Are Cities Losing Innovation Advantages? Online versus Face-to-face Interactions[*]

Ruben Gaetani[a†], Naqun Huang[b‡], Jing Li[c§], Yanmin Yang[b¶]

[a] *University of Toronto, Canada*

[b] *Nanjing Audit University, China*

[c] *Singapore Management University, Singapore*

December 29, 2023

Abstract

How did COVID-19 affect the innovation advantages of dense locations? Using data on the universe of U.S. patent applications, we find that the density premium in the production of novel inventions declined by 18.5%-22.9% in 2020-2021 relative to its pre-pandemic level. Smartphone data on local mobility suggest that the drop in the frequency of local interactions can explain a significant portion of this effect. While COVID-19 resulted in a temporary setback in the innovation advantages of dense locations, the role of urban density in facilitating the exchange and recombination of ideas is unlikely to be persistently replaced by online communication.

Keywords: Innovation, Density, Cities, Novelty

JEL classifications: D83, O18, O31.

[*]We thank Safegraph and IPUMS for providing access to the data. We thank Pierre-Philippe Combes, Yige Duan, Gilles Duranton, Cecile Gaubert, Guangbin Hong, Tomoya Mori, Yanos Zylberberg, and seminar participants at the University of Calgary, Liaoning University, the 2022 SMU Conference on Urban and Regional Economics, the 2023 European Meetings of the Urban Economics Association, the 2023 Asia Meeting of the Econometric Society in China, the 1st Summer Meeting in Urban Economics at Peking University, and the 2023 Asian Meeting of the Econometric Society in East and Southeast Asia. Any remaining errors are our own.

[†]Address: 105 St George Street, Toronto, Canada, M5S 2E8. Phone: +1-905-569-4770. E-mail: ruben.gaetani@utoronto.ca.

[‡]Address: 86 West Yushan Road, Nanjing, Jiangsu, China, 211815. Phone: +86-188-6096-5090. E-mail: naqun.huang.2012@phdecons.smu.edu.sg.

[§]Address: 90 Stamford Road, Singapore 178903. Phone: +65-6808-5454. E-mail: lijing@smu.edu.sg.

[¶]Address: 86 West Yushan Road, Nanjing, Jiangsu, China, 211815. Phone: +86-189-1636-3961. E-mail: yanmin_yang@nau.edu.cn.

"*Physical proximity is important to just about everything we [Silicon Valley companies] do. [...] The level of communication is much higher when you can see each other regularly. You never work on the same level if you do it by telephone and airplane.*"

(Saxenian, 1996)[1]

"*[A]t the [Bell] Labs the experimentalists and theoreticians were encouraged to work together, and that chemists and metallurgists were welcome to join in, too. The interactions could be casual, but the work was a serious matter. [...] It's the interaction between fundamental science and applied science, and the interface between many disciplines, that creates new ideas.*"

(Gertner, 2012)

1 Introduction

An extensive literature in economic growth and geography places cities at the center of the innovation process (Glaeser et al., 1992; Black and Henderson, 1999). This literature emphasizes the role that urban density plays in facilitating random encounters among people from diverse knowledge backgrounds, spurring the creation of novel ideas (Saxenian, 1996; Duranton and Puga, 2001; Gertner, 2012; Packalen and Bhattacharya, 2015).

The onset of the COVID-19 pandemic in early 2020 drastically reduced opportunities for face-to-face interactions in cities. The concurrent adoption of work-from-home (WFH) practices was fueled by a growing availability of tools of online communication, which were designed to provide a viable alternative to in-person interactions. Those tools effectively connected people belonging to the same formal networks (e.g., close colleagues) and opened up new possibilities for collaboration at long distance. However, whether they could generate comparable opportunities for informal knowledge exchange as those provided by urban density—thereby persistently eroding the advantage of cities

[1] Quoting Tom Furlong, former manager of Digital Equipment Corporation (DEC)'s workstation group in Palo Alto.

in the innovation process—is still a widely debated question (e.g., Nathan and Overman, 2020; Emanuel and Harrington, 2021; Yang et al., 2022; Emanuel et al., 2023; Gibbs et al., 2023).

This paper confronts this question empirically. Using a dataset of geolocated U.S. patent applications combined with a text-based measure of patent novelty, we show that the advantage of dense locations in generating the most novel inventions decreased significantly after the onset of the pandemic. Leveraging smartphone data on local mobility patterns, we also show that a substantial part of this decline was due to the fact that dense locations experienced a larger drop in the frequency of face-to-face encounters. These results suggest that the advantage provided by urban density in the innovation process lies in its ability to create opportunities for in-person interactions, which facilitate the exchange of ideas across diverse fields and their recombination in novel ways. While COVID-19 resulted in a temporary setback in the centrality of dense locations in innovation, their role is unlikely to be persistently replaced by online communications.

Our analysis is based on the universe of patent applications at the United States Patent and Trademark Office (USPTO) between 2011 and 2021, that we geolocate at the level of County Sub-Division (CSD). CSDs are the finest partition of the United States to which inventors' locations can be reliably assigned using the information in the application's full text. This level of disaggregation allows us to exploit the large variation in population density that exists within larger geographical units such as metropolitan areas or commuting zones (Berkes and Gaetani, 2021).

Following the approach developed by Arts et al. (2021), we assign to each patent application a measure of novelty defined as the count of new word pairs that appear in the application's text but do not occur in the full text of any of the earlier patents granted since 1976. We show that this novelty measure is predictive of a range of important outcomes of the invention, including its impact and the span of its geographical and technological diffusion. Based on this application-level measure, we define the local degree of innovativeness as the right tail of the novelty distribution for each CSD-year observation (in our baseline specifications, we use the 99$^{\text{th}}$ and 95$^{\text{th}}$ percentiles). We focus on the

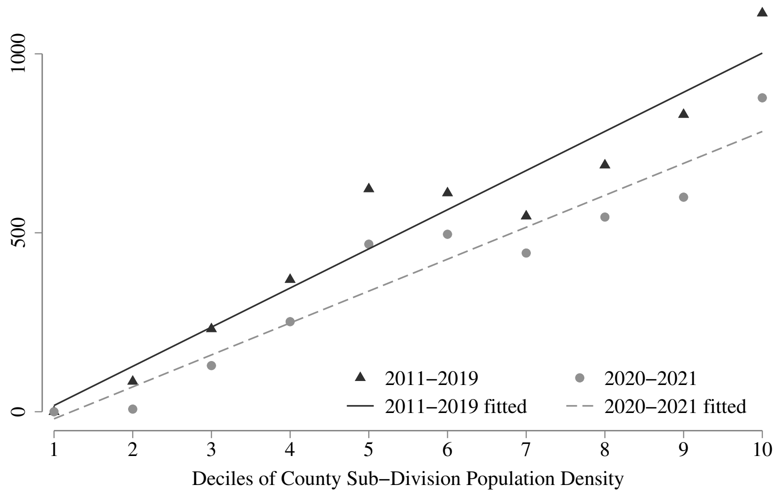

Figure 1: County Sub-Division (CSD) Population Density and Patent Innovativeness

Notes: The local degree of patent innovativeness is defined as the 99th percentile value of the local novelty distribution for patent applications filed in a CSD in a year.

right tail because, consistently with what is widely documented in the literature (e.g., Scherer, 1965; Hall et al., 2005), we find that the distribution of technological value is highly skewed and mostly concentrated among inventions with exceptionally high novelty.

The empirical framework is a difference-in-differences specification at the CSD-year level, where the outcome variable is the local degree of innovativeness, the intensity of treatment is population density prior to the beginning of the sample, and the timing of treatment coincides with the onset of the COVID-19 pandemic in early 2020.

The main results indicate that the outbreak of COVID-19 had a large negative effect on the density premium in innovativeness. Figure 1 summarizes this result. The graph shows the change in the local degree of innovativeness between pre-pandemic years (2011-2019) and the years of the pandemic (2020-2021) for locations at different deciles of the density distribution. More densely populated CSDs experienced a larger drop in local innovativeness relative to less densely populated ones. The difference-in-differences estimates imply that the density premium declined by 18.5% to 22.9% in 2020-2021 relative to its pre-pandemic level. In an event study setting, we find no evidence of pre-trends in the density premium, suggesting that the decline observed in 2020-2021 was likely a direct outcome of the onset of the pandemic.

We test the robustness of these results to a range of alternative specifications. First, we include a set of flexible controls for the number of patents within each CSD-year observation. This allows us to verify that the results are not mechanically confounded by changes in the local patent count. Second, we re-estimate the relationship at a monthly frequency. This allows us to define the treatment period as beginning in March 2020, when restrictions to interactions were first widely adopted in the United States. Third, we redefine units of observation as CSD-month-technological areas. Since COVID-19 likely created uneven opportunities for novel inventions across technological fields, this specification allows us to verify that the main results are not driven by differences across CSDs in the composition of local patenting. Finally, we include a set of time-varying state-level controls to account for spatial differences in the severity of the pandemic and its effects on the local economy. All these alternative specifications deliver estimates that are fully consistent with the baseline.

Using the same empirical framework, we show that the density premium in overall patent quantity was not affected by the onset of COVID-19. However, the density premium in the count of highly novel patents experienced a significant drop in 2020-2021 relative to its pre-pandemic level. Furthermore, the change in the density premium only involved the right tail of the local novelty distribution, and did not affect lower percentiles such as the the median and the 75$^{\text{th}}$ percentile. These results suggest that online communication worked as an effective substitute for the interactions underlying the majority of inventions but not for those underlying the most novel ones.

Having established the main fact, we provide direct evidence for our candidate mechanism underlying the observed drop in the density premium. Using smartphone data on local mobility patterns provided by SafeGraph, we show that densely populated locations experienced a larger drop in the intensity of in-person interactions following the onset of COVID-19. Part of this difference is due to the fact that dense locations initially exhibit a higher frequency of face-to-face interactions and, as a result of the restrictions widely implemented in 2020, experienced a more pronounced fall in the intensity of these interactions.

To investigate to what extent this channel can explain the observed decline in the density premium, we proceed in two steps. First, using the same mobility data, we recover CSD-specific proxies for the drop of in-person interactions both at the workplace and in informal settings (i.e., restaurants and other dining venues). Then, we re-estimate our difference-in-differences specification including these proxies as additional controls. The inclusion of these controls reduces the magnitude of the estimated effects by one-third to two-thirds relative to the baseline. These results suggest that a substantial portion of the decline in the density premium in innovativeness can be attributed to the sharp drop in the intensity of in-person interactions, which disproportionately affected densely populated locations.

Using information on inventors' demographic characteristics, we can rule out other plausible channels behind the observed drop in the density premium. A competing channel stems from the observation that the productivity of female inventors or inventors with young children was adversely impacted by the onset of the pandemic more than other demographic groups (Myers et al., 2020). This fact might confound our interpretation if more affected inventors are more innovative and are more concentrated in dense locations. Another competing explanation lies in the possibility that inventors generating the most innovative ideas might have been more inclined to relocate away from densely populated cities after the onset of the pandemic. In application-level regressions, we show that the estimate of the drop in the density premium in 2020-2021 is not affected by the inclusion of dummy variables for female inventors, young inventors, and inventors who moved to a different location after the outbreak of COVID-19. These findings suggest that those alternative channels are unlikely to play a significant role in explaining the main result.

Overall, our findings support the idea that the ability of dense cities to foster random and unplanned interactions underlies their pivotal role in the innovation process. This explains why the widespread restrictions implemented during the pandemic resulted in a temporary drop in the innovation advantages of dense locations. However, despite the widespread adoption of online communication tools, virtual interactions are unlikely to fully replace the spontaneous nature of face-to-face encounters that urban density has long facilitated. The profound transformations brought about by the COVID-19 pandemic are unlikely to persistently erode cities' position as engines of creativity and groundbreaking inventions.

This paper contributes to the vast literature on the geographical agglomeration of innovation activities (e.g., Carlino et al., 2007; Kerr and Robert-Nicoud, 2020). Many papers within this literature have studied the nature of the agglomeration forces sustaining cities' advantage in the innovation process (e.g., Glaeser et al., 1992; Duranton and Puga, 2001; Packalen and Bhattacharya, 2015; Davis and Dingel, 2019; Berkes and Gaetani, 2021). Another set of papers has provided direct evidence of the role of face-to-face encounters in knowledge creation and diffusion (e.g., Catalini, 2018; Andrews, 2019; Catalini et al., 2020; Atkin et al., 2022; Brucks and Levav, 2022, Koh et al., 2023). Our paper contributes to this literature by providing the first evidence of the effect of a large reduction in the opportunities of in-person interactions on the ability of dense cities to generate novel inventions.

This paper also contributes to the growing literature on remote work and its implications for the spatial organization of economic activities (Delventhal and Parkhomenko, 2020; Gupta et al., 2022; Monte et al., 2022). Existing studies have emphasized the benefits of WFH, such as saving on commuting and increased working time (Barrero et al., 2020; Teodorovicz et al., 2022) and quieter work environment (Bloom et al., 2015), but have also pointed at its limitations, such as lower rates of learning (Emanuel et al., 2023) and information sharing (Yang et al., 2022). Our paper contributes to this debate by showing that there are limits to the extent to which remote work can substitute for the richness of in-person interactions provided by urban density in the innovation process.

The remainder of the paper is organized as follows: Section 2 develops a stylized theoretical framework that derives the main testable hypotheses; Section 3 introduces the data and the measurement of invention novelty and local interactions; Section 4 presents the empirical framework and main results; Section 5 explores possible mechanisms; Section 6 concludes.

2 Theoretical Framework

We present a stylized theoretical framework that illustrates how the intensity of local interactions varies with population density, how this variation gives rise to a density premium in innovativeness, and how restrictions to local interactions, such as those implemented in response to the COVID-19 pandemic, generate a drop in the density premium.

The model focuses on the innovation outcomes of a representative inventor living in a given location. Locations are indexed by n and characterized by a level of density D_n. Each inventor has a unit of time that can be allocated either to in-person interactions (IP) or to online communication (OL), with $\alpha \in [0, 1]$ denoting the share of time devoted to in-person interactions. The two modes of communication translate into different innovation outcomes along two dimensions.

First, the rate of arrival of insights per unit of time devoted to in-person interactions, $B_{IP,n}$, is an increasing function of local density, reflecting the fact that denser locations provide more frequent opportunities for meetings and idea exchange. In particular, we impose $B_{IP,n} = D_n^\gamma$, with $\gamma > 0$. By contrast, the rate of arrival of insights per unit of time devoted to online communication is independent of density and normalized to a constant, B_{OL}.

Second, the degree of novelty of the resulting inventions, denoted by N, is higher for the insights developed through in-person interactions than for those developed via online communications, i.e., $N_{IP} > N_{OL}$. This reflects the intuition, first proposed by Jacobs (1969) and later supported empirically by Glaeser et al. (1992), Packalen and Bhattacharya (2015), and Berkes and Gaetani (2021), among others, that density facilitates the random, informal, and unplanned interactions that catalyze the flow of ideas across disconnected fields and give rise to novel inventions building on atypical combinations of knowledge. In the model, inventors care about novelty because, as we show empirically in the next section, novel inventions are more valuable. Letting $v(N)$ denote the value of an invention of novelty N (with $v(\cdot)$ being an increasing function), this feature implies $v(N_{IP}) > v(N_{OL})$, where we normalize $v(N_{OL}) = 1$ and define $\bar{v} \equiv v(N_{IP})$.

The inventor's problem is to choose the intensity of in-person interactions, α, to maximize the total value of their inventions net of a quadratic cost of in-person interactions:

$$V_n = \max_{\alpha \in [0,1]} (1-\alpha) B_{OL} + \alpha D_n^{\gamma} \bar{v} - \frac{c}{2} \alpha^2, \tag{1}$$

where $c > 0$ controls, among the other things, the time saved on commuting when online communications happen in the context of WFH. This leads to a straightforward solution for the optimal intensity of in-person interactions in city n:

$$\alpha_n^* = \frac{D_n^{\gamma} \bar{v} - B_{OL}}{c}, \tag{2}$$

where we impose sufficient restrictions on the primitives to guarantee that $\alpha_n^* \in (0, 1)$. From this equation, it is immediate to see that the optimal intensity of in-person interactions is an increasing function of local density, from which we derive the first prediction:

Prediction 1 *More densely populated locations are characterized by a higher intensity of in-person interactions.*

The main outcome of interest is the degree of innovativeness in location n, which is defined as:

$$\bar{N}_n = \frac{1}{Q_n} \left[(1 - \alpha_n^*) B_{OL} N_{OL} + \alpha_n^* D^{\gamma} N_{IP} \right], \tag{3}$$

where $Q_n = (1 - \alpha_n^*) B_{OL} + \alpha_n^* D^{\gamma}$ is the local invention rate. Since α_n^* is increasing in density and $N_{IP} > N_{OL}$, from Equation (3) it immediately follows innovativeness is higher in more densely populated locations:

Prediction 2 *More densely populated locations are characterized by a higher degree of innovativeness.*

There are two factors that contribute to this density premium in innovativeness. First, for a given time devoted to in-person interactions, dense locations provide a higher frequency of meetings, which increases the rate of arrival of novel insights. Second, inventors in dense locations endogenously devote more time to in-person interactions, leading to more frequent novel insights.

In the empirical analysis, we explore how the density premium in innovativeness responded to the drastic drop in face-to-face interactions after the onset of the COVID-19 pandemic. To visualize this effect in this simple model, assume that mandated restrictions forced inventors in all locations to reduce the intensity of in-person interactions to a uniform level $\tilde{\alpha}$. As we formally show in Appendix D, if $\tilde{\alpha}$ is sufficiently small the resulting change in the degree of innovativeness is always negative and is larger in absolute value for denser locations, leading to the following:

Prediction 3 *Restrictions bringing in-person interactions down to a low and uniform level induce a drop in the density premium in innovativeness.*

In the remainder of the paper, we provide empirical evidence that is consistent with these predictions. First, we construct a patent-level measure of novelty and show that it correlates with a range of indicators of technological value. We show that the most novel inventions tend to originate from more densely populated areas, giving rise to a density premium in innovativeness. Then, we show that the onset of COVID-19 induced a sharp decline in the density premium. Finally, using data on local mobility and interactions, we show that the drastic reduction in face-to-face interactions in response to the outbreak of COVID-19 explains a significant portion of the drop in the density premium.

3 Data and Measurement

The empirical analysis combines data on georeferenced patent grants and applications from the USPTO, local data on population density from the IPUMS National Historical GIS (NHGIS), and two datasets on local interactions and mobility patterns provided by Safegraph. Our primary units of analysis are County Sub-Divisions (CSDs), which constitute the finest partition of the United States for which the location of inventors can be reliably identified using the information contained in a patent's text (Berkes and Gaetani, 2021). CSDs are significantly smaller than commuting zones and metropolitan areas, and display large variation in population density within those larger geographical units.

