Commuting and Innovation: Are Closer Inventors More Productive?*

Hong Yu Xiao† Andy Wu‡

November 15, 2019

Abstract: Commuting is costly for employees, but is it costly for employers in terms of lost productivity? We examine the direction and size of the causal effects of commuting distance on inventor productivity. We construct a novel panel of U.S. inventors with precisely measured workplace-home distances and a direct measure of productivity via patents. Our identification strategy relies upon within-city firm office relocation events as exogenous shocks to commuting distance. We find a significant negative causal effect from commuting distance on inventor productivity: every ten kilometer increase in distance is associated with a 5% decrease in patent counts per inventor-firm pair per year, and an even greater 7% decrease in patent quality. This effect is economically significant for firms, costing the latter more than \$4,000 per year with conservative estimates of patent values.

JEL Classification Codes: J24, O30, R30, R41

Keywords: Commuting Distance, Inventor Productivity, Firm Relocations, Matched Firm-Inventor Location Data, Patents

---

*Gilles Duranton and Jessie Handbury provided invaluable guidance throughout this project. Fernando Ferreira, Joseph Gyourko, Joseph Harrington, Jean-François Houde, Benjamin Keys, Benjamin Lockwood, Corinne Low, Ryan Peters, Lindsay Relihan, Todd Sinai, and Maisy Wong provided valuable suggestions. The authors acknowledge financial support from the Mack Institute for Innovation Management at the Wharton School of the University of Pennsylvania and the HBS Division of Research and Faculty Development. All opinions expressed herein are those of the authors only; all errors are the responsibility of the authors.

†**Hong Yu Xiao**: The Wharton School of the University of Pennsylvania, Business Economics and Public Policy Department, 3000 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104, USA. E-mail: `hongxiao@wharton.upenn.edu`. Corresponding author.

‡**Andy Wu**: Harvard Business School, Harvard University, Strategy Unit, 15 Harvard Way, Morgan Hall 243, Boston, MA 02163. E-mail: `awu@hbs.edu`.

1 Introduction

Commuting can be costly to both workers and firms. In 2014, 139 million U.S. workers made daily commutes to work, averaging 26 minutes each way (Ingraham, 2016).[1] Total opportunity cost of commuting for workers can exceed their hourly wages (Van Ommeren and Fosgerau, 2009), amounting to thousands of dollars per average worker per year (Perino, 2019), and this before taking into account potential costs on workers' subjective well-being (Kahneman and Krueger, 2006). For firms, the daily commute represents a cost imposed on workers that may affect their effort and productivity. This cost is increasingly salient to employers of skilled knowledge workers, whose time and incremental productivity are especially valuable. As anecdotal evidence, major technology firms are taking notice and are developing strategies to incentivize their employees to live closer to the workplace.[2] Despite this, little is known in the literature about the existence and size of the potential costs of commuting on productivity. We seek to provide one of the first estimates of these costs, focusing on inventors who patent for their employers.

Theoretically, the causal effect of commuting on overall inventor productivity is unclear. On the one hand, there are multiple channels through which commuting can negatively affect productivity. First, longer commutes could simply force inventors to spend less time at work, by imposing schedule constraints such as needing to catch the last commuter train home (Bloom, Kretschmer and Van Reenen, 2009). This would decrease their overall output even if their hourly productivity remained the same. Second, keeping time spent at work fixed, longer commutes could increase inventors' cost of providing effort, and decrease inventor productivity given previous literature showing that more leisure time and slacking off at work are substitutes to each other (Shapiro and Stiglitz, 1984; Zenou, 2002; Ross and Zenou, 2008; Zenou, 2009; Van Ommeren and Gutiérrez-i Puigarnau, 2011). Third, longer commutes

[1] Assumes fifty work weeks in a year, with five work days per week and a round-trip commute each day.

[2] In 2015, Facebook offered employees working at its Silicon Valley headquarters over $10,000 to move closer to the office and avoid the lengthy and time-consuming commute between San Francisco to Menlo Park, CA (Reuters, 2015). Other technology businesses, such as Google, have begun building proximate housing for employees.

may have negative impacts on inventor life quality, which could affect their individual productivity.[3] On the other hand, there could be mitigating health effects, since commuting could reduce stress for people who prefer to separate their work and family lives (Ashforth, Kreiner and Fugate, 2000). A longer commute could provide a good amount of time for brainstorming alone, which has been consistently found to be more effective than group brainstorming (Furnham, 2000). Furthermore, firms could also mitigate the negative effect of longer commutes on effort provision by paying a compensating wage to incentivize effort (Mulalic, Van Ommeren and Pilegaard, 2014). Therefore, estimating the causal effects of commuting on individual-level inventor productivity is an empirical question. Answering this question is especially important for inventors, who are more likely to suffer from long commutes because innovation is geographically concentrated in large cities, and because innovation is important for longer-term firm and economic growth.

There are historically two main empirical challenges in studying the impact of commuting distance on individual-level productivity. The first challenge is data availability, both because productivity is usually not directly measured in administrative data, and because we need to know precisely the address information of both the inventor and the firm he or she works for in order to compute commuting costs. The second empirical challenge is endogenous location choices made by both inventors and firms. Inventors endogenously choose their place of residence based on a long list of factors in addition to commuting costs, such as the quality of local public services, size and price of homes available for sale, accessibility to amenities such as restaurants and movie theaters, etc. Factors that firms consider in their office location decisions include office rent and the neighborhood's productive amenities, in addition to accessibility. We address both empirical challenges.

We solve the empirical data challenge by first using patents as a direct measure of inventor productivity. Patents can be a good measure for inventor productivity because

[3] Kahneman and Krueger (2006) find in a survey that the morning commute came last in terms of generating positive emotions, below even work itself. Extensive public health surveys link commuting with many negative health outcomes, including obesity, high cholesterol, high blood pressure, and depression (e.g., Crabtree, 2010).

they are worth millions of dollars each for firms (Pakes, 1985; Kogan et al., 2017) and produce significant rents for inventors themselves (Toivanen and Väänänen, 2012; Kline et al., 2019). Patents also contain rich information that allow us to compute their scientific and economic value (Pakes, 1986; Hall, Jaffe and Trajtenberg, 2001) individually, which allow us to differentiate between high-worth and average patents in our estimation. In short, patents provide a measure of inventor productivity that is meaningful both scientifically and economically.[4]

In addition to providing a novel measure of individual-level productivity, we leverage a novel combination of several data sources to measure commuting distance for every inventor. While patents do contain information on inventor and establishment locations at the city level, we need precise residential and firm addresses to compute commuting distances. We solve this data challenge by merging patent data with housing transactions for inventor residential locations, and with comprehensive proprietary firm establishment location data obtained from a marketing company. We do so by matching inventor names to homeowner names, and assignee names to firm names. In summary, we construct a unique firm-inventor-year panel of U.S. inventors, with information on their precise productivity and commuting distances.

To solve the second empirical challenge of endogenous location choice by firms and inventors, we exploit workplace relocation events that exogenously shock inventor commuting costs in a stacked difference-in-differences design. We also control for firm-level endogenous location choices with firm-location fixed effects and time-varying firm-level controls. In short, we compare inventors who get shocked farther away to their firm to others who get shocked closer in, all working at the same pre- and post- office location. We thus focus on inventors who retained the same home location and continued to work for the same company before and after a workplace relocation, relying on the fact that many inventors do not re-optimize by moving their residence or changing their job after a workplace relocation due to labor and

[4]Productivity here is defined to be the total output of an inventor over a time period such as a year, which is different from hourly productivity. We believe that our annual measure is the more relevant metric for both firms and inventors, in part because it's unlikely for inventors to be paid using hourly rates.

housing market frictions.[5]

We find a 5% decrease in inventor productivity measured in raw annual patent counts for every 10 km increase in commuting distance, or 0.041 patents per year. For patent quality measures such as average scaled citations, the effect is even larger with a 7% decrease in inventor productivity for every 10 km increase in commuting distance. These results are robust to a variety of commuting cost measures and model specifications, are not driven by any one metropolitan area with particularly long commutes such as the San Francisco Bay Area, and show no evidence of differential pre-trends between inventors who got opposite-direction shocks to their new office location. Furthermore, the measured average effect in our sample is likely a lower bound for the population-level effect since the selected sample of inventors who stay with their firm after office relocation is more likely to have lower commuting costs, and because of attenuation bias due to random measurement errors in the matching process. Taking a conservative lower-bound on the value of an average patent as documented in the prior literature (e.g. Pakes, 1985; Bessen, 2008), this causal effect translates to at least $4000 (in 2010 dollars) in value lost for the employer per year, and could be an order of magnitude larger. This productivity loss is economically significant given that the average pre-relocation commuting distance for inventors in our data is around 20 km, and at least comparable to estimates of the total commuting cost borne by inventors.[6]

Our results are of interest to urban and economic geography scholars. First, we empirically verify the predictions of urban efficiency wage models (Zenou, 2002; Ross and Zenou, 2008; Zenou, 2009) where longer commutes lower worker productivity, unless compensated enough by a sufficiently high wage differential. This effect occurs because longer commutes

[5] For example, Teradyne Inc., a major high-tech producer of electronic component test equipment, moved its headquarters from Boston, MA to North Reading, MA in 2006; for the 10 inventors we identify as working with the company both before and after the relocation, all of them experienced the general impact of the workplace relocation, but half saw their workplace-home distance increase after the workplace relocation, while the other half experienced a decrease in that distance. This within-firm within-inventor but *across-time* variation in workplace-home distance is at the core of our identification strategy.

[6] For a back-of-the-envelope calculation, assume that total commuting costs borne by inventors are equal to their hourly wage for their time spent commuting. They commute 1 hour on average per workday, work 40 hours per week, with a $100,000 salary. Then their annual total commuting costs are $12,500. This number would decrease if the average salary for inventors is less than $100,000.

reduce leisure time at home, thus increasing the cost of effort provision at work and incentivizing workers to shirk. The size of the negative effect suggests that the direct commuting costs for employers are similar in magnitude to the commuting costs for workers modeled in the canonical monocentric city models (Alonso, 1964; Mills, 1967; Muth, 1969; Duranton and Puga, 2015), even after taking into account potential wage compensation (Mulalic, Van Ommeren and Pilegaard, 2014). This effect also suggests that if population density were kept the same, then larger cities should have lower productivity than smaller cities because their average commuting time is higher. This theoretical effect is not observed empirically (Combes et al., 2010) due to mechanisms that produce a positive relationship between city size and productivity, such as a larger labor pool providing better matching opportunities between firms and workers (Helsley and Strange, 1990; Lagos, 2000), worker sorting between cities with higher-ability workers sorting into larger cities (Davis and Dingel, 2014), and knowledge spillovers within a larger population (Jovanovic and Rob, 1989; Glaeser, 1999; Duranton and Puga, 2001). Nevertheless, our results imply that any analysis of agglomeration effects must consider commuting, since the association between higher population density and higher wages could in part be due to higher productivity from less commuting effects.

We contribute to the vast empirical literature evaluating the costs of commuting (Zax, 1991; Van Ommeren and Fosgerau, 2009) and the recent growing literature that uses firm relocations as exogenous shocks to worker commuting costs to estimate plausibly causal effects of commuting on worker outcomes.[7] Prior research in this literature has shown suggestive evidence that longer commutes may lower employee productivity, despite the existence of potential wage compensation, as found by Mulalic, Van Ommeren and Pilegaard (2014) using firm relocations in Denmark. Van Ommeren and Gutiérrez-i Puigarnau (2011) show that an increase in commuting distance is associated with increasing absenteeism from work, and Zax and Kain (1996) show evidence of workers moving closer to the employer's new office location or quitting their job altogether after the employer relocates, suggesting that commuting

[7] For example, Lorenz and Goerke (2016) investigate and find no causal effect of commuting on body weight.

negatively affects productivity.

In addition, this study contributes to a recent literature highlighting the importance of understanding inventor productivity as a consequence of the their spatial organization. Given the increasing importance of innovation for firm performance, we need to understand the role that geography plays in inventor productivity. Recent work continues on a long literature highlighting the importance of spatial proximity (e.g., Breschi and Lenzi, 2016; Carlino and Kerr, 2015). We focus specifically on understanding the understudied dimension of inventor-firm proximity, in contrast to inter-inventor and inter-firm proximity addressed in prior work. Bernstein, McQuade and Townsend (2017) study how regional declines in housing markets hurt effect inventor productivity. Much like their study, we exploit a dataset that links inventor patenting output with their housing records. Other work highlights how geographic variation in governmental policy decisions effects the innovativeness of inventors and where inventors choose to live (e.g., Moretti and Wilson, 2017; Glaeser and Hausman, 2019).

Our results are relevant for firms and policy makers due to the significant impacts innovation has on long-term firm-level outcomes and economic growth. For firms, our findings suggest that they should consider the commuting distance of their employees in their planning decisions. The size of the negative effect of commuting on productivity might also help firms better calculate the benefits and costs of subsidizing or outright building proximate housing for their employees. The benefit of shorter commutes could be especially large for firms in the technology sector dependent on skilled labor to generate intellectual property. For public policy makers, our findings support the importance of considering density for urban planning policy. While workers incur direct time and monetary costs from commuting, we show that commuting potentially imposes a further indirect cost on worker productivity, borne both by the worker and their employer. While Duranton and Turner (2017) find that urban density has only a small effect on total vehicle miles driven, reduced workplace-home distance can still generate welfare improvements through increases in worker productivity.

The rest of the paper is structured as follows. First, we first describe the construction of the inventor-firm-year panel data sample central to this paper. Second, we present our empirical strategy leveraging workplace relocations for causual identification, motivated with a stylized model and discussion of the key endogeneity considerations. Third, we present the empirical findings. We conclude with a discussion of potential future work. ‘

2 Data

Leveraging a novel combination of several data sources, we construct a unique firm-inventor-year panel for U.S. firms and inventors between 1997 and 2012. The data contains precise locations of both the workplace and home of inventors, allowing us to accurately construct various measures of workplace-home distance for each inventor. Moreover, our setting of inventors lends itself directly to measuring individual-level contributions to firm productivity, through measures of patenting output that are linked to both the inventors and their employing firms. This data then allows us to exploit within-city relocations of the firm offices, serving as exogenous shocks to the workplace-home distance for each of a firm's inventors. We now describe the data sources, matching methodology across these data sources, and variable construction from these data.

2.1 Direct Measure of Inventor Productivity: Patents

To identify the firms and the inventors for our sample, as well as construct a measure of individual-level contributions to firm productivity, we start with the universe of utility patents granted by the U.S. Patent Office (USPTO) between 1975 and 2012. Li et al. (2014) provide a disambiguated patent database identifying unique inventors across patents and distinguishing inventors with identical or similar names. Each patent contains the names of the inventors, the firm (i.e., assignee) that owns the patent and most likely employs the inventor(s), and the home city and state of each inventor. The city and state directly provided by the USPTO is insufficiently precise for the analysis we seek to conduct, so we use this information to obtain detailed home address data from another dataset.[8]

We attribute each patent as the output of the listed firm and inventor based on its application year, which is the year when the invention has been initially filed at the USPTO. We only consider granted patents, to ensure that claimed inventions satisfy a minimum quality threshold. However, given the time lag between patent application filings and USPTO decisions on whether to grant patents, the patent database is necessarily incomplete in the years leading up to 2012, since some patent applications had not yet been granted. Assuming that the USPTO's idiosyncratic time lag is invariant within-industry, this sampling consideration should impact inventors working in the same business establishment equally and not bias our estimation results. Between 1997 and 2012, XXX,XXX ultimately granted patents were filed by XXX,XXX inventors living in the 60 CSAs covered in our paper.

Given that all patents we include in our data are eventually granted, they must have passed a minimum threshold of quality as determined by the USPTO.[9] Therefore, we use raw annual patent counts by patent application year as our main measure of inventor productivity. Nevertheless, granted patents do differ in importance and quality (Hall, Jaffe and Trajtenberg, 2005). To adjust for quality, we adopt standard measures used in the innovation literature (Hall, Jaffe and Trajtenberg, 2001; Bessen, 2008). The first patent quality measure is scaled citations, which counts the number of forward citations a patent receives, scaled by year and patent category fixed effects to control for mechanical differences in propensity to cite.[10] Two other measures of patent scientific quality are generality and originality scores, introduced by Trajtenberg, Henderson and Jaffe (1997). Generality scores measure whether a patent is cited by subsequent patents from a wide range of technological categories, and originality scores

[8]We focus on utility patents, which account for more than 90% of all patents granted by the USPTO.

[9]In an investigation of patent applications filed between 1996 and 2005, Carley, Hedge and Marco (2015) find that around 55% of all patent applications are eventually granted, suggesting that granted patents do satisfy some minimum quality threshold.

[10]Citations are a measure of a patent's scientific quality since legally, the cited patent both forms part of existing knowledge that the citing patent builds upon, and constitutes "prior art" that can potentially limit the applicability of claims in the citing patent.

measure whether a patent cites many prior patents from different technological categories.[11] Finally, patents' private value to its owners may be different than its scientific value, so we use the expected number of times a patent is renewed by paying maintenance fees as a measure of its economic value. Because the fee more than doubles for each subsequent renewal, we can plausibly assume that the economic value of a patent is monotonically and positively related to the number of maintenance fee payments (further details provided in the Appendix).

2.2 Data Sources and the Matching Process

To identify the residential home location of inventors in our sample, we use detailed housing transactions data from DataQuick, a leading supplier of real estate data and analytics, to obtain the street address of the inventors.[12] The data covers 60 combined statistical areas (CSA) in 23 states from 1993 to 2012 and includes more than 195 million housing transactions and refinances.[13] For each transaction, we observe both the exact address of each home bought/sold, and the full name of the home buyers and sellers. In addition, we obtain limited demographic information, such as income, by matching transactions against loan application data in the Home Mortgage Disclosure Act (HMDA) files, as described in Ferreira and Gyourko (2015).

To identify the workplace location of inventors in our sample, we use historical business establishment location data from InfoUSA, containing the street addresses of offices of all firms in the US between 1997 and 2012. InfoUSA aggregates business location data from

[11] Mathematically, generality for patent i is:

$$Generality_i = 1 - \sum_{j}^{n_i} s_{ij}^2 \tag{1}$$

where s_{ij} is the share of citations received by patent i that belong to patent category j, out of n_j patent categories. Originality is defined similarly, except it uses citations made by patent i to patent categories j.

[12] DataQuick was acquired by CoreLogic in 2014.

[13] We use combined statistical areas (CSAs), instead of the metropolitan statistical areas (MSAs) that make up each CSA, because CSAs better reflect the possible intra-region commuting flows and economic ties. See Ferreira and Gyourko (2015) for additional information about the construction of the DataQuick sample. The Appendix provides further information about the geographic coverage of the sample.

various public sources, including yellow pages, credit card billing data, company annual reports, etc. InfoUSA also verifies information via phone calls and web research every year (DiNardo and Lee, 2004). The variables associated with each establishment include firm name, street address, NAICS code, employee count, and sales volume. To verify the accuracy of InfoUSA data, we compare it to the County Business Patterns (CBP) data by state, and we find the totals to be quite similar (further detail provided in the Appendix). In 2005, InfoUSA had data on 5.63 million business establishments across the 60 CSAs covered by DataQuick.

