A Online Appendix

A.1 Data Composition and Construction

A.1.1 DataQuick: Sample Description

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. The DataQuick sample does not include information on home renters or on 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 About Here.** ———————

A.1.2 InfoUSA: 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 in Figure A.1 (e.g., Duranton and Turner, 2012). InfoUSA, covering both more establishments and more employees, has more comprehensive establishment coverage than CBP. 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).[32]

——————— **Insert Figure A.1 About Here.** ———————

A.1.3 Matching Procedure

Matching Homeowners to Inventors The first step in the sample construction process is to identify the homes of inventors based on the residential history given in the DataQuick

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

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 initial 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 transactions 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 we match names with home buyers from the previous transaction to determine residential spells. If there are no previous transactions 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 during the sample period. We then eliminate inventors for whom we cannot identify their primary residence, such as inventors 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 that has duplicates. We discuss this in Appendix A.2. 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 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 misspelled 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 that 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 are 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 two 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. The outliers are inventor–firm pairs who satisfy one of the following four criteria in any one year in our sample: distance between home and workplace is greater than 100 km, driving distance is greater than 125 km, change in distance is greater than 57 km, or change in driving distance is greater than 55 km. 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.1.4 Matched vs. Unmatched Inventors

To validate the relevance of our sample of inventors to inventors more generally, we compare the sample of inventors used in the main analysis—the *matched* inventors that remain after the matching processing across data sets—with the *unmatched* inventors intentionally omitted to facilitate the main empirical strategy. In particular, we compare the innovation performance of the matched vs. unmatched inventors to demonstrate that there are observable differences across these sets of inventors.

As an example of the type of situation we would like to rule out, consider a scenario where senior and more-productive inventors were more likely to purchase houses than emerging junior inventors. In this situation, the senior inventors would appear more in the housing data, meaning that our matched sample of inventors would exhibit more upward performance than inventors not in the sample or a universal inventor sample. Because the USPTO data provides comprehensive information about all the inventors with granted patents, we can build innovation performance measures for both matched and unmatched inventors and conduct t-tests to assess whether there are differences in their performance.

Table A.2 shows the t-test results with the three main measures of innovation performance: *Patent Count*, *Scaled Citation*, and *Payment Count*. For *Patent Count*, the magnitude of the estimated difference is small, albeit statistically significant, especially compared to the standard deviation of the underlying data. If we isolate the two samples to only inventors with more than one granted patent, the difference becomes smaller (0.237). This fact suggests that the unmatched inventors could be more likely to be associated with one-off patent records that themselves may have inconsistencies or coding errors, e.g., an incorrectly spelled name. Furthermore, the patent quality measures, *Scaled Citation* and *Payment Count*, show no difference between the two groups, and the unmatched inventors have slightly greater mean values.

—————— **Insert Table A.2 About Here.** ——————

Figure A.2 compares the histograms of each patent measure (logged) by the matched and unmatched inventors.[33] As we would expect given the previous t-test findings, the unmatched inventors tend to have more representation than the matched inventors in the part of the distribution where *Patent Count* is close to one, i.e., a logged value of zero in the histogram). Other than that region, the distributions are largely similar. For the other two quality measures, *Scaled Citation* and *Payment Count*, the distributions in the two groups visually appear nearly identical.

—————— **Insert Figure A.2 About Here.** ——————

In light of these statistical and visual comparisons, we conclude that the set of inventors used in the study are reasonably representative of the broader set of inventors who did not match with housing data, at least in terms of observable innovation performance.

A.2 Duplicate Inventor Names

One important source of measurement error in our regressions is the presence of duplicate names during our matching process. Because 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 is in fact a duplicate. The two first cases are observable in our data, and we mitigate their occurrence

[33] The value of 1 is added to *Payment Count* before taking a log because inventors can have zero payment counts.

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.[34]

A.2.1 Probability of Having Duplicates

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:

$$Q_{ij} = N_j P_{ij}^l P_{ij}^f$$

N_j represents the total population in the city. 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.

P_{ij}^l is the probability that a person has the same last name. Since one's last name often depends on the ethnic group she 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.

P_{ij}^f is the probability that a person has the same first name and middle initial. 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).

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 is identical across all cities. In other words, P_{ij}^l and P_{ij}^f no

[34]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.

longer depend on j, so we can use cross-sectional data without worrying about compositional differences between cities biasing our estimates.

Based on these assumptions, 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.[35] In Figure A.3, the mean probability of finding another buyer with an identical name is plotted against the number of buyers in a city. Even in the largest cities, fewer than 40% of all home buyers over a 20-year period (1992 to 2012) have another buyer with identical name, who could very well be the same person buying another property. Therefore, our results should not suffer too much attenuation bias from duplicate names.

——— Insert Figure A.3 About Here. ———

A.2.2 Weighting Duplicates

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

$$w = \frac{1}{1 + Number\ of\ Duplicates\ in\ CSA}$$

Thus, if an inventor's name only appears once, we assign 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.3. All the estimated coefficients for patent quantity and quality are still negative and significant, consistent with results in the main text.

——— Insert Table A.3 About Here. ———

A.3 Illustrative Model

We modify our inventor sorting model by using skill differences of inventors as the driver of sorting, without additional economic considerations such as income and amenities. We assume that there is a difference in sensitivity to commuting distance with respect to productivity loss, between high- and low-skilled inventors. We explain how this difference in sensitivity between inventors causes inventor sorting.

[35] What we estimate is an upper bound because the same person in the DataQuick database could have made multiple housing purchases.

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: high-skilled workers and low-skilled 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 high-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 low-skilled workers. This general model 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 for the rest of this section that there is a negative total effect of distance on productivity, although derivations with positive total effects hold analogously. We also abstract away from the underlying mechanisms that produce this relationship, and we focus on how worker sorting can interact with this negative correlation between distance and productivity and cause bias in our estimation.[36]

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. Risk-neutral workers maximize:

$$U_i = wl_i - p(h_i)$$

where $p(h_i)$ is the price of housing. Each worker demands the same quantity of housing, so only the latter's price enters into his utility function. All housing stock belongs to absentee owners who maximize revenue, subject to their tenants' participation constraints. To make the model more tractable, we make some further assumptions: the housing stock evenly divides between locations A and B; the worker types are equal in number; and the total

[36] 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.

number of housing units equals the total number of workers so the housing market can clear. There are no moving costs, so workers can costlessly change locations.

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 low-skilled workers declines faster with distance than for high-skilled workers; and $\alpha_h < \alpha_l$, where productivity for low-skilled workers declines slower with distance than for high-skilled workers.

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 high-skilled and low-skilled 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 low-skilled workers living at location A, closer to the city center than high-skilled workers at location B. If $\alpha_h < \alpha_l$, then the separating equilibrium has low-skilled workers living at location B, farther away from high-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.4, when $\alpha_h < \alpha_l$, the observed effect is more negative than either of the true slopes for high-skilled or low-skilled workers, while for the case where $\alpha_h > \alpha_l$ in Figure A.5, the observed effect is less negative than either of the true slopes for high-skilled or low-skilled workers, and could even be positive.

Insert Figure A.4 About Here.

Insert Figure A.5 About Here.