3.1 Patents data and novelty measure

Our primary dataset contains detailed information on granted patents and patent applications released by the USPTO. The information on granted patents is available since 1976. All patent applications (whether granted or not) are published from 2001 onward. For each patent, the USPTO provides the application ID, the filing date, the issue date (if granted), the cities of residence of all inventors, the Cooperative Patent Classification (CPC) code,[2] as well as the full text of the patent, including the title, the abstract, and detailed claims. The full text of granted patents includes a field containing citations to previously granted patents or patent applications. This field is not available in the full text of patent applications not yet granted.[3] Therefore, we only consider citations given by granted patents.

We georeference patent applications at the CSD level based on the cities of residence of all the listed inventors.[4] To georeference patents with multiple inventors living in different CSDs, we follow two complementary approaches. In the first approach, we assign the patent to a CSD as long as at least one of its inventors lives in the CSD. Note that, in this case, one patent can be assigned to multiple CSDs. In the second approach, we assign patents to CSDs fractionally, based on the share of inventors living in the CSD. Using these two approaches, we produce patent counts for each CSD-year. We label the former as the raw counts and the latter as the inventor-weighted counts of patent applications. Summary statistics of both the raw counts and the inventor-weighted counts of patent applications are reported in Table A1.

Our analysis is based on all patent applications filed since 2011. The use of patent applications allows us to have a time horizon that extends well beyond the onset of COVID-19, which, due to long pendency times, would not be possible with patent grants.[5]

[2]CPC is a patent classification system, which has been jointly developed by the European Patent Office (EPO) and the United States Patent and Trademark Office (USPTO). Details can be found at https://www.uspto.gov/web/patents/classification/cpc/html/cpc.html.

[3]For patent applications, the USPTO releases scanned images of the information disclosure statement of the applications that contains references to all prior knowledge. Unfortunately, this information is largely not machine-readable.

[4]See Online Appendix B for details on the nature of CSDs and how the matching is conducted.

[5]The USPTO reports an *average* pendency of about 25 months between filing and final disposition. By contrast, all patent applications are required to be published within 18 months of filing, which

3.1.1 Construction of the novelty measure

In this paper, we confront the hypothesis that local interactions, by offering opportunities for face-to-face meetings among people with diverse knowledge backgrounds, are a critical input to the groundbreaking ideas that push the boundaries of knowledge forward. To test this hypothesis, we study how the onset of COVID-19, which drastically reduced the availability of in-person interactions in densely populated locations, affected the ability of those locations to generate novel inventions.

The main outcome of interest is based on a patent-level measure of novelty constructed following the approach proposed by Arts et al. (2021). In essence, this measure captures the number of new word combinations in a patent's text that do not appear in any of the earlier patent applications.[6]

The construction of the measure occurs in two steps, which we briefly sketch here. Since we closely follow Arts et al. (2021), we refer the reader to the original paper for more details. First, for each filing year t, we construct a knowledge pool made of all the word pairs that have appeared in the full text of the same patent at least once between the beginning of our sample and year $t-1$. In this step, to maximize the time horizon used to build the knowledge pool, we use all patent grants issued since 1976, because the full text of patent applications is only available starting from 2001. Second, for each patent application filed in year t, we define novelty as the number of unique word pairs that do not appear in their corresponding knowledge pool (that is, the knowledge pool built using patents filed up to year $t-1$).[7]

Using this patent-level measure of novelty, we define two CSD-year-level degree of innovativeness as the weighted 99$^{\text{th}}$ and 95$^{\text{th}}$ percentiles of the novelty distribution among

means that our data contain almost all patent applications filed between 2011 and 2020 (data collection took place in August 2022). Note that after 18 months from the filing date, under some special and rare circumstances, a patent application may still be confidential to the patent office and not publicly available. Please refer to the USPTO for examples.

[6]We use a text-based measure of novelty since citation-based measures, such as *originality* (Trajtenberg et al., 1997) and *unconventionality* (Berkes and Gaetani, 2021), would not be feasible in our setting due to the lack of backward citation data for recent patent applications.

[7]As detailed in Section 2.1 of Arts et al. (2021), in constructing the knowledge pool and the novelty measure, we exclude common words such as stop words and words that are unrelated to the technical content of the patent. Arts et al. (2021) constructed five text-based novelty measures. Among these measures, the number of new word combinations demonstrates the strongest discriminatory power to detect novelty. We winsorize the novelty measure at the 99$^{\text{th}}$ percentile of all the patent applications.

all patent applications in each CSD-year. In other words, we think of local innovativeness as the degree of novelty at the right tail of the local distribution. Table A1 reports the summary statistics at the CSD-year level for the period between 2011 and 2021. The average CSD-year specific innovativeness based on the 99$^{\text{th}}$ percentile is 872 and that based on the 95$^{\text{th}}$ percentile is 445.

3.1.2 Novelty predicts impact and span of diffusion

In this subsection, we use citations data to show that this measure of novelty predicts important features of the invention, including its technological impact and the geographical and technological span of its diffusion.[8] This evidence suggests that boosting the creation of more novel ideas also results in inventions that are more impactful and diffuse more broadly, with implications for the welfare effects of policies that affect their supply.

To explore the effect of novelty on impact and span of diffusion, we estimate patent-level regressions of the following form:

$$Y_{ict} = \alpha + \mu_{ct} + \sum_{d \in D} \beta_d \times \mathbb{1}(d(NWC_{ict}) = d) + \varepsilon_{ict}, \tag{4}$$

where the dependent variable, Y_{ict}, represents three different sets of outcomes associated with patent i. First, we consider indicators for whether the patent's technological impact, measured as the number of forward citations received, falls among the top 1%, top 5%, or top 10% among patents filed in the same year t and belonging to the same technology class c. Second, we consider a measure of the span of technological diffusion of patent i as the number of different technology classes (other than its own) from which the patent receives citations. Third, we consider an analogous measure of the extent of spatial diffusion as the mean of pairwise geographic distances between cities of patent i's inventors and those of all the patents citing i.

The explanatory variables are a set of indicators $\mathbb{1}(d(NWC_{ict}) = d)$ of the decile of patent i's novelty (count of new word pairs) among patents filed in the same year t and

[8] Arts et al. (2021) show that patents with high novelty according to this text-based measure are more likely to be linked to a major award, and less likely to be later rejected by both the European and the Japanese patent offices.

belonging to the same technology class c. Since about 30% of the patents in the sample have a novelty score equal to zero, we group the bottom three deciles in one category and use this as the omitted category. We control for technology class-filing year fixed effects, μ_{ct}. Throughout the paper, ε_{ict} denotes an idiosyncratic error term.

In estimating Equation (4), we restrict the sample to patents filed between 2011 and 2015, for which the time horizon to measure forward citations is sufficiently long. We only consider citations by patents granted before 2020 (the treatment year in the main analysis). All variables are summarized in Table A2.

The regression results are reported in Table A3, and the estimated coefficients are plotted in Figure A2. Panels A, B, and C of Figure A2 correspond to the estimated coefficients of the decile of a patent's novelty score when predicting the probability of being a hit patent, the span of diffusion across fields, and the span of diffusion across locations, respectively.

Panel A shows that more novel patents are more likely to be hit patents. The estimated coefficients imply that patents in the 8^{th}, the 9^{th}, and the 10^{th} decile of novelty score are 1.0, 1.4. and 3.1 percentage points more likely to be in the top 5% in terms of forward citations relative to patents in the baseline group. The increase in the probability of being a hit patent is especially pronounced going from the 9^{th} to the 10^{th} decile, highlighting that the most impactful inventions are mostly concentrated in the right tail of the novelty distribution. Similar patterns emerge when defining hit patents as those in the top 1% or the top 10% of citations received.

Panel B shows that more novel patents have a larger span of technological diffusion, as proxied by the number of technology classes from which they are cited. The estimated coefficients imply that patents in the 8^{th} decile of the novelty distribution receive citations from 0.2 more technology classes compared to patents in the baseline group. The span of diffusion increases even further for patents in the 9^{th} decile (0.3) and 10^{th} decile (0.5).

Finally, Panel C presents a positive relationship between a patent's novelty score and the average geographical distance between the patent itself and other grants citing it. The average distance between patents in the 8^{th} decile of the novelty distribution and its citing patents is, on average, 63 kilometers higher than for patents in the baseline group. The increase in the geographical span of diffusion is even larger for patents in the 9^{th} decile (79 kilometers) and the 10^{th} decile (117 kilometers). These results highlight that the largest extent of technological and geographical diffusion is concentrated among patents in the right tail of the novelty distribution.

In summary, novel inventions generate large technological and geographical spillovers. This suggests that shocks affecting their supply (such as the one we will document in the remainder of the paper) can have large effects on welfare, opening up room for policy intervention to improve efficiency.

3.2 Other data sources

We construct CSD-level measures of population density from the 2010 Census, which we obtain from the IPUMS National Historical GIS (NHGIS). Figure A1 displays a histogram of CSD-level log-population density. The mean and the standard deviation of this measure are 618 and 1070 people per square kilometer, as reported in Table A1.

The mechanism analysis of Section 5 leverages data on local mobility and interactions provided by the company Safegraph, which collects GPS pings from a large sample of anonymous mobile devices. We use two datasets provided by Safegraph, which we briefly introduce here and on which we provide more details in Section 5.

The first dataset is the Social Distancing Metrics, containing daily mobility information from January 1^{st}, 2019 to December 30^{th}, 2020, aggregated by the devices' home Census block group. We aggregate this data at the CSD level by assigning each Census block group to a CSD based on where the largest share of its population resides. The measures from January 1^{st} to May 9^{th}, 2020 are constructed following a different methodology from that used from May 10^{th}, 2020, onwards.[9] Hence, to quantify the effect of the onset of COVID-19 on local mobility, we restrict the sample to the period between January 1^{st} and April 30^{th} in 2020, and use the same months in 2019 as a point of comparison.

[9]See the released note in https://docs.safegraph.com/docs/social-distancing-metrics.

The second dataset is the Weekly Pattern data. This dataset reports information on the number of visits and visitors for different consumer points of interest (POIs). In our analysis, we focus on POIs belonging to the North American Industry Classification System (NAICS) code 7225 which stands for restaurants and other eating places, where the largest part of informal meetings are likely to take place. Specifically, NAICS 7225 comprises NAICS 722511 (full-service restaurants), NAICS 722513 (limited-service restaurants), NAICS 722514 (cafeterias, grill buffets, and buffets), and NAICS 722515 (snack and nonalcoholic beverage bars). Using this data, we construct measures of local interactions based on information on hourly visits and weekly visitors to those POIs.

4 The Impact of COVID-19 on the Density Premium

We now turn to the empirical test of our hypothesis. In particular, we test whether the drastic reduction in the frequency of in-person interactions induced by COVID-19 results in a decrease in the novelty content of new inventions. In this section, we show that the density premium in innovativeness, i.e., the degree to which more densely populated locations produce more novel inventions, declined significantly after the onset of the pandemic. In Section 5, we show that a substantial portion of this effect can be accounted for by the fact that more densely populated locations experienced a more severe drop in the frequency of in-person interactions.

4.1 Empirical framework

The research design is a continuous difference-in-differences, where the treatment intensity is given by local population density and the time of treatment coincides with the onset of COVID-19. Specifically, we consider regressions of the following form:

$$I_{lt} = \alpha + \beta \log(Density_l) \times \mathbb{1}(After_t) + \eta_l + \theta_t + \varepsilon_{lt}. \tag{5}$$

In Equation (5), the dependent variable, I_{lt}, represents the degree of innovativeness, defined as either the 99$^{\text{th}}$ or the 95$^{\text{th}}$ percentile value of the novelty distribution of patent applications filed in CSD l and year t. The main coefficient of interest, β, captures the effect of the interaction between log-population density ($Density_l$) and a dummy variable ($\mathbb{1}(After_t)$) that takes value 1 if the filing year is 2020 or later and 0 otherwise. A negative estimate of β indicates that more densely populated locations experienced a more pronounced decline in their degree of innovativeness after the onset of COVID-19 compared to less densely populated locations. The specification includes CSD fixed effects, η_l, capturing time-invariant characteristics that affect local innovativeness, and year fixed effects, θ_t, capturing changes over time in the average degree of innovativeness. The identification assumption is that the density premium in innovativeness would stay the same in the absence of the COVID-19 shock.

4.2 Results

Table 1 reports the estimates of the coefficients in Equation (5). The local degree of innovativeness is measured as the 99$^{\text{th}}$ percentile novelty measure for columns 1 and 2 and as the 95$^{\text{th}}$ percentile for columns 3 and 4.

In columns 1 and 3, we leave out CSD fixed effects. This allows us to identify the density premium before the onset of COVID-19. Higher population density is associated with significantly higher innovativeness. The estimates from columns 1 and 3 of Panel A imply that a 10% increase in population density increases the count of new word combinations at the 99$^{\text{th}}$ percentile by about 22.1, and that at the 95$^{\text{th}}$ percentile by about 7.3. These magnitudes are economically large: moving from a location at the fourth decile of the density distribution (e.g., Armonk, NY, home of the IBM headquarters, with a population density of 192) to a location at the ninth decile of the density distribution (e.g., Redmond, WA, home of the Microsoft headquarters, with a population density of 970) increases the local degree of innovativeness by 1,474.4 and 484.6, respectively, corresponding to about 169% and 109% of its unconditional mean. This finding is consistent with the evidence in Berkes and Gaetani (2021) who, using citation-based measures, show that densely populated locations are more likely to produce inventions building on unconventional combinations of prior knowledge. Our result suggests that the density

Table 1: Impact of COVID-19 on Density Premium of Patent Innovativeness

	(1) 99^{th}	(2) 99^{th}	(3) 95^{th}	(4) 95^{th}
log(pop density)	221.49	-	72.80	-
	(9.19)	-	(5.38)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-40.56	-41.12	-17.28	-16.72
	(6.63)	(6.64)	(5.05)	(5.09)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.07	0.57	0.02	0.39
Obs.	46,041	46,041	46,041	46,041

Notes: All regressions follow Equation (5). The dependent variable is the local degree of innovativeness defined as the 99^{th} percentile value of the local novelty distribution (Columns (1) and (2)) or the 95^{th} percentile value of the local novelty distribution (Columns (3) and (4)). Standard errors clustered at the CSD level are reported in parentheses.

premium in innovativeness is also visible with our text-based measure of novelty.

In columns 2 and 4, we include CSD fixed effects, capturing time-invariant differences in innovativeness across locations. Note that, in this case, the coefficient on population density is not identified. The coefficients of the interaction terms are negative and significant for both measures of innovativeness. In other words, more densely populated CSDs experience a greater loss in innovativeness after the outbreak of COVID-19. The size of the coefficients imply that the density premium decreases by 18.5% (column 2) and 22.9% (column 4) relative to its pre-pandemic level.

A potential concern with the estimates in Table 1 is that, since the outcome variable is defined as the right tail of the local novelty distribution, results might be mechanically driven by a larger drop in the patent count of more densely populated locations after the onset of COVID-19. As we show in Section 4.2.1 below, we find no evidence that the change in the patent count varied systematically with population density. However, to test directly for this potential bias, we augment our baseline specification of Equation (5) by including a range of controls for the local number of patent applications. Results are displayed in Table A5. In Panels A and B we control for a linear and for a quadratic function of patent applications in each CSD-year observation, respectively. The inclusion of these controls has no significant effect on the estimates of our coefficients of interest.

Another potential concern is that, in the yearly specification of Equation (5), the beginning of the treatment period is defined as January 2020, while the restrictions to local interactions in response to COVID-19 in the United States were only widely implemented in March of that year. In Panel A of Table A4, we estimate Equation (5) at a monthly frequency, imposing March 2020 as the beginning of the treatment period. Results are fully consistent with the ones in the yearly analysis.

Furthermore, the distribution of patenting across technological fields varies across locations. The onset of COVID-19 generated unprecedented opportunities for novel ideas, but the arrival of these opportunities was likely heterogeneous across fields. This heterogeneity could have resulted in a drop in the density premium if patenting in denser locations was systematically biased towards fields where COVID-19 created lower opportunities for novel inventions. To address this concern, in Panel B of Table A4, we estimate Equation (5) at the CSD-CPC Section level, again at a monthly frequency.[10] The results are fully consistent with the baseline analysis. This suggests that differences in the composition of local patenting across technological fields are unlikely to be driving the observed drop in the density premium.

A final concern with the main results is that the severity of the pandemic and its effects on the local economy varied significantly across geographical areas, potentially confounding our estimates. To address this concern, in Table A6 we augment our baseline specification with state-level time-varying controls for the severity of the pandemic (measured as the cumulative number of recorded infections per 100k people) and the state of the economy (measured as the unemployment rate). The inclusion of these controls has a negligible effect on our coefficients of interest.

4.2.1 Testing for pre-trends

The interpretation of these coefficients as the effect of COVID-19 on the density premium rests on the assumption that the local trends in innovativeness across locations before 2020 did not vary systematically with population density. To check the validity of this

[10]CPC Sections capture nine broad technological categories assigned by the patent examiner to the patent application based on its contents.