To construct a panel of matched firm-inventor pairs, we start with the inventor and firm information in the patent data. We first match the USPTO inventor names and locations (city and state) with the DataQuick housing transactions data to obtain inventor home street addresses. We then match firm names and locations against business location data. The matching process is described in detail in the paragraphs below.

To obtain inventor home addresses, we first match buyer names exactly against inventor names from the same city. We then match inventor names to seller names in the subsequent transaction to obtain ownership years for each home buyer. To restrict our sample to owner-occupiers, we exclude cases where people with the same names own different addresses in the same city, because we cannot identify their main residence, or whether they are different people with the same name. In other words, we identify homeowners with unique names within a city.[14] Overall, we are able to match around 264,000 inventors, or 47% of all inventors, to their exact home addresses.

We then manually match USPTO firm names against business establishment names in the business location data to obtain inventor workplace addresses. We obtain precise office locations for 36,468 firms that had applied for patents between 1997 to 2012 with matched

[14] While homeowner names are unique within our cleaned dataset, there could be multiple same-name individuals living in the same city who are not in our dataset. This measurement error potentially attenuates our estimates towards zero. As a robustness check, we estimate an alternative specification by weighting each inventor-firm pair inversely proportional to the probability of another person having the same name in the same city. The results are consistent with our main results. See the Appendix for more details.

inventors. To identify an inventor's precise work location if the firm he works for has multiple establishments within the CSA, we select the most likely location based on whether it has by far the largest number of employees,[15] and whether it has a "research laboratory" designation in the corresponding NAICS codes. We drop observations where it was impossible to uniquely identify a main office location within the CSA. This narrows our dataset down to 35,836 single-location firms.

As we will describe later in our empirical strategy section, we use firm relocations as exogenous shocks to commuting distances. In service of this empirical strategy, we identify within-CSA business relocations where the main firm establishment location within the CSA changes from one year to the next. To improve the power of our estimates, we limit our sample to firms making substantial moves of more than one kilometer. We then identify inventors who worked (i.e., patented) for the relocating firm both before and after its relocation. Finally, we eliminate relocations that occurred in 1997 or after 2010, so we have data both before and after the relocation. We also exclude outliers that account for 3% of our total number of observations.[16] (Further details on the matching process in the Appendix)

Our final sample consists of 22,917 inventor-firm-year observations, representing 3,445 unique inventors employed at 1,068 relocating firms.

2.3 Measures of Inventor Workplace-Home Distance

After obtaining the panel of matched inventor-firm-location pairs, we construct a number of commuting cost measures between an inventor's workplace and home. Our primary and most parsimonious measure of workplace-home distance is ***Distance***, which is the geodesic distance, the shortest path between two points on the curved surface of the Earth.[17] This measure of

[15] We only retain establishments that have 5 times more employees than all other establishments of the same firm in the CSA combined

[16] The outliers are inventor-firm pairs who satisfy one of the following five criteria at any one year in our sample: distance between home and workplace is greater than 100km, driving distance is greater than 125km, change in distance is greater than 57km, change in driving distance is greater than 55km, or received more than 10 patents.

[17] We use Vincenty (1975) equations for a mathematical model of the earth.

physical separation between two points is invariant over time (at least for a human lifespan).

We also construct measures of commuting cost based upon the assumption that the inventor might be driving or taking public transit to work, which is presumed to be the case for a large portion of our sample. ***Drive Distance*** is the shortest route for a motor vehicle, i.e. via roads that are legal to drive on, between the inventor's home and workplace. ***Drive Duration*** is the estimated fastest time it takes to drive or take public transit between the inventor's home and workplace, accounting for speed limits and historical traffic conditions. Both of these measures were collected from the Distance Matrix API of the Google Maps platform, which provides travel distance and time for a matrix of origins and destinations (Google LLC, 2018). Due to data limitations, driving-based measures and these other mode-based measures are only based on current transportation infrastructure and does not account for changes in the transportation infrastructure (e.g., new road construction) during the time window of this study. We use geodesic distance as our main measure of commuting cost because it is always fixed and correct over time. Nevertheless, all our results carry over to using ***Drive Distance*** or ***Drive Duration*** as our commuting cost measure.

2.4 Descriptive Statistics

Figure 1 shows that the distribution of workplace-home distance skews towards shorter commutes, and a modal distance around 10 km with substantially fewer observations at greater distances. The majority of inventors in our sample have a commuting distance of less than 18km. This distribution is consistent with studies by the U.S. Department of Transportation (e.g., U.S. Government Bureau of Transportation Statistics, 2003), which suggests that the matching process generated sensible workplaces-homes matches for the inventors. (See more detailed comparisons in the Appendix.) [NEEDS TO CHECK THIS USING NHTS DATA]

——————**Insert Figure 1**——————

Table 3 shows the distribution of our observations between CSAs, focusing on the six largest CSAs by total population. Roughly a third of all inventor-firm pairs in our sample are from the San Jose-San Francisco Bay (CA) area, representing Silicon Valley and the presumed center of the U.S. technology industry. This proportion is similar to the overall proportion of inventors in the Bay Area, showing that our matched inventors are similarly distributed geographically as all U.S. inventors.

——————Insert Table 3——————

Table 1 shows summary statistics for our final sample at the inventor-firm pair level, taking the average over each inventor-firm pair over the sample period. Overall, the average inventor produces 0.524 patents per year cited 0.893 times, well above the median inventor at 0.333 patents per year, cited 0.288 times. This suggests that patent productivity in our sample is heavily skewed with a long tail of very productive inventors, consistent with prior literature (Akcigit, Baslandze and Stantcheva, 2016).

Given that we use firm relocations as exogenous shocks to commuting distances, we then compare pre- and post-relocation distributions of inventor productivity and commuting costs. On the one hand, we find that inventor productivity declines after firm relocation, going from producing 0.858 patents with 1.365 cites per year, to 0.328 patents with 0.633 cites per year. While some of this decline in productivity post-relocation is due to truncation bias, some might be genuine and related to the relocation itself, which we control for in our identification strategy. On the other hand, the mean commuting cost remains roughly constant post-relocation, whether measured in terms of distances or duration.

——————Insert Table 1——————

Table 4 reports the balance test results for inventor-firm-year observations prior to a workplace relocation event, divided by the direction of the commuting distance shock. Columns (4-6) show the results of two-sample t-test of equality of means between these three groups. Prior to the workplace relocation, inventors in the ***Closer*** and ***Farther*** groups have statistically indistinguishable productivity of around 0.85 patents per year. Unsurprisingly, the ***Closer*** and ***Farther*** groups differ significantly in terms of workplace-home distance across measures: inventors who live closer to the workplace are mechanically more likely to experience an increase in their workplace-home distance than for inventors who previously lived further away.

——————**Insert Table 4**——————

We also compare the groups based on other observable characteristics. Table 4 shows that both income and housing prices are somewhat higher for the ***Farther*** than the ***Closer*** group. After further investigation, we find that this difference can be explained by two facts acting together: first, the average establishment tends to be moving away from city center, with the average move being 2.55km away. Second, higher income and more productive inventors are preferentially living in the suburbs, away from the central city. This is shown in Figure 4 where we plot average inventor productivity and income against distance to the central city. Therefore, inventors whom firms on average are moving towards are more likely to live in the suburbs, and they tend to have higher income than inventors they are moving away from. These systematic differences could potentially affect our identification strategy. As discussed in the next section, we alleviate this concern by including firm-inventor fixed effects to absorb time-invariant differences between inventors, and test for differential pre-trends between the ***Farther*** and ***Closer*** groups of inventors.

——————**Insert Figure 4**——————

Turning to firm-level variables in Table 2, we find that establishments move by more than 23 km on average. The mean move is almost three times longer than the median move, meaning that there is both a concentration of establishments that moved relatively little, as well as a long tail of establishments that moved quite far, by tens of kilometers. While the average move is away from the CSA's central city by 2.55km, the standard deviation is almost six times larger, showing that there are many moves towards as well as away from the central city. Most of the establishments have a very small number of matched inventors, but there are some with hundreds of matched inventors. An establishment's average employee count grows post-relocation, suggesting that the average relocating firm is growing over time.

———————Insert Table 2———————

Looking more in-depth at the distribution of commuting distance shocks at the inventor-firm pair level, Figure 3 shows that most distance shocks are relatively small, with the mode of the distribution centered around no change in distance at all. However, a significant minority of inventors experienced a distance shock of more than 10 km. This variation allows the identification of a potential workplace-home distance effect on productivity.

———————Insert Figure 3———————

Finally, in Figure 2, we present a naïve descriptive analysis of our main research question relating workplace-home distance with inventor productivity. Using the full sample of 3,445 matched inventor-firm pairs, Figure 2 shows a clear negative correlation between distance and productivity, which is also statistically significant. The mean annual number of patents for an inventor is slightly less than 0.6 per year for inventors living very close to their offices. The patenting rate declines approximately linearly and steadily with increasing workplace-home distance, down to slightly less than 0.5 patents for inventors with a workplace-home distance of 60km. However, this negative correlation does not yet imply any causal relationship, since it is potentially confounded by endogenous sorting of inventors and firms. The next section details the empirical strategy to identify the causal effect in light of these possible confounding effects.

———————Insert Figure 2———————

3 Empirical Strategy

For estimating the causal effect of commuting distance on inventor productivity, the main challenge is that the location choices of both inventors and firms are endogenously determined. Inventors endogenously choose their place of residence based on a long list of factors in addition to commuting costs, such as the quality of local public services, size and price of homes available for sale, accessibility to amenities such as restaurants and movie theaters, etc. Factors that firms consider in their office location decisions include office rent and the neighborhood's productive amenities, in addition to accessibility. Therefore, a simple regression of inventor productivity on commuting distance would be biased due to sorting. We address both empirical challenges.

We first consider the endogenous location decision of inventors. Theoretically inventor residential sorting could bias our estimated coefficients in either direction in an OLS regression. For example, if skilled inventors choose to live further away from the office than unskilled inventors for any reason, such as a relative preference for good public schools, this sorting pattern would induce an upward bias in our estimated coefficient, and could even result in a positive correlation between commuting distance and inventor productivity. The opposite sorting pattern, with skilled inventors living closer to the firm, would induce a bias in the opposite direction. Consequently, we need an identification strategy that can eliminate the confounding effects of inventor sorting. (See an illustrative model in the Appendix that formalizes one potential mechanism for residential sorting.)

We now demonstrate the assumptions behind a difference-in-differences estimation strategy that can identify the causal effect of distance on productivity. Consider an inventor i working for firm j, living at a distance d_{ij} from the firm. Assume perfectly competitive labor markets where inventors are paid their marginal productivity of labor. Also assume that inventors are heterogeneous in their productivity type θ_i and where inventor productivity l_{ij}

is determined by the following equation:

$$l_{ij} = \theta_{ij} + \beta_i d_{ij} \tag{2}$$

where θ_{ij} is an individual-firm commuting distance-invariant productivity parameter for inventor i and firm j that denotes the quality of the inventor-firm match, and β_i is an individual-specific measure of how distance affects his or her productivity, This formulation allows for heterogeneous distance effects across individuals and time-invariant heterogeneity across both inventors and firms.

We model inventor sorting by skill level and distance. Individual-firm specific productivity parameters are not observed, but they are drawn from the real interval such that the distribution of individual types at each city m at distance x_{mi} from the city center is determined by:

$$\theta_{ij} = \alpha_0 + \alpha_1 x_{mi} + \delta_j + \epsilon_{ij} \tag{3}$$

where $\alpha_1 \neq 0$, ϵ_{ij} is a commuting distance-invariant match quality term and $E(x_{mi}\epsilon_{ij}) = 0$. Thus, there is sorting in city m of inventor types according to distance from city center. If the firm is located at the center of the city, then $x_{mi} = d_{ij}$. To consider a more general case of firm location, assume that firms tend to be located closer to city centers than residents, with a positive correlation between x_{mi} and d_{ij} on average. In this case, inventor i's distance to city center is correlated with distance to the firm, with $x_{mi} = \gamma_i d_{ij} + \mu_{ij}$. Plugging this expression into Equation 3 combined with Equation 2, we get the following expression:

$$\begin{aligned} l_{ij} &= (\alpha_0 + \alpha_1 x_{mi} + \delta_j + \epsilon_{ij}) + \beta_i d_{ij} \\ &= \alpha_0 + \beta_i d_{ij} + \delta_j + \alpha_1(\gamma_i d_{ij} + \mu_{ij}) + \epsilon_{ij} \\ &= \alpha_0 + (\beta_i + \alpha_1 \gamma_i) d_{ij} + \delta_j + \alpha_1 \mu_{ij} + \epsilon_{ij} \end{aligned} \tag{4}$$

Thus, the OLS estimates of β using Equation 2 would be biased if $\alpha_1 \neq 0$ and $\gamma_i \neq 0$, or if δ_j is correlated with d_{ij}. In those cases, the estimated parameter is not a weighted average of β_i, the individual causal relationships of distance on productivity.

We address this endogeneity problem using office relocations as exogenous shocks to commuting distance. In a world with no lumpy moving costs and positive commuting costs for inventors, inventors would re-optimize after firm relocation by changing their residential location. In a world with no job search costs and perfectly competitive labor markets, inventors may re-optimize by changing their job and workplace location. Therefore, the crucial assumption for my estimation is imperfect inventor resorting due to lumpy moving cost c_i, which each inventor i has to pay if he or she wants to move, and the heterogeneous match quality between inventors and firms, which allows for infra-marginal inventors who prefer to stay at the firm even after a negative productivity and wage shock (due to perfectly competitive labor markets). In summary, there are infra-marginal inventors for whom the benefits from lower commuting costs after resorting do not outweigh the moving costs.

Looking at this subsample of inventors who do not re-sort, we subtract their productivity before and after the move to get:

$$\Delta l_{ij} = \beta_i \Delta d_{ij} \tag{5}$$

and the sorting term $\alpha_1 x_{mi}$ drops out. Hence estimating a regression with individual firm-inventor pairs produces a consistent estimate of β_i, and estimating a fixed effect regression on pooled data across all non-re-sorting inventors estimates $\hat{\beta}_{FE}$, a weighted average of individual $\hat{\beta}_i$ (Baum-Snow and Ferreira, 2015). [18]

In this study, we estimate a stacked difference-in-differences regression relating changes

[18]One potential modification to our econometric model would include imperfect labor markets via non-zero job search costs. This would imply a positive wage compensation for commuting distance, which can be a form of efficiency wage to incentivize inventor effort (Ross and Zenou, 2008). In this modified model, Equation 5 would become:

$$\Delta l_{ij} = (\beta_i + ew_{ij}) \Delta d_{ij} \tag{6}$$

where ew_{ij} is the efficiency wage that firm j pays inventor i per unit commuting distance. Given that ew_{ij} should always be the opposite sign of β_i, we are estimating a lower bound for the weighted "pure" commuting effect on inventor productivity in the case of imperfect labor markets. (See further details in the Appendix.)

in commuting distance to changes in inventor productivity: taking into account year fixed effects and firm-inventor pair fixed effects. The precise specification is as follows:

$$Y_{ijt} = \beta d_{ijt} + \alpha_{ij} + \gamma_t + \delta_{jt} + \epsilon_{ijt} \tag{7}$$

where Y_{ijt} is the dependent variable that proxies for productivity, such as number of patents granted or scaled citations received for individual i working for firm j in year t. d_{ijt} is the distance between inventor i's home and firm j's office in year t. α_{ij} is an inventor-firm fixed effect that controls for the inherent productivity differences between individuals, taking into account the matching quality between inventor and firm. γ_t controls for the aggregate yearly trend in patents granted. δ_{jt} controls for the specific firm location before and after the relocation at the ZIP code tabulation area (ZCTA) level, to account for potential time-invariant productive amenity differences, for example due to knowledge spillovers from nearby firms. Finally, ϵ_{ijt} is the error term. Errors are clustered at the inventor-firm pair level.

We now consider the endogenous location choice of firms. Given that firms are profit-maximizing entities, they only decide to relocate when the benefits of doing so are greater than its costs. These benefits can come from direct monetary savings due to lower rent, productivity increases from higher productive amenities offered by the new locations, potentially due to knowledge spillovers from other firms, lower overall commuting distance to its workers including inventors, etc. Some of these factors could both cause differences in average productivity before versus after firm relocation (such as changes in productive amenities) and be correlated with average changes in inventor commuting distance across firms, which would bias our estimation results.

We employ multiple strategies to deal with this concern. First, in the main specification (Equation 7) above, we control for the time-invariant component of productive amenities using firm-location fixed effects at the ZCTA level. However, this still assumes that time-invariant productive amenities are identical across firms within any ZCTA. For a more stringent test, we estimate the robustness specification below:

$$Y_{ijt} = \beta d_{ijt} + \alpha_{ij} + \gamma_t + \delta^{'}_{jt} + \epsilon_{ijt} \tag{8}$$

where $\delta^{'}_{jt}$ is a unique firm-level office location fixed effect that differs before and after the relocation using office location fixed effects. In other words, we control for all time-invariant across-firm variations in firm-level variables such as office rent, productive amenities, average commuting distance to inventors, etc. Thus the only variation we use to estimate β is within-firm across-inventor differential changes in commuting distance.[19]

After controlling for time-invariant determinants of endogenous firm location, our estimates may still be biased if there is reverse causality, where inventors' productivity change over time, and firms are relocating to be closer to inventors whom they anticipate will become more productive in the future. This would violate the parallel trends assumption that inventors who get shocked closer in or farther away should have similar productivity trends pre- and post- the firm relocation year, absent any firm relocations, and bias our coefficient upward. To address this concern, we adopt the standard test for parallel trends by estimating a generalized difference-in-differences specification that looks at differential average patent count year by direction of commuting shock, and finding no significant differential pre-trends in productivity between inventors who get shocked closer in or farther away (see Figure 5).

Another identification assumption is that there are no other events that occur simultaneously with the firm relocation events, that would impact differentially the productivity of inventors whose distance to work increases versus those whose distance to work decreases. This is unlikely, especially since we use a stacked specification; thus, the confounding events must be correlated with firm relocation events in different CSAs across the country and in different years. Balancing test results in Table 1 show that inventors who received a positive distance shock are not significantly different from inventors who received a negative distance

[19]Due to lack of power in some tests of heterogeneity and robustness checks, we do not use Equation 8 as our main specification. Nevertheless, all results estimated using Equation 8 are qualitatively the same as in our main specification.

shock in terms of patent productivity.