Thus, identifying the causal effect of distance on productivity requires 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.3.1 Case I: Faster Productivity Decline for High-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 low-skilled workers live, housing is always priced at what the low-skilled workers can bear. Therefore, in the separating equilibrium with high-skilled workers at location A and low-skilled workers at location B, the housing price at location B is $p^B = wl_l^B$ and low-skilled workers get zero utility which equals the modeled outside option. To determine the housing price at location A, we use the incentive-compatibility constraint for high-skilled workers; for there to be an equilibrium, high-skilled workers must not want to move to location B:

$$wl_h^B - p^B \leq wl_h^A - p^A$$
$$w(1 + \alpha_h)\theta - p^B \leq w\theta - p^A$$
$$w\alpha_h\theta - p^B \leq p^A.$$

Because of profit maximization of the absentee homeowners, the last inequality becomes $p^A = p^B - w\alpha_h\theta$. This price differential satisfies the incentive-compatibility constraint for low-skilled 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 high-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 high-skilled workers living farther away at location B, and low-skilled 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 high-skilled workers because, for there to be an equilibrium, high-skilled workers must not want to move to location A:

$$wl_h^B - p^B \geq wl_h^A - p^A$$
$$w(1 + \alpha_h)\theta - p^B \geq w\theta - p^A$$
$$w\alpha_h\theta - p^B \geq p^A.$$

Because of profit maximization of the absentee homeowners, the last inequality becomes

$p^B = \ p^A + w\alpha_h\theta$. Then, however, the low-skilled 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 high-skilled and low-skilled worker's IC constraints cannot be satisfied at the same time. More specifically, housing price would be equal to the low-skilled worker's willingness-to-pay at both A and B, but this gives high-skilled workers an incentive to move to location A, where they can earn higher utility than in location B.

A.3.2 Case II: Slower Productivity Decline for High-Skilled Workers

In this case, $\alpha_h > \alpha_l$. Therefore, in the separating equilibrium with high-skilled workers at location A and low-skilled workers at location B, the housing price at location B is $p^B = wl_l^B$ and low-skilled 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 high-skilled workers because, for there to be an equilibrium, high-skilled workers must not want to move to location B:

$$wl_h^B - p^B \leq wl_h^A - p^A$$
$$w(1 + \alpha_h)\theta - p^B \leq w\theta - p^A$$
$$w\alpha_h\theta - p^B \leq p^A.$$

Because of profit maximization of the absentee homeowners, the last inequality becomes $p^A = \ p^B - w\alpha_h\theta$. Then, however, the low-skilled 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 high-skilled workers living farther away at location B, and low-skilled 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 high-skilled workers because, for there to be an equilibrium, high-skilled workers must not want to move to location A:

$$wl_h^B - p^B \geq wl_h^A - p^A$$
$$w(1+\alpha_h)\theta - p^B \geq w\theta - p^A$$
$$w\alpha_h\theta - p^B \geq p^A.$$

Because of profit maximization of the absentee homeowners, the last inequality becomes $p^B = p^A + w\alpha_h\theta$. This price differential satisfies the incentive-compatibility constraint for low-skilled workers since they would have negative utility living at B, given that $\alpha_h > \alpha_l$. Therefore, the separating equilibrium with high-skilled workers living farther from the firm is possible.

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

A.3.3 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 low-skilled worker's willingness-to-pay at both locations A and B. In this case, the high-skilled worker would be indifferent about living at either location because she will get the same utility of $wl_h^A - p^A = wl_h^B - p^B = w(1-\beta)\theta$. This unique equilibrium occurs because in both potential separating equilibria high-skilled workers are indifferent about living in either location.

A.4 Patent Maintenance Fee Data

For estimates of the economic value of individual patents, we use patent maintenance fee data from 1980 to 2019, available from the USPTO. Patent maintenance fee data has previously been used to estimate the economic value of individual patents, notably by Pakes (1986) and Bessen (2008).

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.[37] 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.4 shows the distribution of *Payment Count*. Even after taking into account truncation bias, most patents are not renewed a third time.

——————— Insert Table A.4 About Here. ———————

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 fees 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 have only had time to pay the 3.5-year fee end up paying the 7.5- and 11.5-year fees, and the conditional probability that patents that have 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

[37] 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.

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.

A.5 Firm Relocation Decisions and Inventor Performance

We conduct three categories of analyses that assess whether firms relocate based on the performance of their inventors. That is, firms may selectively relocate their offices to get closer to better-performing inventors, which would risk introducing endogeneity into our main estimates. First, we look at whether post-firm-relocation distance *backward* predicts pre-firm-relocation performance. Second, we conduct an analysis restricted to a sample of only large firm establishments. Third, we divide firms by the change in their average commuting distance and compare their pre-firm-relocation innovation performance .

A.5.1 Post-Firm-Relocation Distance vs. Pre-Firm-Relocation Performance

We first investigate whether post-firm-relocation distance *backward* predicts pre-firm-relocation performance. If firms endogenously make their relocation decisions based on their inventor performance, post-firm-relocation distance should be negatively correlated to pre-firm-relocation performance, i.e., shorter commutes would associate with better performance. For each firm, we calculate cumulative and average performance measures before and after firm relocation. These measures follow from the same data that underlie *Patent Count*, *Scaled Citation*, and *Payment Count* in the main analysis.

Table A.5 presents the regression estimates at the firm–inventor level. The post-firm-relocation distance does not have a statistically significant relationship with any of the pre-firm-relocation performance. This result remains unchanged whether we use cumulative- or average-based measures.

——————— Insert Table A.5 About Here. ———————

A.5.2 Large Establishments

We estimate the distance effect in a sample restricted to large firm establishments. If we make some basic assumptions, this test provides another way of addressing the concern of endogenous firm relocation. First, a larger firm is less likely to make its relocation decision based on the location of its inventor because they have a large number of inventors; unless inventors live in residences clustered by their performance, firms cannot selectively reduce the average commuting distance of better-performing inventors. Second, larger firms have significant moving costs, making it less plausible that the expected benefit of relocating towards better-performing inventors would justify the substantial relocation cost, unless the relocation were motivated the more important and common reasons that firms generally relocate, e.g., real estate pricing, need for specific office configuration, etc. Using a subsample of large establishments whose average number of employees are greater than 100 during our sample period, Table A.6 shows that the negative effect of *Distance* on inventor productivity is still statistically significant and even larger in magnitude than in our baseline specification.

—————— Insert Table A.6 About Here. ——————

A.5.3 Validating Underlying Assumptions

We further identify a critical, but likely untrue, assumption underlying whether we need to worry about endogenous firm relocation decisions based on inventor performance. Besides the potentially tenuous starting assumption that firms relocate based on inventor performance, the two key assumptions that would make endogeneity of this form possible are that (1) inventors can be divided into better-performing and worse-performing groups, and (2) their residential locations are spatially segregated. If these two assumptions do not hold, then firms cannot endogenously decide where to relocate based on inventor performance, assuming that firms would want to. To illustrate the implication of the two key assumptions, suppose there are two regions, A and B, and better- and worse-performing inventors are equally mixed in A and B. In this scenario, firms cannot get selectively closer to better-performing inventors by relocating. Hence, these two assumptions are necessary conditions for whether firms could even engage in an endogenous relocation decision.