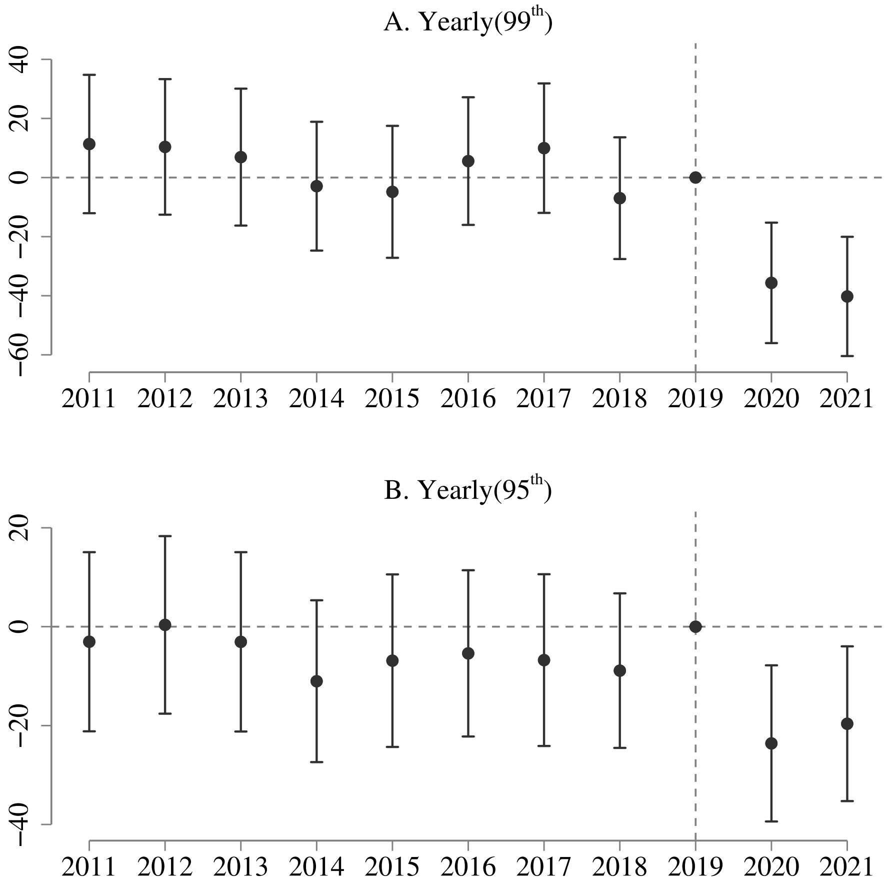

Figure 2: Temporal Changes in Density Premium of Patent Innovativeness

Notes: This figure plots the estimated coefficients and the corresponding 95% confidence intervals from estimating Equation (6). The dependent variable is the local degree of patent innovativeness defined as the 99^{th} percentile value of the local novelty distribution in Panel A or the 95^{th} percentile value of the local novelty distribution in Panel B.

parallel trends assumption, we conduct an event study analysis based on the following specification:

$$I_{lt} = \alpha + \sum_{\tau \in T} \beta_\tau \log(Density_l) \times \mathbb{1}(t = \tau) + \eta_l + \theta_t + \varepsilon_{lt}, \tag{6}$$

where T is the set of years between 2011 and 2021 (with 2019 omitted as the reference year). This specification is analogous to Equation (5), but it includes separate year-specific dummy variables interacted with population density. This allows us to estimate separate coefficients (β_τ) of the effect of density on innovativeness for each year in the sample period.

Figure 2 plots the estimated coefficients, together with the corresponding 95% confidence intervals, when the local degree of innovativeness is measured as the 99^{th} (upper panel) or the 95^{th} (lower panel) percentile value of the local novelty distribution. The density premium of innovativeness was stable before COVID-19, with none of the earlier coefficients being significantly different from that in 2019. Starting from 2020, there is a significant decrease in the density premium. These results confirm that there was no trend in local innovativeness that varied systematically by population density, suggesting that the main results in Table 1 are likely explained by the onset of COVID-19 in early 2020.

Figure A3 replicates this analysis at a monthly frequency, at the CSD-level (upper panel) and the CSD-CPC Section level (lower panel), when the local degree of innovativeness is measured as the 99^{th} percentile value of the local novelty distribution. Both levels of analysis show no evidence of pre-trends before March 2020. Similar results are obtained when the local degree of innovativeness is measured as the 95^{th} percentile value (Figure A4). The drop in the density premium becomes evident starting from June 2020, which is consistent with the existence of a lag between idea inception and the timing of a patent application.

4.2.2 Impact on the density premium in patent quantity

These results raise an immediate question: did COVID-19 have a broader impact on the density premium in patent quantity, or is the effect limited to the most novel inventions? Theoretically, a drop in local interactions may have an ambiguous effect on the overall innovation output. While reducing in-person interactions may curtail idea exchange, the concomitant surge in online communication may boost the production of less novel inventions, for which idea exchange can be more easily coordinated by individuals and organizations. These two effects run in opposite directions, making the impact of COVID-19 on the density premium in total patenting ex-ante ambiguous.

Table A7 shows estimates of regressions analogous to Equation (5), but in which we use the local patent count as the outcome variable (since many patent applications from 2021 have not been published at the time of data collection, we drop 2021 in the regressions discussed here). Columns 1 and 2 use the raw count of patent applications, while columns 3 and 4 use the inventor-weighted count. As a point of reference to interpret magnitudes, columns 1 and 3 show the estimates without CSD fixed effects. Before COVID-19, a 10% increase in population density in a CSD was associated with 4.0 more patent applications per year in the raw count, and 1.9 more applications per year in the inventor-weighted count (Panel A). Columns 2 and 4, which include CSD fixed effects, show that, for both measures, there is no detectable effect of COVID-19 on the density premium. We obtain similar results when conducting the analysis at the monthly level (reported in Panel B) and at the month-by-CPC Section level (reported in Panel C).

Table A8 provides further evidence that the change in the density premium for tail novelty represents a *compositional* change, where patenting in denser locations shifts away from the most novel inventions and towards less novel ones. Panel A (B) shows estimates of regressions analogous to Equation (5), but in which the outcome variable is the count of patents in the top 5% (bottom 95%) of the overall novelty distribution. The density premium drops significantly during the pandemic for the count of the most novel inventions, but it remains stable for the count of the less novel patents. Tables A9 and A10 show consistent results when high-novelty patents are defined as those in the top 1% and 10% of the novelty distribution, respectively.

4.2.3 Impact on the overall novelty distribution

Another important question is whether the effect of COVID-19 only concerned the density premium at the right tail of the novelty distribution, or it involved the overall distribution, affecting, for example, the novelty of the median patent. To address this question, we conduct a set of tests by choosing alternative percentile values from the local novelty distribution—the 90^{th}, 75^{th}, or the 50^{th} percentile—and examine whether the density gradient of those values is also affected by the COVID-19 shock. Table A11 reports the corresponding results, showing that the coefficients associated with the interaction terms are small and not significant.

A possible interpretation of this finding is that, while the local interactions promoted by urban density are crucial in driving the most novel ideas, they play a limited role in generating the majority of inventions. This observation is consistent with the notion that, for the larger part, idea exchange in the innovation process happens within formal networks and organizations, for which in-person interactions can be easily substituted by online communication. The effect of local interactions is only visible for a small set of inventions, although these inventions are significantly more novel and bear disproportionate economic value.

5 Testing for the role of in-person interactions

In this section, we examine possible mechanisms underlying our main findings. We first show that the drop in the density premium in innovativeness can be explained, at least in part, by social distancing during the pandemic which reduced opportunities of in-person idea exchange in dense locations. We then combine our data with information on inventors' demographic characteristics to rule out alternative mechanisms.

Social distancing impacted both formal in-person meetings at the workplace, via the widespread adoption of WFH practices, and informal interactions in other meeting venues, such as restaurants. In principle, both classes of in-person interactions can have an effect on idea exchange and ultimately innovation outcomes (Andrews, 2019; Atkin et al., 2022; Emanuel et al., 2023). In this section, we use Safegraph data in the months surrounding the outbreak of COVID-19 to separately explore these channels. First, we show that in-person interactions (both at the workplace and in informal settings) dropped significantly more in high-density locations relative to low-density ones. Then, we compute location-specific changes in the intensity of both classes of interactions, and show that these changes explain a significant portion of the observed drop in the density premium in innovativeness documented in the previous section.

5.1 Measuring in-person interactions

We use two datasets provided by Safegraph to construct local measures of the intensity of in-person interactions at the workplace and in informal settings.

5.1.1 Measuring in-person interactions at the workplace

Using the Social Distancing Metrics dataset, we compute daily measures of the share of people in a given location who commute to work, and estimate how this share changed with the widespread adoption of WFH during the COVID-19 pandemic. In particular, we construct four alternative measures of the intensity of in-person interactions at the workplace, reflecting different assumptions on how the mobility patterns observed in the data can be informative of the extent of those interactions.

For each CSD and each day, the four measures are defined as (1) the share of cell phone users with a median distance from their residence equal to 0 km, (2) median distance from their residence no greater than 2 km, (3) staying at home full time, and (4) working full time in a location different from their residence.[11] Notice that the first, second, and third measures should increase as WFH becomes more prevalent, while the fourth measure should decrease. Hence, we expect the fourth measure to be negatively correlated with the other three.

In Table A12, Panel A shows the mean and standard errors of the four variables in January to April 2019 and Panel B shows that in 2020. The probability of staying at home is stable in 2019 between January and April, but experiences a sharp increase starting in March 2020, matching the timing of the outbreak of COVID-19 in the U.S. The ratio of devices labeled as completely at home is 36.8% in April, a 56.6% increase compared to 23.5% in January. The ratio of devices labeled as working full-time is 3.4% in April, a 57.5% decrease compared to 8.0% in January. The statistics of all four measures suggest that workers experience sharp declines in in-person meeting opportunities at the

[11] Safegraph categorizes whether the device user is working full time or not by tracing the number of hours spent at a location other than one's home in the daytime. Home is defined as the common nighttime location for the device over a 6-week period where nighttime is between 6pm and 7am. Devices labeled as working full-time are those that spent more than 6 hours at a location other than one's home between 8am and 6pm. For further details, refer to https://docs.safegraph.com/docs/social-distancing-metrics.

workplace.[12]

5.1.2 Measuring in-person interactions in informal settings

We next use the Weekly Pattern dataset to trace the number of visits and visitors to restaurants and other dining places. The data report hourly visits and weekly visitors. The difference between visits and visitors is that a visitor could visit a venue more than once and generate multiple visits. As we intend to capture visits resulting in social interactions, we restrict the attention to venues located in Census block groups within urbanized areas, thereby excluding, for example, venues located along highways. We consider a Census block group as "urban" if the ratio of the population inside urbanized areas, as reported by the NHGIS, accounts for more than 50% of its total population, although our results are robust to alternative definitions of an urban block group (e.g., increasing the cutoff to 90%).[13] Finally, we aggregate visitors and visits from the Census block group level to the CSD level.

We calculate daily visits by aggregating hourly visits, and approximate the number of daily visitors by assuming a constant ratio across seven days within a week based on the reported weekly visitors. Following the same structure as in Table A12, Panels A and B in Table A13 report monthly summary statistics for daily visits and daily visitors from January to April in 2019 and 2020, respectively. Both the daily numbers of visits and the daily number of visitors are stable across four months in 2019, but they decline sharply in March 2020. The summary statistics of the two measures suggest that fewer people are visiting dining places after the onset of COVID-19.

[12]Note that the statistics are not comparable across years because of a methodological difference pointed out by Safegraph in the way data for 2019 and for 2020 are collected. More details can be found at https://docs.safegraph.com/docs/social-distancing-metrics.

[13]For each Census block group, NHGIS reports total population, urban population, the population inside urbanized areas, the population inside urban clusters, and rural population. Urban areas and urban clusters are defined primarily based on residential population density measured at the census tract and census block levels. For detailed information, refer to https://www.census.gov/programs-surveys/geography/about/faq/2010-urban-area-faq.html.

5.2 In-person interactions declined more in dense locations

We now show that, during the COVID-19 pandemic, the frequency of in-person interactions fell significantly more in high-density locations compared to low-density ones. In the coming subsection, we show that this drop in the density gradient of the intensity of local interactions explains a significant portion of the decline of the density premium in innovativeness.

For each of the measures introduced above, we estimate the following specification at a daily frequency for the months of January to April, separately for 2019 and 2020:

$$y_{lt} = \alpha + \sum_{\tau \in T} \beta_\tau \log(Density_l) \times \mathbb{1}(t = \tau) + \eta_l + \theta_t + \varepsilon_{lt}, \tag{7}$$

where y_{lt} represents any of the measures of in-person interactions (at the workplace or in informal settings) for CSD l at date t, and T is a set of dates between January 1$^{\text{st}}$ and April 30$^{\text{th}}$, with February 28$^{\text{th}}$ omitted as the reference date. The terms θ_t and η_l represent date and CSD fixed effects, respectively. The coefficients of interest are β_τ which capture, for each date τ, the density gradient of the frequency of in-person interactions relative to the reference date.

In Figure 3, we plot the estimates of β_τ for the measures of in-person interactions at the workplace. Moving from top to bottom, the panels use estimates using the share of devices with a median travelled distance equal to zero, with a median travelled distance less than 2 km, that are completely at home, and that are full time at work in a location different from the user's residence. The estimates based on the 2019 data are plotted on the left-hand side, and those based on the 2020 data are plotted on the right-hand side. Appendix Figure A5 shows results when weekend days are excluded from the sample. For all four measures, we find that the estimated density gradient of the tendency to stay at home is stable from January to April in 2019. However, in 2020, the gradient increases sharply starting from March (it decreases for the share of devices labeled as full time at work), matching the timing of the widespread lockdowns and the adoption of WFH in response to the outbreak of COVID-19.

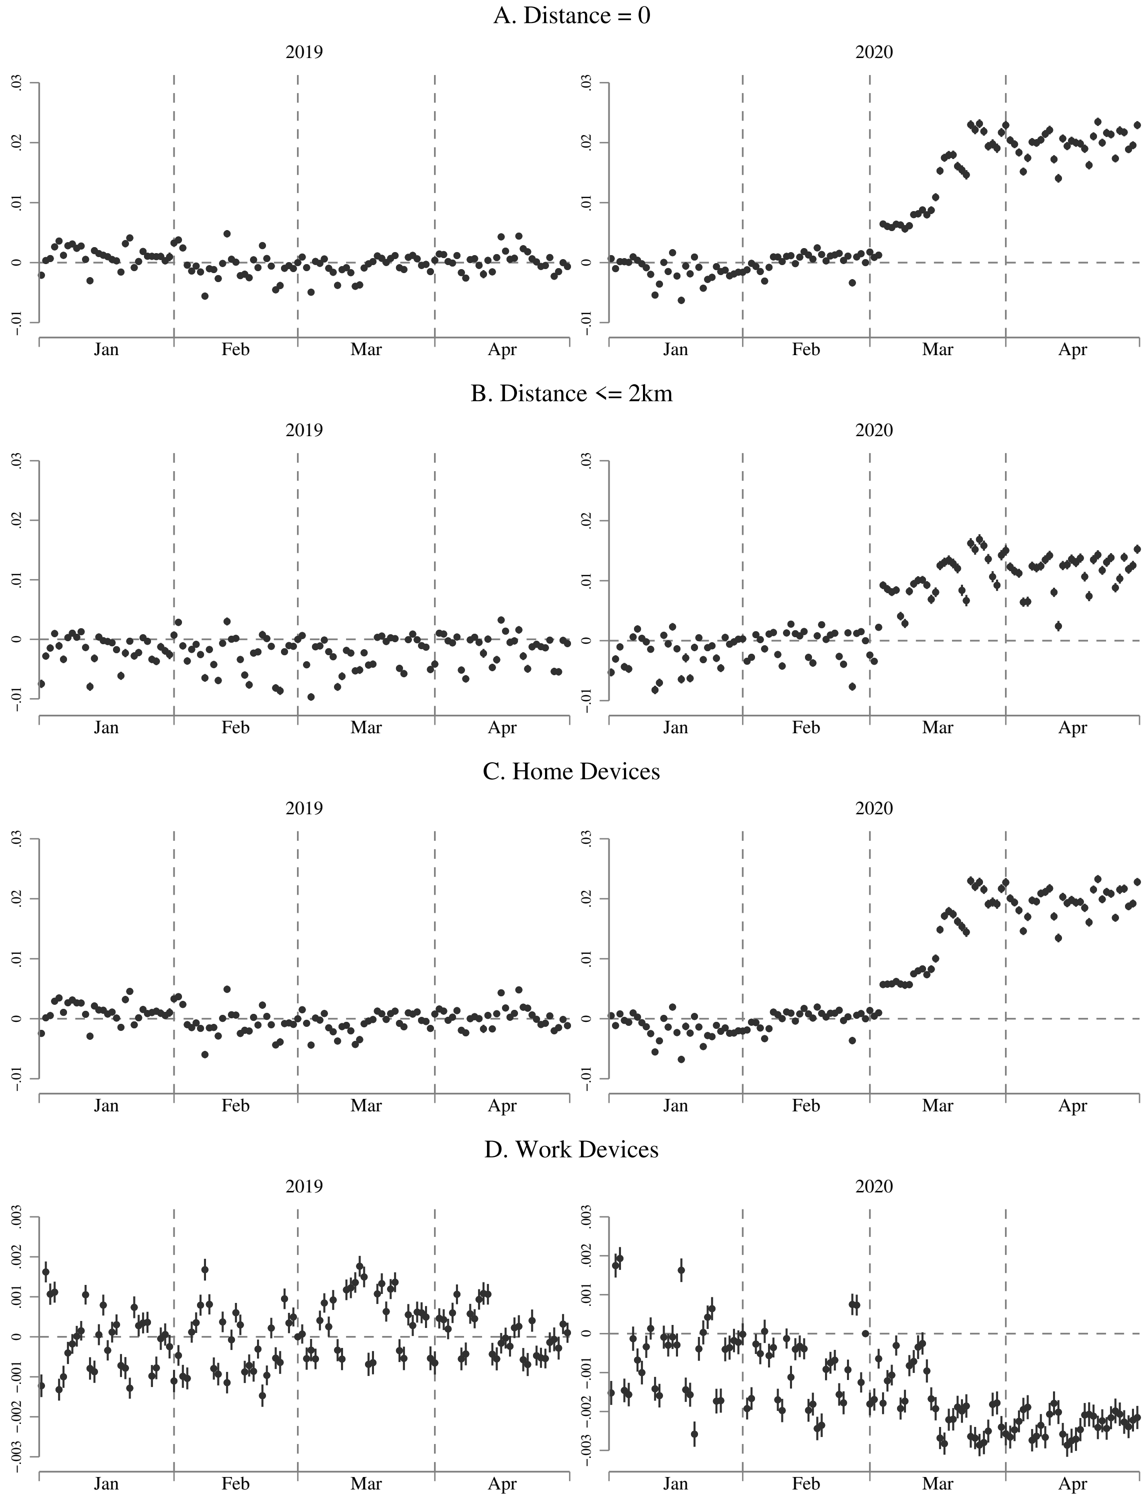

Figure 3: Temporal Changes in Density Gradient of the WFH Intensity

Notes: The empirical specification follows Equation (7). It reports the day-to-day changes in the density gradient of the share of the devices with a median distance traveled equal to 0 in Panel A, or a median distance traveled less than or equal to 2 km in Panel B, that are completely at home in Panel C, and that are full-time at work (a location different from the user's residence) in Panel D. The left-hand side sub-figures present results in 2019 and the right-hand side sub-figures present results in 2020.