Given the discrete dependent variable of patent count and scaled citations, it is arguably better to use models that are more suited to count data, by modeling the dependent variable to follow a Poisson distribution or a negative binomial distribution (Hausman, Hall and Griliches, 1984; Cameron and Trivedi, 2013). Compared to the linear model, both Poisson and negative binomial distributions are intrinsically valued as non-negative integers, which exactly matches outcomes in both the number of patents filed and patent citations. The difference between the two distributions is that Poisson assumes that the distribution mean and variance are identical, while negative binomial allows for overdispersion, with variance higher than the mean. We use both Poisson regression[20] and negative binomial regression[21] as robustness checks.

4 Results

In this section, we first present our main results on both patent quantity and quality. Then we explore tests of heterogeneity that offer suggestive evidence on mechanisms driving our

[20] More specifically, Poisson regression estimates the following model:

$$pr(Y_{ijt}|x_{ijt}) = \frac{e^{-\lambda_{ijt}}\lambda_{ijt}^{Y_{ijt}}}{Y_{ijt}!}$$

where the probability that Y_{ijt} patents are applied for by inventor i for firm j in year t follows the Poisson distribution with λ_{ijt} as the inventor-firm, firm-location and time-specific Poisson parameter, which is obtained as follows:

$$log\lambda_{ijt} = \beta d_{ijt} + \alpha_{ij} + \gamma_t + \delta_{jt} + \epsilon_{ijt}$$

As in the linear regression case, the Poisson regression has errors clustered at the inventor-firm level.

[21] Negative binomial regression generalizes Poisson distribution by adding an unobserved heterogeneity term, which follows the gamma distribution with both shape and rate parameters being θ, to the Poisson regression model, so:

$$pr(Y_{ijt}|x_{ijt}) = \frac{\Gamma(Y_{ijt}+\theta)}{Y_{ijt}!\,\Gamma(\theta)}\Big(\frac{\theta}{\theta+\lambda_{ijt}}\Big)^{\theta}\Big(\frac{\lambda_{ijt}}{\theta+\lambda_{ijt}}\Big)^{Y_{ijt}}$$

where the probability that Y_{ijt} patents are applied for by inventor i for firm j in year t follows the negative binomial distribution where λ_{ijt} is the inventor-firm, firm-location and time-specific Poisson parameter, which is obtained as follows:

$$log\lambda_{ijt} = \beta d_{ijt} + \alpha_{ij} + \gamma_t + \delta_{jt} + \epsilon_{ijt}$$

Standard errors in the negative binomial model are clustered at the inventor-firm level.

results. Finally, we show that our results are robust to alternative specifications.

Table 5 shows estimated coefficients for our main difference-in-differences specifications. Column (1) shows that ***Distance*** is negatively correlated with inventor productivity, as expected given the negative slope in Figure 2. The size of the coefficient, however, triples after controlling for inventor-firm pair fixed effects in Column (2) and becomes more significant. This difference suggests that more skilled inventors endogenously choose residential locations further away from their workplace than less skilled inventors, which biases the OLS coefficient downward.[22] The result remains highly significant when we control for year fixed effects in Column (3) and firm location fixed effects in Column (4), suggesting that endogenous location choice of firms to pursue higher time-invariant productive amenities are not driving our results. Eliminating all between-firm variation and relying solely on within-firm between-inventor differential changes in commuting distance still produces a significant effect in Column (5). Column (4) represents our preferred specification, with every 10 kilometer increase in ***Distance*** causing an average decrease in inventor productivity of 0.041 patents per year. In percentage terms, this represents a 5% decrease in inventor productivity per 10 kilometers, compared with the average 0.86 patents per year per inventor-firm pair before the move. While the prior literature documents a wide range of estimates for the value of a patent, we take a conservative lower-bound of $100,000 on the value of an average patent (Pakes, 1985; Austin, 1993; Barney, 2002; Serrano, 2005; Bessen, 2008; Fischer and Leidinger, 2014). This patent value estimate translates to about $4270 in value generated for the employer per inventor per year per 10km shorter commute.

——————Insert Table 5——————

One remaining endogeneity challenge exists: firms might be moving towards inventors who are becoming more productive, potentially to take advantage of the additional increase in productivity due to shorter commuting distances. We test this possible violation of our

[22]We confirm this by plotting pre-relocation inventor income and home price against commuting distance, and finding a strong positive correlation between the two.

parallel trends assumption in Figure 5, which shows the difference in the predicted average number of patents per inventor-firm pair per year, between pairs whose workplace-home distance increased by more than 1km, versus those who did not, plotted against the number of years pre- and post-relocation. The graph shows that the difference between the productivity of the two groups of inventors is close to zero and displays no significant trends before the firm relocation. Second, productivity for inventors who get shocked farther away falls during the year of the relocation, and afterwards remains lower relative to the productivity of those who get shocked closer in. This is consistent with our hypothesis that the increase in commuting distance permanently affects inventor productivity, and argues against any temporary shocks that occur in conjunction with firm relocations driving our results.

——————Insert Figure 5——————

Turning to measures of patent quality, we find the same negative effect of commuting distance on inventor productivity. Column (1) in Table 6 shows that a 10 km increase in commuting distance causes a 0.094 decrease in patent scaled citations, roughly 7% of the pre-relocation mean. This suggests that the decrease in patent counts is not driven by inventors applying for fewer, but more impactful patents. The generality and originality results in Columns (2) and (3) are consistent with this interpretation: the overall scientific quality and applicability of patents fall with increasing commuting distance, in step with patent counts. Testing for patent economic value using the number of maintenance fee payments shows a potential decrease, even though the coefficient is not significant. This is likely due to the lack of power in this specification, since the last maintenance fee payment is only required 11.5 years after the patent grant date, causing a much more severe truncation problem than the other specifications.

Overall, our results show that increasing commuting distance negatively impacts inventor productivity both in terms of patent quantity and quality. This decrease is economically significant, equaling thousands of dollars per year per inventor using conservative estimates of patent value. It is also permanent and does not reverse over time.

——————Insert Table 6——————

4.1 Tests for Heterogeneity

We explore suggestive evidence for potential underlying mechanisms driving our results. First we examine the possibility of non-linearity in the commuting distance-productivity relationship by adding a squared commuting distance to the estimating equation. On the one hand, the direct opportunity cost of commuting is likely to be linear with distance/time if it directly reduces the amount of time an inventor can spend at work. On the other hand, non-linear effects could arise if some inventors have an optimal distance between home and workplace that is neither too small nor too large and that gives them some amount of separation between their work and family lives. Another reason for non-linearity arises when the cost of providing effort increases non-linearly with longer commutes, such as in urban efficiency wage models.

As shown in Table 7, we find that when the square term is added, in both patent quantity and quality regressions, the main effect becomes insignificant, though its magnitude is similar to our main specification. The coefficient on $Distance^2$, however, is very close to zero and insignificant, suggesting that its causal effect on inventor productivity is linear within the distance range (up to 100 km) covered by our sample. These results are consistent with a linear opportunity cost of lost time at work, and less consistent with models of effort provision where the cost of effort increases non-linearly with commuting distance.

——————Insert Table 7——————

Then, given that the highest-productivity inventors may have a disproportionate impact on a firm's innovation output (e.g., Akcigit, Baslandze and Stantcheva, 2016), we test whether the commuting distance effect is driven by these top inventors. To do so, we estimate a triple-differences specification, by interacting our explanatory variable with an indicator of whether an inventor belongs to the top decile of all inventors by productivity.

Table 8 shows that the negative effect of distance on inventor productivity is largely driven by high-productivity inventors, with their productivity decreasing by 0.159 patents per year per 10 km, which is around 6 times the average effect for the other 90% of less productive inventors, whose coefficient is reduced to an insignificant 0.026 patents per year per 10 km. This large discrepancy is replicated in the patent quality measures. One reason for this difference is that top inventors' mean productivity is higher, so their opportunity cost for every hour lost at work is also higher. Even after taking their higher mean into account, however, top inventors suffer more proportionally than the average inventor: with a 10km increase in ***Distance*** causing a 10% drop in productivity, versus less than 4% for less productive inventors. This large discrepancy suggests that there are potentially different mechanisms driving the effect for the top versus average inventors. For example, perhaps the cost of effort increases more steeply with commuting distance for top versus average inventors.

——————Insert Table 8——————

Estimating the equation separately on shocks to closer versus farther away inventors in Table 9 shows an interesting heterogeneous result where the negative main effect is driven by inventors who get shocked farther away. Indeed, while the effect of changing commuting distance is insignificant for inventors who did not get a large commuting distance shock (Column 2), or who get shocked closer in to their workplace (Column 1), the coefficient for the inventor who get shocked farther away from their office is much larger than in our main specification, with a 10 km increase in commuting distance causing a 0.106 decrease in patent count, equivalent to around 12% of their average pre-relocation productivity. This result suggests that the productivity of inventors who endogenously selected to live closer to their workplace is more sensitive to commuting distance than inventors who selected to live farther away. In fact, there could be inventors in the latter group whose productivity is not affected by commuting distance at all. In summary, this result shows that there is large between-inventor heterogeneity in their elasticity of productivity versus commuting distance.

Insert Table 9

While all inventors named on a patent are legally required to have contributed to the conception of the patented invention, their contributions might differ in importance. An implicit assumption we have adopted up to now is that each inventor on a patent is assumed to have made an equal contribution. However, changes in per-patent contribution could occur in conjunction with a change in commuting distance, potentially biasing our estimates in an uncertain direction versus the "real" effect measured by actual contribution. To account for this, we introduce an alternate patent count that only counts single-authored patents. This count is not subject to the bias above the inventor's contribution to a single-authored patent is always 100%. Because most patents have multiple inventors, however, this reduces our overall patent count by around 80%, which reduces the power in our estimation procedure. Nevertheless, Column (1) in Table 10 show that a 10 km increase in *Distance* causes a significant decrease of 0.012 single-authored patents per year, which corresponds to around 10% of the average pre-relocation inventor productivity in terms of single-authored patents, similar to our main results. It also suggests that commuting costs affect team and solo productivity equally, and is not particularly detrimental to team work even though solo work may be easier to perform with longer commutes, due to the rise of telecommuting.

Insert Table 10

4.2 Robustness Checks

To assess the robustness of our findings relative to the effects of outliers, and to make sure that our results are not driven by a few inventors who live far out from their workplace, we divide the inventor-firm pairs into three categories based on whether their *Distance* got farther, got closer, or remained largely unchanged (change < 1km) after the firm relocation shock. We construct the categorical variable Δ *Distance Direction* that equals 0 to observations before the relocation, 1 for post-relocation observations of pairs that moved farther away, -1 for pairs that moved closer, and 0 for pairs whose distance remained largely unchanged. We then use this categorical variable in place of ***Distance*** in Table 11. These estimates look similar to those in Table 5, our preferred specification. Column (4) shows that a higher Δ ***Distance Direction*** has a negative effect on the number of patents produced per year, while Column (5) again shows that this result is driven by both average- and high-productivity inventors.

———————Insert Table 11———————

We check the robustness of our results using alternative explanatory variables. Table 12 shows that using ***Drive Distance*** or ***Drive Duration*** gives significant negative effects that are similar in magnitude to our main specification using geodesic distances.

———————Insert Table 12———————

Our results are robust to using Poisson and negative binomial estimation. Results are shown in Table 13. The coefficient in Column (1) shows that a 10 km increase in ***Distance*** causes a 4.9% decrease in patent count in the Poisson regression, and the coefficient in Column (3) shows a 5.4% decrease in the negative binomial regression. Both numbers are very similar to our main linear specification.

———————Insert Table 13———————

We have also verified that our results are robust to using name and location weights that take into account the probability of having duplicate names and multiple establishment locations. Our results are also not driven by any one large CSA such as the San Francisco Bay Area, or only by growing or shrinking firms (see the Appendix for further details), etc.

4.3 Selection Issues

We have shown that increasing commuting distance causes decreases in inventor productivity for our sample of non-resorting inventors after a workplace relocation. However, this is a selected sample of the population of all inventors working for the relocating firms, because inventors may *ex post* self-select out of our sample by moving to another job or home location. In this section, we explore the question of how valid our results are accounting for these home "movers" and job "quitters".[23]

Taking "movers" and "quitters" into account, we believe that our estimated commuting elasticity of productivity for non-movers is actually a lower bound to the actual elasticity in the entire population of inventors. Intuitively, we may expect that inventors with relatively high elasticity (i.e. the absolute value of β_i) to be particularly sensitive to negative shocks in commuting distance. Given fixed moving costs c and job search costs s across inventors, and given the same farther out shock in commuting distance, an inventor with high β_i will suffer a proportionately larger loss of productivity and wages than an inventor with low β_i, and thus be more likely to re-sort by moving or quitting, in other words self-selecting out of our final sample. Therefore, we are likely underestimating the average sensitivity to commuting distance among all inventors.

To document whether this possible selection effect exists, and the extent of re-sorting after firm office relocation, we can study the inventor attrition rate both before and after a firm relocates. There are two ways an inventor can re-sort: he can move his location, or he can change his employer. Therefore, we can compute both the raw number and the rate at which inventors are moving and selling their homes, and/or quitting their jobs. We expect to see an increase in both rates for a period after the firm relocates, compared to a period just before the relocation. In order to do so, we define inventors selling their homes as movers, and we define as quitters those who start patenting for a different firm.[24]

Given these possible *ex post* selection effects, we expect that inventors who suffer from a negative distance shock may preferentially self-select out of our sample. Furthermore, the workplace relocation itself may be a signal for inventors to reconsider their residential

[23] We cannot distinguish between voluntary and involuntary job exits, so an equivalent term would be job "separators".

[24] Specifically, we define inventors as having quit their jobs the year they filed a patent for a different firm, or failing that, five years after they last filed a patent for the current firm.

and career options. To partially test for the existence of these ***ex post*** selection effects, we compare ***ex post*** selection for shocks that move inventors closer versus farther away from the firm:

$$q_{it} = \omega_1 post_{it} + \omega_2 post_{it} * direction_{it} + \gamma_i + \delta_t + qe_{it} \tag{9}$$

where q_{it} is a binary variable that equals 1 in the year when inventor i quits his original employer or sells his home, and 0 otherwise. qe_{it} is an individual-level error term, and $post_{it}$ is an indicator function that is equal to one after the company relocates, and 0 otherwise. $direction_{it}$ is a categorical variable equalling 0 prior to a workplace relocation event, and after a relocation event, takes the value of 1 for geodesic distance increases (farther away by $>$ 1km), -1 for distance decreases (closer by $>$ 1km), and 0 when distance is approximately unchanged (changed by less than 1km). ω_1 represents the effect of the firm relocation itself, while ω_2 represents the differential selection effect for inventors shocked farther versus closer to their workplace due to the firm relocation. We would intuitively expect that $\omega_2 > 0$.

Our hypothesis of higher moving and quitting rates for inventors receiving negative distance shocks is consistent with results in Table 14. While we do not detect a significant differential selection for home moves, we do detect a significant post-relocation increase in job quitting rater for inventors shocked farther to their firm versus closer.

———————**Insert Table 14**———————

4.4 Discussions and Interpretations

Finally, the interpretation of our estimated coefficients differs depending on labor market assumptions. In our base econometric model, we assume that the labor market is perfectly competitive with no job search costs, with all inventors are paid at a wage equalling their marginal productivity. In this case, inventor wage does not depend on commuting distance beyond the direct impact that commuting has on productivity, so we are estimating the "pure" causal effect of commuting on inventor productivity.

If job search costs are greater than zero, however, the interpretation changes. In imperfect urban labor markets, firms have market power and pay inventors a wage below their marginal productivity. They may also compensate inventors for longer commutes by paying a somewhat higher wage. This wage compensation incentivizes these negatively shocked inventors to stay at the firm and provide more effort (Ross and Zenou, 2008). This wage compensation has been shown to exist empirically by Mulalic, Van Ommeren and Pilegaard (2014), who study the population of employees in Denmark. If we plausibly assume that wage compensation for commuting also exists for inventors, partly to provide an efficiency wage and incentivize effort, then our results can be interpreted as a total effect, combining the "pure" causal effect of commuting on productivity and the countering effect of higher efficiency wage. This implies that even for non-movers, our estimates represent a lower bound for the "pure" effect of commuting on inventor productivity, without taking into account wage compensation. (See the Appendix for a derivation of the econometric model with wage compensation.)

5 Conclusion

We empirically investigate how commuting distance affects inventor productivity using a stacked difference-in-differences design, comparing inventors who stay at the same firm and residential location before versus after an office relocation. This strategy identifies the causal effect of computing distance on productivity, separated from confounding sorting effects by both firms and inventors. We find strong evidence that commuting negatively affects inventor productivity, with every 10 km increase in commuting distance leading to a 5% or more decrease in patent productivity for the average inventor.

By identifying a causal link between smaller commuting distance and higher inventor productivity, our results have clear policy implications, suggesting that urban planners should aim to increase urban density and firms should encourage their inventors to live closer to the workplace. This would enable both higher levels of firm and economic growth.

Our findings naturally suggest several potential avenues for future research. First, the value of location proximity may change over time. For example, as newer telecommunications technologies, such as broadband and wireless data connections, proliferate, telecommuting and remote work become more common and socially acceptable. These technological changes could decrease the value of living closer to the workplace. This question could be addressed with a longer time series dataset, covering the time period before the broad adoption of a particular technology (e.g., Internet) and the period after its widespread adoption.

Second, further work is needed to explore the exact mechanisms underlying our identified negative elasticity. While we offer suggestive evidence that commuting affects inventor productivity linearly and in equal proportions for both average and top inventors, this finding could still be consistent with multiple potential mechanisms at work. For example, there could be a combination of less time spent at work and being less productive while at work. Additional research disentangling these factors could be possible with more detailed time-use data at the inventor level.

One interesting unexplored mechanism in our paper would be knowledge spillover effects through increased collaboration via social networks (Allen, 1977; Jaffe, Trajtenberg and Henderson, 1993), as workers who live closer to the workplace also tend to live closer together (Bayer, Ross and Topa, 2008; Lindquist, Sauermann and Zenou, 2015). Our paper does not include this mechanism since, to the first order, inventor home locations are largely the same before and after firm relocations. For future work, however, the contribution of knowledge spillovers can be explored through changes in the spatial distributions of observable inventor collaboration networks.

Finally, there needs to be more research done to empirically identify the effect of commuting distance on productivity for other types of workers. Inventors and researchers are knowledge workers, who tend to be highly-educated and earn relatively high incomes. Their jobs also require different skills than other occupations such as manufacturing or construction.

The commuting distance-productivity elasticity may then differ depending on the exact mechanisms at play. In summary, future work could expand on generalizing our particular findings on the link between distance and productivity.

References

Akcigit, Ufuk, Salome Baslandze, and Stefanie Stantcheva. 2016. "Taxation and the International Mobility of Inventors." *American Economic Review*, 106(10): 2930–2981.

Allen, Thomas J. 1977. *Managing the Flow of Technology.* Cambridge, MA:MIT Press.