The setting provides an opportunity for a new analysis that compares firms by their average commuting distance change after relocation. If we assume that firm relocation decisions are endogenous, firms would choose to relocate to a region consisting of more better-performing inventors. If a firm's average commuting distances decrease after its relocation, it means that this firm has more better-performing inventors than worse-performing inventors because we already assume firms relocate towards better-performing inventors. Likewise, a firm's average commuting distances will increase if the firm consists of more worse-performing inventors.

Thus, we investigate whether pre-firm-relocation innovation performance varies across firms based on the changes in their average commuting distances. We divide firms into three groups: firms whose average commutes decrease (i.e., *Decreased*), remain unchanged (i.e., the absolute change is smaller than 1 km, *Unchanged*), and increase (i.e., *Increased*). Table A.7 summarizes three main performance measures when dividing firms into three groups and reports t-tests for differences across firms. There is no observable difference in pre-firm-relocation performance between the three groups: the mean values are about the same in magnitude, and the differences are not statistically significant. Thus, at least on an across-firm level, we do not find evidence of endogenous firm relocation decisions to move closer to (or farther from) better-performing inventors.

—————— Insert Table A.7 About Here. ——————

A.6 Sample Design and Inventor Home Moving

We conduct two additional analyses intended to assess the consequences of the main sample design for the generalizability of our estimates given the intentional exclusion of inventors who move homes. First, we conduct a regression analysis and statistical tests that show that there is no significant difference in *pre-firm-relocation* performance trends between inventors who move their home and inventors who do not move their home. Second, we build an augmented sample that includes inventors who eventually move homes, but after their firm relocates. Using this sample that accounts for these inventors, and in particular their *post-firm-relocation* performance, the analysis generates estimates comparable to our main analysis.

A.6.1 Inventor Performance Trends Prior to Firm Relocation: Regression Analysis

Figure 6 in Section 4.3 provides a descriptive analysis to assess whether inventors who did not move their home locations (i.e., non-moved inventors) tend to perform worse than other inventors who moved their home locations (i.e., moved inventors) before firm relocations. To support the descriptive observations with statistical rigor, we run a complementary regression analysis and report a plot of estimated coefficient plots. For brevity, we only report the results with *Patent Count* as the dependent variable, and the results are similar for the other inventor productivity measures, i.e., *Scaled Citation* and *Payment Count*.

First, we separately estimate for the moved and non-moved inventors their performance trends prior to firm relocation (pre-firm-relocation). The purpose of this to see whether either group experiences any general upward or downward trends in performance before firm relocation that distinguishes one from the other. We estimate the following model separately on the two samples:

$$Inventor\ Productivity_{ijt} = \beta_k \sum_{k=-1}^{-5} Years\ from\ Relocation_{jt-l_j=k} + \alpha_j + \gamma_t + \delta_{jt} + \epsilon_{ijt} \quad \text{(A.1)}$$

where i is an inventor, j is a firm, t is a pre-firm-relocation year, k is the years from firm relocation, l_j is the relocation year of j, α_j is a firm fixed effect, γ_t is a time fixed effect, δ_{jt} is a firm location fixed effect, and ϵ_{ijt} is the error term. As compared to the main model in Equation 5, we include firm fixed effects, rather than inventor–firm fixed effects, because we intend to include the underlying differences across inventors in our estimates rather than excluding them by controlling for the difference. We cluster robust standard errors at the inventor–firm level.

The first and second graphs in Figure A.6 present the coefficient plots for the moved and non-moved inventors, respectively. Although there appears to be a small bump at two years before firm relocation ($k = -2$) in the first graph, the 95% confidence interval of the estimate includes zero and is not statistically significant. In the second graph, the estimated coefficients are close to zero for all the time points. Thus, we do not find any evidence that the moved and non-moved inventors have statistically significant pre-firm-relocation performance trends.

——————— Insert Figure A.6 About Here. ———————

Although both moved and non-moved inventors do not have any pre-trends by themselves, we still need to test whether the time-series *difference* in their performance exhibits any pre-trend. To test for a potential pre-trend in the difference betwen the moved and non-moved inventors, we run a generalized difference-in-differences-style estimation: we add an indicator variable for moved inventors, $Moved\ Inventors_i$, and its interaction term with $\sum_{k=-1}^{-5} Years\ from\ Relocation_{jt-l_j=k}$ to Equation A.1. The third graph in Figure A.6 is the coefficient plot for the interaction term, $\sum_{k=-1}^{-5} Years\ from\ Relocation_{jt-l_j=k} \times Moved\ Inventors_i$. We do not find any statistically significant evidence of a pre-firm-relocation trend in the performance difference between the moved and non-moved inventors. The estimated coefficients are close to zero, and the 95% confidence intervals include zero. In addition, we conduct an F-test to investigate whether the five coefficients of the interaction terms are jointly zero. We fail to reject the null hypothesis that the coefficients are jointly zero, e.g., for *Patent Count* as the dependent variable, $F = 0.64$ and p-value $= 0.6664$. The F-test does not find any statistically significant difference when using any of the three main innovation performance dependent variables. All the regression results support the descriptive analysis in Section 4.3 and reaffirm that there is no observable pre-firm-relocation performance difference between the moved and non-moved inventors.

A.6.2 Augmented Sample with Inventors Moving Homes After Firm Relocation

The previous analysis finds that the moved and non-moved inventors do not have an observable difference in performance *before* firm relocation (pre-firm-relocation). Thus, the remaining concern about the design of the inventor sample is the possibility that the moved inventors are affected by distance changes but only in terms of their performance *after* firm relocations (post-firm-relocation), where their workplace–home distance was unobservable in the dataset after they move homes.

Because a moved inventor can stay in her initial home locations for several years after her firm relocates, there is a "bridge" period of time where the firm has relocated but the moved inventor still resides at her original home. On average, the moved inventors stay in their initial home locations for about three years after their firms relocate. The existence of this bridge period allows us to investigate how the commuting distance shock affects the performance of the moved inventors who change their home locations after firm relocation. We augment the main sample with these moved inventors: for these moved inventors, we use their actual commuting distances during the bridge period and before the firm relocation.

Table A.8 shows the regression results for the non-moved inventors (i.e., inventors in our main sample) and the augmented sample (i.e., both moved and non-moved inventors). We refer to these samples as *Non-Moved* and *Mixed*, respectively. Using *Patent Count* as the dependent variable in Columns (1) and (2), we find that the magnitude of the coefficients for the mixed sample is a bit greater. This finding supports our intuition that our main findings represent at least a lower bound on the true point estimate based on the assumption that moved inventors might be more sensitive to distance changes than non-moved inventors, given that they might have chosen to move because of the firm relocation. When using *Scaled Citation* as the dependent variable in Columns (3) and (4), there is no longer an observable difference in coefficients estimated on the two samples, although the coefficient for the mixed sample has smaller standard errors. We do not find any difference in coefficient sizes and statistical significance when dependent variable is *Payment Count* in Columns (5) and (6). These findings support our intuition that our main estimates at least represent a lower bound on the "true" effect. Moreover, these findings suggest there is no observable difference between moved and non-moved inventors in their sensitivity to distance post-firm-relocation, given that the gap between the estimates on the main sample and the augmented sample appear to be small.