In Figure 4, we plot the corresponding estimates of β_τ for the two measures of in-person interactions in informal settings, namely, the logarithm of one plus the number of daily visits (top panels) and visitors (bottom panels) to restaurants and other dining places. The left-hand side panels present results in 2019 and the right-hand side panels present results in 2020. Appendix Figure A6 shows results when weekend days are excluded from the sample. Consistently with the findings on the intensity of interactions at the workplace, the density gradient of the number of visits and visitors to dining venues is stable from January to April 2019 and it decreases sharply starting in March 2020.

Taken together, these results suggest that opportunities for in-person meetings, both at the workplace and in informal environments, fell significantly more in high-density locations relative to low-density ones. The main reason behind this differential response is that the extent to which the widespread restrictions adopted during COVID-19 limited in-person interactions necessarily depends on the initial intensity of those interactions. Dense locations, where the initial frequency of face-to-face meetings was higher, were more severely affected by those restrictions than less dense locations. However, part of this differential response may also be due to the fact that the adoption of WFH practices and the set of mandated restrictions (as well as the extent to which people complied with them) were not uniform across CSDs, being generally stricter in denser locations (Dingel and Neiman, 2020; Hale et al., 2022).[14]

5.3 In-person interactions and the density premium

In this subsection, we provide evidence that the decline of in-person meeting opportunities explains a significant portion of the drop in the density premium in innovativeness during COVID-19. We do this in two steps.

In the first step, we estimate location-specific reductions in in-person interactions as a result of COVID-19. Specifically, we estimate the following specification at a daily

[14] For example, the percentage of jobs that can be performed entirely at home is higher in denser locations (Dingel and Neiman, 2020). Moreover, areas with a higher share of Democratic voters are denser areas that implemented stay-at-home policies earlier (Allcott et al., 2020), and counties with higher skepticism towards science are low-density locations and recorded lower compliance with the restrictions (Brzezinski et al., 2021).

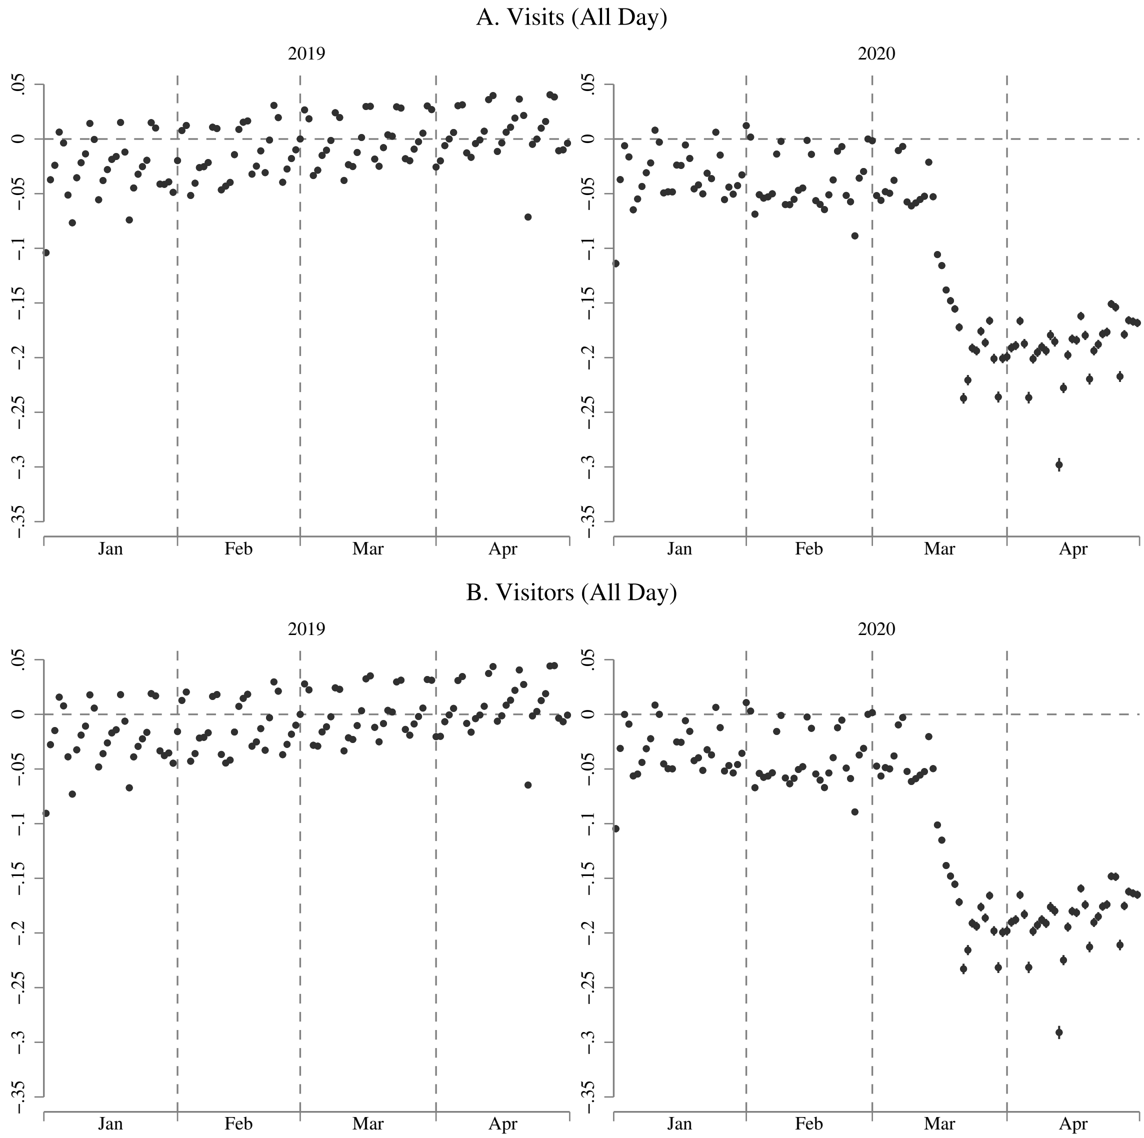

Figure 4: Temporal Changes in Density Gradient of the Intensity of Visits to Restaurants and Other Eating Places

Notes: The empirical specification follows Equation (7). It reports the day-to-day changes in the density gradient of the logarithm of one plus the number of visits to restaurants and other eating places in Panel A and that of visitors to these places in Panel B. The left-hand side sub-figures present results in 2019 and the right-hand side sub-figures present results in 2020.

frequency between January and April 2020:

$$y_{lt} = \alpha + \delta_l \, CSD_l \times \mathbb{1}(After_t) + \eta_l + \theta_t + \varepsilon_{lt}, \tag{8}$$

where CSD_l represents an indicator variable for CSD l, and $\mathbb{1}(After_t)$ is equal to 1 if date t is in March or April of 2020 and equal to 0 if it is in January or February of 2020. We estimate the equation using as outcome variable, y_{lt}, all the measures of in-person interactions at the workplace or in informal settings described above. The location-specific coefficient δ_l captures the magnitude of the drop in face-to-face interactions in CSD l after the onset of COVID-19.

In the second step, we estimate a regression analogous to the baseline specification (Equation 5), in which we include as additional covariate the coefficient $\hat{\delta}_l$ estimated via Equation (8), interacted with the post-COVID-19 dummy variable, $\mathbb{1}(After_t)$:

$$I_{lt} = \alpha + \beta \, \log(Density_l) \times \mathbb{1}(After_t) + \gamma \, \hat{\delta}_l \times \mathbb{1}(After_t) + \eta_l + \theta_t + \varepsilon_{lt}. \tag{9}$$

A drop in the estimate of the coefficient β relative to the baseline would suggest that the decline in in-person interactions was a contributing mechanism behind the drop of the density premium in innovativeness.

Table 2 reports the results for the four measures of in-person interactions at the workplace. Panel A and B show results with the outcome variable defined as the 99$^{\text{th}}$ and 95$^{\text{th}}$ percentile of the novelty distribution, respectively. As a reference, column 1 reports the baseline results as displayed in Table 1. Each location-specific proxy for the underlying mechanism is included separately in columns 2 to 5. Column 6 represents changes when all measures are included as additional covariates. The table reveals two notable patterns.

First, in columns 2, 4, and 5, the coefficients on the interaction between post-COVID-19 and the proxies for staying at home are negative, indicating that the increase in the tendency of staying at home depressed local innovativeness. Consistently, the corresponding coefficient is negative in column 3, where the channel is proxied by the tendency to

Table 2: The Decline in Density Premium of Patent Innovativeness Explained by Changes in the WFH Intensity

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. 99th						
$\mathbb{1}$(after 2020) $\times$						
log(pop density)	-41.12	-27.37	-38.32	-27.54	-36.16	-25.71
	(6.64)	(7.18)	(6.75)	(7.18)	(6.71)	(7.89)
Δ dist $= 0$	-	-937.12	-	-	-	-6429.42
	-	(195.71)	-	-	-	(3563.72)
Δ dist $\leq$ 2km	-	-	-631.83	-	-	350.04
	-	-	(158.04)	-	-	(348.55)
Δ home device	-	-	-	-920.70	-	5470.53
	-	-	-	(195.35)	-	(3526.46)
Δ work device	-	-	-	-	2439.35	1623.58
	-	-	-	-	(468.84)	(536.50)
Adj. R-squared	0.57	0.57	0.57	0.57	0.57	0.57
Obs.	46,041	45,813	45,813	45,813	45,813	45,813
Panel B. 95th						
$\mathbb{1}$(after 2020) $\times$						
log(pop density)	-16.72	-7.60	-14.96	-7.77	-14.02	-5.76
	(5.09)	(5.30)	(5.09)	(5.31)	(5.08)	(5.98)
Δ dist $= 0$	-	-625.49	-	-	-	-6182.05
	-	(150.81)	-	-	-	(2884.27)
Δ dist $\leq$ 2km	-	-	-410.63	-	-	269.81
	-	-	(115.98)	-	-	(246.53)
Δ home device	-	-	-	-610.31	-	5416.71
	-	-	-	(150.61)	-	(2855.89)
Δ work device	-	-	-	-	1379.65	709.04
	-	-	-	-	(352.47)	(390.72)
Adj. R-squared	0.39	0.39	0.39	0.39	0.39	0.39
Obs.	46,041	45,813	45,813	45,813	45,813	45,813

Notes: All regressions follow Equation (9). The dependent variable is the local degree of innovativeness defined as the 99$^{\text{th}}$ percentile value of the local novelty distribution in Panel A and the 95$^{\text{th}}$ percentile value in Panel B. Δ dist $= 0$ represents the CSD-specific change in the share of mobile devices with median distance traveled equal zero. Δ dist $\leq$ 2km represents the CSD-specific change in the share of mobile devices with median distance traveled less than or equal to 2 kilometers. Δ home device represents the CSD-specific change in the share of mobile devices mostly at home. Δ work device represents the CSD-specific change in the share of mobile devices outside the home during work hours. Standard errors clustered at the CSD level are reported in parentheses.

work away from one's own residence.

Second, compared to column 1, the introduction of these controls reduces the estimated coefficient of the interaction between post-COVID-19 and log-density (top-row). This suggests that the main result can be partly explained by location-specific reductions in in-person meeting opportunities. In Panel A, the drop in the main coefficient ranges between 7% of column 3 and 33% of column 2. Results are consistent in Panel B, with the drop in the coefficient ranging between 11% in column 3 and 55% in column 2. In column 6, when all the four proxies are included together, the drop is as large as 37% and 65% in Panels A and B, respectively.

Table 3 reports the corresponding results for the two measures of in-person interactions in informal settings. Again, Panel A and B show results with the outcome variable defined as the 99th and 95th percentile of the novelty distribution, respectively. Column 1 reports the baseline results, columns 2 and 3 introduce the channel as the number of visits and visitors to restaurants, respectively, and column 4 introduces both variables at the same time.

The findings are broadly consistent with the ones obtained using the measures of local interactions at the workplace. In columns 2 and 3, the coefficients on the interaction between post-COVID-19 and the visits or visitors to restaurants are positive, suggesting that the decline in these activities reduced innovativeness after the onset of COVID-19.[15] Moreover, compared to column 1, the magnitudes of the main coefficients become smaller, which indicates that the drop in the density premium is partly explained by location-specific reductions in local interactions in informal settings. The coefficient declines by 32% and 30% in column 4 relative to column 1 in Panels A and B, respectively.

Unsurprisingly, the proxies for the intensity of interactions at the workplace and in informal settings are strongly correlated and are likely to be driven, at least in part, by the same underlying forces (e.g., people working in the office and having lunch in a restaurant). What matters for our interpretation is that, although the two sets of proxies

[15] Visitors and visits are highly correlated, so including them together as covariates in column 4 changes the sign of the estimates. However, the coefficient of $\log(Density_l) \times \mathbb{1}(After_t)$ is stable across columns 2 to 4.

load differently on the two classes of interactions, the ability of these proxies to account for the main result is quantitatively comparable. This suggests that the restrictions to in-person interactions during the pandemic decreased the potential for novel invention in dense locations both by reducing opportunities of idea exchange at the workplace and by making it harder to meet people in person in informal settings.

5.4 Testing for alternative mechanisms

Finally, we use information on inventors' demographic characteristics to test for plausible alternative mechanisms. We briefly present these tests here and leave more details to Appendix C.

We first explore the hypothesis that our results are driven by the fact that female inventors, in particular mothers of young children, were more severely affected by the onset of COVID-19 (Myers et al., 2020). This fact may contribute to explaining our main result if young female inventors are more likely to live in high-density CSDs.

To test this hypothesis, we leverage the information on inventors' gender and the disambiguation of inventors provided by PatentsView. The disambiguation allows us to approximate each inventor's age using their earliest filing year. For each application, we construct indicator variables for whether the application has at least one female inventor, and for whether the application has at least one young inventor, defined alternatively as having their earliest filing in 2005, 2010, or 2015. We then run application-level regressions of the following form:

$$HN_{ilt} = \alpha + \beta \log(Density_l) \times \mathbb{1}(After_t) + \gamma_0 X_i + \gamma_1 X_i \times \mathbb{1}(After_t) + \delta Z_i + \eta_l + \theta_t + \varepsilon_{ilt}, \tag{10}$$

where the dependent variable, HN_{ilt}, is an indicator of high-novelty of application i filed in CSD l at time t, X_i is the indicator for female or for young inventor, and Z_i is a control for team size.

Tables C2 and C3 display the estimates of Equation (10) when the indicator of high-novelty is defined as equal to one if the application's novelty is above the 99$^{\text{th}}$ and 95$^{\text{th}}$

Table 3: The Decline in Density Premium of Patent Innovativeness Explained by Changes in Visits to Restaurants and Other Eating Places

	(1)	(2)	(3)	(4)
Panel A. 99th				
$\mathbb{1}$(after 2020) $\times$				
log(pop density)	-42.02	-27.73	-28.17	-28.34
	(6.89)	(7.93)	(7.94)	(7.95)
Δ visits	-	90.47	-	278.28
	-	(29.21)	-	(270.81)
Δ visitors	-	-	89.59	-195.84
	-	-	(29.80)	(277.36)
Adj. R-squared	0.57	0.57	0.57	0.57
Obs.	45,043	45,043	45,043	45,043
Panel B. 95th				
$\mathbb{1}$(after 2020) $\times$				
log(pop density)	-17.43	-12.58	-12.25	-12.19
	(5.28)	(6.05)	(6.07)	(6.07)
Δ visits	-	30.69	-	-89.92
	-	(22.09)	-	(196.69)
Δ visitors	-	-	33.54	125.78
	-	-	(22.83)	(203.41)
Adj. R-squared	0.39	0.39	0.39	0.39
Obs.	45,043	45,043	45,043	45,043

Notes: All regressions follow Equation (9). The dependent variable is the local degree of innovativeness defined as the 99$^{\text{th}}$ percentile value of the local novelty distribution in Panel A and the 95$^{\text{th}}$ percentile value in Panel B. Δ visits represents the CSD-specific change in the logarithm of one plus the number of visits to restaurants and other eating places. Δ visitors represents the CSD-specific change in the logarithm of one plus the number of visitors to restaurants and other eating places. Standard errors clustered at the CSD level are reported in parentheses.

percentiles, respectively, of all the applications filed in the pre-treatment period (i.e., 2011-2019). The results reveal that applications with at least one female inventor have a higher degree of novelty on average, while the applications with at least one young inventor have a slightly lower degree of novelty (but the difference is not significant when the threshold for high-novelty is defined as the 95$^{\text{th}}$ percentile). The onset of the pandemic had a disproportionately negative effect on the novelty of both categories of applications. However, the inclusion of these controls does not significantly change the estimate of the main coefficient of interest (β), suggesting that the drop in the density premium is unlikely to be explained by the fact that the productivity of women or young inventors with children was more severely affected by the pandemic.

Using the same regression framework, we then examine whether our main finding can be explained by the fact that inventors generating the most novel ideas may have had a higher propensity to relocate from high-density to low-density CSDs after the onset of COVID-19. To this end, we construct for each application an indicator that is equal to one if at least one of the inventors has relocated to a different CSD on or after the year 2020. We then run application-level regressions analogous to Equation (10), in which the term X_i is this indicator for "mover" inventors.

The results are displayed in Table C5. For both measures of high novelty, the inclusion of the control for "mover" inventors does not affect the coefficient on the interaction of log-density and the post-COVID-19 indicator. The stability of the main coefficient of interest to the inclusion of these controls suggests that changing patterns of migration are unlikely to play a significant role in explaining the main result.

6 Conclusion

Opportunities of knowledge exchange via face-to-face interactions are often cited as one of the main reasons why economic activities, and innovation in particular, tend to concentrate in densely populated cities. The COVID-19 pandemic sharply restricted these opportunities of interaction while promoting a growing availability of tools for online communication. Whether these new tools will persistently erode cities' advantage in innovation is an open question that this paper confronted empirically.

We showed that, while the onset of the pandemic did not affect the density premium in patent quantity, it had a large negative effect on the ability of dense locations to generate the most novel ideas, with the drop in face-to-face interactions explaining a sizeable portion of this effect. On the one hand, this implies that dense cities temporarily lost part of their advantages in the innovation process. On the other hand, it suggests that the role that density plays in facilitating knowledge flows and sustaining the creation of novel ideas cannot be effectively replaced by online communication.

These findings bear direct implications for the spatial organization of the workforce within firms in the post-pandemic world. The question of whether the intensity of in-person interactions between workers should revert to pre-pandemic levels has been widely discussed, especially among firms in knowledge-intensive sectors. Elon Musk, who banned remote work at Tesla, SpaceX, and Twitter, has unequivocally expressed his opposition to WFH. Mark Zuckerberg, Meta's CEO, asserted in 2021 that "good work can get done anywhere" while expressing optimism about permanent full-time remote work, "particularly as remote video presence and virtual reality continue to improve." However, he has recently rolled back on the full-time remote work policy, now requiring employees to spend at least part of their time in the office. Our results imply that limits to communication imposed by remote work environments might result in lower innovation potential, thus providing a rationale for advocating work-in-office practices.