Alonso, William. 1964. *Location and Land Use: Toward a General Theory of Land Rent.* Cambridge, MA:Harvard University Press.

Ashforth, Blake E, Glen E Kreiner, and Mel Fugate. 2000. "All in a Day's Work: Boundaries and Micro Role Transitions." *Academy of Management Review*, 25(3): 472–491.

Austin, David H. 1993. "An Event-Study Approach to Measuring Innovative Output: The Case of Biotechnology." *American Economic Review*, 83(2): 253–258.

Barney, Jonathan A. 2002. "A Study of Patent Mortality Rates: Using Statistical Survival Analysis to Rate and Value Patent Assets." *American Intellectual Property Law Association Quarterly Journal*, 30: 317.

Baum-Snow, Nathaniel, and Fernando Ferreira. 2015. "Causal Inference in Urban and Regional Economics." In *Handbook of Regional and Urban Economics*. Vol. 5, , ed. Gilles Duranton, J Vernon Henderson and William Strange, 3–68. Elsevier.

Bayer, Patrick, Stephen L Ross, and Giorgio Topa. 2008. "Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes." *Journal of Political Economy*, 116(6): 1150–1196.

Bernstein, Shai, Timothy McQuade, and Richard R Townsend. 2017. "Does economic insecurity affect employee innovation?"

Bessen, James. 2008. "The Value of US Patents by Owner and Patent Characteristics." *Research Policy*, 37(5): 932–945.

Bloom, Nick, Tobias Kretschmer, and John Van Reenen. 2009. "Work-Life Balance, Management Practices and Productivity." In *International Di erences in the Business Practices and Productivity of Firms*. 15–54. Chicago, IL:University of Chicago Press.

Breschi, Stefano, and Camilla Lenzi. 2016. "Co-invention networks and inventive productivity in US cities." *Journal of Urban Economics*, 92: 66–75.

Cameron, A Colin, and Pravin K Trivedi. 2013. *Regression Analysis of Count Data.* Vol. 53, Cambridge, UK:Cambridge University Press.

Carley, Michael, Deepak Hedge, and Alan Marco. 2015. "What is the probability of receiving a US patent." *Yale JL & Tech.*, 17: 203.

Carlino, Gerald, and William R Kerr. 2015. "Agglomeration and innovation." In *Handbook of Regional and Urban Economics*. Vol. 5, 349–404. Elsevier.

Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, and Sébastien Roux. 2010. "Estimating Agglomeration Economies with History, Geology, and Worker Effects." In *Agglomeration Economics*. , ed. Edward L. Glaeser, Chapter 1, 15–66. Chicago, IL:University of Chicago Press.

Crabtree, Steve. 2010. "Well-Being Lower Among Workers with Long Commutes."

Davis, Donald R, and Jonathan I Dingel. 2014. "The Comparative Advantage of Cities." National Bureau of Economic Research.

DiNardo, John, and David S Lee. 2004. "Economic Impacts of New Unionization on Private Sector Employers: 1984–2001." *Quarterly Journal of Economics*, 119(4): 1383–1441.

Duranton, Gilles, and Diego Puga. 2001. "Nursery Cities: Urban Diversity, Process Innovation, and the Life Cycle of Products." *American Economic Review*, 91(5): 1454–1477.

Duranton, Gilles, and Diego Puga. 2015. "Urban Land Use." In *Handbook of Regional and Urban Economics*. Vol. 5, 467–560. Elsevier.

Duranton, Gilles, and Matthew A Turner. 2012. "Urban Growth and Transportation." *Review of Economic Studies*, 79(4): 1407–1440.

Duranton, Gilles, and Matthew A Turner. 2017. "Urban Form and Driving: Evidence from US Cities." Working Paper.

Ferreira, Fernando, and Joseph Gyourko. 2015. "A New Look at the US Foreclosure Crisis: Panel Data Evidence of Prime and Subprime Borrowers from 1997 to 2012." NBER Working Paper Series.

Fischer, Timo, and Jan Leidinger. 2014. "Testing Patent Value Indicators on Directly Observed Patent Value——An Empirical Analysis of Ocean Tomo Patent Auctions." *Research Policy*, 43(3): 519–529.

Furnham, Adrian. 2000. "The brainstorming myth." *Business strategy review*, 11(4): 21–28.

Glaeser, Edward L. 1999. "Learning in Cities." *Journal of Urban Economics*, 46(2): 254–277.

Glaeser, Edward L, and Naomi Hausman. 2019. "The spatial mismatch between innovation and joblessness." *NBER Working Paper Series*.

Google LLC. 2018. "Developer Guide."

Hall, Bronwyn H, Adam B Jaffe, and Manuel Trajtenberg. 2001. "The NBER patent citation data file: Lessons, insights and methodological tools." National Bureau of Economic Research.

Hall, Bronwyn H, Adam Jaffe, and Manuel Trajtenberg. 2005. "Market value and patent citations." *RAND Journal of economics*, 16–38.

Hausman, Jerry A, Bronwyn H Hall, and Zvi Griliches. 1984. "Econometric Models for Count Data with an Application to the Patents-R&D Relationship." *Econometrica*, 52(4): 909–938.

Helsley, Robert W, and William C Strange. 1990. "Matching and Agglomeration Economies in a System of Cities." *Regional Science and urban Economics*, 20(2): 189–212.

Ingraham, Christopher. 2016. "The Astonishing Human Potential Wasted on Commutes."

Jaffe, Adam B, Manuel Trajtenberg, and Rebecca Henderson. 1993. "Geographic

Localization of Knowledge Spillovers as Evidenced by Patent Citations." *Quarterly Journal of Economics*, 577–598.

Jovanovic, Boyan, and Rafael Rob. 1989. "The Growth and Diffusion of Knowledge." *Review of Economic Studies*, 56(4): 569–582.

Kahneman, Daniel, and Alan B Krueger. 2006. "Developments in the Measurement of Subjective Well-Being." *Journal of Economic Perspectives*, 20(1): 3–24.

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

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological innovation, resource allocation, and growth." *The Quarterly Journal of Economics*, 132(2): 665–712.

Lagos, Ricardo. 2000. "An Alternative Approach to Search Frictions." *Journal of Political Economy*, 108(5): 851–873.

Li, Guan-Cheng, Ronald Lai, Alexander D'Amour, David M Doolin, Ye Sun, Vetle I Torvik, Z Yu Amy, and Lee Fleming. 2014. "Disambiguation and Co-Authorship Networks of the US Patent Inventor Database (1975–2010)." *Research Policy*, 43(6): 941–955.

Lindquist, Matthew J, Jan Sauermann, and Yves Zenou. 2015. "Network Effects on Worker Productivity." CEPR Discussion Paper.

Lorenz, Olga, and Laszlo Goerke. 2016. ""Is your commute really making you fat?": The causal effect of commuting distance on height-adjusted weight." *German Economic Association Annual Conference*.

Mills, Edwin S. 1967. "An Aggregative Model of Resource Allocation in a Metropolitan Area." *American Economic Review*, 57(2): 197–210.

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(7): 1858–1903.

Mulalic, Ismir, Jos N Van Ommeren, and Ninette Pilegaard. 2014. "Wages and Commuting: Quasi-Natural Experiments' Evidence from Firms that Relocate." *Economic Journal*, 124(579): 1086–1105.

Muth, Richard F. 1969. *Cities and Housing: The Spatial Pattern of Urban Residential Land Use.* Chicago, IL:University of Chicago Press.

Pakes, Ariel. 1985. "On Patents, R&D, and the Stock Market Rate of Return." *Journal of Political Economy*, 93(2): 390–409.

Pakes, Ariel. 1986. "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks." *Econometrica: Journal of the Econometric Society*, 755–784.

Perino, Marissa. 2019. "Here's What the Average Person Spends on Their Commute

Annually in Every State.”

Reuters. 2015. “Facebook Gives Employees Big Money to Move Closer to Work.”

Ross, Stephen L, and Yves Zenou. 2008. “Are Shirking and Leisure Substitutable? An Empirical Test of Efficiency Wages Based on Urban Economic Theory.” *Regional Science and Urban Economics*, 38(5): 498–517.

Serrano, Carlos J. 2005. “The Market for Intellectual Property: Evidence from the Transfer of Patents.” *Unpublished manuscript*.

Shapiro, Carl, and Joseph E Stiglitz. 1984. “Equilibrium Unemployment as a Worker Discipline Device.” *American Economic Review*, 74(3): 433–444.

Toivanen, Otto, and Lotta Väänänen. 2012. “Returns to inventors.” *Review of Economics and Statistics*, 94(4): 1173–1190.

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(1): 19–50.

U.S. Census Bureau. 2018. “About This Program.”

U.S. Government Bureau of Transportation Statistics. 2003. “From Home to Work, the Average Commute Is 26.4 Minutes.” *OmniStats*, 3(4).

Van Ommeren, Jos, and Mogens Fosgerau. 2009. “Workers’ marginal costs of commuting.” *Journal of Urban Economics*, 65(1): 38–47.

Van Ommeren, Jos N, and Eva Gutiérrez-i Puigarnau. 2011. “Are Workers with a Long Commute Less Productive? An Empirical Analysis of Absenteeism.” *Regional Science and Urban Economics*, 41(1): 1–8.

Vincenty, Thaddeus. 1975. “Direct and Inverse Solutions of Geodesics on the Ellipsoid with Application of Nested Equations.” *Survey Review*, 23(176): 88–93.

Zax, Jeffrey S. 1991. “Compensation for commutes in labor and housing markets.” *Journal of urban Economics*, 30(2): 192–207.

Zax, Jeffrey S, and John F Kain. 1996. “Moving to the Suburbs: Do Relocating Companies Leave their Black Employees Behind?” *Journal of Labor Economics*, 472–504.

Zenou, Yves. 2002. “How Do Firms Redline Workers?” *Journal of Urban Economics*, 52(3): 391–408.

Zenou, Yves. 2009. *Urban Labor Economics.* Cambridge University Press.

Figures

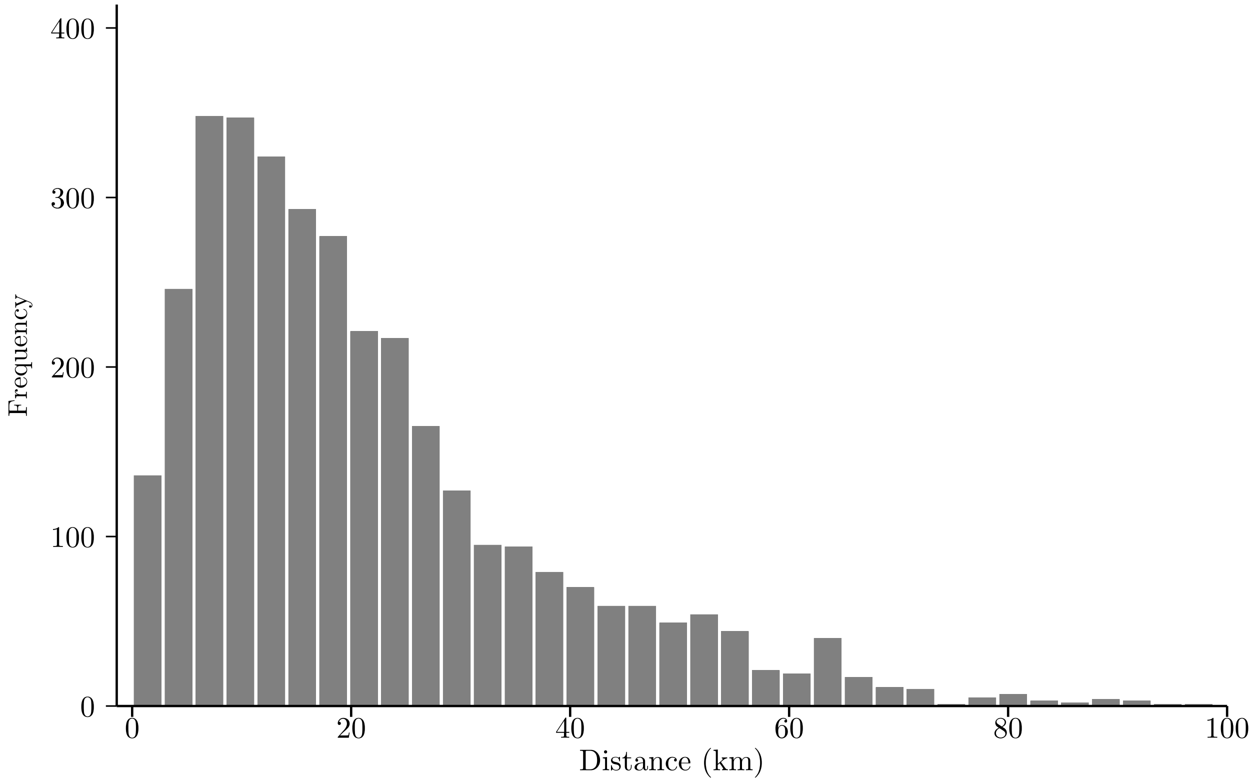

Figure 1: **Distribution of Workplace-Home Geodesic Distance**. This histogram depicts the frequency distribution of geodesic distance between the workplace and home of the inventors in the full sample, prior to any office relocation event; each inventor-firm pair is represented once.

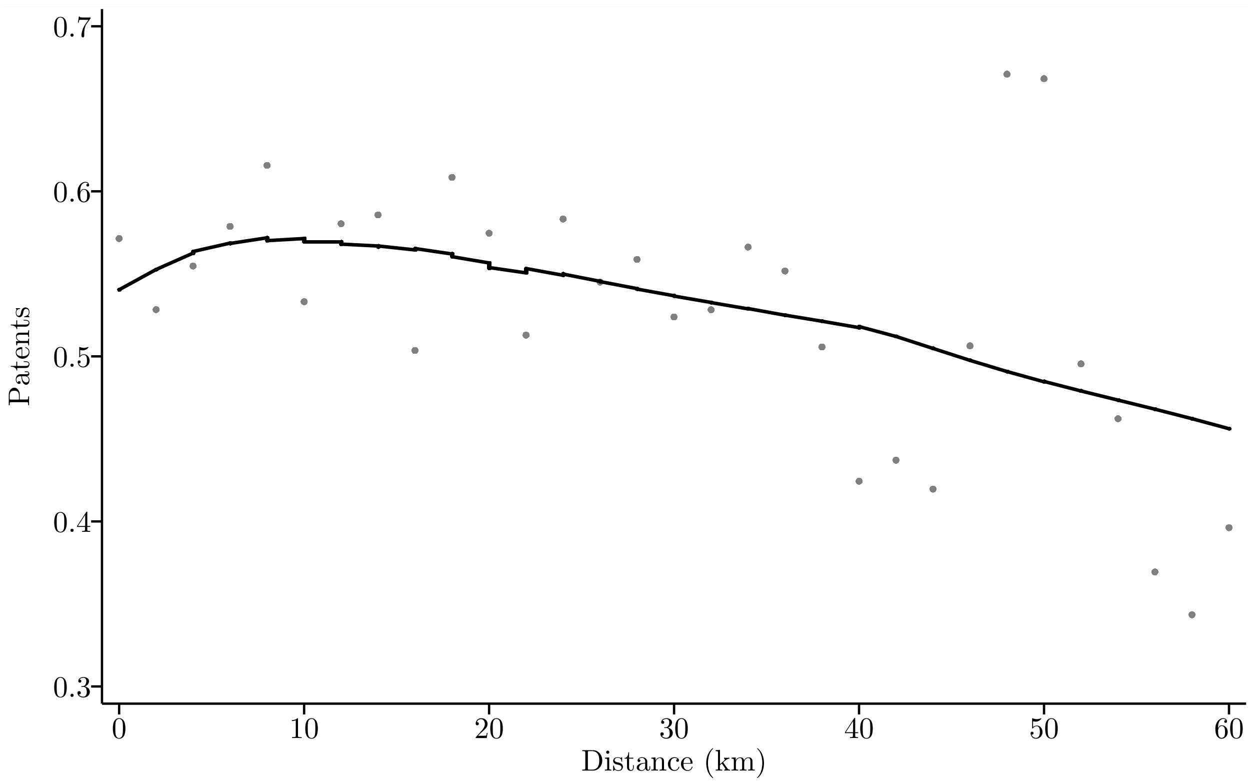

Figure 2: **Descriptive Relationship Between Workplace-Home Geodesic Distance and Patent Grant Output**. This descriptive graph demonstrates the naïve relationship between an inventor's workplace-home geodesic distance (horizontal axis) and an inventor's patent grant output (vertical axis). Each grey dot is the average number of patents per inventor per year for the years prior to a workplace relocation event, where bins formed for each nearest whole kilometer, e.g., the dot for 10km is the average of inventors who have a geodesic distance $[9.5, 10.5)$. Geodesic distance above 60km were omitted from this figure due to insufficient observations.

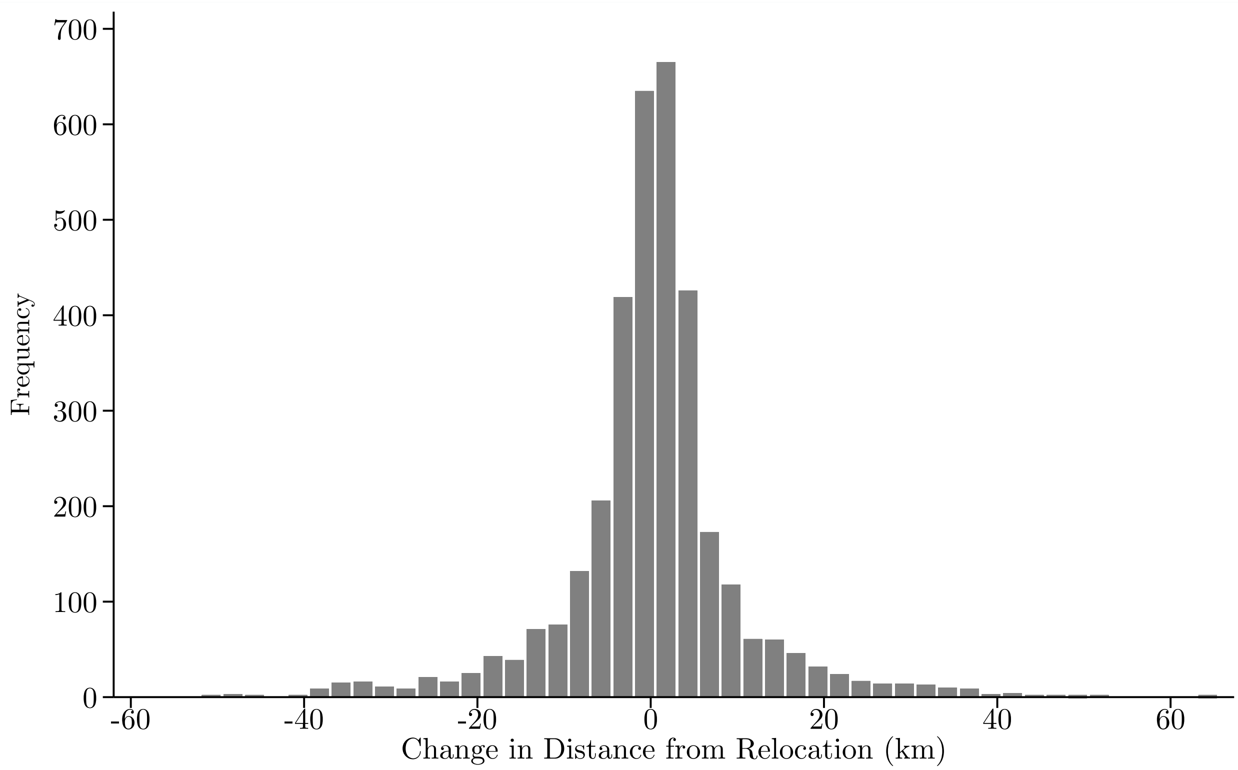

Figure 3: **Distribution of Change in Workplace-Home Geodesic Distance**. This histogram depicts the frequency distribution of changes geodesic distance between the workplace and home of the inventors in the full sample, before and after the corresponding office relocation event; each inventor-firm pair is represented once.