——————— **Insert Table A.8 About Here.** ———————

A.7 Additional Robustness Checks

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

A.7.1 Subsample: Without San Francisco Bay Area

Almost a third of all our observations come from the San Francisco Bay Area, as shown in Table 1 in Section 2.3.1. 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.9. 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.9 About Here.** ———————

A.7.2 Subsample: Distance Shock

Estimating the equation separately on shocks to closer versus farther inventors in Table A.10 shows an interesting heterogeneous result where the negative main effect is driven by inventors who get shocked farther. 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 to their workplace (Column 1), the coefficient for the inventors 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 A.10 About Here.** ———————

A.7.3 Subsample: Limited Distance

Another potential worry for our estimation is that inventor–firm 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.11 show that the negative effect of *Distance* on inventor productivity is almost unchanged and still highly significant, compared against baseline regressions.

——————— **Insert Table A.11 About Here.** ———————

A.7.4 Single-Authored Patents

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 effect on actual contribution. To account for this, we introduce the alternate measure, *Single-Authored Patent Count*, which only counts single-authored patents. This count is not subject to the bias above because 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 (4) in Table A.12 show that a 10 km increase in *Distance* causes a significant decrease of 0.012 *Single-Authored Patent Count* 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.

——————— Insert Table A.12 About Here. ———————

A.7.5 Categorical Distance Variable: Relocation Direction

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* became farther, became closer, or remained largely unchanged (change $<$ 1km) after the firm relocation shock. We construct the categorical variable Δ *Distance Direction* that equals 0 for 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 A.13. These estimates look similar to those in Table 4, 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 highest-performing inventors.

——————— Insert Table A.13 About Here. ———————

A.7.6 Alternative Distance Measures

We check the robustness of our results using alternative explanatory variables. We create two other distance measures based on the assumption that the inventor might be driving or taking public transit to work. *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 measures are 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.

Table A.14 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 A.14 About Here. ———————

A.7.7 Non-Linear Models

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.

Quadratic Term As shown in Table A.15, we find that when the square term, $Distance^2$, is added, the main effect becomes insignificant in both patent quantity and quality regressions. The coefficient on $Distance^2$ is close to zero and insignificant, suggesting that its causal effect on inventor productivity is likely linear at least 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.

——————— Insert Table A.15 About Here. ———————

Poisson and Negative Bionomial Models Furthermore, our results are robust to Poisson and negative binomial estimation. Given the discrete dependent variables, it is arguably better to use models that are better 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). In Table A.16, the marginal effect in Column

(2) shows that a 10 km increase in *Distance* causes a 4.9% decrease in patent count in the Poisson regression, and the marginal effect in Column (4) shows a 5.4% decrease in the negative binomial regression. Both numbers are very similar to our main linear specification.

—————— **Insert Table A.16 About Here.** ——————

Online Appendix: References

Bessen, James. 2008. "The value of US patents by owner and patent characteristics." *Research Policy*, 37(5): 932–945.

Cameron, A Colin, and Pravin K Trivedi. 2013. *Regression Analysis of Count Data.* . 2 ed., Cambridge, UK:Cambridge University Press.

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

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*.

Google LLC. 2018. "Developer guide."

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.

Pakes, Ariel. 1986. "Patents as options: Some estimates of the value of holding European patent stocks." *Econometrica*, 54(4): 755–784.

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).

Online Appendix: Figures

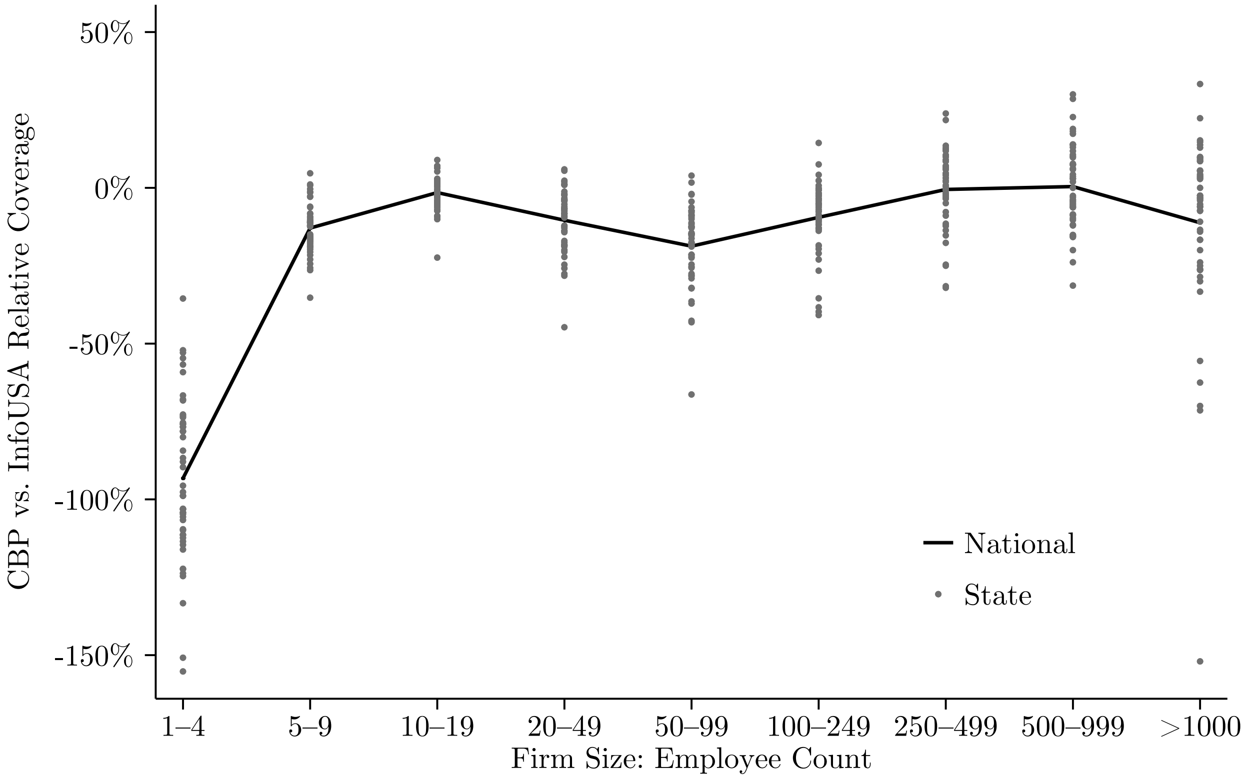

Figure A.1: **Relative Workplace Coverage Comparison Between InfoUSA and County Business Patterns**. Line scatter plot at the state level. This figure depicts the relative coverage of individual business establishments (workplaces) for InfoUSA as compared to County Business Patterns (CBP) for the year 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 black line represents the national-level relative coverage. The gray dots represent the state-level relative coverage.