In addition to the impact on collaboration and idea exchange within companies, the spontaneous encounters that density facilitates are critical in connecting workers across companies and fields with diverse knowledge backgrounds, thereby spurring the creation of highly novel inventions. The large externalities implied by this process open up room for policy to increase coordination and efficiency. On the one hand, novel inventions are particularly valuable from a welfare perspective, as they are more impactful and have a higher span of technological and geographical diffusion than less novel ones. On the other hand, the creation of novel ideas is itself subject to large externalities, and internalizing the value of the random interactions that give rise to those ideas would require cross-firm and cross-field coordination, which is especially hard to achieve with WFH arrangements. As policy discussions develop on the future of cities after the pandemic, it is critical to take this process into account. Whether cities will ultimately regain their central position in the innovation landscape is an important question that we will explore in future research.

References

Allcott, H., L. Boxell, J. Conway, M. Gentzkow, M. Thaler, and D. Yang (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of public economics 191*, 104254.

Andrews, M. (2019). Bar Talk: Informal Social Interactions, Alcohol Prohibition, and Invention. *Working Paper*.

Arts, S., J. Hou, and J. C. Gomez (2021). Natural language processing to identify the creation and impact of new technologies in patent text: Code, data, and new measures. *Research Policy 5*, 104–144.

Atkin, D., K. M. Chen, and A. Popov (2022). The Returns to Face-to-Face Interactions: Knowledge Spillovers in Silicon Valley. *Working Paper*.

Barber, B. M., W. Jiang, A. Morse, M. Puri, H. Tookes, and I. M. Werner (2021). What Explains Differences in Finance Research Productivity during the Pandemic? *Journal of Finance 76*(4), 1655–1697.

Barrero, J. M., N. Bloom, and S. J. Davis (2020). 60 million fewer commuting hours per day: How americans use time saved by working from home. *Working Paper*.

Berkes, E. and R. Gaetani (2021). The Geography of Unconventional Innovation. *The Economic Journal 131*(636), 1466–1514.

Black, D. and V. Henderson (1999). A Theory of Urban Growth. *Journal of Political Economy 107*(2), 252–284.

Bloom, N., J. Liang, J. Roberts, and Z. J. Ying (2015). Does Working from Home Work? Evidence from a Chinese Experiment. *The Quarterly Journal of Economics 130*(1), 165–218.

Brucks, M. S. and J. Levav (2022). Virtual communication curbs creative idea generation. *Nature 605*, 108–112.

Brzezinski, A., V. Kecht, D. Van Dijcke, and A. L. Wright (2021). Science skepticism reduced compliance with covid-19 shelter-in-place policies in the united states. *Nature Human Behaviour 5*(11), 1519–1527.

Carlino, G. A., S. Chatterjee, and R. M. Hunt (2007). Urban density and the rate of invention. *Journal of Urban Economics 61*(3), 389–419.

Catalini, C. (2018). Microgeography and the direction of inventive activity. *Management Science 64*(9), 4348–4364.

Catalini, C., C. Fons-Rosen, and P. Gaulé (2020). How do travel costs shape collaboration. *Management Science 66*(8), 3340–3360.

Davis, D. R. and J. I. Dingel (2019). A spatial knowledge economy. *American Economic Review 109*(1), 153–170.

Delventhal, M. and A. Parkhomenko (2020). Spatial implications of telecommuting. *Working Paper*.

Dingel, J. I. and B. Neiman (2020). How many jobs can be done at home? *Journal of Public Economics 189*, 104235.

Duranton, G. and D. Puga (2001). Nursery cities: urban diversity, process innovation, and the life cycle of products. *American Economic Review 91*(5), 1454–1477.

Emanuel, N., E. Harrington, and A. Pallais (2023). The power of proximity to coworkers: Training for tomorrow or productivity today? *Working Paper*.

Emanuel, N. and N. Harrington (2021). Working remotely? selection, treatment, and the market provision of remote work. *Working Paper*.

Gertner, J. (2012). *The idea factory: Bell Labs and the great age of American innovation.* Penguin Books.

Gibbs, M., F. Mengel, and C. Siemroth (2023). Work from home and productivity: Evidence from personnel and analytics data on information technology professionals. *Journal of Political Economy Microeconomics 1*(1), 7–41.

Glaeser, E. L., H. D. Kallal, J. A. Scheinkman, and A. Shleifer (1992). Growth in cities. *Journal of Political Economy 100*(6), 1126–1152.

Gupta, A., V. Mittal, and S. Van Nieuwerburgh (2022). Work from home and the office real estate apocalypse. *Working Paper*.

Hale, T., N. Angrist, B. Kira, A. Petherick, T. Phillips, and S. Webster (2022). Variation in government responses to covid-19. *Working Paper*.

Hall, B. H., A. Jaffe, and M. Trajtenberg (2005). Market value and patent citations. *RAND Journal of Economics 36*(1), 16–38.

Jacobs, J. (1969). *The economy of cities*. New York: Vintage.

Kerr, W. R. and F. Robert-Nicoud (2020). Tech clusters. *Journal of Economic Perspectives 34*(3), 50–76.

Koh, Y., J. Li, and J. Xu (2023). Subway, collaborative matching, and innovation. *Review of Economics and Statistics*, Forthcoming.

Kruger, S., G. Maturana, and J. Nickerson (2023). How Has COVID-19 Impacted Research Production in Economics and Finance? *The Review of Financial Studies 36*(8), 3348–3381.

Monte, F., C. Porcher, and E. Rossi-Hansberg (2022). Remote work and city structure. *Working paper*.

Myers, K. R., W. Y. Tham, Y. Yin, N. Cohodes, J. G. Thursby, M. C. Thursby, P. Schiffer, J. T. Walsh, K. R. Lakhani, and D. Wang (2020). Unequal effects of the COVID-19 pandemic on scientists. *Nature Human Behaviour 4*(9), 880–883.

Nathan, M. and H. Overman (2020). Will coronavirus cause a big city exodus? *Environment and Planning B: Urban Analytics and City Science 47*(9), 1537–1542.

Packalen, M. and J. Bhattacharya (2015). Cities and Ideas. *Working Paper*.

Saxenian, A. (1996). *Regional advantage: cultural and competition in Silicon Valley and Route 128*. Cambridge, MA: Harvard University Press.

Scherer, F. M. (1965). Firm size, market structure, opportunity, and the output of patented inventions. *The American Economic Review 55*(5), 1097–1125.

Teodorovicz, T., R. Sadun, A. L. Kun, and O. Shaer (2022). How does working from home during covid-19 affect what managers do? evidence from time-use studies. *Human–Computer Interaction 37*(6), 532–557.

Tommar, S. A., O. Kolokolova, and R. Mura (2022). When Paid Work Gives in to Unpaid Care Work: Evidence from the Hedge Fund Industry under COVID-19. *Management Science 68*(8), 6250–6267.

Trajtenberg, M., R. Henderson, and A. Jaffe (1997). University versus corporate patents: A window on the basicness of invention. *Economics of Innovation and new technology 5*(1), 19–50.

Yang, L., D. Holtz, S. Jaffe, S. Suri, S. Sinha, J. Weston, C. Joyce, N. Shah, K. Sherman, B. Hecht, and J. Teevan (2022). The effects of remote work on collaboration among information workers. *Nature Human Behaviour 6*(1), 43–54.

Online Appendix [For Online Publication]

Contents

A Online Appendix Figures and Tables

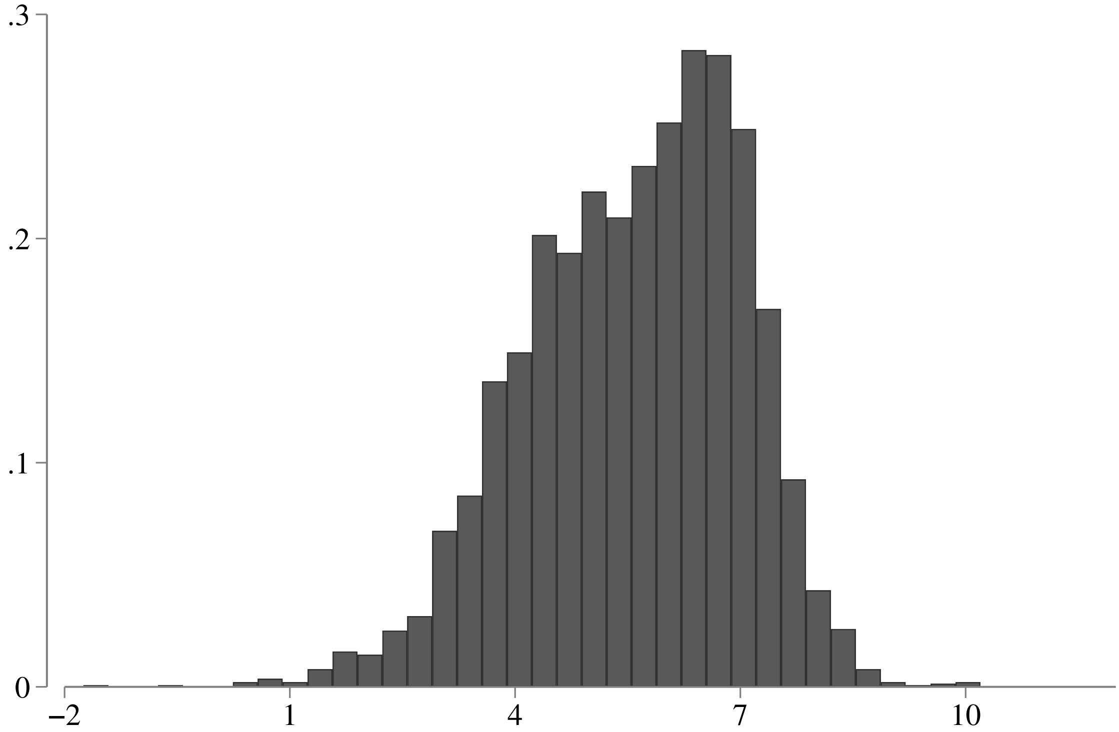

Figure A1: Histogram of CSD-level Log-population Density

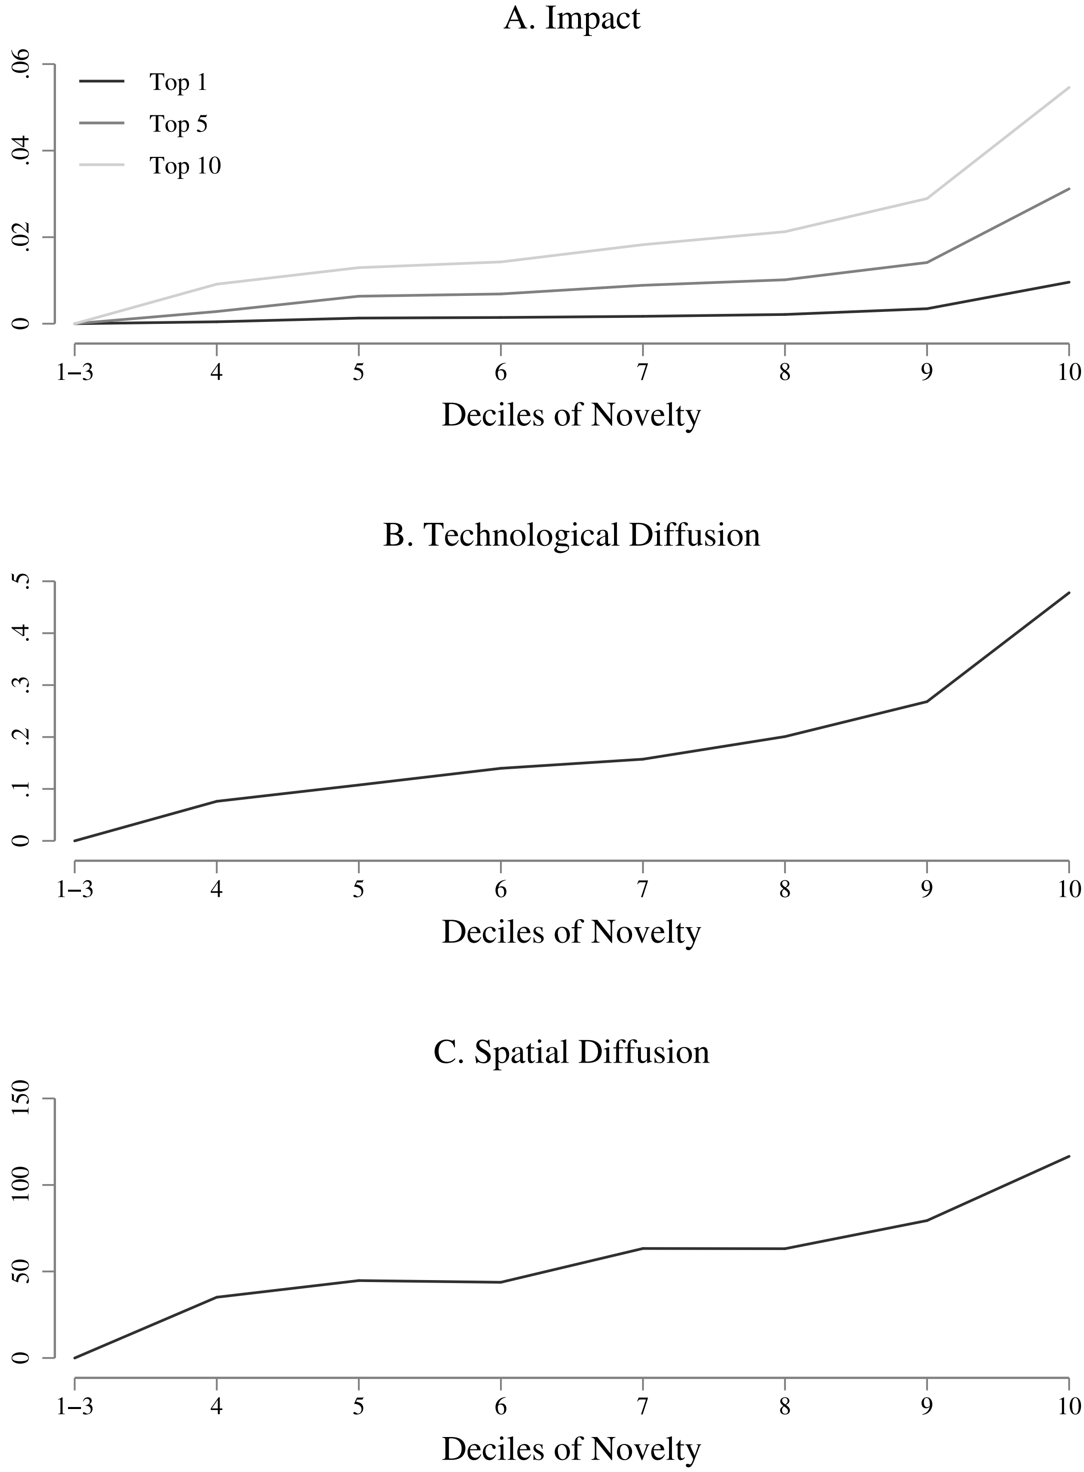

Figure A2: Novelty Predicts Impact and Span of Diffusion

Notes: This figure presents estimated coefficients from Equation (4).

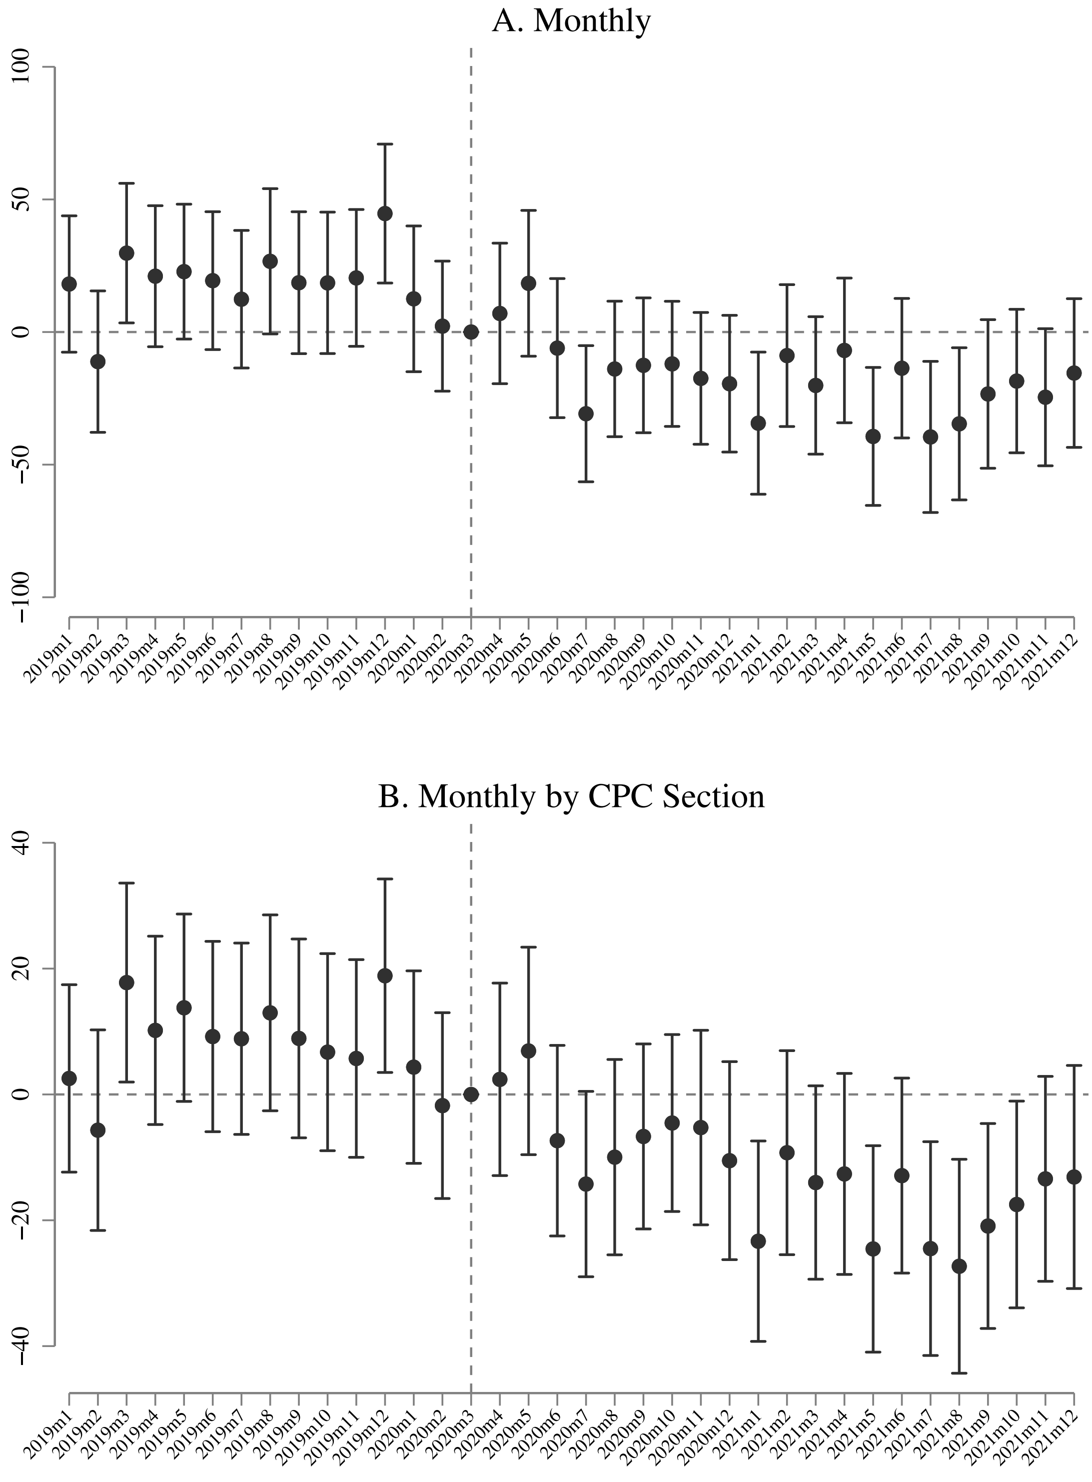

Figure A3: Temporal Changes in Density Premium of Patent Innovativeness (The 99$^{\text{th}}$ Percentile)

Notes: This figure plots the estimated coefficients and the corresponding 95% confidence intervals from estimating Equation (6) at the CSD-month level (Panel A), and the CSD-month-CPC Section level (Panel B). The dependent variable is the local degree of innovativeness measured as the 99$^{\text{th}}$ percentile value of the local novelty distribution.