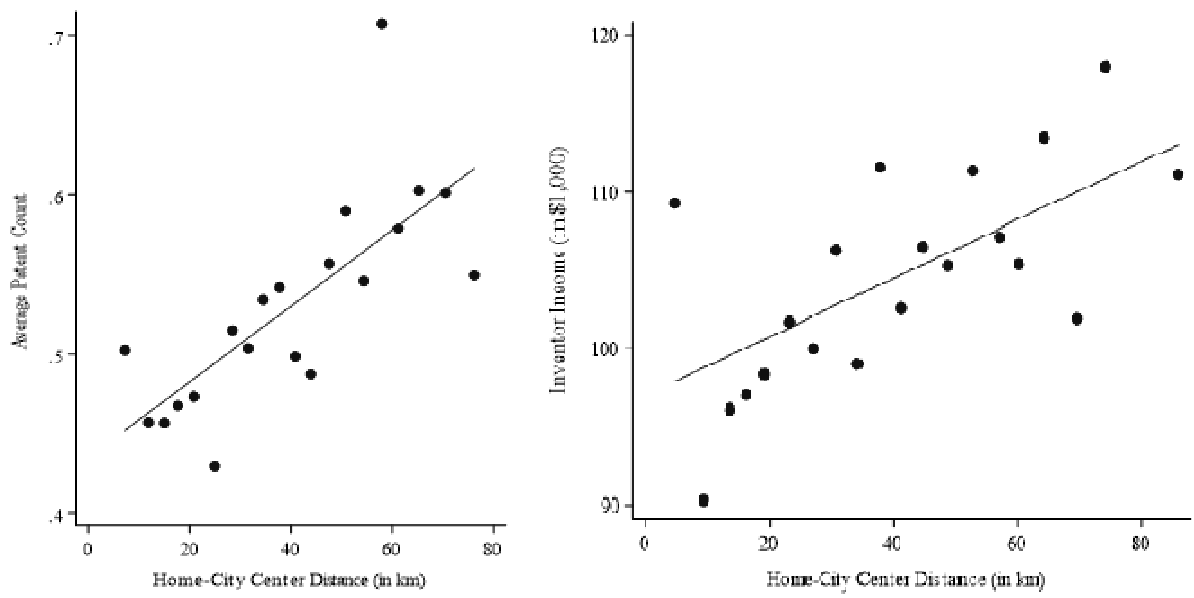

Figure 4: **Residential Sorting by Distance to City Center**. These binscatter plots show the relationship between an inventor's home-central city distance (horizontal axis) and an inventor's productivity and income (vertical axis). Each grey dot is the average number of patents/income for the years prior to a workplace relocation event, where bins formed for each nearest 5 kilometers.

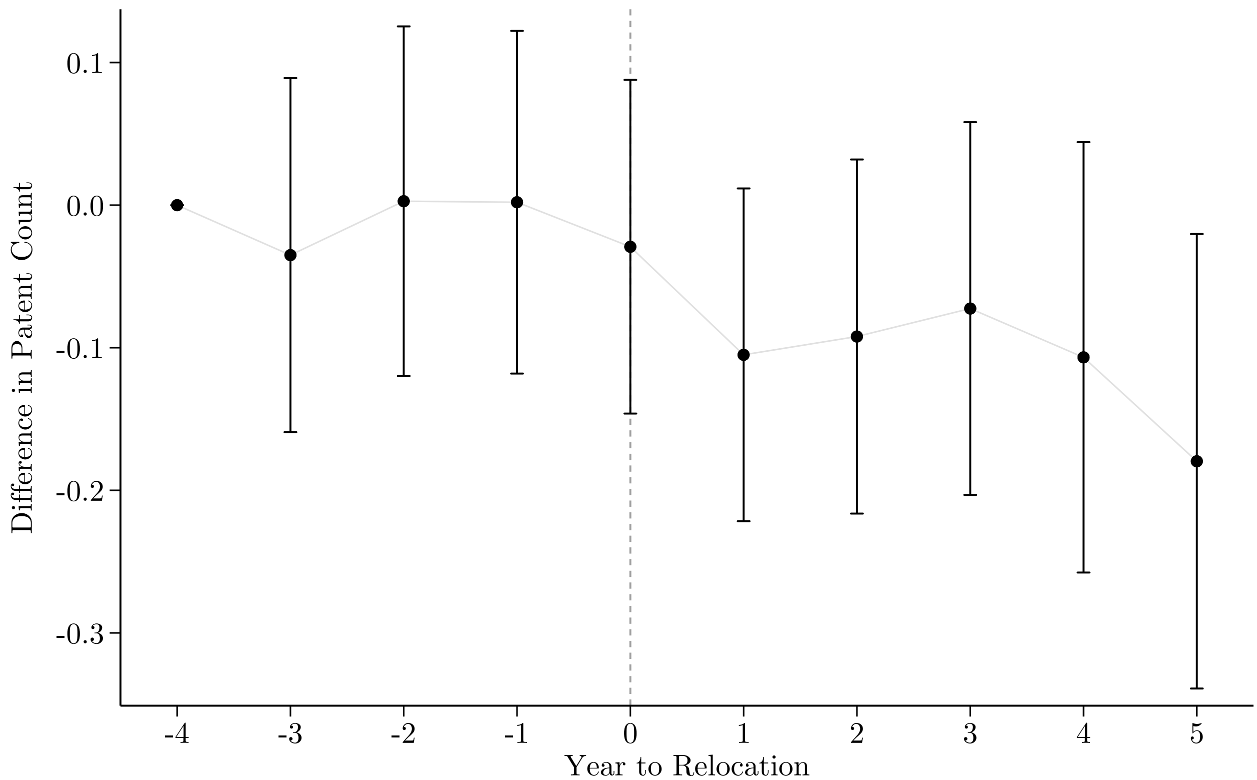

Figure 5: **Effect of Further Commute Relative to Year of Workplace Relocation Event**. This graph shows the effect of a further commute on the number of patents produced by inventors, in each year relative to the year of the workplace relocation event, represented by the vertical gray line at Year 0. Inventors experiencing a closer or identical commute than before the relocation form the control group. The plotted coefficients are $\beta_{ity}^{Further}$ from the following equation for inventor i, firm j, year relative to workplace relocation event y, and year t:

$$\begin{aligned} Y_{ijyt} = \beta_{ity}^{Further} Further\ Commute_{it} * \eta_y \\ + \alpha_{ij} + \eta_y + \gamma_t + \delta_{jt} + \epsilon_{ijt}. \end{aligned} \tag{10}$$

where $Further\ Commute_{it}$ is equal to one for inventor-firm pairs for whom the workplace relocation increased the workplace-home geodesic distance by 1km or more, and 0 otherwise. η_y is a fixed effect, taking the value of 1 for the year $y \in [-4, 5]$ relative to workplace relocation event and 0 otherwise. α_{ij} is an inventor-firm fixed effect, γ_t is a year fixed effect, and δ_{jt} is a firm location fixed effect. The indicator variable for the first year $\eta_{y=4}$ was necessarily omitted for estimation tractability, but shown here as the baseline term, set at a value of 0, i.e., equating the treatment and control groups.

Tables

Table 1: **Summary Statistics for Non-Moving Inventors over Sample Period**. This table presents the mean annual patent count and scaled citations per inventor-firm pair over the full sample period, in the pre-relocation period alone and in the post-relocation period alone. It also presents various measures of commuting distances both before and after the relocation. Income and home price are only available pre-relocation for our sample of non-movers. Income and Home Price are in 1,000 dollars, distances are in kilometers and duration is in minutes.

	Mean	Standard Deviation	Minimum	Median	Maximum	# Inventor-Firm Pairs
Overall						
Patent Count	0.524	0.622	0.000	0.333	6.714	3,445
Scaled Citations	0.893	2.017	0.000	0.288	51.256	3,445
Pre-Relocation						
Patent Count	0.858	0.879	0.000	0.600	8.333	3,445
Scaled Citations	1.365	2.561	0.000	0.491	30.892	3,445
Geodesic Distance	21.456	16.166	0.000	17.184	98.813	3,445
Drive Distance	29.244	21.201	0.000	23.952	119.146	3,445
Drive Duration	26.347	14.207	0.000	23.167	93.150	3,445
Income	125.719	107.599	5.000	105.000	2,920	1,549
Home Price (Buyer)	449.633	349.983	13.500	355.750	3,920	1,596
Post-Relocation						
Patent Count	0.328	0.654	0.000	0.000	9.000	3,445
Scaled Citations	0.633	2.463	0.000	0.000	81.801	3,445
Geodesic Distance	21.607	15.919	0.001	17.463	99.908	3,445
Drive Distance	29.572	21.059	0.000	24.038	120.612	3,445
Drive Duration	26.387	13.997	0.000	22.950	92.433	3,445

Table 2: **Summary Statistics for Relocating Establishments**. This table presents descriptive statistics for relocating establishments at the time of relocation for how far the establishment has moved, the change in establishment's relative distance to central city, and the number of non-moving inventors per establishment. Employee counts are averaged over the entire sample period, and separately over the pre- and post-relocation periods. Distances are in kilometers.

	Mean	Standard Deviation	Minimum	Median	Maximum	Number of Establishments
Overall						
Absolute Relocation Distance	23.731	38.548	1.000	8.186	127.360	1,068
ΔDistance To Central City	2.554	14.006	-59.539	1.234	119.937	1,068
Number of Non–Moving Inventors	3.226	9.267	1.000	2.000	227.000	1,068
Employee Count	185.255	567.495	1.000	50.000	13,331.700	1,068
Pre-Relocation						
Employee Count	167.646	483.593	1.000	37.000	7,600.000	1,068
Post-Relocation						
Employee Count	197.562	640.314	1.000	52.929	15,895.880	1,068

Table 3: **Data Composition by CSA**. This table presents the observation count of *Inventors* and *Firms* by Combined Statistical Area (CSA). The percentages of CSAs in each column are in parentheses. The six largest CSAs by population are shown, ordered by the count of *Inventors*, while the rest are grouped into *Other*.

	Final Sample		All Inventors
Combined Statistical Area	Inventors	Firms	Inventors
San Jose-San Francisco-Oakland	1,168	341	146,631
(CA)	(34%)	(29%)	(26%)
Boston-Worcester-Providence	515	180	67,922
(MA-RI-NH-CT)	(15%)	(15%)	(12%)
Los Angeles-Long Beach	363	126	60,503
(CA)	(11%)	(11%)	(11%)
Chicago-Naperville	205	72	46,188
(IL-IN-WI)	(6%)	(6%)	(8%)
New York-Newark	160	92	84,529
(NY-NJ-CT-PA)	(5%)	(8%)	(15%)
Washington-Baltimore-Arlington	133	65	33,091
(DC-MD-VA-WV-PA)	(4%)	(6%)	(6%)
Other	901	304	124,117
	(26%)	(26%)	(22%)
Total	3,445	1,180	562,981
	(100%)	(100%)	(100%)

Table 4: **Balance Test by Direction of Distance Shocks**. Means presented with standard errors in parentheses and number of inventor-firm pairs in brackets. Closer/Same/Farther columns indicate sub-sample for which commuting distance decreased by more than 1km, approximately stayed the same (i.e., changed by less than 1km), and increased by more than 1km, respectively. Income and home prices are in 1,000 dollars, distance measures are in kilometers and duration is in minutes.

	Closer	Same	Farther	p-value		
Variable	(1)	(2)	(3)	(1 vs 3)	(1 vs 2)	(2 vs 3)
Patent Count	0.850	0.846	0.869	0.547	0.928	0.591
	(0.872)	(0.834)	(0.897)			
	[1348]	[529]	[1569]			
Scaled Citations	1.380	1.272	1.383	0.970	0.415	0.372
	(2.689)	(2.247)	(2.548)			
	[1348]	[529]	[1569]			
Income	131.881	130.837	118.807	0.035	0.918	0.026
	(141.5)	(81.646)	(68.505)			
	[599]	[239]	[711]			
Home Price (Buyer)	454.815	459.404	441.901	0.496	0.863	0.494
	(350.046)	(368.774)	(343.541)			
	[614]	[252]	[730]			
Geodesic Distance	26.47	19.38	17.85	0.000	0.000	0.041
	(16.82)	(15.34)	(14.70)			
	[1,348]	[529]	[1,569]			
Drive Distance	35.42	26.63	24.82	0.000	0.000	0.070
	(21.76)	(20.25)	(19.71)			
	[1,348]	[529]	[1,569]			
Drive Duration	30.49	24.84	23.30	0.000	0.000	0.025
	(14.24)	(13.85)	(13.42)			
	[1,348]	[529]	[1,569]			

Table 5: **Effect of Commuting on Inventor Productivity - Quantity**. The dependent variable ***Patents*** is the count of patents granted to an inventor-firm per year. The independent variable of workplace-home ***Distance*** is geodesic distance measured in 10 kilometers. OLS regression model with robust standard errors clustered at the inventor-firm pair-level shown in parentheses. R^2 includes both within- and between- variation.

Dependent Var.: *Patents*	(1)	(2)	(3)	(4)	(5)
Distance	-0.013*	-0.039**	-0.030**	-0.041**	-0.043**
	(0.007)	(0.016)	(0.014)	(0.019)	(0.019)
Inventor-Firm Pair FE	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Firm Location (ZCTA) FE	No	No	No	Yes	No
Firm X Location FE	No	No	No	No	Yes
R^2	0.000	0.335	0.386	0.415	0.388
Inventor-Firm Count	3,445	3,445	3,445	3,445	3,445
Observations	22,917	22,917	22,917	22,917	22,798

Table 6: **Effect of Commuting on Inventor Productivity - Quality**. The dependent variables are adjusted patent quality measures per inventor-firm pair per year. The independent variable of workplace-home ***Distance*** is geodesic distance measured in 10 kilometers. OLS regression model with robust standard errors clustered at the inventor-firm pair-level shown in parentheses. R^2 includes both within- and between- variation.

Dependent Var.:	*Scaled Citation*	*Generality*	*Originality*	*Payment Count*
	(1)	(2)	(3)	(4)
Distance	-0.094*	-0.025**	-0.013**	-0.047
	(0.049)	(0.011)	(0.006)	(0.042)
Mean of Pre-Relocation D.V.	1.365	0.399	0.206	1.684
Inventor-Firm Pair FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Firm Location FE	Yes	Yes	Yes	Yes
R^2	0.381	0.387	0.392	0.417
Inventor-Firm Count	3,445	3,445	3,445	3,445
Observations	22,863	22,863	22,863	22,863

Table 7: **Effect of Commuting on Inventor Productivity - Nonlinear Effects**. The dependent variables are patent counts and adjusted patent quality measures per inventor-firm pair per year. The independent variable of workplace-home ***Distance*** is geodesic distance measured in 10 kilometers. $Distance^2$ is the square of ***Distance***. OLS regression model with robust standard errors clustered at the inventor-firm pair-level shown in parentheses. R^2 includes both within- and between- variation.

Dependent Var.:	*Patents*	*Scaled Citation*	*Generality*	*Originality*	*Fee Payments*
	(1)	(2)	(3)	(4)	(5)
Distance	-0.061	0.067	-0.006	-0.012	-0.184**
	(0.041)	(0.114)	(0.022)	(0.013)	(0.086)
$Distance^2$	0.003	-0.026	-0.003	-0.000	0.022*
	(0.006)	(0.016)	(0.003)	(0.002)	(0.012)
Mean of Pre-Rel. D.V.	0.858	1.365	0.399	0.206	1.684
Inventor-Firm Pair FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Firm Location FE	Yes	Yes	Yes	Yes	Yes
R^2	0.415	0.381	0.387	0.392	0.417
Inventor-Firm Count	3,445	3,445	3,445	3,445	3,445
Observations	22,863	22,863	22,863	22,863	22,863

Table 8: **Effect of Commuting on Inventor Productivity - Top Inventor**. The dependent variables are patent counts and adjusted patent quality measures per inventor-firm pair per year. The independent variable of workplace-home *Distance* is geodesic distance measured in 10 kilometers. *Top Inventor* are inventors in the top decile in terms of average patent count. OLS regression model with robust standard errors clustered at the inventor-firm pair-level shown in parentheses. R^2 includes both within- and between- variation.

Dependent Var.:	*Patents*	*Scaled Citation*	*Generality*	*Originality*	*Payment Count*
	(1)	(2)	(3)	(4)	(5)
Distance	-0.026	-0.057	-0.022*	-0.009	-0.018
	(0.019)	(0.048)	(0.011)	(0.006)	(0.041)
Distance × *Top Inventor*	-0.159**	-0.407**	-0.039	-0.045**	-0.314**
	(0.073)	(0.172)	(0.038)	(0.022)	(0.149)
Mean of Pre-Relocation	0.781	1.264	0.370	0.189	1.533
Other Inventor D.V.	(0.753)	(2.410)	(0.450)	(0.302)	(1.739)
Mean of Pre-Relocation	1.670	2.421	0.708	0.386	3.276
Top Inventor D.V.	(1.491)	(3.641)	(0.767)	(0.522)	(3.241)
Inventor-Firm Pair FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Firm Location FE	Yes	Yes	Yes	Yes	Yes
R^2	0.416	0.381	0.387	0.392	0.418
Inventor-Firm Count	3,445	3,445	3,445	3,445	3,445
Observations	22,863	22,863	22,863	22,863	22,863

Table 9: **Effect of Commuting on Inventor Productivity - Subsamples by Move Direction**. The dependent variable *Patents* is the count of patents granted to an inventor-firm per year. Closer/Same/Farther columns indicate sub-sample for which commuting distance decreased by more than 1km, approximately stayed the same (i.e. changed by less than 1km), and increased by more than 1km, respectively. The independent variable of workplace-home *Distance* is geodesic distance measured in 10 kilometers. *Top Inventor* are inventors in the top decile in terms of average patent count. OLS regression model with robust standard errors clustered at the inventor-firm pair-level shown in parentheses. R^2 includes both within- and between- variation.

