Effects: Age Structure Measured Using Mean Age

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.019	-0.019**	-0.017	-0.018*
	(-1.50)	(-2.20)	(-1.20)	(-1.74)
Observations	57,385	57,385	57,385	57,385
Adjusted R^2	0.858	0.938	0.822	0.928

Panel D: Firm Fixed Effects: Age Structure Measured Using Young Share

	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	0.707**	0.562***	0.576	0.523**
	(2.35)	(2.60)	(1.62)	(2.00)
Observations	57,385	57,385	57,385	57,385
Adjusted R^2	0.858	0.939	0.822	0.928

Appendix Table 4
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 Commuting Zone-Level Control Variables: Age Structure Measured Using Mean Age				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.031***	-0.035***	-0.040***	-0.044***
	(-3.34)	(-3.73)	(-3.85)	(-4.75)
Observations	14,086	14,086	14,086	14,086
Adjusted R^2	0.435	0.462	0.384	0.414

Panel B: Adding Commuting Zone-Level Control Variables: Age Structure Measured Using Young Share				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	0.871***	0.953***	1.099***	1.206***
	(3.43)	(3.78)	(3.89)	(4.74)
Observations	14,086	14,086	14,086	14,086
Adjusted R^2	0.435	0.462	0.384	0.414

Panel C: Dropping the 50 Most Populous Commuting Zones: Age Structure Measured Using Mean Age				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.031*	-0.029*	-0.033**	-0.028*
	(-1.93)	(-1.92)	(-2.07)	(-1.73)
Observations	2,398	2,398	2,398	2,398
Adjusted R^2	0.428	0.461	0.379	0.408

Panel D: Dropping the 50 Most Populous Commuting Zones: Age Structure Measured Using Young Share				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	0.877**	0.872**	0.921**	0.816*
	(2.06)	(2.12)	(2.15)	(1.94)
Observations	2,398	2,398	2,398	2,398
Adjusted R^2	0.428	0.461	0.379	0.408

Panel E: Firm Age Measured from Founding Date: Age Structure Measured Using Mean Age				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.028***	-0.031***	-0.040***	-0.044***
	(-2.78)	(-2.95)	(-3.62)	(-4.11)
Observations	11,558	11,558	11,558	11,558
Adjusted R^2	0.423	0.448	0.375	0.406

Panel F: Firm Age Measured from Founding Date: Age Structure Measured Using Young Share				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	0.780***	0.854***	1.103***	1.235***
	(2.83)	(2.97)	(3.61)	(4.07)
Observations	11,558	11,558	11,558	11,558
Adjusted R^2	0.423	0.448	0.376	0.406

Panel G: Adding Managerial Age Control Variable: Age Structure Measured Using Mean Age				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.047***	-0.048**	-0.079***	-0.077***
	(-2.63)	(-2.50)	(-3.82)	(-3.89)
Observations	3,591	3,591	3,591	3,591
Adjusted R^2	0.530	0.554	0.487	0.516

Panel H: Adding Managerial Age Control Variable: Age Structure Measured Using Young Share				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	1.282**	1.313**	2.205***	2.144***
	(2.58)	(2.47)	(3.92)	(3.94)
Observations	3,591	3,591	3,591	3,591
Adjusted R^2	0.530	0.554	0.488	0.516

Panel I: Adding Age Structure Dispersion Control Variable: Age Structure Measured Using Mean Age				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	-0.044***	-0.046***	-0.059***	-0.061***
	(-3.99)	(-4.45)	(-4.78)	(-5.59)
Observations	14,086	14,086	14,086	14,086
Adjusted R^2	0.429	0.454	0.375	0.403

Panel J: Adding Age Structure Dispersion Control Variable: Age Structure Measured Using Young Share				
	Dependent variable is ln(1+patents per annum)			
	Patent counts: Next year	Patent counts: Next 5 years	Patent citations: Next year	Patent citations: Next 5 years
Age structure	1.208***	1.271***	1.627***	1.678***
	(4.01)	(4.40)	(4.73)	(5.39)
Observations	14,086	14,086	14,086	14,086
Adjusted R^2	0.429	0.454	0.375	0.403

Appendix Table 5
The Effect of Age Structure on Productivity: Firm-Level Analysis

This table shows the results of regressions of productivity on age structure. The unit of observation is the firm-quinquennial period. The sample and specifications are described in the text. Age structure is measured for the labor force projected based on historical births. 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 Mean Age		
	Dependent variable is total factor productivity	
	Next year	Average of next 5 years
Age structure	-0.008*	-0.007
	(-1.65)	(-1.51)
Observations	9,083	10,199
Adjusted R^2	0.331	0.238

Panel B: Age Structure Measured Using Young Share		
	Dependent variable is total factor productivity	
	Next year	Average of next 5 years
Age structure	0.207*	0.174
	(1.71)	(1.43)
Observations	9,083	10,199
Adjusted R^2	0.331	0.238