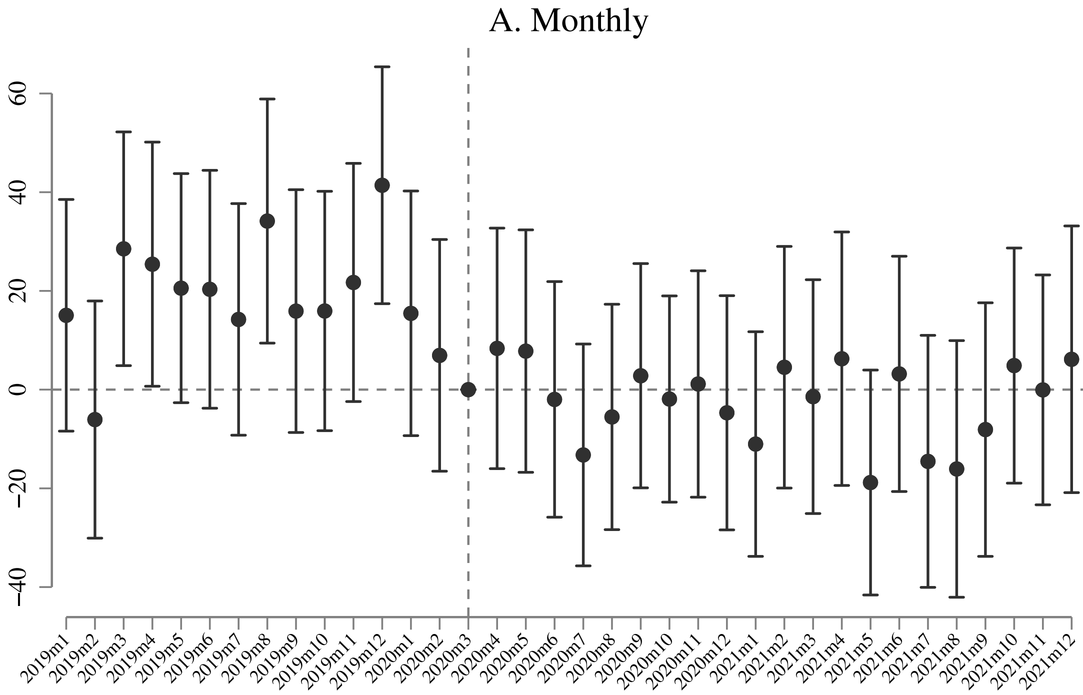

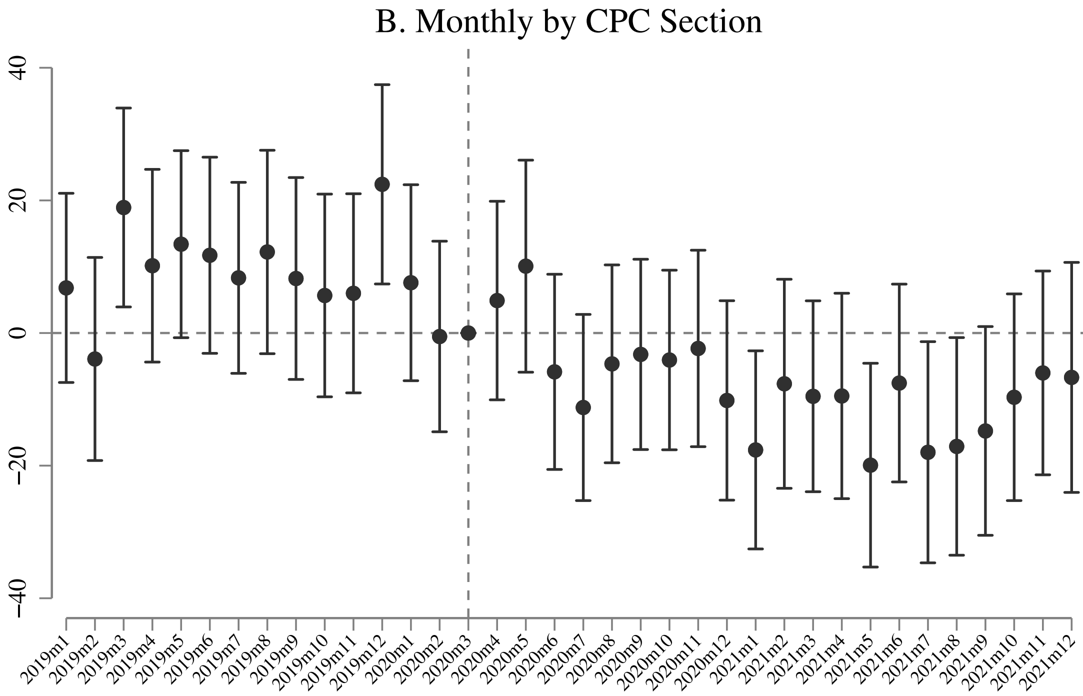

Figure A4: Temporal Changes in Density Premium of Patent Innovativeness (The 95$^{\text{th}}$ Percentile)

Notes: This figure plots the estimated coefficients and the corresponding 95% confidence intervals from estimating Equation (6) at the CSD-month level (Panel A) and the CSD-month-CPC Section level (Panel B). The dependent variable is the local degree of innovativeness measured as the 95$^{\text{th}}$ percentile value of the local novelty distribution.

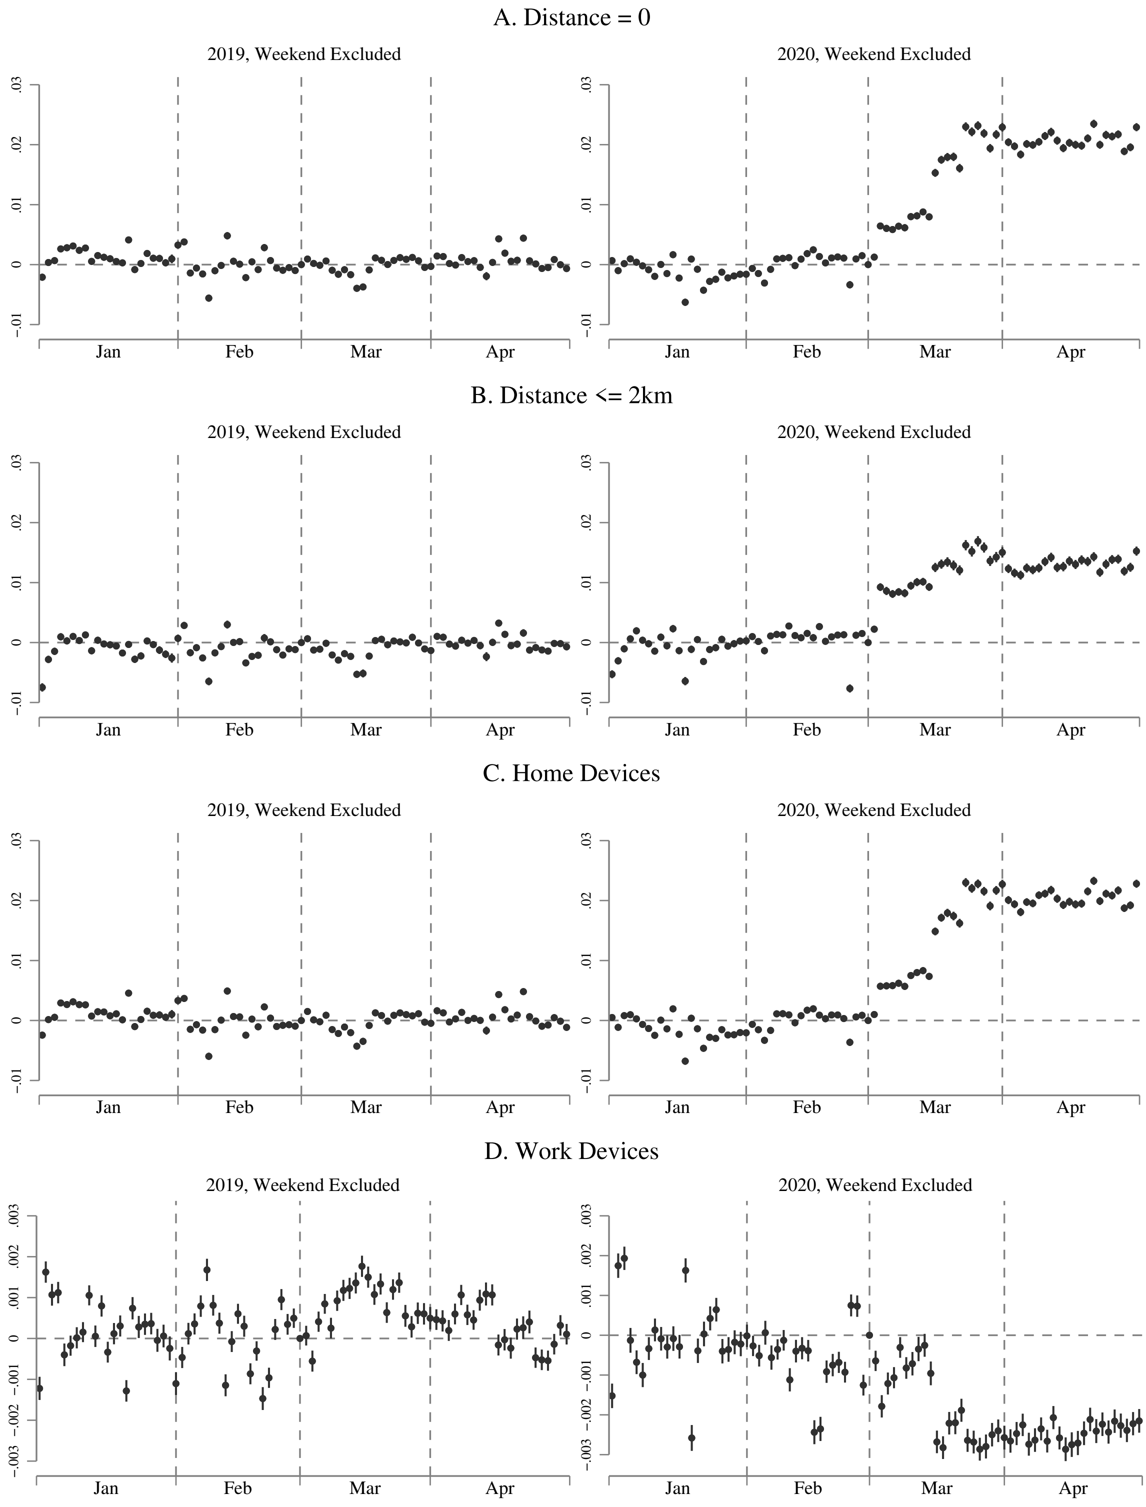

Figure A5: Temporal Changes in Density Gradient of the WFH Intensity (Weekends Removed)

Notes: The empirical specification follows Equation (7). It reports the day-to-day changes in the density gradient of the share of the devices with a median distance traveled equal to 0 in Panel A, or a median distance traveled less than or equal to 2 km in Panel B, that are completely at home in Panel C, and that are full time at work (a location different from the user's residence) in Panel D. The left-hand side sub-figures present results in 2019 and the right-hand side sub-figures present results in 2020. This figure shows results when weekends are excluded.

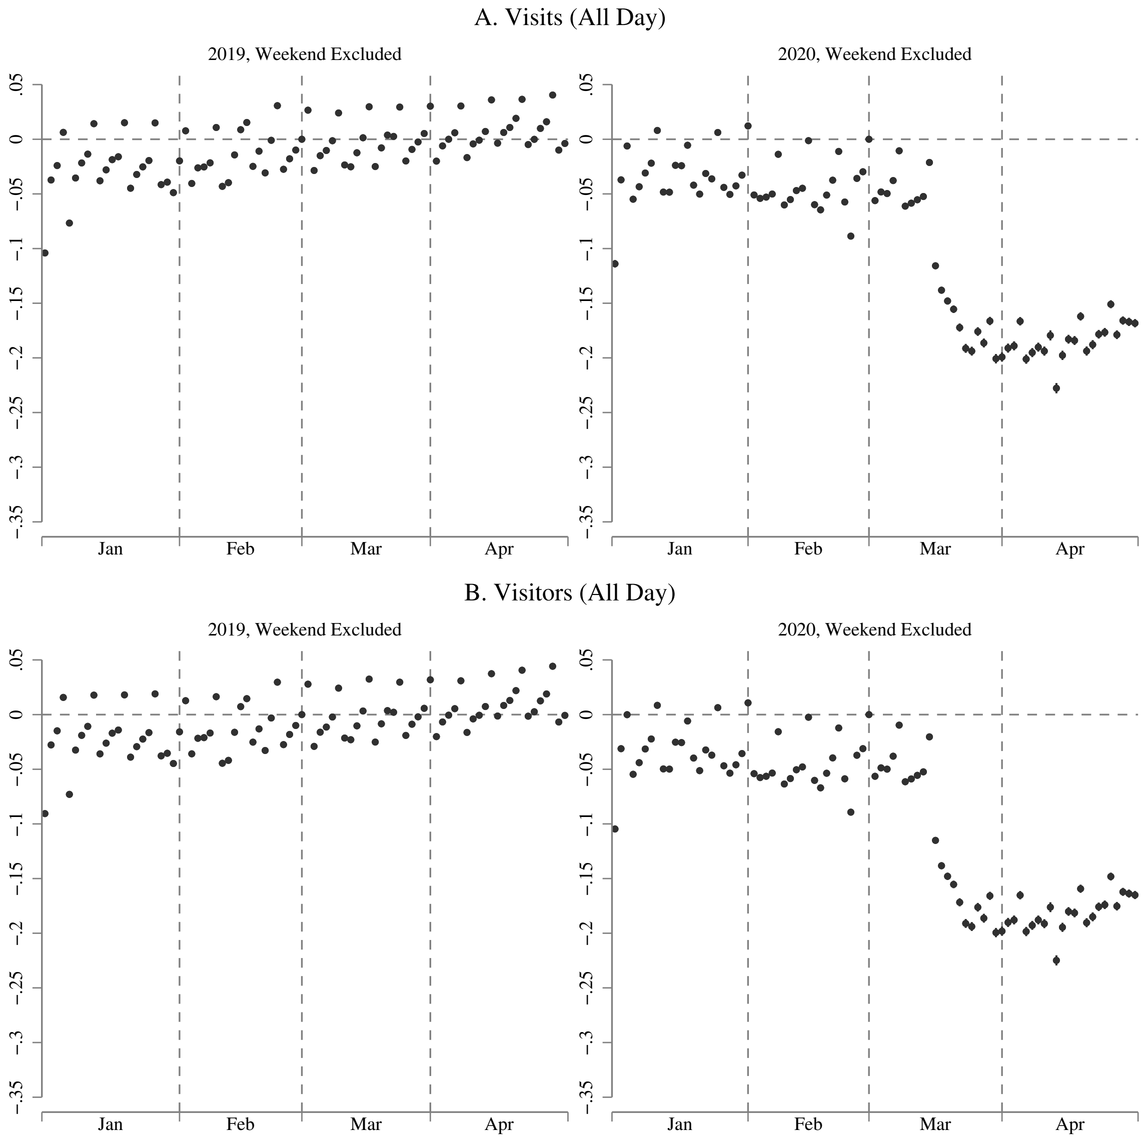

Figure A6: Temporal Changes in Density Gradient of the Intensity of Visits to Restaurants and Other Eating Places (Weekends Removed)

Notes: The empirical specification follows Equation (7). It reports the day-to-day changes in the density gradient of the logarithm of one plus the number of visits to restaurants and other eating places in Panel A and that of visitors to these places in Panel B. The left-hand side sub-figures present results in 2019 and the right-hand side sub-figures present results in 2020. This figure shows results when weekends are excluded.

Table A1: Summary Statistics (County Sub-Division Level)

Variable	Mean	Std. Dev.	Min	Max	Obs.
Density of population ($/km^2$)	618.65	1,070.44	0.17	26,821.9	46,041
Innovation index at the 99$^{\text{th}}$ percentile	872.52	1,181.97	0	3,587	46,041
Innovation index at the 95$^{\text{th}}$ percentile	445.36	742.62	0	3,587	46,041
Annual number of patent applications	70.78	411.44	1	22,261	46,041
Annual weighted number of patent applications	32.69	230.39	0.02	12,755.0	46,041

Notes: This table reports summary statistics at the CSD-year level for the population density, the local degree of innovativeness measured as the 95$^{\text{th}}$ and the 99$^{\text{th}}$ percentile value of the local novelty distribution, the raw counts of patent applications, and the inventor-weighted counts of patent applications.