Dependent Var.: *Patents*	(1)	(2)	(3)	(4)	(5)	(6)
Distance	0.016	0.038	-0.106**	0.011	0.030	-0.062
	(0.049)	(0.778)	(0.049)	(0.050)	(0.779)	(0.049)
Distance × *Top Inventor*				0.072	-0.282	-0.573***
				(0.087)	(0.174)	(0.147)
Sample	Closer	Same	Farther	Closer	Same	Farther
Inventor-Firm Pair FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm Location FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.421	0.415	0.434	0.421	0.415	0.437
Inventor-Firm Count	1,348	535	1,562	1,348	535	1,562
Observations	8,887	3,512	10,417	8,887	3,512	10,417

Table 10: **Effect of Workplace-Home Distance on Inventor Productivity - Solo Patents**. The dependent variable *Solo Patents* is the count of patents granted to an inventor-firm pair per year, where the inventor is the sole inventor. The independent variable of workplace-home *Distance* is geodesic distance measured in 10 kilometers. *Top Inventor* are inventors in the top decile in terms of average patent count. OLS regression model with robust standard errors clustered at the inventor-firm pair-level shown in parentheses. R^2 includes both within- and between- variation. The average pre-relocation single-authored patent count is 0.103 per inventor-firm pair per year.

Dependent Var.: *Solo Patents*	(1)	(2)	(3)	(4)	(5)
Distance	-0.000	-0.010**	-0.009**	-0.012**	-0.009
	(0.002)	(0.004)	(0.004)	(0.006)	(0.006)
Distance × *Top Inventor*					-0.036*
					(0.022)
Inventor-Firm FE	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Firm Location FE	No	No	No	Yes	Yes
R^2	0.000	0.334	0.342	0.375	0.375
Inventor-Firm Count	3,445	3,445	3,445	3,445	3,445
Observations	22,917	22,917	22,917	22,863	22,863

Table 11: **Effect of Distance Direction Change on Inventor Productivity**. The dependent variable *Patents* is the count of patents granted to an inventor-firm per year. The independent variable Δ *Distance Direction* is a categorical variable equalling 0 prior to a workplace relocation event, and after a relocation event, takes the value of 1 for geodesic distance increases (farther away by > 1km), -1 for distance decreases (closer by > 1km), and 0 when distance is approximately unchanged (changed by less than 1km). OLS regression model with robust standard errors clustered at the inventor-firm pair-level shown in parentheses. R^2 includes both within- and between- variation.

Dependent Var.: *Patents*	(1)	(2)	(3)	(4)	(5)
Δ *Distance Direction*	-0.043***	-0.097***	-0.056***	-0.072***	-0.050**
	(0.016)	(0.021)	(0.019)	(0.022)	(0.020)
Δ *Distance Direction* × *Top Inventor*					-0.199*
					(0.110)
Inventor-Firm FE	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Firm Location FE	No	No	No	Yes	Yes
R^2	0.001	0.339	0.396	0.429	0.430
Inventor-Firm Count	3,445	3,445	3,445	3,432	3,432
Observations	19,472	19,472	19,472	19,365	19,365

Table 12: **Effect of Commuting on Inventory Productivity - Alternate Commuting Measures**. The dependent variables are patent counts and scaled citations per inventor-firm pair per year. The independent variable ***Drive Distance*** is the workplace-home driving distance (e.g., by car) in 10 kilometers. The independent variable ***Drive Duration*** is the time it takes to drive or take public transportation between the workplace and home, whichever is less, measured in 10 minutes. OLS regression model with robust standard errors clustered at the inventor-firm pair-level shown in parentheses. R^2 includes both within- and between-variation.

Dependent Var.:	*Patents*	*Scaled Citations*	*Patents*	*Scaled Citations*
	(1)	(2)	(3)	(4)
Drive Distance	-0.038**	-0.074*		
	(0.015)	(0.038)		
Drive Duration			-0.059**	-0.104*
			(0.024)	(0.059)
Inventor-Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Firm Location FE	Yes	Yes	Yes	Yes
R^2	0.416	0.381	0.416	0.381
Inventor-Firm Count	3,445	3,445	3,445	3,445
Observations	22,863	22,863	22,863	22,863

Table 13: **Effect of Commuting on Inventor Productivity - Count Models**. Columns (1) and (2) estimate Poisson regression with conditional inventor-firm pair fixed effects. Columns (3) and (4) estimate negative binomial regression with conditional inventor-firm pair fixed effects. The independent variable of workplace-home ***Distance*** is geodesic distance measured in 10 kilometers. Coefficients are presented as incidence-rate ratios. Robust standard errors clustered at the inventor-firm pair-level are shown in parentheses.

Dependent Var.: *Patents*	Poisson Estimation		Negative Binomial	
	(1)	(2)	(3)	(4)
Distance	-0.038**	-0.049**	-0.041**	-0.054**
	(0.016)	(0.020)	(0.016)	(0.023)
Inventor-Firm Cond. FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Firm Location FE	No	Yes	No	Yes
Pseudo R^2	0.275	0.306	0.207	0.236
Inventor-Firm Count	3,445	3,445	3,445	3,445
Observations	22,917	22,917	22,917	22,917

Table 14: **Effect of Commuting on Inventor Moving and Quitting**. The dependent variable *Inventor Move* takes a value of 1 if the inventor moves homes in a year, and 0 otherwise. The dependent variable *Inventor Quit* takes a value of 1 if the inventor departs the focal firm in a year, and 0 otherwise. The independent variable *Post Relocation* equals 1 for observations after and including the year of a relocation of the inventor's workplace, and 0 before. Δ *Move Direction* is a categorical variable equalling 0 prior to a workplace relocation event, and after a relocation event, takes the value of 1 for geodesic distance increases (farther away by $>$ 1km), -1 for distance decreases (closer by $>$ 1km), and 0 when distance is approximately unchanged (changed by less than 1km). OLS regression model with robust standard errors clustered at the inventor-firm pair-level shown in parentheses.

	Inventor Move		*Inventor Quit*	
Variable	(1)	(2)	(3)	(4)
Post Relocation	0.014***	0.015***	0.017***	0.017***
	(0.003)	(0.003)	(0.006)	(0.006)
Post Relocation		-0.002		0.011*
$\times$*Move Direction*		(0.002)		(0.006)
Inventor-Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
R^2	0.259	0.259	0.274	0.274
Inventor-Firm Count	12,047	12,047	5,158	5,158
Observations	81,563	81,563	29,624	29,624

A Appendix

A.1 Data Sample Composition and Construction

A.1.1 DataQuick

Ferreira and Gyourko (2015) provide more details on the time horizon (i.e., start and end dates) of the sample. The combined statistical areas (CSA) included in the table are presented in Table A.1. Not all parts of each CSA are included.

The DataQuick sample does not include information on home renters, or homeowners who neither moved nor refinanced their mortgages between 1993 and 2012. Thus, we were not able to match these inventors to their home location, and they were not included in our study.

——————Insert Table A.1——————

A.1.2 InfoUSA Data Quality Assessment

To confirm the completeness and validity of InfoUSA establishment data, we compare it to the commonly-used County Business Patterns (CBP) data produced by the U.S. Census Bureau (e.g., Duranton and Turner, 2012).

——————Insert Figure A.1——————

InfoUSA has more comprehensive establishment coverage than CBP, covering both more establishments and more employees. In particular, InfoUSA covers substantively more small establishments, e.g., establishments with between one and four employees, than CBP, while for large establishments and total employees the difference is much smaller. This pattern is consistent with known under-coverage in the CBP data: given its focus on larger firms, the U.S. Census Bureau does not readily update data on smaller multi-unit companies (in terms of employee count), and it may miss establishments for smaller firms and firms that do not respond to their surveys (U.S. Census Bureau, 2018).[25]

A.1.3 Detailed Matching and Sample Construction Procedure

Matching Homeowners to Inventors. The first step in the sample construction process is to identify the homes of inventors based upon the residential history given in the DataQuick data. We match home buyer names to inventor names by city within a given CSA. We exactly match city names and inventor first and last names. Middle names are problematic since some transaction records list middle names in full, some only list middle initials, and others list nothing. Given this inconsistency, we only require that there is no discrepancy in a match. For example, in matching middle initials to middle names, we require consistency between the middle initials and first letter of a middle name. When there are multiple names listed on the same transaction, we consider each buyer separately.

To identify the time period for which a home is owned by an inventor, i.e., a residential spell, we match home buyer names from a home transaction at a specific address to the home seller names in the next transaction of that property. As before, we match first and last names exactly and require consistency in middle names. If none of the buyers from the previous transaction match any of the sellers from the following transaction, we drop the home from our sample because we cannot verify the residential spell. If there are no subsequent transaction at the property, we assume that the inventor residential spell lasts until the end of the sample period (i.e. 2012). Similarly, we match home seller names to inventor names, and match names with home buyers from the previous transaction to determine residential spells. If there are no previous transaction at the property, we assume that the inventor residential spell starts from the beginning of the sample period.

This process identifies residential spells for three groups of inventors: inventors who bought and sold houses during the sample period; inventors who bought and kept their house until the end of the sample period; and inventors who owned houses before and sold them

[25] There are no U.S. Census Bureau estimates of establishment under-coverage.

during the sample period. We then eliminate inventors for which we cannot identify their primary residence, such as inventors with matched to multiple different addresses in the same city in the same year.

Out of a total of around 563,000 inventors, we obtain a sample of 264,287 inventors with home addresses across the 60 CSAs sampled. In this final sample, we match 47% of all inventors to their home address. The three primary limitations to matching were ambiguous homeowner names, non-owning residents, and limited mortgage activity. First, we had to drop many of the observations associated with the more than 20% of all homeowner names in a city have duplicates (see Figure A.4). Second, while more than 30% of the total US population rents, we were not able to identify the location of inventors renting a property. We also could not match family members who live in a home but are not among the listed homeowners. Third, homeowners who have not moved or refinanced their mortgage during our sampling period do not appear in the original DataQuick data.

Matching Workplaces to Inventors. We then obtain the workplace addresses of the inventors, assuming that the inventor is employed by the patent assignee, through fuzzy matching and manually matching patent assignee names to InfoUSA firm names in all CSAs. InfoUSA serves as the source of the workplace addresses to be matched with the assignees named in the patent dataset.

Within these two datasets, we consolidate and standardize firm names within dataset that likely belong to the same firm via fuzzy matching and manual matching. In both the patent data and the InfoUSA data, firm names appear inconsistently, due to misspelling names, abbreviations or acronyms in place of full firm names, and changes in firm names over time due to mergers or other corporate activity. After fuzzy matching, two research assistants then manually and independently verify the accuracy of name groupings using outside public information. To identify firms which changed their name, the research assistants manually search and identify alternative names for all firms with patents in a given CSA, then convert alternative names to the most commonly appearing version of the firm name.

We use the most common assignee name in the patent data as input to the fuzzy matching algorithm. After fuzzy matching the inventor assignees with InfoUSA firms, a team of research assistants manually confirms the matches. We retain all matched firms with patents in any CSA, leaving us with an inventor-workplace matched sample of 36,468 observations.

Identifying Primary Workplace. For firms with multiple office locations in a CSA, we designate the primary office location using two criteria. First, we select the office location with more than five times the number of employees of all the other office locations combined. Second, if the first condition is not met, we select locations designated as a "research laboratory" in its NAICS code, on the assumption that research and development is likely to take place there; this condition is only satisfied if there is a single "research laboratory" location. After assessing these two criteria, we then drop firms where neither of these conditions is satisfied for its workplaces in a CSA. After this process of identifying a primary workplace, we are left with 35,836 single-location firms with matched inventors.

Identifying Workplace Relocations. We define workplace relocations as changes in the primary office location of a firm within a CSA from one year to the next where the geodesic distance between the old and new locations is at least one kilometer. We set the lower relocation distance threshold to increase the power of our estimates by excluding trivial moves between units in the same building or complex. 11,160 firms relocated during our sample period out of a total of 35,836 single-location firms with identifiable matched inventors. We further restrict this sample of relocations for considerations related to our observation time window. We eliminate relocations that last for 2 years or less before going back to the original location, since it is less likely that that these firms undertook an actual physical relocation. To ensure pre- and post-relocation observation periods, necessary for our empirical design, we also only retain observations occurring after the first or before the last year of our sample. These two criteria leave us with 6,944 relocating firms with identifiable inventors.

Final Steps. We match our sample of relocating firms with inventor variables (i.e., inventor-year patent count). To ensure pre- and post-relocation observation periods, we limit our sample of inventors to those who worked at the firm both one year before and one year after the relocation. This leaves us with 3,723 inventor-firm pairs working at 1,220 relocating firms. The last step is to drop outliers with commuting distances that are too large to be believable, and which can bias our results. We also winsorize inventor productivity at the 2% level. This leaves us with a final sample of 22,917 inventor-firm-year observations for 3,445 inventor-firm pairs, employed at 1,068 relocating firms.

A.2 Illustrative Model

Consider a city with firms located at the center. All residents in the city work at the center, and can live at two locations: location A is closer to the firms, while location B is farther away. No matter where they live, all residents commute to the city center for work. There are two types of workers: skilled workers and unskilled workers, with the former being more productive than the latter at both locations.

More specifically, let productivity be $l_h^A = \theta$ and $l_h^B = (1 + \alpha_h)\theta$ for skilled workers living at locations A and B, respectively, as well as $l_l^A = \beta\theta$ and $l_l^B = (\beta + \alpha_l)\theta$ for unskilled workers. Our model is general and incorporates the effects of different underlying mechanisms that link distance with productivity. For example, both commuting and knowledge spillover effects could imply that productivity declines with distance, in which case $\alpha_h, \alpha_l < 0$. If there is an optimal distance that separates work and home, as some of the work-home life separation literature implies, we would have $\alpha_h, \alpha_l > 0$. For simplicity, we assume that, overall, there is a negative total effect of distance on productivity for the rest of this section, though derivations with positive total effects are analogous. We also abstract away from the underlying mechanisms that produce this relationship,[26] and focus on how worker sorting can interact with this negative correlation between distance and productivity and cause bias

[26] Potential mechanisms include having less time at work, lower optimal effort provision with moral hazard, less schedule flexibility, fewer interactions and less knowledge spillover with coworkers, behavioral mechanisms, etc.

in our estimation.

Firms compete in a competitive market and all pay a unit wage w based on performance. So each worker i's take-home wage is wl_i. Workers are risk-neutral and maximize a money-utility:

$$U_i = wl_i - p(h_i) \tag{A.1}$$

where $p(h_i)$ is the price of housing. We assume that each worker demands the same quantity of housing, so only the latter's price enters into his utility function. Assume that all housing stock belongs to absentee owners who try to maximize their revenue, subject to their tenants' participation constraints. To make the model more tractable, we make some further assumptions: the housing stock is evenly divided between locations A and B, the worker types are equal in number, and the total number of housing units equals the total number of workers, so the housing market can clear. There are no moving costs so workers can change locations costlessly.

Given that $\alpha_h, \alpha_l < 0$, there are three potential cases we need to consider. $\alpha_h = \alpha_l$, where productivity for both types of inventors declines at the same rate with distance, $\alpha_h > \alpha_l$, where productivity for unskilled workers declines faster with distance than for skilled workers, and $\alpha_h < \alpha_l$, where productivity for unskilled workers declines slower with distance than for skilled workers.

Then, there are three potential equilibria: two separating equilibria, with each type congregating in one location separately from the other type's location, and one pooling equilibrium with a mixture of both skilled and unskilled workers at both locations. We show later in this section that there is a unique equilibrium for each of the three cases above: if $\alpha_h = \alpha_l$, then the pooling equilibrium occurs. If $\alpha_h > \alpha_l$, then the unique equilibrium is a separating equilibrium with unskilled workers living at location A, closer to the city center than skilled workers at location B. If $\alpha_h < \alpha_l$, then the separating equilibrium has unskilled workers living at location B, farther away from skilled workers living at location A.

In both cases where $\alpha_h \neq \alpha_l$, the observed effect of distance on productivity would be different from the true effect. As shown in Figure A.2, when $\alpha_h < \alpha_l$, the observed effect is more negative than either of the true slopes for skilled or unskilled workers, while for the case where $\alpha_h > \alpha_l$ in Figure A.3, the observed effect is less negative than either of the true slopes for skilled or unskilled workers, and could even be positive.

——————Insert Figure A.2——————

——————Insert Figure A.3——————

This means that to identify the causal effect of distance on productivity, we need an estimation strategy that separates the causal effect from worker sorting based on skill levels.

The rest of this section goes through the proofs to show that there are unique equilibria for each of the three cases in the illustrative model.

A.2.1 Case I: Faster Productivity Decline for Skilled Workers

In this case, $\alpha_h < \alpha_l$. Housing costs at both locations can be derived as follows. First, at any location where some unskilled workers live, housing is always priced at what the unskilled workers can bear. Therefore, in the separating equilibrium with skilled workers at location A and unskilled workers at location B, the housing price at location B is $p^B = wl_l^B$ and unskilled workers get zero utility, which equals the modeled outside option. Then, to determine the housing price at location A, we can use the incentive-compatibility constraint for skilled workers because, for there to be an equilibrium, skilled workers must not want to move to location B:

$$\begin{aligned} wl_h^B - p^B &\le wl_h^A - p^A \\ w(1+\alpha_h)\theta - p^B &\le w\theta - p^A \\ w\alpha_h\theta - p^B &\le p^A \\ p^A &= p^B - w\alpha_h\theta \end{aligned} \tag{A.2}$$

where the last equality comes from profit maximization of the absentee homeowners. This price differential satisfies the incentive-compatibility constraint for unskilled workers since it is higher than the price difference that they can afford, $-w\alpha_l\theta$, given that $\alpha_h < \alpha_l$. Therefore, the separating equilibrium with skilled workers living closer to the firm is possible.

To show that this equilibrium is unique, we show that the other two potential equilibria cannot hold. If there is a separating equilibrium with skilled workers living farther away at location B, and unskilled workers living at A, then housing price at A would be $p^A = wl_l^A$. To determine the housing price at B, we can use the incentive-compatibility constraint for skilled workers because for there to be an equilibrium, skilled workers must not want to move to location A:

$$\begin{aligned} wl_h^B - p^B &\geq wl_h^A - p^A \\ w(1+\alpha_h)\theta - p^B &\geq w\theta - p^A \\ w\alpha_h\theta - p^B &\geq p^A \\ p^B &= p^A + w\alpha_h\theta \end{aligned} \tag{A.3}$$

where the last equality comes from profit maximization of the absentee homeowners. Then, however, the unskilled worker's IC constraint cannot also be satisfied, because he would get a utility gain of $w(\alpha_l - \alpha_h)\theta > 0$ if he moved from location B to A.