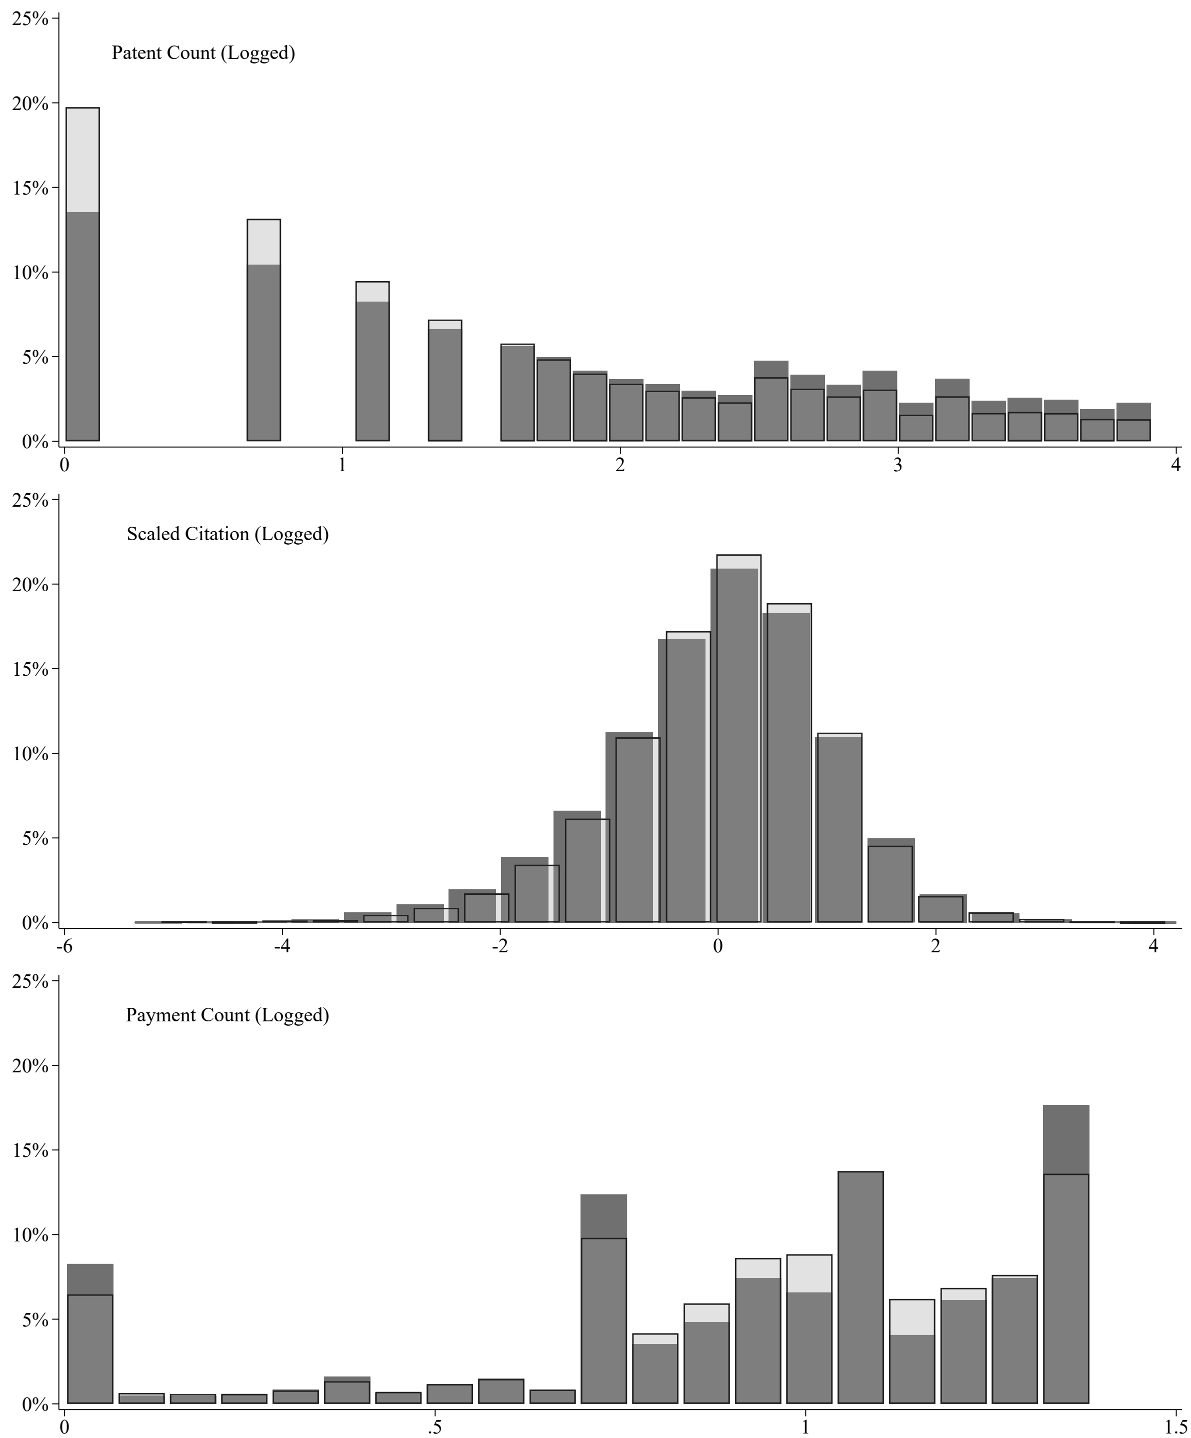

Figure A.2: **Distribution of Innovation Performance for Matched and Unmatched Inventors**. Percent distributions at the inventor level for the three main innovation measures logged. The dark bars represent the inventors in the USPTO data matched to the DataQuick and InfoUSA data, and the gray bars indicate the other unmatched inventors.