Table A2: Summary Statistics (Patent Application Level)

Variable	Mean	Std. Dev.	Min	Max	Obs.
$\mathbb{1}$(Top 1)	0.019	0.138	0	1	794,354
$\mathbb{1}$(Top 5)	0.058	0.234	0	1	794,354
$\mathbb{1}$(Top 10)	0.106	0.308	0	1	794,354
Technological diffusion	1.290	1.691	0	43	644,517
Spatial diffusion	1,362.710	1,081.665	0	5,934.4	578,635
Decile of novelty	5.228	3.134	1	10	794,354

Notes: This table reports summary statistics of indicators for whether the patent's number of citations received falls in the top 1%, top 5%, or top 10% among patents filed in the same year and belonging to the same technology class, a measure of the span of technological diffusion of a patent as the number of different technology classes (other than its own) from which the patent receives citations, a measure of the extent of spatial diffusion as the mean of pairwise geographic distances between cities of a patent's inventors and those of all the patents citing this patent, a set of indicators of the decile of a patent's novelty among patents filed in the same year and belonging to the same technology class.

Table A3: Novelty Predicts Impact and Span of Diffusion

	(1) $\mathbb{1}$(Top 1)	(2) $\mathbb{1}$(Top 5)	(3) $\mathbb{1}$(Top 10)	(4) Technological Diffusion	(5) Spatial Diffusion
4^{th} decile	0.000	0.003*	0.009***	0.076	35.120
	(0.001)	(0.001)	(0.002)	(0.009)	(5.504)
$\mathbb{1}(5^{\text{th}}$ decile)	0.001	0.006	0.013	0.108	44.757
	(0.001)	(0.002)	(0.002)	(0.010)	(6.497)
$\mathbb{1}(6^{\text{th}}$ decile)	0.001	0.007	0.014	0.140	43.780
	(0.001)	(0.002)	(0.003)	(0.011)	(6.483)
$\mathbb{1}(7^{\text{th}}$ decile)	0.002	0.009	0.018	0.157	63.310
	(0.002)	(0.003)	(0.004)	(0.011)	(6.659)
$\mathbb{1}(8^{\text{th}}$ decile)	0.002	0.010	0.021	0.201	63.191
	(0.002)	(0.003)	(0.005)	(0.012)	(6.912)
$\mathbb{1}(9^{\text{th}}$decile)	0.003	0.014	0.029	0.268	79.419
	(0.002)	(0.004)	(0.006)	(0.015)	(8.068)
$\mathbb{1}(10^{\text{th}}$ decile)	0.010	0.031	0.055	0.478	116.536
	(0.003)	(0.006)	(0.009)	(0.029)	(11.400)
Class-year FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.001	0.002	0.003	0.075	0.058
Obs.	794,349	794,349	794,349	644,509	578,627

Notes: This table reports estimated coefficients from Equation (4). Standard errors clustered at the CPC technology class by filing year level are reported in parentheses.

Table A4: The Impact of COVID-19 on Density Premium of Patent Innovativeness (Month Level)

	(1) 99^{th}	(2) 99^{th}	(3) 95^{th}	(4) 95^{th}
Panel A. Month				
log(pop density)	141.03	-	83.94	-
	(7.96)	-	(4.79)	-
$\mathbb{1}$(after 2020m3) $\times$ log(pop density)	-27.87	-35.01	-17.68	-20.40
	(3.00)	(3.05)	(2.75)	(2.76)
CSD FE	No	Yes	No	Yes
Month FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.04	0.39	0.03	0.26
Obs.	352,550	352,550	352,550	352,550
Panel B. Month by CPC Section				
log(pop density)	71.58	-	57.75	-
	(4.26)	-	(3.54)	-
$\mathbb{1}$(after 2020m3) $\times$ log(pop density)	-10.56	-15.33	-10.03	-12.95
	(1.64)	(1.64)	(1.52)	(1.52)
CSD by CPC Section FE	No	Yes	No	Yes
Month FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.02	0.37	0.01	0.33
Obs.	868,763	866,656	868,763	866,656

Notes: All regressions follow Equation (5). The dependent variable is the local degree of innovativeness measured as the 99^{th} percentile value of the local novelty distribution (Columns (1) and (2)) or the 95^{th} percentile value of the local novelty distribution (Columns (3) and (4)). The data in Panel A is at the month level and that in Panel B is at the month-CPC Section level. Standard errors clustered at the CSD level are reported in Panel A. Standard errors clustered at the CSD by CPC Section level are reported in Panel B.

Table A5: The Impact of COVID-19 on Density Premium of Patent Innovativeness (Controlling for Quantity)

	(1) 99^{th}	(2) 99^{th}	(3) 95^{th}	(4) 95^{th}
Panel A. Linear				
log(pop density)	209.89	-	72.23	-
	(10.64)	-	(5.41)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-39.13	-41.02	-17.21	-16.77
	(6.69)	(6.65)	(5.05)	(5.10)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.08	0.57	0.02	0.39
Obs.	46,041	46,041	46,041	46,041
Panel B. Quadratic				
log(pop density)	189.49	-	69.78	-
	(9.24)	-	(5.39)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-36.82	-40.89	-16.93	-16.77
	(6.70)	(6.65)	(5.05)	(5.10)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.10	0.57	0.02	0.39
Obs.	46,041	46,041	46,041	46,041

Notes: All regressions follow Equation (5). The dependent variable is the local degree of innovativeness measured as the 99^{th} percentile value of the local novelty distribution (Columns (1) and (2)) or the 95^{th} percentile value of the local novelty distribution (Columns (3) and (4)). In Panels A and B we control for a linear and for a quadratic function of patent applications in each CSD in a year, respectively. Standard errors clustered at the CSD level are reported in parentheses.

Table A6: The Impact of COVID-19 on Density Premium of Patent Innovativeness (Controlling for State Level Economic Conditions)

	(1) 99^{th}	(2) 99^{th}	(3) 95^{th}	(4) 95^{th}
Panel A. Year (Apr)				
log(pop density)	220.23	-	72.29	-
	(9.18)	-	(5.37)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-39.06	-41.72	-16.62	-17.03
	(6.76)	(6.67)	(5.12)	(5.11)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.07	0.57	0.02	0.39
Obs.	46,041	46,041	46,041	46,041
Panel B. Year (Dec)				
log(pop density)	217.46	-	70.86	-
	(9.19)	-	(5.37)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-40.53	-41.56	-17.54	-17.26
	(6.77)	(6.66)	(5.12)	(5.10)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.08	0.57	0.02	0.39
Obs.	46,041	46,041	46,041	46,041

Notes: All regressions follow Equation (5). The dependent variable is the local degree of innovativeness measured as the 99^{th} percentile value of the local novelty distribution (Columns (1) and (2)) or the 95^{th} percentile value of the local novelty distribution (Columns (3) and (4)). All columns control for the state-year level cumulative number of recorded infections per 100 thousand people and unemployment rate. Panels A and B use the unemployment rate in April and December of each year, respectively. Standard errors clustered at the CSD level are reported in parentheses.

Table A7: The Impact of COVID-19 on Density Premium of Patent Counts

	(1) Count	(2) Count	(3) Weighted Count	(4) Weighted Count
Panel A. Year				
log(pop density)	40.03	-	19.37	-
	(6.02)	-	(3.33)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	0.10	0.10	-0.41	-0.41
	(0.63)	(0.63)	(0.26)	(0.26)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.02	0.99	0.01	0.99
Obs.	42,010	42,010	42,010	42,010
Panel B. Month				
log(pop density)	3.34	-	1.61	-
	(0.50)	-	(0.28)	-
$\mathbb{1}$(after 2020m3) $\times$ log(pop density)	0.01	0.01	-0.03	-0.03
	(0.06)	(0.06)	(0.02)	(0.02)
CSD FE	No	Yes	No	No
Month FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.02	0.95	0.01	0.96
Obs.	504,120	504,120	504,120	504,120
Panel C. Month by CPC Section				
log(pop density)	0.47	-	0.23	-
	(0.04)	-	(0.02)	-
$\mathbb{1}$(after 2020m3) $\times$ log(pop density)	0.00	0.00	-0.00	-0.00
	(0.01)	(0.01)	(0.00)	(0.00)
CSD by CPC Section FE	No	Yes	No	Yes
Month FE	Yes	Yes	Yes	Yes
Adj. R-squared	0.01	0.93	0.01	0.94
Obs.	3,483,240	3,483,240	3,483,240	3,483,240

Notes: All regressions follow Equation (5) except that the dependent variable is the local patent count. Columns (1) and (2) use raw counts of patent applications, and Columns (3) and (4) use inventor-weighted counts of patent applications. The observations in Panels A - C are at the CSD-year level, the CSD-month level, and the CSD-month-CPC Section level, respectively. Standard errors clustered at the CSD level are reported in parentheses in Panels A and B. Standard errors clustered at the CSD by CPC Section level are reported in Panel C.

Table A8: The Impact of COVID-19 on Density Premium of Different Types of Patent Applications (Cutoff at the 95$^{\text{th}}$ Percentile)

	(1) Count	(2) Count	(3) Weighted Count	(4) Weighted Count
Panel A. $\geq$ 95th percentile novelty				
log(pop density)	2.69	-	1.08	-
	(0.33)	-	(0.15)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-0.60	-0.60	-0.19	-0.19
	(0.06)	(0.06)	(0.03)	(0.03)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Mean of DV	4.19	4.19	1.60	1.60
Adj. R-squared	0.03	0.97	0.03	0.96
Obs.	42,010	42,010	42,010	42,010
Panel B. < 95th percentile novelty				
log(pop density)	37.34	-	18.29	-
	(5.75)	-	(3.21)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	0.70	0.70	-0.22	-0.22
	(0.65)	(0.65)	(0.27)	(0.27)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Mean of DV	68.49	68.49	32.02	32.02
Adj. R-squared	0.02	0.99	0.01	0.99
Obs.	42,010	42,010	42,010	42,010

Notes: Panel A (B) shows the coefficients of estimating Equation (5) except that the dependent variable is the count of patents with novelty measure above the 95$^{\text{th}}$ percentile cutoff value based on the novelty distribution before 2020. Columns (1) and (2) use raw counts of patent applications, and Columns (3) and (4) use inventor-weighted counts of patent applications. Standard errors clustered at the CSD level are reported in parentheses.

Table A9: The Impact of COVID-19 on Density Premium of Different Types of Patent Applications (Cutoff at the 99$^{\text{th}}$ Percentile)

	(1) Count	(2) Count	(3) Weighted Count	(4) Weighted Count
Panel A. $\geq$ *99th percentile novelty*				
log(pop density)	0.64	-	0.22	-
	(0.08)	-	(0.03)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-0.17	-0.17	-0.05	-0.05
	(0.03)	(0.03)	(0.01)	(0.01)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Mean of DV	0.98	0.98	0.31	0.31
Adj. R-squared	0.03	0.91	0.02	0.90
Obs.	42,010	42,010	42,010	42,010
Panel B. $<$ *99th percentile novelty*				
log(pop density)	39.38	-	19.15	-
	(5.96)	-	(3.30)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	0.27	0.27	-0.36	-0.36
	(0.63)	(0.63)	(0.27)	(0.27)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Mean of DV	71.70	71.70	33.30	33.30
Adj. R-squared	0.02	0.99	0.01	0.99
Obs.	42,010	42,010	42,010	42,010

Notes: Panel A (B) shows coefficients of estimating Equation (5) except that the dependent variable is the count of patents with novelty measure above the 99$^{\text{th}}$ percentile cutoff value based on the novelty distribution before 2020. Columns (1) and (2) use raw counts of patent applications, and Columns (3) and (4) use inventor-weighted counts of patent applications. Standard errors clustered at the CSD level are reported in parentheses.

Table A10: The Impact of COVID-19 on Density Premium of Different Types of Patent Applications (Cutoff at the 90$^{\text{th}}$ Percentile)

	(1) Count	(2) Count	(3) Weighted Count	(4) Weighted Count
Panel A. $\geq$ *90th percentile novelty*				
log(pop density)	4.92	-	2.10	-
	(0.60)	-	(0.28)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-0.87	-0.87	-0.31	-0.31
	(0.08)	(0.08)	(0.04)	(0.04)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Mean of DV	7.93	7.93	3.25	3.25
Adj. R-squared	0.03	0.98	0.02	0.98
Obs.	42,010	42,010	42,010	42,010
Panel B. $<$ *90th percentile novelty*				
log(pop density)	35.11	-	17.27	-
	(5.50)	-	(3.08)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	0.97	0.97	-0.09	-0.09
	(0.63)	(0.63)	(0.26)	(0.26)
CSD FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Mean of DV	64.75	64.75	30.37	30.37
Adj. R-squared	0.02	0.99	0.01	0.99
Obs.	42,010	42,010	42,010	42,010

Notes: Panel A (B) shows the coefficients of estimating Equation (5) except that the dependent variable is the count of patents above the 90$^{\text{th}}$ percentile cutoff value based on the novelty distribution before 2020. Columns (1) and (2) use raw counts of patent applications, and Columns (3) and (4) use inventor-weighted counts of patent applications. Standard errors clustered at the CSD level are reported in parentheses.

Table A11: The Impact of COVID-19 on Density Premium of Novelty at Other Percentiles

	(1) 90^{th}	(2) 90^{th}	(3) 75^{th}	(4) 75^{th}	(5) 50^{th}	(6) 50^{th}
log(pop density)	22.89	-	-1.83	-	-3.16	-
	(3.28)	-	(1.44)	-	(0.58)	-
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-1.41	-0.90	0.29	0.44	-1.34	-1.33
	(3.85)	(3.89)	(2.15)	(2.22)	(1.28)	(1.34)
CSD FE	No	Yes	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	0.01	0.29	0.00	0.19	0.00	0.12
Obs.	46,041	46,041	46,041	46,041	46,041	46,041

Notes: All regressions follow Equation (5). The dependent variable is the local degree of innovativeness measured as the 90^{th} percentile value of the local novelty distribution (Columns (1) and (2)), the 75^{th} percentile value of the local novelty distribution (Columns (3) and (4)), or the 50^{th} percentile value of the local novelty distribution (Columns (5) and (6)). Standard errors clustered at the CSD level are reported in parentheses.

Table A12: Summary Statistics for Safegraph Data (Devices)

Variable				
Panel A. 2019	*Jan*	*Feb*	*Mar*	*Apr*
Ratio of devices...				
dist = 0	0.36	0.34	0.31	0.30
	(0.09)	(0.09)	(0.09)	(0.08)
dist $\leq$ 2km	0.47	0.46	0.43	0.42
	(0.11)	(0.11)	(0.11)	(0.10)
home device	0.36	0.34	0.31	0.30
	(0.09)	(0.09)	(0.09)	(0.08)
work device	0.06	0.06	0.05	0.06
	(0.04)	(0.04)	(0.04)	(0.04)
Obs.	817,475	738,471	817,696	791,309
Panel B. 2020	*Jan*	*Feb*	*Mar*	*Apr*
Ratio of devices...				
dist = 0	0.23	0.25	0.28	0.37
	(0.08)	(0.08)	(0.11)	(0.10)
dist $\leq$ 2km	0.37	0.39	0.43	0.54
	(0.11)	(0.10)	(0.13)	(0.11)
home device	0.23	0.25	0.28	0.37
	(0.08)	(0.08)	(0.11)	(0.10)
work device	0.08	0.07	0.05	0.03
	(0.05)	(0.05)	(0.04)	(0.02)
Obs.	817,444	764,970	817,774	791,246

Notes: The means and standard errors (in parentheses) of the variables are shown in this table for 2019 in Panel A and 2020 in Panel B. The variables are the share of devices with a median distance traveled equal to 0 km, the share of devices with a median distance traveled less than or equal to two kilometers, the share of devices labeled as completely at home, and the share of devices labeled as full-time working.

Table A13: Summary Statistics for Safegraph Data (Visit and Visitors to Restaurants and Other Eating Places)

Variable				
Panel A. 2019	*Jan*	*Feb*	*Mar*	*Apr*
$log(1 +$ daily number of visits)	1.48	1.51	1.53	1.54
	(2.54)	(2.59)	(2.62)	(2.64)
$log(1 +$ daily number of visitors)	1.41	1.44	1.46	1.47
	(2.45)	(2.48)	(2.52)	(2.53)
Obs.	592,695	535,678	594,720	577,590
Panel B. 2020	*Jan*	*Feb*	*Mar*	*Apr*
$log(1 +$ daily number of visits)	1.57	1.56	1.43	1.31
	(2.67)	(2.66)	(2.47)	(2.28)
$log(1 +$ daily number of visitors)	1.51	1.49	1.36	1.25
	(2.58)	(2.56)	(2.37)	(2.20)
Obs.	596,043	557,737	592,265	569,388

Notes: The means and standard errors (in parentheses) of the variables are shown in this table for 2019 in Panel A and 2020 in Panel B. The variables are logarithm of one plus the daily number of visits and visitors to restaurants and other eating places.

B Matching Inventors' Cities with County Sub-Divisions (CSDs)

This section describes how we match inventors' cities listed in the patent applications to the corresponding CSDs in 2010. We start by georeferencing cities based on their GNIS (Geographic Names Information System) coordinates.[16] GNIS provides a comprehensive list of city names and their corresponding coordinates representing "fixed positions across time, located within the historical, functional center" of these cities.[17] In the analysis, we only focus on geographic places in GNIS whose feature class belongs to populated places. According to GNIS, a populated place "represents a named community with a permanent human population, usually not incorporated and with no legal boundaries, ranging from rural clustered buildings to large cities and every size in between."[18] City names belonging to populated places are preferred as they usually do not include generic terms such as "Town of" or "City of", which make it difficult to derive similarity between two strings.[19]

Each city available in the patent application dataset is assigned the coordinate of the best-matched city among all city names provided by the GNIS within the same state. This is achieved using a fuzzy string-matching algorithm. To be specific, the algorithm is conducted using the Python package *rapidfuzz*, built upon *thefuzz* (originally known as *fuzzywuzzy*). The specific function used is the *extractOne*. For each city name in the patent dataset, this function returns the best-matched geographic name which is the one associated with the highest matching score among all geographic names provided by the GNIS within the same state. The matching score ranges from 0 to 100, where 0 means that two names are not similar at all, and 100 means that two names are identical. The default scorer is *fuzz.Wratio*. It is the weighted ratio built upon *fuzz.ratio* which calculates the Levenshtein distance between two strings.