Similarly, pooling is not an equilibrium either because the skilled and unskilled worker's IC constraints cannot be satisfied at the same time. More specifically, housing price would be equal to the unskilled worker's willingness-to-pay at both A and B, but this gives skilled workers an incentive to move to location A, where they can earn higher utility than in location B.

A.2.2 Case II: Slower Productivity Decline for Skilled Workers

In this case, $\alpha_h > \alpha_l$. Therefore, in the separating equilibrium with skilled workers at location A and unskilled workers at location B, the housing price at location B is $p^B = wl_l^B$

and unskilled workers get zero utility, which equals the modeled outside option. Then, to determine the housing price at location A, we can use the incentive-compatibility constraint for skilled workers because, for there to be an equilibrium, skilled workers must not want to move to location B:

$$
\begin{aligned}
wl_h^B - p^B &\leq wl_h^A - p^A \\
w(1+\alpha_h)\theta - p^B &\leq w\theta - p^A \\
w\alpha_h\theta - p^B &\leq p^A \\
p^A &= p^B - w\alpha_h\theta
\end{aligned}
\tag{A.4}
$$

where the last equality comes from profit maximization of the absentee homeowners. Then, however, the unskilled worker's IC constraint cannot also be satisfied, because he would get a utility gain of $w(\alpha_h - \alpha_l)\theta > 0$ if he moved from location B to A.

If there is a separating equilibrium with skilled workers living farther away at location B, and unskilled workers living at A, then housing price at A would be $p^A = wl_l^A$. To determine the housing price at B, we can use the incentive-compatibility constraint for skilled workers because, for there to be an equilibrium, skilled workers must not want to move to location A:

$$
\begin{aligned}
wl_h^B - p^B &\geq wl_h^A - p^A \\
w(1+\alpha_h)\theta - p^B &\geq w\theta - p^A \\
w\alpha_h\theta - p^B &\geq p^A \\
p^B &= p^A + w\alpha_h\theta
\end{aligned}
\tag{A.5}
$$

where the last equality comes from profit maximization of the absentee homeowners. This price differential satisfies the incentive-compatibility constraint for unskilled workers since they would have negative utility living at B, given that $\alpha_h > \alpha_l$. Therefore, the separating equilibrium with skilled workers living farther from the firm is possible.

Again, pooling is not an equilibrium either because the skilled and unskilled worker's IC constraints cannot be satisfied at the same time. More specifically, housing price would be equal to the unskilled worker's willingness-to-pay at both A and B, but this gives skilled workers an incentive to move to location B, where they can earn higher utility than in location A.

A.2.3 Case III: Same Productivity Decline for Both Types of Workers

In this case, $\alpha_h = \alpha_l$. The pooling equilibrium is possible with housing price equal to the unskilled worker's willingness-to-pay at both locations A and B. In this case, the skilled worker would be indifferent about living at either location, since she will get the same utility of $wl_h^A - p^A = wl_h^B - p^B = w(1-\beta)\theta$. This is the unique equilibrium because in both potential separating equilibria skilled workers are indifferent about living in both locations.

A.3 Name Weights

One important source of measurement error in our regressions is the presence of duplicate names during our matching process. Since we match inventor data to housing transactions data at the city level, there could be multiple people with identical names who live in the same city, which can cause mismatches. Given that neither of our datasets contains a census of all people living in a city, there are three potential ways a mismatch can occur: there could be multiple inventors living in the same city, there could be multiple homeowners participating in a housing transaction during the DataQuick sample period, and there could be a unique inventor matched to a unique homeowner in our sample, but who in fact is a duplicate. The two first cases are observable in our data, and we mitigate their occurrence by dropping them. However, the last case cannot be detected directly, so it is important to gauge the percentage of mismatches there could be in our data, to quantify the amount of noise they introduce in our results. (There could be bias if probability of having an unobservable duplicate is correlated with both patent production and change in commuting distance. We consider this to be unlikely.)

To model how likely it is to encounter duplicate names without knowing it, we assume that the probability of person i living in city j having another person with the same name living in the same city to be Q_{ij}. Then Q_{ij} depends on the following factors:

Total population in the city N_j: The more people living in city j, the higher Q_{ij} becomes, assuming that each additional person has some independent probability of having an identical name.

Probability that a person has the same last name P_{ij}^l: Since one's last name often depends on the ethnic group he belongs to, this factor is the product of two subfactors: the percentage of people in the city belonging to the same ethnic group, and how common the last name is within said ethnic group.

Probability that a person has the same first name and middle initial P_{ij}^f: The choice of a first name is not as strongly linked to one's ethnic group, so we assume that this probability to be proportional to the overall frequency of first names in a given state (or the entire United States).

$$Q_{ij} = N_j P_{ij}^l P_{ij}^f \tag{A.6}$$

Given that we do not observe the entire population, more assumptions are needed to empirically estimate the probability that duplicate names occur. First, we assume that the distribution of first and last names is similar between the observable and unobservable parts of the population; then the only difference in Q_{ij} for person i in city j between what we can estimate and the true value is the number of people.

Second, we assume that the distribution of first and last names are identical across all cities. In other words, P_{ij}^l and P_{ij}^f no longer depend on j, so we can use cross-sectional data without worrying about compositional differences between cities biasing our estimates. Then, we can plot a meaningful graph of the probability of having another person with the same full name living in the same city against total population using cross-city data. Of course, what we estimate is an upper bound, because the same person in the DataQuick database could have made multiple housing purchases.

——————**Insert Figure A.4**——————

The results of this analysis are shown in Figure A.4, where the mean probability of finding another buyer with an identical name is plotted against the number of buyers in a city. This graph shows that, even in the largest cities, less than 40% of all home buyers over a 20-year period (1992 to 2012) have another buyer with identical names, who could very well be the same person buying another property. Therefore, our results should not suffer too much attenuation bias from duplicate names.

Nonetheless, as a robustness check, we estimate our main specifications by weighting observations using the following name weights:

$$w = \frac{1}{1 + Number\ of\ identical\ buyernames\ in\ CSAinDataQuickdata} \tag{A.7}$$

Thus, if an inventor's name only appears once, we assign him or her a weight of one, whereas if there are three more inventors with same first and last names in the CSA, the weight becomes 0.25. Results are shown in Table A.2. All the estimated coefficients for patent quantity and quality are still negative and significant, consistent with results in the main text.

——————**Insert Table A.2**——————

A.4 Additional Robustness Checks

In this section we present additional robustness checks for our main specification.

A.4.1 Subsample: Firm-level Variables

We test the robustness of our results by including more firm-level control variables. In add firm-level financial information, we manually search all the 1,068 firms used in our main analysis in the Compustat data and find financial information for the 405 firms (38%) that have ever gone public. Based on the Compustat data, we control for firm turnover, market value and total assets. From the Infogroup data, we both the number of employees and sales volume at the establishment level. We log transform all the control variables because their distributions are skewed to the right, and 0.01 is added to their original values so that the zero values are not missing in the log transformation. The total number of matched observations is 8,649.

Table A.3 shows the results of patent quantity and quality regressions in this publicly-traded firm subsample. The estimated coefficients of commuting on inventor productivity are negative and significant, consistent with results from our full sample.

——————Insert Table A.3——————

A.4.2 Subsample: Large Establishment

One potential endogeneity concern is reverse causality: small firm establishments could move towards employees who they anticipate will become more productive. In this case, our estimate would confound the causal effect of decreasing distance on inventor productivity, with the anticipated increase in inventor productivity that causes the firm to move in the first place, and our estimates would be biased upward. In the main text, we already show that this is unlikely since firms are not moving towards inventors who are becoming more productive, with no differential pre-trend in productivity between inventors getting shocked in opposite directions. Furthermore, using a stacked difference-in-differences specification means that many firms need to anticipate exactly the year at which their inventors will become more productive in patenting. We believe this is unlikely.

Nevertheless, we further test the robustness of our results by estimating the baseline regressions using a subsample of inventors working at large establishments, which we define to be establishments with an average of more than 100 employees during our sample period. Large establishments are less likely to move in anticipation of increased productivity of any particular employee, due to their own higher moving costs. The results in Table A.5 show that the negative effect of ***Distance*** on inventor productivity in this large establishment subsample is even larger in magnitude than in our baseline specification. [NEED TO CHECK WHETHER THIS IS DUE TO HIGHER MEAN PATENT COUNT FOR THIS SUBSAMPLE]

——————Insert Table A.5——————

A.4.3 Subsample: Bounded Commute Distance

Another potential worry for our estimation is that firm-inventor pairs with higher workplace-home distances could more likely be mismatches. Since inventors who lived farther away before firm relocations are more likely to be moved closer in, this potential measurement error could be correlated with the observed distance and cause bias in our estimation. To check against this potential bias, we estimate the baseline regressions using a subsample of inventor-firm pairs whose maximum distance is always smaller than 50km, which is the 90th percentile of all commuting distances in the United States (U.S. Government Bureau of Transportation Statistics, 2003). The results in Table A.6 show that the negative effect of ***Distance*** on inventor productivity is almost unchanged and still highly significant, compared against baseline regressions.

——————Insert Table A.6——————

A.4.4 Subsample: Without San Francisco Bay Area

Almost a third of all our observations come from the San Francisco Bay Area, as shown in Table 3. We test whether our results are driven by inventors living in the Bay Area by estimating our main specification without them. Results are shown in Table A.4 and we still find a consistently significant negative effect of commuting on productivity. In regressions that are not shown, we verify that our results are robust to the exclusion of inventors from any particular large CSAs.

——————Insert Table A.4——————

A.5 Patent Valuation

Table A.7 summarizes prior estimates of the mean value of a patent in the United States as documented in the literature. These studies utilize a diverse set of methodologies and underlying data, across many years of study. The nominal U.S. dollar estimates of patent value given in each paper are adjusted to 2010 U.S. dollars using the Consumer Price Index (CPI) provided by the U.S. Bureau of Labor Statistics. Indexed to 2010 U.S. dollars, the estimates for the value of a patent range from $102,844 (Serrano, 2005) to $4,225,487 (Pakes, 1985).

From the wide range of existing estimates, we use a conservative assumption of $100,000 USD, a lower bound below all these existing estimates, as the value of an average patent. This assumption can link our estimates of the impact of workplace-home distance on patent productivity to economic value generated for the employer assigned the patent.

——————**Insert Table A.7**——————

A.6 Patent Maintenance Fee Data

This section describes patent maintenance fees in more detail.

For estimates of the economic value of individual patent, we use patent maintenance fee data from 1980 to 2019, available from the USPTO. Patent maintenance fee data have previously been used to estimate the economic value of individual patents, notably by Pakes (1986) and Bessen (2008). From the USPTO website: "Maintenance fees are required to keep in force all utility and reissue utility patents based on applications filed on or after December 12, 1980." In short, to keep patents in force, patent owners need to pay maintenance fees 3.5, 7.5, and 11.5 years after the date of patent grant.

From a profit-maximizing perspective, a patent owner will only pay the maintenance fee if the patent's present value over the next four years (or remaining term for the last fee payment) plus the option value of future renewals is greater than the cost of the maintenance fee. The fee schedule is steeply increasing over the patent term, ensuring that the value of maintained patents monotonically increases over the number of fee payments.[27] Therefore, we expect that firms are willing to pay the maintenance fee longer for the patents that they think are more valuable.

We create a variable for the number of maintenance fee payments for each patent. For all patents applied for and granted between 1980 and 2005, Table A.8 shows the distribution of total maintenance fee payment counts. Even after taking into account truncation bias, most patents are not renewed a third time.

——————**Insert Table A.8**——————

One potential limitation of using the raw number of payments is that newly-granted patents may not have had enough time to pay the additional maintenance fee even if firms consider the patents to be valuable. Specifically, this problem exists for all patents granted since 2007. To correct for this truncation bias, we compute the expected number of payments for each patent. This means we compute the probability that patents that only had time to pay the 3.5-year fee end up paying for the 7.5- and 11.5-year fees, and the conditional probability that patents that only had time to pay 7.5-year fees will pay the 11.5-year fee in the end. The underlying assumption is that the conditional renewal probability of a patent after a previous patent maintenance fee payment is stable over time. The historical data corroborates this assumption. Between 1976–1990, the expected payment count of the paid-once patents is 1.909, and that of the paid-twice patents is 2.622. These expected values are comparable with the values predicted based on the 1991–2005 data: 1.908 for the paid-once patents and 2.623 for the paid-twice patents. We use the expected number of maintenance fee payments as our main dependent variable for patent economic value, but all our results are qualitatively unchanged if we use the raw number of maintenance fee payments instead.

[27] For example, as of October 2019, the current patent maintenance fees are $1,600, $3,600 and $7,400 after 3.5, 7.5 and 11.5 years respectively, for patents assigned to large firms. Before the last fee schedule change in 2013, the fees were $1,130, $2,850 and $4,730 respectively.

A.7 Econometric Model for the Case of Imperfect Labor Market

This is an extension of the econometric model by adding in imperfect labor markets. Assume that firms pay inventors efficiency wage with longer commuting distance to discourage them from shirking (Ross and Zenou, 2008). We define a simplified wage equation as follows:

$$w_{ij} = w_{ij}^m + ew_{ij}d_{ij} \tag{A.8}$$

where w_{ij}^m is the market-clearing wage and ew_{ij} is the per-unit commuting distance efficiency wage that firm j pays to inventor i. Because only the excess wage beyond market-clearing level can affect productivity by preventing inventors from shirking, inventor productivity then becomes:

$$l_{ij} = \theta_{ij} + (\beta_i + ew_{ij})d_{ij} \tag{A.9}$$

Following the same steps as in the main text, Equation 4 now becomes:

$$l_{ij} = \alpha_0 + (\beta_i + ew_{ij} + \alpha_1\gamma_i)d_{ij} + \delta_j + \alpha_1\mu_{ij} + \epsilon_{ij} \tag{A.10}$$

In this case of imperfect labor markets, the existence of non-zero job search cost s acts in conjunction with the existence of infra-marginal inventors due to heterogeneity in inventor-firm match quality ϵ_{ij} to keep some inventors from moving to a different job after firm relocation. Just as before, positive moving costs c_i keep some inventors from moving to a different residential location. Looking at this subsample of inventors who do not re-sort, we subtract their productivity before and after the move to get:

$$\Delta l_{ij} = (\beta_i + ew_{ij})\Delta d_{ij} \tag{A.11}$$

Given that ew_{ij} should always be the opposite sign of β_i, we are now estimating a lower bound for the weighted "pure" commuting effect on inventor productivity.

Appendix: Figures

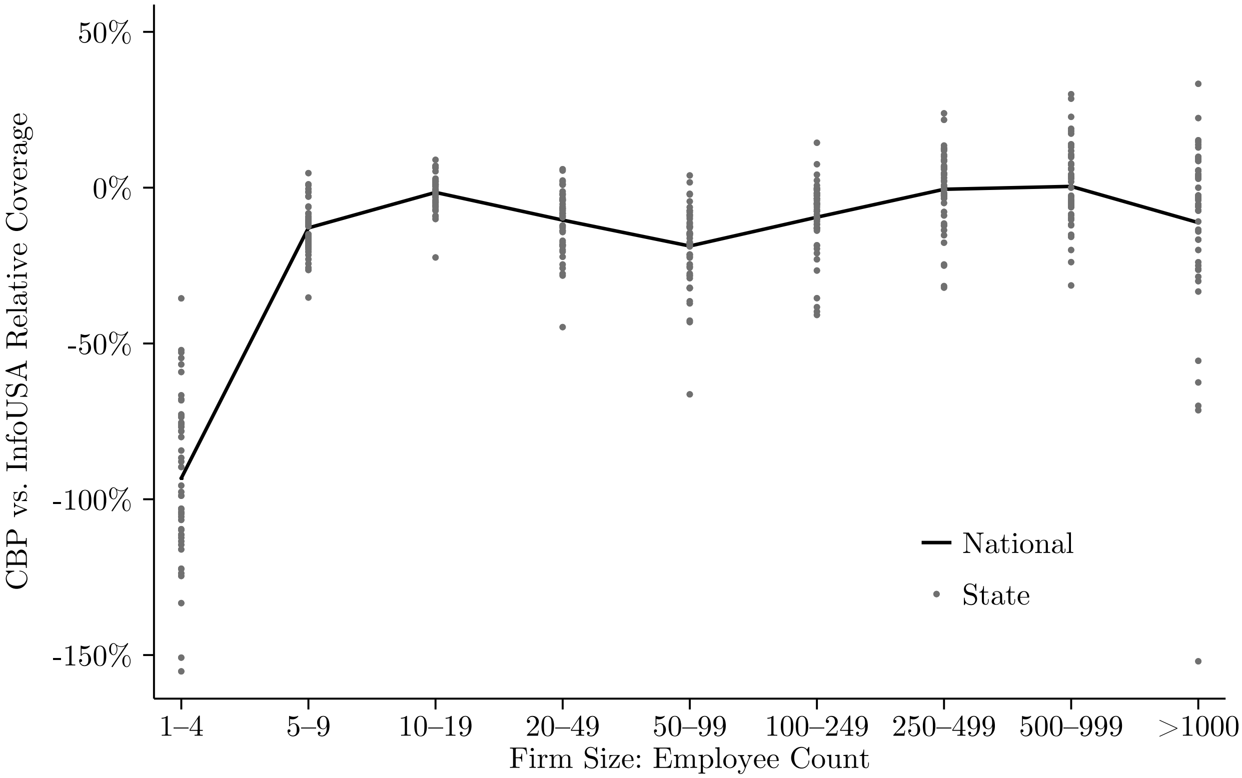

Figure A.1: **Relative Workplace Coverage Comparison Between InfoUSA and County Business Patterns**. This figure depicts the relative coverage of individual business establishments (workplaces) for InfoUSA as compared to County Business Patterns (CBP) for the year 2006. The vertical axis is calculated as (CBP Coverage − InfoUSA Coverage)/InfoUSA Coverage, and the horizontal axis depicts firm size as measured by the count of employees at the establishment. The solid black line represents the national-level relative coverage. The grey dots represent the state-level relative coverage. Overall, the coverage of the two datasets is fairly close, with more establishments covered by InfoUSA. However, small firms with one to four employees are distinctly under-represented in CBP.