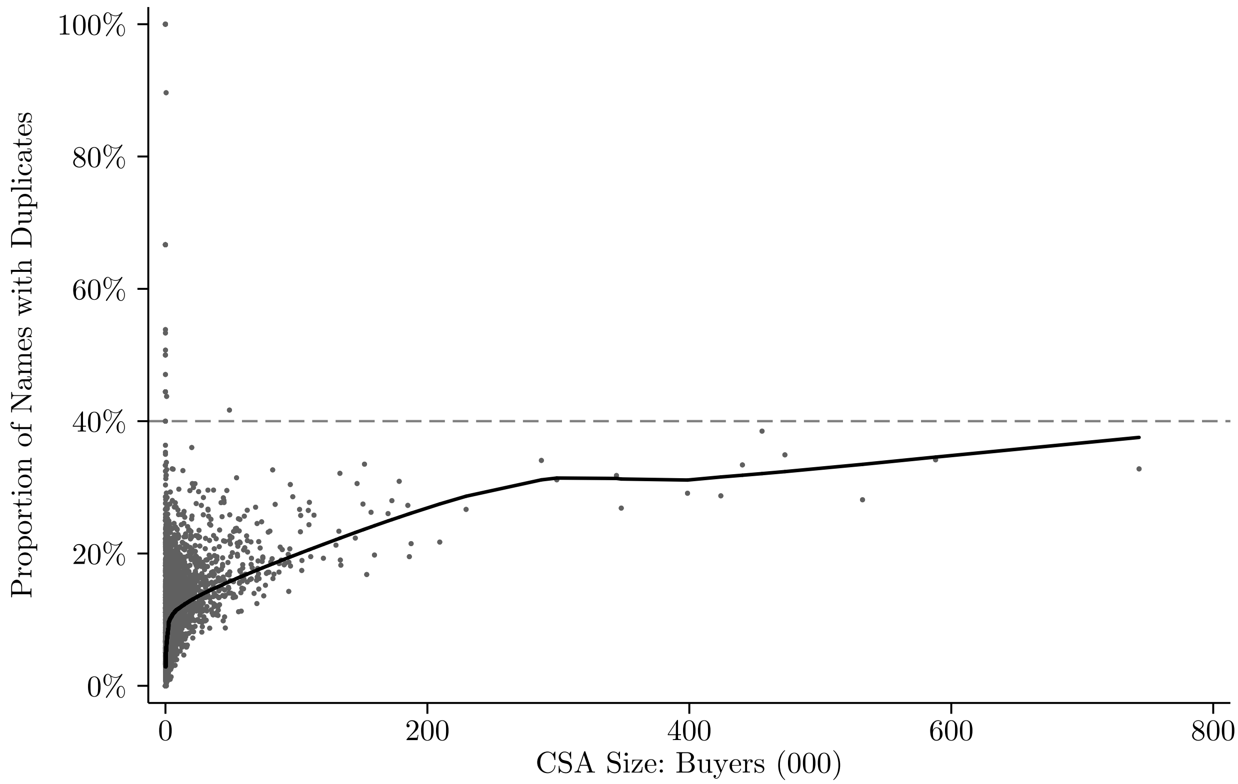

Figure A.3: **Proportion of Names per City that Appear More Than Once**. Line scatter plot at the city level. The vertical axis represents the proportion of unique names that appear more than once. The horizontal axis represents the size of a city in terms of the number of home buyers in the DataQuick sample. Each gray dot represents a particular city, and the black line is a local linear regression model fit onto the city-level points (bandwidth = 0.8).

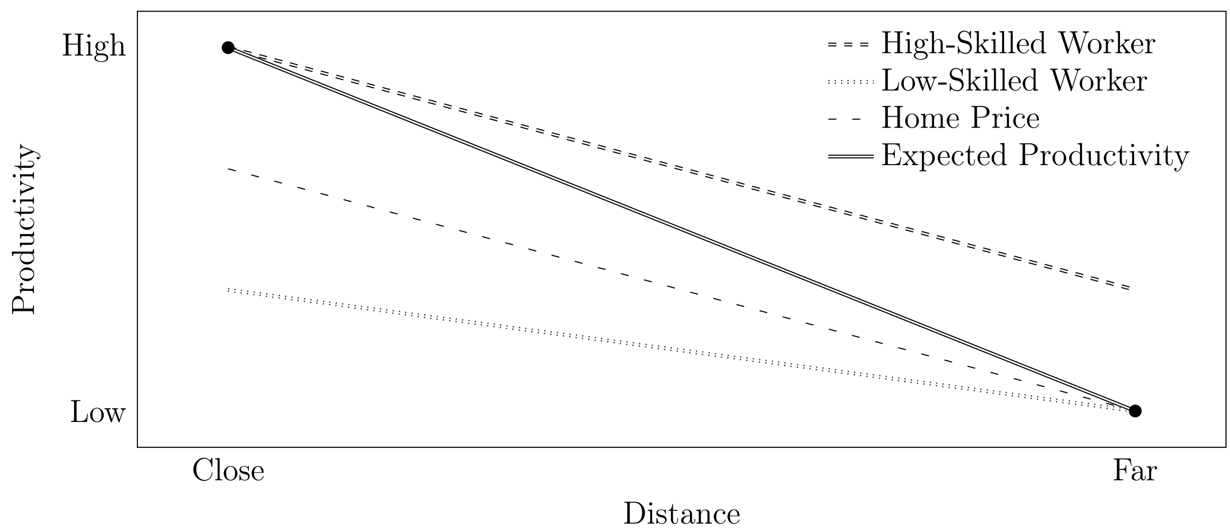

Figure A.4: **High-Skilled Workers Closer, Low-skilled Workers Farther—Amplification Bias**. Hypothetical relationships between home-city center distance and productivity. When high-skilled worker productivity declines with distance more than for low-skilled workers, high-skilled workers live closer to the city center, while low-skilled workers live farther away. In this separating equilibrium, the bias from sorting *amplifies* the naive 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 *High-Skilled Worker* and *Low-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.

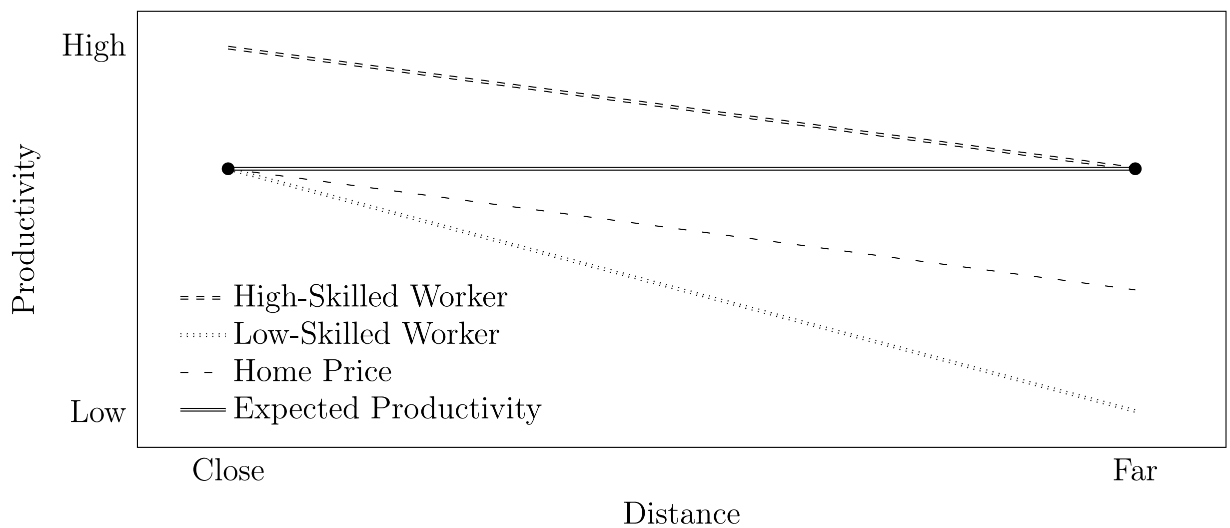

Figure A.5: **High-Skilled Workers Farther, Low-Skilled Workers Closer—Attenuation Bias**. Hypothetical relationships between home-city center distance and productivity. When high-skilled worker productivity declines with distance less than for low-skilled workers, high-skilled workers live farther from the city center, while low-skilled workers live farther away. In this separating equilibrium, the bias from sorting *attenuates* the naive 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 *flatter* line for *Expected Productivity* relative to the slopes of the *High-Skilled Worker* and *Low-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.