In the main analysis, we only include cities that are exactly matched (with a matching score of 100) and cities matched to only one place. Some cities can be matched to multiple geographic places in GNIS sharing the same name. For example, Van T. Walworth

[16]https://www.usgs.gov/u.s.-board-on-geographic-names/download-gnis-data.
[17]https://www.nhgis.org/documentation/gis-data/place-points.
[18]https://www.usgs.gov/us-board-on-geographic-names/how-do-i.
[19]For example, the civil class record is "City of Denver" and the populated place is Denver.

from patent application No. 08/919050 lives in Lebanon, TN, U.S. However, there are four cities named Lebanon in TN according to the GNIS. They are in different counties (Bradley, Hardin, Haywood, Wilson). Therefore, it is not possible to assign the city to any of these places without extra information. By removing inventors' cities that cannot be matched to geographic places in GNIS, we lose about 5.4% of patent applications. Despite this drawback, we prefer using this algorithm to geolocating using Google maps as our method reduces the concern that small cities are likely to be ignored due to the pagerank algorithm of the Google search engine.

After assigning the GNIS coordinates to matched cities in the patent application dataset, we then determine their corresponding CSDs by spatial joining these coordinates with the map of CSDs in 2010. In this way, cities are linked to a unique CSD which the cities' associated coordinates fall in on the 2010 map.

C Additional Analysis Accounting for Inventor Heterogeneity

In this section, we provide more details on the tests for alternative mechanisms summarized in Section 5.4.

The literature has shown that COVID-19 has a disproportionately negative impact on women or parents with young children. For example, while some researchers reported an increase in time devoted to work, female researchers and researchers with young children reported a large decline in research time during COVID-19 (Myers et al., 2020).[20]

The potential diverging impact of COVID-19 on female inventors or inventors with young children may potentially confound our baseline findings if those inventors more negatively affected by COVID-19 are also more innovative and more likely to reside in dense cities. In addition, it is also possible that inventors producing the most novel inventions may have had a higher propensity to relocate from high-density to low-density CSDs after the onset of COVID-19, thereby confounding our baseline results.

To address these concerns, we use information on inventors' demographic characteristics and examine the impact of COVID-19 on the density premium of patent innovativeness while controlling for inventors' exposure to COVID-19 based on their characteristics. We find that those alternative channels are unlikely to play a significant role in driving our baseline results.

C.1 Data and Variables

The analysis relies on the patent application dataset provided by PatentsView.[21] Compared to the patent application files released by the USPTO, PatentsView provides additional information on the inventor's gender and inventor identity. The inventor's identity allows us to track an inventor over time. In this way, we could approximate an inventor's age based on the earliest filing year of patent applications filed by this inventor. We

[20]Similar findings are provided by Barber et al. (2021) who focus on faculty in Finance. Kruger et al. (2023) find that research production in economics and finance (measured by working paper postings) rose significantly following the onset of COVID-19, but women between the age of 35 and 49 experienced no increase. Tommar et al. (2022) show that female hedge fund managers' ability to generate abnormal returns is curbed under COVID-19.

[21]The data is available from https://patentsview.org/download/pg-download-tables, accessed on December 24, 2022.

also track whether an inventor has moved to different CSDs after the year 2020. The inventor-level information allows us to determine whether a patent application's inventor pool includes at least one female inventor, at least one young inventor, or at least one inventor who has relocated.

We construct two novelty indicators at the patent application level based on the text-based novelty measure that we discussed in Section 3. We first link patents from PatentsView and those from the USPTO based on a unique patent application identity to merge the novelty measure with the PatentsView data. We then create two dummy variables indicating whether a patent application is considered as the top 1% or the top 5% in terms of its novelty, respectively. The cutoff value is taken as the values at the 99^{th} and the 95^{th} percentiles of the count of the new word pairs for all applications in our sample before COVID-19 (2011 – 2019).

We report the summary statistics for 1,739,130 patent applications filed from 2011 to 2021 in Table C.2. Among all patent applications, some do not have information on gender. For the remaining 1,635,125 patent applications with information on inventors' gender, 23.6% have at least one female inventor. Regarding young inventors, we explore three definitions: Those whose earliest filling years are after the year of 2005, 2010, or 2015. For example, 54.3% of patent applications have at least one young inventor whose earliest filing year is after 2010. Based on the way in which young inventors are defined, the number of patent applications with young inventors becomes larger mechanically when the cutoff year is set to 2005 compared to 2015. 33.1% of all patent applications have at least one inventor who moved to a different CSD after 2020.

For the application-level novelty indicators, there are 1.0% "top 1% patents" and 4.9% "top 5% patents" among patent applications filed between 2011 and 2021, closely following the definitions with cutoff values based on the pool of applications before 2020. The average team size is about 3 inventors per patent application.

C.2 Empirical Results

We conduct the analysis at the patent application level based on the following specification:

$$HN_{ilt} = \alpha + \beta\, log(Density_l) \times \mathbb{1}(After_t) + \gamma_0\, X_i + \gamma_1\, X_i \times \mathbb{1}(After_t) + \delta\, Z_i + \eta_l + \theta_t + \varepsilon_{ilt}, \tag{11}$$

where the dependent variable, HN_{ilt}, is an indicator of high-novelty of patent application i filed in CSD l at time t, X_i is the indicator for female or young inventor, and Z_i is a control for team size. Standard errors are clustered at the CSD level.

Tables C2 and C3 report findings when dependent variables indicate patents with high novelty, defined as being in the top 1% and top 5%, respectively, of the novelty distribution. Column (1) in Table C2 shows that post-COVID-19, the probability of a patent application achieving top 1% novelty status decreases by 0.001. This decrease corresponds to a 10% reduction based on the mean value of 0.01. Similarly, Column (1) in Table C3 shows that after COVID-19, the probability of a patent application attaining top 5% novelty status decreases by 0.004 or 8.16%, compared to the mean value of 0.049.

Column (2) in both Table C2 and Table C3 reveals that in comparison to patent applications with exclusively male inventors, patent applications with at least one female inventor exhibit higher levels of novelty. However, following the outbreak of COVID-19, there is a decreased likelihood that a patent application with at least one female inventor falls within the top 1% of the novelty. This trend aligns with the existing literature that documents the disproportionate adverse impact of COVID-19 on females.

Next, in addition to controlling for a binary variable indicating gender, we also account for whether a patent application is filed by at least one young inventor who is more likely to have childcare responsibilities. In one of our analyses, we define young inventors as those whose earliest filing years occur after 2010 (exclusive). Assuming that an inventor filed his or her first patent application at the age of 20, given that our patent applications span from 2011 to 2021, young inventors are considered as those aged under 31.

The corresponding results are reported in column (3) of both Table C2 and Table C3, where we include a control for whether a patent application involves at least one young inventor whose first patent application was filed after 2010. Compared to patent applications with all inventors who started their filings before 2010, patent applications with young inventors exhibit a slightly lower level of novelty before COVID-19, albeit not statistically significant. However, after the onset of COVID-19, patent applications with young inventors experienced a substantial and statistically significant reduction in novelty. This observation aligns with the notion that young inventors are more likely to be parents of young children and, consequently, had to allocate more time to childcare and domestic responsibilities during the pandemic. In column (4), we concurrently control for the attributes of a patent application that may involve a female inventor or a young inventor, along with the interaction of those attributes with a binary variable indicating the post-COVID-19 period. The results remain robust across these specifications.

Despite accounting for all the different features, the coefficients associated with the interaction of population density and the post-COVID-19 time dummy variable in columns (2)-(4) resemble that in column (1). The robust coefficients associated with the interaction term suggest that the decrease in patent innovativeness following COVID-19 is unlikely explained by a disproportionate reduction in working time for women or young inventors with children.

While we consider inventors whose earliest patent application was filed after 2010 as proxies for inventors with young children, we also explore alternative cutoff years, namely 2005 and 2015, and report the results in Table C4. Columns (1)-(2) of Table C4 present findings when young inventors are defined as those with earliest filing years after 2005, and columns (3)-(4) report findings for young inventors defined by the earliest filing years after 2015. Once again, our analysis provides robust evidence concerning the extent to which the advantages of big cities in patent innovativeness have declined after the outbreak of COVID-19.

Last but not least, we examine the possibility that our results may be influenced by the potential scenario of the most innovative inventors relocating to less densely populated CSDs. To address this concern, we introduce a separate control for whether a patent application involves “mover” inventors (who have moved to different CSDs after the year 2020), along with its interaction with a binary variable indicating whether the patent is filed after the onset of COVID-19. The results are presented in Table C5. Our baseline findings on the reduced density premium in patent informativeness remain robust.

Table C1: Summary Statistics (Application Level Analysis)

Variable	Mean	Std. Dev.	Min	Max	Obs.
Population density in the densest CSD ($/km^2$)	2,208.856	3,833.554	0.174	26,821.9	1,739,130
Application with female inventors	0.236	0.424	0	1	1,635,125
Application with young inventors whose earliest filing year > 2005	0.742	0.438	0	1	1,739,130
Application with young inventors whose earliest filing year > 2010	0.543	0.498	0	1	1,739,130
Application with young inventors whose earliest filing year > 2015	0.227	0.419	0	1	1,739,130
Application with inventors moved after 2020	0.331	0.471	0	1	1,712,785
Inventor team size	3.102	2.202	1	133	1,739,130
Top 1% patent application	0.010	0.098	0	1	1,739,130
Top 5% patent application	0.049	0.215	0	1	1,739,130

Notes: This table reports summary statistics at the patent application level. The variables include the population per square kilometer of the most populated CSD among all CSDs where inventors reside; the presence of at least one female inventor; whether an application for a patent has at least one young inventor whose first filing year is after 2005, after 2010, or after 2015; whether at least one inventor on a patent application relocated after 2020; team size, which refers to the total number of inventors on a patent application; whether the patent belongs to the top 1% of the novelty distribution and top 5% of the novelty distribution, respectively.

Table C2: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% in Novelty (Controlling for Female and Young Inventors at the Patent Level)

	(1)	(2)	(3)	(4)
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-0.001	-0.001	-0.001	-0.001
	(0.000)	(0.000)	(0.000)	(0.000)
$\mathbb{1}$(female)	-	0.006	-	0.006
	-	(0.001)	-	(0.001)
$\mathbb{1}$(after 2020) $\times$ $\mathbb{1}$(female)	-	-0.004	-	-0.004
	-	(0.001)	-	(0.001)
$\mathbb{1}$(young)	-	-	-0.002	-0.002
	-	-	(0.000)	(0.000)
$\mathbb{1}$(after 2020) $\times$ $\mathbb{1}$(young)	-	-	-0.001	-0.002
	-	-	(0.001)	(0.001)
Team size	Yes	Yes	Yes	Yes
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Mean of DV	0.010	0.010	0.010	0.010
Adj. R-squared	0.020	0.021	0.020	0.021
Obs.	1,739,122	1,635,107	1,739,122	1,635,107

Notes: All regressions follow Equation (11). The dependent variable is a dummy variable indicating whether a patent belongs to the top 1% of the novelty distribution. An inventor is considered young if their first filing year occurs after 2010. Standard errors clustered at the CSD level are reported in parentheses.

Table C3: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 5% in Novelty (Controlling for Female and Young Inventors at the Patent Level)

	(1)	(2)	(3)	(4)
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-0.004	-0.003	-0.004	-0.003
	(0.001)	(0.001)	(0.001)	(0.001)
$\mathbb{1}$(female)	-	0.027	-	0.027
	-	(0.003)	-	(0.003)
$\mathbb{1}$(after 2020) $\times$ $\mathbb{1}$(female)	-	-0.012	-	-0.011
	-	(0.002)	-	(0.002)
$\mathbb{1}$(young)	-	-	-0.001	-0.002
	-	-	(0.001)	(0.001)
$\mathbb{1}$(after 2020) $\times$ $\mathbb{1}$(young)	-	-	-0.008	-0.008
	-	-	(0.002)	(0.002)
Team size	Yes	Yes	Yes	Yes
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Mean of DV	0.049	0.049	0.049	0.049
Adj. R-squared	0.040	0.043	0.040	0.043
Obs.	1,739,122	1,635,107	1,739,122	1,635,107

Notes: All regressions follow Equation (11). The dependent variable is a dummy variable indicating whether a patent belongs to the top 5% of the novelty distribution. An inventor is considered young if their first filing year occurs after 2010. Standard errors clustered at the CSD level are reported in parentheses.

Table C4: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% and 5% in Novelty (Controlling for Female and Young Inventors at the Patent Level with Alternative Young Inventor Proxies)

	(1) Top 1%	(2) Top 1%	(3) Top 5%	(4) Top 5%
$\mathbb{1}$(after 2020) $\times$ log(pop density)	-0.001	-0.001	-0.003	-0.003
	(0.000)	(0.000)	(0.001)	(0.001)
$\mathbb{1}$(female)	0.006	0.006	0.028	0.027
	(0.001)	(0.001)	(0.003)	(0.002)
$\mathbb{1}$(after 2020) $\times$ $\mathbb{1}$(female)	-0.004	-0.004	-0.011	-0.011
	(0.001)	(0.001)	(0.002)	(0.002)
$\mathbb{1}$(young(2005))	-0.003	-	-0.006	-
	(0.000)	-	(0.001)	-
$\mathbb{1}$(after 2020) $\times$ $\mathbb{1}$(young(2005))	-0.002	-	-0.007	-
	(0.001)	-	(0.002)	-
$\mathbb{1}$(young(2015))	-	0.000	-	0.005
	-	(0.001)	-	(0.001)
$\mathbb{1}$(after 2020) $\times$ $\mathbb{1}$(young(2015))	-	-0.003	-	-0.011
	-	(0.001)	-	(0.002)
Team size	Yes	Yes	Yes	Yes
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Mean of DV	0.010	0.010	0.049	0.049
Adj. R-squared	0.021	0.021	0.043	0.043
Obs.	1,635,107	1,635,107	1,635,107	1,635,107

Notes: All regressions follow Equation (11). The dependent variable is a dummy variable indicating whether a patent belongs to the top 1% (Columns (1) and (2)) or the top 5% of the novelty distribution (Columns (3) and (4)). An inventor is considered young if their first filing year occurs after 2005 or after 2015. Standard errors clustered at the CSD level are reported in parentheses.

Table C5: The Impact of COVID-19 on Density Premium of the Likelihood of Being the Top 1% or 5% in Novelty (Controlling for Movers)

	(1) Top 1%	(2) Top 1%	(3) Top 5%	(4) Top 5%
$\mathbb{1}$(after 2020) $\times$ *log*(pop density)	-0.001	-0.001	-0.004	-0.004
	(0.000)	(0.000)	(0.001)	(0.001)
$\mathbb{1}$(mover)	-	0.001	-	0.003
	-	(0.001)	-	(0.001)
$\mathbb{1}$(after 2020) $\times$ $\mathbb{1}$(mover)	-	-0.002	-	-0.006
	-	(0.001)	-	(0.002)
Team size	Yes	Yes	Yes	Yes
CSD FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Mean of DV	0.010	0.010	0.049	0.049
Adj. R-squared	0.020	0.020	0.040	0.040
Obs.	1,739,122	1,712,770	1,739,122	1,712,770

Notes: All regressions follow Equation (11). The dependent variable is a dummy variable indicating whether a patent belongs to the top 1% (Columns (1) and (2)) or the top 5% of the novelty distribution (Columns (3) and (4)). A mover is an inventor who relocated to a different CSD after 2020.

D Derivations

D.1 Derivation of Prediction 3

Let $\tilde{N}_n$ be the degree of innovativeness in location n following the restrictions, and let $\Delta\bar{N}_n = \tilde{N}_n - \bar{N}_n$ be the resulting change in innovativeness. Notice that

$$\frac{\partial \Delta\bar{N}_n}{\partial D_n^\gamma} = \frac{\partial \tilde{N}_n}{\partial D_n^\gamma} - \frac{\partial \bar{N}_n}{\partial D_n^\gamma} = \frac{\partial \tilde{N}_n}{\partial D_n^\gamma} - \left[\frac{\partial \bar{N}_n}{\partial D_n^\gamma}\Big|_{\alpha_n^*} + \frac{\partial \bar{N}_n}{\partial \alpha_n^*}\frac{\partial \alpha_n^*}{\partial D_n^\gamma} \right].$$

Prediction 3 states that $\frac{\partial \Delta\bar{N}_n}{\partial D_n^\gamma} < 0$ provided that $\tilde{\alpha}$ is sufficiently small. Since $\frac{\partial \bar{N}_n}{\partial \alpha_n^*}\frac{\partial \alpha_n^*}{\partial D_n^\gamma} > 0$, it is sufficient to show that $\frac{\partial \tilde{N}_n}{\partial D_n^\gamma} < \frac{\partial \bar{N}_n}{\partial D_n^\gamma}\Big|_{\alpha_n^*}$.

Expanding the two terms:

$$\frac{\partial \tilde{N}_n}{\partial D_n^\gamma} = \frac{\tilde{\alpha}(1-\tilde{\alpha})B_{OL}(N_{IP}-N_{OL})}{[(1-\tilde{\alpha})B_{OL}+\tilde{\alpha}D_n^\gamma]^2}$$

$$\frac{\partial \bar{N}_n}{\partial D_n^\gamma}\Big|_{\alpha_n^*} = \frac{\alpha_n^*(1-\alpha_n^*)B_{OL}(N_{IP}-N_{OL})}{[(1-\alpha_n^*)B_{OL}+\alpha_n^*D_n^\gamma]^2}.$$

Simple algebra shows that $\frac{\partial \tilde{N}_n}{\partial D_n^\gamma} < \frac{\partial \bar{N}_n}{\partial D_n^\gamma}\Big|_{\alpha_n^*}$ is equivalent to the following condition:

$$B_{OL}^2\,(1-\tilde{\alpha})\,(1-\alpha_n^*) > D_n^{2\gamma}\,\tilde{\alpha}\,\alpha_n^*. \tag{12}$$

This condition is always satisfied for $\tilde{\alpha} = 0$ and, by continuity, for positive but sufficiently small values of $\tilde{\alpha}$. In particular, since α_n^* is increasing in D_n, condition (12) is always satisfied if it holds for $D_n = D_{\max}$. Hence, a sufficient condition on $\tilde{\alpha}$ is the following:

$$\tilde{\alpha} < \frac{B_{OL}^2\,(1-\alpha_{\max}^*)}{B_{OL}^2\,(1-\alpha_{\max}^*) + D_{\max}^{2\gamma}\,\alpha_{\max}^*}.$$