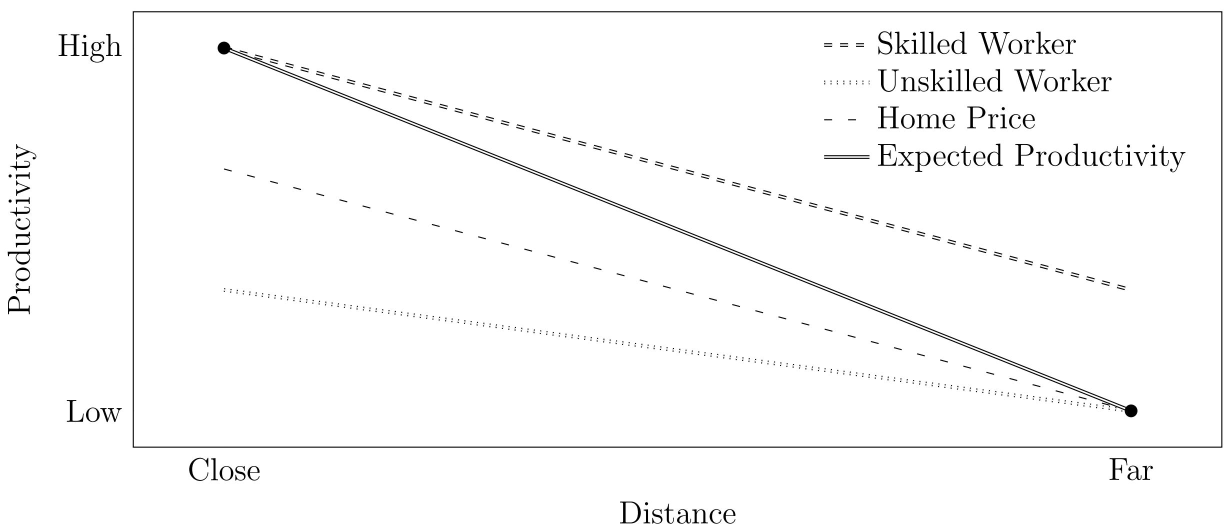

Figure A.2: **Skilled Workers Closer, Unskilled Workers Further: Amplification Bias**. When skilled worker productivity declines with distance more than for unskilled workers, skilled workers live closer to the city center, while unskilled workers live further away. In this separating equilibrium, the bias from sorting ***ampli es*** the naïve expected effect of workplace-home distance upwards relative to the "true" effect, i.e., it makes the negative effect even more negative, as shown by the ***steeper*** line for ***Expected Productivity*** relative to the slopes of the ***(Un)Skilled Worker*** lines. The model is originally based on two discrete points, ***Close*** and ***Far***, but we relax that simplifying assumption and plot (expected) continuous lines to provide a better visual guide

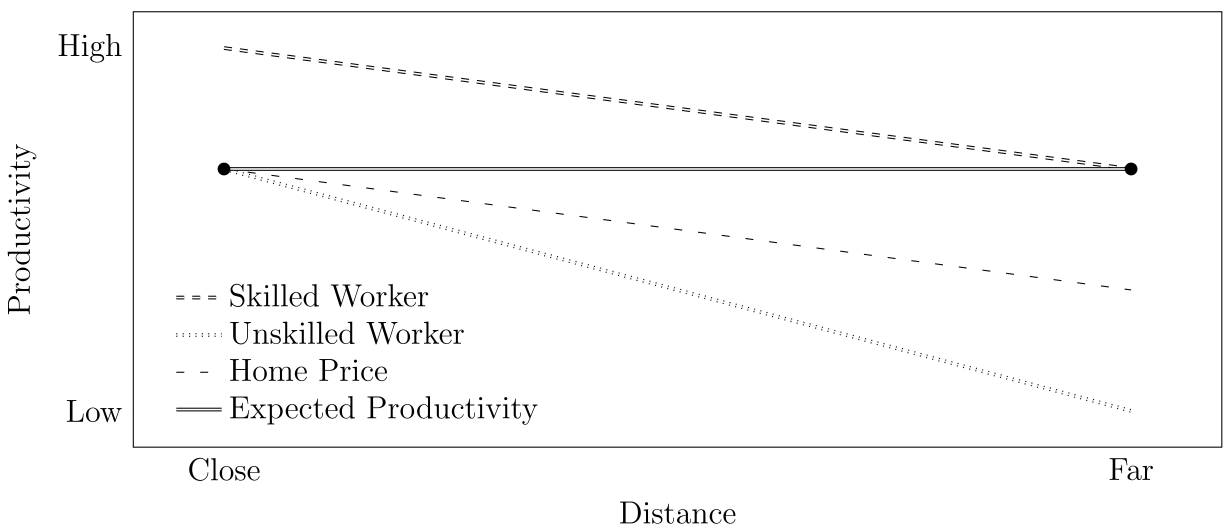

Figure A.3: **Skilled Workers Further, Unskilled Workers Closer: Attenuation Bias**. When skilled worker productivity declines with distance less than for unskilled workers, skilled workers live further from the city center, while unskilled workers live further away. In this separating equilibrium, the bias from sorting *attenuates* the naïve expected effect of workplace-home distance upwards relative to the “true” effect, i.e., it makes the negative effect less negative, as shown by the *atter* line for *Expected Productivity* relative to the slopes of the *(Un)Skilled Worker* lines. The model is originally based on two discrete points, *Close* and *Far*, but we relax that simplifying assumption and plot (expected) continuous lines to provide a better visual guide

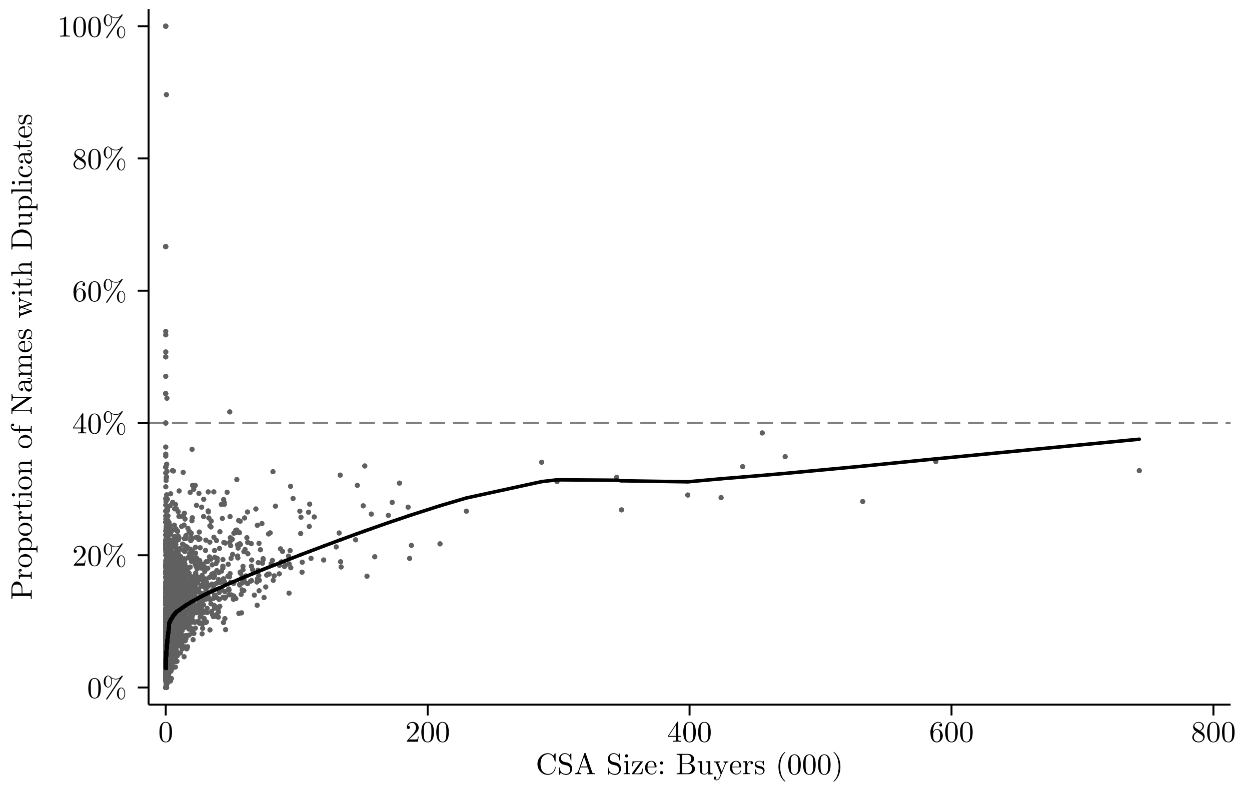

Figure A.4: **Proportion of Names per City that Appear More Than Once**. The vertical axis represents the proportion of unique names that appear more than once. The horizontal axis represents the size of a city in terms of the number of home buyers in the DataQuick sample. Each grey dot represents a particular city, and the black line is a local linear regression model fit onto the city-level points (Bandwidth= .8).

Appendix: Tables

Table A.1: **Combined Statistical Areas (CSA) available in DataQuick**. Further details on this sample are discussed in Ferreira and Gyourko (2015). Some CSAs are not covered in full.

	Combined Statistical Area		Combined Statistical Area
1	New York-Newark (NY-NJ-CT-PA)	31	Bakersfield (CA)
2	Los Angeles-Long Beach (CA)	32	Modesto-Merced (CA)
3	Chicago-Naperville (IL-IN-WI)	33	Springfield-Greenfield Town (MA)
4	Washington-Baltimore-Arlington (DC-MD-VA-WV-PA)	34	Spokane-Spokane Valley-Coeur d'Alene (WA)
5	San Jose-San Francisco-Oakland (CA)	35	Colorado Springs (CO)
6	Boston-Worcester-Providence (MA-RI-NH-CT)	36	Lakeland-Winter Haven (FL)
7	Miami-Fort Lauderdale-Port St. Lucie (FL)	37	Visalia-Porterville-Hanford (CA)
8	Detroit-Warren-Ann Arbor (MI)	38	Reno-Carson City-Fernley (NV)
9	Seattle-Tacoma (WA)	39	Palm Bay-Melbourne-Titusville (FL)
10	Phoenix-Mesa-Scottsdale (AZ)	40	Pensacola-Ferry Pass-Brent (FL)
11	Cleveland-Akron-Canton (OH)	41	Santa Maria-Santa Barbara (CA)
12	Denver-Aurora (CO)	42	Salinas (CA)
13	Tampa-St. Petersburg-Clearwater (FL)	43	Peoria-Canton (IL)
14	Orlando-Deltona-Daytona Beach (FL)	44	Tallahassee-Bainbridge (FL)
15	Portland-Vancouver-Salem (OR-WA)	45	Eugene (OR)
16	Sacramento-Roseville (CA)	46	Gainesville-Lake City (FL)
17	Columbus-Marion-Zanesville (OH)	47	Ocala (FL)
18	Las Vegas-Henderson (NV-AZ)	48	Fort Collins (CO)
19	Cincinnati-Wilmington-Marysville (OH-KY-IN)	49	San Luis Obispo-Paso Robles-Arroyo Grande (CA)
20	Jacksonville-St. Marys-Palatka (FL-GA)	50	Crestview-Fort Walton Beach-Destin (FL)
21	Hartford-West Hartford (CT)	51	Yakima (WA)
22	Oklahoma City-Shawnee (OK)	52	Redding-Red Bluff (CA)
23	Memphis-Forrest City (TN-MS-AR)	53	Chico (CA)
24	Tulsa-Muskogee-Bartlesville (OK)	54	Prescott (AZ)
25	Fresno-Madera (CA)	55	Bellingham (WA)
26	Cape Coral-Fort Myers-Naples (FL)	56	Yuma (AZ)
27	Honolulu (HI)	57	Panama City (FL)
28	Dayton-Springfield-Sidney (OH)	58	Grand Junction (CO)
29	Tucson-Nogales (AZ)	59	Flagstaff (AZ)
30	North Port-Sarasota-Bradenton (FL)	60	Pittsfield (MA)

Table A.2: **Effect of Commuting on Inventor Productivity - Name Weights**. OLS regression model with observations weighted by the inverse of the frequency of an inventor's name in DataQuick real estate data. The dependent variables are patent counts and patent quality measures per inventor-firm pair per year. The independent variable of workplace-home ***Distance*** is geodesic distance measured in 10 kilometers. *Top Inventor* are inventors in the top decile in terms of average patent count. Robust standard errors clustered at the inventor-firm pair-level are shown in parentheses. R^2 includes both within- and between-variation.

	Patents		*Scaled Citation*		*Fee Payments*	
Variable	(1)	(2)	(3)	(4)	(5)	(6)
Distance	-0.046**	-0.035*	-0.109**	-0.092*	-0.060	-0.034
	(0.020)	(0.019)	(0.050)	(0.050)	(0.045)	(0.044)
Distance×		-0.153*		-0.233		-0.346
Top Inventor		(0.084)		(0.176)		(0.171)
Inventor-Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm Location FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.413	0.413	0.385	0.385	0.416	0.417
Inventor-Firm Count	3,445	3,445	3,445	3,445	3,445	3,445
Observations	22,861	22,861	22,861	22,861	22,861	22,861

Table A.3: **Effect of Commuting on Inventor Productivity - Firm-Level Variables**. OLS model with subsample of observations matched to Compustat. The dependent variables are patent counts and adjusted patent quality measures per inventor-firm pair per year. The independent variable of workplace-home *Distance* is geodesic distance measured in 10 kilometers. *Top Inventor* are inventors in the top decile in terms of average patent count. Robust standard errors clustered at the inventor-firm pair-level are shown in parentheses. R^2 includes both within- and between- variation.

	Patents		*Scaled Citation*		*Payment Count*	
Variable	(1)	(2)	(3)	(4)	(5)	(6)
Distance	-0.060**	-0.038	-0.077	-0.036	-0.111	-0.058
	(0.026)	(0.025)	(0.058)	(0.058)	(0.068)	(0.066)
Distance×		-0.262**		-0.467**		-0.600**
Top Inventor		(0.114)		(0.186)		(0.280)
Log Turnover	-0.009	-0.009	-0.007	-0.008	0.017	0.016
	(0.027)	(0.027)	(0.088)	(0.088)	(0.054)	(0.054)
Log Market Value	0.031	0.031	-0.015	-0.016	0.050	0.049
	(0.032)	(0.032)	(0.115)	(0.115)	(0.073)	(0.074)
Log Assets	0.093**	0.094**	0.249*	0.251*	0.149	0.151
	(0.042)	(0.042)	(0.133)	(0.133)	(0.095)	(0.095)
Log Employee Count	-0.012	-0.013	-0.002	-0.002	-0.026	-0.026
	(0.010)	(0.010)	(0.025)	(0.025)	(0.023)	(0.023)
Log Sales Volume	-0.009**	-0.009**	-0.021**	-0.020**	-0.024***	-0.024***
	(0.004)	(0.004)	(0.010)	(0.010)	(0.008)	(0.008)
Inventor-Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm Location FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.421	0.422	0.445	0.445	0.408	0.409
Inventor-Firm Count	1,493	1,493	1,493	1,493	1,493	1,493
Observations	9,218	9,218	9,218	9,218	9,218	9,218

Table A.4: **Effect of Commuting on Inventor Productivity - No Bay Area**. OLS model with subsample of observations excluding inventors in the San Francisco Bay Area. The dependent variables are patent counts and adjusted patent quality measures per inventor-firm pair per year. The independent variable of workplace-home ***Distance*** is geodesic distance measured in 10 kilometers. *Top Inventor* are inventors in the top decile in terms of average patent count. Robust standard errors clustered at the inventor-firm pair-level are shown in parentheses. R^2 includes both within- and between- variation.

	Patents		*Scaled Citation*		*Payment Count*	
Variable	(1)	(2)	(3)	(4)	(5)	(6)
Distance	-0.053**	-0.037	-0.091*	-0.057	-0.076	-0.043
	(0.024)	(0.023)	(0.051)	(0.059)	(0.048)	(0.045)
Distance×		-0.130		-0.316*		-0.267
Top Inventor		(0.087)		(0.191)		(0.177)
Inventor-Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm Location FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.400	0.401	0.396	0.356	0.411	0.411
Inventor-Firm Count	2,277	2,277	2,277	2,277	2,277	2,277
Observations	15,099	15,099	15,099	15,099	15,099	15,099

Table A.5: **Effect of Workplace-Home Distance on Inventor Productivity - Large Establishments**. The dependent variable *Patents* is the count of patents granted to an inventor-firm pair per year, where the establishment the inventor works for has on average more than 100 employees during the sample period. The independent variable of workplace-home *Distance* is geodesic distance measured in 10 kilometers. *Top Inventor* are inventors in the top decile in terms of average patent count. OLS regression model with robust standard errors clustered at the inventor-firm pair-level shown in parentheses. R^2 includes both within- and between- variation.

Dependent Var.: *Patents*	(1)	(2)	(3)	(4)	(5)
Distance	-0.031***	-0.088***	-0.070***	-0.060**	-0.040
	(0.009)	(0.024)	(0.021)	(0.028)	(0.027)
Distance × *Top Inventor*					-0.173*
					(0.099)
Inventor-Firm FE	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Firm Location FE	No	No	No	Yes	Yes
R^2	0.002	0.342	0.389	0.416	0.416
Inventor-Firm Count	2,042	2,042	2,042	2,042	2,042
Observations	13,786	13,786	13,786	13,762	13,762

Table A.6: **Effect of Workplace-Home Distance on Inventor Productivity - Limited Distance Subsample**. Subsample includes inventor-firm pairs whose workplace-home geodesic distances are less than 50km for all years of employment. *Top Inventor* are inventors in the top decile in terms of average patent count. OLS regression model with robust standard errors clustered at the inventor-firm pair-level shown in parentheses. R^2 includes both within- and between- variation.

Dependent Var.: *Patents*	(1)	(2)	(3)	(4)	(5)
Distance	-0.014	-0.056***	-0.033**	-0.049**	-0.023
	(0.010)	(0.018)	(0.017)	(0.025)	(0.024)
Distance × *Top Inventor*					-0.273**
					(0.107)
Inventor-Firm FE	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	Yes	Yes
Firm Location FE	No	No	No	Yes	Yes
R^2	0.000	0.334	0.385	0.414	0.415
Inventor-Firm Count	3,118	3,118	3,118	3,118	3,118
Observations	20,788	20,788	20,788	20,740	20,740

Table A.7: **Patent Valuations in Prior Literature**. This table summarizes prior estimates of the mean value of a patent documented in the literature. These studies generate estimates using different methodologies and underlying data. The nominal U.S. dollar estimates of patent value given in each paper are adjusted to the 2010 U.S. dollars using the Consumer Price Index (CPI) provided by the U.S. Bureau of Labor Statistics.

Article	Valuation Methodology	Nominal Value (USD)	Nominal Year	Value in 2010 (USD)
Pakes (1985)	Stock Market Returns	810,000	1972	4,225,487
Austin (1993)	Stock Market Returns	1,504,000	1991	2,407,902
Kogan et al. (2017)	Stock Market Returns	3,200,000	1982	8,330,000
Serrano (2005)	Patent Maintenance Fee	86,782	2003	102,844
Bessen (2008)	Patent Maintenance Fee	78,168	1992	121,490
Fischer and Leidinger (2014)	Marketplace Transactions	104,781	1992	162,852

Table A.8: **Patent Maintenance Fee Payment Counts**. This table summarizes the raw number of patents by maintenance fee payment counts, for patents applied for between 1980 and 2005.

Payment Count	Number of Patents	Percentage
0	852,309	28%
1	535,687	18%
2	493,703	16%
3	1,174,096	38%
Total	3,055,795	100%