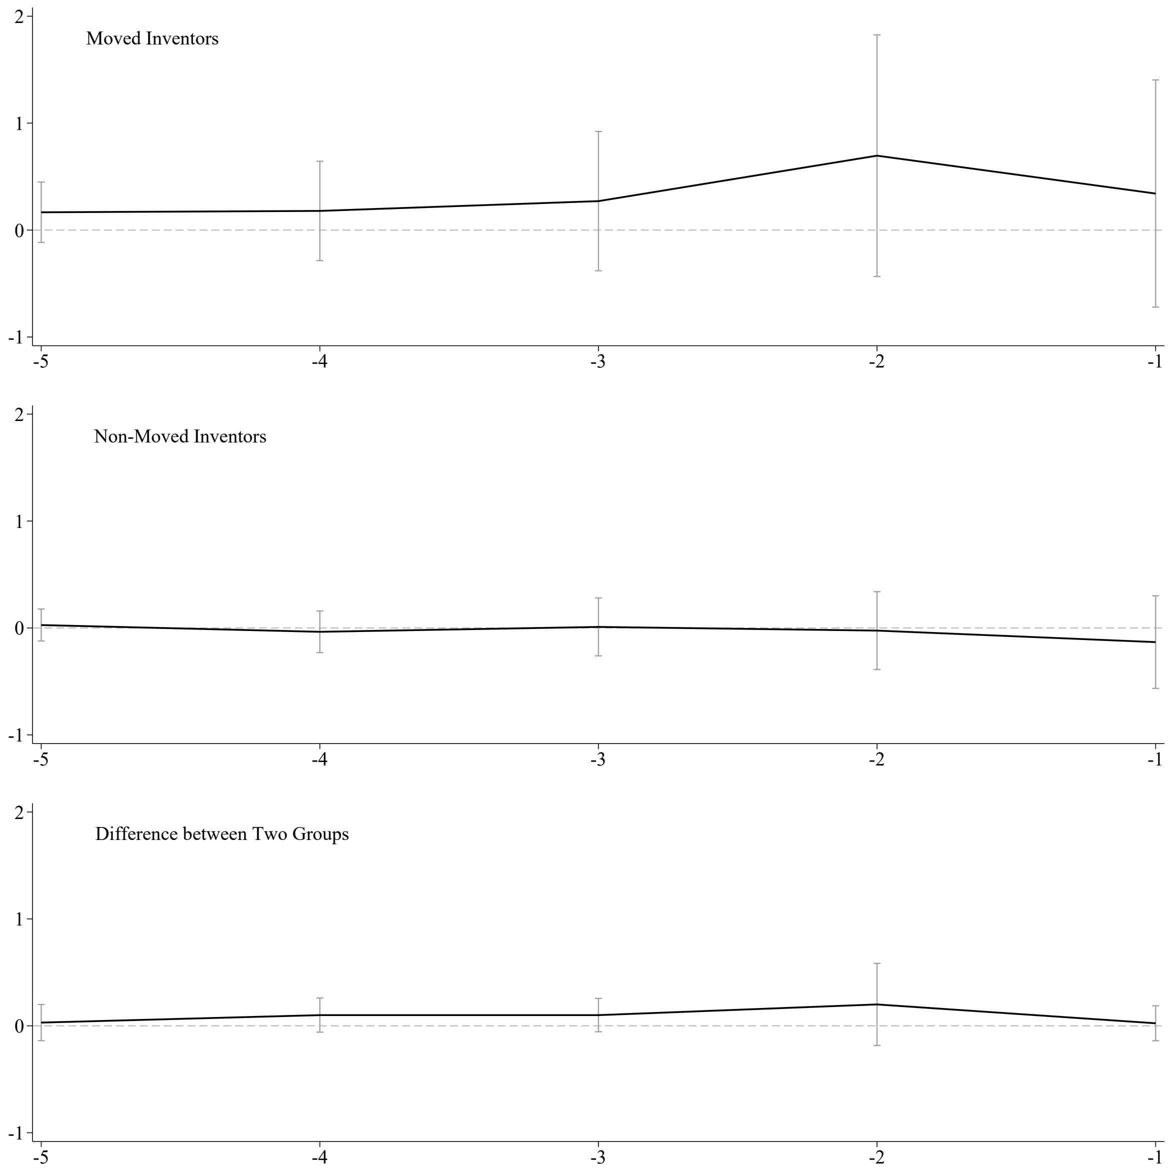

Figure A.6: **Trends in Innovation Performance For Moved and Non-Moved Inventors Prior to Firm Relocation—Regression Estimates**. Plot of estimated coefficients at the inventor–firm–year level with 95% confidence intervals. The horizontal axes denote years before firm relocations up to five years (i.e., from -1 to -5). The coefficients in the top graph denote the difference in the number of granted patents produced by inventor–firm per year compared to time points other than the five pre-firm-relocation years within the inventors who moved their home locations after firm relocations (i.e., *Moved Inventors*). The coefficients in the middle graph indicate the same difference among the inventors who did not move their home locations after firm relocations (i.e., *Non-Moved Inventors*). The coefficients in the bottom graph show the difference in the patent count between these two inventor groups relative to time points other than the five pre-firm-relocation years.

Online 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: **Performance of Matched and Unmatched Inventors**. Summary statistics and t-test results are provided at the inventor level. Among inventors identified in the USPTO data, *Matched* group indicates those who are matched to both the DataQuick and InfoUSA data, and *Unmatched* group represents the others. *Difference* is calculated as the *Matched* mean minus the *Unmatched* mean. Standard deviations are shown in parentheses, and the numbers of observations are shown in brackets.

Variable	Matched	Unmatched	Difference	p-value
Patent Count	2.429	2.141	0.287	0.000
	(2.074)	(1.851)		
	[264,287]	[268,861]		
Scaled Citation	1.460	1.462	-0.002	0.7085
	(2.148)	(2.129)		
	[264,287]	[268,861]		
Payment Count	1.728	1.730	-0.002	0.3471
	(0.943)	(0.986)		
	[264,287]	[298,694]		

Table A.3: **Effect of Commuting on Inventor Productivity — Name Weights**. Fixed-effects OLS regressions at the inventor–firm–year level weighted by the inverse of the frequency of an inventor's name in the DataQuick data. Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
D.V.:	*Patent Count*		*Scaled Citation*		*Payment Count*	
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 × *Top Inventor*		-0.153*		-0.233		-0.346
		(0.084)		(0.176)		(0.171)
Mean of D.V. (Pre-Relocation):						
All Inventors	0.858		1.365		1.816	
Ordinary Inventors		0.781		1.264		1.533
Top Inventors		1.670		2.421		3.276
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.4: **Patent Maintenance Fee Payment Counts**. Patent shares by the number of maintenance fee payments made by a firm to renew a patent (i.e., *Payment Count*). Patents applied between 1980 and 2005 are included.

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%

Table A.5: **Post-Firm-Relocation Distance and Pre-Firm-Relocation Performance**. Fixed-effects OLS regressions at the inventor–firm level. The dependent variables are the three main innovation measures aggregated at the inventor–firm level before firm relocations. The aggregation is conducted by taking cumulative sum (i.e., Cumulative) and taking average (i.e., Average). The independent variable, *Post-Relocation Distance*, represents the commuting distance after firm relocations at the inventor–firm level. Firm FE is included, instead of Inventor–Firm FE. Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
D.V. (Pre-Relocation):	*Patent Count*		*Scaled Citation*		*Payment Count*	
	Cumulative	Average	Cumulative	Average	Cumulative	Average
Post-Relocation Distance	-0.014	0.001	0.181	0.050	-0.060	-0.013
	(0.064)	(0.014)	(0.142)	(0.037)	(0.131)	(0.029)
Mean of D.V.	2.831	0.741	4.418	1.190	5.491	1.439
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm Location FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.273	0.299	0.379	0.422	0.276	0.330
Observations	2,791	2,791	2,791	2,791	2,791	2,791

Table A.6: **Effect of Workplace–Home Distance on Inventor Productivity — Large Establishments**. Fixed-effects OLS regressions at the inventor–firm–year level with subsample observations whose firms have more than 100 employees on average during the sample period. Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.V.: *Patent Count*	(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.7: **Firm-Level Innovation Performance by Average Commuting Distance Changes**. Summary statistics and t-test results are provided at the firm level. Firms are divided by three groups: firms whose average commutes decrease (i.e., *Decreased*), remain unchanged (i.e., the absolute change is smaller than 1 km, *Unchanged*), and increase (i.e., *Increased*). Standard deviations are shown in parentheses.

	(1)	(2)	(3)	p-value		
Avg. Commute:	Decreased	Unchanged	Increased	(1) vs. (3)	(1) vs. (2)	(2) vs. (3)
Patent Count	0.720	0.705	0.754	0.418	0.873	0.607
	(0.646)	(0.685)	(0.684)			
Scaled Citation	1.269	1.405	1.300	0.835	0.691	0.733
	(2.432)	(2.719)	(2.182)			
Payment Count	1.443	1.569	1.546	0.291	0.556	0.920
	(1.500)	(1.899)	(1.599)			
Observations	481	59	528			

Table A.8: **Augmented Sample with Inventors Moving Homes After Firm Relocation**. Fixed-effects OLS regressions at the inventor–firm–year level with two subsamples. The first subsample is the inventors who did not move their home locations (i.e., *Non-Moved*). The second subsample is the first subsample plus the inventors who moved their home locations with observation years before their actual moving years (i.e., *Mixed*). Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
D.V.:	*Patent Count*		*Scaled Citation*		*Payment Count*	
Distance	-0.041**	-0.042**	-0.094*	-0.094**	-0.047	-0.047
	(0.019)	(0.019)	(0.049)	(0.046)	(0.042)	(0.041)
Sample	Non-Moved	Mixed	Non-Moved	Mixed	Non-Moved	Mixed
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	0.415	0.490	0.381	0.414	0.417	0.476
Inventor–Firm Count	3,445	5,255	3,445	4,612	3,445	5,255
Observations	22,863	31,099	22,863	31,099	22,863	31,099

Table A.9: **Effect of Commuting on Inventor Productivity—Bay Area Excluded**. Fixed-effects OLS regressions at the inventor–firm–year level with subsample observations excluding inventors in the San Francisco Bay Area. Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
D.V.:	*Patent Count*		*Scaled Citation*		*Payment Count*	
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 × *Top Inventor*		-0.130		-0.316*		-0.267
		(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
Mean of D.V. (Pre-Relocation):						
All Inventors	0.858		1.365		1.816	
Ordinary Inventors		0.781		1.264		1.533
Top Inventors		1.670		2.421		3.276
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.10: **Effect of Commuting on Inventor Productivity—Subsamples by Distance Shock**. Fixed-effects OLS regressions at the inventor–firm–year level with three subsamples. The subsamples are the inventors whose commuting distance decreased by more than 1 km (i.e., *Closer*), increased by more tham 1 km (i.e., *Farther*), and changed by less than 1 km (i.e., *Same*). Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.V.: *Patent Count*	(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 $\times$ *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 A.11: **Effect of Workplace–Home Distance on Inventor Productivity—Limited Distance Subsample**. Fixed-effects OLS regressions at the inventor–firm–year level with subsample including inventor–firm pairs whose workplace–home *Distance* is less than 50 km for all years of employment. Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.V.: *Patent Count*	(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.12: **Effect of Workplace–Home Distance on Inventor Productivity—Single-Authored Patents**. Fixed-effects OLS regressions at the inventor–firm–year level. The dependent variable *Single-Authored Patent Count* counts patents with the sole inventor granted to an inventor–firm pair per year. The average pre-firm-relocation *Single-Authored Patent Count* is 0.103. Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.V.: *Single-Authored Patent Count*	(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 $\times$ *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 A.13: **Effect of Distance Direction Change on Inventor Productivity**. Fixed-effects OLS regressions at the inventor–firm–year level. The independent variable Δ *Distance Direction* is a categorical variable equalling 0 prior to a workplace relocation event; after a relocation event, it takes the value of 1 for geodesic distance increases (farther away by $>$ 1 km), -1 for distance decreases (closer by $>$ 1 km), and 0 when distance is approximately unchanged (changed by less than 1 km). Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D.V.: *Patent Count*	(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 A.14: **Effect of Commuting on Inventory Productivity—Alternate Commuting Measures**. Fixed-effects OLS regressions at the inventor–firm–year level. The independent variable *Drive Distance* is the workplace–home driving distance (e.g., by car) in 10 km, and *Drive Duration* is the time it takes to drive or take public transportation between the workplace and home, whichever is less, measured in tens of minutes. Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)
D.V.:	*Patent Count*	*Scaled Citation*	*Patent Count*	*Scaled Citation*
Drive Distance	-0.038**	-0.074*		
	(0.015)	(0.038)		
Drive Duration			-0.059**	-0.104*
			(0.024)	(0.059)
Mean of D.V. (Pre-Relocation)	0.858	1.365	0.858	1.365
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 A.15: **Testing Non-Linear Effects—Square Term of Distance**. Fixed-effects OLS regressions at the inventor–firm–year level. Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)
D.V.:	*Patent Count*	*Scaled Citation*	*Payment Count*
Distance	-0.061	0.067	-0.184**
	(0.041)	(0.114)	(0.086)
Distance^2	0.003	-0.026	0.022*
	(0.006)	(0.016)	(0.012)
Mean of D.V. (Pre-Relocation)	0.858	1.365	1.684
Inventor–Firm Pair FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Firm Location FE	Yes	Yes	Yes
R^2	0.415	0.381	0.417
Inventor–Firm Count	3,445	3,445 5	3,445
Observations	22,863	22,863	22,863

Table A.16: **Testing Non-Linear Effects—Poisson and Negative Binomial Estimation**. Conditional fixed-effects Poisson and negative binomial regressions at the inventor–firm–year level. Reported coefficients are marginal effects. Robust standard errors clustered at the inventor–firm level are shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)
D.V.: *Patent Count*	Poisson Estimation		Negative Binomial	
Distance	-0.038**	-0.049**	-0.041**	-0.054**
	(0.016)	(0.020)	(0.016)	(0.023)
Inventor–Firm Conditional 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